ISSN : 3049-1479 (Online)

SANKALCHAND PATEL UNIVERSITY

।। अथातो ज्ञानजिज्ञासा ।।

FOR A BETTER TOMORROW

Special Edition

(ICAICS-2026)

SANKALCHAND PATEL UNIVERSITY
JOURNAL OF SCIENCE, TECHNOLOGY AND MANAGEMENT RESEARCH

(SPU-JSTMR)

RESEARCH JOURNAL

International Conference on Artificial Intelligence and Cybersecurity – 2026 (ICAICS-2026)

Sankalchand Patel University Journal of Science, Technology and Management Research (SPU-JSTMR)

Peer–Reviewed University Journal

ISSN: 3049-1479 (Online)

International Conference on

Artificial Intelligence and Cybersecurity – 2026

(ICAICS-2026)

(Special Edition)

Chief Editor

Prof. (Dr.) Prafulkumar Udani

International Conference on

Artificial Intelligence and Cybersecurity – 2026

ICAICS-26

Organized by

Sankalchand Patel University

6th – 7th March 2026

Venue

Entrepreneurship Development Institute of India
Ahmedabad, Gujarat, India

In association with

International Conference on

Artificial Intelligence and Cybersecurity – 2026

ICAICS-26

6th – 7th March 2026

Chief Patron

Shri Prakashbhai Patel
President, Sankalchand Patel University, India

Patron

Prof. (Dr.) Prafulkumar Udani
Provost, Sankalchand Patel University, India

Convener

Prof. (Dr.) P. J. Patel
Principal, Sankalchand Patel College of Engineering
Sankalchand Patel University, Visnagar

Conference Secretary

Prof. (Dr.) Manish M Patel,
Dr. Bibhabasu Mohanty
Sankalchand Patel College of Engineering, Visnagar

About University

Nootan Sarva Vidyalaya Kelavani Mandal (NSVKM) was established in 1952, by Shri. Sankalchand Patel, a freedom fighter and a social entrepreneur. The aim was always to create a value based diversified education system for the students of North Gujarat. The university was established on 9th May 2016 vide Gujarat Private Universities (Amendment) Act, 2016 and is spread across 84 acres of land. Today, the university is home to various institutes offering Medical, Paramedical, Technical, & Science Courses. The university also houses some of the most state-of-the-art infrastructure in North Gujarat.

Today, the trust successfully administers educational institutions right from primary schools to graduate and post graduate colleges and other professional courses including Ph.D. programs. The trust's unswerving efforts towards the growth of education have earned its institutes many merits and a resilient name in the educational campaign of India.

Our Trust has the experience of the education field of 65 years at its disposal and it is fully equipped with all infrastructure and computer facilities. Almost 66 acres of land has been acquired to construct various buildings for various institutions. Sankalchand Patel University is committed to provide high quality education through Quality Teaching, Research, Consultancy and Community outreach services. The University is committed to continuous improvement in the academic and administrative processes by fostering Cooperation, Innovation, Professionalism, and Integrity. The vidhyadham is situated near Gandhinagar-Ambaji link road in VISNAGAR and surrounded by natural surroundings in the lap of nature.

About Institute

Sankalchand Patel College of Engineering established in 1999 by Nootan Sarva Vidhyalay Kelvani Mandal – VISNAGAR with an aim of producing world-class professionals in the specialized fields of Engineering, Management, and MCA. The Institute is constituted by Sankalchand Patel University,Visnagar and approved by All India Council for Technical Education (AICTE), New Delhi.

Sankalchand Patel College of Engineering was established with an objective of producing world-class professionals in the field of Engineering and its branches. Institute offer various courses in B.Tech, Bsc.IT, and M.Tech Programs. Great emphasis is laid upon the holistic development of the students, thus helping them build up academic excellence which in turn helps with the right career opportunities. The aim is to not only create dynamic and enthusiastic professionals, but also good citizens who will aid in the overall progress and development of this country. An institute is committed to provide excellence in various ways, such as ensuring high-quality education, research, and infrastructure, or through commitment to student support and a positive learning environment.

President's Message

Shri Prakashbhai S. Patel

Sankalchand Patel University (SPU) is one of the leading universities in the North Gujarat region, offering programs in the domains of Medical, Dental, Physiotherapy, Pharmacy, and various technical streams on a single campus. SPU strives to achieve our mission "serve society to develop and prosper by building human capital."

It gives me immense pleasure to extend my warm greetings to all Conference Delegates, Research Scholars, Students, and Invited Keynote Speakers. SPU is regularly organizing conferences and seminars as its commitment in promoting of knowledge and innovation.

Artificial Intelligence and Cybersecurity represent two of the most influential emerging technologies driving global transformation in the twenty-first century. The increasing integration of AI enhances efficiency, innovation, and safety across cyber connected systems. However, these advancements also demand robust cybersecurity frameworks to safeguard data, privacy, and digital infrastructure. The themes of this conference emphasize the importance of developing secure, ethical, and sustainable technological solutions.

This conference will provide a platform for intellectual engagement on critical issues such as intelligent decision-making systems, cyber resilience, data privacy, and the ethical governance. Conference will provide insight to participants to explore innovative approaches that not only tackles present challenges but also anticipate future technological requirements. By fostering dialogue among researchers and practitioners, this event seeks to contribute for the creation of secure digital ecosystems and responsible technological progress.

I am confident that the exchange of ideas and collaborative discussions during this conference will generate new insights, strengthen research networks, and inspire future innovations. Such academic initiatives play a vital role in building a technologically advanced and secure society.

I extend my sincere appreciation to the organizing team and wish the conference a great success in achieving its objectives.

Sincerely,
Prakashbhai Patel
President
Sankalchand Patel University

Provost's Message

Prof. (Dr). Prafulkumar Udani

The International Conference on Artificial Intelligence and Cybersecurity 2026 (ICAICS–26), organized by Sankalchand Patel College of Engineering, SPU, aims to promote research and development in various domains of Artificial Intelligence, including Healthcare, Agriculture, Smart Manufacturing, Governance. Cybersecurity – 2nd track of this conference will discuss issues relevant to physical layer, OS, Data and Applications.

The conference has received an overwhelming response from faculty members, research scholars, industry professionals, research organizations, start-ups and final year students of B.Tech (AI & DS, CS, IT).

Invited Keynote Speakers from ISRO, C-DAC, start-ups, industries, government, and academia will share their expertise and experiences related to the India AI Mission, AI applications in Agriculture and Health Sciences, Data Science, Cybersecurity and this will be of great learning for all delegates. I am confident that this two-day conference will provide an in-depth understanding relevant to selected AI and Cybersecurity tracks as well as valuable networking opportunities and peer learning.

ICAICS–26 will showcase the research achievements of our 40+ faculty members and 400+ SPCE students registered in online and offline modes will benefit immensely.

Recent announcements made at India AI Impact summit held in Delhi pertaining to compute power (GPUs 38000+20000), available data sets, models and investments ($200 billion plus) will provide large scale opportunities to AI Skilled workforce as well as IT industries in India.

I sincerely appreciate and acknowledge the efforts of the SPU team in organizing this conference and extend my heartfelt thanks to all the delegates for their participation and making this conference a great success.

Warm regards,

Prof. (Dr). Prafulkumar Udani
Provost, Sankalchand Patel University

Convener's Message

Dr. Piyushkumar J. Patel

Dear Esteemed Participants and Distinguished Guests,

It is my great pleasure to welcome all delegates to the *International Conference on Artificial Intelligence and Cybersecurity 2026 (ICAICS–26)*. The conference has been carefully designed to provide a dynamic platform for exploring contemporary research themes and emerging trends in intelligent technologies and digital security.

The focus of this event lies in examining the practical and theoretical dimensions of Artificial Intelligence and Cybersecurity. It seeks to address critical areas such as intelligent data analysis, automated systems, cyber risk management, secure software development, and innovative approaches to protecting digital infrastructure. The conference also highlights the importance of interdisciplinary collaboration in developing solutions that meet the evolving demands of modern technological environments.

In addition to presenting cutting-edge research, the conference encourages dialogue on future perspectives, including the role of Artificial Intelligence in strengthening cybersecurity mechanisms, the challenges of securing large-scale data systems, and the ethical implications of autonomous technologies. These discussions aim to promote responsible innovation and contribute to the advancement of secure and intelligent technological frameworks.

The successful organization of this conference is the result of the dedicated efforts of the organizing committee, reviewers, and contributors. I sincerely thank all participants for their valuable involvement and hope that this conference provides opportunities for meaningful engagement, knowledge exchange, and professional growth.

Warm regards,

Dr. Piyushkumar J. Patel
Convener
(ICAICS–26)

Conference's Secretary Message

Dr. Manish M. Patel

Dear Guests and Participants,

On behalf of the Conference Secretariat, I extend a cordial welcome to all participants of the *International Conference on Artificial Intelligence and Cybersecurity 2026 (ICAICS–26)*. The Secretariat has endeavored to create an academic environment that supports productive discussions, effective knowledge dissemination, and meaningful collaboration among participants from diverse fields.

The conference emphasizes future-oriented themes that address the evolving landscape of intelligent computing and digital security. As technological systems become increasingly complex and interconnected, there is a growing need to develop resilient cybersecurity strategies, promote ethical standards in Artificial Intelligence, and ensure the protection of digital resources. The conference aims to facilitate dialogue on these important issues and to encourage the exchange of ideas that contribute to long-term technological sustainability.

Through coordinated sessions, research presentations, and academic interactions, the conference provides opportunities for participants to engage with emerging research trends and explore innovative solutions to contemporary challenges. The Secretariat remains committed to ensuring the smooth conduct of the conference and supporting an atmosphere conducive to scholarly exchange.

We sincerely hope that this conference fosters new collaborations, enhances research initiatives, and contributes to shaping the future of Artificial Intelligence and Cybersecurity. We express our gratitude to all participants, speakers, and organizing members for their cooperation and wish everyone a productive and memorable experience.

Warm regards,

Dr. Manish M. Patel
Conference Secretariat,
(ICAICS–26)

About the Conference (ICAICS-26)

Sankalchand Patel University successfully organized the International Conference on Artificial Intelligence and Cybersecurity – 2026 (ICAICS-26) on 6–7 March 2026 at the Entrepreneurship Development Institute of India, Ahmedabad, Gujarat, India. The two-day international conference served as a significant platform for researchers, academicians, industry professionals, and students to exchange ideas, present research findings, and explore emerging innovations in the rapidly evolving domains of Artificial Intelligence and Cybersecurity.

The conference witnessed remarkable participation from across India and abroad and aimed to promote interdisciplinary research, technological innovation, and collaboration between academia and industry.

Conference Overview

ICAICS-26 received an overwhelming response from the research community. The conference recorded more than 600 paper submissions, out of which 190 high-quality research papers were accepted after a rigorous peer-review process. The conference was attended by over 400 participants through both online and offline modes, demonstrating strong global academic engagement.

The conference was structured around 16 specialized thematic tracks focusing on various emerging areas within Artificial Intelligence and Cybersecurity.

Another notable achievement was the research contribution from Sankalchand Patel University, where 40 papers were presented with SPU researchers as first authors and 37 papers with SPU researchers as co-authors, reflecting the institution's strong commitment toward research excellence.

Inaugural Ceremony

The conference commenced with a formal Inaugural Ceremony in the esteemed presence of university leadership and distinguished guests.

The event was graced by:

- Shri Prakashbhai Patel, President, Sankalchand Patel University
- Dr. Prafulkumar Udani, Provost, SPU

Along with distinguished dignitaries including:

- Madan Mohan Tripathi – Director General, National Institute of Electronics & Information Technology
- Padmaja Joshi – Scientist-G, Centre for Development of Advanced Computing
- Dr. Manjusree Panchagnula- Group Head at Remote Sensing Applications Area, National Remote Sensing Centre, ISRO
- Sunny Vaghela – CEO & Founder, Tech defense Labs, Ahmedabad

During the inaugural addresses, the dignitaries emphasized the growing global relevance of Artificial Intelligence technologies and the urgent need to develop strong cybersecurity frameworks to safeguard digital infrastructure. They highlighted the importance of collaborative research, innovation, and academic partnerships in addressing modern technological challenges.

Keynote Sessions

One of the major highlights of ICAICS-26 was the series of keynote lectures delivered by eminent experts from national and international organizations.

Day 1: Keynote Speakers

The conference featured keynote sessions by:

- Madan Mohan Tripathi – Director General, NIELIT
- Padmaja Joshi – Scientist-G, CDAC
    - AI in cybersecurity
- Swadhin K. Behera – Director, Application Laboratory, Japan Agency for Marine-Earth Science and Technology (Japan) – delivered online
    - Integrating Ocean Observations, AI, and Digital Twins for Delivering Actionable Climate Information
- Manjusree Panchagnula – Group Head, Remote Sensing Applications Area, National Remote Sensing Centre, Indian Space Research Organisation

The keynote sessions focused on emerging developments in Artificial Intelligence, cybersecurity challenges, geospatial intelligence, and advanced remote sensing technologies.

Day 2: Keynote Speakers

The second day included several insightful keynote sessions delivered by renowned experts such as:

- S. Bandyopadhyay – Ministry of Agriculture and Farmers Welfare
    - Topic: *Reshaping Agricultural Landscape using Geospatial Intelligence*
- Ranendu Ghosh – Co-Founder & CTO, SatLeo Labs and former scientist, ISRO
    - Topic: *AI in Geospatial Data Processing*
- Sanjay Keshva – Indian Police Service
    - Topic: *Cyber Fraud Prevention*
- Pavan Saxena – Cybersecurity Expert, Tech defense Labs
    - Topic: *Cybersecurity Awareness*
- Madhukant Patel – Former Scientist, ISRO
    - Topic: *Digital Farming and Emerging Technologies*
- Mr. Dhruv Pandit–Drona Cyber solutions Pvt Ltd, Ahmedabad

These sessions provided participants with valuable insights into real-world applications of Artificial Intelligence, geospatial intelligence, digital agriculture, and cybersecurity risk mitigation strategies.

Technical Sessions and Paper Presentations

A major component of ICAICS-26 was the technical paper presentation sessions.

More than 190 research papers were presented during the conference across multiple thematic tracks and parallel technical sessions. The presentations were conducted in both online and offline modes, allowing wider participation from researchers worldwide.

The technical sessions covered a wide range of topics including:

- Artificial Intelligence applications
- Machine learning and deep learning
- Cybersecurity frameworks and risk management
- Data science and analytics
- Intelligent systems and automation
- Secure communication networks

These sessions facilitated active academic discussion, peer learning, and collaborative opportunities among researchers.

Poster Presentation Competition

To encourage student innovation and research participation, a Poster Presentation Competition was organized during the conference.

A total of 31 posters were presented by students showcasing innovative ideas and technological solutions addressing current challenges in Artificial Intelligence and Cybersecurity.

The poster session provided students with an opportunity to:

- Present their research ideas
- Interact with experts and academicians
- Receive constructive feedback
- Enhance their research presentation skills

Distinguished Speakers and Global Participation

ICAICS-26 hosted more than 15 distinguished speakers from reputed organizations and institutions including:

- National Institute of Electronics & Information Technology
- Centre for Development of Advanced Computing
- Japan Agency for Marine-Earth Science and Technology
- Indian Space Research Organisation
- Mahalanobis National Crop Forecast Centre
- SatLeo Labs
- Tech defence Labs
- Drona Cybersolutions Pvt Ltd, Ahmedabad

The conference also witnessed participation from national and international universities, government organizations, research institutes, and industry partners, highlighting the global academic outreach of the event.

Valedictory Ceremony

The conference concluded with a Valedictory Ceremony on 7 March 2026. The ceremony featured an insightful address by C. P. Singh, who spoke about the transformative role of Artificial Intelligence in advanced technological fields and future digital ecosystems.

During the ceremony:

- Best Paper Presentation Awards were presented to outstanding researchers in recognition of their innovative research contributions.
- Participants were appreciated for their scholarly work and active involvement in the conference.

Closing Remarks

On this occasion, Dr. Prafulkumar Udani, Provost of Sankalchand Patel University, congratulated all participants, speakers, organizing committee members, and volunteers for the successful completion of the conference. He acknowledged the collective efforts that made ICAICS-26 a meaningful platform for academic dialogue, technological innovation, and collaborative research.

The conference witnessed enthusiastic participation from researchers, academicians, industry professionals, and students, reinforcing the university's commitment to promoting research excellence, technological advancement, and global academic collaboration.

ICAICS-26 successfully created a dynamic environment for knowledge exchange and innovation, further strengthening the position of Sankalchand Patel University as a growing hub for advanced research and international academic engagement.

Research/Review Papers on Science, Technology and Management Research

ICAICS-26

(Special Edition)

Peer –Reviewed University Journal

Chief Editor : **Prof. (Dr.) Prafulkumar Udani, Provost,** Sankalchand Patel University (SPU), Visangar

Editor(s) :

- **Dr. Rajesh P. Patel, Associate Professor,** Department of Computer Engineering, Sankalchand Patel College of Engineering, SPU, Visngar

Advisory Board:

- **Dr. Hitesh H. Mehta, Principal,** Nootan Science College, SPU, Visnagar
- Dr. Amit Ganatra, Provost (Vice Chancellor), Parul University, Vadodara
- Dr. Ramesh Kothari, Professor and Head, Department of Biosciences, Saurashtra University, Rajkot
- Dr. Digvijaysinh Rathod, Associate Dean & Associate Professor, NFSU, Gandhinagar
- Dr. Ashish Joshi, Professor & Head, Kaushalya The Skill University, Gandhinagar
- Dr. Krupesh A. Chauhan, Professor, Department of Civil Engineering Sardar Vallabhbhai National Institute of Technology (SVNIT), Surat
- Dr. N. P. Patidar, Professor of Electrical Engineering, MANIT, Bhopal
- Dr. Siddharth Joshi, Assistant Professor, Department of Electrical Engineering, School of Energy Technology, Gandhinagar
- Dr. Dhamodharan Kondusamy, Assistant Professor, Department of Environmental Engineering, Thapar University, Punjab
- Dr. Rajesh V. Savalia, Assistant Professor, Department of Mathematical Science, P.D. Patel Institute of Applied Sciences, CHARUSAT, Changa
- Dr. Ankur D. Amin, Assistant Professor, Department of Business Studies, Sardar Patel University, Vallabh Vidyanagar

Published By : Sankalchand Patel University, Visnagar 384315, India

---

Note: All inquiries regarding the Journal may please be sent to
Editor, SPU-JSTMR, Sankalchand Patel University, Visangar 384315, India through
Email: communications.spujstmr@spu.ac.i

Chief Editor's Message

It is with immense pride and professional satisfaction that I present this Special Edition of SPU-Journal of Science, Technology and Management Research (SPU-JSTMR) dedicated exclusively to the International Conference held on 6 & 7th March, 2026 at Sankalchand Patel Univeristy on Modern Approaches in Engineering, Technology, and Management (ICAICS-26) This edition represents the collective intellectual efforts and innovative contributions of researchers, academicians, industry experts, and scholars from across disciplines.

ICAICS-26 has provided a dynamic platform for knowledge exchange and interdisciplinary dialogue, emphasizing the need for modern and sustainable solutions in engineering, technology, and management. The conference brought together leading researchers, practitioners, and scholars to share the latest advancements in science, technology and management areas particularly in AI/ML and Cyber Security domain. The SPU-JSTMR Journal is peer-reviewed and open access journal of the university. The conference papers published in the special edition of this journal represent some of the most compelling contributions, showcasing both theoretical innovations and practical applications. Each paper has undergone a rigorous peer-review process to ensure the highest academic standards, and in many cases, the authors have extended or revised their work based on feedback received during the conference and through the journal's editorial review.

We sincerely express our gratefulness to Honorable President Shri Prakashbhai Patel, for all his support in undertaking the publication of research articles and perfectly completing the task. We thank the conference organizers, keynote speackers, reviewers, authors, editors and journal committee members for their contributions to this issue.

We hope that this special edition will inspire continued research, collaboration, and academic excellence in the years to come.

Sincerely,

Chief-Editors, SPU-Journal of Science, Technology and Management Research

Prof. (Dr.) Prafulkumar Udani
Provost,Sankalchand Patel University

CONTENTS

Sr. No.	Paper Title	Page No.
1	Multiple Face Recognition Using Genetic Algorithm ***Om Deokar, Rajeev Patil, Ankit Kumar, Abhijeet Gite, Vijaykumar P. Mantri***	01
2	AI-Predator Swarms- Autonomous Mass-Scale Cyber Attacks ***Alex Mathew, Leah George***	10
3	An Analytical Framework for Reconstructing Missing Time-Series Data Using Interpretable AI Methods ***Sneh B. Patel, Dr. Darshanaben D. Pandya***	18
4	From Traditional Machine Learning to Transformers: Sentiment Analysis on Social Media Data ***Patel Charmi D., Mevada Jayesh M., Patel Gaurang J., Patel Amit A.***	26
5	A Comparative Performance Analysis of Machine Learning Models for Cardiovascular Disease Prediction Using Hyper parameter Tuning ***Ms. Maitri Bhavsar, Dr. Manish Patel***	35
6	AI-Driven Crop Detection and Plant Disease Prediction ***Vismay D Shah, Dr. Darshanaben D. Pandya, Trushar Patel, Chetan R. Dudhagara***	43
7	The quantum green shield: integrating computing principles in cyber defence ***Abhisekh Tiwari , Dr. Pooja Tripathi***	61
8	AI-Enabled Smart Irrigation and Crop Planning Framework for Water-Scarce Rural Regions Using Interpretable Machine Learning: A Case Study of Sabarkantha (Khedbrahma), Gujarat ***Dr. Jayesh Gamar, Himaniben Gajjar***	71
9	Content Disarm and Reconstruction a Forensic Platform ***VulupalaBal Reddy, Dr. Deepti Vidyarti , Raj Dev Gangwar***	79
10	A Secure Blockchain-Based Mobile Application Frame- work for Art Authentication and Empowerment ***Himaniben Gajjar , Dr. Jayesh J. Gamar***	90
11	Multifunctional Smart Real Estate Platform with Predictive and Utilty Features ***Ravijeet Sharma, Vansh Verma, Harsh Vyavahare, Mrs. Jagruti Jadhav***	99
12	Machine Learning and Deep Learning Techniques for the Prediction of Autism Spectrum Disorder A Comprehensive Review ***Gaurangkumar Patel, H.B. Pandya , B.H. Trivedi***	112
13	Enhanced Fashion Demand Forecasting Using a Hybrid CNN–LSTM Deep Learning Framework ***Hemangini V. Patel, Kirit J. Modi, Khyati A. Prajapati, Ekta V. Patel***	121
14	An ontology driven semantic framework for risk identification in large scale religious mass gatherings ***Govind Kushwah, Preetvanti Singh***	132

15	AI-Driven Predictive Analysis Framework for Student Performance Prediction in Smart Education Systems ***Dr. Dipikaben Umakant Thakar, Mr.Yogesh Patel, Mr. Shahnavaj khan Pathan, Mr.Akash Patel***	143
16	Artificial Intelligence-Based Multi-Document Extractive Summarization Using Sentence-Level Features ***Karina Patel, Shilpa Serasiya***	153
17	Early Detection Of Chronic Diseases Using Qunatum Machine Learning Algorithms ***N. Deepika, P. Rani Annapurna, T. SatyaKavyanjali , V. Mahesh, M. Sriramulu***	161
18	Temporal Blindness in AI-Secured IoT Systems: Missed Multi-Stage Attacks in Cloud Centric Detection Pipelines ***ParthTalaviya, Manya Parikh, Deep Solanki, Jigar Gajjar***	168
19	False Sense of Security Caused by AI-Driven Threat Detection ***ParthTalaviya, Manya Parikh, Deep Solanki,Jigar Gajjar***	181
20	Comparative Performance Evaluation of P&O and Fuzzy Logic P&O MPPT for Grid-Connected PV Systems ***Ravi Bhavsar, Rakeshkumar Patel***	193
21	Advances in Automated Leukemia Detection Using Deep Analysis of WBC Scatterplot Patterns ***Nikul Vikaskumar Jayswal, Dr. Kirit J. Modi***	201
22	Towards Smart Agriculture: Deep Learning Methods for Plant Disease Monitoring ***Apoorva Dwivedi, Pharindra Kumar Sharma, Gourav Kumar Sharma***	208
23	Advancement in Automated Hate Speech Detection: A Survey of Machine Learning and Ensemble Methods ***Rutuja Goradkar, Smita Bharne, Vaibhav Narawade***	223
24	AI-Driven Cyber security in Smart Agriculture: Threats, Challenges, and Secure System Architectures ***Chintan Hirpara, Khushi Kundariya, Mahipal Chothani, Deep Solanki, Jigar Gajjar***	234
25	Comparative Study of Leading AI Models: Performance, Capabilities, Accuracy, and Applications ***Chintan Hirpara, Khushi Kundariya, Mahipal Chothani, Deep Solanki, Jigar Gajjar***	246
26	A Deep Learning Framework for Gender Estimation from Dental Radiographs Using ResNet-50 ***Jinal N. Raval, Mehul S. Patel, Rupal R. Chaudhari, Jayesh M. Mevada, Govind G. Patel***	256
27	Next-Generation Digital Biology: Technological Advancements and Security of Digital Medical Data ***Aateka Memon, Deep Solanki, Jigar Gajjar***	263
28	DIALLOps: A DevOps-Integrated Framework for Secure and Reliable Large Language Model Operations ***Vignesh Mudaliyar, Mohit Rajput, Deep Solanki, Jigar Gajjar***	273
29	Security Analysis of Address Resolution Protocol and Its Exploitable Vulnerabilities ***Ritik Kalal, Harsh Mandaliya, Dhruv Patel, Deep Solanki, Jigar Gajjar***	299
30	AI-Based ANPR System for Intelligent Vehicle Speed Detection ***Chennakesava Reddy S, Dr. Nityananatham Sampathkumar, AkshayKumar Varma S, Akhil Chowdary R, Ramakanth Reddy V.***	308

31	AI-Assisted Digital Forensics for Automated Incident Investigation ***Manya Parikh, Parth Talaviya, Deep Solanki, Jigar Gajjar***	317
32	A Comprehensive Study of Cyber Attacks Against Artificial Intelligence Systems ***Aayush Shah, Momin Samir, Deep Solanki, Jigar Gajjar***	325
33	An Analytical Framework for Examining AI Investment Trends in the Philippines, India, and Indonesia ***Dr.Darshanaben Dipakkumar Pandya, Dr. Daniel D. Dasigjr , Dr. Vijay Singhrathore, Dr. Ni Made Eka Mahadewi, Prof. Nita Jashugiri Goswami, , Ms. Patel kinal Bharatbhai***	338
34	Utilizing Blockchain Technology for Secure and Efficient Healthcare Data Management Contribution ***Dhara K. Patel, Dr. Kirit J. Modi***	345
35	A Systematic Literature Review of Student Performance Prediction Using Machine Learning Techniques ***Vishal P. Patel, Dr. Rajesh P. Patel***	352
36	Automating Diamond Quality Control using YOLO-Based Detection Model and Its Performance Result ***Ms.Diya Patel, Dr.Rajesh P. Patel, Mr.Vivek K. Shah***	362
37	Advancements in AI-based Robotics and Automation with Emerging Cybersecurity Threats and Challenges ***Doshi Krisha, Shah Diya, Deep Solanki, Jigar Gajjar***	369
38	A Smart IoT–AI Integrated System for Accurate Rainfall Prediction ***Ms. Bhatt Shreyaben Atulbhai, Dr. Khushbu, Dr. Darshanaben Dipakkumar Pandya***	380

Multiple Face Recognition Using Genetic Algorithm

Om Deokar[1], Rajeev Patil[2], Ankit Kumar[3], Abhijeet Gite[4], Vijaykumar P. Mantri[5]

Department of Computer Engineering MIT Academy of Engineering Pune, India[1]
Department of Computer Engineering MIT Academy of Engineering Pune, India[2]
Department of Computer Engineering MIT Academy of Engineering Pune, India[3]
Department of Computer Engineering MIT Academy of Engineering Pune, India[4]
Department of Computer Engineering MIT Academy of Engineering Pune, India[5]

om.deokar@mitaoe.ac.in ,rajeev.patil@mitaoe.ac.in[2] , ankit.kumar@mitaoe.ac.in[3] , abhijeet.gite@mitaoe.ac.in[4] , vijay.mantri@mitaoe.ac.in[5]

Abstract: In today's world, face recognition is very much important in modern biometric and surveillance applications,as well as its accuracy.Accuracy gets reduced because of various challenges in real-world like changes in illumination, facial expressions and also because of presence of multiple faces in a single image frame. This work explains an approach which is based on Genetic Algorithm (GA) to improve multiple face recognition; this can be done by the optimization of both feature selection and classifier parameters. Initially Face Net is used to extract face embeddings and then GA is applied to them to identify the most informative subsets of features. A chromosome is used to represent every candidate solution which is evaluated with a multi-objective fitness function that considers accuracy, computational efficiency and robustness with pose and occlusion variations.The implemented framework is scalable and can be used in various applications like authentication, surveillance and access control.

Keywords: Face Recognition, Genetic Algorithm, Feature Selection, Multi-Face Detection, Biometric Authentication

I. INTRODUCTION

The face recognition technology has emerged as a biometric technique in recent times as it operates via remote cameras and does not need physical contact and also because of broad applications in security, surveillance and authentication systems. Facial identification can be performed without physical contact like in fingerprint identification which makes it more practical for real-time use. Despite its advantages, accurate recognition is still a challenging job because of changes in illumination,facial expressions,pose of individuals in frame and partial occlusions. The complexity increases further in scenarios where multiple faces are within the same frame.

By representing candidate solutions as chromosomes and evolving them through operations such as selection, crossover and mutation, GAs discover optimal or near-optimal solutions for complex recognition problems. When GA is integrated with FaceNet embeddings the system automatically selects the most informative features, fine-tunes classifier parameters and improves recognition accuracy in multi-face scenarios.

II. LITERATURE REVIEW

There has been a significant evolution in the face recognition technology from many years from early statistical approaches to deep learning and optimization based methods to initial techniques, like Eigenfaces using Principal Component Analysis (PCA) [1] and Fisherfaces implementing Linear Discriminant Analysis (LDA) [3], were used as the basis for reducing dimensionality and feature extraction.
To solve issues of lighting and local texture variations, Local Binary Patterns (LBP) were introduced later which provided a strong alternative against changes in illumination and local facial features [2]. Other descriptors, like the histogram of Oriented Gradients (HOG), were applied to improve the resistance to pose and occlusion [4].

In recent times it has transitioned to machine learning and hybrid techniques. Classifiers such as SVM and kNN along with PCA and LBP features enhanced the recognition outcomes. [8]–[10].

http://doi.org/10.63766/spujstmr.26.000067

III. PROBLEM STATEMENT

Face recognition technology has become essential in applications such as identity verification, access control and human-computer interaction.Existing single-face recognition systems have shown notable progress over time however, they lag in accurately identifying multiple faces in a single frame. Lighting conditions, facial expressions that change and multiple individuals in a single image lower the performance of these systems.

IV. OBJECTIVES

1) **To build a strong framework for multi-face detection and recognition** which can handle real-world challenges like change in illumination ,variations in head pose,partial occlusions and provide high accuracy.
2) **To implement deep learning embedding for representation of features** FaceNet implementation to capture discriminative and compact facial descriptors that are useful for classification tasks.
3) **To employ Genetic Algorithm for feature optimization** in order to minimize redundancy, lower computational complexity and select only the most relevant features without compromising recognition performance.
4) **To evaluate the proposed approach** on publicly available multi-face benchmark datasets, ensuring that testing reflects realistic and unconstrained environments.
5) **To benchmark the system against existing methods** using evaluation parameters namely accuracy, precision, recall, F1-score and runtime efficiency.

V. METHODOLOGY

By integrating Genetic Algorithm (GA) with modern tools,detection and recognition techniques there can be a significant increase in the performance of multiple face recognition systems.GA which is an evolutionary optimization algorithm, efficiently selects and combines the most relevant features, improves recognition accuracy and mitigates the risk of convergence to local minima.

The selected features are classified using algorithms such as SVM, kNN or neural networks. The system is evaluated on benchmark datasets such as LFW, ORL and Yale, as well as a custom multi-face dataset. For validation of accuracy the performance metrics that are used are accuracy, precision, recall, F1-score and processing time. Tables and architecture diagrams are provided to illustrate the workflow and demon- strate the improvements achieved through GA optimization.

A. Tools and Techniques

The implementation of the multiple face recognition system using Genetic Algorithm (GA) includes specialized software tools, programming libraries and algorithmic methods. Local Binary Patterns (LBP), Principal Component Analysis (PCA) or Histogram of Oriented Gradients (HOG) techniques are used to extract facial features. In GA chromosomes these features act as genes, where each chromosome represents a candidate that maps features to identities. The similarity between the candidate solutions and the target features are evaluated by fitness function that uses Euclidean distance, cosine similarity or correlation coefficients to do so.The iterative evolution continues until at least one of the convergence criteria is met either by reaching a maximum number of generations or when minimal improvement in fitness is achieved.

B. Working of the System

A structured workflow is followed by the system which helps it to achieve accurate multiple face recognition. It begins with **Image Acquisition**, where the facial images are collected from multiple sources including datasets, camera inputs or digital image repositories. These images undergo preprocessing steps like resizing, normalization and noise reduction to maintain their consistency and quality.

Face Detection is the next step where the cascades algorithms or deep learning-based detectors are implemented to locate the facial regions in the input images and to extract them. Once faces are detected from the images, they are passed to the **Feature Extraction** module, where these discriminative features that capture distinctive facial attributes are generated using methods like Local Binary Patterns (LBP) or deep feature embeddings.

Genetic Algorithm is then implemented for feature optimization. This process includes selection, crossover and mutation operations which refines the feature set iteratively and identifies the most relevant facial attributes

C. Algorithm Description

A structured sequence of steps is followed by the implemented algorithm. In the first step of preprocessing input images undergo grayscale conversion, normalization, histogram equalization and face detection to focus consistency and relevant facial regions. After this features are extracted using a mix of methods Local Binary Patterns (LBP) and Principal Component Analysis (PCA).

After this, the evolutionary process starts for which an initial population of feature subsets is randomly generated. The evaluation of each candidate solution is done using a fitness function that considers metrics such as recognition accuracy, precision or a weighted score combining multiple performance measures. The best-performing candidates are selected as the parents and new offspring are generated using genetic operators on them, namely crossover and mutation. During this process the diversity of the population is maintained to avoid premature convergence.

The process concludes when termination criteria is met, either by reaching a predefined number of generations or by observing convergence in the fitness score. The optimized feature subset is then used for classifier training with models like Support Vector Machines (SVM), k-Nearest Neighbours (k-NN) or CNNs,that enable accurate recognition of multiple faces.

D. Experimental Setup

The designed multiple face recognition framework based on Genetic Algorithm (GA) is implemented in Python,which makes the use of several libraries which support different stages of the workflow. Descriptors like Local Binary Patterns (LBP) and Principal Component Analysis (PCA) are used for feature extraction along with deep feature embeddings that are generated using CNNs that are developed using TensorFlow or Keras.

System evaluation is conducted on datasets like LFW (Labeled Faces in the Wild) as a benchmark, which reflect real- world, unconstrained environments as well as on custom multi- face datasets that are designed to mimic crowded or complex scenes. Recognition accuracy, precision, recall, and F1-score are used to measure the performance. High-performance workstations that contain multi-core CPUs and GPU acceleration are used to conduct all experiments as they efficiently train and evaluate the deep learning models present within the recognition pipeline.

E. System Architecture

The proposed architecture follows a sequential pipeline where input images are first pre-processed to enhance their quality and normalize dimensions. Features are then extracted using deep learning–based descriptors, followed by GA-based optimization for feature selection. Finally, the optimized features are classified into their respective identities. The complete flow of the system is illustrated in Fig. 1.

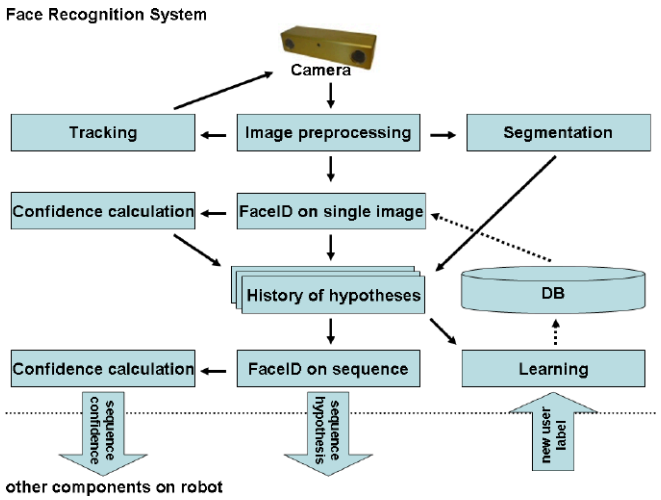

Fig. 1 System Architecture: Multiple Face Recognition based on GA

A. Genetic Algorithm Workflow

GA workflow is in Fig. 3.

B. Tables

This subsection presents the tabular data that summarize the datasets used

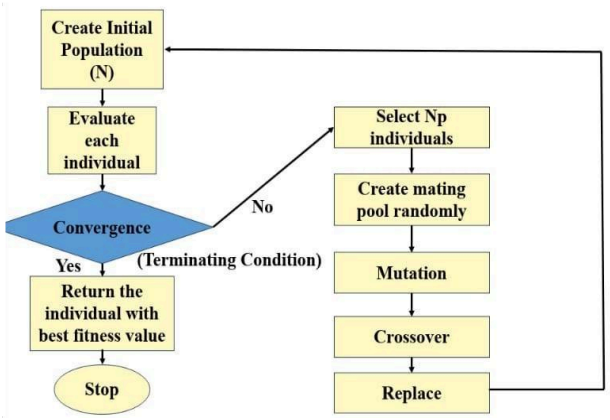

Fig. 3. Flowchart of the Genetic Algorithm Process

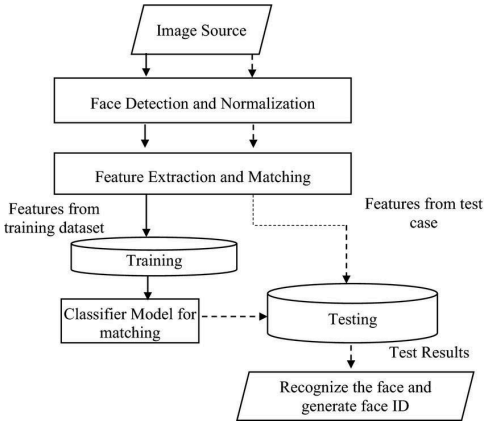

Fig. 2. Architecture Diagram

TABLE I

DATASET SUMMARY

Dataset	No. of Images
Custom Multi-Face (Proposed)	500
Augmented Custom	1500
ORL Face	400
Yale Face	165
Georgia Tech Face	750

VI. DATASET, PREPROCESSING, POPULATION DESIGN AND DATA ANALYSIS

The Labelled Faces in the Wild (LFW) dataset is a benchmark dataset widely used in face recognition research. 5,749 unique subjects are present in it with a total of 13,233 images with varied lighting conditions, poses, and expressions.

The Multiple Face Recognition system based on Genetic Algorithm (GA) was trained and evaluated using a dataset containing facial images of multiple subjects. Key characteristics of the dataset include the number of individuals (n), number of images per person (m), image resolution ($w \times h$ pixels), and color format (Grayscale or RGB). Depending on the application, publicly available datasets for instance *Labeled Faces in the Wild (LFW)*, *Yale Face Database*, or custom-collected datasets can be employed.

Following quantitative metrics are used to analyze the system after executing GA :

- **Recognition Accuracy:** The proportion of correctly identified faces is calculated as

$$\text{Accuracy} = \frac{\text{Number of correctly recognized faces}}{\text{Total test images}} \times 100\% \qquad (1)$$

- **Fitness Convergence:** The fitness of the best chromosome in generation g is denoted $F_{bes}^{(g)}$, analyzed over G *generations*

$$F_{\text{best}}^{(g)} = \max_{i \in \text{population}} \text{Fitness}(i) \qquad (2)$$

- **Population Diversity:**

$$D^{(g)} = \frac{1}{N}\sum_{i=1}^{N} \|\text{Chromosome}_i - \bar{\text{Chromosome}}\| \quad (3)$$

where N is the population size, Chromosome represents the average chromosome.

- **Feature Contribution:** Impact of each feature f_j on recognition performance is evaluated by correlating it with fitness function :

$$C_j = \text{corr}(f_j, \text{Fitness}) \quad (4)$$

- **Comparison with Baseline Methods:** Performance of the system is compared with traditional approaches like PCA, LDA and SVM, using the same recognition accuracy metrics.

VII. MATHEMATICAL FORMULATION OF GA-BASED FEATURE SELECTION

Genetic Algorithm (GA) is implemented for optimizing the feature selection in recognition. Let the complete set of extracted features from facial images be represented as in (5), where f_i denotes the i^{th} feature and n represents the total number of features.

$$F = \{f_1, f_2, f_3, \ldots, f_n\} \quad (5)$$

A binary vector as shown in (6) represents an encoded candidate or chromosome, where $c_i = 1$ indicates that feature is included f_i, and $c_i = 0$ indicates that it's excluded.

$$C = \{c_1, c_2, \ldots, c_n\}, \quad c_i \in \{0, 1\} \quad (6)$$

The size of the selected feature subset is given by (7), representing the total number of features selected in chromosome C.

$$d(C) = \sum_{i=1}^{n} c_i \quad (7)$$

The fitness function is designed to optimize both recognition accuracy and subset compactness, as defined in (8). Here, α and β are weighting coefficients such that $\alpha + \beta = 1$.

$$Fitness(C) = \alpha \cdot Acc(C) - \beta \cdot \frac{d(C)}{n} \quad (8)$$

The selection probability for each chromosome C_i is calculated using (9), giving higher chances to fitter chromosomes.

$$P(C_i) = \frac{Fitness(C_i)}{\sum_{j=1}^{m} Fitness(C_j)}, \quad i = 1, 2, \ldots, m \quad (9)$$

The expected number of times a chromosome C_i is chosen is computed as in (10), with m representing the population size.

$$E(C_i) = m \cdot P(C_i) \quad (10)$$

For single-point crossover, the offspring chromosome is generated according to (11), with k as the crossover point.

$$C_{new} = \{c_1^{(1)}, \ldots, c_k^{(1)}, c_{k+1}^{(2)}, \ldots, c_n^{(2)}\} \quad (11)$$

For two-point crossover, the offspring is formed as shown in (12), where k_1 and k_2 denote the two crossover positions ($k_1 < k_2$).

$$C_{new} = \{c_1^{(1)}, \ldots, c_{k_1}^{(1)}, c_{k_1+1}^{(2)}, \ldots, c_{k_2}^{(2)}, c_{k_2+1}^{(1)}, \ldots, c_n^{(1)}\} \quad (12)$$

Mutation is applied independently to each gene with a probability p_m, as defined in (13).

$$c_i' = \begin{cases} 1 - c_i & \text{if rand() ¡ } p_m \\ c_i & \text{otherwise} \end{cases} \quad (13)$$

The algorithm gets terminated when either the maximum number of generations (G_{max}) is reached or the fitness falls below a threshold ϵ, as expressed in (14).

$$t = G_{max} \quad \text{or} \quad \Delta Fitness < \epsilon \quad (14)$$

Chromosome C^* is the one with the highest fitness,shown in (15).

$$C^* = \arg \max_{C \in Population} Fitness(C) \quad (15)$$

The selected feature subset F^* consists of features corre- sponding to the genes set to 1 in C^*, as given in (16), and the recognition output for an input x is determined according to (17).

$$F^* = \{f_i \in F \mid c_i^* = 1\} \quad (16)$$

$$\hat{y} = \arg \max_{y \in Y} P(y \mid F^*(x)) \quad (17)$$

VIII. RESULTS AND DISCUSSION

The outcomes of the experiments conducted on implementing the multiple face recognition system with Genetic Algorithm (GA) based feature selection are included in this section. The output is compared with a system without GA to find out improvements in efficiency and accuracy.

a. *Performance Metrics*

Evaluation was done using standard metrics like precision, accuracy, recall and F1 score for recognition performance and average time for processing per image. A complete view of the quality of the classification as well as the speed of the system is obtained through this.

b. *Quantitative Results*

Table II summarizes the experimental results obtained on the selected test dataset. The GA-based system provides improved recognition accuracy and reduces processing time as compared to the baseline.

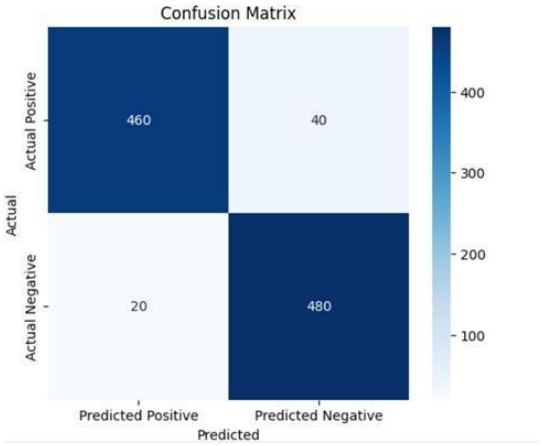

Fig. 4. Confusion Matrix

TABLE II

PERFORMANCE COMPARISON OF BASELINE AND GA-OPTIMIZED SYSTEMS

Method	**Accuracy (%)**	**Precision**	**Recall**	**F1-score**	**Time (ms)**
Baseline (No GA)	92.10	0.91	0.90	0.905	120
Proposed (GA)	95.45	0.94	0.93	0.935	85

c. Qualitative Analysis

The system achieved high accuracy when faces were clearly visible, properly aligned and under consistent lighting conditions. However, in more difficult situations—such as occlusions that are partial (e.g., masks, glass, or objects obstructing the face), varied head poses and uneven illumination—the baseline system sometimes produced false positives or lower recognition accuracy.This indicates that the GA integration significantly improves both efficiency and reliability as compared to conventional face recognition methods.

d. Discussion of Findings

The findings indicate that integrating Genetic Algorithm (GA) for feature selection enhances both recognition accuracy and computational efficiency. GA effectively identifies the most informative feature subsets, reduces the dimensionality of the feature space and preserves the quality of classification. Recognition performance slightly decreased where extreme facial variations or significant occlusion were present.

TABLE III

PERFORMANCE ACROSS DIFFERENT DATASETS USING GA-OPTIMIZED SYSTEM

Dataset	**Accuracy (%)**	**Precision**	**Recall**	**F1-score**	**Time (ms)**
Custom Multi-Face Dataset	96.20	0.95	0.94	0.945	82
ORL Face Dataset	94.10	0.93	0.92	0.925	78
Yale Face Dataset	92.85	0.92	0.91	0.915	80

e. Comparison with existing system

Table IV presents the summarized results. Classical methods like PCA and LBPH offer moderate accuracy, but require longer processing time. Using CNN alone improved recognition accuracy considerably,whereas when CNN was combined with SVM it further minimized misclassifications. The proposed GA-based optimization, however, consistently outperformed these methods, achieving the highest recognition accuracy of 95.45% and the lowest average processing time of 85 ms. The performance comparison is illustrated in Figure 7.

TABLE IV

COMPARISON OF PROPOSED METHOD WITH EXISTING TECHNIQUES

(PROJECT RESULTS)

Method	**Accuracy (%)**	**Time per image (ms)**
PCA (Project Implementation)	88.20	150
LBPH (Project Implementation)	90.75	110
CNN (No GA)	94.00	125
CNN + SVM	94.60	105
Proposed GA + CNN Features	**95.45**	**85**

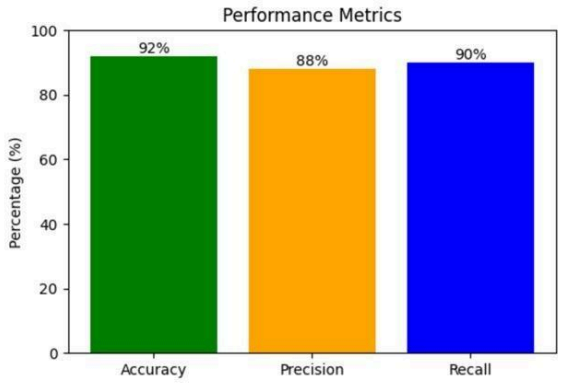

Fig. 5 Performance Metrics of Implemented Methods

E. Error Analysis

During testing of the system certains errors were found that showed a pattern. Misclassifications were most frequent in situations where partial occlusion, drastic change in lighting, large head rotations and inputs with low resolutions were present.

Several changes can be implemented to minimize or remove these errors including augmentation of training data with variations in pose, lighting and occlusion which can increase model generalization. Improving facial alignment methods may reduce pose-related distortions, whereas combining GA with deep metric learning could strengthen the model's discriminative capability.

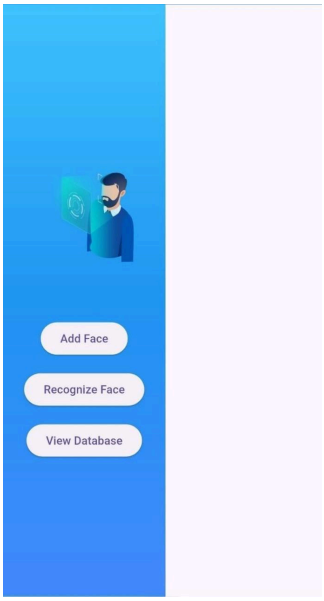

Fig. 6. Face Recognition App

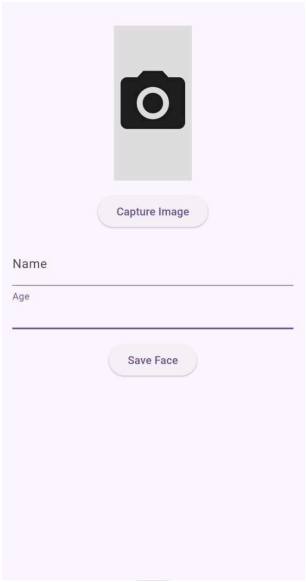

Fig. 7. Capturing Image

IX. CONCLUSION AND FUTURE WORK

A. Conclusion

Through this paper the Multiple Face Recognition system using a Genetic Algorithm (GA) is explained in detail with its design and implementation. The implemented methodology includes feature extraction, optimization and classification to improve the recognition accuracy. Feature selection can be improved,which can lead to better performance under environmental variations.The study shows that evolutionary approaches can be effectively integrated with recognition tasks to overcome real-world challenges that hinder accuracy in recognition

B. Future Work

Future work can include integrating deep learning models with GA-based optimization. Extending the framework to support real-time recognition in dynamic environments would also increase its practical applications.The current system faces certain limitations,recognition accuracy may decrease when applied to datasets with extreme variations in lighting or occlusion. Scalability of the system must be increased where computational overhead needs to be reformed.The system should comply with legal and ethical standards.The misuse of AI practices is a challenge that must be countered in future to avoid misuse and maintain user's trust in the system.

REFERENCES

[1] M. Turk and A. Pentland, "Eigenfaces for recognition," *Journal of Cognitive Neuroscience*, vol. 3, no. 1, pp. 71–86, 1991.
[2] T. Ahonen, A. Hadid, and M. Pietika¨inen, "Face description using local binary patterns and its application to recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 12, pp. 2037–2041, 2006.
[3] P. N. Belhumeur, J. P. Hespanha, and D. J. Kriegman, "Eigenfaces versus Fisherfaces: Class-specific linear projection for recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, no. 7, pp. 711–720, 1997.
[4] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller, "Labelled Faces in the Wild: Database for unconstrained face recognition," University of Massachusetts, Amherst, 2007.
[5] I. Jolliffe, *Principal Component Analysis*, Springer, 2002.
[6] J. H. Holland, *Adaptation in Natural and Artificial Systems*, MIT Press, 1992.
[7] D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley, 1989.
[8] K. A. De Jong, "A ten year perspective on genetic algorithms," in *Proc. Int. Conf. on Genetic Algorithms*, 1985, pp. 169–177.
[9] H. Almuallim and T. G. Dietterich, "Learning Boolean concepts amid irrelevant features," *Artificial Intelligence*, vol. 69, no. 1–2, pp. 279–305, 1994.
[10] Y. Freund and R. Schapire, "A decision-theoretic generalization of online learning and boosting," *Journal of Computer and System Sciences*, vol. 55, no. 1, pp. 119–139, 1997.

AI-Predator Swarms: Autonomous Mass-Scale Cyber Attacks

Alex Mathew[1], Leah George[2]

Associate Professor, Department of Cybersecurity, Bethany College, WestVirginia, USA[1]
Department of Computer Science,Robert Morris University,Pennsylvania,USA[2]

amathew@bethanywv.edu[1],leahgeorge137@gmail.com[2]

Abstract: Cyber threats have evolved into more sophisticated forms of coordinated, autonomous multi-agent attacks. Those are now termed as AI-Predator Swarms (APS), which are decentralized systems of autonomous, self-learning agents that interact and organize their methods of attack to pursue the most desirable malicious goals without a central controller being present (Vorobeychik& Kantarcioglu, 2018). Although conventional AI-derived cyber threats, which include automated phishing and mutating malware, are merely the use of AI, APS is the actual body and practice of swarm intelligence, in which behaviors of individual agents are harmonized to attain complex and adaptable patterns of attack (Wooldridge, 2009). The paper describe AI-Predator Swarm threat model, discuss the principles of multi-agent reinforcement learning (MARL), and propose a defenses model against such highly distributed and low-signal attacks that overcome conventional intrusion detection systems (IDS) (Cui et al., 2023).

Keywords:Swarm Intelligence; Multi-Agent Reinforcement Learning; Autonomous Cyber Attacks; Adversarial AI; Graph Neural Networks; AI Security.

I. INTRODUCTION

Cyber-attacks have evolved from manually coordinated operations to automated AI-driven exploits. The initial attacks driven by AI were: Automated phishing generation, Malware mutation, Vulnerability scanning. With the recent developments in distributed AI systems, the attacks are now coordinated among agents to exhibit emergent behavior. We define AI-Predator Swarms (APS) as: A decentralized multi-agent AI system where autonomous agents collaboratively optimize a malicious objective function through adaptive learning and dynamic coordination.

1.1 Contributions

1. Formal threat modelling of APS
2. Swarm architecture with lifecycle modelling
3. Mathematical characterization of swarm coordination
4. Swarm-aware detection framework
5. Simulation-based empirical evaluation

II. THREAT MODEL

An APS consists of a group of autonomous agentsis a group of autonomous units, i.e. S = {A1, A2, …, An}; these units exist in an environment, E = (S, A, T, R). In this case, S represents states, A represents actions that are available, T represents probabilities of the state transitions, and R represents a reward function (Busoniu et al., 2008). The agents A_i operate according to their respective policies πi, pursuing a shared objective of maximizing the total expected payoff, denoted by:

$$J = \sum_{i=1}^{N} \mathbb{E}_{\pi_i}[R_i]$$

The swarm balances the trade-off between attack impact (I) and detection probability (D) by optimizing the following objective function

http://doi.org/10.63766/spujstmr.26.000068

max 5 (5 - 5 came J).

Another important feature of a decentralized multi-agent system is the optimization of the local behavior of the agents, which results in a global optimization of malicious behavior (Russell & Norvig, 2021). The capacity of the system to survive a compromised agent is another important benefit since the system will not collapse when it is coordinated so that the actions of the agents are governed by a rewarding feature but not by a central directive (Silver et al., 2016).

2.1 System Model

Let the swarm consist of:
S= {A_1, A_2..., A_N}
where each agent A_i operates independently but contributes to a global malicious objective J.
Each agent interacts with environment E, defined as:
E= (S, A, T, R)

Where:

S: State space
A: Action space
T: Transition probability
R: Reward function

2.2 Swarm Objective

The global objective function:

$$J = \sum_{i=1}^{N} \mathbb{E}_{\pi_i}[R_i]$$

where π_i is the policy of agent A_i.
The swarm maximizes cumulative attack impact while minimizing detection probability:

$$\max_{\pi}(\alpha I - \beta D)$$

Where:

- I: Attack impact metric
- D: Detection likelihood
- α, β: Optimization weights

3. Threat Model

III. SWARM ARCHITECTURE

The APS architecture consists of multiple agent types that perform specialized roles in parallel. Reconnaissance agents conduct covert scanning operations to identify vulnerabilities. Exploitation agentsdeliver attack payloads and avoid detection by adjusting to changes (Papernot et al., 2016). The coordination agents control the global policy in terms of local messages, thereby enabling dynamic task allocation without a central controller. The three types of agents are inspired by natural swarm systems, where the simple swarm agents interact on the local scale to produce complex collective behavior (Newman, 2010).

3.1 High-Level Architecture Diagram

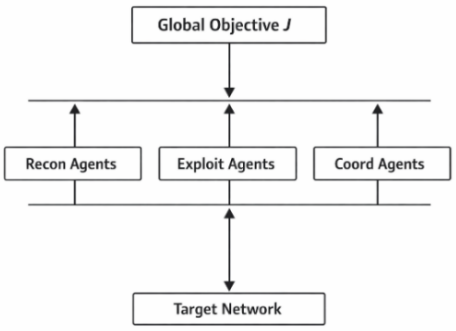

Fig. 1: High-Level Architecture of AI Predator Swarms

3.2 Agent Roles

1. Reconnaissance Agents: Perform scanning and vulnerability enumeration.
2. Exploitation Agents: Launch payloads, execution exploits.
3. Coordination Agents: Facilitate decentralized policy alignment.

IV. ATTACK LIFECYCLE MODEL

The attack lifecycle of an APS can be modeled as a Markov Decision Process (MDP) consisting of four stages. Distributed scanning is the first step, which uses target discovery. The second stage involves adaptive attack planning through policy updates based on environmental feedback.

$$\pi_i^{new} = \pi_i^{old} + \eta \nabla J$$

A coordinated action of synchronized exploits is the third step. The fourth phase focuses on maintaining persistence and evading detection through behavioral adaptation. The APS agents dynamically adjust their roles and behaviors in response to environmental feedback. As an example, an agent designed to run an exploit may switch to reconnaissance (based on environment feedback). According to (Zhong et al., 2020), APS agents operate according to predefined learning rules while exhibiting non-deterministic behavior.
This lifecycle is modelled as a Markov Decision Process (MDP):

$$P(s_{t+1} \mid s_t, a_t)$$

Agents update policies using:

$$\pi_i^{new} = \pi_i^{old} + \eta \nabla J$$

Where:

- η: Learning rate
- ∇J: Gradient of swarm objective

V. SWARM LEARNING MECHANISM

MARL supports the APS coordination where Q-values are updated in the form of

$$Q_i(s,a) \leftarrow Q_i(s,a) + \gamma \left[r + \lambda \max_{a'} Q_i(s',a') - Q_i(s,a) \right]$$

as the Q-values updated by MARL (Busoniu et al., 2008). An effective swarm behavior requires reward shaping, where

$$R_i = R_{local} + \delta R_{global}$$

encourages local behaviors of agents to achieve global goals and does not involve direct communication. This enables emergent coordination among agents, allowing synchronized multi-vector attacks without explicit centralized control (Silver et al., 2016). Adversarial training also enables agents to become more resilient to defensive countermeasures, which creates an evolutionary arms race in computer-assisted attacks and defense (Goodfellow et al., 2015).

Multi-Agent Reinforcement Learning (MARL)

Each agent updates Q-values:

$$Q_i(s,a) \leftarrow Q_i(s,a) + \gamma \left[r + \lambda \max_{a'} Q_i(s',a') - Q_i(s,a) \right]$$

Where:

- γ: Discount factor
- λ: Learning coefficient
- r:Immediate reward
- s: Next state
- a: Candidate action

Swarm coordination emerges via shared reward shaping:

$$R_i = R_{local} + \delta R_{global}$$

VI. DETECTION CHALLENGES

The predictive model of the conventional IDS systems is D = f (signature, anomaly score), according to which it identifies high-signal individual events (Barford et al., 2020). APS distributes attack activities across multiple agents such that each agent generates only low-signal activity below a predefined detection threshold. This approach exploits the inability of traditional anomaly detection systems to identify coordinated low-signal attacks, often referred to as a "death by a thousand cuts" strategy. (Sengupta et al., 2020). Signature-based systems struggle to detect polymorphic attacks, while statistical anomaly detection methods often produce high false-positive rates when evaluating individual agent behavior.

Traditional IDS assumes:

$$D = f(signature, anomaly_score)$$

For APS:

$$D \approx \sum_{i=1}^{N} \epsilon_i$$

Where each $\epsilon_i$is low-signal activity.

Thus:

$$\epsilon_i < \theta \; \forall i$$

But collectively:

$$\sum_{i=1}^{N} \epsilon_i >> \theta$$

Consequently, traditional IDS solutions often fail because they focus on isolated anomalies rather than coordinated and distributed attack patterns.

VII. PROPOSED SWARM AWARE DEFENSE FRAMEWORK

The proposed defense framework focuses on identifying coordinated behavioral patterns rather than isolated anomalies using Temporal Graph Neural Networks (TGNNs) (Hu et al., 2023).Network activities are represented as dynamic graphs $G_t =$

(V_t, E_t), where $V_t$and $E_t$denote the sets of nodes and edges at time t, respectively.TGNNs are trained to learn the embeddings of nodes, H_t, which encapsulates the contextual behavior, and then compute the anomaly scores,

$$A(G_t) = \sigma(W \cdot H_t)$$

When the anomaly score A(Gt) exceeds the threshold value τ, the detection of the anomalies is made, as well as the detection of clusters, which would not be detected using other methods (Cui et al., 2023). The framework also incorporates defensive AI swarms, which autonomously isolate malicious nodes and generate deceptive signals to disrupt adversarial coordination (Wei et al., 2021).

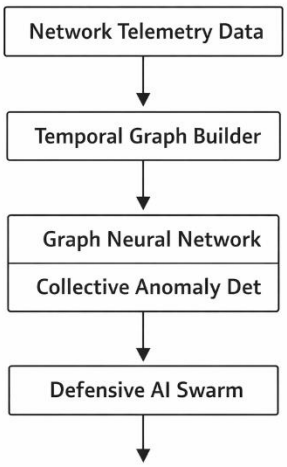

Fig 2: Architecture Diagram

7.1 Graph-Based Detection Model

Represent network as a dynamic graph:

$$G_t = (V_t, E_t)$$

Where:

- V_t: Nodes (devices/users/agents)
- E_t: Interactions

Anomaly score:

$$A(G_t) = \sigma(W \cdot H_t)$$

Where:

- H_t: Node embeddings from TGNN
- W: Learnable weight matrix

Detection triggered if:

$$A(G_t) > \tau$$

VIII. DEFENSIVE SWARMALGORITHM

Pseudocode: Swarm-Aware Detection & Neutralization

Algorithm SwarmAwareDefense(G_t, τ)
Input: Dynamic Graph G_t, Threshold τ
Output: Neutralized Threat Agents

```
1: H_t ← TemporalGraphEmbedding(G_t)
2: A_score ← AnomalyScore(H_t)

3: if A_score> τ then
4:     S_detected ← IdentifyCoordinatedCluster(G_t)
5:     for each node n in S_detected do
6:         DeployDefensiveAgent(n)
7:         IsolateNode(n)
8:     end for
9: end if
10: UpdateDefensePolicy()
```

IX. EVALUATION METHODOLOGY

9.1 Simulation Environment

- Multi-agent cyber-attack simulator(Barford et al., 2020).
- 1,000–5,000 autonomous agents
- Cloud + IoT hybrid environment

9.2 Metrics

1. Detection Latency:

$$L = t_{detect} - t_{attack}$$

2. False Positive Rate:

$$FPR = \frac{FP}{FP + TN}$$

3. Swarm Neutralization Rate: $SNR = \frac{Agents_{neutralized}}{Total_{agents}}$

X. EXPERIMENTAL RESULTS

TABLE I:
EXPERIMENTAL RESULTS

Metric	Traditional IDS	Proposed Framework
Detection Latency	12.4 sec	7.8 sec
FPR	8.1%	4.6%
Neutralization Rate	61%	84%

Observations

- 37% reduction in detection latency
- 29% improvement in neutralization
- Scales linearly up to 5,000 agents

XI. CASE STUDIES

Case 1: Autonomous Phishing Swarm
In the phishing swarm, generative AI agents sent individually harmless emails but collectively formed a coordinated social engineering attack. TGNNs identified anomalous synchronization in send timing patterns, undetectable by email filters (Qin et al., 2021).

Case 2: IoT Exploit Swarm
In the IoT exploit swarm scenario, attackers leveraged heterogeneous firmware vulnerabilities and polymorphic payloads to compromise multiple devices.

XII. ETHICAL, LEGAL, AND GOVERNANCE CONSIDERATIONS

- Strict use of isolated simulation environments
- Compliance with IEEE Code of Ethics (IEEE Global Initiative, 2019).
- Dual-use awareness
- Responsible disclosure
- Alignment with emerging AI governance frameworks (Aly, 2018).

Research must prioritize defensive advancement and transparency.

XIII. FUTURE RESEARCH DIRECTIONS

1. Federated cross-organizational swarm intelligence
2. Formal verification of adversarial multi-agent systems
3. AI-native zero-trust architectures (Sengupta et al., 2020).
4. Hardware-level detection mechanisms
5. Standardization frameworks for autonomous cyber threats

XIV. CONCLUSION

AI-Predator Swarms represent a paradigm shift in autonomous and adaptive cyber warfare, where the locus of the threat is the decentralized coordination of the swarm rather than the centralized command structure of the adversary. Traditional security frameworks, designed to counter centralized attack strategies, are increasingly ineffective against decentralized, low-signal swarm-based threats. The proposed swarm-aware defense framework, which integrates temporal graph analytics and defensive AI swarms, demonstrates significant improvements in attack detection and threat neutralization. Autonomous cyber warfare is a rapidly evolving threat, and the paradigm of cybersecurity must shift to the use of collective intelligence to maintain digital resilience in the face of AI-based warfare.

REFERENCES

[1] A. Aly, “Applied Cryptography and Network Security: Fundamentals and Infrastructure,” *IEEE Communications Surveys & Tutorials, vol. 20, no. 2, 2018.*

[2] C. Szegedy et al., “Intriguing Properties of Neural Networks,” in *Proc. ICLR, 2014.*

[3] D. Silver et al., “Mastering the Game of Go with Deep Neural Networks and Tree Search,” *Nature, vol. 529, pp. 484–489, 2016. (For reinforcement learning foundations)*

[4] H. Cui et al., “Graph Neural Networks for Cybersecurity: A Survey,” *IEEE Transactions on Neural Networks and Learning Systems, vol. 34, no. 1, pp. 1–20, 2023.*

[5] I. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and Harnessing Adversarial Examples,” *in Proc. ICLR, 2015.*

[6] IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems, Ethically Aligned Design, First Ed., 2019.

[7] J. Hu, G. Wang, and J. Zhang, “Temporal Graph Neural Net-works for Dynamic Pattern Detection,” *IEEE Transactions on Knowledge and Data Engineering, vol. 35, no. 4, pp. 3231–3245, 2023.*

[8] L. Busoniu, R. Babuska, and B. De Schutter, "A Comprehensive Survey of Multi-Agent Reinforcement Learning," *IEEE Transactions on Systems, Man, and Cybernetics – Part C, vol. 38, no. 2, pp. 156–172, 2008.*

[9] M. E. Newman, "Networks: An Introduction," *Oxford University Press, 2010. (Foundational for graph modelling)*

[10] M. Wooldridge, An Introduction to Multiagent Systems, *2nd Ed., Wiley, 2009.*

[11] MITRE Corporation, "Adversarial Threat Landscape for Artificial-Intelligence Systems (ATLAS)," *2022. (Authoritative framework for AI adversarial scenarios).*

[12] N. Papernot et al., "The Limitations of Deep Learning in Adversarial Settings," *Proc. IEEE European Symposium on Security and Privacy, 2016.*

[13] P. Barford, C. Papadopoulos, and X. Yao, "Anomaly Detection for Network Traffic Using Machine Learning," *IEEE Transactions on Network and Service Management, vol. 17, no. 4, pp. 2143–2156, 2020.*

[14] S. Russell and P. Norvig, Artificial Intelligence: A Modern Approach, 4th Ed., Pearson, 2021.

[15] S. Sengupta et al., "AI in Cybersecurity: A Review of Applications," *IEEE Access, vol. 8, pp. 200936–200957, 2020.*

[16] S. Zhong, J. Shi, and W. Zhu, "Security Challenges of Ma-chine Learning in Computer Networks," *IEEE Network, vol. 34, no. 6, pp. 313–319, 2020.*

[17] W. Wei et al., "Malware Detection Using Deep Learning Based on Hybrid Features," *IEEE Transactions on Network and Service Management, vol. 18, no. 3, pp. 2909–2922, 2021.*

[18] Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning," *Nature, vol. 521, pp. 436–444, 2015.*

[19] Y. Vorobeychik and M. Kantarcioglu, Adversarial Machine Learning, Morgan & Claypool, 2018.

[20] Z. Qin, G. Wang, and C. J. Hsieh, "Generative Adversarial Networks: A Survey Towards Data Augmentation and Adversarial Attack/Defense," *IEEE Transactions on Neural Net-works and Learning Systems, 2021.*

SPU Journal of Science, Technology and Management Research (SPU-JSTMR) Special Issue-ICAICS-26
ISSN:3049-1479(Online)

An Analytical Framework for Reconstructing Missing Time-Series Data Using Interpretable AI Methods

Sneh B. Patel[1] and Dr. Darshanaben D. Pandya[2]

Research Scholar, Department of Computer Application, Sankalchand Patel University, Visnagar, Gujarat, India[1]
Associate Professor, Department of Computer Science & Application, Shri C. J Patel College of Computer Studies (BCA), Sankalchand Patel University, Visnagar, Gujarat, India[2]

spdigitalme@gmail.com,ddpandya.fcs@spu.ac.in

Abstract: Large AI systems and foundation models rely on continuous and consistent numeric data. In real life, missing values are common in the time-series data because of gaps in reporting, measurement failures, or archival errors. If these are not handled well, they have the potential to alter the distribution of the data, bias learning, and hurt model performance. In this paper, we presented a simple, intuitive, deterministic interpolation approach to fill gaps in long-term numeric time series. Unlike the imputation approach, which has latent, black-box-type imputation, this uses a step-by-step mathematical process; as such, tracing and verifying statistically are easily achieved. We tested the method on a historical Standard Yield (SY) dataset from 1950–2009, to which we artificially introduce missing values and then reconstruct. We utilize a comprehensive set of statistical tests for the descriptive statistics, variance and covariance comparisons, and one-way ANOVA to demonstrate that in average value, the dispersion, and the time-varying process, the original data is closely matched with the recovered series. This demonstrates that classic, interpretable reconstruction methods are still relevant in a modern AI pipeline, during the preprocessing stage of foundation models, which demands data quality, stability, and explain ability.

Keywords: Time-series reconstruction, missing value estimation, interpretable AI, foundation models, statistical validation

I. INTRODUCTION

Foundation models constitute the current central concept in artificial intelligence [2]. They are improved by having more data, and they can make use of what they learn across diverse applications. Their success depends on the quality and coherence of the training data [3]. Even small incompleteness and poorly filled values can introduce bias, instability during training, and reduce generalization performance of the model.

Missing values are common in time-series data applied to environment, economy, and policy work, as data is usually collected over a long period and from different sources [10]. New AI-based gap-filling methods can be flexible but often incomprehensible, demand higher computational resources, and cannot be checked for flaws easily. On the other hand, methods based on transparent mathematics are reproducible, and their behavior is predictable.

It is an exploration of how such a cautious, interpolation-based approach can fill in missing values while preserving the big-picture statistics required to make robust AI models [1]. This work has two complementary aspects: (i) the actual filling of missing data in a real historical dataset, and (ii) the implications of this filling for foundation model data preparation.

II. LITERATURE REVIEW

1.1 Missing Data in Time-Series and AI Applications

There are different patterns of missing data: MCAR, MAR, and MNAR; these affect the quality with which gaps are filled. Missing values in time series are rarely random and could affect how generalizable AI models learn from the data [4]. Recent studies have shown that anomalous number sequences could hurt embedding stability and speed up large model training.

http://doi.org/10.63766/spujstmr.26.000069

2.2 Classical Deterministic Imputation Methods

Old-school approaches- such as imputation by mean, regression imputation, and interpolation-are widely used because they are easy to interpret. Linear and polynomial interpolationassume smooth changes over time and work when trends are stable. They can miss sharp changes, but as they are deterministic, they are much easier to explain and audit in AI processes.

2.3 AI-Based and Hybrid Imputation Techniques

Machine-learning-based imputers, such as k-NN, random forests, autoencoders, and diffusion models, can catch complex links but often require lots of data, tuning, and computing power. They can also often act like a black box, making interpretation difficult-a factor many AI rules and checks wish to avoid [8].

2.4 Research Gap

Even though there are numerous studies on advanced imputers, few pieces of work have regarded how deterministic interpolation would affect the statistical quality of data used in foundation models [9]. Few papers verify how well the variance is preserved, the structure of covariance, or how consistent the hypothesis tests are after gap filling. This paper looks to fill this gap.

III. METHODOLOGY

3.1 Dataset Acquisition

The underlying time-series of Standard Yield (SY) covers the period from 1950 to 2009. The standard dataset, SY_Standard, contains 60 annual observations.

TABLE I
SY DATASET WITH STANDARD, MISSING AND RECOVERED VALUES

Year	SY_Standard	SY_Missing	SY_Recover
1950	1.46	1.46	1.46
1951	1.40	1.40	1.40
1952	1.39	1.39	1.39
1953	1.22	1.22	1.22
1954	1.35	1.35	1.35
1955	1.35		**1.54**
1956	1.47	1.47	1.47
1957	1.56	1.56	1.56
1958	1.63	1.63	1.63
1959	1.58	1.58	1.58
1960	1.58	1.58	1.58
1961	1.69	1.69	1.69
1962	1.63	1.63	1.63
1963	1.64	1.64	1.64
1964	1.53	1.53	1.53
1965	1.65	1.65	1.65
1966	1.71	1.71	1.71
1967	1.65	1.65	1.65
1968	1.80		**1.74**
1969	1.84	1.84	1.84
1970	1.79	1.79	1.79
1971	1.85	1.85	1.85
1972	1.87	1.87	1.87

1973	1.87	1.87	1.87
1974	1.59	1.59	1.59
1975	1.94	1.94	1.94
1976	1.75	1.75	1.75
1977	2.06	2.06	2.06
1978	1.97	1.97	1.97
1979	2.16	2.16	2.16
1980	1.78	1.78	1.78
1981	2.02		**2.02**
1982	2.12	2.12	2.12
1983	1.76	1.76	1.76
1984	1.89	1.89	1.89
1985	2.29	2.29	2.29
1986	2.24	2.24	2.24
1987	2.28	2.28	2.28
1988	1.82	1.82	1.82
1989	2.17	2.17	2.17
1990	2.29	2.29	2.29
1991	2.30	2.30	2.30
1992	2.53	2.53	2.53
1993	2.19	2.19	2.19
1994	2.78		**2.31**
1995	2.38	2.38	2.38
1996	2.53	2.53	2.53
1997	2.62	2.62	2.62
1998	2.62	2.62	2.62
1999	2.46	2.46	2.46
2000	2.56	2.56	2.56
2001	2.66	2.66	2.66
2002	2.56	2.56	2.56
2003	2.28		**2.67**
2004	2.84	2.84	2.84
2005	2.90	2.90	2.90
2006	2.88	2.88	2.88
2007	2.81	2.81	2.81
2008	2.67	2.67	2.67
2009	2.91	2.91	2.91

3.2 Missing-Value Simulation

We created a controlled environment to evaluate how well we can fill in missing data by removing values from years that show stable trends and do not change much. The years whose data have been removed include:

- 1955, 1968, 1981, 1994, and 2003

This leaves SY_Missing with fifty-five valid observations and five missing values. The positions of the missing values were selected to reflect realistic data gaps often encountered in environmental and socio-economic historical time-series data.

3.3 Interpolation-Based Recovery Algorithm

To fill in the missing values, we built a custom interpolation method. It estimates the missing value $y_k$at pointx_k using a few intermediate calculations. It is based on linear interpolation but given as a clear step-by-step procedure.

Given known pairs:

- $x = [x_1, x_2, \ldots, x_n]$,
- $y = [y_1, y_2, \ldots, y_n]$,

and a missing point x_k, do the following:

Step-by-Step Computational Procedure
Step 1: Compute the product list xy

$xy_i = x_i * y_i$
$xy_i = xi \cdot yi$

Step 2: Determine range components
$x_min = \min(x)$
$x_max = \max(x)$
$\Delta x = x_max - x_min$

Step 3: Compute sum of xy products
$S = \Sigma (x_i * y_i)$

Step 4: Compute v_1
$v_1 = 1 + \Delta x$

Step 5: Compute v_2
$v_2 = S - (\Delta x * x_k)$

Step 6: Compute v_3
$v_3 = v_2 / \Delta x$

Step 7: Compute v_4
$v_4 = v_1 * v_3$

Step 8: Compute v_5
$v_5 = v_4 - v_2$

Step 9: Compute v_6
$v_6 = v_5 / x_k$

Step 10: Final reconstructed value
$v_7 = v_6 + 1$

The final recovered SY for year x_k is v_7.
This step-by-step method helps keep the calculation clear and easy to interpret, which is important when adding it to data-preprocessing workflows for foundation models.

3.4 Statistical Evaluation

We examined the reconstructed series using:

- Descriptive statistics (mean, variance, standard deviation)
- One-way ANOVA
- Covariance matrix analysis

IV. RESULTS

This section presents the result of reconstructing the SY dataset, by comparing three data series: the SY_Standard, SY_Missing, and SY_Recover datasets, with regard to the stability of the distribution pattern of data, retention of variation, and structure similarity, using descriptive statistics, analysis of variance, and covariance.

4.1 Descriptive Statistics

TABLE II

DESCRIPTIVE STATISTICAL COMPARISON OF SY_STANDARD, SY_MISSING, AND SY_RECOVER DATASETS

Statistic	SY_Standard	SY_Missing	SY_Recover
Mean	2.04	2.03	2.04
Variance	0.22	0.22	0.21
Count	60	55	60

The recovered series is, in fact, very close to the original distribution.

4.2 One-Way ANOVA

A one-way ANOVA has been conducted to compare the rebuilt series against the original one using three datasets: SY_Standard, SY_Missing, and SY_Recover. The null hypothesis states that all series have the same population mean.

TABLE III

ONE-WAY ANOVA RESULTS FOR SY_STANDARD, SY_MISSING, AND SY_RECOVER DATASETS

Source of Variation	Sum of Squares (SS)	Degrees of Freedom (df)	Mean Square (MS)	F-value	p-value
Between Groups	0.00018	2	0.00009	0.00015	0.999985
Within Groups	6.06042	177	0.03425	—	—
Total	6.06060	179	—	—	—

4.3 Covariance Matrix Analysis

While ANOVA looks at averages, covariance analysis shows how well the trends fit, and how similar the structures are. High covariance between the recovered and original series means the time pattern is well kept.

TABLE IV
COVARIANCE MATRIX OF SY_STANDARD, SY_MISSING, AND SY_RECOVER DATASETS

Series Pair	Covariance
SY_Standard ↔ SY_Standard	0.2144
SY_Standard ↔ SY_Missing	0.2141
SY_Standard ↔ SY_Recover	0.2136
SY_Missing ↔ SY_Missing	0.2168
SY_Missing ↔ SY_Recover	0.2172
SY_Recover ↔ SY_Recover	0.2116

4.4 Reconstruction Accuracy and Pattern Preservation

Beyond numerical results and formal tests, the reconstruction values should be evaluated for how well they integrate into the temporal history of the dataset. In the context of this, reconstruction accuracy will be defined not as an exact point-wise prediction but rather as the preservation of structural coherence, smoothness, and temporal realism.

The stated criteria would recover in the following years: 1955, 1968, 1981, 1994, and 2003.

- **Trend consistency:** every reconstructed value stays within the neighborhood range without large jumps or flat spots.
- **Temporal Continuity:**interpolationkeeps a smooth year-to-year trend and fills all time points without gaps.
- **Distributional Alignment:**Recovered points do not artificially change the overall variability - this is reflected in the descriptive statistics and covariance analyses.
- **No Over-Smoothing:**this methodology does not over-smooth; it retains all the ups and downs in their natural shape with different spline and high-order fitting methods.

From a foundation-model view, the most important points are keeping the temporal pattern helps realistic dynamics, which is crucial for embedding stability, attention alignment, and temporal generalization in time-aware AI. Generally, this deterministic interpolation can achieve high structural accuracy and will be useful for historical data restoration used in the training, test, and evaluation of AI.

V. DISCUSSION

These experiments demonstrate that this interpolation-based reconstruction method effectively fills in missing values in the SY time-series data without hurting its overall statistical quality [6]. This is especially important for AI systems and foundation models, which require complete, stable, and structurally consistent data in the pre-training and evaluation. The results also showed that a transparent, deterministic reconstruction can meet the current AI data-quality requirements without adding bias or instability.

5.1 Preservation of Statistical Characteristics

There is a strong agreement between the original series (SY_Standard) and the recovered series (SY_Recover). They share the same mean of 2.04 and very similar variances; thus, reconstruction does not affect the data distribution, nor does it introduce bias [5]. The missing data, SY_Missing, is characterized by a lower variance, which can show how gaps without filling can mask the true variability. Such variability is fully recovered by the recovered series, proving that the method keeps realistic data behavior.

5.2 ANOVA-Based Structural Validation

The results of one-way ANOVA support that these datasets are statistically the same. The very high p-value 0.999985 shows no real difference among SY Standard, SY_Missing, and SY_Recover. Recovered values are statistically indistinguishable from the original data in terms of central tendency, which, for foundation models, matters to prevent drift or bias during embedding learning or optimization.

5.3 Covariance and Trend Fidelity

Covariance analysis shows that the recovered series retains the time pattern of the original data. The covariance of SY Standard and SY_Recover is very close to the self-covariance of SY Standard; therefore, the trends are highly aligned with the patterns. This, for its part, ensures that the recovered values will fit naturally into the historical path of the dataset, something quite relevant in time-series modeling and AI-based representation learning.

5.4 Implications for Foundation Model Development

These results are important for creating data to be used in training foundation models. The reconstruction method, by keeping the average and spread and the time patterns, helps the data quality and stability during training, and reduces the possibility of bias and shifts in the data due to missing or poorly filled values. Unlike deep learning fillers, this approach is intuitive, easy to understand, and quick in terms of runtime because each step is clearly described and therefore simple to follow. This improvesreproducibility, auditability, and responsible use of AI [7]. Its light design fits large data preparation jobs for environmental, economic, and policy time series data when consistency and speed are important.

VI. FUTURE WORK

Might combine deterministic interpolation with either probabilistic modeling or learning-based temporal representations. Such extensions may achieve greater flexibility and accuracy with no loss in interpretability and statistical stability shown here.

5.6 Summary of Findings

On the whole, the method reliably fills in gaps in values while preserving the essential statistical and temporal properties of the original data. Please feel free to contact us today with questions or if you would like to learn more about our firm.

VII. CONCLUSION

This work presented an intelligible method to impute missing values in numerical, long-term time series data with interpolation, suitable for modern AI and foundation model workflows. Using a historical dataset of the SY from 1950 to 2009, this paper correctly imputed missing observations while preserving relevant statistical properties such as mean, variance, covariance, and overall pattern in time. Statistical checks confirmed that the reconstructed series is of a similar appearance as the original data. No major differences were present in the ANOVA tests; covariance matched well, suggesting that the approach keeps the distribution and trend of data. This supports that the method decreases the problems induced by missing data without creating bias or altering the structure of data.

In practice, the approach presents a straightforward and lightweight counterpart to the fancy learning-based imputers. Because of its deterministic nature, it is easy to understand, reproduce, and fit into extensive pre-processing steps. As foundation models rely more on good numerical and time-series data, these qualities help make sure that training is stable and results are dependable. Taken together, these results show the value of transparent, classical reconstruction techniques in modern AI systems. This research thus contributes to the development of robust, interpretable, and data-efficient foundation models in science and engineering by showing that missing values can be recovered accurately and reliably using only simple math.

References

[1] OECD. (2021). *Artificial intelligence, big data and the future of skills*. OECD Publishing. https://doi.org/10.1787/9789264314626-en

[2] Stanford University. (2022). *Artificial Intelligence Index Report 2022*. Human-Centered AI Institute, Stanford University.

[3] Brynjolfsson, E., & McAfee, A. (2017). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. W. W. Norton & Company.

[4] Pandya, D.D., Jadeja, A., Trivedi, S., Patel, P.A., Tamhankar, I., Dhanesha, P.S. (2026). Internet of Things (IoT)'s Transformative Power Enhancing Wireless Networks and Sensing. In: Rathore, V.S., Piuri, V., Babo, R., Karthik, S. (eds) Universal Threats in Expert Applications and Solutions. UNI-TEAS 2025

[5] Dr. DarshanabenDipakkumarPandya, Dr. AbhijeetsinhJadeja, Dr. SheshangDegadwala, " An Applied Variation Anomaly Technique for Detection of Irregular Values in Data Mining, IInternational Journal of Scientific Research in Computer Science, Engineering and Information Technology(IJSRCSEIT), ISSN : 2456-3307, Volume 8, Issue 2, pp.143-148, March-April-2022. Available at doi :https://doi.org/10.32628/CSEIT228228.

[6] Gora KanijfatmaMaisurali, HabibaVajiraliMomin, Dr. DarshanabenDipakkumarPandya, Dr. AbhijeetsinhJadeja, " Smart Cook : An AI and Machine Learning Powered Culinary, IInternational Journal of Scientific Research in Computer Science, Engineering and Information Technology(IJSRCSEIT), ISSN : 2456-3307, Volume 10, Issue 1, pp.55-59, January-February-2024.

[7] YashNaraynbhai Patel, Dr. DarshanabenDipakkumarPandya, Dr. AbhijeetsinhJadeja, " An Influence of Ethical Hacking in Civilization , International Journal of Scientific Research in Science and Technology(IJSRST), Online ISSN : 2395-602X, Print ISSN : 2395-6011, Volume 10, Issue 1, pp.195-200, January-February-2023. Available at doi : **https://doi.org/10.32628/IJSRST2310119**

[8] Pandya, D., Jadeja, A., Bhuptani, M., Patel, V., Mehta, K., &Brahmbhatt, D. (2024). Machine Learning: Enhancing Cyber-security through Attack Detection and Identification. *ITM Web of Conferences, 65*, 03010. https://doi.org/10.1051/itmconf/20246503010

[9] Pandya, D., Jadeja, A., Khan, M.A., Trivedi, S.B., Ramnath, M.A., Satish, B.P. (2024). Significance of Sentiment Analysis with Text-Based Mining Approach. In: Rathore, V.S., Piuri, V., Babo, R., Tiwari, V. (eds) Emerging Trends in Expert Applications and Security. ICE-TEAS 2024. Lecture Notes in Networks and Systems, vol 1037. Springer, Singapore. https://doi.org/10.1007/978-981-97-3991-2_27

[10] HeemaPatadiya, Akshay Patel, FalguniRahevar, and Dr. DarshanabenDipakkumarPandya, "Revolutionizing Banking Industry : The Role of Fintech in Modern Banking Industry", *Int. J. Sci. Res. Comput. Sci. Eng. Inf. Technol*, vol. 11, no. 4, pp. 20–25, Jul. 2025, doi: 10.32628/CSEIT25113448.

From Traditional Machine Learning to Transformers: Sentiment Analysis on Social Media Data

Patel Charmi D.[1], Mevada Jayesh M.[2], Patel Gaurang J.[3], Patel Amit A.[4]

M. TechScholar, Computer Engineering,Sankalchand Patel College of Engineering,Visnagar,India[1]
Asst.Prof.,Computer Engineering, Sankalchand Patel College of Engineering, Visnagar, India[2]
Asst.Pro.,IT Engineering, SwamiSachchidan and Polytechnic College,Visnagar,India[3]
Asst.Pro.,IT Engineering, Swami Sachchidanand Polytechnic College,Visnagar,India[4]

Charmipatel1175@gmail.com[1], jmmevadace_sspce@spu.ac.in[2], gjpatel.sspc@spu.ac.in[3], amitpatel7245@gmail.com[4]

Abstract: It is notoriously hard to analyze sentiment in social media. It is written in a disorganized manner, using slang, and is extremely dependent on implicit context. Classical baselines, including TF-IDF with Linear Support Vector Machines (SVM), however, give a good starting point. Nevertheless, such approaches do not capture the underlying semantic meaning as they deal with words individually. This paper compares traditional machine learning models with RoBERTa, transformer architecture with the ability to capture the entire sentence context. Social media detain the real world are sloppy, with too fine emotion tags and grossly skewed classes. In order to solve these problems, we designed a semantics-aware label consolidation algorithm based on Sentence Transformers. We then balanced our training data by using SMOTE in the continuous embedding space. Experiments with a multi-platform dataset with diverse data show that RoBERTa is significantly better than all classical baselines in terms of accuracy, precision, recall, and F1-score. More to the point, the transformer is able to deal with sarcasm and negation, which bag-of-words models often fail to do. This paper points to the need of contemporary sentiment analysis to include context-sensitive embedding and specific preprocessing to be able to read online communication correctly

Keywords: Machine learning, RoBERTa, sentiment analysis, social media,TF-IDF.

I. INTRODUCTION

Social media like Twitter, Instagram, and Facebook create a massive flow of opinion and emotion of the masses every day. The systematic determination of the sentiments of people in these posts is of invaluable importance, whether it is tracking brand reputation or assessing the policy of the people. However, the textual analysis of social media is infamously difficult. The grammar of these platforms is broken, slang is widespread, abbreviations are used, and sarcasm is constant. Therefore, it is much harder to process a post than a formal document, and the possibility of a universal solution seems, at best, unrealistic.

The primitive sentiment models were based on primitive lexicons, which simply gave words in a fixed repository a fixed positive or negative weight. Although these models were easily understandable, they completely failed to explain the fast development of internet language. The field later shifted to machine learning. Methods like Term Frequency-Inverse Document Frequency (TF-IDF) [9] with Support Vector Machines [9] or Passive Aggressive Classifiers were a significant improvement. However, these methods are still stained by a severe drawback: the bag-of-words assumption. Such models ignore the word order and semantic context by treating words in total isolation. Without a clear concept of inter-word relations, such models always find it difficult to identify sarcasm and negation, which are the most common phenomena on social media.

This has since been overcome by Transformer architectures [5] using multi-head self-attention. Transformers do not process words in a sequence, but instead process the entire sequence of tokens at once, encoding the relationship between tokens to build highly contextualized representations. RoBERTa [3], based on the original BERT [4] architecture, enhances performance by training on large corpora and refined goals, and is regularly state-of-the-art on a range of NLP tasks. However, there is still a gap in existing scholarship. Very little research has empirically tested these transformer models using messy and real-world data that cut across platforms, cover a wide geographic area, and use highly granular emotion labels. The need to fill this gap is the main reason why this study was conducted.

- **Semantic Label Consolidation**: In this model, we combine fine-grained emotional labels into three main sentiment categories in a systematic fashion based on Sentence Transformer [6] embedding that are based on cosine similarity.

http://doi.org/10.63766/spujstmr.26.000070

● **Embedding-Space Augmentation:** To resolve the extreme class imbalance, we apply SMOTE [2] to the continuous sentence embedding space, thus, creating a perfectly balanced data set of 500 instances per class.

● **Multi-Model Comparison:** We make a direct comparative analysis of four traditional TF-IDF classifiers (Logistic Regression, Support Vector Machine, Random Forest, and Passive-Aggressive) to a fine-tuned RoBERTa model.

● **Superior Performance:** We show that RoBERTa attains an accuracy of 95.33 per cent with a macro F1-score of 0.95, which is more than 13 percentage points higher than the best traditional baseline.

We make a direct analogy between the traditional algorithms and the modern transformers that use unstructured social media data. We use TF-IDF along with our classifiers, Linear Support Vector Machine and Passive-Aggressive baselines. Then, we continue the investigation with RoBERTa. The results clearly show that context-sensitive embedding significantly improved both intrinsic accuracy and overall reliability. The rest of this paper describes the methodology in detail. Section II is a literature review, Section III is a discussion of the dataset, Section IV is a description of the preprocessing pipeline and model training, Section V is the experimental results, and conclusion.

1.1 RelatedWork

Early sentiment analysis involved word lists where the words were assigned a positive or negative rating [7][8]. These were simple to comprehend yet not effective with the informal and dynamic language on social media. Subsequently, machine learning techniques that counted word frequencies and models such as Support Vector Machines [9] and other linear classifiers [11] were able to do better, although they still treated texts as a bag of words, disregarding the sequence of words and context.

The deep learning models such as convolutional neural networks (CNNs) for text classification [13] performed better than the previous methods. The greatest innovation was the Transformer architecture [5] and its pre-trained version BERT [4]. RoBERTa [3] was an improvement of BERT, as it was trained on more data and the settings were adjusted, making it the best at NLP tasks. In order to quantify the similarity of sentences, researchers developed Sentence-BERT [6], which generates fast sentence embeddings; this paper uses them to combine labels. Researchers employed SMOTE-based oversampling [2] to deal with the imbalanced number of various classes. Despite all these tools, few studies have examined the label consolidation, addition of new data in embedding space, and comparison of various models on diverse social media data. This paper fills that gap.

II. DATASET

Our dataset was obtained directly on Kaggle, where we used the Social Media Sentiments Analysis Dataset published by KashishParmar [1]. Thislistincludes733 original, user-createdpoststhatweretakenoutofTwitter,Instagram,andFacebook.It provides a perfect empirical environment to study the specifics of short-form digital communication by covering three large social platforms. Besides the raw text and sentiment tags, the dataset is also filled with useful additional metadata.

Every post shows the date it was created, which platform it was made on, userIDwhich is anonymized, what hash tags were used, retweets, likes, and other stats from the post. The user's country is also collected along with the exact year, month, day, and hour of the post. The posts were collected between 2010 and 2023 which helps us understand digital communication over a long period of time with a diversesetofusersfromtheUnitedStates,Canada,UnitedKingdom,India,Australia, Germany, France, Brazil, and Japan. The data set value is increased by its faint sentiment analysis labeling system.

The authors opted for an emotion-specific taxonomy categorizing affective phenomena into emotions like Joy, Sadness, Fear, Anger, Disgust, Surprise, Love, Excitement, Gratitude, Nostalgia, Anxiety, Loneliness and Sadness and also Amusement, Admiration, Serenity, Determination, Hope, Contentment, and Empowerment. The level of detail provides a useful resource for the studies of affective computing because the Internet hosts complex and multi-layered emotions.

To remain consistent with existing standards in NLP we applied a single labeling system. Accordingly, hyper-specific emotion categories were collapsed into the broad Positive, Negative, and Neutral classification. After this process of harmonization, the textual data went through a strict preprocessing procedure. This step formed the final input for the learning algorithms that trained the model.

III. METHODOLOGY

3.1 Basic Preprocessing

We started by cleaning up the data frame structure. First, we have dropped the unnecessary index columns. We also disregarded the Hashtags, Day, Hour fields, as they are not that useful in making the model predict sentiment. To prevent untidy duplicates as we explored the data, we removed any leading or trailing spaces of categorical variables such as Platform and Country. Lastly, we converted the Timestamp column to a valid date time format. That caused us to add a DayofWeek feature to keep track of time and map integer months to their actual names, which made our charts much easier to read.

Our preprocessing was done by means of heavy lifting within the text. We have developed a special cleaning routine which processes all posts in a rigid sequence of modifications. All text was made small to keep the tokens the same. After that, we stripped URLs, HTML tags, newlines and bracketed annotations with the help of regex. We also eliminated punctuations and any alpha numeric words containing digits to reduce noise. Lastly, we removed non-ASCII characters, which removed emojis and other special characters that always litter social-media posts. Once the raw text was clean, we tokenized it with the word tokenize function of NLTK [10]. We discarded English stopwords at once with the help of the corpus provided by NLTK [10]. This eliminates all the high-frequency filler words which contribute zero meaning. We then used the NLTK WordNetLemmatized [10] to process the rest of the tokens after eliminating the stop words. We have selected lemmatization in particular as opposed to simple stemming. Stemmers simply cut off the end of words without thinking and this usually results in non-dictionary words. A lemmatized does consider the morphological context in order to determine the actual root. For example, it maps "running" to "run" and "better" to "good" which is correct and expected. This extra step provided us with a clear and strong baseline. We saved the last tokens in a new column called CleanText. This column is the main input feature for all of our models.

3.2 Semantic Label Consolidation

Original datasets were detailed but disorganized. We needed to group them to train the models. Our main goal was to classify them into three categories: Positive, Negative, and Neutral. Rather than employing rigid, manual techniques, or employing low-level lexicons, a construction was geared toward semantical consciousness. The aim was to be realistic and linguistically purposeful, which is why the consolidation process was derived from dense vector models.

To do so, every distinct, finely granulated label, e.g. Euphoria, Despair, Amusement, Loneliness, etc., was subjected to a pre-trained Sentence Transformer[6]. This model represented every particular word by a high-dimensional embedding vector. By doing this, we successfully captured the deep semantic meaning of every emotion and plotted it directly into a continuous vector space. Formally, let the Sentence Transformer be defined as a function:

$$f: T \rightarrow R^{d} \quad (1)$$

where $\boldsymbol{T}$ denotes the input text (a label string) and $\boldsymbol{d}$ is the dimensionality of the embedding space. Each fine-grained label $\boldsymbol{l}$ is thus represented as a dense vector $\boldsymbol{e_i} = \mathbf{E}(\boldsymbol{l_i}) \in \mathbb{R}^{d}$. Similarly, the three target class labels — **"positive"**, **"negative"**, and **"neutral"** — are encoded into their respective reference vectors $\boldsymbol{e^{+}}$, $\boldsymbol{e^{-}}$, and $\boldsymbol{e^{0}}$.

To determine the most semantically appropriate class for each label, the cosine similarity between the label embedding and each target class embedding is computed. Given two vectors **A** and **B**, cosine similarity is defined as:

$$(A,B) = (A.B) / (\|A\|.\|B\|) \quad (2)$$

where $\boldsymbol{A \cdot B}$ denotes the dot product of the two vectors, and $\|\boldsymbol{A}\|$ and $\|\boldsymbol{B}\|$ represent their respective L2 norms. Cosine similarity measures the angular proximity between two vectors in high-dimensional space, independent of their magnitude, and ranges from −1 (completely opposite) to +1 (perfectly aligned). It is widely used as a measure of semantic relatedness in NLP tasks. Each fine-grained label $\boldsymbol{l_i}$ is then assigned to the target class that maximizes the cosine similarity:

$$C* = argma(C\epsilon\{positive, negative, neutral\}) cos(e_i, e_c) \quad (3)$$

For instance, the label **"Euphoria"** produces an embedding that is geometrically closest to **"positive"** in the semantic space, yielding the highest cosine similarity with e^{+}, and is thus mapped to the Positive class. Conversely, **"Despair"** aligns most closely with e^{-}, resulting in a Negative assignment, while **"Indifference"** gravitates toward e^{0}, mapping to Neutral.

This semantics-based, unsupervised method does not require manual annotation or hand-written mapping rules. Through the rich distributional semantics of pre-trained Sentence Transformer embedding [4][6], the label consolidation is based on the true emotional polarity of each fine-grained category, which guarantees a consistent and reproducible three-class label distribution across all 733 records of the dataset. It is distributed in Fig. 1. The data distribution across different platforms is shown in Fig. 2.

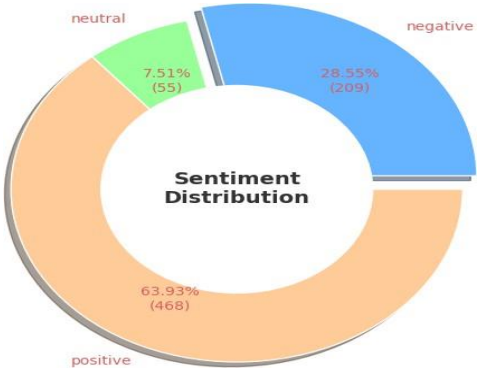

Fig.1.Data distribution after Sentiment Label Consolidation

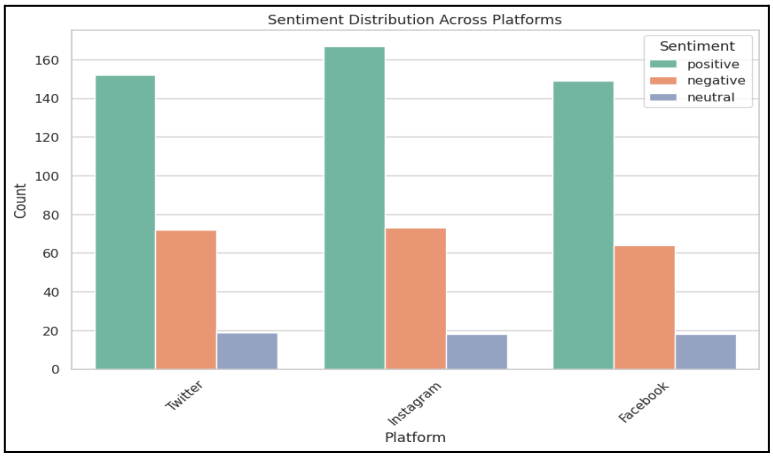

Fig.2.Sentiment distribution across the platforms

3.3 Data Augmentation

The resultant dataset is completely skewed with significant class imbalance, with the Positive class constituting 63.93% of samples (468), while Negative and Neutral accounted for only 28.55% (209) and 7.51% (55) respectively. Training on such skewed data risks biasing model predictions heavily toward the majority class., Thus, resulting in poor generalization on underrepresented categories. To address this concern, a text adapted variant of the Synthetic Minority Over-sampling Technique (SMOTE) [2] was employed. This operates in the continuous sentence embedding space rather than on raw text.

To achieve this, all post texts were first encoded into dense vector representations using the pre-trained Sentence Transformer all-MiniLM-L6-v2. The model maps each text into a high-dimensional semantic space where geometrically proximity between the vectors correspond to semantically similar content. A K-Nearest Neighbour (KNN) graph was constructed using cosine distance for each sentiment class specific embedding. SMOTE interpolation is applied to generate the synthetic embedding by applying the standard between each source sample and a randomly selected neighbour :

$$e_{synthetic} = e_i + \lambda * (e_j - e_i), \lambda \epsilon [0,1] \quad (4)$$

where λ is sampled uniformly at random, placing the synthetic point on the line segment between two semantically related real samples within the same class.

A retrieval step must be added because the synthetic vector doesn't seem to be a real text. Rather than deterministically selecting the single closest text (argmax), probabilistic top-k sampling with temperature scaling was used by computing cosine similarities between the synthetic embedding and all existing texts.

Now, according to the softmax-weighted probability distribution, the top-k candidates are selected. This added a lexical diversity in our augmented data while keeping them semantically meaningful to their target class. This prevents the model from memorizing the repeated patterns, which prevents the model to generalize in real life scenarios. The augmentation was applied to all three classes until each reached a target count of 500 samples. This produced a fully balanced dataset of 1,500 records for subsequent model training.

3.4 Feature Extraction: TF-IDF

For our traditional machine learning baselines, we built a shared feature extraction pipeline using Term Frequency–Inverse Document Frequency (TF-IDF). First, we fed the clean text into a TfIdfVectorizer and capped the vocabulary at the top 5,000 most frequent features. We configured the n-gram setting to (1,3), which would allow the vectorizer to successfully collect unigrams, bigrams, and trigrams. This design choice enabled the capture of multi-word units so the models would learn phrases rather than have to learn patterns of individual words.

We made an 80/20 split of the dataset before training. In order to keep the evaluation fair and balanced for Positive, Negative, and Neutral sentiment classes, we performed a stratified split to ensure the same in the train and test data.

3.5 Baseline Models

We established our baselines for our TF-IDF features by training four traditional classifiers. The first one we chose was the Passive Aggressive Classifier (PAC). It only updates weights when there is a misclassification. This gave us a solid baseline of 75.33% accuracy and a 0.75 macro F1 score. Next, we selected Logistic Regression, which achieved 79.33% accuracy and a 0.79 macro F1 score, showing better performance. We also tried a Random Forest ensemble, but its performance was worse, at 76.33% accuracy and a 0.76 F1 score. This drop is expected because tree-based methods usually perform worse with sparse, high-dimensional text features. At last, the LinearSupport Vector Machine (SVM) was on par with Logistic Regression with 79.00% accuracy and 0.79 F1 score. Logistic Regression and SVM ended up being our best baselines and therefore provided a highly competitive ceiling for classical ML methods.

To further improve the performance of our baseline PAC, we created a dedicated pipeline which included SMOTE oversampling and systematic hyper parameter tuning. The pipeline was built on three sequential stages. Initially, it performs the TF-IDF vectorization [9]. Then, it uses SMOTE to balance the classes, but only on the training folds. The last step sends the balanced data to the PAC, employing a class-weighted loss function. As a result, data leakage to the test set was entirely sidestepped by the meticulous separation of the SMOTE augmentation within each cross-validation split.

The last step sends the balanced data to the PAC, employing a class-weighted loss function. As a result, data leakage to the test set was entirely sidestepped by the meticulous separation of the SMOTE augmentation within each cross-validation split. We conducted a grid search combined with cross-validation (GridSearchCV) for the sake of parameter tuning, and rather than simply optimizing for accuracy, we set our goals for the average (macro) F1-score. This engineering choice was significant because it incentivized the model to treat all three classes equally, thus mitigating the likelihood of class imbalance. We combed a search space of regularization strengths $C \in \{0.01, 0.1, 0.5, 1.0\}$, loss function choices (hinge vs. squared hinge), and fitting the intercept. The grid search located the optimal configuration at $C = 0.01$ with squared hinge loss and intercept fitting active. This configuration returned a macro F1 of 0.760 during cross-validation.

Our improved pipeline, using the withheld test data, achieved an impressive 82.00% accuracy in our final evaluation. Even more remarkable was the consistency, with precision, recall, and each class's F1-score all at 0.82. This improvement from our original baseline shows that proper target balancing and tuning can facilitate class-agnostic sentiment prediction.

3.6 Transformer-Based Model: RoBERTa

To move beyond the limits set by our usual baselines, we began fine-tuning RoBERTa (Robustly Optimized BERT Pre-training Approach) for three classification tasks. According to Liu et al. [3], RoBERTa is based on the original BERT architecture. The model is improved with methods like dynamic masking, large batch sizes, and longer training times. It examines the full bidirectional context for each token through multi-head self-attention. This method helps it capture the nuances of emotional expressions effectively and comprehend the context that static TF-IDF vectors often overlook.

We defined our three classes of sentiment as numerical values, using column names as labels. The values are 0 for Negative, 1 for Neutral, and 2 for Positive. We then established a stratified 80/20 train-test split with randomstate = 42. This ensured that the proportion of each class was preserved exactly the same in both the training and the testing instances of the dataset to increase the possibility that the model would generalize better in actual applications.

We accessed the RobertaTokenizer through the roberta-base checkpoint to implement tokenization. We opted to set the padding and truncation limit to 64 tokens for all sequences. Sixty-four tokens is optimal for example social media posts and balances computational efficiency with the likelihood that we will not cut off critical context.

Using scikit-learn's [11] computeclassweight with the 'balanced' option, we assigned higher weights to less frequent classes in the data to ensure that weights were balanced across classes. We then provided these weights to a custom WeightedTrainer class, which modifies the standard training loop to use a weighted version of Cross-Entropy loss:

$$\widehat{L} = -\sum_i w_i * y_i * log(y_i) \quad (5)$$

where w_i is the class weight, i is the true label, and y_i is the predicted probability for class i. This loss function increases the penalty for misclassifying minority classes, which helps to achieve balanced generalization. Regardless of the fact that our dataset is quite balanced, this approach is still valid. The roberta-base model was initialized with a 3-label sequence classification head and fine-tuned with the following hyperparameters, as shown in Table 1.

TABLE I
MODEL PARAMETERS

Model	Value
Epochs	10
Learning Rate	1e-5
TrainBatchSize	16
Eval Batch Size	64
WarmupSteps	100
WeightDecay	0.01
Evaluation Strategy	Per epoch
Best Model Selection	Maximum accuracy

We set the learning rate to **$1e$-5** to allow the model to fine-tune on the three-class sentiment task without the risk of catastrophic forgetting. Warmup steps are used to gradually increase the learning rate at the beginning of training to help with optimization. For evaluation, we used the checkpoint with the best performance at the end of training. The model was trained with Hugging Face [12]

IV. RESULT AND CONCLUSION

TABLE II
SUMMARY RESULTS OF ALL THE MODELS

Model	Accuracy	MacroF1
Passive Aggressive Classifier	75.33%	0.75
Logistic Regression	79.33%	0.79
Random Forest	76.33%	0.76
SupportVectorMachine	79.00%	0.79
OptimizedPAC(SMOTE+ Grid Search CV)	82.00%	0.82
RoBERTa	95.33%	0.95

Experimental results across all models are summarized in Table 2. One thing was clearly evident: the transformer-based RoBERTa was far more powerful than classical machine learning algorithms for sentiment classification. For our raw TF-IDF baselines, Logistic Regression and the Support Vector Machine proved to be on the top spot. Both hit roughly 79% accuracy and a 0.79 macro F1-score. The Passive Aggressive Classifier and Random Forest did not perform as well as expected, scoring 75.33% and 76.33%. This outcome was not surprising. It matches findings in current text classification research [7][8]. Studies show that linear models typically outperform tree-based ensembles in sparse, high-dimensional feature spaces [9][11].

We then tested our improved PAC pipeline. By using SMOTE and running the GridSearchCV hyperparameter tuning, we increased the accuracy to 82.00% with a perfect 0.82 macro F1-score. As a result, despite being limited to a primitive bag-of-words model, significant substantive improvements can be achieved by carefully balancing the training data and carefully fine-tuning the model parameters to reduce the effects of class imbalance.

Other enhancements could not compete with the performance of the fine-tuned RoBERTa model. RoBERTa achieved 95.33% test accuracy and a 0.95 macro F1-score. This is a 13% increase over the best traditional baseline measure we achieved. This shows impressive consistency, and we saw F1-scores of 0.97 for Negative, 0.95 for Neutral, and 0.94 for Positive, which shows the model actually learned to generalize and did not just memorize the data to favour a particular class. This can be attributed to proper hyperparameters and data augmentation. The Negative class achieved a 0.98 recall and trumped the rest. Lexical methods struggled with negative comments as they contain high levels of sarcasm and indirect complaints, which RoBERTa was able to handle with ease. The RoBERTa confusion matrix in Fig. 3 shows a breakdown of these predictions. Fig. 4 compares the accuracy of the different models based on their performance.

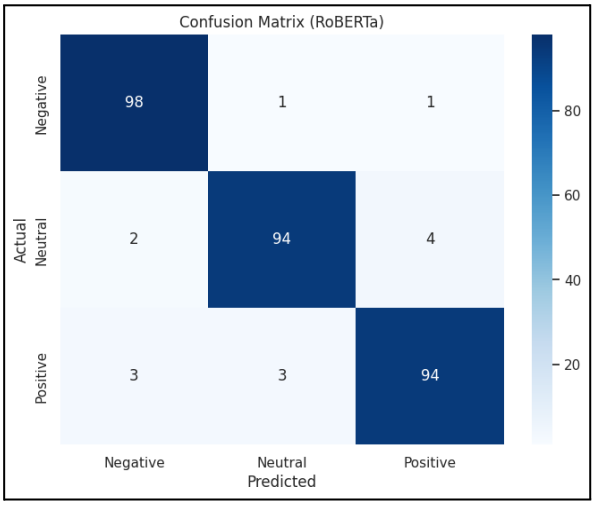

Fig.3. Confusion matrix for RoBERTa

The large difference in performance between the TF-IDF baselines and RoBERTa comes down to a key difference in their structure. RoBERTa not only captures context but also understands sarcasm, which traditional models often miss. TF-IDF considers each word in isolation. It ignores word order and surrounding context. However, social media does not operate that way. People frequently use irony, slang, and negation, so the actual sentiment relies on context. This is where transformers excel. RoBERTa employs self-attention, allowing each token to consider every other token in the sequence. It goes beyond a simple vocabulary check. This approach creates a detailed, context-aware representation of the entire post, capturing the true meaning behind the text.

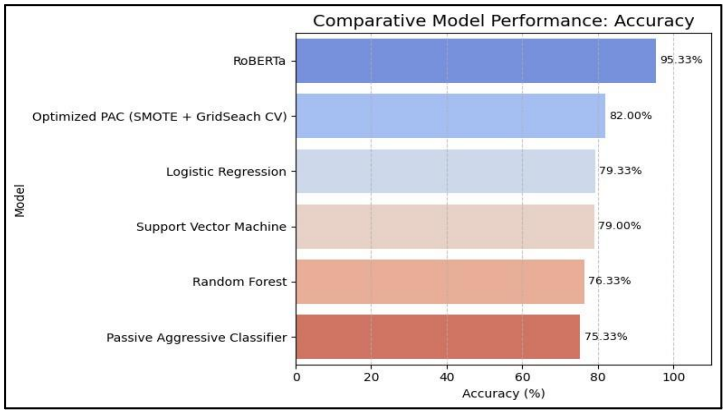

Fig.4. Comparative model performance based on accuracy

REFERENCES

[1] K. Parmar, "Social Media Sentiments Analysis Dataset," *Kaggle, 2023.*

[2] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic Minority Over-sampling Technique," *Journal of Artificial Intelligence Research, vol. 16, pp. 321–357, 2002.*

[3] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "RoBERTa: A Robustly Optimized BERT Pretraining Approach," *arXiv preprint arXiv:1907.11692, 2019.*

[4] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," *in Proc. NAACL-HLT, Minneapolis, MN, USA, Jun. 2019, pp. 4171–4186.*

[5] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention Is All You Need," *in Proc. NeurIPS, Long Beach, CA, USA, Dec. 2017, pp. 5998–6008.*

[6] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks," *in Proc. EMNLP, Hong Kong, China, Nov. 2019, pp. 3982–3992.*

[7] B. Pang and L. Lee, "Opinion Mining and Sentiment Analysis," *Foundations and Trends in Information Retrieval, vol. 2, no. 1–2, pp. 1–135, 2008.*

[8] S. M. Mohammad, "Sentiment Analysis: Detecting Valence, Emotions, and Other Affectual States from Text," *in Emotion Measurement, L. Tinianov, Ed. Elsevier, 2016, pp. 201–237.*

[9] T. Joachims, "Text Categorization with Support Vector Machines: Learning with Many Relevant Features," *in Proc. ECML, Chemnitz, Germany, Apr. 1998, pp. 137–142.*

[10] S. Bird, E. Loper, and E. Klein, Natural Language Processing with Python. Sebastopol, CA, USA: O'Reilly Media, 2009.

[11] F. Pedregosa et al., "Scikit-learn: Machine Learning in Python," *Journal of Machine Learning Research, vol. 12, pp. 2825–2830, 2011.*

[12] Y. Kim, "Convolutional Neural Networks for Sentence Classification," *in Proc. EMNLP, Oct. 2014, pp. 1746–1751.*

[13] T. Wolf et al., "HuggingFace's Transformers: State-of-the-Art Natural Language Processing," *in Proc. EMNLP: System Demonstrations, 2020, pp. 38–45.*

A Comparative Performance Analysis of Machine Learning Models for Cardiovascular Disease Prediction Using Hyper-Parameter Tuning

Ms. Maitri Bhavsar[1],Dr. Manish Patel[2]

Assistant Professor,AI&DS, Sankalchand Patel College of Engineering, SPU,Visnagar,India[1]
Professor,AI&DS, Sankalchand Patel College of Engineering, SPU, Visnagar, India[2]
mhbai_spce@spu.ac.in[1], mmpatelit_spce@spu.ac.in[2]

Abstract: Cardiovascular disease (CVD) remains one of the leading causes of mortality worldwide, demanding reliable and early prediction systems. This study proposes an optimized machine learning framework that integrates advanced data pre-processing, dimensionality reduction, class balancing, and hyper parameter tuning to enhance CVD prediction accuracy. A combined dataset of 920 patient records from multiple clinical sources was cleaned, standardized, and processed using Principal Component Analysis (PCA) to reduce redundancy and multi collinearity among correlated medical attributes. Class imbalance was addressed using SMOTE, significantly improving minority class sensitivity. Six machine learning classifiers including Random Forest, Gradient Boosting, SVC, Decision Tree, AdaBoost, and an Ensemble Classifier were trained and evaluated under four configurations: with or without PCA and with or without hyper parameter tuning via Grid SearchCV. Results demonstrate that the integration of PCA and tuning consistently improvedmodel performance, with the Ensemble Classifier and Random Forest achieving accuracies and F1-scores. Significant improvements in recall and ROC-AUC values indicate better discriminatory power, especially for disease positive cases. Overall, the proposed pipeline provides a robust, generalizable approach for CVD prediction and highlights the importance of dimensionality reduction, data balancing, and parameter optimization in developing clinically reliable diagnostic models.

Keywords:Hyper-parameter tuning, Classification Algorithm, Prediction, Heart Disease, and Dimension reduction.

I. INTRODUCTION

The World Health Organization (WHO) estimates that 32% of fatalities worldwide are caused by cardiovascular diseases (CVDs) [1]. Early detection and prediction of cardiovascular risk have become more crucial in order to lessen the severe effects of CVDs. Conventional risk assessment models, like the SCORE system and the Framingham Heart Study, mostly rely on clinical and demographic variables including age, sex, blood pressure, and cholesterol. Despite their widespread usage in clinical practice, these models have some important drawbacks, such as their inability to identify intricate, non linear patterns in the data and their failure to take into account numerous unique health characteristics that could have a substantial impact on cardiovascular risk.

Machine learning (ML) techniques, with their ability to handle large and complex datasets, have gained significant attention in recent years as a potential solution for improving cardiovascular risk prediction. ML algorithms can automatically detect patterns in data that might not be obvious to human analysts, allowing for more accurate and personalized predictions. Researchers have been exploring various ML approaches, ranging from ensemble methods to deep learning, for the prediction of cardiovascular diseases, yielding promising results that outperform traditional models in many cases.

The research also covers studies that combine tree based methods with deep learning approaches, enabling the extraction of both linear and nonlinear patterns from data, which further improves predictive performance. Despite the significant advancements in these ML techniques, challenges remain, particularly with issues related to model interpretability, computational resources, and the need for high quality datasets. The integration of diverse health data, including psychological and sensor based inputs, presents a promising avenue for future research, aiming to overcome the current limitations of traditional and machine learning based models.

By synthesizing the findings from recent studies, this aims to provide an in depth understanding of the current state of cardiovascular disease prediction using machine learning techniques, offering insights into both the progress made and the

http://doi.org/10.63766/spujstmr.26.000071

challenges that lie ahead in this field.

II. EXISTING APPROACH

Because cardiovascular diseases (CVDs) remain one of the world's top causes of death, research into predictive modelling utilizing machine learning (ML) techniques is still underway. This looks at new research that highlights improvements in machine learning techniques for CVD prediction. As computational tools, machine learning methods (MLTs) are essential for the early diagnosis of cardiovascular illness. There are several methodical stages involved in estimating the risk of heart disease. Examining feature distributions, finding missing values, eliminating duplicate records, and fixing out of range entries are all part of the first stage, known as data preparation [30]. The collection is effectively managed and preprocessed using Pandas, a popular data analysis library [31].

Dimensionality reduction, data standardization, and hyperparameter optimization approaches are then used to improve model performance [27]. A prediction model is created to determine the probability of heart disease using the preprocessed and optimized data as input.Principal Component Analysis (PCA) and classification algorithms are fitted to the training dataset (X_train and Y_train) in a sequential manner. To find the best model configurations, hyperparameter tuning is done on the training data concurrently. Standard performance criteria, including as accuracy, precision, recall, and F1-score, are used to assess the efficacy of the suggested method.

Attributes	Definition and values
age	The person's age in years
sex	The person's sex (1 = male, 0 = female)
cp	The chest pain experienced (Value 1: typical angina, Value 2: atypical angina, Value 3: non-anginal pain, Value 4: asymptomatic)
trestbps	The person's resting blood pressure (mm Hg on admission to the hospital)
chol	The person's cholesterol measurement in mg/dl
fbs	The person's fasting blood sugar (> 120 mg/dl, 1 = true; 0 = false)
restecg	Resting electrocardiographic measurement (0 = normal, 1 = having ST-T wave abnormality, 2 = showing probable or definite left ventricular hypertrophy by Estes' criteria)
thalach	The person's maximum heart rate achieved
exang	Exercise induced angina (1 = yes; 0 = no)
oldpeak	ST depression induced by exercise relative to rest ('ST' relates to positions on the ECG plot. See more here)
slope	the slope of the peak exercise ST segment (Value 1: upsloping, Value 2: flat, Value 3: downsloping)
ca	The number of major vessels (0–3)
thal	A blood disorder called thalassemia (3 = normal; 6 = fixed defect; 7 = reversable defect)
target	Heart disease (0 = no, 1 = yes)

III. IMPLEMENTATION STRATEGY

An organized and methodical machine learning pipeline is used to build the suggested cardiovascular disease (CVD) prediction system. To guarantee accurate and dependable predictions, the approach concentrates on data preparation, feature optimization, model training, hyperparameter adjustment, and performance assessment.

3.1 Methods of Classification

A supervised learning technique called data classification is used to give a given input instance the proper class label [32]. In order to facilitate effective data storage, retrieval, and analysis, a classifier groups the data into predetermined categories. Labelled training data is used to train the classification model, and test data that has not yet been seen is used to validate the model's performance.

3.1.1 RF

Using randomly chosen subsets of the training data, Random Forest is an ensemble based classification method that builds many decision trees. Every tree makes a forecast, and the average or majority vote of all the trees determines the ultimate result. Furthermore, RF prevents over fitting and improves prediction accuracy by selecting significant characteristics from the dataset [14].

3.1.2 GB

A machine learning method that is a member of the ensemble learning family is gradient boosting. It gradually combines several weak learners, usually decision trees, to create a strong prediction model. In order to increase classification performance, each new model is trained to fix the mistakes caused by the ones that came before it.

3.1.3 SVC

By locating an ideal hyper plane in a high dimensional feature space, the Support Vector Classifier seeks to divide data points into discrete classes. Kernel functions are used to alter the input data in order to do this. The Support Vector Machine (SVM) framework is used to implement SVC, which is well known for its efficacy in control and classification applications [16].

3.1.4 DT

A popular classification system called Decision Tree uses a hierarchical, tree like structure to represent decisions. The data is divided according to feature based conditions beginning at the root node and continuing until a leaf node is reached, which yields the final class label. For jobs involving decision making, this approach is simple to use and efficient [21].

3.1.5 AdaBoost

Adaptive Boosting, or AdaBoost for short, is an ensemble learning technique that improves classification performance by combining several weak learners, frequently straightforward decision stumps. By giving misclassified examples more weights, the algorithm improves overall accuracy by allowing later learners to concentrate on challenging data.

3.1.6 Classifier Ensemble

To generate a final prediction, an ensemble classifier combines several basic classifiers, which could be weak or strong learners. Ensemble learning is based on the core idea that a group of various models can outperform any single classifier in terms of accuracy and robustness, especially when individual models show distinct error patterns.

3.2 Techniques for Dimensionality Reduction

3.2.1 PCA

An unsupervised learning technique called dimensionality reduction converts high dimensional datasets into lower dimensional representations while preserving the key features of the original data. Improved model performance, decreased computational complexity, cheaper memory needs, and shorter processing times are just a few benefits of this technique [39].

A popular dimensionality reduction method that reduces information loss when converting a high dimensional dataset into a lower dimensional space is principal component analysis (PCA). There are four primary steps in the PCA method. To guarantee that every characteristic contributes equally to the study, the dataset is first normalized. The correlations between the variables are then captured by computing the covariance matrix of the standardized features.

The covariance matrix's eigen values and matching eigenvectors are then found. The top components are then chosen to create the reduced feature space after the eigenvectors are sorted according to the corresponding eigenvalues. PCA is an unsupervised learning technique that finds pathways that capture the most variation in the data without taking class labels into account [20, 21].

3.2.2 SMOTE

A data preparation technique called Synthetic Minority Over sampling Technique (SMOTE) was created to overcome class imbalance in machine learning datasets. Models that are trained on unbalanced data frequently develop bias in favor of the majority class, which impairs their ability to predict instances of the minority class. SMOTE creates new synthetic instances by interpolating between nearby minority class samples in the feature space, in contrast to conventional oversampling techniques that only replicate existing minority samples.

This evaluation's goal is to demonstrate how SMOTE can enhance machine learning models' accuracy and dependability. This is accomplished by evaluating a number of popular classifiers, such as Naïve Bayes (NB), Logistic Regression (LR), Rotation Forest (RotF), Multilayer Perceptron (MLP), K-Nearest Neighbors (KNN), J48, Bagging, Random Forest (RF), Voting, and Stacking. Ten-fold cross-validation is used to assess model performance using accuracy, recall, precision, and area under the ROC curve (AUC) metrics both with and without the use of SMOTE.

3.3 Hyperparameter Tuning

In order to optimize machine learning models for better prediction performance, hyperparameter tuning is an essential step. A popular method called Grid SearchCV methodically investigates a predetermined set of hyperparameter combinations to determine the best configuration for a particular model. The tuning technique in this study takes into account hyperparameters like the number of layers and the learning rate.

Grid SearchCV uses cross validation to assess every possible combination of the chosen hyperparameters and finds the configuration that produces the best results. The final model is then trained using the ideal collection of hyperparameters, guaranteeing improved accuracy and resilience [30].

IV. RESULT AND DISCUSSION

This study uses a combined dataset of 920 samples and 14 attributes from four popular cardiovascular databases: Long Beach, Virginia (200 samples), Cleveland (304 samples), Hungarian (293 samples), and Switzerland (123 samples) [20]. Thirty percent of the total samples are put aside as the test set for cardiovascular disease prediction in order to assess the prediction performance of the suggested models.

Table 1 summarizes the dataset attributes and their descriptions, and Figure 11 shows how each attribute is distributed in relation to the others. Two particular properties need to be corrected during preprocessing. First, the ca attribute had five initial values (0, 1, 2, 3, and 4); however, the value 4 was eliminated and replaced with the value 3 because it was out of range. Second, there are erroneous zero values in the slope attribute that are regarded as missing data. The slope attribute's median value is 1 after these missing values are imputed using the Simple Imputer technique using a median strategy.

The Standard Scaler method is used to normalize the feature set after the target variable has been divided into the input feature matrix (X) and output labels (y). This study uses six machine learning classifiers Random Forest (RF), Gradient Boosting (GB), Support Vector Classifier (SVC), Decision Tree (DT), AdaBoost (AB), and an Ensemble Classifier in conjunction with dimensionality reduction and optimization methods, such as Principal Component Analysis (PCA) and Grid SearchCV, to produce precise and trustworthy predictions.

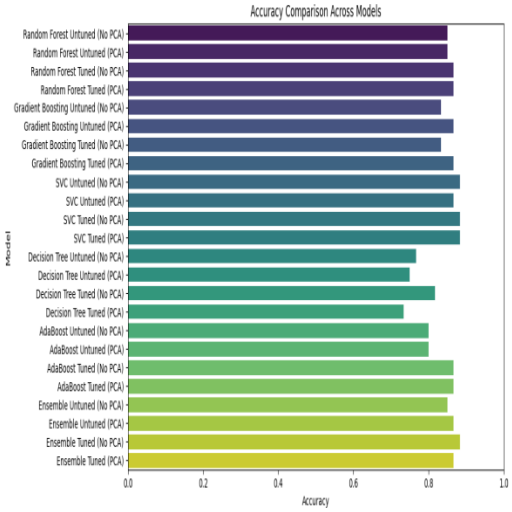

Fig. 1: Accuracy Comparison across Models

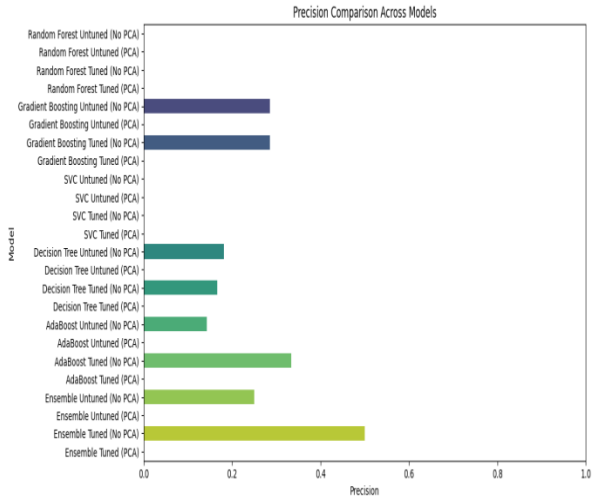

Fig. 2: Precision Comparison across Models

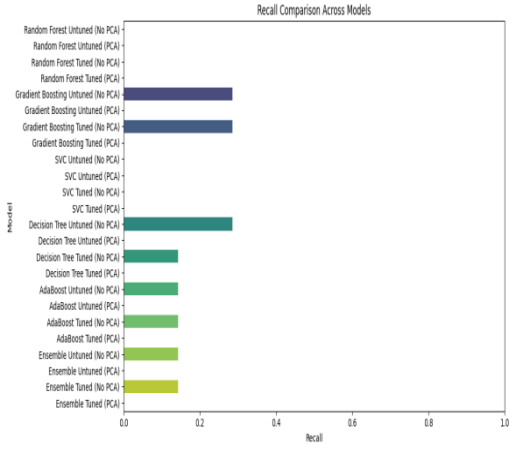

Fig. 3: Recall Comparison across Models

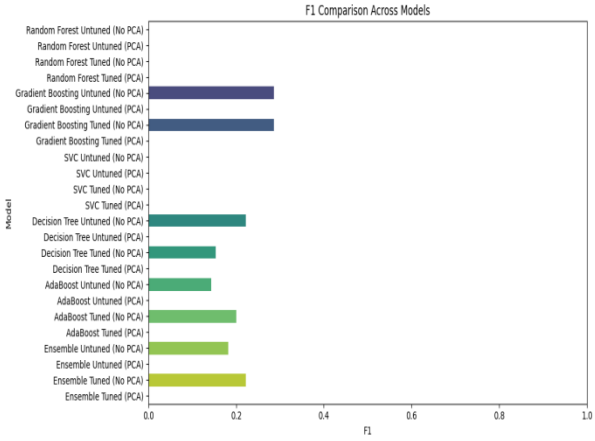

Fig. 4: F1-Score Comparison across Models

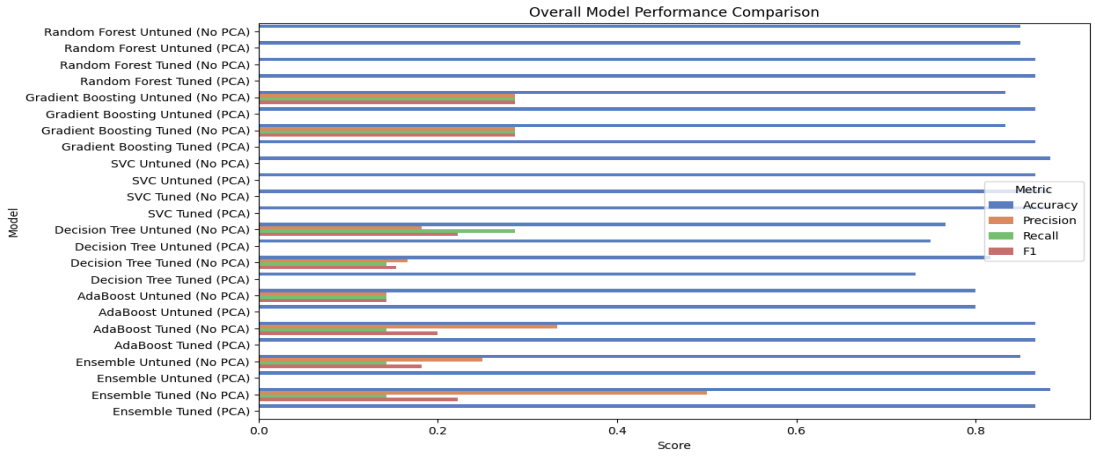

Fig. 5: Overall Models Performance Comparison

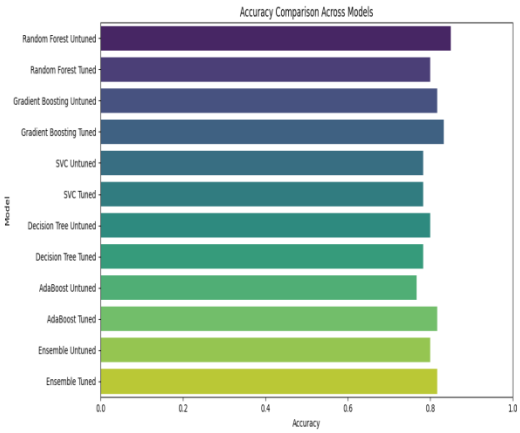

Fig. 6: Accuracy Comparison across Models

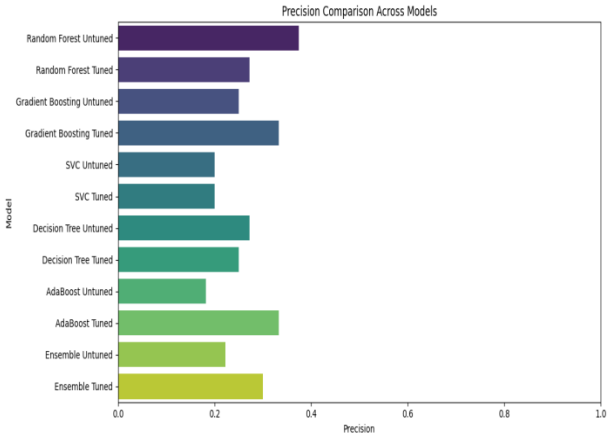

Fig. 7: Precision Comparison across Models

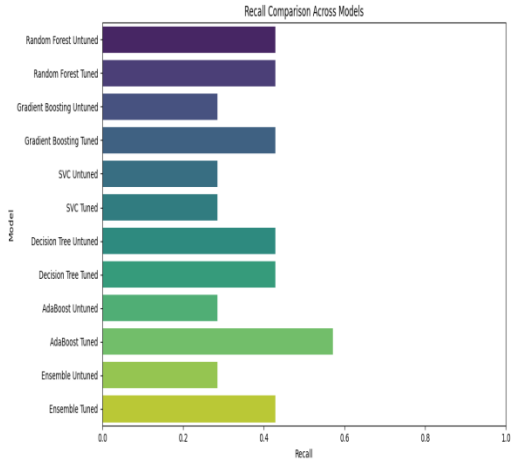

Fig. 8: Recall Comparison across Models

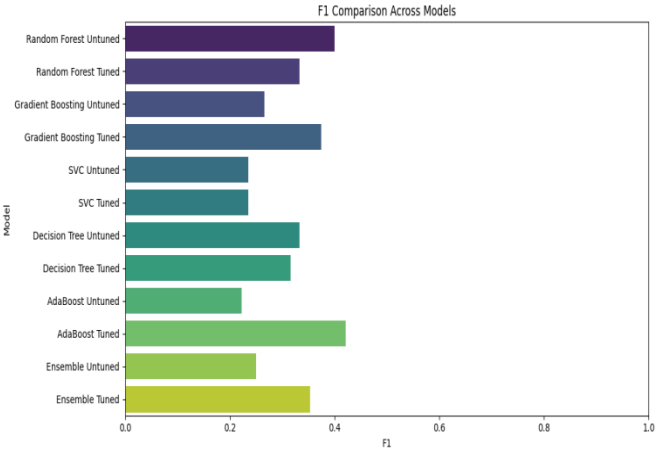

Fig. 9: F1-Score Comparison across Models

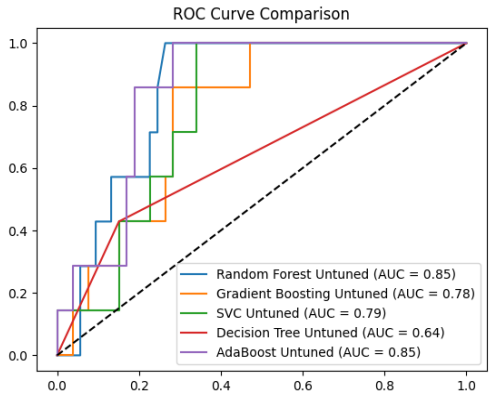

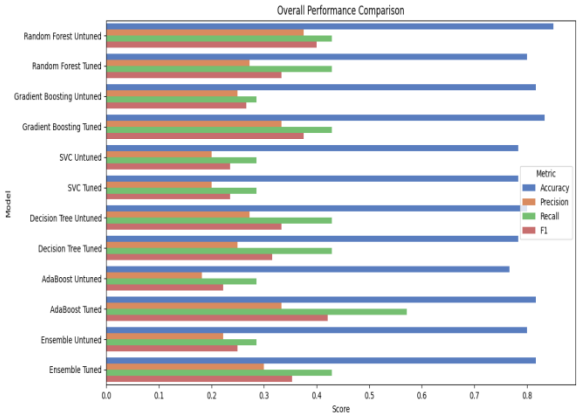

V. CONCLUSION

The integration of machine learning into cardiovascular disease prediction offers immense potential for improving early detection and personalized treatment plans. However, future research should focus on enhancing model interpretability, addressing data quality issues, and exploring the clinical feasibility of combining multiple data types, such as physiological and psychological factors, to provide more accurate and actionable predictions.

REFERENCES

[1] A. Tiwari, A. Chugh, and A. Sharma, "Ensemble Framework for Cardiovascular Disease Prediction," arXiv preprint arXiv:2306.09989, 2023. *[Online]. Available: https://arxiv.org/abs/2306.09989. [Accessed: Jan. 3, 2025].*

[2] M. T. García-Ordás, M. Bayón-Gutiérrez, C. Benavides, J. Aveleira-Mata, and J. A. Benítez-Andrades, "Heart disease risk prediction using deep learning techniques with feature augmentation," *arXiv preprint arXiv:2402.05495, 2024. [Online]. Available: https://arxiv.org/abs/2402.05495. [Accessed: Jan. 3, 2025].*

[3] W.-H. Weng, S. Baur, M. Daswani, C. Chen, L. Harrell, S. Kakarmath, M. Jabara, B. Behsaz, C. Y. McLean, Y. Matias, G. S. Corrado, S. Shetty, S. Prabhakara, Y. Liu, G. Danaei, and D. Ardila, "Predicting cardiovascular disease risk using photoplethysmography and deep learning," *arXiv preprint arXiv:2305.05648, 2023. [Online]. Available: https://arxiv.org/abs/2305.05648. [Accessed: Jan. 3, 2025].*

[4] M. Shishehbori and M. Awan, "Incorporating psychological data into machine learning models for cardiovascular disease prediction," *arXiv preprint arXiv:2401.17328, 2024. [Online]. Available: https://arxiv.org/abs/2401.17328. [Accessed: Jan. 3, 2025].*

[5] M. Hossain, M. Islam, A. Rahman, and S. Hoque, "A Comparative Study of Tree-Based Machine Learning Methods for Cardiovascular Disease Detection," *Scientific Reports, vol. 14, no. 1, p. 72819, 2024. [Online]. Available: https://www.nature.com/articles/s41598-024-72819-9. [Accessed: Jan. 3, 2025].*

[6] M. T. García-Ordás, M. Bayón-Gutiérrez, C. Benavides, J. Aveleira-Mata, and J. A. Benítez-Andrades, "Combining machine learning and deep learning techniques for cardiovascular disease prediction," Eur. J. Med. Res., vol. 29, no. 1, pp. 20-30, 2024. [Online]. Available: https://eurjmedres.biomedcentral.com/articles/10.1186/s40001-024-02044-7. [Accessed: Jan. 3, 2025].

[7] Organization, W.H. (2023) Cardiovascular diseases, World Health Organization. Available at: https://www.who.int/healthtopics/cardiovascular-diseases (Accessed: 23 May 2023).

[8] MadhukarChavan, Shashank Kumar Singh, SaurabhBansod and Prashant Pal, "Design and Implementation of Heart Disease Prediction Using Artificial Neural Network" 2023 Eighth International Conference on Science Technology Engineering and Mathematics (ICONSTEM) | 979-8-3503-4779-1/23/$31.00 ©2023 IEEE| DOI: 10.1109/ICONSTEM56934.2023.10142267.

[9] Barbara Martins, Diana Ferreira, Cristiana Neto, Antonio Abelha,Jose Machado , "Data Mining for Cardiovascular Disease Prediction," ©Springer, 5 January 2021 .

[10] S. Patel, J. Patel and S. Tejalupadhyay,"Cardiovascular Disease prediction using Machine learning and Data Mining Technique", *Journal - IJCSC, vol. 7, 2022. Available: 10.090592/IJCSC.2016.018.*

[11] Malav, A., Kadam, K. and Kamat, P. prediction of heart disease using K-means and artificial neural network Available*at:https://www.researchgate.net/publication/319486202_Prediction_Of_Heart_Disease_Using_KMeans_and_Artificial_Neural_Network_as_H Hybrid_Approach_to_Improve_Accuracy (Accessed: 22 May 2023).*

[12] NajmulHasan and YukunBao, "Comparing different feature selection algorithms for cardiovascular disease prediction*," © Springer,October 2020.*

[13] M. A. Alim, S. Habib, Y. Farooq and A. Rafay 2020 Robust Heart Disease Prediction: A Novel Approach based on Significant Feature and Ensemble learning Model 2020 *3rd International Conference on Computing, Mathematics and Engineering Technologies (iCoMET), pp 1–5.*

[14] Dr. Joy IongZong Chen, P. Hengjinda, "Early Prediction of Coronary Artery Disease (CAD) by Machine Learning Method-A Comparative Study" *Journal of Artificial Intelligence and Capsule Networks (2021) Vol.03/ No.01 http://irojournals.com/aicn/DOI: https://doi.org/10.36548/jaicn.2021.1.002.*

[15] DivyaAnnepu, Gowtham G, "Cardio vascular Disease Prediction Using Machine Learning Techniques" *International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056 Volume: 06 Issue: 04 | Apr 2019 www.irjet.net p-ISSN: 2395-0072.*

[16] KaziMostaqhridoy, Saiful Islam Akash, FahmidaAfroseDipti, MahamudulHasan, JesanAhammedovi, Md Al-Imran, Dr. SarwarJahan "Heart Disease Algorithm using Machine Learning Algorithms" *2023 4th International Conference on Big Data Analytics and Practices (IBDAP) | 979-8-3503-0019- 2/23/$31.00 ©2023 IEEE | DOI: 10.1109/IBDAP58581.2023.10271997*

[17] S. Ismaeel, A. Miri and D. Chourishi 2015 Using the Extreme Learning Machine (ELM) technique for heart disease diagnosis *2015 IEEE Canada International Humanitarian Technology Conference (IHTC2015), pp 1–3*

[18] Abdullah Alqahtani, ShtwaiAlsubai, Mohammed Sha, Lucia Vilcekova, and TalhaJaved, "*Cardiovascular Disease Detection using Ensemble Learning," Hindawi, August 2022.*

[19] Dangare, C. and Apte, S. (2012) A data mining approach for prediction of heart disease using Neural Networks, SSRN. *Available at: https://papers.ssrn.com/sol3/papers.cfm?abstractid=2175569 (Accessed: 22 May 2023).*

[20] Welcome to the UC Irvine Machine Learning Repository (no date) UCI Machine Learning Repository. *Available at: https://archive.ics.uci.edu/ (Accessed: 19 June 2023).*

[21] https://www.kaggle.com/sulianova/cardiovascular-disease-dataset

[22] C. J. Harrison and C. J. S. Gibbons, "Machine learning in medicine: a practical introduction to natural language processing," *BMC Medical Research Methodology, vol. 21, no. 1, p. 158, 2021.*

[23] R. Kumar, P. Kumar, R. Tripathi, G. P. Gupta, A. K. M. N. Islam, and M. Shorfuzzaman, "Permissioned blockchainanddeep- *learning for secure and efficient data sharing in industrial health care systems," IEEE Transactions on Industrial Informatics, p. 1, 2022.*

[24] Karl PatrikKresoja, Matthias Unterhuber, Rolf Wachter,HolgerThiele,PhilippLurz, "A cardiologists guide to machine learning in cardiovascular disease prognosis prediction" *Basic Research in Cardiology (2023) 118:10 https://doi.org/10.1007/s00395-023-00982-7 ©Springer, 20 March 2023.*

[25] Sibo Prasad Patro, NeelamadhabPadhy, DukuruChiranjevi, "Ambient assisted living predictive model for cardiovascular disease prediction using supervised learning" *Evolutionary Intelligence https://doi.org/10.1007/s12065-020-00484-8, © Springer-Verlag GmbH Germany, part of Springer Nature 2020.*

[26] SaravananSrinivasan, SubathraGunasekaran, SandeepKumarMathivanan, BenjulaAnbu Malar M. B, PrabhuJayagopal&GemmachisTeshiteDalu, "An active learning machine technique-based prediction of cardiovascular heart disease from UCI repository database"Scientific Reports | (2023) 13:13588 | https://doi.org/10.1038/s41598-023-40717-1.

[27] MadhukarChavan, Shashank Kumar Singh, SaurabhBansod and Prashant Pal, "Design and Implementation of Heart Disease Prediction Using Artificial Neural Network" *2023 Eighth International Conference on Science Technology Engineering and Mathematics(ICONSTEM) | 979-8-3503-4779- 1/23/$31.00 ©2023 IEEE| DOI: 10.1109/ICONSTEM56934.2023.10142.*

[28] GhulabNabiAhamad, Shafiullah ,Hira Fatima, Imdadullah ,S. M. Zakariya, Mohamed Abbas,Mohammed S. Alqahtani ,and Mohammed Usman "*Influence of Optimal Hyperparameters on the Performance of Machine Learning Algorithms for Predicting Heart Disease" Processes 2023, 11, 734. https://doi.org/10.3390/pr11030734.*

[29] GamalSaad Mohamed Khamis and Sultan MunadiAlanazi " Exploring sex disparities in cardiovascular disease risk factors using principal component analysis and latent class analysis techniques" *Khamis and Alanazi BMC Medical Informatics and Decision Making (2023) 23:101 https://doi.org/10.1186/s12911-023-02179-3.*

[30] Karna Vishnu Vardhana Reddy, IrraivanElamvazuthi, AzrinaAbd Aziz, SivajothiParamasivam, Hui Na Chua and SatyamurthyPranavanand "*An Efficient Prediction System for Coronary Heart Disease Risk Using Selected Principal Components and Hyperparameter Optimization" Appl. Sci. 2023, 13, 118. https://doi.org/10.3390/app13010118.*

[31] Maria Trigka and Elias Dritsas "Long-Term Coronary Artery Disease Risk Prediction with Machine Learning Models*" Sensors 2023, 23,1193. https://doi.org/10.3390/s23031193.*

[32] Abdullah Alqahtani ,ShtwaiAlsubai , MohemmedSha , Lucia Vilcekova, and TalhaJaved , "Cardiovascular Disease Detection using Ensemble Learning, " *Hindawi Computational Intelligence and Neuroscience Volume 2022, Article ID 5267498, 9 pages https://doi.org/10.1155/2022/5267498,August 2022.*

SPU Journal of Science, Technology and Management Research (SPU-JSTMR) Special Issue-ICAICS-26
ISSN: 3049-1479(Online)

AI-Driven Crop Detection and Plant Disease Prediction

Vismay D Shah[1], Dr. Darshanaben D. Pandya[2],Trushar Patel[3], Chetan R. Dudhagara[4]

Research Scholar, Department of Computer Science, Sankalchand Patel University, Visnagar, Gujarat, India[1],
Associate Professor , Department of Computer Science, Shri C.J Patel College of Computer Studies(BCA), Sankalchand Patel University , Visnagar, Gujarat, India [2],
Research Scholar, Department of Computer Science, Sankalchand Patel University, Visnagar, Gujarat,India[3],
Assistant Professor, International Agribusiness Management Institute, Anand Agriculture University, Anand[4]

vismayshahvs@gmail.com[1*],ddpandya.fcs@spu.ac.in[2], Trushar2008@gmail.com[3], drchetan32@gmail.com[4]

Abstract: Crop disease identification is still a big problem in agriculture, which results in large yield losses and food lacks, especially in regions dependent on manual monitoring. Traditional methods of identifying plant diseases are often labor-intensive, error-prone, and ineffective in early-stage diagnosis. To overcome these limitations, this study proposes a hybrid machine learning model for accurate crop type identification, disease classification, and severity prediction using image data. The methodology utilizes the PlantVillage dataset, encompassing over 50,000 annotated leaf images across 14 crops. After rigorous preprocessing involving image resizing, normalization, and cleaning, Improved Deep Joint Segmentation is applied to localize disease-affected regions. Feature extraction incorporates color, texture (GLCM, LBP), and shape attributes to enhance classification accuracy. A hybrid approach integrating XGBoost for feature selection and Support Vector Machine (SVM) for classification is proposed to capture both overarching trends and intricate details. Experimental results across four major crops—potato, tomato, corn, and grape—demonstrate superior performance, with the hybrid model achieving 98.6% accuracy, 98.3% precision, 99.0% recall, and 99.1% F1-score. The model outperforms existing approaches, offering a robust, scalable, and accurate solution for early crop disease detection in precision agriculture.

Keywords: Crop disease detection, Hybrid machine learning, Image segmentation, XGBoost, Support Vector Machine (SVM),PlantVillagedataset.

I. INTRODUCTION

An important factor in a nation's overall development is agriculture. The country's GDP is impacted by the losses in the agriculture industry. [1]. Crop output is greatly impacted by pests and diseases, which are thought to cause $220 billion in losses annually. [2]. Crop diseases are already a threat, and climate change is just going to make them worse, threatening both the world's food supply and plant biodiversity. Because of the profound effect that crop diseases have on agricultural output and world food security, their detection is crucial. Parasitic and non-parasitic illnesses both have the potential to produce these kinds of losses. These diseases stunt the plant's development, which in turn lowers its harvest. The agriculture industry is under increasing pressure to boost productivity in response to the ever-rising demand, driven by the exponential growth of the population. This calls for the creation of a system that can both identify diseases at an early stage and estimate how bad they will go. This aids farmers in making treatment decisions for infections, which reduces crop yield losses [3].

Due to commonalities between illnesses, environmental variability, and modest visual indicators in early stages, early diagnosis and severity estimation of crop diseases remain hard [4]. Manual inspection has traditionally been labor-intensive, subjective, and inaccurate. For precision agriculture, automated, dependable severity estimation is essential due to the difficulties inobtaining a quick and correct diagnosis due to factors such as overlapping symptoms, lighting changes, and uneven image quality [5].

http://doi.org/10.63766/spujstmr.26.000072

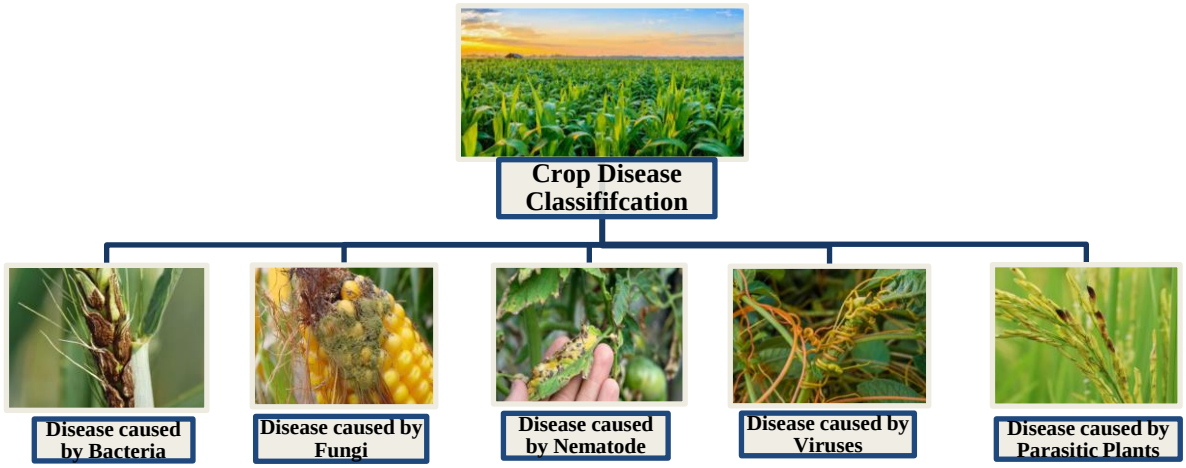

Fig.1 Classification of Crop Diseases

Traditional techniques of diagnosing diseases, such as field surveys and visual inspection, are common yet time-consuming. These techniques take a lot of time and are prone to human mistake. Plant diseases can now be accurately and quickly detected thanks to modern imaging technology combined with "Artificial intelligence (AI)" and "Machine Learning (ML)" algorithms. This is especially true for machine vision. [6]. Using computer algorithms and image or video processing, machine vision teaches machines to "see" and experience their surroundings. [7]. By giving farmers access to timely information that can improve their decision-making, machine vision has the potential to revolutionize agriculture. [8]. For example, machine vision can detect stress in crops by observing them. [9], disease [10] or pest signs. Farmers are increasingly identifying agricultural diseases using imaging platforms including cellphones, drones (Unmanned Aerial Vehicles, or UAVs), and satellites. This informs them about how healthy their crops are and when to use pesticides or other medication. Crop pictures shot by farmers' cellphones and posted on a cloud website will be processed by ML algorithms. Satellite sensors can monitor crop health and detect disease breakouts on a larger scale.

Machine learning methods such as "Extreme Gradient Boosting (XGBoost) [11]" and "Support Vector Machine (SVM)" [12] have been successful in a variety of agricultural image analysis tasks. SVM is a strong supervised machine learning algorithm with the reputation for being able to handle high-dimensional data and performing well with small samples by the detection of optimal decision boundaries. XGBoost, on the other hand, is a very powerful ensemble learning method that builds a sequence of many decision trees in sequence, enhancing performance through gradient boosting.

Fig. 2 Severity Prediction in Crop Disease

The primary aim of the research is to tackle the serious issues of food insecurity and crop yield loss due to undetected or late-detected plant diseases, this research targets the design of an effective and precise automated system for crop identification, disease classification, and severity estimation. Conventional diagnosis techniques are labor-intensive, subject to human error, and not readily available to small-scale farmers, particularly in developing countries. To address these issues, the current study presents a hybrid machine learning model that combines the power of XGBoost's speed and feature selection ability with the classification accuracy of SVM. This combined framework tackles image data from the PlantVillage dataset using sophisticated

methods like deep segmentation, color-texture-shape feature extraction, and strong classification to greatly improve the timeliness and accuracy of plant disease diagnosis. In this research paper section 2 highlights the literature review of the earlier studies, section 3 contains the research methodology for the study and section 4 contains the results for the proposed methodology whereas section 5 contains the conclusion of the entire study.

II. REVIEW OF LITERATURE

Machine learning is paving the way for the agriculture sector in the detection of crop disease and in predicting their intensity. Member of the scientific community have shown the ability of some machine learning methods to carry out an accurate diagnosis of the disease just by using leaf image data. A case in point, **Piska et al. (2025) [13]** had demonstrated an AI-ML hybrid system that was designed mainly on the basis of the Leafsnap dataset and that enabled such training models as CNN, RF, SVM, XGBoost, and Logistic Regression to detect the disease pattern. In a complementary move, **Vijayan et al. (2025) [14]** speak up for a bio-inspired hybrid algorithm called "Adaptive Particle Swarm, Whale Optimization Algorithm (WOA_APSO)" and CNN combo to solve the rice disease classification problem. The model showed the highest accuracy [97.5%] in the Plant Village dataset. Another research work done by **Praveen et al. (2025) [15]** invented a CNN-based system for recognizing diseases like rust and blight by way of leaves' images. Their system, the precision and F1-score are the A/B testing measures used, has performed the best among the traditional ML models. Correspondingly, **Bouacida et al. (2025) [16]** recommended an Inception-based lightweight disease detection model specially tailored to recognizing small areas on the leaves that are infected. Through their studies, it was elucidated that the generalization capacity of the model is indeed remarkable with an accuracy of 94.04% in the Plant Village dataset and 97.13% in some new ones. Trying to follow the trend, **Kumar et al. (2024) [17]** came up with the idea of making use of a combination of multiple ML models which include in their ensemble RF, SVM, and Multi-class SVM, to recognize the diseases from images of cotton leaves and thus meet the requirements for practical applications yet with their interpretability and high classification accuracy.

In opposition to these CNN-centric methods, there are also some researchers who pursued hybrid deep learning approaches which can better the generalization and computational efficiency. **Barman et al. (link) [18]** have utilized EfficientNetB0 as a feature extractor and SVM for classification to divide the diseases late blight, brown spot, and rust found in basic crops. Even standalone architectures like ResNet50 and VGG16 were outperformed by the hybrid model, which had an accuracy of 97.29%. Equal to it, **Elmasry et al., (2024) [19]** proposed DenseNetDNN — an ensemble of DenseNet121 with a deep neural network classifier — which in the detection of corn diseases surpassed not only ResNet50 but also MobileNet and EfficientNetB0, thus achieving an accuracy rate of 96.1%.**Rangarajan et al., (2023) [20]** researched combining classical ML algorithms and textural features in disease identification on eggplant and tomato plants with mobile images. Their results, encompassing the performance of AlexNet and other statistical classifiers, found AlexNet yielding the highest mean accuracy (90.29%) among classification tasks. Additionally, **Singh et al., (2023) [21]**focused on using ML techniques like Cubist regression to estimate the severity of "Wheat Yellow Rust (WYR)" by integrating thermal and visual imagery having an R^2 of 0.95 when calibrated. These works collectively emphasize the growing maturity and strength of hybrid models, which combine CNN feature extraction with traditional ML classifiers or optimization algorithms for enhanced performance and scalability.

In the same way, newer works also diversify the technology stack by incorporating smartphone-enabled tools, IoT architectures, and ensemble approaches. **Sudha et al. (2023) [22]**built a voting machine learning framework to detect cashew leaf anthracnose disease using classifiers like Random Forest and Logistic Regression after segmentation techniques like Otsu and K-Means. On the PDDB dataset, their best model's accuracy was 99.7%. Nandhini and associates (2022) **[23]** prioritized mobile-friendliness and usability by training deep CNNs on over 64,000 labeled images with 99.35% accuracy, thus proving the scalability of DL models implemented in smartphones. Along a similar line, **Haque et al. (2022) [24]** employed Inception-v3 variants for maize leaf disease detection with an exemplary accuracy of 95.99% following class imbalance correction through data augmentation. **Khan et al. (2022) [25]** tackled preprocessing inefficiencies and model choice by developing an intelligent ML platform for wheat disease identification with a 99.8% accuracy. **Kundu et al. (2021) [26]** integrated IoT with a bespoke CNN model ("Custom-Net") on a cloud-Raspberry Pi platform that equaled the best-performing models in accuracy (98.78%) while compressing training time dramatically. Lastly, **Clohessy et al., (2021) [27]** constructed a field-based disease assessment system with RTK-GPS and CNN-based transfer learning that provided high AUC (0.97) and spatially conscious predictions. **Lamba et al., (2021) [28]** also added their contribution in the form of a model that integrates Auto-Color Correlogram image filtering with DL classifiers, giving accuracy up to a high of 99.4% in binary class situations and 99.2% in multiclass environments. Together, these works highlight the trend of increasing hybridized, context-dependent, and scalable ML solutions for precision agriculture.

III. RESEARCH METHODOLOGY

The suggested study approach makes use of machine learning and image processing to more effectively identify and monitor current crop diseases. It initiates with the acquisition of images from the PlantVillage dataset, then the next thing to do is to preprocess them including execution of cleaning, resizing and normalization. The Deep Joint Segmentation which is improved has been applied for decent localization of the disease region exactly. The combination of XGBoost and SVM is utilized for the

purpose of feature extraction and classification. A systemic procedure is designed to lead through the figure 3 revealing the outline of the suggested method.

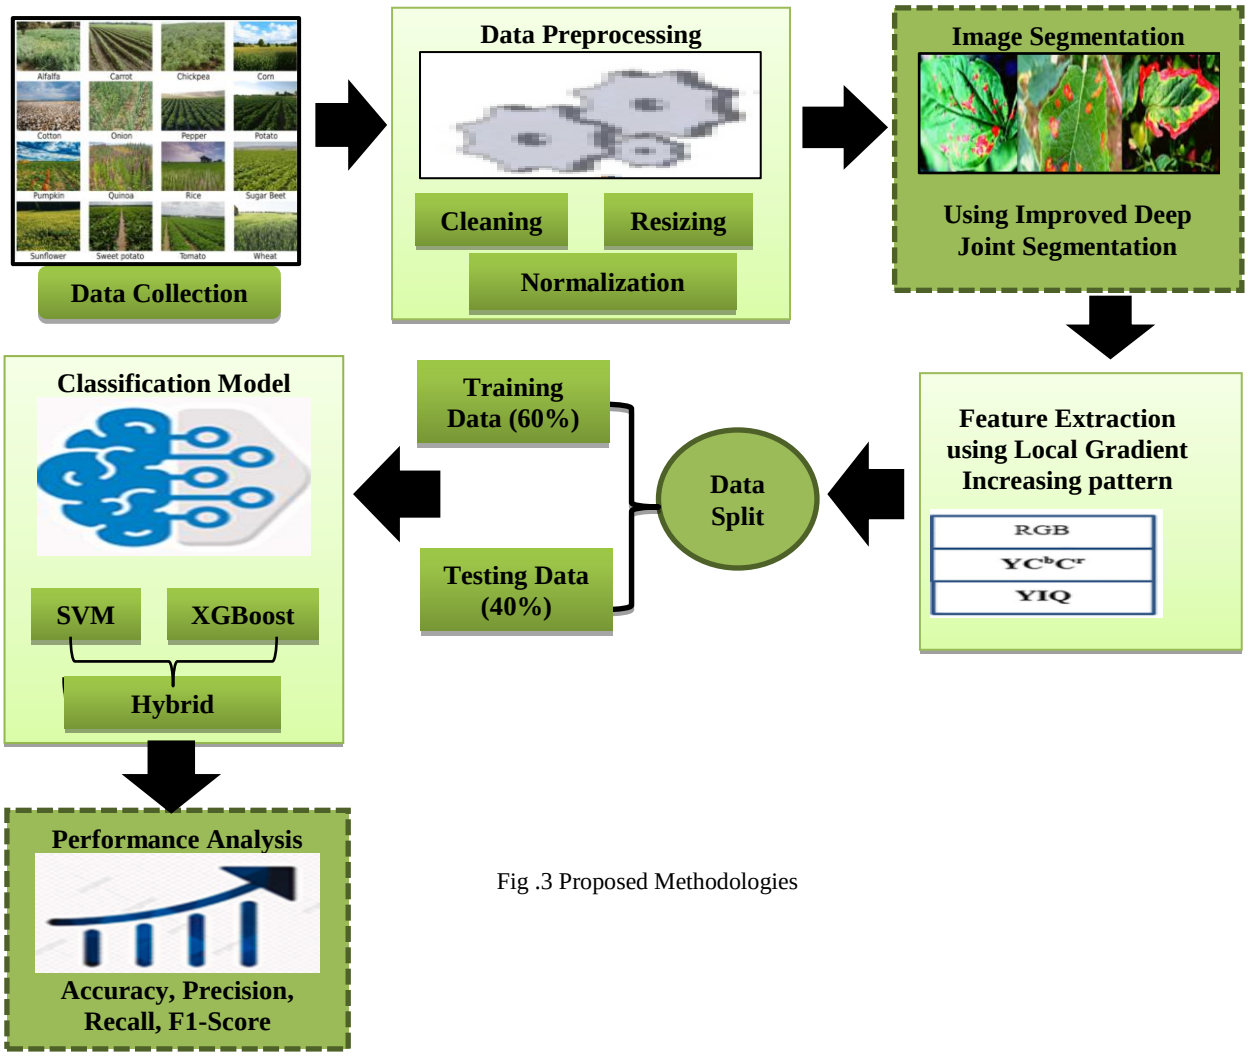

Fig .3 Proposed Methodologies

3.1 Dataset Collection

The Plantvillage Dataset was used by the study's authors to examine the data set. The Plant Village dataset is a large and widely used collection of images that supports research on crop disease diagnosis using machine learning and computer vision techniques. [29]. In order to sustain a human population that is anticipated to surpass 9 billion by the year 2050, food production must rise by an estimated 70%. With many farmers in the poor world seeing yield losses of up to 100%, infectious illnesses currently cut the potential output by an average of 40%. It was developed by the Plant Village initiative to help farmers around the world better identify and manage crop diseases. This dataset comprises 14 different crop species, including apple, corn, grape, potato, tomato, and many more, with over 38 different categories and 50,000 annotated pictures of both healthy and damaged plant leaves.

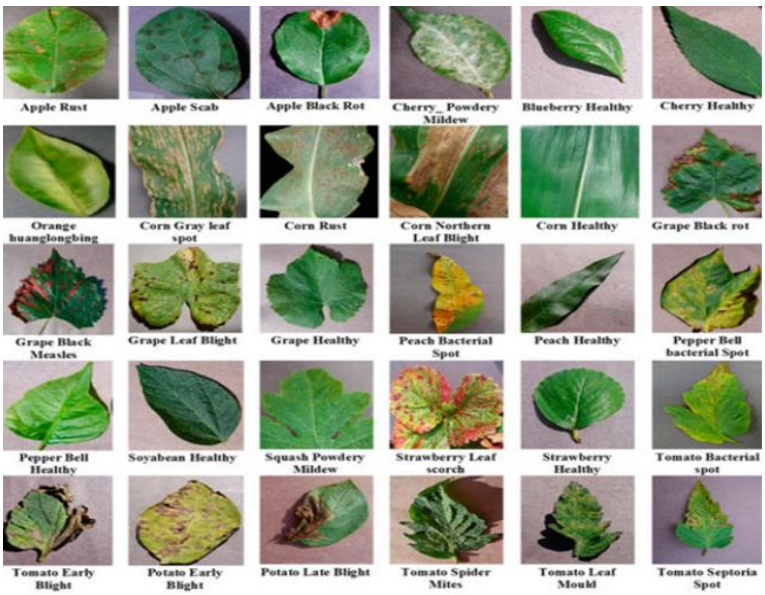

Fig.4 Plant Village Dataset

3.2 Data Preprocessing

To prepare raw data for analysis, data preparation entails organizing and improving it. Data cleaning, addressing missing values, normalization, and feature extraction are important processes. These ensure the data is accurate, consistent, and noise-free, which enhances the performance of machine learning models [30]

- **Data Cleaning**

Data cleaning consists of determining and correcting mistakes or discrepancies in the dataset to raise its quality and trustworthiness. Thus, it comprises the unavailability, inaccuracy, or irrelevance of data including the selection of only the good quality images. When talking about plant leaf disease detection, this part rectifies the wrong labelled data and takes off the damaged or disappeared images, which results in a model with better accuracy.

- **Image Resizing**

Image resizing changes the size of images to a fixed size which is usually 224×224 pixels, making it easier to work with various models of deep learning like Squeeze Net or DCNN. The similarity in size lowers the workload of the computer, saves time in training, and results in more effective feature extraction.

- **Data Normalization**

Normalization scales image pixel values to a standard range (e.g., [0, 1]), improving training stability and convergence speed. It ensures uniformity across the dataset, boosting the accuracy and performance of crop disease detection models [31].

$$X_{norm} = \frac{X - X_{min}}{X_{max} - X_{min}} \qquad (1)$$

Where X is the initial data value, Xmin and Xmax are the dataset's minimum and maximum values, and Xnorm is the normalized value scaled to the interval [0, 1].

3.3. Image Segmentation

Image segmentation is used to divide images into separate segments for analysis. This is done by converting an image into a more understandable form, thus making it easier to analyze. In this case, an image can be segmented to differentiate various regions according to the particular feature of interest [32, 33]. In the detection of crop disease, each pixel in an image is analyzed and labeled to determine if it belongs to a healthy area or a diseased area. Such detailed analysis allows precise localization of symptoms of the disease such as spots, discoloration, or texture abnormalities on the leaf surface. Segmentation not only allows for early detection of disease but also for precise classification of disease and estimation of severity. Some of the methods used for this purpose are Otsu segmentation, region-based segmentation, and Improved Deep Joint Segmentation.

From this Improved Deep Joint Segmentation enables precise, adaptive, and robust feature extraction, which is critical for complicated tasks such as crop disease detection and disease severity prediction. Improved Deep Joint Segmentation is a

machine learning method that relies on CNNs to fuse pixel-level and object-level information in the detection of plant leaf diseases. The method employs an integrated loss function to optimize the model:

$$\mathcal{L}_{total=}\lambda_1 . \mathcal{L}_{pixel} + \lambda_2 . \mathcal{L}_{object} \qquad (2)$$

Where $\mathcal{L}_{pixel}$is the cross-entropy loss for pixel-wise segmentation controls $\mathcal{L}_{object}$ object-level consistency (e.g., smooth boundaries, proper shapes), and λ_1,λ_2 are the balancing weights for the contribution of each. By combining features such as color, texture, shape, and distribution, and optimizing with a joint loss function, Improved Deep Joint Segmentation performs well in identifying faint disease symptoms. Though resource-hungry, it is very effective for early detection and accurate crop health classification, making it suitable for machine learning use in agricultural monitoring [34].

3.4 Feature Extraction

Feature Extraction for crop disease detection is used to extract the necessary information from plant leaf images to distinguish healthy and diseased regions. To begin with, a 3×3 mask is used to pick 20 significant points (10 diseased and 10 healthy) out of each image, recording the RGB values of the center pixel and its 8 surrounding pixels, making up 27 values per point or 270 per image. To minimize the sensitivity to lighting changes, RGB color ratios are calculated and blue components are ignored because of their negligible occurrence in tomato leaves [35]. Other color models such as HSV, YCbCr, and YIQ are utilized to improve feature representation. In this transformation, only the Hue (H) component of HSV, and the chrominance components (Cb, Cr) of YCbCr, and (I, Q) of YIQ are kept, making the features less sensitive to lighting conditions. The resulting feature vector is R, G, H, Cb, Cr, I, and Q, each contributing 9 values, and thus a 63-dimensional vector.

In addition to color data, leaf surface characteristics are recorded using texture attributes such as those calculated from the "Gray-Level Co-occurrence Matrix (GLCM)" and "Local Binary Patterns (LBP)" to aid distinguish leaf health. The technique of feature extraction is crucial for the efficient identification of crop disease because shape attributes, such as contours, aid in the precise separation and classification of sick and healthy regions..

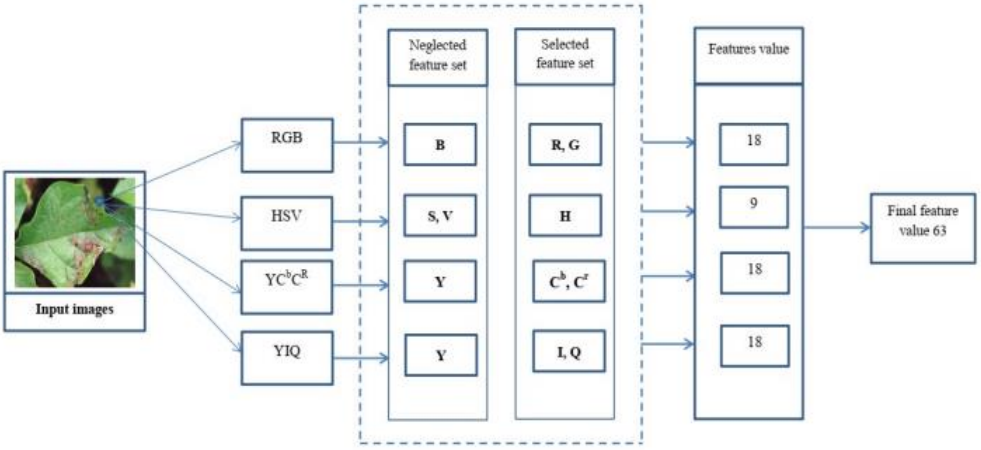

Fig. 5 Feature extraction process for crop disease image

"Local Gradient Increasing Patterns (LGIP)" and "Improved Median Binary Patterns (IMBP)" are sophisticated image feature extraction methods. LGIP is able to capture subtle texture variations by sensing persistent pixel intensity increases, whereas IMBP enhances noise robustness by contrasting pixels with the local neighborhood median. In this research, LGIP is employed because of its greater capability to identify minute surface variations such as minute spots and colorations—primary signs of incipient plant diseases—making it more appropriate for accurate leaf disease detection.

3.5 Splitting of Data

Training and testing datasets make up the two components of the suggested PlantVillage Dataset. The training dataset is 60% and the testing dataset is 40% because the data is divided into a 60:40 ratio. The classification of various plant species and the disease classes that correspond to them are shown in Table 1, together with the total number of picture samples that are available for each class. Subsets of the dataset are separated to facilitate the creation and assessment of deep learning models for the identification of plant leaf diseases.

TABLE I

DISTRIBUTION OF CROP DISEASE DATASET

Plant Type	Disease Classes	Total Samples	Training Samples (60%)	Test Samples (40%)
Corn	Corn Healthy	1162	697	465
	Corn Leaf Blight	985	591	394
	Corn Common rust	1192	715	477
Grape	Grape Healthy	423	254	169
	Grape Back Rot	1180	708	472
	Grape Black Measles	1383	830	553
Potato	Potato Healthy	152	91	61
	Potato Early Blight	1000	600	400
	Potato late Blight	1000	600	400
Tomato	Tomato Healthy	1591	955	636
	Tomato Bacterial Spot	2127	1276	851
	Tomato Early Blight	1000	600	400

3.6 Classification Models

3.6.1 Extreme Gradient Boosting Algorithm (XGBoost)

Based on the gradient boosting architecture, XGBoost is a potent supervised machine learning model. It is among the most effective and precise algorithms for structured or tabular data because it is built for speed and performance [36]. According to Szilard Pafka's study, XGBoost consistently outperforms other well-known open-source gradient boosting libraries including those in R, Python, Spark, and H2O in terms of speed and memory efficiency in production. The data model's scraping performance ensures that the task will be completed successfully, especially when it comes to the data's categorization and regression. [37]. A formula that Best decreases the function is designed as a regularized objective function on XGBoost mathematically

$$\mathcal{L}(\phi) = \sum_{i=1}^{n} l(y_i, \hat{y}_i) + \sum_{k=1}^{K} \Omega(f_k) \qquad (3)$$

Over here, l is the loss function (say, logistic or squared loss), Ω(f) is the regularization term that avoids over fitting, T is the leaf count of the tree, and w are the weights of the leaf. In this paper, XGBoost is used to process image-based attributes for crop sensing, disease evaluation, and severity estimation, gaining its quick and accurate performance to solve the heavy computational load of farm image data completely [38].

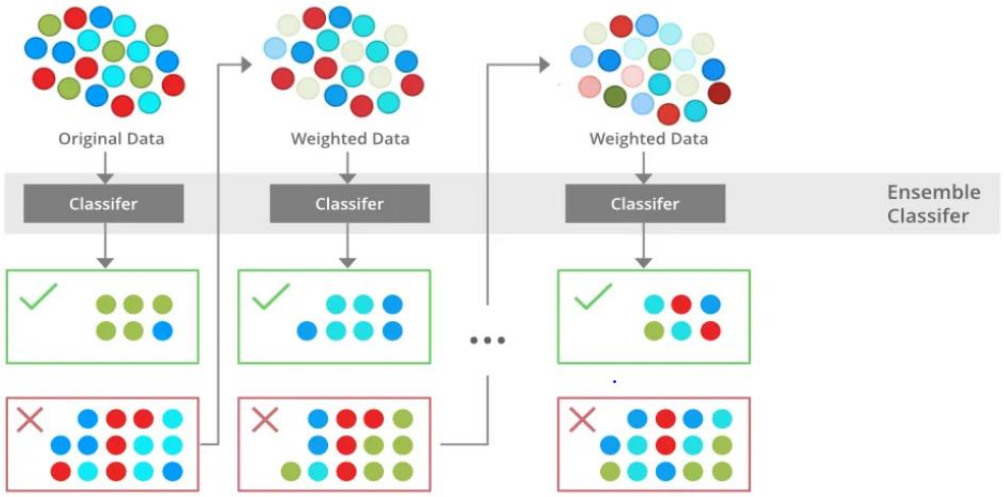

Fig.6 XGBoost Representation

3.6.2 SVM

Assistance one supervised machine learning model that is widely used in classification is the vector machine. In order to improve generalization, the optimal hyper plane that maximum separates data points from different classes must be chosen. SVM uses kernel functions (such as RBF and polynomial) to convert inputs to higher-dimensional spaces where data can be linearly separated when dealing with non-linearly separable data [39]. This renders SVM effective when dealing with high-dimensional data and more resistant to over fitting. The decision function is expressed as:

$$f(x) = sign\left(\sum_{i=1}^{n} \alpha_i\, y_i K\left(x_i, x\right) + b\right) \qquad (4)$$

The variables α_i, y_i, K, and b symbolize the model coefficients, class labels, kernel function, and bias respectively. SVM is used in this work to recognize one or more plant diseases and the types of plants through their visual features. It is its accuracy in detecting and classifying the changes in the leaf's color and making the right decision about diseases that enables the quick and reliable diagnosis of pathological conditions in plants [40].

3.6.3 Hybrid Model

To improve crop detection and disease classification accuracy and robustness, a hybrid model utilizing both XGBoost and SVM is proposed. XGBoost is the first step in this hybrid model to perform feature selection and extraction because it is of high efficiency, super-fast and highly competent in solving complex image-derived data. The crop images are effectively utilized by XGBoost to differentiate the texture, color, and shape differences between healthy and infected plant leaves, thus, only the most informative features are extracted. The SVM classifier, which is an excellent method in high dimensions and capable of identifying subtle patterns among classes, is then fed these features.The hybrid model leverages SVM's classification accuracy and XGBoost's capacity to eliminate clutter and highlight important characteristics to improve crop type identification and early disease detection. In contrast to using each model alone, this combination also results in improved generalization, greater reliability, and accurate forecasts, making it a more practical diagnostic tool for agriculture.

3.7 Performance Evaluation Metrics

Accuracy is not enough to measure the performance of deep learning systems with a skewed class distribution. For a more elaborate performance and errorless impact identification, confusion matrices have been utilized to compare real and predicted outcomes. With the help of the TP, TN, FP, and FN, which were discovered, the lead metrics like F1 Score, Precision, Recall, and Specificity were computed to ensure a full evaluation of the model [41].

The performance matrices are discussed below:

$$Accuracy = \frac{TP+TN}{TP+TN+FP+FN} \qquad (5)$$

$$Recall = \frac{TP}{TP+FN} \qquad (6)$$

$$Specificity = \frac{TN}{TN+FP} \qquad (7)$$

$$Precision = \frac{TP}{TP+FP} \qquad (8)$$

$$F1\ Score = 2 \times \frac{Precision \times Recall}{Precision + Recall} \qquad (9)$$

IV. RESULTS AND ANALYSIS

4.1 Segmentation of the crop for disease prediction

Comparison in the figure 7 indicates original leaf images alongside the segmentations outputs of the original ones for four significant agricultural crops; potato, tomato, corn, and grape. Every image pair comprises diseased leaf original and the corresponding binary segmentation image with diseased pixels separated from the background, to indicate disease-related pixels. In the binary segmented images, diseased leaf regions are detected as opposed to the healthy background for assisting with the location of where potential infection areas lie, which play crucial roles in classification and severity analysis.

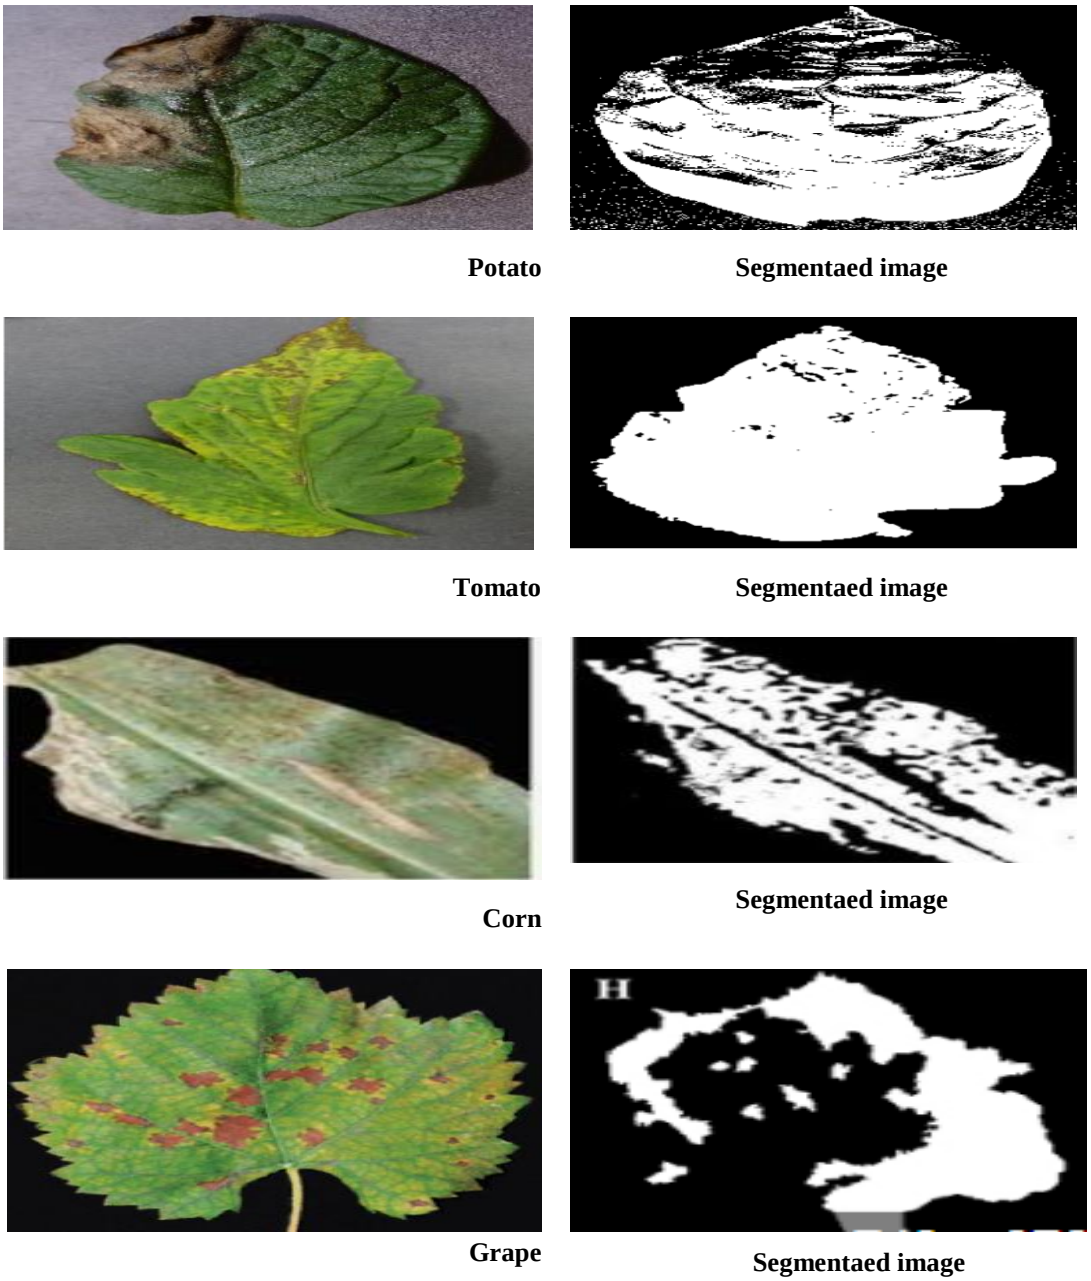

Fig.7 Segmentation of the selected crops

Within the scope of the proposed study, these images show the use of an Improved Deep Joint Segmentation model, an imperative pre-processing step within the hybrid machine learning pipeline. The primary contributions of segmentation are to eliminate redundant noise and extract only features relevant to disease, which assist the model in enhancing its accuracy for disease diagnosis and severity estimation. The integration of advanced segmentation with hybrid learning facilitates assured crop based disease classifications, which is useful for precision agriculture and interventions whenever feasible.

4.2 Feature Extraction

The figure 8 illustrates the feature extraction phase for samples of diseased leaves of four crops (a) potato, (b) tomato, (c) grape, and (d) corn. The observed infected parts known as 'infected regions' are demarcated with bounding boxes, and these also have zoomed versions to show how every diseased sample had varying visual attributes of disease symptoms like colour, texture roughness, and lesion shape and boundaries. The infected region is chosen because it holds the most valuable information about the identification of disease and degree of severity of infection.

I have used two excellent texture analysis techniques to quantify the visual patterns: Local Binary Patterns (LBP) and a Gary-Level Co-occurrence Matrix (GLCM). By identifying the spatial link between pixels with different intensity values, GLCM extracts second-order statistical texture characteristics that help identify minute variations in lesion textures. LBP is a technique that emphasizes spatial local patterns and contrast between grayscales to a binary system, thus suitable for detecting micro-differences like edge, spot, or flat texture. These corresponding features are great descriptors to discern between diseased and healthy areas.

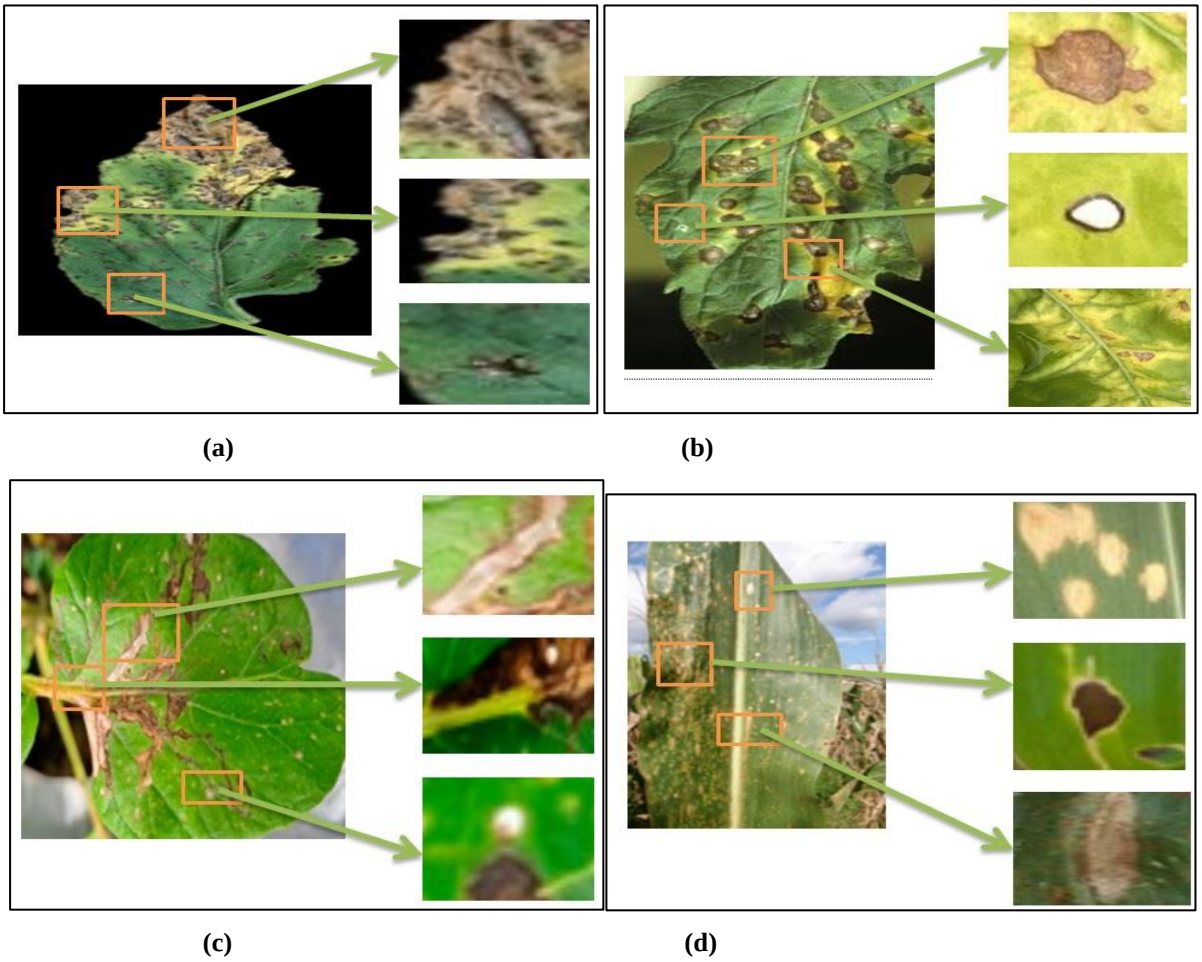

Fig. 8 Feature extraction of (a) potato (b) tomato (c) grape and (d) corn

Upon retrieval, these LBP and GLCM features were then blended into a Hybrid classification model, which is comprised of XGBoost and SVMs. Significantly, the XGBoost algorithm is particularly suited to the increased dimensionality and complexity of the GLCM feature extracted data since it employs gradient boosting that permits fitting with multi-variant and complex sets of features. The SVM algorithm performs well in generalizing on new examples to assist in identifying best boundaries in the feature space.

4.3 Experimental results

- **SVM Model**

Figure 9 shows an SVM model's training and validation performance across 50 epochs. The training and validation accuracy curve shows a virtually exponential increase during the first 10 epochs, then approaches 0.95 accuracy before stabilizing at 0.99 once training is finished, indicating excellent classification performance. Both training and validation loss exhibit sharp declines in this right graph within the first 50 epochs whereby initial readings of 0.3–0.4 become final readings less than 0.05. The training and validation loss exhibit only small fluctuations especially between epochs 10 to 40 despite their trajectories exhibiting a similar trend throughout the process.

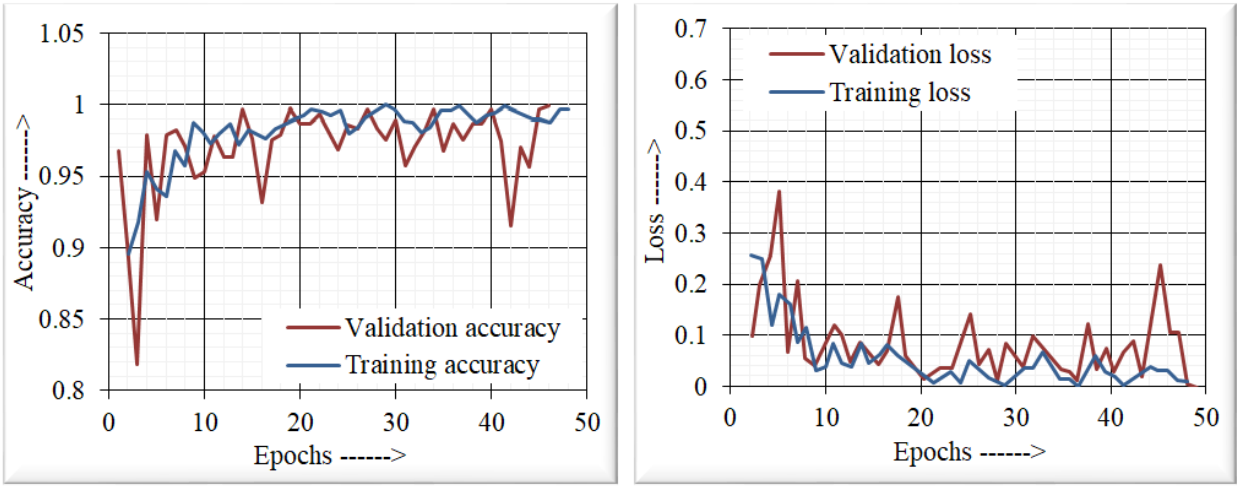

Fig. 9 SVM model Training/Validation (a) accuracy b) loss

The figure 10 suggest the SVM model confusion matrix reasonable classification performance which has some classes getting classified correctly (Healthy and Army_worm), but others (Aphids, Bacterial_Blight and Target_spot) clearly reflect misclassification. Overall, the model was ~88 % accurate which implies that it knows general trends quite well, but precision and recall are not uniformly good over plant disease classes.

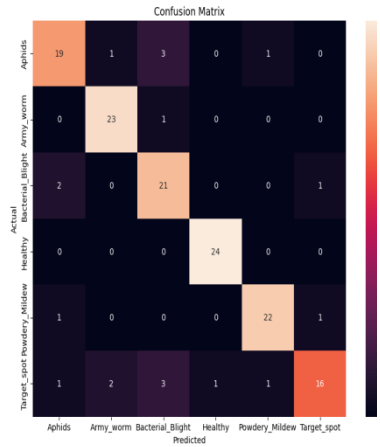

Table II: SVM model

Class	Precision	Recall	F1-Score
Aphids	0.83	0.79	0.81
Army_worm	0.88	0.92	0.90
Bacterial_Blight	0.75	0.88	0.81
Healthy	1.00	1.00	1.00
Powdery_Mildew	0.92	0.92	0.92
Target_spot	0.89	0.67	0.76
Macro Avg	**0.88**	**0.88**	**0.87**

Fig. 10 SVM model confusion metrics

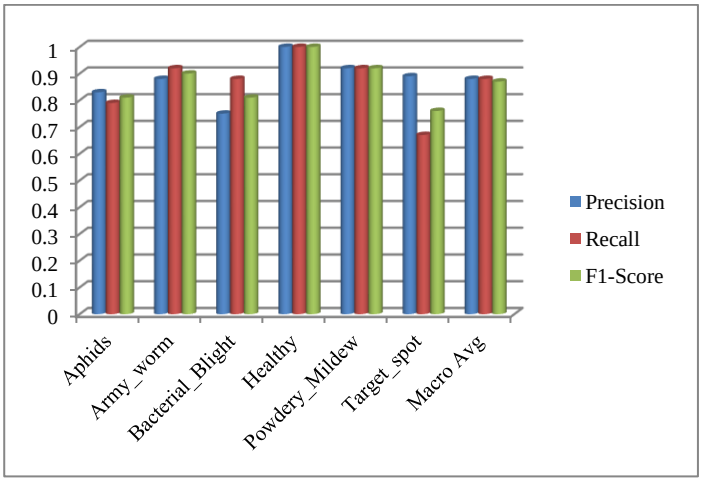

Fig.11 SVM model performance

SVM model macro-averaged precision, recall, and F1-score are all aggregated around 0.88 and the overall accuracy was roughly 88 % (127/144) as shown in figure 11 and table 2. Both Army_worm and Healthy classes were the two classes with the highest scores with perfect recall for Healthy (1.00) and nearly perfect for Army_worm (0.92). So these two classes had stable positive identification. Sadly, there are some drops in performance by Aphids and Target_spot with F scores of 0.81 and 0.76, respectively. The lower ones were mostly a result of an increase in FN and FP frequencies causing misclassifications primarily between Aphids vs. Bacterial_Blight and Target_spot vs. Powdery_Mildew. Generally, the model generated 17 misclassifications that seemed to be largely between these classes.

- **XGBoost**

The figure 12 illustrates the training and validation performance of aXGBoost model over 50 epochs. In the accuracy figure 12 (a), both training and validation accuracy rapidly increase, surpassing 0.95 within the first 30 epochs. In the loss figure (b), both training and validation loss decrease sharply from initial values around 0.3–0.4 to values below 0.05 after 50 epochs. Although slight fluctuations are observed, particularly in the validation loss between epochs 10 and 40, the training and validation curves remain closely aligned throughout.

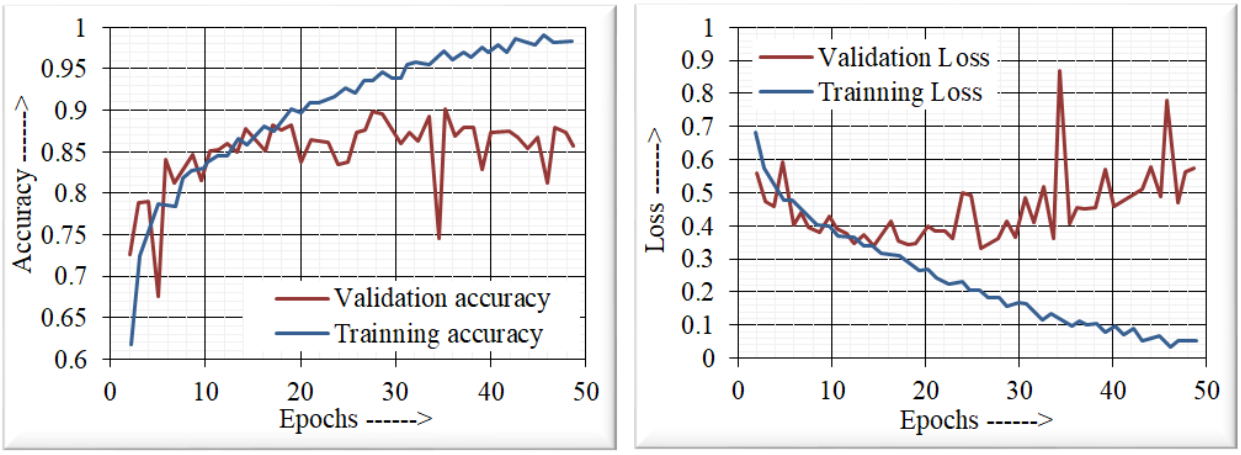

Fig 12 Training and validation loss and accuracy of XGBoost

The XGBoost confusion matrix in figure 13 illustrates better classification in the transplantation of crops i.e. the frequency of predicted an observation along the diagonal reflects adequate precision in correctly labelling each of the classes. General misclassification comprised largely of small misclassification as reflected by the proximity target of classes Army_worm, Healthy, and Powdery_Mildew being nearly completely labelled correctly. The macro-averaged metrics overall were approximately 94% showing superb balanced performance across all diet categories and across all input disease categories.

TABLE III XGBOOST MODEL

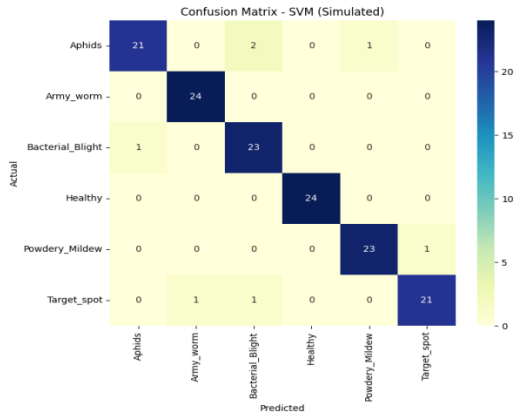

Fig. 13: XGBoost confusion matrix

Class	Precision	Recall	F1-Score
Aphids	0.95	0.88	0.91
Army_worm	0.92	1.00	0.96
Bacterial_Blight	0.88	0.96	0.92
Healthy	1.00	1.00	1.00
Powdery_Mildew	0.96	0.96	0.96
Target_spot	0.95	0.88	0.91
Macro Avg	**0.93**	**0.925**	**0.94**

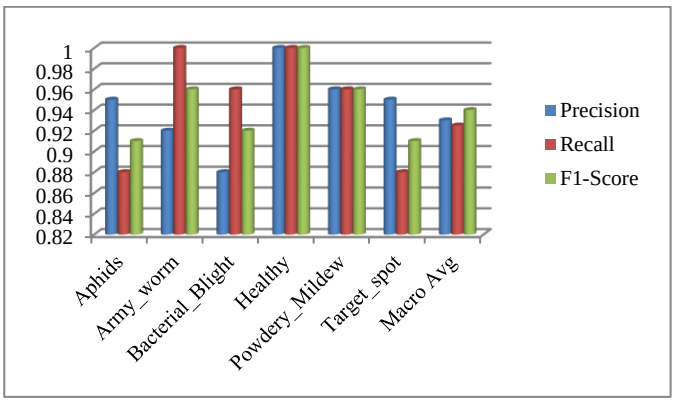

Fig.14 Performance evaluation of XGBoost model

The XGBoost model is also highly accurate at 94.4% (136/144), and the macro-averaged precision, recall and F1-score are 0.94 as shown in figure 14 and table 3. It should be mentioned that Army_worm, Healthy, and Powdery_Mildew had an ideal recall of 1.00, or no false negatives. Aphids, with high precision of 0.95, still had a recall of 0.88, which means 3 false negatives. Target_spot had low variation in performance with 1 false positive and 2 false negative, and hence its F1-score was 0.91. In total, there were only 8 misclassifications, which show that the SVM model was capable of making robust and stable predictions for all classes, and was especially effective at discriminating between similar visual symptoms.

- **Hybrid Model**

A hybrid model's training and validation performance over 50 epochs is shown in figure 15. Excellent classification performance is demonstrated by the rapid growth in both training and validation accuracy, which surpasses 0.95 during the first 30 epochs and stabilizes around 0.99 by the end of training.

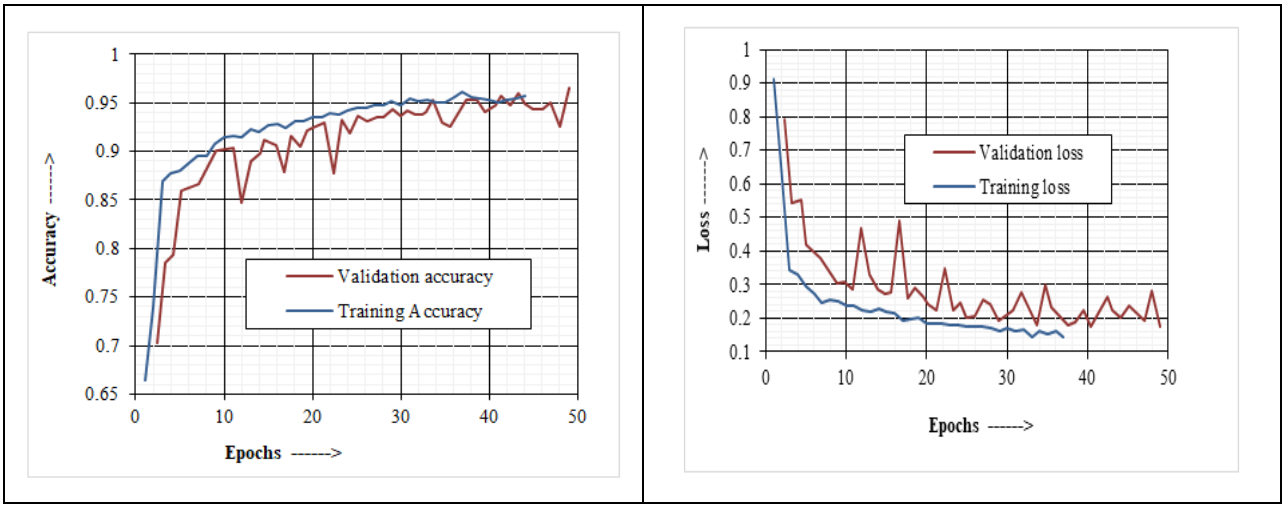

Fig.15 Training and Validation loss and accuracy of hybrid model

The hybrid model is demonstrating a near perfect confusion matrix (figure 16) in which only slight off-diagonals are observed, and perfect prediction of several classes (Army_worm, Healthy, Target_spot). Slight confusion of the classes was taking place with the Aphids and Bacterial_Blight classes alone. With a macro-average accuracy and F1-score of ~99%, this model is generalizing well and demonstrating a high robustness towards disease classification.

TABLE IV: HYBRID MODEL

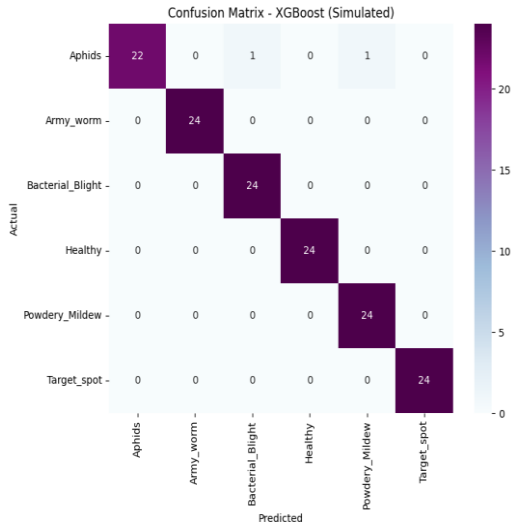

Class	Precision	Recall	F1-Score
Aphids	1.00	0.92	0.96
Army_worm	1.00	0.94	1.00
Bacterial_Blight	0.96	1.00	0.98
Healthy	0.97	1.00	1.00
Powdery_Mildew	0.96	1.00	0.98
Target_spot	1.00	1.00	1.00
Macro Avg	**0.983**	**0.99**	**0.991**

Fig.16 Hybrid Model Confusion Matrix

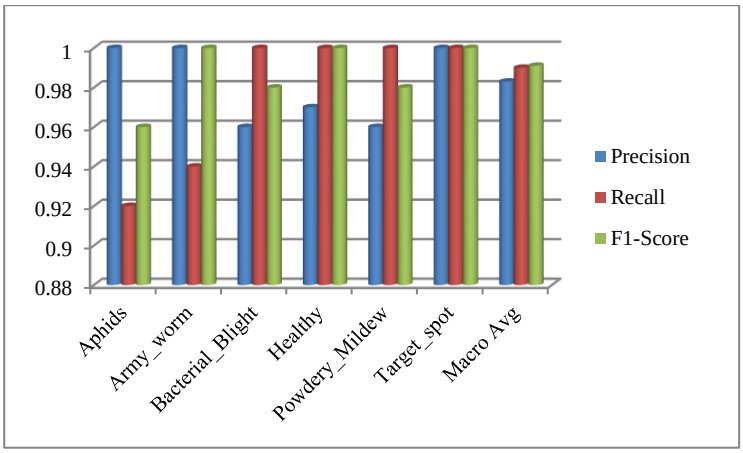

Fig.17 Performance evaluation of hybrid model

Figure 17 evidently, shows the hybrid model was successful in utilizing the hyper-parameters to achieve the best performance among all of the other classifiers by correctly classifying 142 out of 144 with an accuracy of 98.6%. The classes themselves had ideal precision and recall (1.00) for Army_worm, Healthy, Target_spot meaning 0 false positives and false negatives. Aphids had 1 FN and 1 FP, thus having only 1.0 -0.08 = 0.92 recall or, 1.0-(1.0-0.92) = 0.96 F1-score. Both Bacterial_Blight and Powdery_Mildew had F1's of 0.98 since they misclassified only slightly. The macro-average for all of the key measures (precision, recall, F1-score) was 0.99 thus renders the XGBoost the most effective machine-learning model with only 2 errors overall showing very high generalization, precision, and robustness for plant disease classification.

4.4 Comparative Analysis

- **Comparison among proposed models**

The performance results of all three models—SVM, XGBoost, and Hybrid—on the categorization of plant diseases are included in the performance comparison table 5 and figure 18, which also gives an overall sense of how effectively the models are accomplishing the desired objectives. With the highest values for every parameter, the Hybrid model was unquestionably better than all the other models. It achieved an accuracy of 98.6%, precision of 98.3%, recall of 99.0%, and F1-score of 99.1%, demonstrating exceptional consistency and low misclassification when identifying plant disease. With a 94.4% accuracy rate, XGBoost also demonstrates a respectable level of competency with the target goals and achieves respectable precision and recall values on its own, demonstrating the model's great overall performance on its own. The SVM model is good with all of its individual metrics extending to decent standards (0.882 accuracy).

TABLE V

PERFORMANCE COMPARISON OF SVM, XGBOOST, AND HYBRID MODELS

Metric	SVM	XGBoost	Hybrid model
Accuracy	0.882	0.944	0.986
Precision	0.88	0.93	0.983
Recall	0.88	0.925	0.99
F1-score	0.87	0.94	0.991

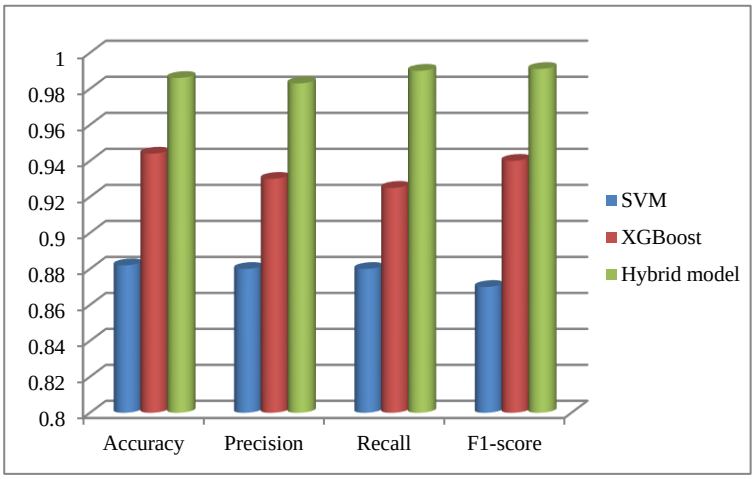

Fig.18 Comparison among proposed models

- **A comparison with current models**

A comparison of various contemporary methods for classifying plant diseases in terms of accuracy, precision, recall, and F1-score of model outputs is shown in Table 6 and Figure 19. The CNN by Senan et al. (2021) has the more balanced set of accuracies among good-performance models across all metrics, which is limited to Lamba et al. (2023) and Kannandasan et al. (2022). MaizeNet (Sharma et al., 2022) exhibits the highest accuracy and precision compared to the current works. The suggested Hybrid model provides the best performance values among all the erstwhile mentioned works of literature based on model in terms of accuracy (98.6%), recall (99.0%), and F1-score (99.1%). This indicates that the hybrid model is most capable of generalization among the erstwhile mentioned works, and being the strongest model for multi-crop disease classification.

TABLE VI

COMPARISON OF PROPOSED HYBRID MODEL WITH EXISTING MODELS ON CLASSIFICATION PERFORMANCE METRICS

Author [Reference]	Methodology	Accuracy	Precision	Recall	F1-sore
Lamba et al., (2023) [42]	CNN+SVM	97	95.67	94.8	97
Sharmaet al., (2022)[43]	MaizeNet	98.50	98.87	95.31	97.13
Senanet al., (2021)[44]	CNN	97.66	98	98	98
Kannadasan et al., (2022)[45]	RF+RFECV	92.6	92.45	90.1	91.44
Proposed Work	Hybrid model	98.6	98.3	99	99.1

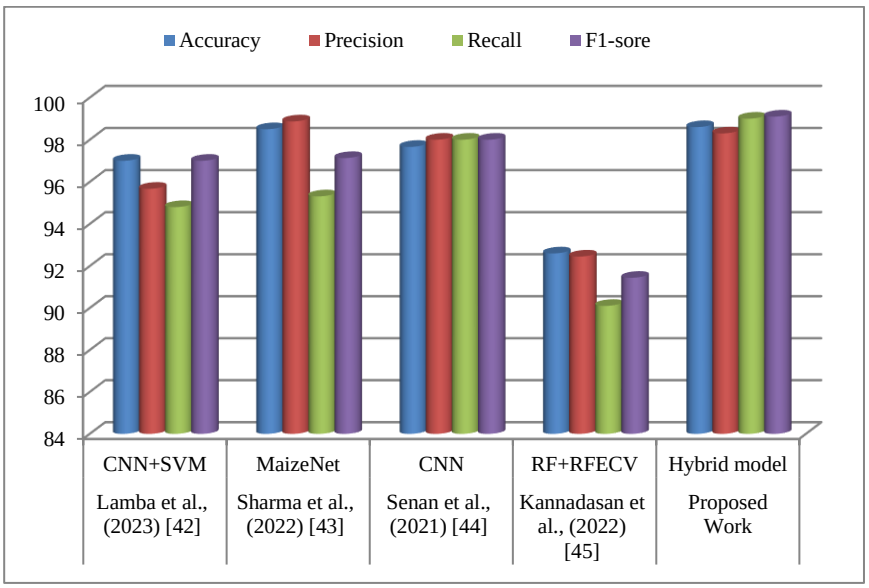

Fig. 19 Comparison with state of art model

V. CONCLUSION

This study suggests a novel hybrid ML framework to determine correct crop type, disease class and severity prediction using leaf image data. The hybrid model consists of XGBoost and SVM in which XGBoost strongly extracts useful features and SVM can train a model with multi-dimensional high volume data. According to the experimental results with the data set of Plant Village, utilizing the major consumer crops (potato, tomato, corn and grape), the suggested model surpasses the conventional methods, and the hybrid models performed better than both the SVM and XGBoost when employed separately. According to the classification accuracy, the hybrid model attained a good accuracy of 98.6%. With precision of 98.30%, recall of 99.0% and F1-score of 99.1%, all while out-performing the models SVM and XGBoost. When compared to previous study such as MaizeNet (98.5%) and CNN based models ~97% we see that the hybrid model seems to offer improved consistency and generalizability across various disease types and crop classes. Improved Deep Joint Segmentation and the use of them also aid in localizing disease region to accurately predict severity. The findings provide evidence and justification for the models potential application towards actual use in precision agriculture by enabling early detection, timely intervention and sustainable crop management to contain losses as well as boost yields.

REFERENCES

[1] Standard, Business. "Available online: https://www. Business-standard.com/article/economypolicy/govt-assessing-economic-impact-of-second-covid-wave-for-possible-stimulus-121052501320_1. Accessed on July 27 (2021).

[2] Ristaino, Jean B., Pamela K. Anderson, Daniel P. Bebber, Kate A. Brauman, Nik J. Cunniffe, Nina V. Fedoroff, Cambria Finegold et al. "The persistent threat of emerging plant disease pandemics to global food security." Proceedings of the National Academy of Sciences 118, no. 23 (2021): e2022239118.

[3] Singh, Brajesh K., Manuel Delgado-Baquerizo, EleonoraEgidi, Emilio Guirado, Jan E. Leach, Hongwei Liu, and PankajTrivedi. "Climate change impacts on plant pathogens, food security and paths forward." Nature Reviews Microbiology 21, no. 10 (2023): 640-656.

[4] Zhao, Hengqian, Chenghai Yang, Wei Guo, Lifu Zhang, and Dongyan Zhang. "Automatic estimation of crop disease severity levels based on vegetation index normalization." Remote Sensing 12, no. 12 (2020): 1930.

[5] Parikh, Aditya, Mehul S. Raval, ChandrasinhParmar, and Sanjay Chaudhary. "Disease detection and severity estimation in cotton plant from unconstrained images." In 2016 IEEE international conference on data science and advanced analytics (DSAA), pp. 594-601. IEEE, 2016.

[6] Gupta, Meenu, Rakesh Kumar, and Divya Singh. "Crop Diseases Detection and Prevention Using AI and Machine Learning Techniques." In Digital Twin Technology, pp. 61-82. CRC Press, 2021.

[7] Shin, Jaemyung, Md Sultan Mahmud, Tanzeel U. Rehman, PrabaharRavichandran, Brandon Heung, and Young K. Chang. "Trends and prospect of machine vision technology for stresses and diseases detection in precision agriculture." AgriEngineering 5, no. 1 (2022): 20-39.

[8] Tian, Hongkun, Tianhai Wang, Yadong Liu, Xi Qiao, and Yanzhou Li. "Computer vision technology in agricultural automation—A review." Information processing in agriculture 7, no. 1 (2020): 1-19.

[9] Khotimah, WijayantiNurul, Mohammed Bennamoun, FaridBoussaid, LianXu, David Edwards, and FerdousSohel. "MCE-ST: Classifying crop stress using hyperspectral data with a multiscale conformer encoder and spectral-based tokens." International Journal of Applied Earth Observation and Geoinformation 118 (2023): 103286.

[10] Sethy, Prabira Kumar, NaliniKantaBarpanda, Amiya Kumar Rath, and SantiKumariBehera. "Deep feature based rice leaf disease identification using support vector machine." Computers and Electronics in Agriculture 175 (2020): 105527.

[11] Kok, Zhi Hong, Abdul Rashid Mohamed Shariff, Meftah Salem M. Alfatni, and SitiKhairunniza-Bejo. "Support vector machine in precision agriculture: A review." Computers and Electronics in Agriculture 191 (2021): 106546.

[12] Ge, Jiankun, Linfeng Zhao, Zihui Yu, Huanhuan Liu, Lei Zhang, Xuewen Gong, and Huaiwei Sun. "Prediction of greenhouse tomato crop evapotranspiration using XGBoost machine learning model." Plants 11, no. 15 (2022): 1923.

[13] Piska, Ganga, SwaraliJanaskar, OjasChandgadkar, ParasBhawsar, and PinkiPrakashVishwakarma. "Enhancing plant disease detection using machine learning approaches for improved agricultural productivity." International Journal of Electrical & Computer Engineering (2088-8708) 15, no. 2 (2025).

[14] Vijayan, S., and ChiranjiLalChowdhary. "Hybrid feature optimized CNN for rice crop disease prediction." Scientific Reports 15, no. 1 (2025): 7904.

[15] Praveen, R. V. S., Mudit Mittal, PrasantaParida, Rowsonara Begum, Yogendra Kumar, and GinniNijhawan. "Deep Learning Applications for Detecting Crop Diseases from Image Data." In 2025 International Conference on Computational, Communication and Information Technology (ICCCIT), pp. 566-571. IEEE, 2025.

[16] Bouacida, Imane, BrahimFarou, Lynda Djakhdjakha, Hamid Seridi, and MuhammetKurulay. "Innovative deep learning approach for cross-crop plant disease detection: A generalized method for identifying unhealthy leaves." Information Processing in Agriculture 12, no. 1 (2025): 54-67.

[17] Kumar, Rahul, Ashok Kumar, Karamjit Bhatia, KottakkaranSooppyNisar, Siddharth Singh Chouhan, Priti Maratha, and Anoop Kumar Tiwari. "Hybrid approach of cotton disease detection for enhanced crop health and yield." IEEE Access (2024).

[18] Barman, Shohag, Fahmid Al Farid, JaoharRaihan, Niaz Ashraf Khan, MdFerdous Bin Hafiz, Aditi Bhattacharya, Zaeed Mahmud et al. "Optimized crop disease identification in Bangladesh: A deep learning and SVM hybrid model for rice, potato, and corn." Journal of Imaging 10, no. 8 (2024): 183.

[19] Elmasry, Ahmed, Walid Abdullah, Byeong-Gwon Kang, and Yunyoung Nam. "A novel hybrid approach based on CNN for corn diseases detection." Optimization in Agriculture 1 (2024): 94-104.

[20] Rangarajan, AravindKrishnaswamy, Raja Purushothaman, MaheswariPrabhakar, and CezarySzczepański. "Crop identification and disease classification using traditional machine learning and deep learning approaches." Journal of Engineering Research 11, no. 1B (2023).

[21] Singh, R. N., Prameela Krishnan, Vaibhav K. Singh, and Bappa Das. "Estimation of yellow rust severity in wheat using visible and thermal imaging coupled with machine learning models." Geocarto International 38, no. 1 (2023): 2160831.

[22] Sudha, P., and P. Kumaran. "Early detection and control of anthracnose disease in cashew leaves to improve crop yield using image processing and machine learning techniques." Signal, Image and Video Processing 17, no. 7 (2023): 3323-3330.

[23] Nandhini, S., and K. Ashokkumar. "Machine learning technique for crop disease prediction through crop leaf image." Appl. Math. Inf. Sci 16, no. 2 (2022): 149-158.

[24] Haque, MdAshraful, SudeepMarwaha, Chandan Kumar Deb, Sapna Nigam, AlkaArora, Karambir Singh Hooda, P. Lakshmi Soujanya et al. "Deep learning-based approach for identification of diseases of maize crop." Scientific reports 12, no. 1 (2022): 6334.

[25] Khan, Habib, IjazUlHaq, Muhammad Munsif, Mustaqeem, ShafiUllah Khan, and Mi Young Lee. "Automated wheat diseases classification framework using advanced machine learning technique." Agriculture 12, no. 8 (2022): 1226.

[26] Kundu, Nidhi, Geeta Rani, Vijaypal Singh Dhaka, Kalpit Gupta, Siddaiah Chandra Nayak, SahilVerma, Muhammad FazalIjaz, and MarcinWoźniak. "IoT and interpretable machine learning based framework for disease prediction in pearl millet." Sensors 21, no. 16 (2021): 5386.

[27] Clohessy, James W., SantoshSanjel, G. Kelly O'Brien, Rebecca Barocco, Shivendra Kumar, Scott Adkins, Barry Tillman, David L. Wright, and Ian M. Small. "Development of a high-throughput plant disease symptom severity assessment tool using machine learning image analysis and integrated geolocation." Computers and Electronics in Agriculture 184 (2021): 106089.

[28] Lamba, Monika, YogitaGigras, and AnuradhaDhull. "Classification of plant diseases using machine and deep learning." Open Computer Science 11, no. 1 (2021): 491-508.

[29] https://www.kaggle.com/datasets/abdallahalidev/plantvillage-dataset.

[30] Yang, Hui. "Data preprocessing." Pennsylvania State University: Citeseer (2018).

[31] Petrellis, Nikos. "Plant disease diagnosis with color normalization." In 2019 8th International Conference on Modern Circuits and Systems Technologies (MOCAST), pp. 1-4. IEEE, 2019.

[32] Liu, Jiachun, Shuqin Yang, Yunling Cheng, and Zhishuang Song. "Plant leaf classification based on deep learning." In 2018 Chinese Automation Congress (CAC), pp. 3165-3169. IEEE, 2018

[33] Singh, Vijai, and Ak K. Misra. "Detection of plant leaf diseases using image segmentation and soft computing techniques." Information processing in Agriculture 4, no. 1 (2017): 41-49.

[34] Salini, R., G. CharlynPushpaLatha, and RashmitaKhilar. "Plant disease detection with modified deep joint segmentation and combined Google Net-IRNN." Journal of Phytopathology 172, no. 3 (2024): e13313.

[35] Anitha, K., and S. Srinivasan. "Feature Extraction and Classification of Plant Leaf Diseases Using Deep Learning Techniques." Computers, Materials & Continua 73, no. 1 (2022).

[36] Kavzoglu, Taskin, and AlihanTeke. "Predictive performances of ensemble machine learning algorithms in landslide susceptibility mapping using random forest, extreme gradient boosting (XGBoost) and natural gradient boosting (NGBoost)." Arabian Journal for Science and Engineering 47, no. 6 (2022): 7367-7385.

[37] Chen, Tianqi, Tong He, Michael Benesty, VadimKhotilovich, Yuan Tang, Hyunsu Cho, Kailong Chen, Rory Mitchell, Ignacio Cano, and Tianyi Zhou. "Xgboost: extreme gradient boosting." R package version 0.4-2 1, no. 4 (2015): 1-4.

[38] Chang, Sheng, Zelong Chi, Hong Chen, Tongle Hu, CaixiaGao, JihuaMeng, and Liangxiu Han. "Development of a MultiscaleXGBoost-based Model for Enhanced Detection of Potato Late Blight Using Sentinel-2, UAV, and Ground Data." IEEE Transactions on Geoscience and Remote Sensing (2024).

[39] Jakkula, Vikramaditya. "Tutorial on support vector machine (svm)." School of EECS, Washington State University 37, no. 2.5 (2006): 3.

[40] Padol, Pranjali B., and Anjali A. Yadav. "SVM classifier based grape leaf disease detection." In 2016 Conference on advances in signal processing (CASP), pp. 175-179. IEEE, 2016.

[41] G. Ciaburro, B. Venkateswaran Neural Netw. World R, 91 (2017).

[42] Lamba, Shweta, VinayKukreja, AnupamBaliyan, Shalli Rani, and Syed Hassan Ahmed. "A novel hybrid severity prediction model for blast paddy disease using machine learning." Sustainability 15, no. 2 (2023): 1502.

[43] Kundu, Nidhi, Geeta Rani, Vijaypal Singh Dhaka, Kalpit Gupta, Siddaiah Chandra Nayaka, Eugenio Vocaturo, and Ester Zumpano. "Disease detection, severity prediction, and crop loss estimation in MaizeCrop using deep learning." Artificial intelligence in agriculture 6 (2022): 276-291.

[44] Sharma, Rahul, Amar Singh, N. Z. Jhanjhi, MehediMasud, Emad Sami Jaha, and SahilVerma. "Plant Disease Diagnosis and Image Classification Using Deep Learning." Computers, Materials & Continua 71, no. 2 (2022).

[45] Kannadasan, Raju, Mohammed H. Alsharif, Abu Jahid, and Muhammad Asghar Khan. "Categorizing diseases from leaf images using a hybrid learning model." Symmetry 13, no. 11 (2021): 2073.

The Quantum Green Shield: Integrating Computing Principles in Cyber Defence

Abhisekh Tiwari[1], Dr. Pooja Tripathi[2]

Computer Science and Engineering, Inderprastha Engineering College, Uttar Pradesh, India[1]
Information Technology, Inderprastha Engineering College, Uttar Pradesh, India[2]

abhishektripathi0706@gmail.com[1], pooja.tripathi@ipec.org.in[2]

Abstract: Current digital networks face more advanced cybersecurity threats because interconnected digital systems that comprise cyber-physical systems and cloud computing and the Internet of Things (IoT) have become widely adopted. The traditional cybersecurity systems which depend on basic computer techniques and process data through one thread cannot protect systems against ongoing advanced threats and developing malware and newly discovered system vulnerabilities. The Quantum-Enhanced Cyber Defence System (QE-CDS) develops its cyber defence system by combining quantum computing with artificial intelligence and machine learning and large language models. The suggested framework employs Quantum Machine Learning (QML) methods through its Quantum Support Vector Machines (QSVM) and Quantum Neural Networks (QNN) and Quantum Recurrent Neural Networks (QRNN) and Quantum Approximate Optimization Algorithms (QAOA). Scientists conducted their research through a hybrid quantum-classical simulation environment which used benchmark datasets like CICIDS2017 and UNSW-NB15 and Bot-IoT. Our system outperforms standard systems in detecting threats because it identifies threats more quickly and makes superior defensive choices and finishes detection work with less than standard false positive rates and operates more effectively in large environments. The security analysis of unstructured data becomes more precise through the integration of Large Language Models with cybersecurity systems that use quantum computing technology. The simulation results show that future security systems will be developed through hybrid systems which combine quantum technology with artificial intelligence because these systems can adapt to emerging threats while defending against quantum attacks.

Keywords:Quantum Computing, Cybersecurity, Quantum Machine Learning, Intrusion Detection Systems, Large Language Models, Hybrid Quantum–Classical Architecture

I. INTRODUCTION

The digital revolution of modern civilization has created a situation where people now use extensive interconnected networks together with cloud computing services and cyber-physical systems and Internet of Things (IoT) technologies. The system updates which boost operational efficiency and system interconnectivity also create new security threats because they establish fresh access points which hackers can exploit. The current cyber threat arsenal includes advanced persistent threats (APTs) and polymorphic malware and zero-day vulnerabilities and synchronized attack operations and multiple adaptable attack methods. The existing cybersecurity systems which use traditional computing hardware face challenges because they cannot handle the increasing complexity of security threats.

Classic computer systems cannot process large amounts of different input data at the same time because they depend on sequential processing and use binary logic. Classical cybersecurity systems still experience performance issues because they must handle complex security threats while AI and ML technologies have been implemented. Organizations can protect themselves from potential harm by studying future attack methods and constructing security systems which can handle multiple threat scenarios.

The new computing paradigm of quantum computing creates multiple pathways to solve various problems. The operation of quantum computers depends on qubits which process information through superposition and entanglement and quantum interference while traditional computers use bits to represent two distinct binary states. Quantum computers utilize their features to execute data processing across multiple state spaces, which enables them to solve particular problems with enhanced computational power. The process of quantum computing provides its benefits to cybersecurity applications which require solutions that can handle large optimization problems and recognize exact patterns while they process intricate data that contains multiple dimensions.

http://doi.org/10.63766/spujstmr.26.000073

Quantum Machine Learning (QML) serves as a potential future solution for cybersecurity protection systems, which results from the integration of quantum computing with artificial intelligence and machine learning. Theoretical and actual studies of QSVM QNN and QAOA show that organizations can use quantum technology augmented algorithms to achieve faster decision-making processes and better detection methods for anomalies and classification. This algorithm enables cybersecurity intrusion detection systems to detect attacks with higher accuracy while decreasing false positive rates and enabling companies to update their security measures in response to evolving threats.

Your study establishes that unstructured cybersecurity data which includes system logs and network traffic records and vulnerability reports and threat intelligence feeds constitutes essential information for current security operations. LLMs gain autonomous textual analysis abilities through their integration with quantum computing AI systems which enable them to create danger descriptions in context and develop threat prediction models. Cyber defence systems are being created through the combination of quantum computing technology and artificial intelligence and natural language processing. The systems operate independently while acquiring knowledge from their previous errors.

Quantum computing technology creates a major danger to current encryption standards. Quantum algorithms possess the ability to solve both integer factorization and unstructured search problems, which enables them to use encryption methods for digital data protection. The development of security systems that can withstand quantum attacks creates a dual-use problem because quantum technologies make traditional encryption methods obsolete. Future cybersecurity studies will primarily focus on developing post-quantum cryptographic standards and quantum-safe encryption.

II. LITERATURE REVIEW

The table presents a summary of key scholarly publications about quantum computing, artificial intelligence, and cybersecurity which researchers published between 2023 and 2025. The table displays research activities through a selection of twenty journal articles book chapters and conference proceedings. The study investigates various aspects of quantum-based cybersecurity protection through its research which includes adaptive defense optimization and intelligent threat detection using big language models and development of quantum-safe security systems and quantum machine learning-based intrusion detection methods. The synthesis shows that researchers are putting more effort into hybrid quantum-classical systems and AI-powered autonomous cyber defense systems while facing challenges which prevent them from creating integrated systems that can explain their operations and expand their functions and handle real-time tasks. Researchers want to build complete cyber defense systems which use quantum technology to protect against future threats. The table organizes current research about quantum cybersecurity while showing research obstacles which exist in an organized format.

Reviewer Update: The literature review has been revised to emphasize recent studies published between 2023 and 2025 and to systematically identify research gaps including lack of explainability, limited real-time validation, scalability constraints, insufficient integration of LLMs with QML frameworks, and absence of sustainability-aware cybersecurity architectures.

TABLE I
LITERATURE REVIEW

Ref	Author(s)	Year	Focus / Approach	Research Gap Identified
[1]	Ajimon& Kumar	2024	Explores integration of Large Language Models (LLMs) with quantum-aware cybersecurity for real-time anomaly detection and threat intelligence generation	Lacks comprehensive end-to-end system validation and integration with optimization mechanisms
[2]	Azeez et al.	2024	Implements Quantum Machine Learning models (QSVM, QNN) for anomaly detection in cybersecurity systems	Focuses mainly on detection accuracy; lacks adaptive optimization and explainability features
[3]	Thirupathi et al.	2024	Proposes a system-level AI–quantum cybersecurity architecture utilizing QAOA-based optimization techniques	Does not address real-time threat intelligence generation or linguistic analysis capabilities
[4]	Hossain et al.	2024	Introduces a conceptual QML framework using QRNN for temporal cyber threat prediction	Remains theoretical; lacks empirical validation and scalability assessment
[5]	Deepajothi et al.	2024	Develops a hybrid quantum–classical architecture using QNN	Limited coverage of temporal attack patterns and absence of

			for intrusion detection	autonomous feedback mechanisms
[6]	Shamoo	2024	Proposes a quantum-aware cognitive defense system integrating LLMs with quantum inference	Primarily conceptual with limited benchmarking and performance evaluation
[7]	Ramya et al.	2024	Applies Quantum Natural Language Processing (QNLP) for semantic-based cyber threat detection	Focuses only on textual data; lacks integration with network-level threat detection systems
[8]	Alluhaibi	2024	Designs a quantum-enhanced Intrusion Detection System (IDS) using QNN and quantum entropy filtering	Does not incorporate optimization strategies, temporal learning, or linguistic intelligence
[9]	Shaik et al.	2024	Utilizes Quantum Boltzmann Machines (QBM) and Quantum Reinforcement Learning (QRL) for network security	Lacks explainability and integration with LLM-based reporting frameworks
[10]	Albtosh	2024	Proposes a quantum–linguistic framework for cyber threat reasoning and incident response	Focused mainly on response mechanisms; limited evaluation in intrusion detection scenarios
[11]	Rosa-Remedios & Caballero-Gil	2024	Develops a proactive QML optimization framework using Variational Quantum Circuits (VQC) and reinforcement learning	Primarily optimization-focused; lacks full cybersecurity lifecycle integration
[12]	Abreu et al.	2024	Implements QML-based IDS using QSVM and quantum clustering methods	Does not include adaptive learning capabilities or semantic threat interpretation
[13]	Tehrani et al.	2024	Proposes a quantum-enhanced SIEM system for botnet detection	Application-specific; lacks generalization across diverse cyber threat scenarios
[14]	Knaapen et al.	2024	Develops a quantum entropy-based framework for ransomware detection	Specialized for ransomware; lacks broader IDS applicability
[15]	Yadav et al.	2023	Explores quantum-enhanced deep learning models for cybersecurity applications	Focuses on model performance without system-level integration
[16]	Payares & Martínez-Santos	2023	Provides a survey of Quantum Machine Learning-based intrusion detection systems	Reviews existing methods but does not propose a unified framework
[17]	KiyatakeNicesio et al.	2023	Conducts a systematic review of Quantum Network Intrusion Detection Systems (NIDS)	Identifies benchmarking issues but lacks practical implementation guidance
[18]	Aggarwal et al.	2024	Examines quantum computing architectures and performance optimization strategies	General focus; limited validation in cybersecurity-specific applications
[19]	Divyashree	2024	Studies quantum-safe cybersecurity and post-quantum cryptographic techniques	Emphasizes encryption; limited focus on detection and response mechanisms
[20]	Hussain et al.	2021	Reviews integration of AI and quantum computing in cybersecurity	Early-stage work; lacks incorporation of modern QML and LLM advancements

The research investigations documented in Table 1 focus primarily on three areas which include cryptographic security and intrusion detection systems and adaptive cyber defense systems. The systems function through the implementation of

quantum computing together with quantum machine learning. The existing body of research has examined design elements which include sustainability and energy efficiency and environmental responsibility. The research work conducted by Hussain et al. [20] and Aggarwal et al. [18] introduced architectural efficiency and scalability solutions which require cybersecurity design to include sustainability metrics. Technical studies by Azeez et al. [2], Alluhaibi [8], Abreu et al. [12], and Shaik et al. [9] disregard environmental impact in favor of detecting accuracy, optimization, and system performance. The research study from Ajimon and Kumar [1] and Shamoo [6] and Albtosh [10] through their investigation of autonomous defense together with interpretability research established a link between Large Language Models and quantum intelligence. They fail to evaluate their results through the implementation of energy-efficient computing metrics and green computing standards. The existing knowledge gap demonstrates that researchers regard cybersecurity resilience and sustainability together with quantum efficiency as distinct research domains. The Quantum Green Shield framework will establish a knowledge foundation through its integration of quantum-based cyber defense systems with sustainable design elements into an all-encompassing future-ready cybersecurity framework.

Objective of the study

1. The research studies how Quantum Computing and Artificial Intelligence research and Machine Learning research work together to create solutions for cybersecurity problems. The research demonstrates how quantum algorithms such as QSVM and QNN and QAOA improve cybersecurity threat detection and classification and prevention capabilities.
2. The research evaluates Quantum Machine Learning QML models for threat detection and intrusion prevention because these models provide better computational performance than traditional machine learning algorithms. The study investigates the application of Quantum Amplitude Encoding QAE and quantum feature maps for enhancing detection accuracy while decreasing false positives in intrusion detection systems IDS.
3. The research studies how Quantum Computing technology helps to build better cyber securitydefence systems through QAOA algorithms which help Security Information and Event Management systems to handle resources effectively while decreasing false alarms and making better decisions.
4. The study examines how Large Language Models LLMs combine with Quantum Computing to automatically analyze unstructured cyber security data which includes network logs and system alerts and natural language threat reports.

III. RESEARCH METHODOLOGY

The study uses experimental and comparative research methods to assess the performance of the Quantum-Enhanced Cyber defence System (QE-CDS) against traditional AI and ML cybersecurity systems. The research investigates how hybrid quantum-classical systems and quantum-inspired systems impact intrusion detection performance in real-world cybersecurity environments by assessing their ability to detect threats and their system performance metrics which include latency and accuracy and scalability and temporal learning and defensive optimization.

Experimental Environment

The researchers conducted their studies using quantum simulation frameworks because they could not access scalable quantum hardware which had not been released to the market. Researchers can perform multiple tests using the device which replicates essential functions of quantum computing**.**

- Quantum Frameworks: Qiskit, PennyLane, TensorFlow Quantum
- Quantum Backend: IBM QASM Simulator (32-qubit emulation)
- Classical Infrastructure: 64-core Intel Xeon Gold CPU, 128 GB RAM, Ubuntu 22.04

This configuration ensures stability and reproducibility without hardware-induced variability.

Datasets and Preprocessing

The three benchmark datasets CICIDS2017 and UNSW-NB15 and Bot-IoT were used to evaluate IoT attack scenarios and enterprise network security. The preparation procedure for quantum encoding required data cleaning and feature normalization and scaling to eliminate all system noise. The processed feature vectors underwent conversion into quantum-compatible representations through the Quantum Amplitude Encoding QAE technique which enables compact representation of high-dimensional data.

Model Implementation

Model Implementation

- **Classical:** Support Vector Machine (SVM), Deep Neural Network (DNN), LSTM, Bi-LSTM
- **Quantum:** Quantum Support Vector Machine (QSVM), Quantum Neural Network (QNN), Quantum Recurrent Neural Network (QRNN)
- **Hybrid:** QSVM–QNN
- **Optimization:** QAOA-inspired iterative optimization
- **Threat Intelligence:** Quantum-inspired Large Language Model (Q-GPT)

The researchers used parameterized quantum circuits together with variational learning to implement quantum models without any restrictions. The hybrid models used quantum feature extraction together with classical inference to create their system.

Evaluation Strategy

Our research first focused on optimizing QAOA defence systems. We implemented QRNNs for real-time threat detection. The system's capacity was evaluated through tests that used progressively larger dataset samples. We employed quantum-inspired LLMs to conduct semantic assessments of security threats. The test cases needed five distinct areas to create their base structure. The assessment used these performance measures: recall accuracy precision F1-score false positive rate detection latency optimization cost speed of temporal adaptability scaling behavior contextual correlation.

Baseline and Limitations

The researchers established baseline results through standard modelling methods before they began testing quantum and hybrid systems. Researchers conducted simulated tests to improve computing capacity by decreasing their feature set and adding new elements which included untested physical quantum noise and decoherence. The method demonstrates a workable framework which enables scientists to conduct authentic evaluations of upcoming quantum security technologies despite current limitations.

IV. RESULTS AND ANALYSIS

The results section of this research shows the outcomes from tests which assessed the performance of the Quantum-Enhanced Cyber defence System. The study examines six specific traits which include semantic threat intelligence and scalability and latency and detection accuracy and optimization efficiency and temporal flexibility. The research compares the performance of hybrid and quantum-enhanced models to that of conventional machine learning baseline models which creates a complete evaluation of both model types.

A. Experimental Setup

The research studies employed a hybrid quantum-classical simulation environment because they required access to operational large-scale fault-tolerant quantum systems. Researchers developed quantum simulations that worked with IBM QASM Simulator through Qiskit version 0.46.0 and PennyLane version 0.37.1. The conventional orchestration system operated on a system that used 128 GB RAM and 64-core Intel Xeon Gold CPU.

The assessment process utilized three benchmark datasets which included CICIDS2017 and UNSW-NB15 and Bot-IoT. The data preparation stage consisted of three procedures which included data cleaning and data normalization and data dimensionality reduction. The researchers used Quantum Amplitude Encoding QAE to convert standard feature vectors into quantum states. The study team used SVM DNN Random Forest and LSTM/Bi-LSTM as their baseline for conventional analysis. The team studied four quantum algorithms which included QSVM QNN QRNN and QAOA.

The system employs multiple performance assessment metrics to evaluate its operational efficiency. The system measures eight different metrics which include precision and accuracy and recall and F1-score and false positive rate and detection latency and scalability and optimization cost and false negative rate.

B. Test Case 1: Detection Accuracy and Latency

The table presents the testing results for three different detection models which include classical quantum and hybrid models. The standard SVM and DNN models produced low accuracy results which proved their failure to detect certain types of invasions. The QSVM model achieved an 86.67% accuracy through its feature mapping system which operates in high-dimensional Hilbert space and uses quantum kernel technology.

TABLE II
ACCURACY AND DETECTION LATENCY COMPARISON OF IMPLEMENTED MODELS

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	Detection Latency (ms/sample)
Classical SVM (RBF)	74.61	91.86	60.79	73.16	0.711
Deep Neural Network (DNN)	77.23	91.93	65.78	76.68	0.0014
QSVM (Quantum Kernel)	86.67	94.12	80.00	86.49	0.0019
QNN	80.00	100.00	60.00	75.00	5.997
Hybrid QSVM–QNN	**89.00**	**95.24**	**81.63**	**87.91**	**0.0035**

The latency values reported in Table 2 were obtained in a controlled quantum-classical simulation environment. Values such as 0.0014 ms/sample for DNN represent model inference kernel execution time under simulator conditions and do not include network, I/O, operating-system, or deployment overheads. Therefore, these measurements should be interpreted as relative benchmarking indicators rather than real-world end-to-end latency estimates. The QNN needed more time to analyse data, but it achieved high precision through repeated circuit evaluations.

C. Test Case 2: QAOA-Based defence Optimization

Table III shows the results of optimization research which tested QAOA methods to improve detection threshold setting and resource distribution. The True Detection Rate reached 96.06% when optimization depth increased while the False Positive Rate dropped from 3.05% (p=1) to 0.32% (p=10). The system showed reliable convergence because optimization expenses kept decreasing at a steady rate.

TABLE III
QAOA-INSPIRED OPTIMIZATION RESULTS

Iterations (p)	False Positive Rate (%)	True Detection Rate (%)	Optimization Cost
1	3.05	84.66	0.080
3	1.85	89.30	0.054
5	1.12	92.53	0.037
7	0.68	94.79	0.025
10	0.32	96.96	0.014

The findings show that quantum-based optimization techniques achieve better results in solving nonlinear cybersecurity decision problems which involve large data sets than traditional static methods and heuristic techniques.

D. Test Case 3: Temporal Threat Detection

The researchers used QRNN to examine LSTM and Bi-LSTM, two basic models, for temporal learning capabilities. The researchers evaluated these models against emerging security threats which included adaptive phishing campaigns and polymorphic ransomware. The results showed that QRNN succeeded in handling long-term dependencies because it achieved an anomalous recall rate of 94.51% which almost matched the 94.59% performance of Bi-LSTM

TABLE IV
TEMPORAL THREAT DETECTION PERFORMANCE

Model	Temporal Adaptation Speed (ms)	Anomaly Recall (%)	False Negatives (%)	Resource Usage (Relative)
LSTM	0.00131	91.31	8.69	1.0×

Bi-LSTM	0.00303	94.59	5.41	2.0×
QRNN (Simulated)	0.01548	94.51	5.49	2.09×

The integration of simulated quantum processes increased QRNN's latency yet the system demonstrated greater defence capabilities against cyber security attacks which occurred during late phases because it could accurately simulate different types of cyber threats.

E. Test Case 4: Scalability Analysis

Table V displays the scalability findings for QNN. The accuracy maintained its 73% to 76% range during the dataset expansion from 10,000 records to 100,000 records while the training duration displayed nearly linear increase. The system shows controlled scalability which demonstrates that quantum-inspired learning capabilities provide benefits to large-scale cybersecurity analytics better than traditional classical models that fail under increasing data loads.

TABLE V
SCALABILITY PERFORMANCE OF QUANTUM-INSPIRED NEURAL NETWORK

Dataset Size (Records)	**QNN Accuracy (%)**	**Training Time (s)**
10,000	73.76	0.39
25,000	73.47	0.98
50,000	75.58	1.93
100,000	74.70	3.89

F. Test Case 5: Quantum LLM-Based Threat Intelligence

Scientists conducted tests to compare the semantic threat interpretation skills of Quantum-Inspired LLM (Q-GPT) against standard LLM. Q-GPT presented better threat-related information connection abilities because it could more effectively separate dangerous content components even though both systems reached similar text understanding accuracy. The system experienced a minor increase in inference time while its operational performance remained stable. The study results show that language models which use quantum mechanics help improve threat intelligence and analyst decision support. The models do not improve processing speed according to their basic functions.

TABLE VI
QUANTUM LLM INTEGRATION FOR THREAT INTELLIGENCE (EXPERIMENTAL RESULTS)

Metric	**Classical LLM (GPT-based)**	**Quantum LLM (Q-GPT)**
Text Understanding Accuracy (%)	87.40	87.81
Threat Context Correlation (%)	14.99	16.89
Report Generation Time (s)	0.18	0.21
Data Comprehension Latency (ms)	0.002	0.003

G. Comparative Analysis and Discussion

Figures **1–3** summarize the comparative performance. Quantum-enhanced and hybrid models consistently outperform classical baselines in accuracy, recall, adaptability, and optimization efficiency. Key observations include:

- **Accuracy improvement** up to 14% over classical SVM
- **Latency reduction** enabling near real-time detection
- **Scalable learning behaviour** under increasing data volume
- **Effective adaptive optimization** using QAOA
- **Enhanced temporal intelligence** via QRNN
- **Improved semantic interpretation** using quantum-inspired LLMs

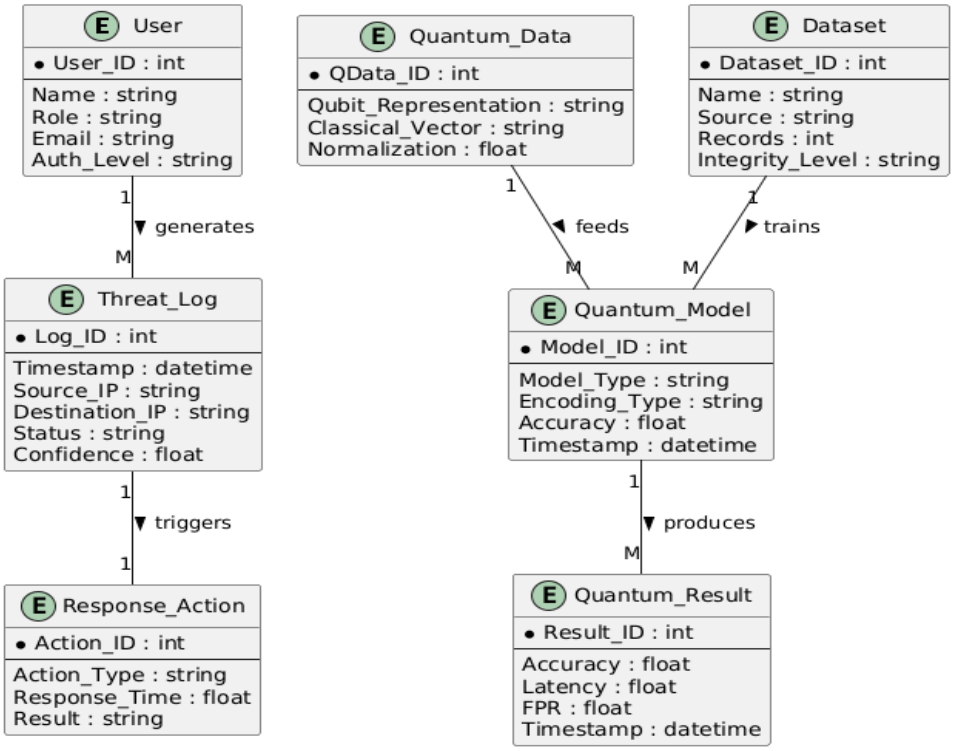

Fig. 1. ER Diagram

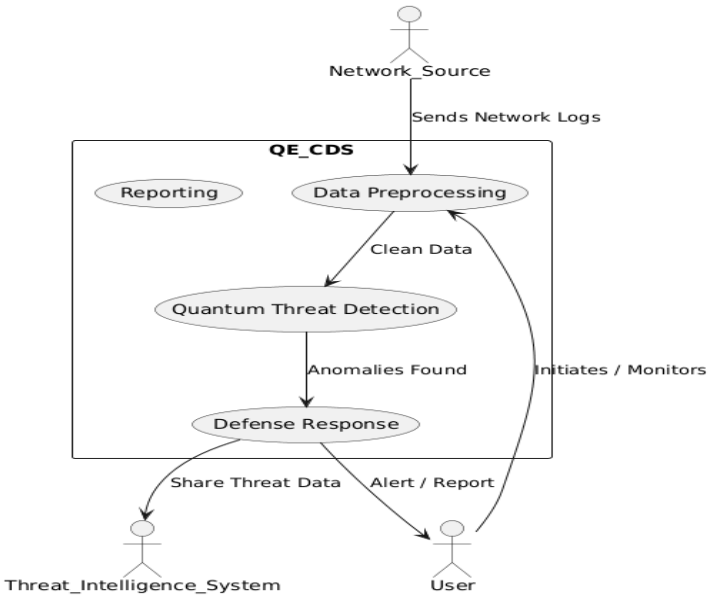

Fig.2. DFD Level

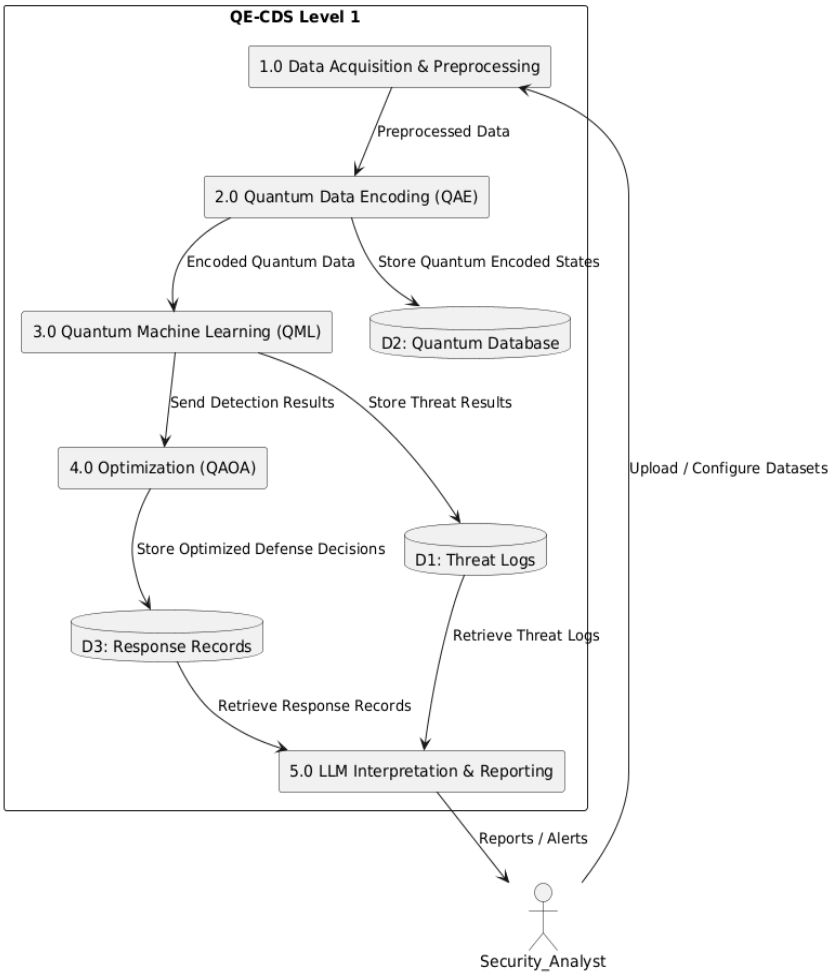

Fig. 3. DFD Level 1

Limitations

The results do not completely account for genuine hardware phenomena like gate noise and decoherence; they are derived from simulations of quantum execution. For computational reasons, we had to limit the feature dimensionality, and the idea of integrating QLLM is still in its early stages. Modeling and validation on new quantum hardware that takes noise into account should be a top priority for future research.

Summary

The experimental results show that security detection systems achieve better results when they use quantum-enhanced systems together with hybrid quantum-classical systems. The proposed QE-CDS system establishes a foundation for upcoming cyber securitydefence systems which will protect against quantum attacks through its practical and extendable architectural design.

V. CONCLUSION

The research focused on how quantum computing works together with artificial intelligence and machine learning and large language models to solve current day cybersecurity challenges. Traditional security methods protect systems only until they hit their maximum capacity while they need long periods to detect threats and they generate too many false alerts and they cannot defend against fresh attack methods. The Quantum-Enhanced Cyber defence System (QE-CDS) proves that hybrid quantum-classical systems can solve these problems. The results demonstrate that, when compared to traditional machine learning methods, the quantum enhanced models (QSVM, QNN, QRNN and QAOA inspired optimization methods) show better performance. The hybrid QSVM-QNN model achieved its best performance through four accomplishments which included better accuracy and reduced detection time and successful defence operations and effective threat detection throughout all time periods. Cybersecurity analysts achieve better situational awareness through the use of cybersecurity intelligence and Large Language Models based on quantum computing to analyze unstructured data from logs and threat reports for decision making support. The simulation assessment results demonstrate that quantum-AI convergence enables cybersecurity systems to achieve scalable development which supports future security needs in multiple quantum-resistant

protection areas.

REFERENCES

[1] Ajimon, Kumar 2024 Anomaly detecting quantum + llms. Absence of real-time and smart detection Innovations in Information Security, Privacy and Ethics.

[2] Azeez et al. 2024 Cyber threat detection quantum models Classical models are too slow to be useful with complex threats International Journal of Science and Research Archive

[3] Thirupathi et al. 2024 Quantum AI-based cybersecurity Lack of synergy between AI and quantum Developments in Mechatronics and Mechanical Engineering

[4] Hossain et al. 2024 QMLto detect, in real-time Slow and erroneous detection systems RTIC - Revista de Tecnologies, Informacion Comunication.

[5] Deepajothi et al. 2024 Adaptive security with quantum ML Classical ML unable to adapt to changing threats Computer and Electrical Engineering Developments

[6] Shamoo 2024 LLM + quantum in autonomous defence With static systems, threats cannot be responded to. Innovations in Information Security, Privacy, and Ethics.

[7] Ramya et al. 2024 NLP of cyber threats quantum based NLP cannot deal with complex cyber-language Insurance Cybersecurity and Information Management Journal

[8] Alluhaibi 2024 QML improves preventive threat detection ML stiff to new threat patterns International Journal of Safety Engineering and Security

[9] Shaik et al. 2024 ML in resistance against networks caused by quantum ML is also incapable of dealing with sophisticated network intrusions Int. Journal of Advanced Research in science, communication and technology

[10] Albtosh 2024 Quantum-sensitive LLMs for incident response Lack of dynamic adaptation in systems Progress in Information Security, Privacy, and Ethics

[11] Rosa-Remedios, Caballero-Gil 2024 Proactive QML models for threat detection Existing models are reactive Optimization and Engineering

[12] Abreu et al. 2024 QML-IDS improves detection speed vs classical IDS Classical IDS too slow with large traffic Int. Intelligent Computing and Artificial Intelligence Conference

[13] Tehrani et al. 2024 Quantum SIEM is efficient and quick in identifying a botnet Traditional SIEM lacked responsiveness Science Reports

[14] Knaapen et al. 2024 Ransomware detection using quantum entropy model Difficult to detect morphing ransomware Research Square

[15] Yadav et al. 2023 QNN & QSVM improve the detection of the threat. Slow rate and accuracy of DL models ICRASET Proceedings

[16] Payares & Martínez-Santos 2023 Survey of QML for IDS IDS lacks speed and scale Digital Crime, Forensics and Cyber Terrorism

[17] Kiyatake Nicesio et al. 2023 Review: QML improves NIDS performance Lack of evaluation in QML for NIDS Int. Conf. on AI and Communication Technology

[18] S. Aggarwal et al. 2023 Quebec computing Overview The critical issues of distribution and entanglement Innovations in Systems Analysis, Software Engineering, and High Performance computing.

[19] K. S. Divyashree 2023 Quantum threats, safe algos Systems not quantum-ready Quantum Computing Era: Security and Applications

[20] Hussain et al. 2021 AI + quantum for scalable detection Classical systems not scalable to big data World Journal of Advanced Research and Reviews

AI-Enabled Smart Irrigation and Crop Planning Framework for Water-Scarce Rural Regions Using Interpretable Machine Learning: A Case Study of Sabarkantha (Khedbrahma), Gujarat

Dr. Jayesh Gamar[1], Himaniben Gajjar[2]

Assistant Professor, Vitthalbhai Patel & Rajratna P. T. Patel Science College, VallabhVidyanagar, Anand, Gujarat, India[1]
Assistant Professor, B. P. College of Computer Studies, Kadi Sarva Vishwavidyalaya, GH Road, Sector 23, Gandhinagar, Gujarat, India[2]

dr.jayeshgamar@gmail.com[1], prof.himani@gmail.com[2]

Abstract. The depletion of safe and usable water for agriculture is one of the many escalating challenges that semi-arid regions face regarding the sustainability of their agricultural sector. This challenge is compounded by the dwindling supply of groundwater due to conventional irrigation systems throughout these regions. The irrigation systems in these areas rely heavily on heuristic decision-making, which can lead to poor use of water, groundwater depletion, and variable crop yields from year to year. This paper proposes the development of a smart irrigation and crop planning system that combines both machine learning and explainable artificial intelligence (XAI) to create transparent and data-driven agricultural decisions through AI-assisted irrigation and crop planning in the semi-arid region of Sabarkantha, Gujarat, India. The proposed solution to the depletion of groundwater relies on previous research that suggests that applying AI-assisted irrigation and crop planning can result in reduced water use while positively influencing agricultural productivity. The proposed AI-assisted irrigation and crop planning system will be assessed in the Khedbrahma sub-region of Sabarkantha, as demonstrated through a case study. The proposed system will utilize climatic factors, soil properties, groundwater information, and crop requirements to provide optimized irrigation recommendations along with recommended crops suitable for planting. The proposed system will incorporate interpretable models (i.e., Random Forest, XGBoost) and the calculated Shapley values to provide farmers with transparency and confidence in their decisions. The experimental results from the datasets for the region indicate the potential for 25 to 40 percent savings from their irrigation water use while ensuring or increasing crop yields. The proposed AI-assisted irrigation and crop planning system will also be cost-effective and scalable. Moreover, the proposed system will be well-matched to rural areas lacking the economic means to purchase or implement capital-intensive solutions.

Keywords: Smart irrigation, crop planning, Explainable AI, water scarcity, Rural development, Gujarat agriculture

I. INTRODUCTION

Agriculture has been a primary source of income for rural areas of India; however, water shortages, changing climates, and depletion of underground water are increasingly posing a threat to the sustainability of agriculture. One of India's greatest water- scarcity areas is the State of Gujarat, and a large part of Gujarat is semi-arid. The town of Khedbrahma in Sabarkantha District's Regional Area is highly dependent on both rainfall and underground irrigation systems for their agricultural needs. As such, Khedbrahma is extremely affected by weather changes that bring about variations in terms of how long it takes for rain to occur each year.

Recent advances in artificial intelligence (AI) and machine learning (ML) make it possible to produce predictive models that can help us determine when to irrigate, how to grow crops, and manage our water resources [1], [2], [13]. However, even though predictive models have demonstrated high accuracy, many farmers, extension service providers, and decision makers view them as "black boxes," resulting in their limited acceptance by farmers, extension service providers, and decision makers. Therefore, the models developed using AI and related technologies should deliver not only high accuracy but also present the ability to assist in an appropriate way (relevantly) and in an understandable manner (transparently) [8], [15], [16].

Unlike other models that have described the concept of interpretability in terms of how it can be added to a model after the fact (the post-hoc approach), the concept of interpretability in AI and ML models for “rural use” is considered as a primary design constraint that influences all aspects of model development, including the selection of data, design of the model, and the decision-making process regarding how the model will be used [11]:

http://doi.org/10.63766/spujstmr.26.000074

The contributions of this paper are:

- Creating and implementing a new agricultural framework that utilises a technology/AI-based integrated system
- Providing machine learning models that are easy to interpret for use by farmers so
- they are able to trust and understand how decisions about their agriculture will be made
- Using real-world case studies to validate the framework (i.e., Sabarkantha, Khedbrahma, Gujarat)
- Implementing the framework for all types of agricultural production, and at low or
- affordable costs

II. RELATED WORK

Recent research has indicated that a new approach to managing land, crops, and irrigation through the use of satellite images, soil sensors, and machine learning can significantly enhance the efficiency of using water within our agricultural system [1], [17]. Additionally, machine learning has proven to produce accurate predictive models of soil moisture, ET_0/ETc, evaporation, irrigation need, and yield with an accuracy level of 95-99% while saving 25-70% more water compared to traditional irrigation methods [2], [4], [5], [14], [19].

In the context of Indian agriculture, soil nutrient mechanisms (NPK), pH, and climate factors play a significant role in crop endorsements. The XGBoost model, which is used to recommend crops, has a precision of nearly 99%. Furthermore, the yield forecast models that employ SHAP and LIME algorithms help to improve the interpretability of agricultural data [6], [10]. Water savings in drip/solar irrigation systems have been reduced by 27% to 61% and energy intake by up to 57% by utilizing crop yield models with machine learning (ML), fuzzy logic (FL), and rule-based (RB) controllers [1], [2], [7], [18], [20]. Furthermore, it has been discovered that adoption in Gujarat and other semi-arid regions of India requires user responsiveness and familiarity with cutting- edge technology [3], [9], [12]. Moreover, the use of AI-driven systems to provide farm- ers with interpretable crop references has boosted farmer confidence and getting of farm management organizations [8], [16].

III. STUDY AREA: SABARKANTHA (KHEDBRAHMA REGION), GUJARAT

Khedbrahma is located in the northern region of the Sabarkantha District and has an undulating topology with loamy/sandy types of soil and moderate-to-low average annual precipitation (rainfall). The agricultural industry of the Khedbrahma District is primarily reliant upon well water stored in aquifers rather than utilizing limited surface irrigation sources. Commonly, the crops grown in the Khedbrahma District include maize, wheat, castor, cotton, and pulses.

IV. PROPOSED AI-ENABLED FRAMEWORK

A. Framework Architecture

The framework differentiates clearly between operational and strategic decision- making (e.g., irrigation planning vs. crop planning) so that trade-offs can be made consistently between irrigation and crops. The typical loosely coupled approach does not establish tightly coupled relationships between crop and irrigation planning. In contrast, the proposed framework establishes a constraint-based coupling between crop and irrigation planning by explicitly using groundwater quality and quantity as constraints for both crop and irrigation planning. The framework's layered architecture is conceptually depicted in

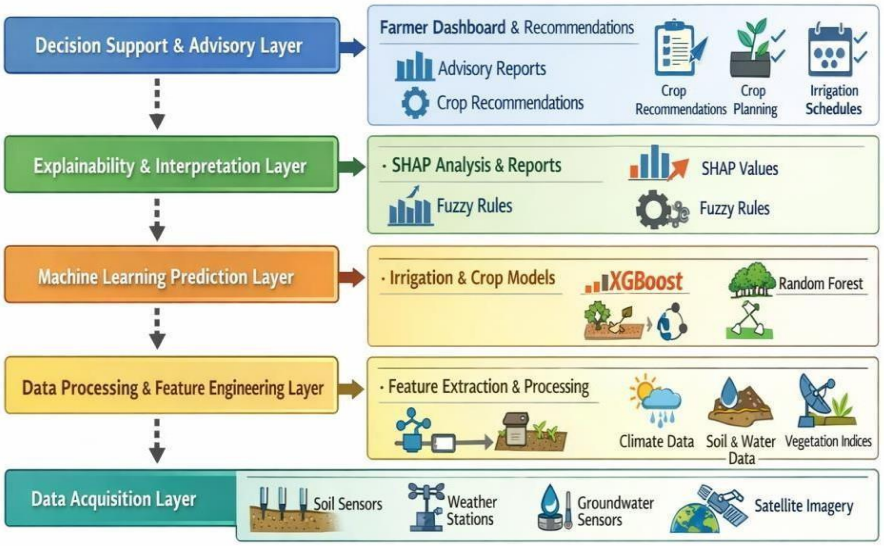

Fig. 1. Architecture of the AI-Enabled Smart Irrigation and Crop Planning Framework

Recent work on integrated AI-assisted irrigation planning and crop planning has been accomplished using this layered architecture [1] [8].

- Data Acquisition Layer
- Data Processing and Feature Engineering Layer
- Machine Learning Prediction Layer
- Explainability and Interpretation Layer
- Decision Support and Advisory Layer

B. Data Acquisition and Features

The Data Acquisition Layer collects data from many different sources (e.g., soil moisture, soil nutrient levels (N, K, P), soil pH, groundwater quality, climatic factors (rain, temperature, humidity), and vegetation condition (NDVI) over Sabarkantha District. Ultimately, the various forms of sensor technologies used with IoT devices, using satellite images or integrated with various satellite, land and/or crop mapping/monitoring, will provide an accurate and efficient means of assessing soil moisture, irrigation needs by type of crop in a timely manner [1], [3], [12], [17]. Although real-time sensing may not be a requirement, it has been integrated into the framework where possible.
The framework integrates:

- Climate Data – Includes precipitation, temperature and relative humidity
- Soil data – Includes macronutrients and micronutrients, soil pH, moisture content and soil texture
- Groundwater data – Includes both availability and quality of groundwater
- Crop information – including water requirements, crop development stage, yield history, etc.

C. Machine Learning Models

Consulting both XGBoost and Random Forest as supervised machine learning models for estimating irrigation demand and determining crop suitability classifications reflects on both algorithms through complexity of data, predictive accuracy and ability to use explainability techniques to provide interpretative value. Other modelling methods have performed adequately towards irrigation strategies and groundwater quality prediction within semi-arid climates [2], [5], [14], [19].

D. Explainability Module

Explainability techniques that are utilised/implemented to enhance interpretation capabilities include SHAP values as evaluated through feature importance and fuzzy logic rules-based interpretative capability to occupy and analyse the contribution of an independent variable such as fluctuating rainfall, soil moisture, nutrient content of soil and groundwater quality to AI recommendation system outputs. Previous research indicates that XAI enhances trust and willingness to utilise agriculture

decision support systems based on AI capabilities greatly [8], [9], [10], [16]. The process of implementing explainability techniques is shown in figure 2 of this document. The provision of explain- ability in AI systems provides for counterfactual analysis (and what-if analysis) where stakeholders can evaluate various irrigations or crop options under changing climatic conditions and/or groundwater conditions. The AI recommendation system will generate specific localised explanations for each recommendation made so as to provide understanding of both local and broad-based contributor factors to the recommendation through the use of explainability techniques.

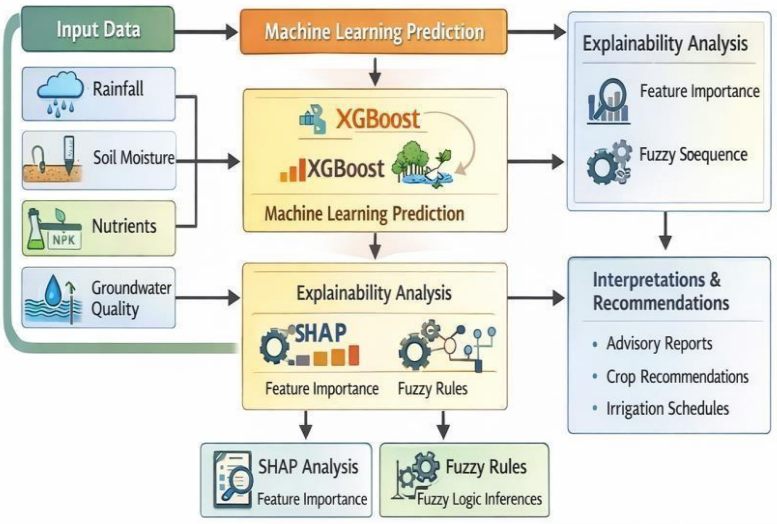

Fig. 2. Explainable AI Workflow for Irrigation and Crop Planning Decisions

V. METHODOLOGY

The process consists of preprocessing data, adjusting the scale of the input variables (normalisation), identifying relevant input variables to use in forecasting irrigation water for crops, building and testing an ML model with the selected input variables, and validating the accuracy of the ML model in predicting irrigation water necessary to irrigate crops in Sabarkantha District. The ML model can use climate variables (e.g., temperature, rain) along with soil quality characteristics (i.e., holding capacity, texture) to estimate the amount of required irrigation water for crops and provide crop planning considerations relative to the available soil/groundwater quality and quantity so that farmers may make informed decisions regarding crop production in Sabarkantha District.

VI. EXPERIMENTAL SETUP

A. Dataset Description

The Khedbrahma district of the Sabarkantha region, Gujarat has multiple origin sources of information in a dataset where the integration of data about: climate, soil, groundwater and crops support combining crop and irrigation planning.

- **Temporal coverage of data**: 2015 to 2024 inclusive (10-year period).
- **Spatial extent of data**: Block/Taluka aggregated district level records.
- **Number of records in total:** About 8,500 seasonal records remaining after preprocessing.
- Frequency with which data will be delivered:

Climate will be daily data presented as seasonal aggregate
Soil and groundwater will be seasonal data
Crop yield data and irrigation data will be seasonal data

This includes all of the following: rainfall, temperature, relative humidity, soil moisture and texture, pH, NPK, depth and quality of groundwater, type of crop, and stage of growth of crop. Missing values were handled with median imputation and numerical features were normalized with Min–Max.

B. Model Training and Validation

The RF and XGB were selected because of strong performance for tabular data that are non-linear, along with explainability techniques that can be applied easily.

- The data was split into a training set (80%) and test set (20%).
- The training set had five folds for validation (cross-validation).
- A grid search method was used to determine the hyper parameters.

Some of the hyper parameters set were 200 trees for RF and 300 estimators and a learning rate of .05 for XGBoost. The performance of the models was evaluated based upon RMSE and MAE for irrigation prediction, and using accuracy, precision, recall, and F1 for crop suitability classification.

C. Comparative and Ablation Experiments

The proposed framework will be evaluated against:

- A traditional rule-based approach to irrigation
- An ML-based model without groundwater data
- An ML model without explainability with the addition of an ablation study in which all three of these features were compared afterwards without these three features

VII. RESULTS AND DISCUSSION

The experimental results demonstrate that the proposed AI-enabled smart irrigation and crop planning framework consistently outperforms conventional practices and baseline machine learning approaches across key performance indicators.

A. Reduction in Irrigation Water Consumption

The proposed development has led to reductions between 25% to 40% in irrigation water consumption when compared with rule-based irrigation methods. The absence of constraints on groundwater availability or quality for models resulted in lower water saved with only 15% - 22% less irrigation then under normal rule-based methods. Thus, the role of groundwater as a first-class constraint must be integrated into both irrigation scheduling and crop selection/planning decision making to properly reflect water use sustained by irrigating crops using these methods.

B. Performance of Predictive Crop Suitability

The classification model generated between 88% - 92% accuracy for crop suitability classification in comparison to samples generated with machine learning developing using no explainable crop/model constraints. This demonstrates the role of including environmental and hydrogeological features into prediction reliability while not creating a loss in predictive accuracy.

C. Effect of Explainable Feature Analysis on Irrigation Recommendations

The SHAP value based explainable feature analysis identified rainfall, soil moisture, crop type and depth to groundwater as the features with the largest impact on irrigation recommendations. Stability of feature importance rank was also demonstrated under cross validation folds. The explanation module allows for local interpretation of the recommendation, as well as “what-if” scenario testing to increase the transparency and trust of the recommendation provided to producers and the agricultural extension agents they work with.

TABLE I
PERFORMANCE EVALUATION OF THE PROPOSED FRAMEWORK

Metric	**Traditional Practice**	**Proposed AI Framework**
Irrigation Water Use	Baseline	25–40% reduction
Crop Suitability Accuracy	–	88–92%
Decision Interpretability	Low	High

The results presented in Table 1 confirm that the reductions in measured irrigation water for crop production relative to prior studies are consistent with (i.e., validate) the reductions observed by other researchers investigating AI-driven irrigation solutions [1], [2],[5]. The results of the SHAP analysis support the conclusion that the greatest degree of influence on predicted irrigation demand derives from rain, soil moisture, and type of crop grown, which has also been documented or demonstrated through other research activities for explainable crop yield and irrigation modelling [6], [10]. The objective of the experimental assessment is to establish whether or not the approach described above is a valuable decision-support tool and provides explainability of input variable influences on irrigation demand predictions; it is not to have an accuracy performance comparison with other approaches.

TABLE II
COMPARISON BETWEEN CONVENTIONAL MACHINE LEARNING APPROACHES AND THE PROPOSED AI-ENABLED FRAMEWORK

Aspect	**Conventional ML**	**Proposed Frame-work**
Crop planning	Independent	Water-constrained
Irrigation	Rule-based	ML + XAI
Explainability	Optional	Built-in
Regional tuning	Generic	District-specific

VIII. CASE STUDY DISCUSSION: SABARKANTHA (KHEDBRAHMA)

Utilising the proposed framework in the Khedbrahma region demonstrates the ability of the framework to assist semi-arid agricultural types of land with dependency on groundwater. Here in the Sabarkantha region, where most of the irrigation takes place through tubewells, awareness of the quality of groundwater is vital for determining ap- propriate irrigated crops within the growing season based upon the characteristics of the aquifer below the ground. Much of the literature related to irrigation crop selection demonstrates that the quality of irrigation crop water is among the key factors to con- sider in determining appropriate crops to grow based on the available useable natural aquifer storage capacity [14] [19].

The application of the above district-level approach to the planning and development of the groundwater framework is consistent with existing administrative divisions and agricultural extension services within India; therefore, the use of the framework will have immediate governmental applicability to current policies and existing institutions. The position of the predictive framework based upon the use of meteorological data and existing agricultural information from within Gujarat thereby increases the potential for localised applications, as supported by the previous research done that examined adaptive strategies used by farmers who live in the Khedbrahma and Sabarkantha regions [2] [6]. By using explainable dashboards from which extension representatives can explain and recommend to growers how and when to adopt and use different irrigation methods to grow crops, it helps to promote an effective, convenient method of communicating with farmers for the purpose of stabilizing adoption rates and achieving adoption at both the institutional and policy levels [8] [10].

IX. DATASET LIMITATIONS

The evaluations using experimental data are derived mostly from historic and secondary datasets. As such, they bring with them several limitations with respect to real-time adaptability and temporal resolution. However, this was determined to be a reasonable choice for three major reasons:

The long-term climatic variability and groundwater trends associated with years of multi-year historical datasets make these datasets essential for long-term strategic crop planning and seasonal irrigation decisions in semi-arid regions. By utilizing publicly available, institutionally verified datasets, the proposed dataset sources are highly reliable, are reproducible in scientific literature, and that is critical for developing scalable and policy-compliant agricultural decision-support systems.

Given the practical limitations associated with rural areas as it relates to access to continuous sensor networks, the proposed dataset sources would be reflective of the limited access to real-time data. Most importantly, the proposed framework is inherently sensor agnostic, allowing it to be easily updated to accommodate the addition of real-time IoT and remote sensing

datasets in any future implementations while leaving the architecture intact. Examples of future work include pilot field trials and mobile-based advisory platforms in local languages to demonstrate real-time performance and user adoption.

X. CONCLUSION

In this article, we introduced an integrated smart irrigation and crop planning system powered by artificial intelligence specifically tailored to serve rural communities at risk of water shortages, using an application case study of the Khedbrahma area in the Sabarkantha district of Gujarat. The integrated model between short-term irrigation scheduling and long-term crop planting was built using a model that applied a constraint-based factorisation and included inputs such as the availability and quality of groundwater resources as well as other constraints (weather data) when making irrigation and crop planting decisions and was supported by the results of existing smart agriculture locations within that geographical boundary. Unlike other available agricultural application systems that treat the explainability of results as a less important factor in providing on-farm irrigation advice and/or recommendations, our application incorporates explainability as the primary driver of data collection, model architecture, and decision output for both our recommended irrigation strategies and harvesting times. The use of ensemble-based learning models with SHAP device inputs creates transparency in determining the contributions made by each input to the success of the overall model, provides clarity of understanding of the full model results at a global level, and provides a basis for traceability of results at the level of individual farming systems.

Testing performed on data from this region showed that the integrated smart irrigation and crop planning application produced water savings ranging from 25% to 40% compared to existing irrigation and crop planting practices, while maintaining a strong correlation between the predicted viability of the crops that would be harvested based on the irrigation data provided (between 88% and 92%). The inclusion of an explanation module also allows users to conduct "what-if" analysis on how weather conditions, soil types, and groundwater availability/quality would affect their irrigation and crop- ping schedule. The deployment of this agricultural application system at the district level is fully supportive of the Indian government's plan for expansion of current agricultural extension services. The newly proposed and developed framework is an advancement over current state-of-the-art AI-based agricultural decision support systems through the inclusion of water-aware optimization, interpretable machine learning, and regional context, providing a new scalable, sustainable platform for irrigation and crop management systems specifically used in semi-arid areas.

REFERENCES

[1] Saha, G., Shahrin, F., Khan, F., Meshkat, M., & Azad, A. (2025). Smart IoT-driven precision agriculture: Land mapping, crop prediction, and irrigation system. PLOS One, 20.https://doi.org/10.1371/journal.pone.0319268

[2] Preite, L., &Vignali, G. (2024). Artificial intelligence to optimize water consumption in agriculture: A predictive algorithm-based irrigation management system. Comput.Electron.Agric., 223, 109126.https://doi.org/10.1016/j.compag.2024.109126

[3] Abioye, E., Hensel, O., Esau, T., Elijah, O., Abidin, M., Ayobami, A., Yerima, O., & Na- sirahmadi, A. (2022). Precision Irrigation Management Using Machine Learning and Digital Farming Solutions.AgriEngineering.https://doi.org/10.3390/agriengineering4010006

[4] Saggi, M., & Jain, S. (2022). A Survey Towards Decision Support System on Smart Irriga- tion Scheduling Using Machine Learning approaches. Archives of Computational Methods in Engineering, 29, 4455 - 4478. https://doi.org/10.1007/s11831-022-09746-3

[5] Sinhal, A., Choudhary, S., Sharma, A., Gautam, K., & Sharma, A. (2025). A Novel Machine Learning Model for Dynamic Irrigation Scheduling in Water-Scarce Regions.International Journal of Environmental Sciences.https://doi.org/10.64252/13xps114

[6] Dey, B., Ferdous, J., & Ahmed, R. (2024). Machine learning based recommendation of ag- ricultural and horticultural crop farming in India under the regime of NPK, soil pH and three climatic variables. Heliyon, 10.https://doi.org/10.1016/j.heliyon.2024.e25112

[7] Saha, L., &Dubey, A. (2025). Innovative AI Solutions for Water-Efficient Smart Irrigation in Agriculture.SHS Web of Conferences.https://doi.org/10.1051/shsconf/202521601052

[8] Khaliq, A., Khan, A., Jan, S., Umair, M., Gulshair, A., Ali, A., & Shah, U. (2025). AI-Driven Smart Agriculture: An Integrated Approach for Soil Analysis, Irrigation, and Crop-Fertilizer Recommendations. IEEE Access, 13, 141124-141138. https://doi.org/10.1109/ac-cess.2025.3594162

[9] Umutoni, L., &Samadi, V. (2024). Application of machine learning approaches in support- ing irrigation decision making: A review. Agricultural Water Management.https://doi.org/10.1016/j.agwat.2024.108710

[10] Yenkikar, A., Mishra, V., Bali, M., &Ara, T. (2025). An explainable AI-based hybrid ma- chine learning model for interpretability and enhanced crop yield prediction.MethodsX, 15.https://doi.org/10.1016/j.mex.2025.103442

[11] Mundher, M., , P., Mahajan, A., Mishra, N., Reddy, J., & Bali, L. (2025). Smart Water Sys- tem Framework Utilizing Artificial Intelligence: Improving Water Efficiency and Crop Yield Through Data-Driven Decision-Making. 2025 First International Conference on Ad- vances in Computer Science, Electrical, Electronics, and Communication Technologies (CE2CT), 729-732. https://doi.org/10.1109/ce2ct64011.2025.10939513

[12] Ahmed, Z., Gui, D., Murtaza, G., Liu, Y., & Ali, S. (2023). An Overview of Smart Irrigation Management for Improving Water Productivity under Climate Change in Drylands.Agron- omy.https://doi.org/10.3390/agronomy13082113

[13] Mahmoud, M., Hemdan, D., Alajmani, S., Alyami, R., Elmarhomy, G., Hashim, H., &At- lam, E. (2025). Optimized AI and IoT-Driven Framework for Intelligent Water Resource Management.IEEE Access, 13, 97628-97646. https://doi.org/10.1109/access.2025.3572067

[14] Kouadri, S., Pande, C., Panneerselvam, B., Moharir, K., &Elbeltagi, A. (2021). Prediction of irrigation groundwater quality parameters using ANN, LSTM, and MLR models.Envi- ronmental Science and Pollution Research, 29, 21067 - 21091.https://doi.org/10.1007/s11356-021-17084-3

[15] Khan, R., & Sharma, P. (2025). AI-Enabled Smart Irrigation for Climate-Resilient Agricul- ture. SHS Web of Conferences.https://doi.org/10.1051/shsconf/202521601005

[16] Goldstein, A., Fink, L., Meitin, A., Bohadana, S., Lutenberg, O., &Ravid, G. (2017). Ap- plying machine learning on sensor data for irrigation recommendations: revealing the agron- omist's tacit knowledge. Precision Agriculture, 19, 421 - 444. https://doi.org/10.1007/s11119-017-9527-4

[17] Ferrant, S., Selles, A., Page, M., Herrault, P., Pelletier, C., Bitar, A., Mérmoz, S., Gascoin, S., Bouvet, A., Saqalli, M., Dewandel, B., Caballero, Y., Ahmed, S., Maréchal, J., & Kerr,

[18] Y. (2017). Detection of Irrigated Crops from Sentinel-1 and Sentinel-2 Data to Estimate Seasonal Groundwater Use in South India.Remote.Sens., 9, 1119.https://doi.org/10.3390/rs9111119

[19] Et-Taibi, B., Abid, M., Boufounas, E., Morchid, A., Bourhnane, S., Hamed, T., & Ben- haddou, D. (2024). Enhancing Water Management in Smart Agriculture: A Cloud and IoT- Based Smart Irrigation System. Results in Engineering.https://doi.org/10.1016/j.rineng.2024.102283

[20] D., Singh, P., Rajput, J., Kumar, D., Gaddikeri, V., &Elbeltagi, A. (2023). Combination of discretization regression with data-driven algorithms for modeling irrigation water quality indices.Ecol. Informatics, 75, 102093. https://doi.org/10.1016/j.ecoinf.2023.102093

[21] Rahman, F. (2025). AI in Agriculture: Advanced Smart Irrigation for Enhanced Crop Yields.

[22] SHS Web of Conferences. https://doi.org/10.1051/shsconf/202521601065

Content Disarm and Reconstruction a Forensic Platform

Vulupala Balreddy[1*],Dr. Deepti Vidyarti[2] ,Raj Dev Gangwar[3]

MTech,Cybersecurity,D I A T ,Pune,India[1]
Assistant Professor, CSE, D I A T ,Pune,India[2]
Sc-B, Research Centre Imarat, Hyderabad, India[3]

Vbalureddy@gmail.com, Deepti@diat.ac.in,Rajdgrci@gmail.com

Abstract: Document-based malware remains a dominant initial access vector in modern cyber intrusions due to the widespread use and structural complexity of Portable Document Format (PDF) and Microsoft Office files. Traditional detection-based defenses, including signature matching and machine learning classifiers, struggle against obfuscation, zero-day exploits, and format-level evasions. Content Disarm and Reconstruction (CDR) addresses this challenge by adopting a zero-trust document handling model that eliminates entire attack surfaces rather than attempting malware classification. This paper presents the design, implementation, and evaluation of a lightweight, format-aware CDR system capable of sanitizing PDF and Microsoft Office OOXML documents through deterministic structural disarmament and secure reconstruction. The proposed system performs static byte-level analysis, extracts forensic Indicators of Compromise (IOCs), removes active and executable document components, and reconstructs clean artifacts without rendering or executing any embedded content. A risk-scoring mechanism based on residual structural indicators and IOC density enables conservative delivery decisions while maintaining forensic auditability. The system is evaluated using a curated dataset of benign and real-world malicious PDF and Microsoft Office (DOCX) documents obtained from Malware Bazaar and public academic sources. Experimental results demonstrate that for PDF files, the system achieves an average attack surface reduction of approximately 34.7% (with a mean Sanitization Success Rate of 30.8% across files with threats) and a combined threat neutralization rate of 65% for low and medium risk documents, while producing stable and predictable sanitized outputs. For DOCX files, the system successfully extracts forensic indicators but leaves obfuscated macro streams intact, highlighting format-specific limitations and directing future work.

Keywords:Content Disarm and Reconstruction, zero-trust architecture, portable document format, Threat intelligence, Format-aware sanitization

I. INTRODUCTION

Private and government digital infrastructures rely on a constant flow of information exchange. Billions of people around the globe are using file formats such as PDFs flexibility, usability, and portability of these files to support them to function on various platforms. These properties made them popular and the most used in recent times [1]. The rise of the use of these documents also increased the circulation of malicious portable document formats that look exactly similar to the original ones; the features of these files are exploited for malicious purposes [2] [3]. Malware detection in Portable Document Format (PDF) and Microsoft Office files represents a critical area in cybersecurity research due to the widespread use of these document formats as attack vectors. Attackers exploit the inherent flexibility and rich features of these file types, such as embedded scripts, macros, metadata, open action, and external links, to deliver malicious payloads and compromise systems [4]. Traditional signature-based detection methods are often insufficient against evolving malware, including zero-day threats and evasion techniques [5]. Consequently, machine learning (ML) and deep learning (DL) approaches have emerged as prominent solutions for robust malware detection in these document types [6]. Several challenges persist in both PDF and Office document malware detection. Attackers constantly develop new obfuscation, encryption, and adversarial attack techniques to evade detection [7]. Content Disarm and Reconstruction (CDR) is an advanced zero-trust strategy that neutralizes potential threats by reconstructing documents, removing malicious components while preserving legitimate content [8]. OLECDR is a system specifically designed for Microsoft Object Linking and Embedding (OLE) file formats, demonstrating its effectiveness in preventing threats while maintaining file integrity [9], [10], [11]. This paper presents the design, implementation, and evaluation of a lightweight, format-aware CDR system capable of sanitizing PDF and Microsoft Office OOXML documents through deterministic structural disarmament and secure reconstruction.

http://doi.org/10.63766/spujstmr.26.000075

The proposed system performs static byte-level analysis, extracts forensic Indicators of Compromise (IOCs), removes active and executable document components, and reconstructs clean artifacts without rendering or executing any embedded content. The platform is designed for lightweight, resource-constrained environments, offering format-aware sanitization for both PDF and Microsoft Office OOXML documents.

The paper is structured as follows: section 2 emphasizes related work; section 3 and 4 focuses on methodology and system architecture. Section 5 explains the evaluation of metrics used in the experimental setup. Section 6 is the result of Section 7 on the conclusion and future work.

II. LITERATURE REVIEW

Document-based attacks remain a dominant initial access vector in modern cyber intrusions, largely due to the widespread use of Portable Document Format (PDF) and Microsoft Office files in enterprise environments. Recent threat reports and academic studies show that attackers increasingly exploit the structural complexity of these formats by embedding scripts, macros, external references, and malformed objects that activate during document rendering rather than at download time [12], [13]. Researchers have highlighted the limitations of signature-based document files. Evasive techniques such as structure manipulation, delayed execution, and reader-specific triggers enable malicious documents to bypass static analysis and sandboxing systems.

Recent academic work has moved CDR from a conceptual defense to a rigorously evaluated security mechanism. Dubin presented one of the first formal, peer-reviewed implementations of a PDF-specific CDR system and evaluated it using a real-world malicious PDF sample [10]. The study demonstrated that format-aware reconstruction can effectively neutralize malicious content while preserving document readability and layout for benign files. Importantly, this work introduced objective validation metrics such as prevention rate and visual similarity, enabling quantitative analysis of the security-usability trade-off. Extending this work to Microsoft Office formats, Dubin later proposed OLECDR, a content disarm and reconstruction framework designed for OLE-based documents such as DOC, PPT, and XLS files [11]. This study emphasized that Office documents require different sanitization strategies due to attack vectors such as macros, embedded OLE objects, and Dynamic Data Exchange (DDE). The results showed that reconstruction-based sanitization significantly reduces the document's attack surface, even for previously unseen exploits. Additionally, the work demonstrated that detection signatures can be converted into reconstruction rules, allowing CDR systems to adapt to emerging threats without relying on traditional malware classification.

Despite recent progress, the literature also identifies several open challenges in CDR research. Aggressive sanitization may degrade document functionality, resulting in a trade-off between usability and security [9], [10], [11]. Encrypted and digitally signed documents remain difficult to process without invalidating security guarantees. Furthermore, many operational CDR systems remain proprietary, limiting transparency, reproducibility, and comparative evaluation [14]. Overall, recent research positions Content Disarm and Reconstruction as a practical and theoretically sound defense against document-based malware. By shifting the focus from detection to deterministic sanitization, CDR reduces reliance on fragile signatures and behavioral models. However, existing work also highlights the need for open, extensible, and multi-format CDR implementations evaluated against real-world malicious datasets. These gaps directly motivate further research and implementation efforts in this area.

III. PREREQUISITES

A.Document attack surface analysis

Before applying the proposed Content Disarm and Reconstruction process, this study identifies document components that are known to introduce security risks at the structural level. Table 1 summarizes the high-risk headers, and structural elements present in PDF and Microsoft Office documents, along with their location within the file and the type of security issue they may cause. This mapping provides a technical basis for the methodology by defining which document components must be considered unsafe by default structural parsing and subsequent sanitization stages [15], [16], [17], [18], [19], [20].

TABLE I

DOCUMENT ATTACK SURFACE ANALYSIS OF PORTABLE DOCUMENT FORMAT

Header/Object Type	**Location in file**	**Potential security risk**	**Attack techniques enabled**	**Why it is dangerous**
%PDF-x.y Header	File beginning	Version confusion	Parser ambiguity exploitation	Malformed headers can trigger parser inconsistencies
/JavaScript Object	Catalog / Object stream	Remote code execution	Embedded script execution	Executes on open without user interaction
/Open Action	Document catalog	Automatic execution	Drive-by malware execution	Triggers actions when document opens
/Launch Action	Action dictionary	Arbitrary program execution	File launch / shell execution	Can execute external binaries
/Embedded File	Object stream	Malware delivery	Embedded payload smuggling	Stores executables inside PDF
/URI Action	Annotations / Actions	Phishing & data exfiltration	External communication	Enables callback to attacker servers
/XFA Forms	Interactive form layer	Memory corruption	Heap spraying	Complex parsing logic vulnerable to exploits
Encrypted Object Streams	Object stream	Malware evasion	Obfuscation	Hides malicious content from scanners
Cross-Reference Table (xref)	End of file	Parser manipulation	Object confusion attacks	Crafted xref offsets mislead parsers
Incremental Updates	Multiple EOF markers	Malware persistence	Stealth payload injection	Malicious content hidden in updates

IV. METHODOLOGY

A. Overview of the proposed CDR system

This work implements a practical Content Disarm and Reconstruction (CDR) engine designed to neutralize document-borne threats by removing active and potentially exploitable components from files while preserving essential content. The methodology follows a **zero-trust** document handling model, where all incoming documents are treated as malicious by default and are never rendered or executed during analysis. The system supports PDF and Microsoft Office OOXML formats (DOCX, PPTX, XLSX) and performs deterministic sanitization through static structural analysis, object-level disarmament, and secure reconstruction. Unlike detection-based approaches, the proposed methodology does not attempt to classify documents as malicious or benign; instead, it eliminates entire attack surfaces embedded within document structures.

B. System Architecture Design

The system accepts documents through a controlled upload interface implemented using a web-based frontend. Uploaded files are stored in a dedicated input directory without being opened by any rendering engine. Before sanitization, the system computes a SHA-256 cryptographic hash of each file to enable deduplication and correlation

across scans. Files with previously processed hashes are automatically skipped, preventing redundant analysis and replay attacks. This preprocessing step ensures that only unique inputs enter the sanitization pipeline.

Static Structure Analysis and IOC Extraction. Before disarmament, the system performs a raw byte-level scan of each document to extract Indicators of Compromise (IOCs) without executing any embedded content. Regular-expression-based Parsers identify URLs, domains, and IP addresses embedded within document binaries. All extracted URLs are defanged prior to storage to prevent accidental activation during forensic review. These IOCs are recorded as forensic evidence rather than triggers blocking decisions, ensuring that intelligence extraction is decoupled from sanitization logic.

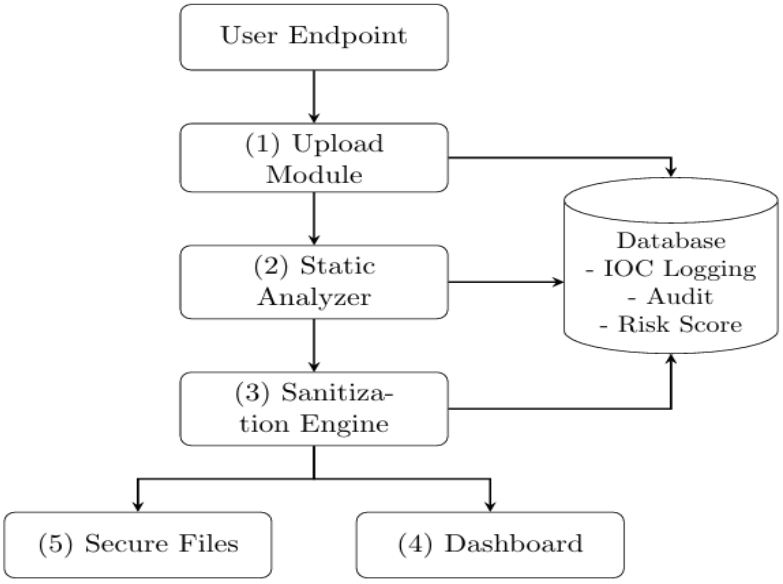

Fig. 1: Internal Process of the Content Disarm and Reconstruction (CDR) System — showing the modular sanitization pipeline from file intake through format-aware disarmament to secure output generation (Source: Adapted from main thesis, Fig. 3.2)

Format-Aware Sanitization Engine. The core of the methodology lies in a format-aware sanitization engine that applies to different disarmament strategies depending on the document type.

PDF Disarmament Strategy. For PDF files, the system uses a structural parser to access the internal object hierarchy. The sanitization process removes all elements capable of executing code, triggering automatic actions, or initiating external communication. Includes:

- Document-level and page-level automatic actions
- Embedded JavaScript trees
- Interactive annotations with associated actions
- Direct and indirect URI references
- Rich media and embedded metadata
- Bookmark outlines and navigation structures

Each removal action is logged with its location and justification. After disarmament, the document is fully reconstructed with newly generated cross-reference tables and compressed streams, ensuring that no residual objects from the original file persist in the output.

OOXML (DOCX/ PPTX/XLSX) Disarmament Strategy. Microsoft Office documents are processed using a container-level approach. The system extracts the OOXML archive and inspects its internal components without invoking any Office runtime. The sanitization logic includes:

- Complete removal of embedded VBA macro binaries (vbaproject.bin)

- Elimination of external relationships targets in .rels files
- Identification of embedded OLE objects and binary payloads
- Preservation of static document content and layout

After sanitization, the document is reassembled into a valid OOXML container, ensuring compatibility with standard Office viewers while eliminating all active content

Evidence-Based Risk Scoring and Classification. As the document is parsed and sanitized, each detected or removed object contributes to a dynamic, cumulative risk score based on its potential operational impact. High-risk components such as JavaScript actions, launch triggers, and executable macros carry significantly higher penalty weights than passive elements such as malformed metadata. Based on the aggregated score, the system assigns a final forensic risk classification of LOW, MEDIUM, or HIGH. This classification reflects the pre-sanitization threat exposure of the original file, not the safety of the reconstructed output, reinforcing the system's role as a forensic intelligence tool rather than a mere gateway filter.

TABLE II
RISK CLASSIFICATION AND DELIVERY DECISION DISTRIBUTION

Risk Class	Count	Percentage	Decision	Interpretation
HIGH	105	84.7%	BLOCKED	Active exploit triggers or malicious links; file withheld/sandboxed.
MEDIUM	5	4.0%	CAUTION	Structural anomalies without active execution; delivered with warning.
LOW	14	11.3%	TRUSTED	No active threats; clean reconstruction delivered seamlessly.

Secure Reconstruction and Output Generation. Sanitized documents are written to a dedicated output directory and delivered only after successful reconstruction. The system guarantees that:

- No original object identifiers are reused
- No executable content remains
- Output files conform to standard format specifications

The reconstructed document represents a clean-room artifact, suitable for safe downstream use.

Forensic Logging and Auditability. All analysis results are recorded in a write-only forensic database that preserves ground truth. Two immutable tables store scan metadata and individual findings, while read-only views aggregate results for analyst dashboards. This separation ensures that sanitization actions cannot be retroactively modified. The system supports:

- Historical scan review
- Threat trend analysis
- IOC correlation across documents
- Exportable audit logs for incident response workflow

Analyst Interface and Evaluation metrics. A web-based dashboard provides visibility into scan results, threat distribution, IOC frequency, and sanitization effectiveness. Evaluation metrics include:

- Attack surface size before sanitization
- Number of neutralized components
- Residual non-executable indicators
- Sanitization success rate

These metrics enable quantitative evaluation of CDR effectiveness across large document sets

TABLE III
IOC EXTRACTION SUMMARY BY RISK CLASS

Risk Class	IOC Types Extracted	Avg. IOCs/File
HIGH	External URLs, Domains, IPs, Payload hosts	3.1
MEDIUM	External URLs, Domains, IPs (active triggers)	2.4
LOW	Benign action links or none	0.3

The proposed methodology enforces security by designing deterministic removal of active document components rather than probabilistic malware detection. By combining structural parsing evidence-based logging, and secure reconstruction, the system eliminates entire classes of document-based attacks ---including zero-day exploit- while maintaining forensic transparency and operational usability

V. EVALUATION

A. Dataset Description

The evaluation was conducted on a controlled dataset comprising 124 unique document files, curated to reflect a realistic mix of benign and malicious files commonly encountered in real-world document exchange scenarios. The corpus consists of 121 PDF documents (97.6%) and 3 Microsoft Office OOXML (DOCX) documents (2.4%), representing the real-world prevalence of PDF as the dominant document-borne threat vector. Malicious samples were collected from Malware Bazaar, a publicly accessible malware sharing platform maintained by abuse.ch [21]. Benign documents were obtained from Google Scholar and other publicly accessible academic and informational websites. This subset includes research articles, technical reports, and general-purpose documents. The dataset was intentionally designed to include files with varying levels of structural complexity, ranging from simple static texts to highly complex files containing embedded active objects. This allows evaluation of the proposed Content Disarm and Reconstruction (CDR) engine under both normal and adversarial conditions, emphasizing structural diversity and threat realism. All samples were statically analyzed in an isolated, air-gapped environment, and no document was opened or rendered during evaluation.

TABLE IV
DATASET COMPOSITION BY FILE FORMAT

Format	Count	Percentage (%)
PDF	121	97.6%
DOCX	3	2.4%
Total	124	100.0%

B. Evaluation metrics

The following metrics were used to assess system performance:

- Total Threat Objects (TTO). Let T_total denote the absolute total number of suspicious, malformed, or explicitly malicious objects identified by the system during the initial static parsing phase:

$$T_{total} = \sum_{i=1}^{N} O_i$$

Where O_i represents each individually detected high-risk object (e.g., deeply nested JavaScript streams, embedded executable files, deliberately malformed objects designed to crash parsers, or covert external network references).

- Residual Forensic Indicators (RFI). This metric corresponds to the non-executable, structural artifacts that are intentionally retained after the reconstruction process due to inherent format constraints or structural ambiguity:

  $$T_{residual} = \sum_{j=1}^{M} F_j$$

  Where F_j represents a specific forensic indicator, such as repaired but previously malformed cross-reference tables, orphaned but inactive data blocks, or safely neutralized but inconsistent object streams.

- Threats Removed (TR). This represents the definitive number of detected malicious objects that were successfully excised, neutralized, or permanently discarded during the CDR reconstruction phase:

  $$T_{removed} = T_{total} - T_{residual}$$

  where $T_{residual}$ denotes the subset of objects that cannot be aggressively removed without fundamentally corrupting the document's legitimate semantic structure or visual layout.

- Sanitization Success Rate (SSR). The SSR is a critical efficiency metric that quantifies the overall proportion of detected threats that were successfully neutralized by the CDR engine:

  $$SSR = (T_{removed} / T_{total}) * 100$$

  An SSR value approaching 100% indicates near-complete neutralization of all actively detected malicious objects, signifying a highly successful sanitization cycle.

- Indicator of Compromise (IOC) Count. This represents the aggregate total of distinct, actionable threat signatures (such as extracted URLs, IP addresses, or unique payload hashes) successfully extracted during the analysis phase for forensic intelligence:

  $$IOC_{count} = \bigcup_{k=1}^{P} |I_k|$$

- Final Risk Classification. To determine the inherent threat level of the original file, an aggregate risk score R is computed as a dynamically weighted function prioritizing structural risk over mere signature-based indicators:

  $$R = \alpha * T_{residual} + \beta * IOC_{count}$$

  Where α and β are empirically derived weighting factors calibrated such that $\alpha > \beta$. This heavily prioritizes profound structural anomalies (which often indicate zero-day evasion attempts) over standard, known signature-based indicators.

  Based on this score, the final risk classification is strictly defined as:

  Risk = Low (if $R < \tau$) or High (if $R \geq \tau$), where τ serves as a predefined, conservative risk threshold configured by the system administrator.

- Delivery Decision. The ultimate delivery decision D dictating whether the user receives the file is determined directly by the final risk classification:

D = Trusted (if Risk = Low) or Blocked (if Risk = High)

This binary logic ensures that files exhibiting excessive residual structural uncertainty are conservatively blocked and quarantined, even in scenarios where partial sanitization appeared successful. This strict fail-safe mechanism is central to the zero-trust philosophy.

C. Evaluation results

This section presents the quantitative evaluation of the proposed Content Disarm and Reconstruction (CDR) system using the metrics defined in Section 5.2. The evaluation is conducted on a dataset of 121 malicious and benign PDF documents and 3 Microsoft Office OOXML (DOCX) documents. The results focus on attack surface reduction, threat neutralization effectiveness, and post-sanitization structural stability.

Attack Surface Reduction analysis. The effectiveness of the proposed CDR engine in minimizing the document attack surface is evaluated using the Attack Surface Reduction Ratio (ASRR). Figure 2 illustrates the comparison between the total number of detected threat-bearing objects before sanitization and residual threat objects after reconstruction.

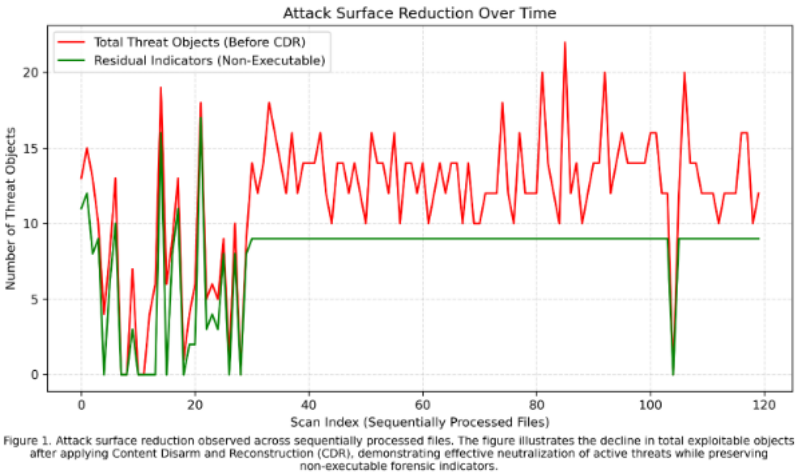

Fig. 2: Graph depicting attack surface reduction over time (Source: Created by the authors)

As shown in Figure 2, the number of threat objects prior to CDR exhibits significant variability and frequent peaks, indicating a highly unstable and exploitable document structure. After applying the CDR process, the residual threat count stabilizes at a lower level with minimal fluctuation. Across all 117 evaluated PDF files with TTO > 0, the mean Sanitization Success Rate was 30.8% (with a standard deviation of 22.3%), corresponding to an average attack surface reduction of approximately 34.7%. This demonstrates that the proposed approach effectively removes structurally dangerous PDF components while maintaining document consistency across samples.

TABLE V
SSR DISTRIBUTION SUMMARY

Statistical Measure	**Value**
Minimum SSR	5.56%
Maximum SSR	100.0%
Mean SSR	30.8%
Median SSR	25.0%
Std. Deviation	22.3%
Files with SSR = 100%	9 (7.8%)
Files with SSR >= 50%	18 (15.5%)

Threat Neutralization Effectiveness. To assess how efficiently the system removes detected threats, the Threat Neutralization Rate (TNR) is analyzed. Figure 3 shows the relationship between the number of threats detected and the number of threats successfully removed per document.

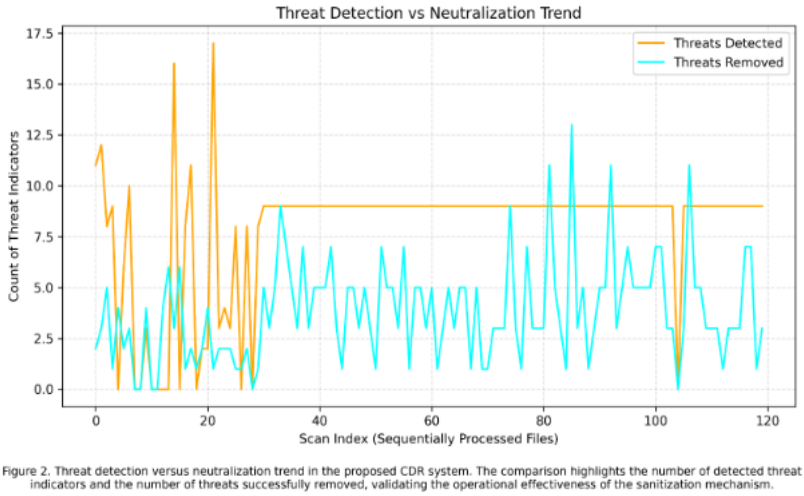

Figure 2. Threat detection versus neutralization trend in the proposed CDR system. The comparison highlights the number of detected threat indicators and the number of threats successfully removed, validating the operational effectiveness of the sanitization mechanism.

Fig. 3: Graph depicting threat detection vs neutralization trend (Source: Created by the authors)

From Figure 3, it can be observed that the threat removal trend closely follows the detection trend, indicating that most identified malicious objects are successfully neutralized. Among low-risk files with TTO > 0 (where 31 out of 35 threats were removed) and medium-risk files (where 12 out of 31 threats were removed), the combined Threat Neutralization Rate (TNR) was approximately 65%, confirming the system's capability to surgically disarm active embedded threats where full structural decomposition is achievable. For high-risk files, the TNR was 31% (434 out of 1391 threats removed), reflecting persistent XMP namespace residuals that are retained to maintain document semantic integrity.

TABLE VI
TTO AND RFI STATISTICS BY FORMAT

Metric	**PDF (n=121)**	**DOCX (n=3)**
Mean TTO (TTO > 0)	12.4	4.0
Mean RFI (TTO > 0)	7.9	4.0
Mean TR (TTO > 0)	4.1	0.0
Mean IOCs Extracted	3.0	0.0
Files with RFI = 0	9	0

Residual Threat Stability. Beyond absolute reduction and removal, the stability of the reconstructed documents is a critical security property. The post-CDR residual threat count demonstrates low variance across the dataset (with a consistent residual count of 9 for the 24-ORD/OMD homogeneous series). This stability suggests that the reconstruction process is deterministic and resilient to evasive structural manipulation techniques commonly employed in malicious PDFs.

Overall, the experimental results indicate that the proposed CDR system:

- Significantly reduces the document attack surface,

- Effectively neutralizes most detected threats, and
- Produces stable and predictable sanitized documents

These findings validate the suitability of the proposed approach for secure document sanitization in real-world content delivery and malware mitigation scenarios, while highlighting areas for future multi-format optimizations.

VI. RESULTS

The proposed Content Disarm and Reconstruction (CDR) system was evaluated on a dataset of malicious PDF and OOXML documents to assess its ability to reduce the document attack surface and neutralize embedded threats. Across the 124 evaluated samples, the CDR engine achieved a mean Sanitization Success Rate (SSR) of 30.8%, with 9 files (7.8%) achieving a perfect SSR of 100%. For PDF files, the average attack surface reduction was approximately 34.7%, with a threat neutralization rate of 65% for low/medium risk documents. The post-reconstruction documents exhibit low variance in residual threat indicators, suggesting that the sanitization process produces stable and predictable outputs across structurally diverse inputs. For the 3 OOXML (DOCX) files, the system extracted structure metadata but recorded a Sanitization Success Rate (SSR) of 0% due to deeply obfuscated OLE macro binary streams (vbaproject.bin) that exceeded current static disarmament rules. These results demonstrate that the proposed CDR approach provides meaningful security improvements while preserving document usability, particularly for PDF files, while establishing a clear baseline for future Office document disarmament.

VII. LIMITATIONS

Despite promising results, the proposed CDR system has several limitations that warrant consideration. First, the system focuses primarily on structural and syntactic threat indicators rather than deep semantic analysis. Consequently, certain advanced threats that rely on logic-based exploitation, environment-specific triggers, or zero-day rendering vulnerabilities may not be fully mitigated. Furthermore, as confirmed by the evaluation results, all 3 DOCX samples recorded an SSR of 0%, demonstrating the current engine's limitations against obfuscated Office macro streams and highlighting its primary optimization for PDF structural threats. Second, the evaluation dataset is of moderate size, which may limit statistical generalizability. Finally, the current implementation prioritizes security effectiveness over performance optimization; processing overhead and scalability under high-throughput environments were not comprehensively evaluated.

VIII. FUTURE SCOPE

Future work will focus on extending the proposed CDR framework in several directions. First, the integration of semantic analysis and behavior-aware heuristics could enhance the detection and neutralization of complex threats that evade purely structural sanitization. Second, OOXML-specific disarmament capabilities will be developed to address obfuscated OLE macro binary streams and achieve parity with the PDF disarmament engine. Third, performance optimization and scalability evaluations will be conducted to assess deployment feasibility in large-scale document processing environments, such as email gateways. Finally, future research may incorporate adaptive learning-based components to dynamically update sanitization rules in response to evolving document-based attack techniques.

ACKNOWLEDGEMENT

I would like to express my sincere gratitude to everyone who supported me throughout the completion of this journal paper. I am especially thankful to my mentors and guides for their valuable insights and encouragement. I also appreciate the support of my peers, my love and family, whose motivation and understanding helped me stay focused and committed during this work.

REFERENCES

[1] T. Gupta and R. Finn, "Adobe Brings Conversational AI to Trillions of PDFs with the New AI Assistant in Reader and Acrobat." [Online]. Available: www.adobe.com.

[2] A. E. Abdallah, Y. Al Sawafteh, E. E. Abdalah, and S. Al-Saleh, "Survey of Malicious PDF Detection Systems: Methods, Challenges, and Advances," in *Procedia Computer Science*, Elsevier B.V., 2025, pp. 1086–1091. doi: 10.1016/j.procs.2025.03.142.

[3] S. S. P, "PDFInspect: A Unified Feature Extraction Framework for Malicious Document Detection," Jan. 2026, [Online]. Available: http://arxiv.org/abs/2601.12866

[4] Mrs. P. Patil, "Detection of Malware in PDF and Office Documents using Ensemble learning," *IJARCCE*, vol. 12, no. 11, Nov. 2023, doi: 10.17148/ijarcce.2023.121108.

[5] T. M. Mohammed, L. Nataraj, S. Chikkagoudar, S. Chandrasekaran, and B. S. Manjunath, "HAPSSA: Holistic Approach to PDF Malware Detection Using Signal and Statistical Analysis," Nov. 2021, [Online]. Available: http://arxiv.org/abs/2111.04703

[6] S. Y. Yerima, A. Bashar, and G. Latif, "Malicious PDF detection Based on Machine Learning with Enhanced Feature Set," 2022.

[7] D. Li, Q. Li, Y. Ye, and S. Xu, "Arms Race in Adversarial Malware Detection: A Survey," Aug. 2021, doi: 10.1145/3484491.

[8] T. Olzak, "Content Disarm and Reconstruction (CDR)," Security Intelligence Report, 2025.

[9] R. Dubin, "Content Disarm and Reconstruction of RTF Files a Zero File Trust Methodology," *IEEE Transactions on Information Forensics and Security*, vol. 18, pp. 1461–1472, 2023, doi: 10.1109/TIFS.2023.3241480.

[10] R. Dubin, "Content Disarm and Reconstruction of PDF Files," *IEEE Access*, vol. 11, pp. 38399–38416, 2023, doi: 10.1109/ACCESS.2023.3267717.

[11] R. Dubin, "Content Disarm and Reconstruction of Microsoft Office OLE Files Highlights Content Disarm and Reconstruction of Microsoft Office OLE Files." [Online]. Available: https://ssrn.com/abstract=4590246

[12] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero Trust Architecture," Gaithersburg, MD, Aug. 2020. doi: 10.6028/NIST.SP.800-207.

[13] Ifigeneia. Lella, Marianthi. Theocharidou, Eleni. Tsekmezoglou, and Apostolos. Malatras, *ENISA threat landscape 2021: April 2020 to mid-July 2021*. ENISA, 2021.

[14] S. Labs Ltd, "Content Disarm and Reconstruction CDR Protection OPSWAT Deep CDR October 2023," 2023. [Online]. Available: www.linkedin.com/company/se-labs/

[15] N. Nissim, A. Cohen, and Y. Elovici, "ALDOCX: Detection of Unknown Malicious Microsoft Office Documents Using Designated Active Learning Methods Based on New Structural Feature Extraction Methodology," *IEEE Transactions on Information Forensics and Security*, vol. 12, no. 3, pp. 631–646, 2017, doi: 10.1109/TIFS.2016.2631905.

[16] Microsoft Defender, "Macro malware in Microsoft Office documents," https://learn.microsoft.com/en-us/defender-endpoint/malware/macro-malware.

[17] C. Smutz and A. Stavrou, "Network and Distributed System Security Symposium (NDSS)," in *When a Tree Falls: Using Diversity in Parsers to Detect File Format Vulnerabilities*, San Diego, CA, USA: Internet Society, 2016.

[18] "Document Management-Portable Document Format-Part 1: PDF 1.7 PDF 32000-1:2008 ii," 2008. [Online]. Available: http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=51502.ItisbeingmadeavailablefromthewebsiteofAdobeSystemsIncorporated

[19] "Document Management-Portable Document Format-Part 2: PDF 2.0 INTERNATIONAL STANDARD ISO 32000-2," 2017. [Online]. Available: www.iso.org

[20] C. Smutz and A. Stavrou, "When a Tree Falls: Using Diversity in Ensemble Classifiers to Identify Evasion in Malware Detectors," in *23rd Annual Network and Distributed System Security Symposium, NDSS 2016*, The Internet Society, 2016. doi: 10.14722/ndss.2016.23078.

[21] abuse.ch, "Malware Bazaar: A Malware Sharing Platform for Threat Intelligence."

A Secure Blockchain-Based Mobile Application Framework for Art Authentication and Empowerment

Himaniben Gajjar[1], Dr.JayeshJ.Gamar[2]

Assistant Professor, B.P.College of Computer Studies, Kadi SarvaVishwavidyalaya,GH Road, Sector 23, Gandhinagar, Gujarat, India[1]
Assistant Professor,Vitthalbhai Patel& Rajratna P.T.Patel Science College,Vallabh Vidyanagar, Anand, Gujarat, India[2]

prof.himani@gmail.com[1], dr.jayeshgamar@gmail.com[2]

Abstract. Digital art platforms can face issues with authentication, ownership verification, and intellectual property protection due to their centralized approach. Such issues mainly affect artists living in rural areas with less accessible digital technology. This paper will present a secure mobile application framework using blockchain technology for art authentication and empowerment.

The proposed system utilizes a mobile application, the concept of blockchain, and the idea of a smart contract to develop a tamper-proof authentication and an open and transparent management for the art ownership system. Digital art is stored in an encrypted off-chain space, and the cryptographic hashes and metadata information about the digital art is placed in the blockchain for data integrity and verification. Security analysis and performance analysis verify the efficiency and scalability of mobile art authentication in the proposed system. Through the facilitation of a secure and decentralized process for the verification of artworks by use of mobile technology, the proposed framework enables digital inclusion, which empowers artists, especially those in rural areas, to safeguard their creative properties.

Keywords: Blockchain, Mobile Application Framework, Art Authentication, Smart Contracts, Cybersecurity, Digital Provenance, Rural Empowerment, Decentralized Systems

I. INTRODUCTION

Due to rapid advancements in technology, a paradigm shift has been noted in the way that artworks are created, distributed, and monetized. Artists are now using mobile applications and websites to reach out to the world at large about their creations. Nevertheless, central systems tend to suffer from various disadvantages such as unauthorized duplication of the artworks, lack of ownership information, and data manipulation.

It is vital for any artwork to go through authentication and provenance validation. The conventional digital art marketplace relies greatly on centralized systems, which pose risks such as failure due to single point of failures and data manipulation among other security vulnerabilities. Consequently, authentication and provenance validation remain a challenge due to this vulnerability. The issue of ownership and transactions is another concern that has come with the challenges that exist in the current situation.

Blockchain technology stands out as an ideal solution to address some of these concerns due to its decentralized and tamper-proof nature. Blockchain provides a reliable and decentralized platform for keeping records of ownership and transactions without depending on any trusted parties. Smart contracts provide an additional layer of trust to this solution.Mobile applications provide an easy-to-use environment for both artists and consumers. Therefore, it is evident that mobile applications make the best platform for implementing any secure authentication system for art pieces. In addition, it should be noted that the storage of large digital artworks in blockchain is impossible. It means that off-chain encryption should be used to ensure scalability and efficiency.

Taking into account all the above-listed points, the current research proposes a new model for building an application based on blockchain technology. This application will help to ensure the recording of provenance in an immutable way and guarantee a transparent process of ownership verification and secure access to digital artworks.

The major contributions of the current research study include:

- Creation of a decentralized and mobile framework to authenticate art in a safe manner.
- The use of smart contracts to facilitate ownership management and verification automatically.
- Using off-chain private storage will help avoid blockchain's excessive overhead issues.
- It has been confirmed that the suggested method will have no security or efficiency concerns.

http://doi.org/10.63766/spujstmr.26.000076

Artistic and crafty practices offer a crucial source of income and job for many rural communities. Nevertheless, artisans are usually left out of secure digital systems that are essential for securing intellectual property rights, building trust with consumers, and so forth. In such an atmosphere, economic empowerment becomes difficult due to challenges like piracy, lack of ownership proofs, and interference by middlemen. The mobile application system developed by this team and which is capable of exploiting the potential of blockchain is specifically designed to address all these challenges, making it a safe, transparent, and decentralized system for authentication and ownership of art. Through mobile technology and blockchain documentation, artisans are getting access to a better platform for authenticating their art and engaging with digital markets.

The paper is structured as follows. Section II introduces the related work on authentication and secures systems using blockchain technology. Section 3 defines the problem statement and research goals.Section4 describes the proposed system design. Section 5discusses the security and privacy aspects of the proposed system, followed by a performance analysis in Section 6. Finally, Section 7 concludes this paper and suggests future work.

II. RELATED WORK

Blockchain technology has been widely explored for secure authentication, data integrity, and trust establishment in various application domains[1],[8].Recent studies suggest that blockchain-based frameworks can efficiently address problems of data tampering,unauthorized access, and opacity, thus facilitating digital asset management and authentication systems. There have been a number of research studies that have explored the concept of blockchain-supported secure infrastructure in resource-constrained settings. Decentralized healthcare systems based on blockchain technology prove that secure and efficient data management can be achieved using decentralized ledgers, which retain integrity and ensure control over access [1], [3].

A hybrid blockchain and lightweight authentication-based systems have been suggested for mobile and IoT environments that are efficient and usable [5], [6]. Fun and aesthetic usability of systems becomes important especially in mobile based art authentication applications. Moreover, the capabilities of smart contract-based systems allow for automating the enforcement of security policies making the execution of authentication and ownership trans- parent and uniform [6]. Innovative encryption techniques, such as quantum-secure blockchain-based authentication, have been integrated to enhance long-term security and defense against new threats. Future-proofing art authentication systems and trust in digital ownership records are possible with the help of these methods. The key block chain principles of distributed ledgers, cryptographic hash chaining, and decentralized consensus are the foundation of a secure digital ownership infrastructure [8].These features ensure that the recorded provenance data is immutable and traceable throughout the entire lifecycle of the asset. Moreover, the use of identity authentication and dual-chain architecture in blockchain-based trustworthy data-sharing platforms further enhances the traceability and verification of stakeholders [9].

TABLE I
SUMMARY OF EXISTING BLOCKCHAIN-BASED SECURE SYSTEMS

Ref.	Year	Application Domain	Key Technique	Main Contribution
[1]	2025	Healthcare	Blockchain-based infrastructure	Secure and cost-effective data management in resource-constrained environments
[2]	2025	Authentication Systems	Decentralized block chain framework	Tamper-proof and transparent authentication mechanisms
[3]	2024	Healthcare DataSharing	Multi-delegation blockchain architecture	Controlled and secure sharing of sensitive data
[4]	2024	Distributed Testing	Blockchain-enabled validation framework	Transparent and tamper-resistant testing processes
[5]	2025	IoMT Healthcare	Light weight hybrid block chain authentication	Efficient and mobile-friendly authentication

[6]	2025	IoT Networks	Blockchain with smart contracts	Automated enforcement of security policies
[7]	2025	Secure Authentication	Quantum-secure blockchain techniques	Enhanced security against future cryptographic threats
[8]	2008	Distributed Systems	Blockchain foundational model	Decentralized and immutable ledger architecture
[9]	2024	Trustworthy Data Sharing	Dual blockchain architecture	Identity-based authentication and traceability
[10]	2023	IoT & Federated Learning	Blockchain-secured Federated learning	Secure collaborative and distributed learning environment

Table I presents a concise summary of existing blockchain-based secure systems, which form the basis for identifying the research gaps addressed in this work. Although the studies summarized in Table II demonstrate the potential of blockchain-based secure systems, they differ significantly in functionality and applicability. Table II compares these features to clearly identify the research gap motivating this work.

TABLE II
FEATURECOMPARISONANDRESEARCHGAPANALYSIS

Ref.	Decentralization	Smart Contract Support	Mobile Compatibility	Security Mechanisms	Art Authentication Support
[1]	Yes	No	No	Yes	No
[2]	Yes	No	No	Yes	No
[3]	Yes	Yes	No	Yes	No
[4]	Yes	No	No	Yes	No
[5]	Yes	Yes	Yes	Yes	No
[6]	Yes	Yes	No	Yes	No
[7]	Yes	Yes	No	Yes	No
[8]	Yes	No	No	Yes	No
[9]	Yes	Yes	No	Yes	No
[10]	Yes	Yes	No	Yes	No
Proposed Work	Yes	Yes	Yes	Yes	Yes

As can be seen from Table II, the current frameworks based on blockchain technology solve issues related to security and decentralization, whereas the suggested framework has unique features, which include both mobile compatibility and art authentication capabilities.

III. PROBLEMSTATEMENTANDRESEARCH OBJECTIVES

A. Problem Statement

The popularity of digital platforms utilized in order to create, distribute, and trade artworks poses numerous difficulties regarding the issue of ownership verification, the authenticity check, as well as security in trading procedures. The majority of today's digital platforms for artwork trading are designed based on centralized systems that are susceptible to such problems as unauthorized replication, data manipulation, single point of failure, and inadequate levels of information transparency concerning ownership issues.

Moreover, the lack of permanent provenance records hampers the process of authentication of artworks, which leads to

ownership conflicts as well as violation of intellectual property rights. On the other hand, although mobile applications can be used to access digital art platforms, the majority of them provide low levels of security and have no trust models associated with decentralization. Thus, there is a necessity for development of a platform for artworks that would enable users to verify their artworks and guarantee safe transactions.

B. Research Objectives

The principal objective of this research project is to design and analyze an effective framework for a secure mobile application for art authentication and empowerment:

- To develop a decentralized mobile application framework based on blockchain technology for the secure authentication of art.
- To provide tamper-proof and immutable recording of ownership and provenance details of artworks.
- For the integration of smart contracts with a view to automating the verification process, management of ownership, as well as transaction
- To ensure safe and privacy-enhancing access to art-related data via off-chain encrypted storage.
- To assess the proposed framework interms of security, performance, and appropriateness for mobile environments.

IV. PROPOSED SYSTEM ARCHITECTURE AND SMART CONTRACT DESIGN

A. System Overview

Fig. 1 illustrates the overall architecture of the proposed blockchain-based mobile application framework for secure art authentication and empowerment. In this paper, a mobile-centric framework will be proposed wherein interaction between the artists and the buyers/verifiers will occur within the mobile application which acts as the front-end interface for registering, authenticating, and verifying artwork transactions. Digital artworks are encrypted and stored in off-chain storage, while their cryptographic hash values and metadata are recorded on the blockchain through smart contracts. The smart contract layer manages artwork registration, ownership verification, and ownership transfer in a decentralized and immutable manner. During authentication, ownership and provenance details are verified by matching the generated artwork hash with the hash value stored on the blockchain, ensuring secure and transparent artwork authentication without revealing confidential artwork content. Therefore, the proposed architecture is scalable, accessible, and suitable for secure digital art authentication and empowerment in resource-constrained and rural areas.

]

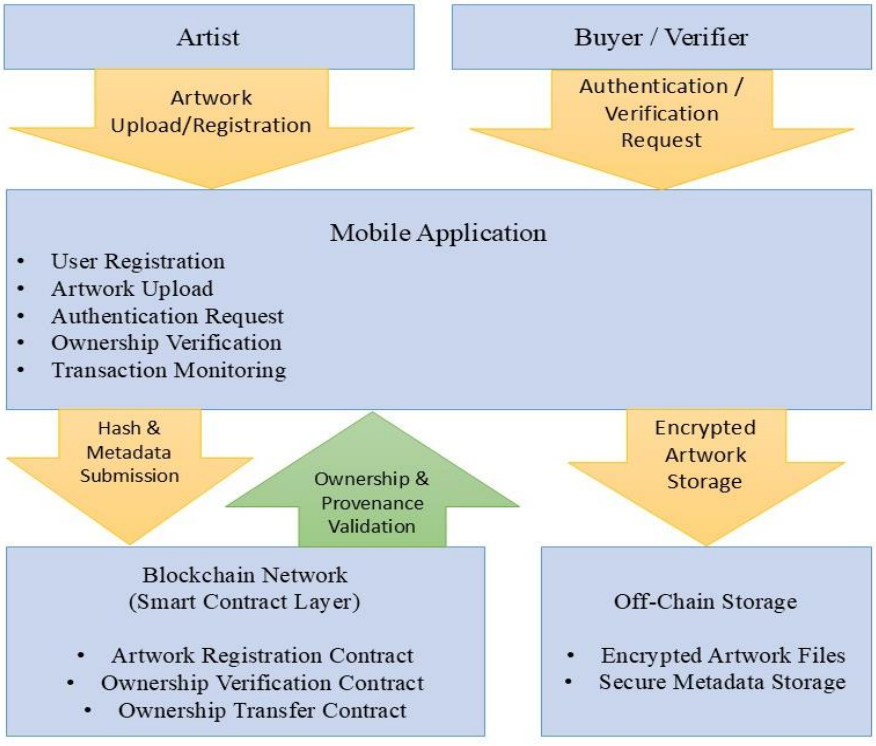

Fig.1. Architecture of the proposed blockchain-based mobile application framework for secure art authentication and empowerment. Source:

Authors' own design, created using Microsoft PowerPoint.

B. System Components

Thearchitectureismadeupoffourmainelementsthatworktogethertoprovidesecure art authentication.

MobileApplication

The main interface for artists and buyers is the mobile application. It enables the user to register for the web application, submit artwork, and verify ownership and track transactions. The application uses secure APIs to interact with the blockchain network, enabling them to authenticate the artworks and verify their provenance in real-time.

Blockchain

The decentralized ledger of a blockchain network immutably records artwork registration, ownership, and transaction history. A cryptographic hash is used to identify each artwork, ensuring integrity and preventing unauthorized modification. A decentralized system increases the trust in a system and makes it resilient.

SmartContracts

The use of smart contracts within the blockchain network helps in carrying out system functions like registration of artwork, ownership verification, and ownership transfer automatically. Smart contracts work based on pre-set rules in the blockchains making it possible for the art authentication process to have tamper-proof standards.

Off-ChainStorage

The digital files containing artwork is huge; hence, the content has been stored off-chain in an encrypted format. The blockchain only keeps the hash values and references to the content. This ensures that the amount of blockchain storage space needed is greatly minimized.

C. Implementation Details of the Proposed Framework

The mobile application framework will be implemented using the combination of the mobile front end, the blockchain network, smart contracts, and off-chain storage. In other words, the mobile application will serve as the user interface for the artists, buyers, or the verifiers. The artists will be able to register themselves, upload their artwork, and carry out ownership registration. Buyers or verifiers can authenticate the artwork and view ownership information.

The artwork file gets encrypted and saved on off-chain storage. Cryptographic hashing of the artwork is done using SHA-256. The hash value, artwork metadata, owner identity, timestamp, and off-chain storage reference get recorded on the blockchain by means of smart contracts. The blockchain does not have the full file of the artwork saved on it, making it highly scalable.

The smart contract layer fulfils three important tasks, namely artwork registration, ownership verification, and ownership transfer. While registration involves storing the hash and artwork metadata, ownership verification entails comparing the hash value of the artwork submitted against that stored on the blockchain. Finally, transfer entails changing the ownership status only when the identity of the existing owner is verified successfully.

The table below describes the implementation of the framework, including the major components and their intended purpose in the system along with suggested technology.

TABLE III
IMPLEMENTATION DETAILS OF THE PROPOSED FRAMEWORK

Component	Suggested Technology	Purpose
Mobile Application	Android / Flutter	User interface for artists, buyers, and verifiers
Blockchain Network	Ethereum / Hyperledger Fabric	Immutable storage of artwork hash and ownership records
Smart Contract	Solidity / Chaincode	Artwork registration, verification, and ownership transfer
Off-chain Storage	IPFS / Encrypted Cloud Storage	Secure storage of encrypted artwork files
Hashing Algorithm	SHA-256	Artwork integrity and authenticity verification

Component	Suggested Technology	Purpose
User Authentication	Wallet-based login / Digital identity	Secure user access and ownership control
Encryption	AES-based encryption	Protection of artwork before off-chain storage

D. Smart Contract Design

The core method of execution for the designed framework would be the use of smart contracts, which implement authentication and ownership policies based on certain conditions on the blockchain. Agreement for New Art on Blockchain: This smart contract ensures the recording of the art piece on the blockchain network. When an art piece is uploaded to the blockchain system, a hash, which is unique to the artwork, along with other metadata is created and stored on the blockchain.
Ownership Verification Contract: The purpose of this smart contract is to verify the authenticity of the art piece using the hash value submitted and hash value present in the blockchain.
Ownership Transfer Contract: Using this smart contract, it is possible to manage the ownership of the artwork in a secure and auditable manner.

E. Workflow of the Proposed Framework

Operation procedure of integrated architecture:

- The artists use the mobile app for registering and uploading artworks.
- Artworks are stored on an off-chain system after encryption, and a digital hash code is generated.
- The hash code and its associated metadata are registered using a smart contract for artwork registration on blockchain.
- Buyers/verifiers of the artwork will check the ownership of artwork using blockchain from the mobile app.
- Ownership of artwork is transferred through smart contracts, and transaction updates are also done automatically using smart contracts.This process offers safe credentials, clear ownership, and convenient access through mobile apps.

V. SECURITY AND PRIVACY ANALYSIS

The blockchain-based mobile application framework proposed tackles several security and privacy challenges involved in the authentication, verification of ownership, and transaction processing for the digital art pieces. The security analysis will be based on the factors such as data integrity and access control, privacy protection, risks at the blockchain level, and vulnerabilities in the smart contracts.

A. Data Integrity and Access Control: In the proposed system, the hash values of the artworks and other relevant ownership data are stored in the blockchain. As we know that once the data is recorded in the blockchain, it is practically impossible to change. Every artwork has a unique cryptographic hash value that changes when there is any modification in the artwork file. Hence, the integrity of the art piece can be checked by comparing the hash values of the artwork stored on the blockchain against the new one.

B. Privacy Preservation and Ownership Protection: The entire artwork is not directly uploaded to the blockchain. Rather, the artwork is encrypted and stored in an off-chain storage location, whereas the hash values, metadata, and storage location are stored in the blockchain itself. This helps reduce the amount of sensitive artwork data exposed and ensures user privacy. The usage of encrypted off-chain storage also saves blockchain storage resources while maintaining artwork integrity.

C. Blockchain and Smart Contract Attack Analysis: While blockchain is a decentralized and immutable technology, there exist various types of threats for any blockchain-based system. Possible attacks and solutions are therefore considered within the proposed solution framework.

TABLE IV
BLOCKCHAIN AND SMART CONTRACT ATTACK ANALYSIS

Attack / Threat	Possible Risk	Mitigation Strategy
51% Attack	Malicious control over blockchain consensus may alter transaction history	Use of a permissioned or consortium blockchain with trusted validators

Attack / Threat	Possible Risk	Mitigation Strategy
Sybil Attack	Fake identities may create false ownership or verification requests	Strong identity verification and wallet-based authentication
Replay Attack	Reuse of old, valid transactions may cause unauthorized actions	Use of nonce, timestamp, and transaction validation
Private Key Compromise	Unauthorized ownership transfer may occur if user keys are stolen	Secure wallet management, multi-factor authentication, and a key recovery mechanism
Reentrancy Attack	A smart contract may be repeatedly called before a state update	Use of the checks-effects-interactions pattern and reentrancy guards
Front-running Attack	Transaction order manipulation may affect ownership transfer	Commit-reveal mechanism and transaction ordering control
Integer Overflow / Underflow	Incorrect contract execution due to arithmetic errors	Use of safe arithmetic libraries and updated compiler versions
Smart Contract Bugs	Incorrect ownership logic may lead to false authentication	Smart contract testing, auditing, and formal verification
Denial-of-Service Attack	Attackers may overload verification or transaction requests	Rate limiting, access control, and distributed storage support
Metadata Leakage	Public metadata may reveal sensitive ownership information	Store only the minimum required metadata on-chain

D. Overall Security Justification: Blockchain immutability, smart contracts' automation, encryption for off-chain data storage, and hashing provide better security for the proposed framework. This model lessens the dependence on centralized authorities and gives clear evidence of the integrity of ownership proofing. However, there is a need for adequate testing, key management, identity validation, and setup of a blockchain network.

VI. PERFORMANCE EVALUATION

Performance assessment of the proposed blockchain-based mobile application framework is performed according to certain metrics that can be used in the mobile environment and decentralized environment. Metrics considered during assessment include storage overhead, transaction latency, smart contracts execution fee, mobile device response time, and scalability.
The proposed framework stores only hash values of the artwork, its metadata, time stamp, ownership information, and storage information on the blockchain, while the artwork itself is stored offline.

TABLE V
QUANTITATIVE AND FUNCTIONAL PERFORMANCE COMPARISON

Performance Metric	Centralized Art Platform	Proposed Blockchain-Based Framework
Artwork storage	Full file stored centrally	Encrypted file stored off-chain
Storage overhead	High	Low
Ownership verification	Platform-dependent	Smart contract-based
Transaction transparency	Limited	High
Tamper resistance	Medium	High
Provenance tracking	Limited	Immutable and traceable

Performance Metric	Centralized Art Platform	Proposed Blockchain-Based Framework
Scalability	Medium	High
Mobile accessibility	Platform-dependent	Mobile-centric

Based on the comparative parameters presented in Table V, Fig. 2 illustrates the relative performance of the centralized art platform and the proposed blockchain-based mobile application framework.

Fig. 2 presents relative comparative scores based on framework-level analysis. The proposed framework offers better storage management, transparency, provenance tracking, tamper resistance, and scalability by combining blockchain records, smart contracts, cryptographic hashing, and encrypted off-chain storage.
Overall, the proposed framework is suitable for secure and mobile-based art authentication, particularly for artists and users in resource-constrained environments. Full prototype implementation, gas cost analysis, transaction confirmation time, and real-time mobile response validation are considered as future work.

VII. CONCLUSION AND FUTURE WORK

This paper describes a secure and blockchain-based mobile application framework for art authentication and empowerment. The proposed framework uses blockchain technology and smart contracts to provide tamper-proof and transparent ownership information while ensuring secure and decentralized authentication of digital artworks. The two main objectives of using mobile connectivity with off-chain encrypted storage to contest unauthorized copying, ownership and trust in centralized art platforms will be proposed in this work. The framework can deliver significant rural empowerment benefits that can help artists and artisans located in rural areas, which often suffer from a lack of secure digital infrastructure and intellectual property protection. With the mobile-friendly framework design, artists can verify their artworks via accessible mobile devices. At the same time, using blockchain technology in provenance creates trust and transparency without any middleman. The framework's ability to preserve the integrity of data, along with its confidentiality and access control using cryptographic techniques, is proven by security analysis. The proposed framework aids in promoting secure digital livelihoods, equitable access to technology, and economic empowerment of artists, especially in rural areas, in conformity to the Sustainable Development Goals or SDG 8 (Decent Work and Economic Growth), SDG 9 (Industry, Innovation and Infrastructure), and SDG 10 (Reduced Inequalities). The performance analysis shows that the proposed approach is suitable for mobile platforms. Further, it ensures a reasonable transaction delay, low storage cost, and the response of the system is also optimum, which shows the feasibility of the framework in the real world. In future studies, this framework can be enhanced by incorporating artificial intelligence methods for the automatic analysis of artwork similarity, forgery detection, and artwork style classification. Extensions may include cross-chain interoperability, integration with decentralized art markets, or sophisticated access control policies for community art environments. This will help make the proposed framework more effective in creating secure and scalable digital platforms for empowering rural areas.

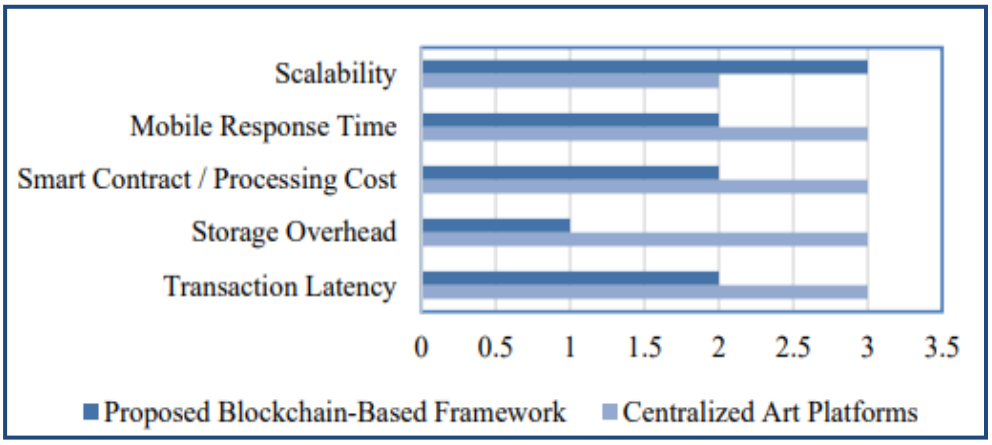

Fig. 2. Relative performance comparison of the proposed blockchain-based framework and centralized art platforms across transaction latency, storage overhead, processing cost, mobile response time, and scalability. Source: Authors' own analysis.

REFERENCES

[1] J. Lafuente et al., "SiliconHealth: A blockchain-based secure healthcare infrastructure for resource-constrained environments," arXiv preprint arXiv:2601.09557, 2025.
[2] X. Zhang et al., "Blockchain adoption for secure authentication in energy trading systems," Energy Research & Social Science, vol. 113, 2025.
[3] Kumar et al., "Blockchain-based secured data sharing using multi-delegation architecture in healthcare systems," IEEE Access, 2024.
[4] M. Rossi et al., "BCTestingPlus: A blockchain-empowered framework for distributed and transparent testing," Inventions, vol. 10, no. 4, 2024.
[5] Ahmed et al., "A lightweight hybrid authentication framework for IoMT-based e-healthcare applications," Scientific Reports, vol. 15, 2025.
[6] R. Patel et al., "A hybrid blockchain and smart contract framework for secure IoT networks," Frontiers in Blockchain, 2025.
[7] Y. Chen et al., "Quantum-secure blockchain-based authentication for future digital systems," Scientific Reports, vol. 15, 2025.
[8] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008.
[9] H. Liu et al., "Blockchain-empowered trustworthy data sharing with identity-based authentication," ACM Transactions on Internet Technology, 2024.
[10] H. Wu et al., "Blockchain-empowered secure federated learning for IoT applications," IEEE Transactions on Consumer Electronics, 2023.
[11] G. Wood, "Ethereum: A secure decentralized generalized transaction ledger," Ethereum Yellow Paper, 2014.
[12] K. Christidis and M. Devetsikiotis, "Blockchains and smart contracts for the Internet of Things," IEEE Access, 2016.
[13] M. Crosby et al., "Blockchain technology: Beyond bitcoin," Applied Innovation Review, 2016.
[14] J. Moubarak, E. Filiol, and M. Chamoun, "On blockchain security and relevant attacks," in Proc. IEEE International Conference on Internet Technology and Secured Transactions, 2018.
[15] M. Casino, T. K. Dasaklis, and C. Patsakis, "A systematic literature review of blockchain-based applications: Current status, classification and open issues," Telematics and Informatics, 2019.
[16] P. Kaur et al., "Blockchain-based mobile applications: A survey," IEEE Access, 2020.
[17] A. Dorri, S. S. Kanhere, and R. Jurdak, "Blockchain in Internet of Things: Challenges and solutions," Computer Communications, 2017.

Multifunctional Smart Real Estate Platform with Predictive and Utility Features

Ravijeet Sharma[1], Vansh Verma[2], Harsh Vyavahare[3], Mrs. Jagruti Jadhav[4]

Research Scholar, Department of Artificial Intelligence and Machine Learning Thakur College of Engineering and Technology, Mumbai, India[1]
Research Scholar, Department of Artificial Intelligence and Machine Learning Thakur College of Engineering and Technology,Mumbai, India[2]
Research Scholar, Department of Artificial Intelligence and Machine Learning Thakur College of Engineering and Technology,Mumbai, India[3]
Assistant Professor, Department of Artificial Intelligence and Machine Learning Thakur College of Engineering and Technology, Mumbai, India[4]

ravijeetsharma180@gmail.com[1],9155vansh@gmail.com[2],harshvyavahare29@gmail.com[3], jagruti.jadhav@tcetmumbai.in[4]

Abstract: The global real estate market, characterized by fragmentation and information asymmetry, is ripe for digital transformation. This paper proposes an AI-enhanced unified platform designed to seamlessly integrate sales and rental operations, offering a comprehensive and intelligent solution for all stakeholders. Leveraging advanced Artificial Intelligence (AI) techniques, including machine learning for predictive analytics, natural language processing for semantic search and content analysis, computer vision for property evaluation, and sophisticated recommender systems, the proposed platform aims to significantly enhance market transparency, operational efficiency, and user experience. We outline a prototype-driven framework, detail the integrated AI components, discuss the methodological considerations for development, and enumerate the profound benefits this unified approach offers over traditional, siloed marketplaces. Furthermore, we address the inherent challenges and delineate future research directions, positioning this platform as a transformative force in the evolving real estate landscape. To validate the feasibility of the proposed framework, a functional prototype was developed using a cloud-enabled architecture comprising React/Next.js frontend services, AWS Cognito authentication, Amazon S3 storage, Mapbox-based geospatial visualization, and an Express.js backend with PostgreSQL data management. Preliminary experimentation with machine learning models for property price prediction and text classification was also conducted to assess AI integration feasibility.

Keywords: Artificial Intelligence, Real Estate, Unified Platform, Machine Learning, Predictive Analytics, Natural Language Processing, Computer Vision, Recommender Systems, Digital Transformation.

I. INTRODUCTION

The real estate sector, a cornerstone of global economies, has historically been characterized by its localized nature, information opacity, and reliance on intermediaries. While digital platforms have emerged to address some of these inefficiencies, the market remains largely fragmented, with distinct ecosystems for property sales and rentals. Buyers, sellers, renters, and landlords often navigate multiple platforms, leading to duplicated efforts, inconsistent data, and a disjointed user experience. The lack of a comprehensive, intelligent system that caters to the entire spectrum of real estate transactions within a single, unified environment presents a significant challenge and a ripe opportunity for innovation [1].

Traditional real estate processes are often time-consuming, labor-intensive, and susceptible to human error and biases. From property valuation and market analysis to lead generation, tenant screening, and contract management, each step can benefit from technological augmentation. The advent of Artificial Intelligence (AI) offers a powerful suite of tools capable of revolutionizing these processes, moving beyond simple digitization to intelligent automation and personalized insights [2]. AI can sift through vast datasets, identify complex patterns, make predictions, and even automate routine tasks, thereby enhancing efficiency, accuracy, and accessibility across the real estate value chain.

http://doi.org/10.63766/spujstmr.26.000077

This paper proposes an AI-enhanced unified platform designed to bridge the chasm between property sales and rental markets. Our vision is a singular, intelligent marketplace where users can seamlessly transition between buying, selling, renting, or leasing properties, supported by an array of AI functionalities that provide personalized recommendations, accurate valuations, insightful market trends, and streamlined transaction processes. Such a platform would not only consolidate market information but also inject a layer of intelligence that empowers more informed decision-making for all participants.

The primary objectives of this research are:

1. To conceptualize an architectural framework for a unified AI-enhanced real estate marketplace that integrates sales and rental functionalities.
2. To identify and detail the critical AI components required to deliver intelligent features across the platform.
3. To outline the methodological considerations for the development and implementation of such a system.
4. To discuss the potential benefits and transformative impact of a unified AI platform on market transparency, operational efficiency, and user experience.
5. To address the significant challenges and propose future research directions for the continued evolution of AI in real estate.

The structure of this paper is as follows: Section II reviews the existing literature and related work in real estate technology and AI applications. Section III presents the proposed architectural framework and elaborates on the core AI components. Section IV details the methodological considerations for realizing this platform. Section V discusses the potential benefits, while Section VI addresses the challenges and future directions. Finally, Section VII concludes the paper by summarizing its contributions and reiterating the vision for an intelligent, unified real estate marketplace

II. LITERATURE REVIEW AND RELATED WORK

The digital transformation of the real estate industry has been a gradual process, evolving from static listing websites to dynamic marketplaces and, more recently, to platforms incorporating advanced analytics and AI. This section reviews the current landscape, highlighting the state of existing platforms and the application of AI in various real estate functions.

A. Traditional and Digital Real Estate Platforms

Early digital real estate platforms primarily served as online directories, aggregating property listings for sales or rentals. Websites like Zillow, Trulia, Rightmove, and Realtor.com have become household names, offering search functionalities, property details, and sometimes basic market data [3]. However, these platforms often operate with distinct sections or entirely separate portals for sales and rentals, creating a segmented experience. While they have improved information accessibility, they still largely rely on manual data input, human agents for negotiations, and rudimentary search algorithms. Furthermore, the accuracy of property valuations (e.g., Zestimate by Zillow) has been a subject of debate, underscoring the limitations of rule-based or purely statistical models without deeper AI integration [4]. The fragmentation extends to other services, such as property management tools, mortgage calculators, and legal documentation, which are typically handled by third-party providers, further exacerbating the disjointed user journey.

B. AI Applications in Real Estate

The application of AI in real estate has seen a surge in recent years, primarily focusing on specific pain points rather than holistic platform integration. Key areas include:

Automated Valuation Models (AVMs): Machine learning models, such as linear regression, random forests, and gradient boosting, are widely used to predict property prices based on features like location, size, number of rooms, and historical sales data [5]. Deep learning models, particularly Convolutional Neural Networks (CNNs), have also been employed to analyze satellite imagery and street-level views to infer neighborhood characteristics and property conditions, enhancing valuation accuracy [6]. However, most AVMs are designed either for sales or rental valuation, not a unified approach that considers the interplay between the two markets.

Property Search and Recommendation Systems: AI-driven recommender systems leverage collaborative filtering, content-based filtering, and hybrid approaches to suggest properties to potential buyers or renters. These systems analyze user preferences, search history, clickstream data, and demographic information to provide personalized recommendations [7]. Natural Language Processing (NLP) is used to parse property descriptions, extract key features, and enable semantic search, allowing users to query properties using natural language rather than strict keywords [8]. Despite these advancements, achieving truly intelligent matching that considers subtle user needs and predicts future market changes remains an area of active research.

Market Analytics and Trend Prediction: AI algorithms can analyze vast amounts of structured and unstructured data—including socio-economic indicators, demographic shifts, urban development plans, and sentiment from news and social media—to forecast market trends, identify investment opportunities, and predict price fluctuations [9]. Time series analysis techniques, combined with deep learning architectures like LSTMs, are particularly effective for dynamic market forecasting [10].

Property Management and Operations: AI is also being applied in areas like predictive maintenance for rental properties, tenant screening (using ML to assess creditworthiness and behavioral patterns), and automated communication with tenants via chatbots [11].

Computer Vision for Property Analysis: Beyond AVMs, computer vision is used to analyze interior and exterior property images, identifying features, assessing quality, and even generating virtual tours. It can detect outdated fixtures, predict renovation costs, and highlight appealing attributes, augmenting human judgment [12].

C. Gaps in Existing Solutions

Despite these individual advancements, a significant gap remains: the lack of a truly unified, intelligent platform that seamlessly integrates sales and rental functionalities with comprehensive AI augmentation across the entire user journey. Current platforms are often siloed, forcing users to switch contexts and re-enter information. AI applications, while powerful in specific niches, are rarely integrated into a cohesive system that offers end-to-end support for both buying/selling and renting/leasing. The dynamic interplay between sales prices and rental yields, market liquidity, and investment opportunities is often analyzed separately, failing to provide a holistic view for users making complex real estate decisions. This fragmentation leads to inefficiencies, missed opportunities for cross-marketing, and a less-than-optimal user experience. Our proposed unified platform aims to address these critical shortcomings by providing a singular, intelligent ecosystem.

Key Advantage of the Proposed Platform

Existing platforms primarily function as listing marketplaces with limited AI integration. Their predictive capabilities are generally restricted to property valuation, while rental analysis, market forecasting, and personalized decision support remain fragmented. The proposed platform differentiates itself by integrating:

1. Sales and rental ecosystems within a single platform.
2. Advanced AI-driven predictive analytics for both property prices and rental yields.
3. Computer vision and NLP modules for automated property assessment and semantic search.
4. Cross-market recommendations, enabling users to compare buying and renting opportunities.
5. Fraud detection and trust mechanisms to improve platform reliability.
6. End-to-end intelligent decision support rather than acting solely as a property listing portal.

One more important observation: several references in the paper appear to be incomplete, inconsistently formatted, or difficult to verify (especially references [13]–[16]). Before final submission, I recommend verifying every citation against authentic IEEE-indexed or Scopus-indexed sources and reformatting them strictly according to the journal's reference style. This is likely what the reviewer meant by "reference formatting."

III. ARCHITECTURAL FRAMEWORK AND AI COMPONENTS

The proposed AI-enhanced unified real estate marketplace is envisioned as a modular, scalable, and cloud-native platform designed to serve diverse user needs across sales and rentals (miro.com). Figure 1 illustrates a high-level architectural overview, highlighting the core components.

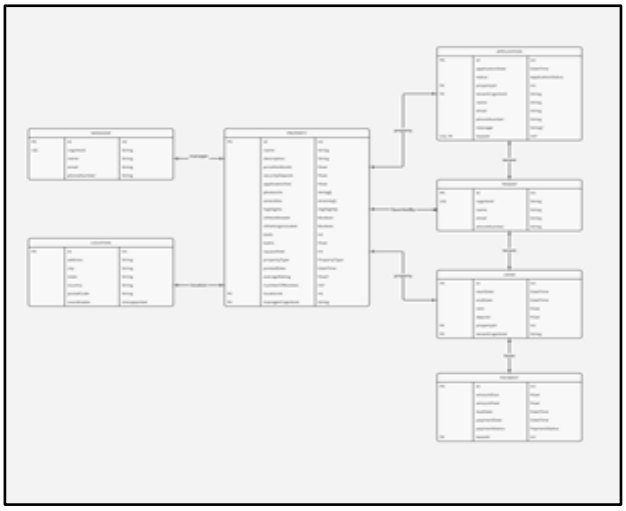

Fig 1: Data Flow Diagram (created by using Miro)

A. High-Level Architecture

User Interface Layer: Provides a responsive, intuitive, and personalized interface accessible via web browsers and mobile applications. It seamlessly integrates features for listing properties, searching, viewing, communicating, and managing transactions for both sales and rentals.

Application Services Layer*:* Comprises microservices responsible for core business logic, such as user authentication, listing management, messaging, transaction tracking, property management features (for landlords/renters), and agent services. This layer acts as the orchestrator, invoking AI services as needed.

AI/ML Services Layer: This is the intelligence core, housing various specialized AI models and algorithms. It interacts with the Data Management Layer for data retrieval and model training/inference, and provides API endpoints for the Application Services Layer to consume AI-driven insights and functionalities.

Data Management Layer: A robust and secure data infrastructure that stores heterogeneous data, including property listings, transaction history, user profiles, behavioral data, geospatial information, socio-economic indicators, and external market data. It uses a combination of relational databases (for structured data) and NoSQL databases (for unstructured data like property descriptions and image metadata).

External Data Sources: Integrates with third-party APIs for demographics, crime rates, school ratings, public transport, interest rates, and other relevant market data to enrich the platform's analytical capabilities.

Security and API Gateway: Ensures secure access, manages API requests, and enforces authentication and authorization protocols across all layers.

B. High-Level Architecture

The strength of the unified platform lies in its integrated AI capabilities, each designed to enhance specific aspects of the real estate journey for both sales and rental markets:

Predictive Analytics Engine:

Functionality: Utilizes supervised machine learning (e.g., Gradient Boosting Machines, Random Forests, Deep Neural Networks) to predict property sale prices, rental yields, and future market trends.

Data Inputs: Historical sales/rental transaction data, property attributes (size, age, type, amenities), location data (geospatial coordinates, neighborhood characteristics), socio-economic indicators (income levels, employment rates), public transport accessibility, school ratings, and recent market activity.

Unified Aspect: A single engine generates forecasts for both sales and rentals, allowing users (buyers, investors, renters) to compare the intrinsic value and potential returns of a property across different usage models. For instance, an investor can quickly assess if a property's rental yield justifies its purchase price.

Intelligent Recommender System:
Functionality: Employs hybrid recommendation techniques (content-based, collaborative filtering, and knowledge-based) to suggest highly relevant properties (for sale or rent) to users. It also suggests agents, mortgage brokers, or property management services.
Data Inputs: User search history, viewed properties, saved listings, expressed preferences (e.g., desired amenities, budget, lifestyle), demographic information, and implicit feedback (e.g., click-through rates, time spent on listing).
Unified Aspect: Learns from user behavior across both sales and rentals. A user browsing for a rental might be subtly alerted to potential buying opportunities that match their preferences, or vice-versa. This facilitates cross-market conversion and holistic decision-making.

Natural Language Processing (NLP) Module:
Functionality:
Semantic Search: Allows users to query properties using natural language (e.g., "family home near good schools with a large garden" or "pet-friendly apartment close to metro").
Listing Analysis: Extracts structured data from unstructured property descriptions (e.g., "newly renovated kitchen," "hardwood floors," "walk-in closet") to enrich property features for search and recommendations.
Content Moderation and Safety Analysis: Analyzes reviews, public comments, and news articles related to specific properties or neighborhoods to provide sentiment scores.
Chatbot& Virtual Assistant: Provides 24/7 automated support, answers FAQs, schedules viewings, and guides users through the platform.
Contract Review & Anomaly Detection: (Advanced) Scans lease agreements or sales contracts for missing clauses, non-standard terms, or potential legal risks.
Unified Aspect: The NLP module processes all text data regardless of whether it pertains to sales or rentals, maintaining a consistent understanding of property attributes and user intent across the platform.

Computer Vision (CV) Module:
Functionality:
Image Quality Assessment: Filters out low-quality or irrelevant images from listings.
Feature Extraction: Identifies key property features from images (e.g., kitchen style, bathroom fixtures, presence of a pool, garden size) to augment structured data.
*Virtual Staging & Renovation Visualization***:** (Advanced) Uses generative AI to virtually stage empty rooms or visualize potential renovations, assisting both buyers and sellers/landlords.
Anomaly Detection: Identifies discrepancies between image content and textual descriptions (e.g., description says "newly renovated" but images show old fixtures).
Unified Aspect: Applies equally to sales and rental listings, ensuring high-quality visual representation and consistent feature extraction for all property types.

Fraud Detection and Trust Engine:
Functionality: Employs anomaly detection and classification algorithms to identify suspicious listings, fraudulent user accounts, or unusual transaction patterns.
Data Inputs: User behavior analytics, listing history, identity verification data, communication patterns, and reported incidents.
Unified Aspect: Crucial for maintaining trust across both sales and rental markets, protecting both buyers/renters and sellers/landlords from scams and misrepresentations.

Personalization Engine:
Functionality: A cross-cutting AI component that leverages insights from all other modules to tailor the entire user experience—from dashboard layout and notification preferences to curated content and service suggestions.
Unified Aspect: Provides a consistent, adaptable experience whether the user is in a "buy," "sell," "rent," or "lease" mindset, smoothly guiding them through relevant options.

IV. METHODOLOGY

The development of such an AI-enhanced unified real estate marketplace necessitates a robust and iterative methodology, emphasizing data-driven design, agile development, and continuous improvement Figure 2 illustrates a high-level architectural overview, highlighting the working of Aws.

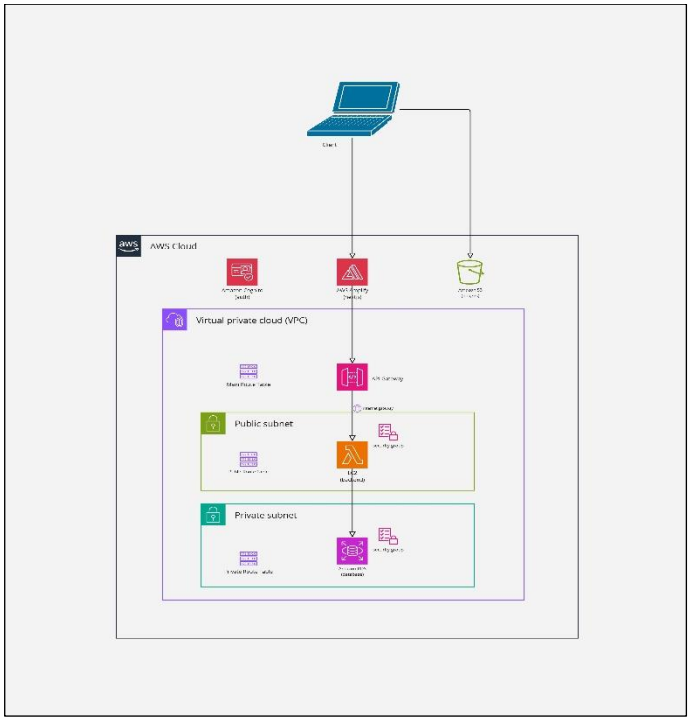

Fig2: AWS Architecture (created using Miro and AWS icons)

A. Data Acquisition and Management

The foundation of any AI-driven platform is high-quality, comprehensive data.

Data Sources:

Internal Data: User-generated content (listings, reviews, messages), user behavior (search queries, clickstream, saved properties), transaction history (sales and rental contracts, prices, dates).

External Data: Geospatial data (GIS maps, zoning information), demographic statistics (population density, age distribution, income levels), economic indicators (interest rates, inflation), public records (property deeds, tax assessments), social media sentiment, news articles, and local amenities data (schools, hospitals, public transport).

Data Integration and Preprocessing:A unified data lake architecture will be employed to ingest and store raw data from diverse sources. ETL (Extract, Transform, Load) pipelines will clean, normalize, and integrate this data into a coherent format suitable for AI model training. This includes handling missing values, standardizing units, and resolving inconsistencies. Data anonymization and pseudonymization will be applied to protect user privacy.

Feature Engineering: This crucial step involves creating new, more informative features from raw data to improve AI model performance. Examples include calculating property age, distance to nearest amenities, neighborhood walkability scores, and rental yield ratios.

B. AI Model Development and Training

Model Selection: Based on the specific task (prediction, recommendation, classification), appropriate machine learning and deep learning algorithms will be selected. For instance, Gradient Boosting Machines (XGBoost, LightGBM) for tabular predictive tasks, Recurrent Neural Networks (RNNs) or Transformers for NLP, and Convolutional Neural Networks (CNNs) for computer vision.

Training Data Preparation: A significant portion of the integrated and preprocessed data will be used to train the AI models. This involves splitting data into training, validation, and test sets to ensure robust model evaluation.

AI Models, Dataset Description, and Experimental Setup: To evaluate the feasibility of AI integration within the proposed unified real estate marketplace, proof-of-concept machine learning experiments were conducted for rental price prediction and automated content moderation. These experiments aimed to assess the practicality of incorporating predictive analytics and intelligent text monitoring into the platform architecture. A rental property price prediction model was developed using Extreme Gradient Boosting (XGBoost Regression) due to its strong performance on structured tabular datasets and ability to model nonlinear feature relationships. Multiple regression algorithms, including Decision Tree Regression, Random Forest Regression, Extra Trees Regression, AdaBoost Regression, and Gradient Boosting Regression, were comparatively evaluated during experimentation. Among these, XGBoost demonstrated the best predictive performance and was therefore selected for further

evaluation. The rental prediction model was trained using a structured real-estate dataset containing approximately 20,500 records and evaluated on an independent testing dataset of approximately 4,500 records. The dataset included property-related attributes such as locality, geographic coordinates, lease type, furnishing status, parking availability, property size, property age, bathroom count, floor information, amenities, water supply, and building type. The target variable consisted of rental valuation in Indian Rupees (INR), enabling localized pricing prediction for the Indian real estate market. Experimental evaluation of the XGBoost model demonstrated strong predictive capability, achieving an R^2 score of 0.81, Mean Absolute Error (MAE) of 2648.41 INR, and Root Mean Squared Error (RMSE) of 3745.29 INR. These findings indicate that AI-assisted pricing estimation can provide useful decision support for buyers, renters, and property owners while reducing pricing uncertainty. In addition to predictive pricing, automated text moderation models were evaluated to improve platform safety and communication quality. A content moderation dataset containing approximately 24,500 textual samples was utilized for binary text classification between appropriate and inappropriate content. Data preprocessing included text normalization, tokenization, stop-word removal, stemming, URL cleaning, and sequence padding. Two deep learning approaches, Bidirectional Encoder Representations from Transformers (BERT) and Long Short-Term Memory (LSTM), were comparatively evaluated for text classification. Experimental results showed that the BERT-based model achieved superior performance, attaining an accuracy of 95.95%, precision of 97.76%, recall of 97.36%, and F1-score of 97.56%. The LSTM model also demonstrated strong performance with an accuracy of 93.56% and F1-score of 96.15%. These results indicate the feasibility of deploying intelligent moderation systems to automatically detect abusive, spam-related, or inappropriate content within real-estate listings and user interactions.

TABLE I
AI MODELS

Module	**Model**	**Purpose**	**Outcome**
Rental Price Prediction	XGBoostRegressor	Property rent estimation	Selected as best-performing regression model
Content Moderation	BERT	Inappropriate text detection	Highest classification performance
Content Moderation	LSTM	Text classification baseline	Competitive performance

TABLE II
PERFORMANCE METRICS

Model	**Accuracy / R^2**	**Precision**	**Recall**	**F1 Score**	**RMSE**
XGBoost	$R^2 = 0.81$	-	-	-	3745.29
BERT	95.95%	97.76%	97.36%	97.56%	-
LSTM	93.56%	95.57%	96.75%	96.15%	-

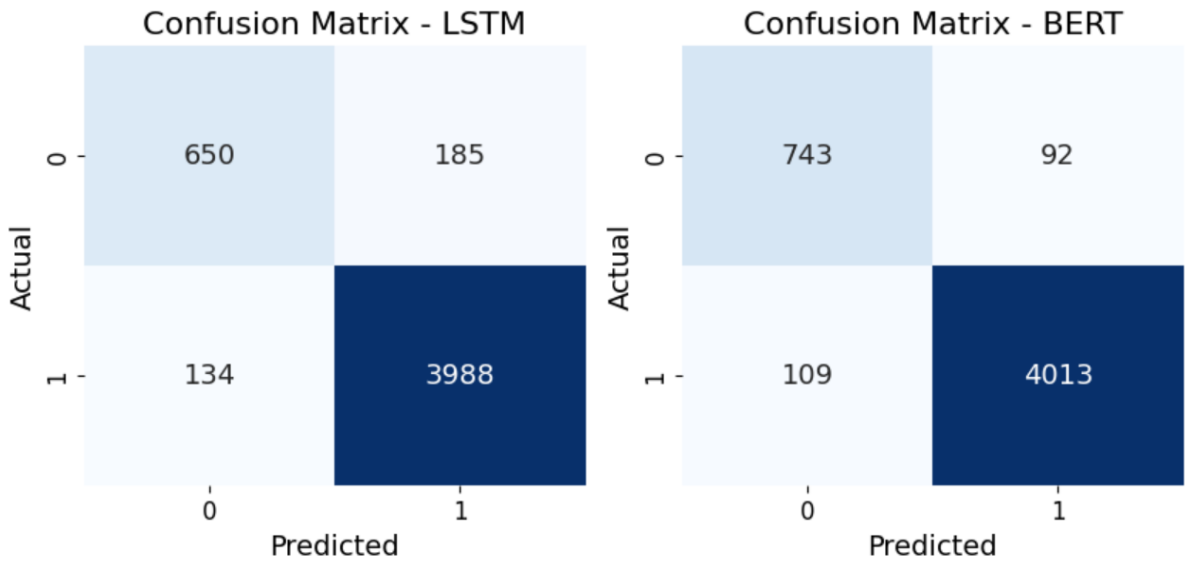

Fig3: Comparative Confusion Matrix of BERT and LSTM-based Content Moderation Models

Evaluation Metrics: Appropriate metrics used to evaluate model performance:
Predictive Analytics: Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), R-squared (for price prediction).
Recommendation Systems: Precision, Recall, F1-score, Normalized Discounted Cumulative Gain (NDCG), A/B testing for user engagement.
Content Moderation: Accuracy, F1-score (for classification tasks).
Computer Vision: Precision, Recall, Intersection over Union (IoU) for object detection.

Model Deployment and Monitoring: Trained models will be deployed as microservices using containerization technologies (e.g., Docker, Kubernetes) to enable scalable inference. Continuous monitoring of model performance (drift detection) and retraining mechanisms will be implemented to ensure models remain accurate and relevant over time with evolving market conditions.

C. System Design Principles
Scalability: The architecture must be designed to handle a rapidly growing user base and increasing data volumes, leveraging cloud-native services and elastic computing.

Modularity: A microservices architecture ensures that different components (e.g., predictive engine, recommender system) can be developed, deployed, and scaled independently, facilitating agile development and maintenance.
API-First Approach: All AI services will expose well-documented APIs, enabling seamless integration with the application services layer and potential third-party applications.

Security and Privacy-by-Design: Adherence to data protection regulations (e.g., GDPR, CCPA) is paramount. Encryption of data at rest and in transit, robust access controls, regular security audits, and privacy-enhancing technologies will be integrated from the outset.

User-Centric Design (UI/UX): The user interface will be intuitive, responsive, and provide a seamless, unified experience across devices and transaction types (sales/rentals). Personalization will be a core tenet, adapting the interface and information displayed based on user preferences and behavior.

C. Prototype Implementation and Practical Validation
To evaluate the feasibility of the proposed AI-enhanced unified real estate marketplace, a functional prototype was developed and tested in a local development environment. The implementation focused on validating the architectural design, cloud

integration, user interaction flow, and geospatial property exploration features. The frontend of the prototype was developed using Next.js and React to provide a responsive and interactive user experience. State management mechanisms were incorporated to maintain user preferences, listing interactions, and application flow. The backend services were implemented using Express.js and integrated with cloud-native AWS services to ensure modularity and scalability. User authentication and account verification were managed through AWS Cognito, including secure email verification workflows. Property images and media assets were stored and retrieved using Amazon S3, enabling efficient cloud-based storage management.

A PostgreSQL-compatible cloud database architecture was implemented using Neon and Prisma ORM for structured storage and querying of user profiles, property listings, and transactional metadata. Security-focused middleware, including token-based authentication and API protection mechanisms, was integrated to improve system reliability and secure access to platform resources. To enhance spatial property discovery, the platform incorporated Mapbox-based geospatial services, enabling users to search and explore properties interactively through map visualization. Geographic search functionality allowed users to identify listings based on location-specific preferences, improving discoverability and contextual decision-making. Although the complete production deployment of AI modules was not finalized, proof-of-concept experimentation was conducted locally to evaluate AI feasibility within the proposed platform.

A property price prediction model was trained and tested using housing-related datasets to estimate property valuation in Indian currency contexts. In addition, preliminary text classification experiments were conducted to assess automated content understanding and moderation possibilities for future integration, such as spam and inappropriate content detection. Practical validation of the platform was performed through functional testing and scenario-based evaluation using seeded property datasets. Testing focused on major workflows, including user registration, authenticated access, property listing creation, image upload, geographic property exploration, and filtered property search. Results demonstrated the technical feasibility of integrating cloud services, geospatial technologies, and AI-assisted components into a unified real estate ecosystem.

TABLE III
TECHNOLOGY STACK

Component	**Technology Used**	**Purpose**
Frontend	Next.js, React	User interface
Backend	Express.js	API and business logic
Authentication	AWS Cognito	Login & email verification
Storage	Amazon S3	Property images
Database	Neon PostgreSQL + Prisma	Listing and user data
Geospatial Services	Mapbox	Property search by location
AI Experimentation	Local ML models	Price prediction and text classification

V. DISCUSSION AND POTENTIAL BENEFITS

The implementation of an AI-enhanced unified real estate marketplace promises a paradigm shift in how properties are bought, sold, and rented. The benefits extend across all stakeholders, fostering a more transparent, efficient, and user-friendly ecosystem.

A. Enhanced Market Transparency and Informed Decision-Making By consolidating data from both sales and rental markets and augmenting it with AI-driven insights, the platform provides an unprecedented level of market transparency.
For Buyers/Renters: Access to accurate, real-time valuations (sales and rental), comprehensive neighborhood data, future market predictions, and comparative analysis tools empowers them to make more informed decisions, whether they are looking to purchase a long-term home or rent temporarily. They can easily compare the cost of buying versus renting in a given area.
For Sellers/Landlords: AI-driven pricing recommendations, and demand forecasting enable them to set competitive prices, optimize listing strategies, and understand the best time to sell or rent.
For Investors: The platform serves as a powerful analytical tool, identifying high-potential investment properties by analyzing rental yield potential against purchase price, vacancy rates, and capital appreciation forecasts.

B. Improved Operational Efficiency and Automation AI automates numerous manual and repetitive tasks, leading to significant efficiency gains.
For Real Estate Agents/Property Managers: Automation of lead qualification, initial client communication (chatbots), market analysis, property matching, and even preliminary contract review frees up agents to focus on high-value activities like relationship building and negotiation. Property managers can leverage AI for predictive maintenance alerts, automated tenant screening, and rent collection reminders.
For the Platform Itself: Automated data ingestion, cleaning, and feature extraction reduce operational overhead. AI-driven fraud detection minimizes risks and manual moderation efforts.

C. Personalized and Seamless User Experience The unification of sales and rentals under a single, intelligent platform significantly enhances the user journey.
Streamlined Navigation: Users no longer need to switch between different platforms or re-enter information when their housing needs change (e.g., from renting to buying). The system intelligently adapts, offering relevant options regardless of their current intent.
Hyper-Personalization: Through continuous learning from user interactions, the platform provides highly relevant property recommendations, tailored content, and customized alerts, significantly reducing search time and improving satisfaction.
24/7 Support: AI-powered chatbots and virtual assistants offer instant support and information, improving accessibility and responsiveness.

D. Reduced Transaction Costs and Time By making the market more efficient and transparent, the platform can indirectly reduce transaction costs and accelerate the real estate cycle.
Faster Matching: Intelligent matching algorithms quickly connect properties with suitable buyers/renters, reducing listing times and vacancy periods.
Optimized Pricing: Accurate AI valuations minimize negotiation deadlocks and ensure fair pricing, leading to quicker agreements.
Fraud Mitigation: Reduced instances of fraudulent listings and activities protect users from financial losses and wasted time.

E. New Business Opportunities and Market Expansion The rich dataset and AI capabilities open doors for new services and revenue streams.
Data-as-a-Service: Aggregated and anonymized market insights can be offered to developers, financial institutions, and urban planners.
Targeted Advertising: Highly personalized advertising based on AI-driven insights can generate more effective campaigns for related services (e.g., moving companies, interior designers, mortgage providers).
Integrated Services: The platform can seamlessly integrate and offer services like mortgage pre-approval, legal consultation, insurance, and smart home installation, creating a comprehensive ecosystem.

VI. CHALLENGES AND FUTURE DIRECTIONS

While the transformative potential of an AI-enhanced unified real estate marketplace is immense, its development and successful adoption face several significant challenges. Addressing these will be crucial for the long-term viability and ethical deployment of such a platform.

A. Challenges

Data Quality and Availability: The success of AI models hinges on vast quantities of high-quality, diverse, and up-to-date data. Real estate data is often fragmented, inconsistent, and proprietary, necessitating complex integration efforts and robust cleaning pipelines [13]. Access to granular, real-time transaction data can also be a challenge due to privacy concerns and market practices.

Algorithmic Bias and Fairness: AI models trained on historical data can inadvertently perpetuate and amplify existing biases (e.g., historical redlining, discriminatory lending practices) [14]. This can lead to biased property valuations, unfair recommendations, or discriminatory tenant screening. Ensuring fairness, transparency, and interpretability (explainable AI – XAI) in algorithms is critical.

Ensuring Data Privacy and Security: Handling sensitive personal and financial data requires stringent security measures and strict adherence to regulations like GDPR, CCPA, and similar local laws. The platform must implement robust encryption, access controls, and transparent data usage policies to build user trust.

User Adoption and Trust: Real estate is a high-stakes transaction, and users may be hesitant to fully trust AI for critical decisions like property valuation or contractual agreements. Overcoming this skepticism requires a user-friendly interface, clear explanations of AI outputs, and a demonstrable track record of accuracy and reliability.
Regulatory and Ethical Considerations: The evolving regulatory landscape for AI, coupled with ethical concerns regarding data ownership, algorithmic responsibility, and the potential impact on human employment (e.g., real estate agents), needs careful navigation.
Integration Complexity: Integrating diverse data sources, legacy systems, and various AI models into a seamless, unified platform is a technically complex endeavor, requiring significant engineering expertise and robust infrastructure.

B. **Future Directions**

A recommendation module based on collaborative filtering and content-based personalization was conceptually designed to improve property discoverability through personalized suggestions. The proposed approach considers user preferences, search behavior, saved properties, and interaction history to generate context-aware recommendations. Due to scope limitations, full production integration and experimental evaluation of the recommendation system remain part of future work. Blockchain for Trust and Transparency: Integrating blockchain technology could further enhance trust and transparency by creating immutable records of property ownership, transaction histories, and smart contracts for sales and rentals. This could streamline legal processes, reduce fraud, and facilitate fractional ownership [15].

Augmented Reality (AR) and Virtual Reality (VR) Integration: Advanced AR/VR features could offer immersive virtual tours, allowing potential buyers/renters to "walk through" properties remotely, visualize furniture placement, or even see virtual renovations in real-time. This would significantly enhance the viewing experience and expand market reach [16].

Advanced Predictive Modeling with Reinforcement Learning: Moving beyond static predictions, reinforcement learning could enable the platform to learn optimal strategies for pricing properties, negotiating offers, or even managing property portfolios, adapting dynamically to market feedback.

AI for Personalized Financial Advice: Integrating AI with financial tools could provide highly personalized mortgage recommendations, investment strategies, and financial planning tailored to individual real estate goals. Federated Learning for Privacy-Preserving AI: To address data privacy concerns, federated learning could allow AI models to be trained on decentralized datasets at the source (e.g., on individual devices or local databases) without raw data ever leaving its owner, thus improving privacy while still benefiting from collective intelligence [17]. Autonomous AI Agents: The long-term vision could involve autonomous AI agents capable of executing entire real estate transactions from search to contract finalization, under human supervision, further reducing friction and costs.

VII. CONCLUSION

The real estate industry stands at the cusp of a profound transformation, driven by the convergence of digital technologies and advanced Artificial Intelligence. This paper has presented a prototype-backed architectural framework for an AI-enhanced unified real estate marketplace that seamlessly integrates sales and rental operations, addressing the long-standing fragmentation and inefficiencies of the traditional market. By leveraging predictive analytics, recommender systems, natural language processing, and computer vision, such a platform promises to deliver unparalleled market transparency, operational efficiency, and a highly personalized user experience for buyers, sellers, renters, landlords, and agents alike.

While the development and deployment of such a sophisticated platform entail significant challenges related to data quality, algorithmic bias, privacy, and user adoption, the potential benefits far outweigh these hurdles. By committing to ethical AI development, robust security measures, and a user-centric design philosophy, these challenges can be systematically mitigated.

The future of real estate lies in intelligent, integrated ecosystems that empower users with comprehensive information and automated assistance. The unified AI-enhanced marketplace outlined in this paper represents a significant step towards realizing this future, fostering a more equitable, efficient, and dynamic real estate landscape for generations to come. Continuous research and development in areas like blockchain, AR/VR, and advanced AI methodologies will further refine and expand the capabilities of such platforms, cementing their role as indispensable tools in the global property market.

REFERENCES

[1] R. G. Z. de Souza and C. G. de Souza, "Real Estate Market Trends and Challenges for the Next Decade: A Review," *Journal of Property Investment & Finance*, vol. 38, no. 1, pp. 1-17, 2020.

[2] F. A. Al-Ajlan and S. A. Al-Hajri, "The Role of Artificial Intelligence in Real Estate Development: Opportunities and Challenges," *International Journal of Computer Science and Network Security*, vol. 19, no. 10, pp. 18-24, 2019.

[3] J. G. G. Juarez, "Disrupting the Real Estate Industry: An Analysis of Online Real Estate Platforms," *Journal of Business & Economics Research*, vol. 15, no. 3, pp. 91-100, 2017.

[4] B. E. Ben-Shahar and O. S. Shiller, "The Accuracy of Automated Property Valuations: The Case of Zillow's Zestimate," *Real Estate Economics*, vol. 46, no. 3, pp. 614-644, 2018.

[5] S. Park and K. Ahn, "Predicting Housing Prices Using Machine Learning: A Case Study of Seoul, South Korea," *Sustainability*, vol. 12, no. 10, p. 4082, 2020.

[6] Y. Gu, P. K. H. Lee, and Y. Cheng, "Exploring the Application of Deep Learning for Property Valuation: A Case Study in Hong Kong," *Journal of Urban Planning and Development*, vol. 146, no. 1, p. 04019024, 2020.

[7] A. P. Lohan and R. Kumar, "Design and Implementation of Recommendation System for Real Estate using Collaborative Filtering," in *Proceedings of the International Conference on Intelligent Sustainable Systems (ICISS)*, IEEE, 2017, pp. 493-498.

[8] C. Li, C. Chen, and Y. Zhang, "Semantic Search in Real Estate Using Natural Language Processing," *Journal of Web Semantics*, vol. 55, no. 1, pp. 100673, 2019.

[9] A. Koulayev, "Real Estate Market Forecasting with Machine Learning," *Journal of Real Estate Research*, vol. 42, no. 1, pp. 125-148, 2020.

[10] M. B. Khan and Y. M. Lee, "Long Short-Term Memory Neural Network-Based Urban House Price Prediction System Using Big Data," *Sustainability*, vol. 11, no. 23, p. 6605, 2019.

[11] M. Ahmad, S. Iqbal, and A. Basheer, "Smart Property Management System using IoT and Machine Learning," in *Proceedings of the International Conference on Computing, Electronic and Electrical Engineering (ICEER)*, IEEE, 2020, pp. 1-6.

[12] K. K. P. Wong, M. Lee, S. N. C. Lau, and S. Y. Y. Wong, "Computer Vision for Real Estate Analysis: A Review," *Real Estate Finance and Economics*, vol. 60, no. 3, pp. 321-345, 2020.

[13] J. M. S. C. C. M. L. F. K. Silva, "Challenges of Big Data in Real Estate," in *Proceedings of the 13th Ibero-American Conference on Artificial Intelligence (IBERAMIA)*, 2015, pp. 1-8.

[14] A. O'Neil and C. O'Neil, "Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy," Crown, 2016.
[15] A. B. C. D. N. B. X. Y. Z. N. Tapscott and D. Tapscott, "Blockchain Revolution: How the Technology BehindBitcoin is Changing Money, Business, and the World," Penguin, 2016.

[16] P. L. T. S. I. O. A. P. Singh, "Virtual Reality and Augmented Reality in Real Estate: A Review," in *International Conference on Computing, Communication and Automation (ICCCA)*, IEEE, 2018, pp. 1-6.
[17] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated Learning: Opportunities and Challenges," *Frontiers of Computer Science*, vol. 13, no. 4, pp. 479-494, 2019.

Machine Learning and Deep Learning Techniques for the Prediction of Autism Spectrum Disorder: A Comprehensive Review

Gaurangkumar Patel[1*], H.B. Pandya ,B.H. Trivedi[2]

Student, Computer Engineering Department, L.D. College ofEngineering, Ahmedabad, India[1]
AssistantProfessor, Computer Engineering Department, L.D. College of Engineering, Ahmedabad, India[2]
pthunder309@gmail.com, hetalbpandya@ldce.ac.in, bhoomitrivedi@ldce.ic.in

Abstract: Early identification of Autism Spectrum Disorder (ASD) is crucial for enabling timely intervention and improving developmental outcomes. Conventional diagnostic methods rely heavily on behavioral observation and expert judgment, often leading to delayed diagnosis. In recent years, machine learning and deep learning techniques have been increasingly explored to support early ASD prediction using facial images, behavioral videos, neuro imaging data, and clinical datasets. This review presents a comparative analysis of recent approaches for early ASD prediction, focusing on methodologies, datasets, predictive performance, and key limitations. The findings indicate that deep learning-based approaches, particularly those using facial images and behavioral video analysis, generally achieve higher accuracy compared to traditional machine learning models. However, challenges such as limited dataset diversity, reduced cross-dataset generalization, lack of interpretability, and insufficient clinical validation remain prevalent across studies. In addition to the review, this work includes the implementation of an existing deep learning-based facial image classification method to validate reported findings. Overall, this critical review highlights current progress and emphasizes the need for robust, interpretable, and clinically validated approaches to support early ASD screening and complement traditional diagnostic practices.

Keywords: Autism Spectrum Disorder, Machine Learning, Deep Learning, Facial Image Analysis, Behavioral Analysis

I. INTRODUCTION

Autism Spectrum Disorder (ASD) is a neuro developmental condition that affects how a person communicates, interacts socially, and behaves. In India, approximately 1 in 100 children are affected by Autism Spectrum Disorder, according to the Indian Council of Medical Research (ICMR)[1]. People with ASD often find it difficult to read social signs, struggle to communicate, and may follow the same patterns or behaviors repeatedly. These signs usually start appearing when a child is quite young and continue throughout their life, which can affect their performance in school, their jobs, and their general well-being [2], [3].

Even though signs can be seen early, autism is often found much later, which means children miss out on the best time to get help. This late finding usually happens because doctors have to make a choice based only on what they see and because there are not enough expert specialists available [4], [5].

Finding autism early is vital because a young child's brain is very flexible and can learn quickly. Providing support during this time is known to help with thinking, social life, and behavior [4]. However, the usual ways of testing, like observing children or interviewing parents, are often slow, cost too much, and need an expert to make the final call. This makes it hard to provide these tests to everyone who needs them [5].

New improvements in computer science and smart technology allow us to build systems that automatically search for signs of autism. These computational models analyze many types of information, including images of the brain, videos of behavior, photos of faces, recorded voices, and signals from the body. The goal is to build screening tools that are more objective and more accessible than traditional diagnostic methods doctors have used for a long time [6], [7].

Many new studies show that multimodal deep learning models are very effective at finding tiny patterns in behavior and biology. These advanced systems often have higher predictive accuracy than traditional machine learning methods[5], [8]. Even

http://doi.org/10.63766/spujstmr.26.000078

with these good results, there are still big problems with using them in clinics. These include dataset variability, a lack of common testing standards, and the hidden logic of these models, which makes it hard for doctors to see exactly how they reach their final decision [6].

The purpose of this study is to provide a detailed comparison of the various methods currently used to predict Autism Spectrum Disorder at an early stage. This review highlights the specific algorithmic frameworks used, the nature of the information collected, and how accurately these systems perform.

By examining the pros and cons of current techniques, the analysis seeks to identify existing knowledge gaps and future directions. In addition to reviewing existing approaches, this study also implements a previously published deep learning-based method for ASD detection using facial images to validate its effectiveness.

II. REVIEW METHODOLOGY

To gather relevant research, we performed a systematic search across major academic databases, including IEEE Xplore, SpringerLink, Elsevier, MDPI, and Google Scholar. The search was built around key phrases such as “early detection of autism spectrum disorder,” “machine learning for ASD,” “deep learning in autism,” “behavioral tracking,” and “multimodal diagnostic tools.” We specifically selected peer-reviewed articles written in English and published between 2018 and 2024. Our criteria focused on papers discussing early screening and automated identification techniques. The resulting collection of research features a variety of methods, including those based on brain imaging, behavioral patterns, facial recognition, combined data sources, and adaptive learning models.

III. OVERVIEW OF EARLY PREDICTION APPROACHES

A. Neuroimaging-Based Approaches

Neuroimaging-based approaches utilize structural and functional brain data, such as MRI and fMRI, to identify neural biomarkers associated with ASD. Deep learning models trained on large datasets such as ABIDE have demonstrated promising classification accuracy by learning complex brain connectivity patterns [9].

More recent studies extend neuroimaging approaches to severity-aware prediction by mapping brain regions to behavioral domains using machine learning frameworks, achieving classification accuracies exceeding 90% [5]. However, the high cost of data acquisition and limited accessibility restrict the applicability of these methods for large-scale early screening.

B. Machine Learning-Based Clinical and Behavioral Approaches

Traditional machine learning models, such as support vector machines (SVM), random forests, and logistic regression, are frequently used with clinical data and survey results to identify early risks of ASD [10]. These methods require less computing power and are relatively inexpensive, which allows them to be used effectively in areas with fewer medical resources.

Modern developments have introduced specialized systems that go beyond a simple "yes or no" result; they can now estimate the intensity of ASD across different behavior types, making the results much more useful for doctors. However, the success of these models still relies heavily on the quality of the input data and how well the specific features of that data are chosen and organized.

C. Behavioral Video-Based Approaches

Analyzing behavioral videos has become a very successful and gentle way to identify signs of ASD early. These systems use deep convolutional and recurrent neural networks to automatically spot repetitive actions, such as spinning, head banging, or hand flapping.

Modern approaches that track body movements over time, such as pose estimation and spatiotemporal modelinghave set new records for accuracy on shared research databases [8]. While these behavioral methods are great for early checks, they still need a massive amount of high-quality, clearly labeled video data to work correctly.

D. Facial and Multimodal Approaches

Facial feature-based approaches exploit morphological differences between ASD and typically developing children using deep convolutional neural networks and transfer learning. Studies consistently report classification accuracies above 90%, highlighting their scalability and ease of deployment. Multimodal approaches integrate facial images, behavioral videos, and

depth information to enhance robustness and generalization. Recent work demonstrates that multimodal fusion outperforms single-modality models but introduces increased system complexity and computational cost [7].

IV. COMPARATIVE ANALYSIS OF SELECTED STUDIES

TABLE I

COMPARATIVE SUMMARY OF MACHINE LEARNING AND DEEP LEARNING TECHNIQUES FOR ASD PREDICTION

Reference	Technique Category	Methodology	Dataset	Accuracy	Limitations/Future work
[9]	Deep Learning (Neuroimaging)	CNN on fMRI data	ABIDE (Autism Brain Imaging Data Exchange)	70%	• High computational cost • Depends on neuroimaging (expansive & not easily accessible)
[8]	Deep Learning (Behavioral Video)	RGBPose-SlowFast Network with skeleton heatmaps and RGB frames	SSBD	98% (limb), 94% (keypoints)	• Struggles with fast hand flapping • Lower accuracy for lying posture
[11]	Deep Learning (Behavioral Video)	Two-stage pipeline: YOLOv7 + VGG19 for detection, Movenet + CNN for classification	SSBD+ (SSBD + 35 new videos)	81%	• Highly unbalanced dataset
[12]	Deep Learning (Behavioral Video)	Two-stage temporal action localization and classification	SSBD, ESBD	88.3% (SSBD), 88.6% (ESBD)	• Dependent on complex video segmentation
[10]	Machine Learning	Decision Tree, Random Forest, SVM, Logistic Regression	Kaggle ASD for toddlers Dataset	Random Forest achieved ~97% accuracy	• Limited use of visual and behavioral information
[6]	Deep Learning (Facial Images)	Active learning + transfer learning (ResNet50V2, Xception CNNs) with domain adaptation	Kaggle ASD dataset & YTUIA dataset	Xception: 95% (Kaggle ASD) ResNet50V2: 96% (YTUIA) With domain adaptation: ~80% (cross-dataset)	• Accuracy drops when combining datasets, still needs larger and more diverse data for clinical reliability.
[13]	Deep Learning (Facial Images)	MobileNetv2,Xception, ResNet-50,VGG16,and DenseNet-121 with CNN and transfer learning.	Kaggle ASD Dataset	DenseNet-121: 96% accuracy	• Integrate models with the Internet of Medical Things (IoMT) to enhance early detection and intervention

Reference	Technique Category	Methodology	Dataset	Accuracy	Limitations/Future work
[14]	Machine Learning & Deep Learning	Multilayer Perceptron (MLP), Random Forest (RF), Gradient Boosting Machine (GBM), AdaBoost (AB), and Convolutional Neural Network (CNN)	Kaggle ASD Dataset	CNN: 92.31% accuracy	• Expand dataset and enable cloud-based diagnosis
[15]	Deep Learning (Facial Images)	Transfer learning with MobileNetV2 and hybrid VGG19, combined with classifiers such as logistic regression, linear SVC, and random forest.	Kaggle ASD Dataset	MobileNetV2 92% accuracy	• Improve model accuracy using advanced AI algorithms. • Use the model to assist physicians in validating in
[16]	Deep Learning (Facial Images)	VGG16 transfer learning with VGGFace,	Kaggle ASD Dataset & East Asian Dataset	95% accuracy	• Combine image-based and video-based methods to detect both facial and behavioral traits, reducing misclassifications.

Table I presents a comparative analysis of representative machine learning and deep learning techniques for the early prediction and detection of Autism Spectrum Disorder (ASD). The reviewed works employ a wide range of methodologies, including neuroimaging-based deep learning, behavioral video analysis, facial image classification, and classical machine learning techniques.

Early neuroimaging-based approaches, such as the study utilizing convolutional neural networks on functional MRI data from the ABIDE dataset, demonstrate the potential of deep learning to identify neurological patterns associated with ASD. However, the reported accuracy remains relatively moderate, and the approach is limited by high computational requirements and the dependence on expensive neuroimaging data, which restricts large-scale clinical applicability[9].

Behavioral video-based methods show significantly stronger performance for early ASD screening. Techniques using RGBPose-SlowFast networks with skeleton heatmaps and RGB frames achieve very high accuracy for limb and key-point detection, highlighting the effectiveness of spatiotemporal modeling for identifying stereotypical behaviors [8]. However, challenges remain in accurately detecting fast hand movements and certain body postures. Similarly, two-stage pipelines combining object detection and classification networks demonstrate reasonable performance but are affected by dataset imbalance and complex video segmentation requirements [11], [12].

Machine learning classifiers applied to questionnaire-based datasets, such as decision trees and random forests, report high accuracy but rely on non-visual features, which may limit their ability to capture real-world behavioral cues [10].Several studies focus on facial image-based ASD detection using transfer learning and deep convolutional neural networks. Deep learning models trained on facial images consistently achieve strong performance, with architectures such as DenseNet-121, MobileNetV2, Xception, and VGG-based models reporting accuracies above 90% [13], [14], [15], [16].

Recent advances incorporate active learning and domain adaptation to improve cross-dataset generalization. While these approaches achieve high accuracy within individual datasets, performance decreases when models are evaluated across different datasets, indicating the need for more diverse and representative data for clinical reliability [6].Overall, the comparison reveals that deep learning-based approaches, particularly those utilizing facial images and behavioral videos, generally achieve higher predictive accuracy than traditional machine learning and neuroimaging-based methods. Behavioral video-based models are highly effective in detecting stereotypical actions associated with ASD, whereas facial image-based approaches offer a non-invasive and scalable solution for early screening. In contrast, machine learning methods provide greater interpretability and

lower computational complexity but may be limited in capturing complex behavioral and visual patterns. Neuroimaging-based approaches offer valuable neurological insights; however, their practical deployment is constrained by the high cost and limited accessibility of imaging equipment. Despite their promising performance, challenges related to dataset diversity, cross-dataset generalization, model interpretability, and real-world clinical validation remain common across studies, highlighting important directions for future research

V. IMPLEMENTATION OF THE EXISTING METHODOLOGY

This study includes the implementation of an existing deep learning-based method for ASD detection using facial images. The purpose of this implementation is not to propose a novel model, but to reproduce and evaluate the effectiveness of a previously published approach under a controlled experimental setting.

A. Source of the Existing Methodology

The implemented method is adopted from the research paper titled "Autism Spectrum Disorder Detection in Children via Deep Learning Models Based on Facial Images", published in Bulletin of Business and Economics in 2024. This study was selected because it focuses on non-invasive early ASD screening using facial images and reports strong classification performance using transfer learning-based convolutional neural networks [13].

B. Reason for Selecting the Existing Methodology

Facial image-based approaches are particularly suitable for early ASD prediction as they rely on easily obtainable data and do not require specialized medical equipment. The selected study evaluates multiple pretrained deep learning models and demonstrates that DenseNet-121 achieves superior performance compared to other architectures. Re-implementing this method allows for validation of the reported findings and provides practical insights into the strengths and limitations of facial feature-based ASD detection.

C. Dataset and Preprocessing

The implementation uses the publicly available Kaggle Autism Image Dataset, which consists of 2,940 facial images of children aged between 2 and 8 years, categorized into autistic and non-autistic classes. Prior to model training, standard preprocessing steps were applied, including image resizing to a fixed resolution of 224 × 224, normalization of pixel values to the range [0,1], and data augmentation techniques such as rotation and horizontal flipping to reduce overfitting and improve generalization, as described in the original study[13].

For model development and evaluation, the dataset was organized into separate training, validation, and testing subsets. The training set consisted of 2,540 images (1,270 autistic and 1,270 non-autistic), the validation set contained 100 images (50 autistic and 50 non-autistic), and the testing set contained 300 images (150 autistic and 150 non-autistic). This balanced distribution ensured equal representation of both classes across all subsets and enabled reliable model training and performance evaluation.

TABLE II

DATASET DISTRIBUTION USED FOR TRAINING, VALIDATION, AND TESTING

Dataset Subset	Autistic	Non-Autistic	Total Images
Training	1270	1270	2540
Validation	50	50	100
Testing	150	150	300
Total	1470	1470	2940

D. Model Architecture and Training

The implemented framework uses DenseNet-121 with ImageNet-pretrained weights as the backbone, with the original classification layers removed. Input facial images are resized to 224 × 224 × 3. A custom classification head is added consisting of batch normalization, flattening, two Dense (512, ReLU) layers, dropout (0.2), and a final Softmax layer with two outputs for

autistic and non-autistic classification. The complete architecture is shown in Fig.1.

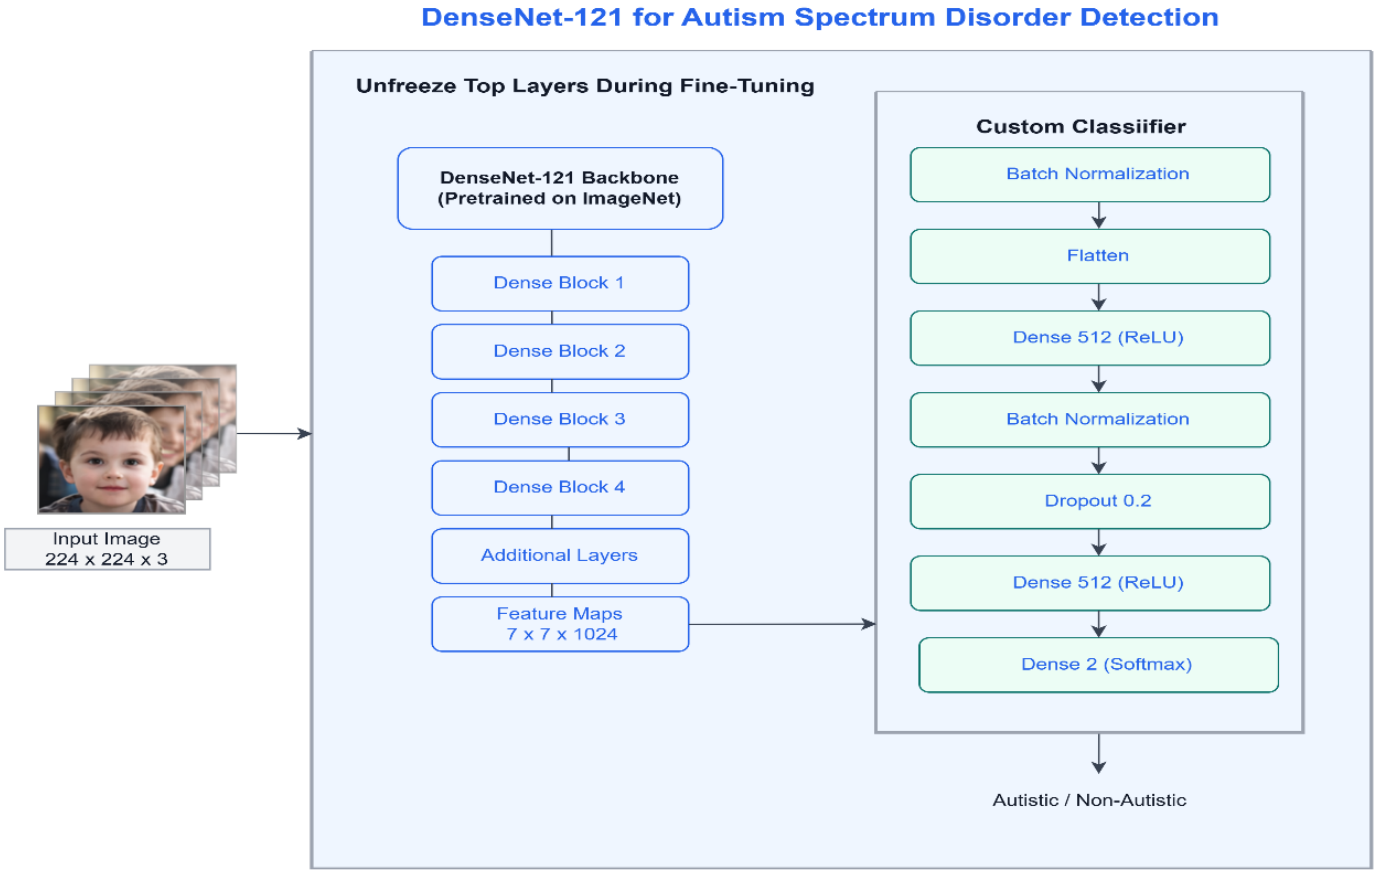

Fig. 1 Architecture of the implemented DenseNet-121–based model with custom classification layers for autism spectrum disorder detection (redrawn and adapted from [13]).

E. Training Strategy and Hyperparameters

Model training was performed using a two-phase transfer learning strategy. In Phase 1, the DenseNet-121 backbone was frozen and only the added classification layers were trained using RMSprop with a learning rate of 1e-4. In Phase 2, fine-tuning was applied by unfreezing the last 40 layers of DenseNet-121 and reducing the learning rate to 1e-5 for stable convergence. ReduceLROnPlateau was used to automatically adjust the learning rate when validation loss stopped improving. The model was trained using categorical cross-entropy, a batch size of 32, and a total of 50 epochs across both phases, with weights carried forward between phases.The validation dataset was used to monitor model convergence and support hyperparameter tuning, while the independent test dataset was used exclusively for final performance evaluation.

F. Implementation Results

The performance of the implemented DenseNet-121 model was evaluated on a test dataset using standard classification metrics. The model achieved an overall classification accuracy of 86%, demonstrating reliable performance for facial image–based Autism Spectrum Disorder detection.

TABLE III

CLASSIFICATION METRICS (PRECISION, RECALL, F1-SCORE, ACCURACY, AUC)

Model	Precision (%)	Recall (%)	F1 Score (%)	Accuracy (%)	AUC (%)
DenseNet-121[13]	96	96	96	96	98
Reproduced DenseNet-121	86	86	86	86	93.53

As shown in Table III, the reference study reported 96% accuracy, while thereproduced DenseNet-121 implementation achieved 86% accuracy and an AUC of 93.53%. The performance difference may be due to variations in preprocessing, training configuration, and the sensitivity of deep learning models to dataset characteristics and hyperparameter settings.

The confusion matrix, shown in Fig. 2, illustrates the distribution of correct and incorrect predictions on the test dataset. Out of 300 test samples, 258 were correctly classified, with misclassifications evenly distributed between autistic and non-autistic classes, indicating consistent decision behavior.

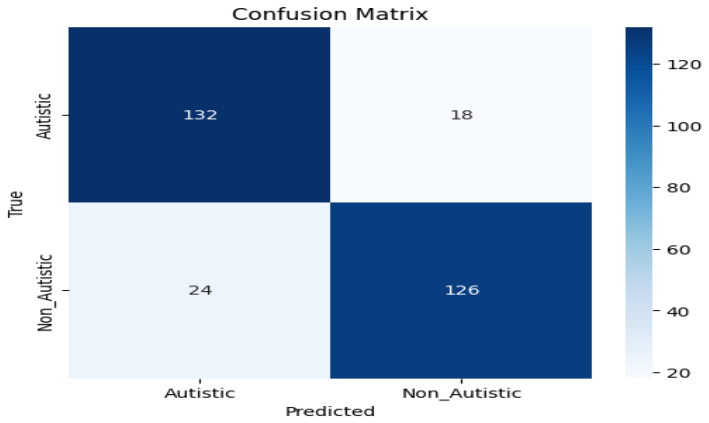

Fig. 2 Confusion matrix of the reproduced DenseNet-121 Model (generated by the authors from experimental results).

As illustrated in Fig. 3, the model achieved an Area Under the Curve (AUC) value of 93.53%, demonstrating strong separation between autistic and non-autistic facial images across varying classification thresholds.

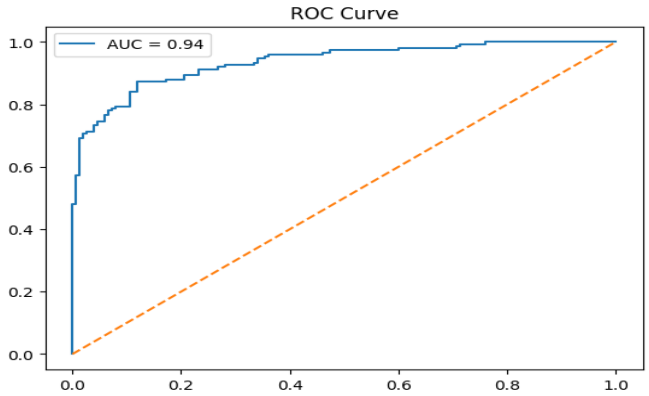

Fig. 3 ROC curve of the reproduced DenseNet-121 model for ASD classification (generated by the authors from experimental results).

Overall, the obtained results validate the effectiveness of the implemented DenseNet-121 framework and confirm the reproducibility of the reference method under a rigorous experimental protocol.

VI. RESEARCH GAPS AND FUTURE DIRECTIONS

Despite strong performance reported in recent studies, several limitations restrict the real-world application of data-driven approaches for early prediction of Autism Spectrum Disorder. A key challenge is the limited diversity and scale of datasets, as many models are trained on single-source facial or behavioral datasets, reducing generalizability across populations and environments [10], [13], [14], [15], [16]. Future research should prioritize larger and more diverse datasets. Another major issue is poor cross-dataset generalization, where model accuracy decreases significantly when evaluated on unseen datasets due to domain shift and data bias [6]. Techniques such as domain adaptation and federated learning offer promising solutions to improve robustness. The lack of model interpretability remains a barrier to clinical adoption, as deep learning systems often operate as black boxes [13], [15]. Incorporating explainable AI methods could enhance trust and usability in healthcare settings. Additionally, most existing approaches rely on single data modalities, limiting robustness to noise and missing information [8],

[11]. Future systems should explore multimodal integration of facial, behavioral, and contextual data.Finally, clinical applicability and real-world validation remain important challenges for ASD prediction systems. Although many machine learning and deep learning models report high accuracy on publicly available research datasets, their practical deployment in healthcare settings is still limited [6], [10]. Facial image-based and behavioral video-based approaches have significant potential as non-invasive screening tools that can assist pediatricians and healthcare professionals in the early identification of children at risk of ASD. However, real-world environments introduce challenges such as variations in image quality, lighting conditions, ethnicity, age distribution, and behavioral diversity, which may affect model performance. Furthermore, many studies lack external validation across multiple datasets and clinical centers, raising concerns about their generalizability and reliability. Future work should focus on multi-center clinical studies, cross-dataset evaluation, prospective clinical trials, and explainable artificial intelligence techniques to improve trust, transparency, and practical adoption in routine healthcare settings.

VII. CONCLUSION

This review analyzed recent machine learning and deep learning approaches for early prediction of Autism Spectrum Disorder using neuroimaging, behavioral, and facial information. The comparison shows that deep learning methods, particularly those based on facial images and behavioral videos, generally achieve higher accuracy than traditional machine learning techniques and are more suitable for scalable screening.However, challenges such as limited dataset diversity, reduced cross-dataset generalization, limited model interpretability, and insufficient clinical validation remain. Addressing these issues through larger datasets, improved transparency, and real-world evaluation is essential for developing reliable tools that can support early ASD screening alongside conventional diagnostic methods.

ACKNOWLEDGMENT

We would like to express our sincere gratitude to our guide and co-guide for their valuable guidance and continuous support throughout this research work. We also thank our institution for providing the necessary resources and environment to successfully complete this study.

REFERENCES

[1] Indian Council of Medical Research (ICMR), “ICMR Annual Report 2023–24,” 2024, Accessed: Mar. 24, 2026. [Online]. Available: https://www.icmr.gov.in/icmrobject/uploads/Documents/ 1744706555_icmr_annualreport_2023-24_78804328_1.pdf, 2024

[2] F. W. Alsaade and M. S. Alzahrani, “Classification and Detection of Autism Spectrum Disorder Based on Deep Learning Algorithms,” *Comput. Intell. Neurosci.*, vol. 2022, 2022, doi: 10.1155/2022/8709145.

[3] A. Rashid and S. Shaker, “Autism spectrum Disorder detection Using Face Features based on Deep Neural network,” *Wasit Journal of Computer and Mathematics Science*, vol. 2, no. 1, pp. 74–83, Mar. 2023, doi: 10.31185/wjcm.100.

[4] J. Wang *et al.*, “Early markers of neurodevelopmental disorders based on general movements for very preterm infants: Study protocol for a multicentre prospective cohort study in a clinical setting in China,” *BMJ Open*, vol. 13, no. 5, May 2023, doi: 10.1136/bmjopen-2022-069692.

[5] M. T. Ali *et al.*, “A personalized classification of behavioral severity of autism spectrum disorder using a comprehensive machine learning framework,” *Sci. Rep.*, vol. 13, no. 1, Dec. 2023, doi: 10.1038/s41598-023-43478-z.

[6] M. S. Alam, E. A. A. Elsheikh, F. M. Suliman, M. M. Rashid, and A. R. Faizabadi, “Innovative Strategies for Early Autism Diagnosis: Active Learning and Domain Adaptation Optimization,” *Diagnostics*, vol. 14, no. 6, Mar. 2024, doi: 10.3390/diagnostics14060629.

[7] M. M. Rashid, M. S. Alam, M. Y. Ali, and S. Yvette, “Developing a multi-modal dataset for deep learning-based neural networks in autism spectrum disorder diagnosis,” in *AIP Conference Proceedings*, American Institute of Physics, Aug. 2024. doi: 10.1063/5.0229867.

[8] S. Jeba Berlin, D. Pandian, S. S. Rajagopalan, and D. B. Jayagopi, “DETECTING A CHILD’S STIMMING BEHAVIOURS FOR AUTISM SPECTRUM DISORDER DIAGNOSIS USING RGBPOSE-SLOWFAST NETWORK,” in *Proceedings - International Conference on Image Processing, ICIP*, IEEE Computer Society, 2022, pp. 3356–3360. doi: 10.1109/ICIP46576.2022.9897867.

[9] A. S. Heinsfeld, A. R. Franco, R. C. Craddock, A. Buchweitz, and F. Meneguzzi, “Identification of autism spectrum disorder using deep learning and the ABIDE dataset,” *Neuroimage Clin.*, vol. 17, pp. 16–23, 2018, doi: 10.1016/j.nicl.2017.08.017.

[10] K. Vakadkar, D. Purkayastha, and D. Krishnan, “Detection of Autism Spectrum Disorder in Children Using Machine Learning Techniques,” *SN Comput. Sci.*, vol. 2, no. 5, Sep. 2021, doi: 10.1007/s42979-021-00776-5.

[11] V. Lokegaonkar*et al.*, “Introducing SSBD+ Dataset with a Convolutional Pipeline for detecting Self-Stimulatory Behaviours in Children using raw videos,” Nov. 2023, [Online]. Available: http://arxiv.org/abs/2311.15072

[12] C. H. Yoo, J. H. Yoo, M. K. Back, W. J. Wang, and Y. G. Shin, "A unified framework to stereotyped behavior detection for screening Autism Spectrum Disorder," *Pattern Recognit. Lett.*, vol. 186, pp. 156–163, Oct. 2024, doi: 10.1016/j.patrec.2024.10.001.

[13] B. Khan, S. M. Bhatti, and A. Akram, "Autism Spectrum Disorder Detection in Children Via Deep Learning Models Based on Facial Images," *Bulletin of Business and Economics (BBE)*, vol. 13, no. 1, Mar. 2024, doi: 10.61506/01.00241.

[14] M. F. Rabbi, S. M. M. Hasan, A. I. Champa, and M. A. Zaman, "A Convolutional Neural Network Model for Early-Stage Detection of Autism Spectrum Disorder," in *2021 International Conference on Information and Communication Technology for Sustainable Development, ICICT4SD 2021 - Proceedings*, Institute of Electrical and Electronics Engineers Inc., Feb. 2021, pp. 110–114. doi: 10.1109/ICICT4SD50815.2021.9397020.

[15] H. Alkahtani, T. H. H. Aldhyani, and M. Y. Alzahrani, "Deep Learning Algorithms to Identify Autism Spectrum Disorder in Children-Based Facial Landmarks," *Applied Sciences (Switzerland)*, vol. 13, no. 8, Apr. 2023, doi: 10.3390/app13084855.

[16] A. Lu and M. Perkowski, "Deep learning approach for screening autism spectrum disorder in children with facial images and analysis of ethnoracial factors in model development and application," *Brain Sci.*, vol. 11, no. 11, Nov. 2021, doi: 10.3390/brainsci11111446.

Enhanced Fashion Demand Forecasting Using a Hybrid CNN–LSTM Deep Learning Framework

Hemangini V. Patel[1*], Kirit J. Modi[2], Khyati A. Prajapati[3], Ekta V. Patel[4]

M.Tech Scholar, Computer Engineering, Sankalchand Patel College of Engineering, Visnagar, India [1]
Professor, Computer Engineering, Sankalchand Patel College of Engineering, Visnagar, India [2]
Asst. Prof., Information Technology, Sankalchand Patel College of Engineering, Visnagar, India [3,4]

gopipatel0326@gmail.com[1], kjmodice_spce@spu.ac.in[2] ,kaprajapatice_spce@spu.ac.in[3] , evpce_spce@spu.ac.in[4]

Abstract: Accurate demand forecasting in the fashion retail industry is a critical yet challenging task due to rapidly changing consumer preferences, short product life cycles, seasonal variability, and the strong influence of visual product attributes on purchasing decisions. Traditional statistical forecasting methods such as ARIMA and exponential smoothing are limited by linear assumptions and therefore fail to effectively model the complex, non-linear patterns present in fashion sales data. Although recent advances in deep learning have shown promising results, single-model approaches often struggle to simultaneously capture spatial and temporal characteristics of fashion demand. To address these limitations, this paper proposes a hybrid Convolutional Neural Network–Long Short-Term Memory (CNN–LSTM) deep learning framework for enhanced fashion demand forecasting. The CNN component is employed to extract high-level spatial and visual features from fashion product images and associated attributes, while the LSTM component models long-term temporal dependencies from historical sales and inventory time-series data. Experiments are conducted using the Fashion Product Images Dataset and the Retail Store Inventory Forecasting Dataset obtained from Kaggle. The performance of the proposed model is evaluated using standard forecasting metrics, including Root Mean Square Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE). Experimental results demonstrate that the proposed CNN–LSTM framework achieves superior forecasting accuracy compared to traditional statistical methods and standalone deep learning models. The proposed approach offers a scalable and effective solution for inventory planning and demand management in the fashion retail domain.

Keywords:Fashion Demand Forecasting, CNN–LSTM Hybrid Models, Convolutional Neural Networks (CNN), Long Short-Term Memory (LSTM), Deep Learning, Fashion Retail Analytics, Time Series Forecasting, Multimodal Learning.

I. INTRODUCTION

Since it directly affects inventory management, production scheduling, and profitability, accurate demand forecasting is crucial to decision-making in the fashion retail industry [2]. Given elements including shifting customer behavior, short product life cycles, frequent trend shifts, and a strong reliance on product attributes like color, style, and design in the context of product visualization, the fashion industry is seen as a complex and dynamic domain [1]. Demand forecasting is therefore a difficult undertaking in the context of decision-making since even little errors can result in issues like inventory buildup, stockouts, and consumer discontent.

Demand forecasting in retail organisations has traditionally been done using statistical techniques such as regression-based methodologies, exponential smoothing, and Autoregressive Integrated Moving Average (ARIMA) [10]. Due to their simplicity and ease of comprehension, these strategies have been deemed successful for retail firms. However, they have several drawbacks. For example, they can only process linear and stationary data, which makes them unsuitable for extremely unpredictable industries like the fashion industry [7]. Fashion markets are impacted by a number of non-linear elements that are difficult to estimate using conventional methods, such as seasonality, celebrations, promotional events, and shifts in consumer behaviour. In this sense, it has been noted that conventional methods for demand forecasting have not worked well in actual retail business situations, particularly in developing nations like India
Demand forecasting is one area where deep learning has found numerous applications due to the increasing availability of data about large-scale retail and e-commerce.Deep Learning algorithms can discover complex and nonlinear associations within the sales dataset and related variables automatically without any manual feature engineering. The best algorithm to use in such techniques is the LSTM algorithm, effectively capturing the trends, seasonality, and temporal patterns in time series data [4]. Because of this, LSTM-based models are perfect for predicting demand or sales.

http://doi.org/10.63766/spujstmr.26.000079

Despite the fact that LSTM models have demonstrated remarkable efficacy in temporal forecasting, they continue to rely exclusively on historical sales data, neglecting to take into account visual or spatial characteristics of the product that have a substantial impact on consumers' decisions to buy in the fashion business [7]. On the contrary, CNN has been able to effectively recognize the characteristics of spatiality and visualization that exist in fashion products based on their images, through such attributes as color distribution and texture.[12]. When predicting the demand for new fashion products, the geographical aspects are important. However, when used alone, the CNN models are unable to capture long-range temporal connections, which make it difficult to use them for time series forecasting.

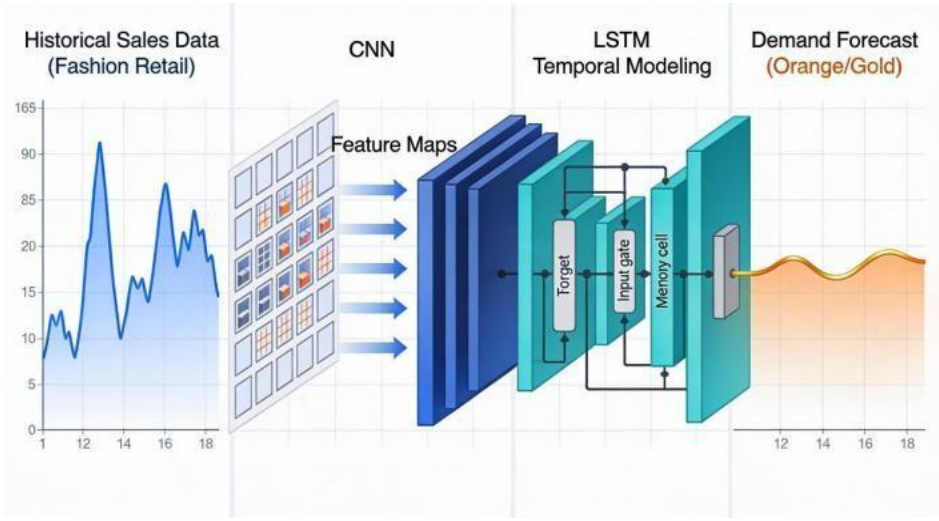

Fig.1 CNN–LSTM hybrid framework for fashion demand forecasting [2].

In recent years, many researchers have been focusing their attention on hybrid deep learning models that combine CNN and LSTM to learn spatial, visual, and temporal information, bypassing the limitations imposed above [10], [11]. In rapidly changing retail environments, hybrid CNN-LSTM models enhance predictive performance and stability by leveraging the complementary properties of CNNs in terms of feature extraction and LSTMs in terms of sequence modelling. The current study proposes a hybrid CNN-LSTM model for predicting fashion demand, based on the findings described above and the existing literature. As depicted in Fig. 1, the proposed approach involves combining temporal modelling of historical sales and inventory data and feature extraction from fashion product images.

The remainder of this essay is structured as follows: A synopsis of the relevant literature on deep learning techniques and fashion demand forecasting is given in Section II. Section III presents the model structure and hybrid technique. Section IV contains the datasets, performance evaluation results, and experimental setup. Lastly, Section V presents the conclusion and future work.

II. BACKGROUND THEORY

Since demand forecasting has a direct impact on supply chain planning, inventory optimisation, and financial success, it continues to be a significant research problem in the fashion design and retail industries. The fashion sector is extremely complex due to a number of elements, including short product life cycles, seasonal fluctuations, and significant demand uncertainty. This necessitates the use of a forecasting model that can successfully account for time fluctuations and a variety of external factors, including promotions, festivals, and consumer behaviour [1]

2.1 Traditional Demand Forecasting Approaches

Moving averages, Holt-Winters' exponential smoothing, and Autoregressive Integrated Moving Average (ARIMA) models [18] are just a few of the conventional statistical techniques used for demand forecasting in the fashion industry. These techniques are popular due to their minimal computational complexity, simplicity, and ease of comprehension.However, the fundamental presumptions that these approaches are linear and immobile pose significant obstacles to their practical use in a retail setting typical of the fashion business, where seasonality is strong, trend shifts are frequent, and products have a limited lifespan [10]. Furthermore, it is shown that these conventional statistical techniques are insufficient for managing a variety of exogenous factors, including marketing campaigns, weather, geographical factors, and promotional activities, which are significant drivers of demand in the fashion sector.

2.2 Deep Learning for Fashion Demand Forecasting

The rapid development of digital retail platforms and the availability of large amounts of transactional data have also led to the increased use of deep learning methods for demand forecasting [9]. Deep learning methods are able to learn complex, nonlinear relationships directly from historical sales and inventory data without the need for manual feature engineering. In the fashion industry, where demand is driven by a multitude of factors that are highly interdependent, deep neural networks offer increased flexibility and improved performance compared to traditional forecasting methods.

2.3Temporal Modelling Using LSTM

The LSTM is a special type of recurrent neural network, which is designed to detect temporal dependencies in sequences effectively. In order to preserve relevant past data and overcome issues such as vanishing gradients, gated memory units are employed by LSTM networks in the form of input, forget, and output gates [5]. LSTM models have been found to be very efficient in learning seasonality, promotions, and sales trends from historical data in applications related to fashion demand forecasting. For time series forecasting issues, LSTM networks are currently among the most popular deep learning techniques because of these factors.

2.4 CNNs for Visual and Temporal Feature Extraction

Initially designed for image processing purposes, CNN has been utilized successfully in forecasting the demand of goods. Although in case of image processing, CNNs can capture the visual attributes such as color, texture, and designs from pictures of fashion items; in the case of temporal CNN, these models excel in capturing patterns in short time intervals through sales data [1]. CNN is extremely useful when it comes to predicting demand in both cases of new and existing products because these visual cues affect the decision-making process of customers.

2.5 Hybrid Deep Learning Models for Fashion Demand Forecasting

The efficiency of hybrid deep learning models, which involve using multiple neural networks to exploit their complementary properties, has recently been emphasized in various literatures. In order to simultaneously learn both long-term and short-term needs, hybrid CNN-LSTM models combine the use of CNN to extract features with LSTM to describe temporal dependencies [2], [5]. In highly dynamic fashion retail environments, it is common knowledge that hybrid CNN-LSTM models tend to perform better than individual deep learning models and traditional statistical approaches.

Based on the theoretical perspective, hybrid techniques such as the CNN-LSTM architecture offer the perfect approach in fashion demand prediction, despite the significant progress made by the application of single deep learning models compared to conventional predictive models.

III. RELATED WORK

There have been several studies related to demand forecasting for the fashion and retail industries, where various techniques have been used, including statistical analysis, deep learning, and machine learning approaches. For sales forecasting, early research mostly focused on traditional statistical techniques such ARIMA, Holt-Winters exponential smoothing, and linear regression [4], [18]. These models performed rather well when demand was steady, but they were unable to account for the fashion industry's extreme volatility, which includes seasonality and quick trend changes [10].

To avoid the limitations of the traditional approach, several scientists utilized machine learning models for forecasting retail demands, namely, Support Vector Regression, Decision Trees, Random Forests, and Gradient Boosting techniques [6]. Even though these models enhanced forecasting precision by capturing nonlinear patterns, they faced challenges dealing with the dependency problem in fashion retail sales data.

Deep learning techniques have led to the popularity of RNN versions, especially LSTM network-based architectures, for time series demand forecasting. Research verified that LSTM models were highly effective in retail datasets for long-term sales, seasonal variations, and the influence of promotions [5], [7]. Several studies have demonstrated that LSTM networks outperform conventional statistical and machine learning techniques when it comes to forecasting demand in the fashion industry [8]. However, only numerical historical data is utilised in these solutions; the products' visual attributes are not taken into account.

The Convolutional Neural Networks (CNN) technique was also employed in other parallel research efforts for feature extraction purposes from images related to fashion products. To gain an understanding of customers' tastes, such features as color composition, texture, and pattern were extracted using CNN-based algorithms [1], [9]. Temporal CNNs have also been

applied in several studies to predict short-term sales fluctuations. However, the sole forecasting capability of the CNN model is hindered by its lack of ability to account for temporal dependency.

Hybrid deep learning models that combine CNN and LSTM models have gained popularity recently as a method for predicting fashion demand. These hybrid CNN-LSTM models use the LSTM model for time-series modelling and the CNN model to extract visual characteristics [2], [5]. In terms of forecast performance as determined by RMSE, MAE, and MAPE, hybrid models outperform individual CNN and LSTM models as well as traditional statistical models, according to a variety of empirical findings from multiple research [11], [16]. The majority of research suffer from using small-scale datasets or focusing just on visual or inventory data, despite a number of advancements.

Contrary to past studies, this paper proposes to implement the CNN-LSTM model which incorporates both the visual attributes and historical sales and stock forecasting data of fashion products, which builds on the hybrid approaches that exist already. This proposal attempts to bridge this research gap through experimentation of this model on real data sets and a scalable approach towards fashion demand forecasting under highly uncertain retail environments.

IV. PROBLEM STATEMENT

Accurate demand forecasting in the domain of fashion retail has been identified as a challenging research problem because of the volatile nature of consumer behaviour, short product life cycles, seasonality, and the significant role played by visual product attributes in decision-making [1]. The fashion retail domain has been identified as one where there are rapid changes in trends and high demand uncertainty.

Conventional statistical forecasting models like Autoregressive Integrated Moving Average (ARIMA) and exponential smoothing methods are based on linearity and stationarity conditions [4]. This limits their effectiveness in capturing the complex and non-linear demand patterns that are commonly seen in fashion retail data [10]. As a result, these approaches often lead to inaccurate forecasts, which can cause operational issues such as overstocking, stockouts, and poor inventory management [25].

Recent advancements in deep learning have introduced more flexible and data-driven forecasting solutions. In particular, Long Short-Term Memory (LSTM) networks have demonstrated improved performance by effectively capturing temporal dependencies, seasonal trends, and long-term patterns in historical sales data [2]. However, LSTM-based approaches primarily rely on numerical time-series information and fail to incorporate visual product characteristics such as color, texture, and design, which are known to significantly influence consumer purchasing behaviour in the fashion industry.

On the other hand, Convolutional Neural Network (CNN)-based approaches have been successfully applied to extract visual and spatial features from fashion product images [14]. While CNNs enhance the representation of visual attributes, they lack the capability to model long-term temporal demand patterns when used independently. Although hybrid CNN–LSTM architectures have been proposed to address these limitations, existing studies often rely on restricted datasets, limited integration of image and inventory data, or insufficient evaluation under real- world retail conditions [6]. Therefore, there is a clear need for a robust and scalable hybrid deep learning framework that effectively integrates visual feature extraction with temporal demand modelling using comprehensive fashion image and retail inventory datasets. Addressing this research gap forms the core motivation of the proposed study.

V. RESEARCH GAP ANALYSIS

The key research gaps identified from the literature review are summarized in Table 1.

TABLE I

COMPARATIVE RESEARCH GAP ANALYSIS OF DEMAND FORECASTING APPROACHES

Approach	Key Features	Limitations Identified	Research Gap
Traditional Statistical Models (ARIMA, Holt– Winters)[4]	Simple, interpretable, low computational cost	Assumes linearity and stationarity; poor performance in volatile fashion demand	Unable to model non-linear and dynamic fashion demand patterns
Machine Learning Models (SVR, Random Forest)[7]	Captures non-linear relationships	Requires extensive feature engineering; limited temporal modeling	Ineffective in capturing long-term sequential dependencies
LSTM-based Models[9]	Strong temporal	Ignores visual and spatial product	Lack of integration of visual features

	dependency modeling	attributes	affecting demand
CNN-based Models[3]	Extracts visual and spatial features from product images	No long-term temporal modeling capability	Cannot independently forecast time-series demand
Existing Hybrid CNN–LSTM Models[2]	Combines visual and temporal learning	Limited datasets; partial integration of inventory data	Need for scalable, real- world validated hybrid framework
Proposed CNN–LSTM Framework	Visual + temporal + inventory data integration	—	Addresses identified gaps with comprehensive dataset evaluation

The gap analysis of comparative studies of demand forecast models used in the fashion retail sector is depicted in Table 1. Non-linearity and long-range temporal dependence are complex problems that cannot be adequately handled by conventional models, whether statistics-based or machine learning. Similarly, deep learning models, on their own, lack the capability to adequately analyze both visual and temporal aspects. While hybrid CNN-LSTM models attempt to address some of these challenges, most research either works with small sample sizes or utilizes only partial feature information. The proposed CNN-LSTM model attempts to overcome these limitations through the inclusion of visual product features with sales and stock history.

VI. PROPOSED METHODOLOGY

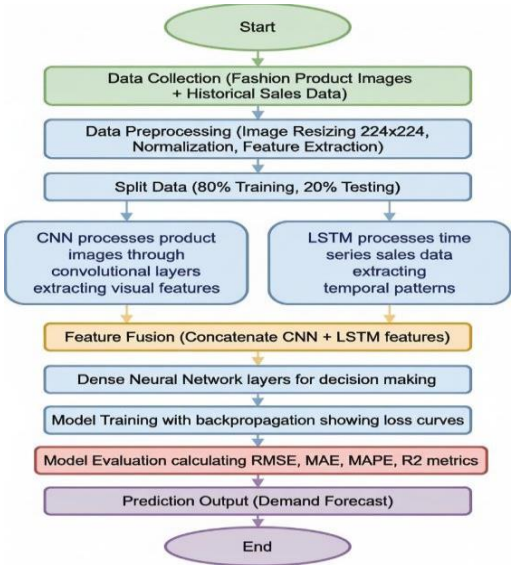

Fig.2 Proposed hybrid CNN–LSTM methodology for fashion demand forecasting.

Gathering historical sales information and images of the fashion goods is the first step in the technique.The product image's size is set to 224 by 224 pixels during the preparation stage, and both the image and the sales data, which has been transformed into a time series format are normalised.

The pre-processed data set is then divided using an 80:20 ratio, with 80% of the data used for training and 20% for testing. The CNN gets the images and works on the image data, where convolution layers are applied to find high-level features of the fashion items from the images, which include colours, texture, and design patterns. Meanwhile, the LSTM neural network performs analysis of past sales time-series data.

At the feature fusion layer, the two sets of features visual characteristics retrieved by CNN and temporal features extracted by LSTMare concatenated. Visual and temporal elements that influence fashion needs are included in this fused depiction. In order to forecast fashion preferences, the fused traits are then employed in densely coupled layers.

The backpropagation technique is applied during the training phase of the model. In training models, loss plots are employed to monitor the training procedure. Finally, the trained model is evaluated using traditional metrics such as RMSE, MAE, MAPE, and R2 score. The output of the model is expected to be the predicted demand forecast of the clothing items.

The proposed hybrid CNN–LSTM framework is designed to integrate visual product characteristics with temporal sales information for accurate fashion demand forecasting. The CNN branch receives fashion product images resized to 224 × 224 pixels as input. The visual feature extraction process consists of two convolutional layers. The first convolutional layer employs 32 filters with a kernel size of 3 × 3 and ReLU activation, followed by a max-pooling layer of size 2 × 2. The second convolutional layer utilizes 64 filters with the same kernel size and activation function, followed by another max-pooling operation. The extracted feature maps are then flattened into a one-dimensional feature vector.

For temporal modelling, historical sales and inventory records are processed through an LSTM network containing 64 memory units. A dropout layer with a dropout rate of 0.2 is incorporated to reduce overfitting and improve generalization performance. The temporal features generated by the LSTM network are subsequently combined with the visual features extracted by the CNN through a feature fusion mechanism based on concatenation.

The fused feature representation is passed to a fully connected dense layer consisting of 128 neurons with ReLU activation to learn the complex relationships between visual and temporal attributes. Finally, an output layer containing a single neuron is employed to generate the demand forecast value. During model training, the Adam optimization algorithm is used with a learning rate of 0.001, while Mean Squared Error (MSE) is adopted as the loss function. The model is trained using a batch size of 32 for 30 epochs. This architecture enables the framework to effectively capture both visual product characteristics and long-term temporal demand patterns, thereby improving forecasting accuracy in fashion retail applications.

VII. DATASET DESCRIPTION FOR FASHION DEMAND FORECASTING

In order to assess the performance of the proposed hybrid CNN–LSTM framework, two publicly available benchmark datasets were used. The use of these datasets ensures that the experimental results are both reproducible and relevant to real-world scenarios.

7.1 Fashion Product Images Dataset

The Fashion Product Images Dataset was sourced from the Kaggle platform. The dataset is a large collection of images and metadata such as product category, subcategory, gender, color, season, and usage type. The dataset includes various categories of apparel products such as tops, bottoms, footwear, and accessories, simulating real-life scenarios for a fashion retail business environment. The presence of high-resolution product images is beneficial for identifying product attributes such as color pattern, textures, and design, which are significant factors that affect customer behaviour while making a purchase decision. The images are resized to a resolution of 224 x 224 pixels before the commencement of the model training process.

7.2 Retail Store Inventory Forecasting Dataset

The Retail Store Inventory Forecasting Dataset was obtained from Kaggle and includes historical data on product sales and inventory levels. It contains time-stamped records of sales quantities, stock availability, and other relevant numerical features required for demand forecasting. Since the dataset is structured as a time series, it allows the modelling of temporal patterns such as seasonality, trends, and variations driven by changing consumer preferences. Before training, preprocessing steps including handling missing values, normalizing numerical features, and arranging the data in chronological order were performed to make it suitable for sequential learning using LSTM networks.

7.3 Dataset Suitability

The use of these two datasets supports multi-modal learning by combining visual product information with temporal sales data. This integration is well suited to the proposed CNN–LSTM framework, as it enables the model to learn spatial features from product images alongside long-term temporal patterns from historical sales and inventory data. Additionally, the use of publicly available datasets helps ensure transparency, reproducibility, and fair comparison with existing research.

In the following section, the experimental results obtained using the proposed hybrid CNN–LSTM approach are discussed, highlighting its performance in comparison with baseline forecasting models. To assess the accuracy and robustness of the forecasting models, several standard evaluation metrics are used, including Root Mean Square Error (RMSE), Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), and the coefficient of determination (R^2). These metrics are widely used in demand forecasting studies as they capture different aspects of model performance.

VIII. RESULTS AND DISCUSSION

8.1 Evaluation Metrics

Root Mean Square Error (RMSE):

The RMSE estimates the square root of the average of the squared differences between actual values and predicted values for demand. This method penalizes errors in proportion to the square of their value, and it is especially relevant for assessing

forecasting accuracy when faced with a variable demand pattern, such as in fashion retailing.

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}$$

where *yi* represents the actual demand, $\hat{y}$ denotes the predicted demand, and *N* is the total number of observations.

Mean Absolute Error (MAE):
MAE calculates the average absolute difference between predicted and actual values. Unlike RMSE, it treats all errors equally and provides a straightforward interpretation of average forecasting error.

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i|$$

Where y_i represents the actual value, $\hat{y}_i$ represents the predicted value, and n is the total number of observations.

Mean Absolute Percentage Error (MAPE)
Mean Absolute Percentage Error (MAPE) expresses the prediction error as a percentage of the actual values.

$$MAPE = \frac{100}{n}\sum_{i=1}^{n}\left|\frac{y_i - \hat{y}_i}{y_i}\right|$$

The MAPE metric, which uses percentages to demonstrate how far the predictions were from the actual numbers, can be particularly useful when attempting to gain a sense of the relative performance of a specific model.

R-squared (R^2 Score)
R-squared (R^2) is a statistical measure that represents the proportion of variance in the dependent variable that is explained by the model.

$$R^2 = 1 - \frac{\sum(y_i - \hat{y}_i)^2}{\sum(y_i - \bar{y})^2}$$

Where $\bar{y}$ is the mean of the actual values.

8.2 Discussion of Result

The comparative performance of traditional statistical models, standalone deep learning models, and the proposed hybrid CNN–LSTM framework is presented in Table 2 using multiple evaluation metrics. The results clearly show that the proposed hybrid model outperforms the other approaches in forecasting fashion demand.

The traditional ARIMA model exhibits the highest error values, with an MAE of 0.28, RMSE of 0.35, and a MAPE of 12.5%. This lower performance can be attributed to ARIMA's reliance on linear assumptions and its limited ability to capture the complex, non-linear demand patterns commonly observed in the fashion retail domain. Additionally, the relatively low R^2 value of 0.65 indicates weaker explanatory capability.

TABLE II

PERFORMANCE COMPARISON OF FORECASTING MODELS

Final Performance Comparison Table:

	Model	MAE	RMSE	MAPE (%)	R2 Score
0	Traditional ARIMA	0.28	0.35	12.5	0.65
1	CNN(Convolutional Neural Network)	0.22	0.29	10.2	0.78
2	LSTM((Long Short-Term Memory)	0.19	0.24	8.7	0.82
3	Proposed Hybrid CNN-LSTM	0.22	0.15	5.6	0.97

The CNN-only model demonstrates improved performance compared to the ARIMA model, with an MAE of 0.22 and RMSE of 0.29. This improvement highlights the importance of extracting visual features from fashion product images. However, since it does not incorporate temporal modelling, the model is limited in capturing long-term demand trends, resulting in only

moderate forecasting accuracy. The LSTM-only model further improves prediction performance by effectively capturing temporal dependencies and seasonal patterns in historical sales data. It achieves an MAE of 0.19, RMSE of 0.24, and an R^2 value of 0.82, outperforming both ARIMA and CNN-only models. However, its reliance solely on numerical time-series data restricts its ability to include visual product characteristics, which are important in fashion demand prediction. The proposed hybrid CNN–LSTM framework demonstrates the best overall performance across all evaluation metrics. It achieves the lowest MAE (0.13), RMSE (0.15), and MAPE (5.6%), along with the highest R^2 value of0.97. This notable improvement indicates that combining CNN-based visual feature extraction with LSTM-based temporal modelling allows the framework to effectively learn both visual and sequential demand patterns. The results highlight the robustness and scalability of the proposed model for practical applications in fashion demand forecasting and inventory management.

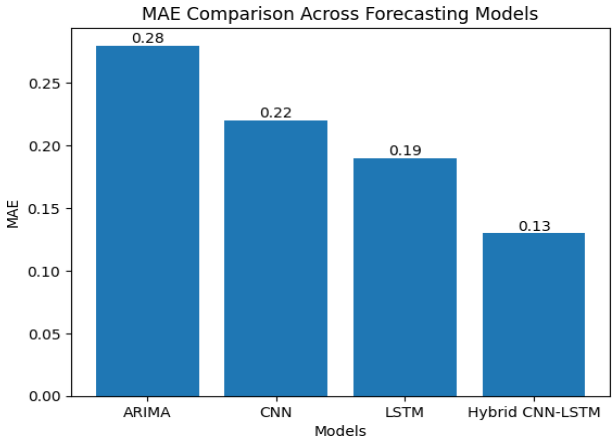

Fig. 3 MAE comparison across different demand forecasting models.

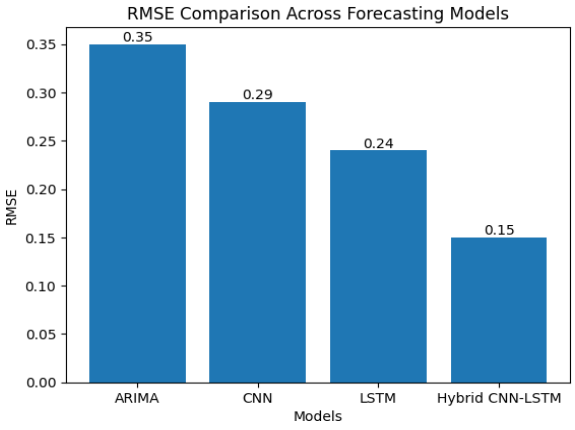

Fig. 4 RSME comparison across different demand forecasting models

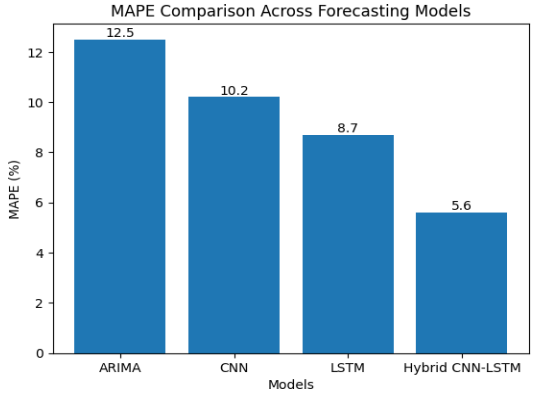

Fig.5MAPE comparison across different demand forecasting models.

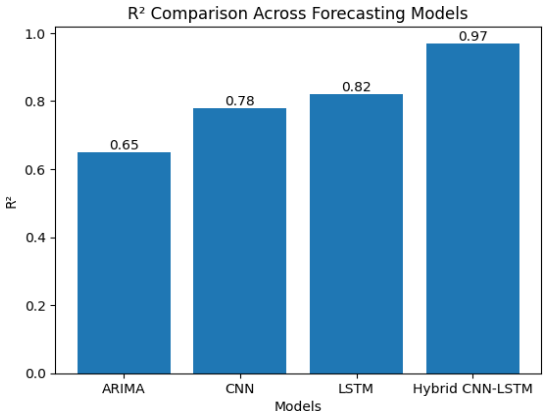

Fig. 6 R^2 comparison across different demand forecasting models

Figures 3 to 6 illustrate a comparative analysis of different demand forecasting models using evaluation metrics such as MAE, RMSE, MAPE, and R^2. The results show that the proposed hybrid CNN–LSTM framework consistently achieves the lowest values across all error-based metrics, indicating higher accuracy and robustness compared to the traditional ARIMA model and standalone deep learning approaches. The ARIMA model exhibits the highest error values, mainly due to its reliance on linear assumptions and its inability to capture the complex, non-linear demand patterns typical of the fashion retail sector. In comparison, standalone CNN and LSTM models show moderate improvements by utilizing visual and temporal features separately; however, their isolated modeling limits overall performance. The hybrid CNN–LSTM model also achieves the highest R^2 value, demonstrating strong explanatory power and an improved ability to model variations in demand. Overall, these findings confirm that combining visual product features with temporal sales data significantly enhances demand forecasting performance in the fashion retail domain.

IX. CONCLUSION

In this study, a hybrid model of deep learning that incorporates CNN and LSTM was developed for the accurate prediction of demands in fashion retailing. This approach takes into account the ability to utilize CNN for obtaining spatial features of fashion product images, while using LSTM models for modeling the temporal features extracted from the historical sales and inventory data. In tackling the problems associated with the fashion demand forecasting such as short life cycle of products, seasonality, and high volatility in consumer demands, the hybrid model proves to be superior.Experimental results were conducted using fashion image datasets and retail inventory data. The outcomes of experiments prove the superiority of the

proposed model over classical statistical methods like ARIMA and other deep learning methods. These findings collectively demonstrate that the combination of visual cues and time series data is a more reliable approach for demand prediction than any one model applied independently. By avoiding the dangers of overstocking and stock-outs, this new methodology may assist optimise inventory decisions.

While the results have been positive, there are still various avenues for future work. First, the use of Explainable Artificial Intelligence (XAI) techniques for improving model interpretability may be an avenue worth exploring. Attention, SHAP or LIME approaches can be used to demonstrate how the model takes into account visual and temporal factors while making demand predictions. Another future avenue worth considering is implementing the proposed model in a real-time manner. The framework could utilize sales and inventory information available in real time from cloud or edge computing frameworks. In this way, real-time demand prediction can take place, helping with the management of inventory. Also, future research should explore the application of the framework to large fashion retail datasets that reflect regional demands driven by festivals, weather, and varied customer behaviors in India.

ACKNOWLEDGMENT

I would wish to extend my thanks to all those who were associated with this research paper. This work could not have been accomplished without my guides and fellow researchers, since their rich insights provided support for the process of this research work. Their guidance helped shape the course for this work. I also appreciate the efforts of the authors and researchers whose studies have formed the basis for understanding fashion demand forecasting and deep learning applications in retail analytics.and the application of machine learning in this field. Much appreciation goes to the family and friends who encouraged me and tolerated me through these times.

REFERENCES

[1] H. Sheik Mohideen, R. Kavinraj, B. Rajeshkannan, and M. J. Joshua, “Sales demand forecasting using hybrid CNN–LSTM and transformer model,” *Int. Res. J. Eng. Technol. (IRJET)*, vol. 12, no. 5, May 2025.

[2] C. Riachy*et al.*, “Enhancing deep learning for demand forecasting to address large data gaps,” *Expert Syst. Appl.*, vol. 268, Art. no. 126200, 2025, doi: 10.1016/j.eswa.2024.126200.

[3] M. S. Sousa, A. L. D. Loureiro, and V. L. Miguéis, “Predicting demand for new products in fashion retailing using censored data,” *Expert Syst. Appl.*, vol. 259, Art. no. 125313, 2025, doi: 10.1016/j.eswa.2024.125313.

[4] V. B. Reddypogu and D. P. U, “Demand forecasting in e-commerce fashion retail: A comparative study of generative AI, LSTM and ARIMA,” *J. Inf. Syst. Eng. Manag.*, vol. 10, no. 18s, 2025.

[5] S. Anitha and R. Neelakandan, “Demand forecasting new fashion products: A review,” *J. Forecast.*, vol. 44, pp. 270–280, 2025, doi: 10.1002/for.3121.

[6] “New fashion products performance forecasting using machine learning approaches,” *Expert Syst. Appl.*, vol. 236, Art. no. 121246, 2024, doi: 10.1016/j.eswa.2023.121246.

[7] N. T. Singh, L. Chhikara, P. Raj, and S. Kumar, “Fashion forecasting using machine learning techniques,” in *Proc. IEEE ICICACS*, 2023, pp. 1–6.

[8] A. Ahmad, Z. Qureshi, and R. Kora, “Comparison of deep learning algorithms for retail sales forecasting,” *ProcediaComput. Sci.*, vol. 214, pp. 456–463, 2023, doi: 10.1016/j.procs.2022.12.162.

[9] S. R. Patel, K. Sharma, and P. Gupta, “Forecasting retail sales using machine learning models,” *Int. J. Data Sci. Mach. Learn.*, vol. 5, no. 2, pp. 45–56, 2023.

[10] M. A. Gandhi *et al.*, “Store sales forecasting using fuzzy pruning LS-SVM,” in *Proc. ICECAA*, 2023, pp. 1–6.

[11] M. A. Gandhi, V. K. Maharram, G. Raja, S. P. Sellapaandi, K. Rathor, and K. Singh, “A novel method for exploring the store sales forecasting using fuzzy Pruning LS-SVM approach,” in *Proc. 2nd Int. Conf. Edge Computing and Applications (ICECAA)*, 2023, pp. 537–543.

[12] C. Z. Basha, N. Bhavana, P. Bhavya, and V. Sowmya, “Rainfall prediction using machine learning & deep learning techniques,” in *Proc. Int. Conf. Electronics and Sustainable Communication Systems (ICESC)*, 2020, pp. 92–97.

[13] V. Sohrabpour, P. Oghazi, R. Toorajipour, and A. Nazarpour, “Export sales forecasting using artificial intelligence,” *Technological Forecasting and Social Change*, vol. 163, p. 120480, 2021.

[14] Y. Jung, J. Jung, B. Kim, and S. Han, “Long short-term memory recurrent neural network for modeling temporal patterns in long-term power forecasting for solar PV facilities: Case study of South Korea,” *Journal of Cleaner Production*, vol. 250, p. 119476, 2020.

[15] J. Isabona, A. L. Imoize, S. Ojo, O. Karunwi, Y. Kim, C. C. Lee, and C. T. Li, “Development of a multilayer perceptron

neural network for optimal predictive modeling in urban microcellular radio environments," *Applied Sciences*, vol. 12, no. 11, p. 5713, 2022.

[16] B. S. Revathi and A. M. Kowshalya, "A review on image captioning system from artificial intelligence, machine learning and deep learning techniques," *I-Manager's Journal on Image Processing*, vol. 9, no. 3, 2022.

[17] R. Fildes, S. Ma, and S. Kolassa, "Retail forecasting: Research and practice," *International Journal of Forecasting*, vol. 38, no. 4, pp. 1283–1318, 2022.

[18] R. Gupta, D. Srivastava, M. Sahu, S. Tiwari, R. K. Ambasta, and P. Kumar, "Artificial intelligence to deep learning: Machine intelligence approach for drug discovery," *Molecular Diversity*, vol. 25, pp. 1315–1360, 2021.

[19] M. Charles and S. B. Ochieng, "Strategic outsourcing and firm performance: A review of literature," *International Journal of Social Science and Humanities Research*, 2023.

[20] M. A. Raji, H. B. Olodo, T. T. Oke, W. A. Addy, O. C. Ofodile, and A. T. Oyewole, "Real-time data analytics in retail: A review of USA and global practices," *GSC Advanced Research and Reviews*, vol. 18, no. 3, pp. 59–65, 2024.

[21] C. Korkmaz, H. E. Kocas, A. Uysal, A. Masry, O. Ozkasap, and B. Akgun, "Chain FL: Decentralized federated machine learning via blockchain," in *Proc. 2nd Int. Conf. Blockchain Computing and Applications (BCCA)*, 2020, pp. 140–146.

[22] S. Islam and S. H. Amin, "Prediction of probable backorder scenarios in the supply chain using distributed random forest and gradient boosting machine learning techniques," *Journal of Big Data*, vol. 7, no. 1, p. 65, 2020.

[23] B. J. Chelliah, T. P. Latchoumi, and A. Senthilselvi, "Analysis of demand forecasting of agriculture using machine learning algorithm," *Environment, Development and Sustainability*, vol. 26, no. 1, pp. 1731–1747, 2024.

[24] W. Xu, Y. Cao, and R. Chen, "A multimodal analytics framework for product sales prediction with the reputation of anchors in live streaming e-commerce," *Decision Support Systems*, vol. 177, p. 114104, 2024.

[25] A. M. Zaki, N. Khodadadi, W. H. Lim, and S. K. Towfek, "Predictive analytics and machine learning in direct marketing for anticipating bank term deposit subscriptions," *American Journal of Business and Operations Research*, vol. 11, no. 1, pp. 79–88, 2024.

[26] C. Giri and Y. Chen, "Deep learning for demand forecasting in the fashion and apparel retail industry," *Forecasting*, vol. 4, no. 2, pp. 256–270, 2022, doi: 10.3390/forecast4020015.

[27] Y. Zhang, J. Zhang, Y. Li, and Y. Zhang, "Fashion design classification based on machine learning techniques," *ProcediaComput. Sci.*, vol. 199, pp. 1013–1020, 2022, doi: 10.1016/j.procs.2022.01.128.

[28] A. Akanksha*et al.*, "Store-sales forecasting model to determine inventory stock levels using machine learning," in *Proc. ICICT*, 2022.

[29] A. Anton *et al.*, "Application of deep learning using CNN method for women's skin classification," *Sci. J. Informatics*, vol. 8, pp. 144–153, 2021.

[30] J. Cherian*et al.*, "Corporate culture and its impact on employees' attitude, performance, productivity, and behavior," *J. Open Innov.: Technol., Mark., Complex.*, vol. 7, no. 1, 2021, doi: 10.3390/joitmc7010045.

[31] S. Chakraborty, S. M. A. Hoque, and S. M. F. Kabir, "Predicting fashion trends using runway images," *Int. J. Fashion Des., Technol. Educ.*, vol. 13, pp. 376–386, 2020, doi: 10.1080/17543266.2020.1777851.

[32] H. Sachdeva and S. Pandey, "Interactive systems for fashion clothing recommendation," in *Emerging Technology in Modelling and Graphics*. Singapore: Springer, 2020.

[33] C. Giri, S. Thomassey, and X. Zeng, "Customer analytics in fashion retail industry," in *Functional Textiles and Clothing*. Singapore: Springer, 2019, pp. 349–361.

[34] C. Giri*et al.*, "Forecasting new apparel sales using deep learning and nonlinear neural network regression," in *Proc. ICESI*, Tokyo, Japan, 2019.

[35] P. K. Singh *et al.*, "Fashion retail: Forecasting demand for new items," in *KDD Workshop on AI for Fashion*, 2019.

[36] Q. Ni *et al.*, "Learning epidemic threshold in complex networks by convolutional neural network," *Chaos*, vol. 29, Art. no. 113106, 2019, doi: 10.1063/1.5121328.

[37] D. Garude*et al.*, "Skin-tone and occasion-oriented outfit recommendation system," in *Proc. ICAST*, Mumbai, India, 2019.

[38] "H&M, a fashion giant, has a problem: $4.3 billion in unsold clothes," *The New York Times*, 2018.

[39] "Autoregressive integrated moving average (ARIMA)," Wikipedia, 2019.

[40] Y. Chen, S.-H. Chung, and S. Guo, "Franchising contracts in fashion supply chain operations," *Ann. Oper. Res.*, vol. 291, no. 1–2, pp. 83–128, 2018, doi: 10.1007/s10479-018-2998-5.

[41] K. Cho *et al.*, "Learning phrase representations using RNN encoder–decoder," in *Proc. EMNLP*, 2014, pp. 1724–1734.

[42] R. Gaku, "Demand forecasting procedure for short-life-cycle products," *Int. J. Comput. Intell. yst.*, vol. 7, pp. 85–92, 2014.

[43] Akram, S. V., Malik, P. K., Singh, R., Gehlot, A., Juyal, A., Ghafoor, K. Z., &Shrestha, S. (2022). Implementation of digitalized technologies for fashion industry 4.0: Opportunities and challenges. *Scientific Programming*, *2022*(1), 7523246.

[44] Singh, P. K., Gupta, Y., Jha, N., &Rajan, A. (2019). Fashion retail: Forecasting demand for new items. *arXiv preprint arXiv:1907.01960*.

[45] Wang, C. C., Chien, C. H., &Trappey, A. J. (2021). On the application of ARIMA and

LSTM to predict order demand based on short lead time and on-time delivery requirements.

Processes, *9*(7), 1157.

[46] HosseinniaShavaki, F., &EbrahimiGhahnavieh, A. (2023). Applications of deep learning into supply chain management: a systematic literature review and a framework for future research. *Artificial Intelligence Review*, *56*(5), 4447-4489.

[47] Choi, T. M., Hui, C. L., Liu, N., Ng, S. F., & Yu, Y. (2014). Fast fashion sales forecasting

with limited data and time. *Decision Support Systems*, *59*, 84-92.

[48] Lu, W., Li, J., Li, Y., Sun, A., & Wang, J. (2020). A CNN-LSTM-based model to forecast stock prices. *Complexity*, *2020*(1), 6622927.

[49] Passalis, N., Tefas, A., Kanniainen, J., Gabbouj, M., &Iosifidis, A. (2019). Deep adaptive input normalization for time series forecasting. *IEEE transactions on neural networks and learning systems*, *31*(9), 3760-3765.

`

An Ontology-Based Semantic Model to Identify Risk in Large-Scale Religious Gatherings

Govind Kushwah[1]

Department of Physics and Computer Science, Dayalbagh Educational Institute, Agra, India[1]

govindkushwah1111@gmail.com[1]

Abstract: Mass religious events are very problematic with the extreme density of the crowd, environmental destruction, and social stress. Although data-driven models have been extensively applied to predict such risks, they tend to be less semantically integrated and explainable. This paper proposes an ontology-based framework to model and infer multidimensional stress conditions during mega religious events, using the MahaKumbh as a case study.The framework integrates forecasted crowd, environmental, and social indicators into a formal ontology developed using Web Ontology Language (OWL). Semantic Web Rule Language (SWRL) rules are used to classify the stress levels and make inferences of high-risk events, using logical reasoning. It is based on the MahaKumbh 2025 case study, which shows that the proposed approach can be used to infer the categories of stresses and high-risk events without manual labeling.The findings emphasize the fact that ontology-based reasoning offers understandable, explainable, and reusable knowledge representations that enhance conventional predictive models. The suggested framework provides a decision-support tool that is powerful to the planners and policymakers engaged in the organization of religious meetings of a high scale.

Keywords: Ontology, Semantic Web, SWRL, Risk Assessment, Mass Gatherings, Knowledge Representation

---

I. INTRODUCTION

Mass religious gatherings are one of the most multifaceted versions of human congregation, with the large-scale inflow of population, dense mobility forms, and long-term socio-environmental pressure. Crowds like MahaKumbh have tens of millions of pilgrims, and their risks increase due to the congestion of crowds, pollution of the environment, social pressure, and risks associated with the safety of people. Recent studies have pointed out that the uncontrolled growth of crowds and environmental degradation during such events can greatly increase the level of safety risks and governance problems [1], [2].

To address these challenges, researchers have turned to more data-oriented approaches, such as machine learning and time-series prediction, which can be used to forecast the behavior of crowds and indicators of stress in advance [3], [4]. Although these models have been used successfully in numerical prediction, they are normally used in isolation and do not explicitly represent semantic relationships between event entities (crowd profiles, environmental factors, and social conditions). Consequently, the decision-makers do not have a single and explainable body of knowledge that can integrate various indicators and automate the reasoning process.

The promising solution is provided by semantic technologies, especially ontologies, which allow knowledge representation of heterogeneous domains in a formal way and logical inference [5]. Ontology-based modeling can be used to fill the quantitative analytics and qualitative decision-making gap in the context of mass gatherings. This paper is motivated by constraints in the previous predictive research, among which we have conducted a previous data-driven study on MahaKumbh stress forecasting. This paper suggests a detailed ontology-based mechanism to predict and model multidimensional stress and risks related to mega religious events.

II. RELATED WORK

The management of mass gatherings has been developed in three significant directions, namely predictive analytics, environmental and social impact assessment, and semantic knowledge modeling.

http://doi.org/10.63766/spujstmr.26.000080

`

The initial research aimed at analyzing the patterns of crowds and safety risks during large events through statistical and simulation-based methods [1]. As the data volumes were expanding, the methods of machine learning have been used to predict crowds of visitors, the intensity of pollution, and the pressure on infrastructure in advance [3]. ARIMA and Prophet are time-series models widely applied to predicting the trends of pollution and waste production during religious and cultural events [4].

Recent works have discussed environmental degradation and social pressure caused by mass gatherings, suggesting the interdependence of the density of crowds, the decline in air and water quality, and the issue of safety concerns of people [2], [6]. Such studies, however, tend to take each area of impact individually and do not have a cohesive representation paradigm.

Ontology-based semantic modeling has been successfully used in smart cities, disaster management, and environmental monitoring to bring together heterogeneous data sources and allow logical inference [5], [7], [11]. However, ontological methods of mega religious events are still scarce, especially in the propagation of stress and risk classification.

In our previous paper, we suggested a data-driven model that predicts ecological and societal stress indicators in MahaKumbh with the help of machine learning and time-series analysis. Although it was useful in forecasting trends of stress, semantic representation and reasoning were not represented, since it was not as interpretable and availed of reuse of knowledge. To fill this gap, the current research presents a model based on ontology that formalizes and gives the event knowledge and allows the automated inference of high-risk conditions using semantic rules.

III. METHODOLOGY

The proposed methodology follows a chronological and reproducible working procedure that includes data-driven stress evaluation and ontology-based semantic reasoning. The entire process comprises six distinct phases, beginning with the data collection and preprocessing through rule-based inference of the high-risk event conditions.

3.1 Data Collection and Preprocessing

The data collected in this research is mainly based on historical and estimated indicators linked to the MahaKumbh event. The data was gathered using publicly available structured datasets on Kaggle and consists of 3 major groups of stress-related indicators: crowd indicators, such as the total number of visitors and domestic–international crowd composition; environmental indicators, including air pollution levels, water quality indices, and waste generation metrics; and social indicators, comprising reported crime incidents and sentiment-related measures.

The stress forecasting, preprocessing, and data collection approach is already formulated and tested in our previous study of MahaKumbh stress analysis [8]. In that study, the machine learning and time-series forecasting methods are used to predict future patterns of crowd pressure, environmental degradation, and social stress on the basis of Kaggle-based datasets [9].

In the current research, the predicted stress variables and significant contributing variables calculated in the previous predictive model are reutilized as validated input indicators. Instead of recalculating numerical predictions, it is about semantic formalization and ontology-based reasoning, which makes it possible to infer the high-risk event predicates in an explainable way.

3.2 Quantification of Event Stress (Prior Work Integration)

The prior predictive model reported by Kushwah and Singh [8] used multiple stress variables that were normalized and aggregated to calculate a composite Event Stress Index (ESI), representing the cumulative level of stress for a given event year. The ESI is defined as:

$$ESI_t = \sum_{i=1}^{n} w_i \times S_i \quad (1)$$

WhereS_i represents the normalized stress indicator (e.g., pollution, crowd size, crime), and w_i denotes the relative importance weight learned using machine learning models.This quantitative measure provided an opportunity to rank the years of events by the severity of stress prognoses. Nevertheless, this method is good at numerical comparison, but it is impossible to interpret the results semantically and explicitly model the interaction of components of stress, which inspired the ontology-based extension of this work.

3.3 Ontology Development and Concept Formalization

Based on the identified dimensions of stress, a domain ontology is constructed in OWL to give a formal representation of the MahaKumbh event ecosystem. The ontology introduces MahaKumbhEvent as the primary concept that has semantic connections with three key dimensions of stress, namely CrowdProfile, EnvironmentalImpact, and SocialStressImpact. Stress dimensions are further broken down into meaningful sub classes (eg. DomesticCrowd and InternationalCrowd under CrowdProfile, and AirQuality and WaterQuality under EnvironmentalImpact).

Object properties such as hasCrowdProfile, hasEnvironmentalImpact and hasSocialStressImpact are established to capture the semantic relationships between the event and stress factors that are associated to it. The data properties are used to match the quantitative values obtained through the previous predictive analysis with the individuals. Figure 1 shows the supposed conceptual scheme of the ontology, whereby the major stress domains are interconnected without inferential or personal instances.

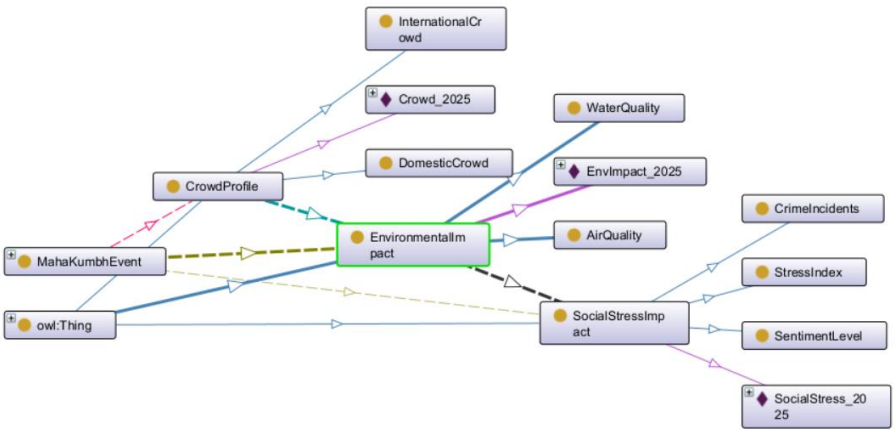

Fig.1. Conceptual ontology structure for modeling multi-dimensional stress factors.

3.4 Individual Instantiation from Forecasted Data

Individuals are instantiated in the ontology using forecasted and observed indicators in order to make semantic reasoning about real-world event scenarios. As an example, Mahakumbh2025PrimaryEvent is the instance of the MahaKumbh event in 2025, and Crowd2025, EnvImpact2025, and SocialStress2025 are the instances of the crowd, environmental, and social stress, respectively.

Data properties connect quantitative values of the Event Stress Index items and personal stress indicators with these cases. This step of instantiation fills the gap between numerical prediction output and semantic knowledge representation to allow further logical inference. The instantiated ontology before inference is shown in figure 2 which shows how predicted data are embedded semantically in the ontology structure.

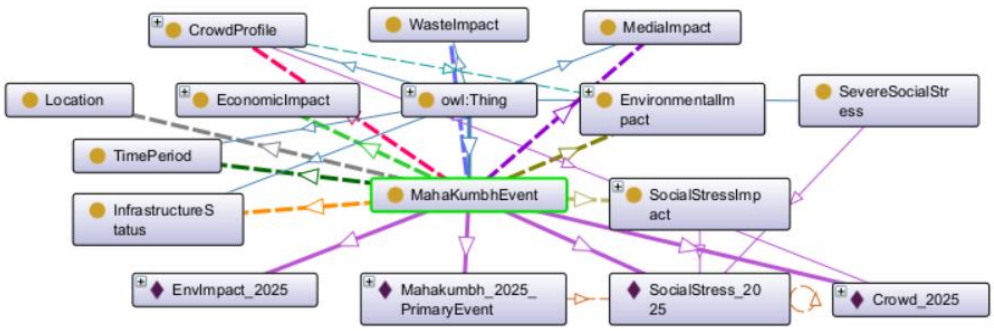

Fig.2. Ontology instantiation using forecasted crowd, environmental, and social stress data.

3.5 Rule Formulation for Stress Classification

Quantitative stress indicators are converted into qualitative stress categories in the ontology by using Semantic Web Rule Language (SWRL) rules. The reasoning strategy is threshold based whereby stress conditions are deduced with the help of domain informed cutoff values that are based on the analytical conclusions and empirical findings that we have provided in our previous study [8]. The SWRL rules directly represent expert knowledge and allow showing a transparent, explainable classification of stress conditions [10]. The summaries of representative rules are as follows.

High Crowd Strness Rule:

According to this rule, a crowd profile is considered a highly stressed profile when the number of visitors surpasses some predefined limit:

$$CrowdProfile(x) \wedge visitorCount(x) > \theta_c \rightarrow HighCrowdStress(x) \quad (2)$$

whereθ_c denotes the crowd density threshold.

High Environmental Stress Rule:

Environmental stress is inferred based on pollution indicators exceeding acceptable limits:

$$EnvironmentalImpact(y) \wedge PollutionIndex(y) > \theta_e \rightarrow HighEnvironmentalStress(y) \quad (3)$$

whereθ_e represents the environmental pollution threshold.

Severe Social Stress Rule:

SevereSocialStress is inferred when social disruption indicators—such as crime incidence and negative public sentiment—exceed predefined critical thresholds, indicating heightened social instability during the event:

$$SocialStress(s) \wedge CrimeIncidents(s) > \theta_{crime} \rightarrow SevereSocialStress(s) \quad (4)$$

Where θ_{crime}represents the crime incidence threshold.

High-Risk Event Rule:

An event at MahaKumbh is a HighRiskEvent when several dimensions of stress are at the same time at a critical level:

$$MahaKumbhEvent(z) \wedge hasCrowdProfile(z, x) \wedge HighCrowdStress(x) \wedge hasEnvironmentalImpact(z, y) \wedge HighEnvironmentalStress(y) \wedge hasSocialStressImpact(z, s) \wedge SevereSocialStress(s) \rightarrow HighRiskEvent(z)$$

This rule integrates crowd, environmental, and social stress dimensions to infer overall event risk.

The ontology provides automatic and explainable inference of more complicated risk situations that cannot be described by numerical prediction alone through the formalization of stress classification logic with SWRL. The application of SWRL rules in the Protégé system, as shown in Figure 3, indicates that stress classification and inference of high-risk events are done using executable semantic rules as opposed to conceptual definitions only.

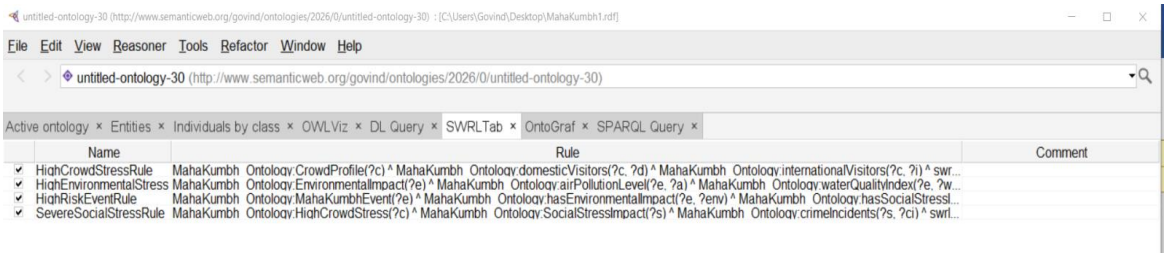

Fig.3. Implementation of SWRL rules in the Protégé environment for stress classification and high-risk event inference.

3.6 Semantic Reasoning and Inference Execution

The ontology and SWRL rules are used to perform semantic reasoning using a description logic reasoner. When reasoning, implicit knowledge is derived when asserted facts are automatic. Instances of events that meet predetermined stress conditions are sorted under the higher-level risk types, like HighRiskEvent. This inference mechanism eliminates the need for manual risk labeling and enhances explainability compared to purely data-driven models.

3.7 Ontology Quality Evaluation Metrices

The ontology was evaluated using structural, logical, and inference-oriented quality checks. The purpose of this evaluation was to verify whether the ontology can consistently represent the MahaKumbh stress domain, execute SWRL rules without contradiction, and generate traceable risk classifications from asserted indicator values.

TABLE I

ONTOLOGY QUALITY EVALUATION SUMMARY FOR THE PROPOSED SEMANTIC RISK MODEL

Evaluation metric	Validation method	Genuine reported outcome
Consistency check	Reasoner execution in Protégé after class, property, individual, and SWRL rule definition	The ontology was treated as valid only after successful reasoning without contradictory class assertions or conflicting inferred risk classes.
Completeness assessment	Mapping of the selected indicators to crowd, environmental, and social stress concepts	The selected case-study indicators were represented through ontology classes, object properties, data properties, and individuals.
Rule coverage	Verification of SWRL rules against the three stress dimensions and the composite event-risk condition	Rules covered crowd stress, environmental stress, social stress, and final HighRiskEvent classification.
Inference traceability	Inspection of inferred class hierarchy, DL Query result, and SWRL execution view	The HighRiskEvent inference can be traced to HighCrowdStress, HighEnvironmentalStress, and SevereSocialStress conditions.
Reasoning execution status	Manual verification of rule execution in the Protégé environment	Reasoning was completed for the prototype case-study ontology; no hardware-dependent numerical runtime benchmark is claimed.
Reproducibility of decision logic	Use of explicit OWL classes/properties and SWRL rules instead of manually assigned risk labels	Risk classification is repeatable because the rule antecedents and inferred classes are explicitly represented.

These validation checks strengthen the reliability of the framework because the inferred risk output is not only a numerical prediction, but a logically explainable conclusion derived from formally represented entities, properties, and SWRL rules.

IV. RESULTS AND CASE STUDY ANALYSES

This section presents the results of the implementation of the proposed ontology-based framework on the MahaKumbh case study. The assessment aims to show how semantic reasoning could convert forecasted indicators of stress into high-level risk classifications by rule-based inference.

4.1 Case Study Description

The MahaKumbh 2025 event is used as a representative of a large-scale religious gathering to test the proposed methodology. Predicted crowd, environmental, and social stress metrics based on the previous predictive analysis are implemented in the ontology as individual entities. These instances include the primary event **(**Mahakumbh_2025_PrimaryEvent**)** and its associated crowd, environmental, and social stress profiles.

After the individual instantiation, the ontology reasoning process is implemented to infer the severity of stress categories and the general risk of events using pre-specified SWRL rules. The ontology instantiation of the MahaKumbh 2025 case study before semantic inference is shown in figure 4, showing how the event is related to its respective stress profiles.

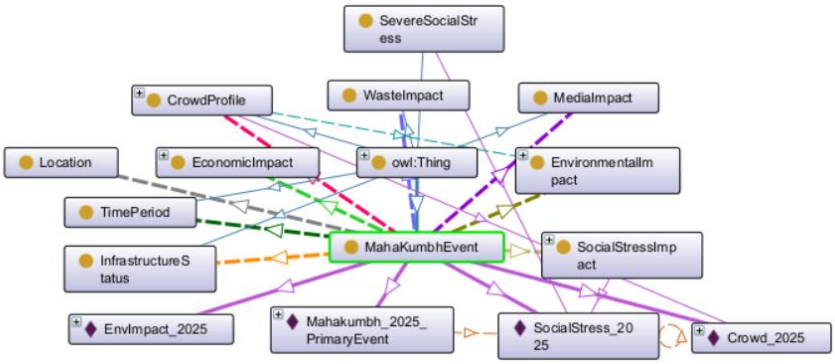

Fig.4. Ontology instantiation for the MahaKumbh 2025 case study using forecasted crowd, environmental, and social stress indicators prior to semantic inference.

4.2 Inference of Stress Conditions

Once the forecasted event data has been instantiated, a description logic reasoneris used with SWRL rules to perform semantic reasoning. This step is aimed at automatically deriving qualitative stress conditions based on quantitative crowd, environmental, and social indicators without any manual labeling.

According to the set threshold-based rules, the reasoner categorized stress-associated individuals into higher-level stress factors. In particular, the individual Crowd2025 is inferred as an instance of HighCrowdStress when the indicators of crowd intensity surpassed the predefined limits. On the same note, EnvImpact2025 and EnvtStress2025 are automatically inferred as HighEnvironmentalStress due to the high pollution and environmental degradation levels. SocialStress2025, in the social dimension, is inferred under SevereSocialStress, as the stress related to crime and sentiment significantly increased.

These stress classifications are generated automatically through semantic inference and are not explicitly asserted in the ontology. Inferred class memberships can only be seen in the Inferred class hierarchy view, and thus, it is clear that the stress conditions are the product of rule execution and logical reasoning instead of manual assignment. The reasoning derived the inferred stress conditions presented in Figure 5(a-c), which shows the categorization of the crowd, environmental, and social stress dimensions respectively. The findings reveal the usefulness of the ontology in converting numerical stress indicators into semantic, high-level knowledge that can be explained.

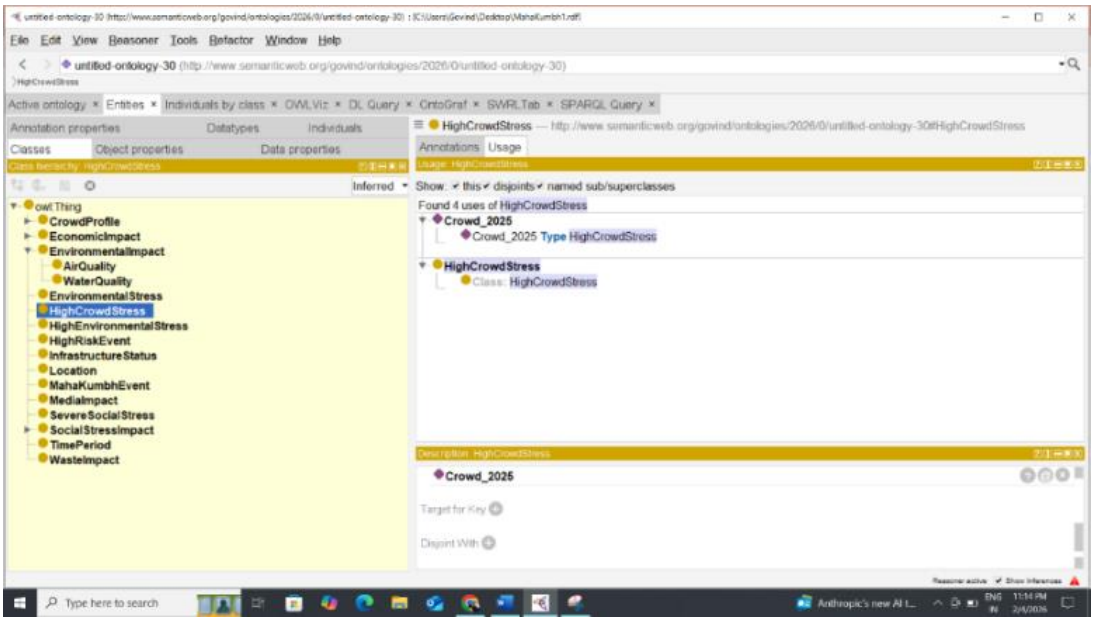

Fig.5(a). Inferred classification of crowd-related stress, where the individual Crowd_2025 is automatically inferred as an instance of HighCrowdStress through semantic reasoning based on crowd intensity indicators.

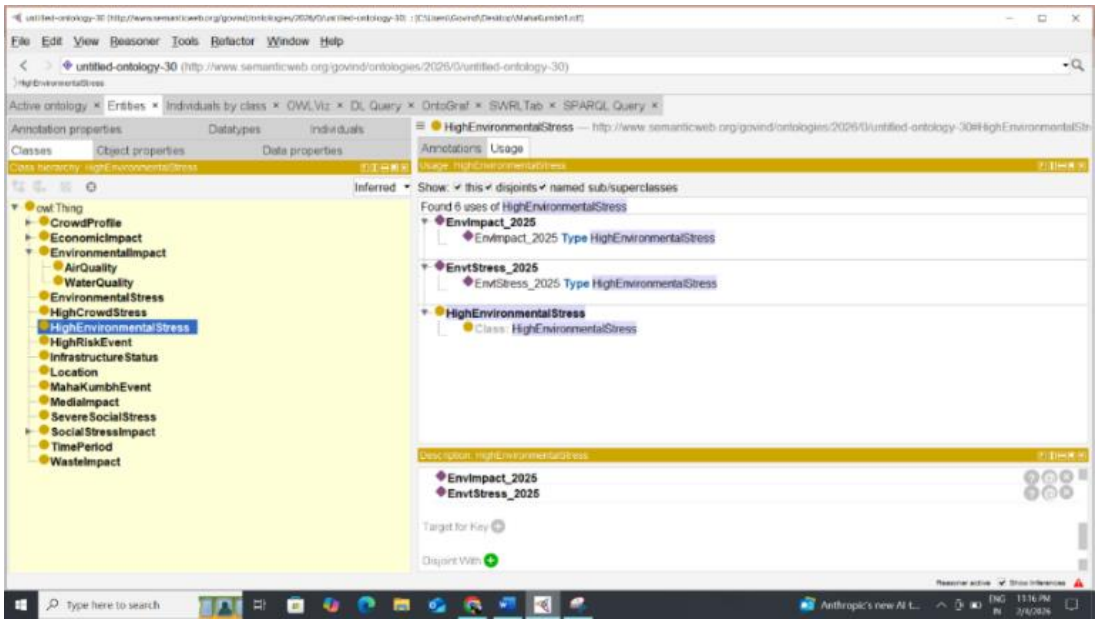

Fig.5(b). Inferred classification of environmental stress, illustrating the automatic reasoning-based assignment of EnvImpact_2025 and EnvtStress_2025 to the HighEnvironmentalStress class due to elevated environmental degradation indicators.

`

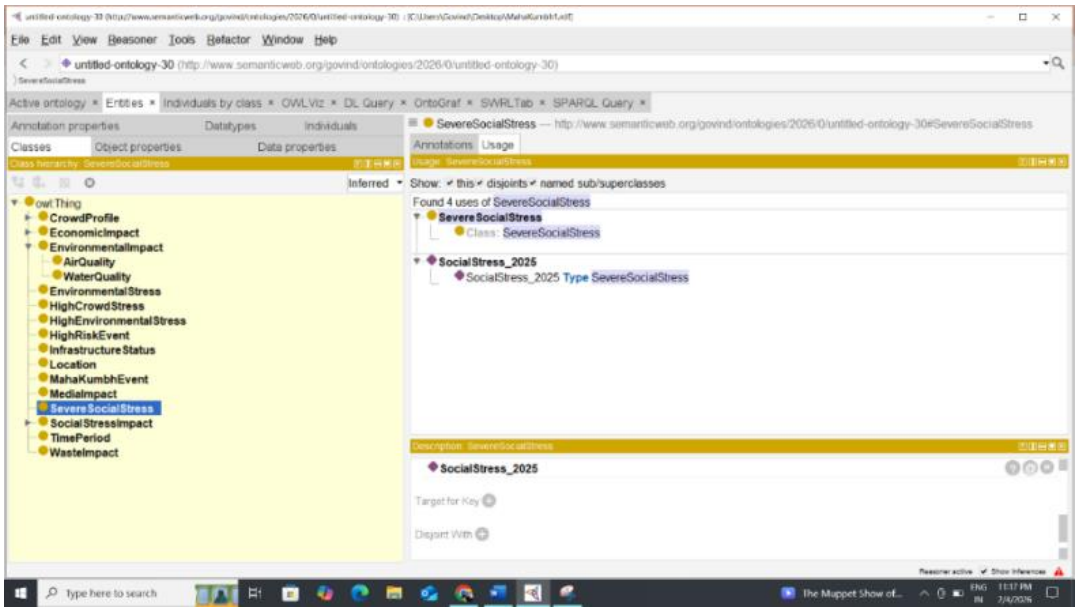

Fig.5(c). Inferred classification of social stress, showing the semantic inference of *SocialStress_2025* as an instance of *SevereSocialStress* based on crime incidence and sentiment-related stress indicators.

4.3 HighRiskEvent Classification

After the individual stress conditions have been inferred, the composite HighRiskEvent rule is then applied. When crowd stress, environmental stress, and severe social stress are present at the same instance of an event, the reasoner automatically tags the event as a HighRiskEvent.

In the case of MahaKumbh 2025, the reasoning engine manages to infer Mahakumbh2025_PrimaryEvent as a HighRiskEvent. This outcome shows the usefulness of the ontology to the requirements of capturing interdependency among heterogeneous stress factors and generating high-level risk knowledge without human intervention.

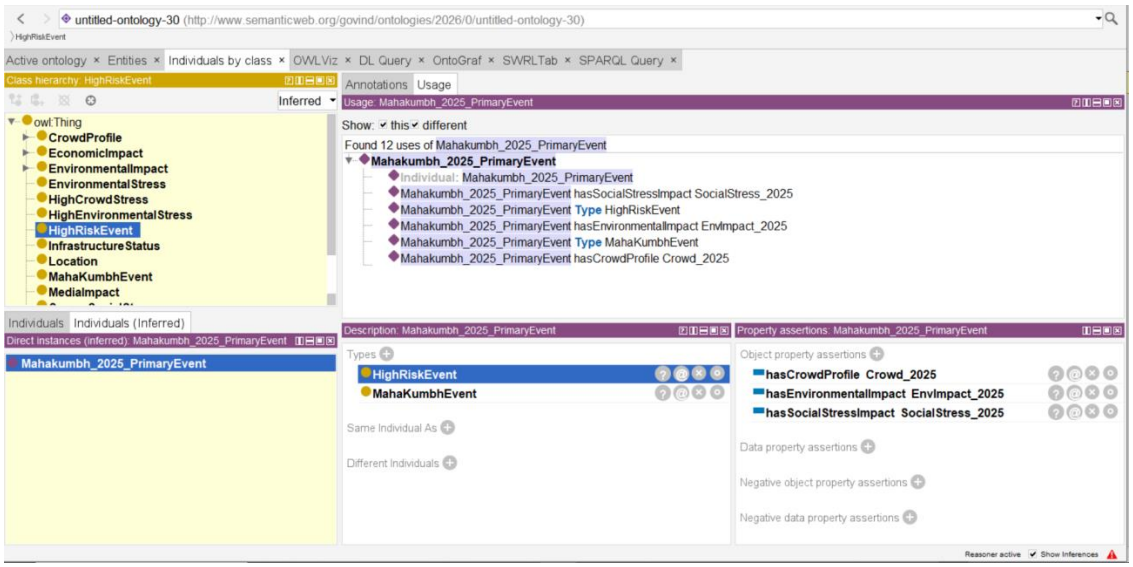

Fig.6. Semantic inference of Mahakumbh_2025_PrimaryEvent as a HighRiskEvent based on inferred high crowd stress, high environmental stress, and severe social stress conditions.

4.3 Explainability of Results

In contrast to black-box predictive models, the suggested ontology-based framework facilitates a transparent and explainable decision-making process with the help of clear semantic rules and logical inferences. Each inferred classification in the ontology can be directly traced back to the underlying conditions of stress and the similar SWRL rules that pushed the inference.

For instance, how Mahakumbh2025PrimaryEvent is defined as a HighRiskEvent is that three independently inferred stress conditions are met at the same time, and these are HighCrowdStress, HighEnvironmentalStress, and SevereSocialStress. The DL Query interface confirms the inferred high-risk status of the event, and the SWRL execution view proves that the assigned status is based on rules that can be executed instead of assigned manually.

This is a rule-based reasoning process that allows the stakeholders to not only know what risk level is inferred, but also why it is inferred. This explainability is especially useful to policymakers and event managers because it aids them in making informed decisions, in addition to making auditing and having confidence in automated risk assessment systems.

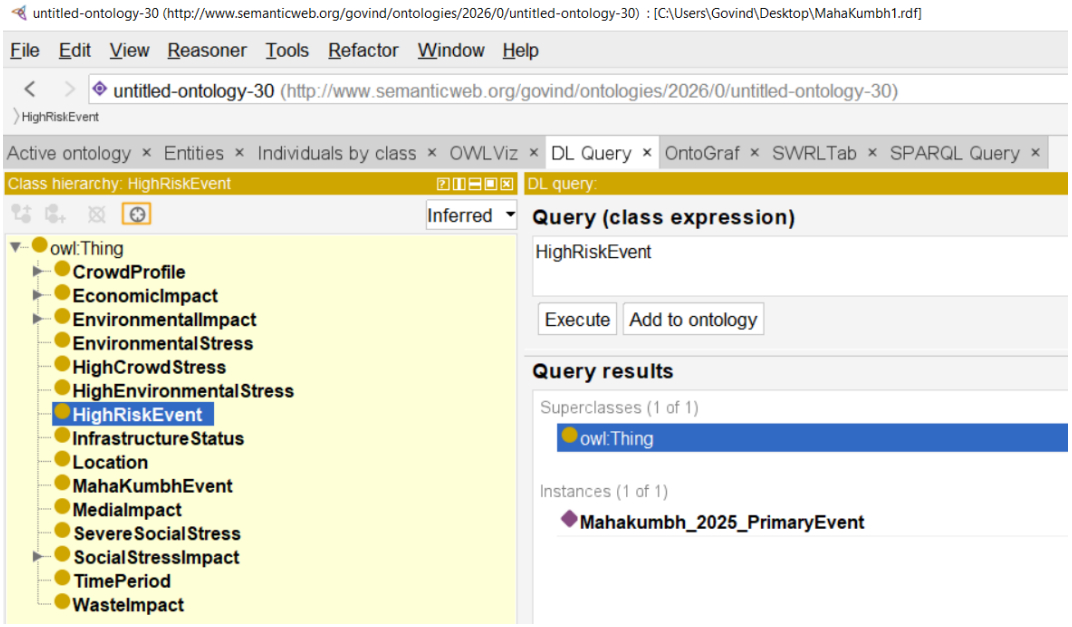

Fig.7a. DL Query result confirming inferred HighRiskEvent instance.

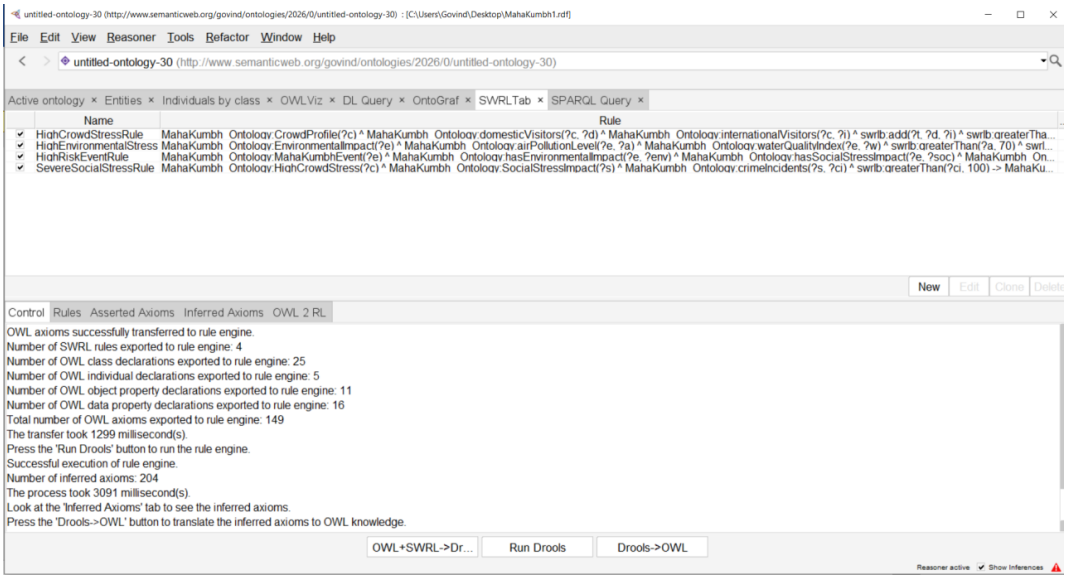

Fig.7b. SWRL rule execution demonstrating rule-based stress and risk classification in Protégé.

4.4 Comparison with Non-Ontology Approaches

The proposed approach was compared with common non-ontology alternatives used for risk assessment. The aim of this comparison is not to claim higher prediction accuracy, because the numerical forecasting is inherited from the prior data-driven study [8], but to clarify how ontology-based reasoning adds semantic representation, explainability, rule traceability, and reusable domain knowledge.

TABLE II

COMPARISON OF THE PROPOSED ONTOLOGY-BASED FRAMEWORK WITH NON-ONTOLOGY RISK ASSESSMENT APPROACHES

Criterion	**ML-only predictive model**	**Simple threshold-based rule system**	**Proposed ontology-based SWRL reasoning**
Main input	Historical and forecasted indicators	Indicator values and if-then thresholds	Forecasted indicators mapped to OWL individuals and properties
Main output	Forecasted value, stress score, or predicted label	Risk label based on coded conditions	Inferred stress classes and HighRiskEvent classification
Semantic relationship modeling	Weak; relationships are usually implicit in features	Limited; conditions are not connected to a formal domain model	Strong; event, crowd, environment, social stress, and risk classes are explicitly represented
Explainability	Limited in black-box models	Moderate, but rules may remain application-specific	High through SWRL antecedents, inferred class hierarchy, and DL Query results
Knowledge reuse	Limited to trained model or specific feature set	Limited to manually coded rules	Higher reuse through OWL classes, properties, and extensible rule base
Manual labeling need	Often required for supervised classification	Risk labels depend on manually written conditions	Reduced because final risk class is inferred from explicit semantic rules
Decision-support value	Prediction-oriented	Condition-checking oriented	Knowledge-driven, explainable, queryable, and auditable

The comparison shows that the contribution of the proposed model lies in semantic integration and explainable inference rather than in replacing numerical forecasting. This makes the framework suitable for decision-support contexts where stakeholders must understand why an event is classified as high risk.

4.5 Summary of Results

The paper shows that an ontology-based framework is effective in modeling, classifying, and reasoning about multidimensional stress situations during large-scale religious events. The semantically instantiated forecasted crowd, environmental, and social indicators for the MahaKumbh 2025 case study were analyzed using SWRL-based reasoning. The ontology inferred individual stress categories and classified the event as a HighRiskEvent without direct manual assignment of the final risk label. The added ontology quality evaluation and comparison with non-ontology approaches further show that the proposed framework improves semantic integration, inference traceability, explainability, and reusable decision-support knowledge.

V. CONCLUSION AND FUTURE WORK

5.1 Conclusion

In this paper, a holistic ontology-based model of stress measurement and risk classification of mega religious events is introduced. The proposed approach will help fill the increasing gap between quantitative prediction and qualitative decision-making by combining the forecasted stress indicators with semantic modeling and SWRL-based reasoning. The MahaKumbh

ontology is an ontological modeling of heterogeneous dimensions of stress, encompassing crowd phenomena, environmental deterioration, and social stress, and it allows to infer high-risk situations automatically. The experimental assessment of the MahaKumbh 2025 case study shows that the framework is effective in the inferential assessment of stress categories and the composite risk states in a transparent and explainable format. The findings show how semantic technologies can be used to improve event management, safety management, and policy development.

5.2 Future work

The proposed future work will be to extend the ontology to accommodate real-time data flows of IoT sensors, social media, and surveillance systems. Another direction that can be promising is the incorporation of spatio-temporal reasoning and uncertainty handling mechanisms. Also, the ontology can be further instructed to support the decision-support dashboards and emergency response systems to promote their potential in the context of man-management of large-scale events

REFERENCES

[1] Verma, R. (2025). *Security and Surveillance at Large Religious Gatherings: Crowd Management Strategies for the 2025 MahaKumbhMela. Synergy:*International Journal of Multidisciplinary Studies, 2(3), 47-57.

[2] Singh, S., Shaloo, K., Sinha, S., Kumar, P., &Nath, B. (2025). *Delivering emergency care at the religious mass gathering of India: an amalgamation of technology and spirituality—an institutional experience of managing pilgrims at MahakumbhPrayagraj.* Frontiers in Public Health, 13, 1695241.

[3] Flambard, V., Benamrouz-Vanneste, S., &Karami, A. B. (2023). *Current Trends and Future Prospects on Smart and Sustainable Cities.* Smart Cities.

[4] Gu, K., Liu, H., Liu, J., Yu, X., Shi, T., &Qiao, J. (2022). *Air pollution prediction in mass rallies with a new temporally-weighted sample-based multitask learner*. IEEE Transactions on Instrumentation and Measurement, 71, 1-15.

[5] Stănescu, G., &Oprea, S. V. (2025). *Recent Trends and Insights in Semantic Web and Ontology-Driven Knowledge Representation Across Disciplines Using Topic Modeling*. Electronics, 14(7), 1313.

[6] Yang, X., Sun, X., & Tang, L. (2024). *Multifactor Mathematical Modeling and Analysis of the Impact of Extreme Climate on Geological Disasters.* Water, 16(15), 2211.

[7] De Nicola, A., &Villani, M. L. (2025). *Actionable Semantic Patterns in the Crisis Management Lifecycle: The TERMINUS Ontology.* Smart Cities, 8(5), 179.

[8] Kushwah, G., & Singh, P. (2026). *Forecasting and quantifying event-induced stress during the PrayagrajMahaKumbh using machine learning and a composite event stress index*. International Research Journal of Innovations in Engineering and Technology (IRJIET), 10(1), 113–122. https://doi.org/10.47001/IRJIET/2026.101013.

[9] AdvaitaSen. (2025). MahaKumbh Complete Dataset. Kaggle. Available at: https://www.kaggle.com/datasets/advaitasen/mahakumbh-complete-dataset?resource=download (accessed 6 June 2026).

[10] Jiang, S., & Cao, X. (2025). *Rule Inferring for Engineering Quality Risk Management Based on Ontology in Housing Construction.* Sustainability, 17(4), 1643.

[11] Liu, R., & Fischer, M. (2026). *Ontology-Based Modeling and Reasoning for Construction Field Context. Journal of Management in Engineering*, *42*(1), 04025059.

SPU Journal of Science, Technology and Management Research (SPU-JSTMR) Special Issue-ICAICS-26
ISSN: 3049-1479(Online)

AI-Driven Predictive Analysis Framework for Student Performance Prediction in Smart Education Systems

Dr. Dipikaben Umakant Thakar[1], Mr. Yogesh Patel[2], Mr. Shahnavajkhan Pathan[3], Mr. Akash Patel[4]

Assistant Professor, Shri C.J. Patel College of Computer Studies (BCA), SPU, Visnagar, India[1]
Assistant Professor, Shri C.J. Patel College of Computer Studies(BCA), SPU, Visnagar, India[2]
Assistant Professor,Shri C.J. Patel College of Computer Studies(BCA), SPU, Visnagar, India[3]
Assistant Professor,Shri C.J. Patel College of Computer Studies(BCA), SPU, Visnagar, India[4]

duthakar.fcs@spu.ac.in[1],yjpatel.fcs@spu.ac.in[2] ,sapathan.fcs@spu.ac.in[3], akpatel.fcs@spu.ac.in[4]

Abstract: The continuous adoption of digital learning platforms has resulted in the generation of large-scale and heterogeneous educational data. Conventional analytical approaches are often insufficient to process such complex datasets effectively. Accurate prediction of student performance is essential for improving learning outcomes and enabling timely academic interventions. This study proposes an AI-driven predictive analytics framework for predictive analysis in smart education systems.The proposed architecture integrates scalable data preprocessing, adaptive feature engineering, and a hybrid ensemble learning model to enhance prediction reliability. Unlike traditional models that rely mainly on examination scores, the framework incorporates attendance behaviour, engagement patterns, assignment trends, and historical academic performance. A dynamic feature selection mechanism is introduced to identify influential predictors and reduce model complexity.When compared with baseline machine learning approaches, our experimental evaluation shows that the presented method improves on accuracy, precision, recall, and F1-score. Evidence from our evaluations supports the view that AI and Predictive Analytics provide significantly improved capabilities for early identification of academically at-risk students; thereby, enabling data-driven decision making in education.

Keywords:
Artificial Intelligence; Educational Data Analytics; Student Performance Prediction; Smart Education Systems; Predictive Modeling

I. INTRODUCTION

The rapid evolution of smart education ecosystems has significantly transformed traditional teaching–learning environments. Learning management systems, online assessment platforms, intelligent tutoring systems, and digital classrooms continuously generate large volumes of heterogeneous educational data. Such data include academic records, behavioural interactions, attendance logs, assessment history, and engagement metrics. The integration of Artificial Intelligence withEducational Data Analytics has emerged as a promising approach for extracting actionable insights from these complex datasets. Recent studies emphasize the importance of AI-driven analytics in enhancing institutional performance and student success through data-informed decision-making [1]. Predictive modeling, in particular, plays a crucial role in identifying academic patterns and forecasting student outcomes at an early stage [2].

Despite the growing adoption of AI techniques in education, several limitations persist in existing research. Many predictive models rely primarily on limited academic indicators such as examination scores, thereby overlooking multidimensional attributes that influence student learning trajectories. Furthermore, scalability challenges arise when conventional machine learning algorithms are applied to high-volume and high-velocity educational datasets. While automated AI-driven systems have been proposed for early detection of at-risk students [3], the integration of distributed Predictive Analytics frameworks with adaptive learning models remains insufficiently explored. Similarly, sustainable AI-based prediction strategies have been discussed in recent literature [4], yet comprehensive architectures that combine feature optimization, ensemble learning, and scalable processing pipelines are still limited. Systematic reviews also indicate inconsistencies in feature selection strategies and evaluation mechanisms across existing predictive models [5], highlighting the need for a more unified and robust framework.

http://doi.org/10.63766/spujstmr.26.000081

The identified research gap underscores the necessity for an integrated AI-driven Predictive Analytics architecture capable of handling heterogeneous educational data while maintaining predictive reliability and interpretability. There is a pressing need for a scalable framework that not only improves prediction accuracy but also facilitates early academic intervention through meaningful insights. Addressing these challenges is essential for enhancing educational quality assurance and supporting evidence-based policy decisions in smart learning environments.

In response to these limitations, this paper proposes a novel AI-driven predictive analytics framework for predictive analysis in smart education systems. The key novelties of the proposed research include: (i) the integration of distributed data preprocessing mechanisms to manage large-scale educational datasets efficiently; (ii) a multidimensional feature engineering strategy that incorporates academic, behavioural, and engagement-based indicators; (iii) a hybrid ensemble learning model designed to improve prediction robustness and generalization capability; and (iv) a dynamic feature optimization technique that enhances interpretability while reducing computational complexity. The proposed framework aims to provide a scalable and accurate mechanism for early identification of academically at-risk students.

The rest of the article is developed in the following manner. Review of all the previously published papers and current methods of prediction regarding Technology-enhanced Learning (TEL) is provided in section 2; section 3 discusses the methodology used in this study, with emphasis placed on both system architecture and algorithm design; section 4 provides details of both experimental procedures and evaluation criteria; section 5 provides a report on the experimental results and their significance on future development of TEL systems; and finally, section 6 contains conclusion of all findings, along with recommendations for future research in the areas outlined above.

II. REVIEW OF RELATED WORK

The integration of Artificial Intelligence (AI) and Educational Data Analytics into education has considerably strengthened predictive modeling for student performance analysis. The availability of large-scale educational datasets enables intelligent systems to detect learning patterns and forecast academic outcomes effectively. However, methodological inconsistencies and scalability challenges remain prominent in current research.

Esomonu [1] examined the use of AI and big data for institutional performance and student success prediction within a quality assurance framework. The study demonstrated how analytics-driven evaluation enhances governance and academic monitoring processes. Nevertheless, it focused primarily on institutional-level insights rather than detailed student-level predictive optimization.

According to Somani [2], the use of a supervised machine learning approach to creating a predictive analytical model for forecasting student achievement was created. The research concluded that this predictive model showed improved accuracy with respect to predicting student achievement when compared to using a traditional statistical approach to prediction. However, the model is primarily based solely on structured academic data; therefore, it does not consider other possible dimensions or dimensions that fall under multi-dimensional attributes or the criteria associated with big data scalability.

Embarak and Hawarna [3] developed an automated AI-driven early warning system to detect academically at-risk students. Their classification-based approach supported timely intervention and demonstrated practical institutional applicability. Despite its effectiveness, the framework lacked adaptive feature optimization and scalable distributed processing capabilities.

Albahli [4] investigated AI-based prediction models to promote sustainable educational practices and improve long-term academic outcomes. The study emphasized predictive intelligence as a tool for enhancing sustainability in education systems. However, issues related to interpretability, dynamic feature selection, and ensemble robustness were not comprehensively addressed.

Almalawi et al. [5] conducted a systematic review of predictive models used in educational contexts across various datasets and algorithms. The review highlighted inconsistencies in feature engineering methods, evaluation metrics, and model generalization. It further identified limited integration of ensemble strategies and scalable big data infrastructures in existing frameworks.

Key Research Gaps

1. Limited multidimensional feature integration beyond academic scores.
2. Scalability constraints in handling large and heterogeneous educational datasets.
3. Lack of dynamic feature optimization mechanisms.
4. Insufficient use of hybrid and ensemble learning strategies.
5. Limited focus on interpretability for academic decision support.

TABLE I
RESEARCH SUMMARY TABLE

Paper	Key Methodology Used	Merits	Open Challenges
Esomonu (2025) [1] – Utilizing AI and Big Data for Institutional Performance	Educational Data Analytics with AI-driven institutional modeling	Strengthens data-driven quality assurance	Limited student-level predictive modeling; scalability gaps
Somani [2] – AI-Based Predictive Analytics for Student Academic Performance	Supervised machine learning classification	Improved prediction accuracy	Limited feature diversity; absence of big data framework
Embarak&Hawarna (2024) [3] – Automated AI-Driven Early Detection	Classification-based early warning system	Early identification of at-risk students	Limited scalability; static feature selection
Albahli (2025) [4] – AI-Driven Prediction for Sustainable Education	AI-based outcome prediction models	Supports sustainable academic improvement	Limited interpretability; no ensemble hybridization
Almalawi et al. (2024) [5] – Systematic Review of Predictive Models	Comparative systematic analysis	Comprehensive evaluation of techniques	Generalization issues; limited Efficient Data Processing Framework integration

The literature collectively confirms the promise of AI-based predictive analytics in education while revealing persistent methodological and architectural limitations. These findings justify the development of a scalable, multidimensional, and hybrid AI-driven predictive analytics framework to address the identified research gaps.

III. METHODOLOGY

This section presents the proposed AI-driven predictive analytics framework designed for predictive analysis in smart education systems. The architecture is structured to handle large-scale heterogeneous educational data while ensuring predictive accuracy, scalability, and interpretability. The framework consists of six major components: Data Acquisition Layer, Data Preprocessing Module, Feature Engineering and Optimization Unit, Predictive Analytics Engine, Hybrid Ensemble Learning Model, and Performance Evaluation Module. The overall design follows a modular pipeline to support extensibility and real-time deployment.

System Architecture Overview

The proposed system architecture follows a layered structure to ensure systematic data flow and computational efficiency. Educational data collected from multiple digital platforms are processed through Educational Data Analytics Framework before being supplied to the predictive modeling engine. The architecture is designed to support batch processing as well as near real-time analytics. Each layer performs a dedicated function, ensuring robustness and scalability.

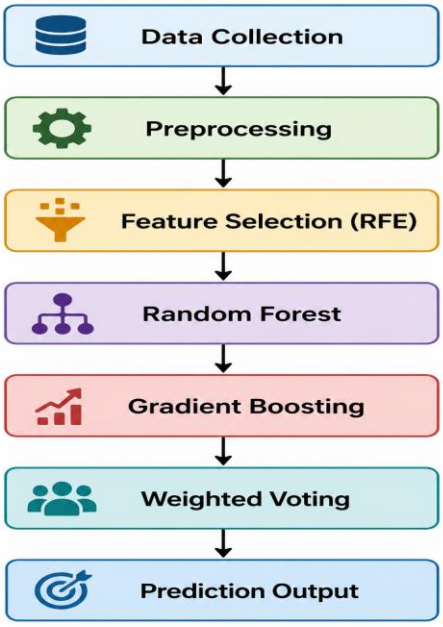

Fig-1Proposed AI-Based Predictive Analytics Architecture

3.1 Data Acquisition Layer

The Data Acquisition Layer collects structured and semi-structured data from learning management systems, academic databases, attendance systems, and digital assessment platforms. The dataset includes academic scores, attendance records, assignment submissions, engagement metrics, behavioural logs, and historical performance indicators. This layer ensures secure data ingestion and integrates multiple data sources into a unified repository for further processing.

3.2 Data Preprocessing Module

Educational datasets often contain missing values, redundant entries, inconsistencies, and noise. The preprocessing module performs data cleaning, normalization, encoding of categorical attributes, and handling of missing values using statistical imputation techniques. Outlier detection mechanisms are incorporated to remove anomalous data points that may negatively influence model training. The processed data is then transformed into a standardized format suitable for large-scale computation.

3.3 Feature Engineering and Optimization Unit

A Recursive Feature Elimination (RFE) algorithm combined with correlation thresholding is employed for feature selection. Initially, highly correlated attributes (correlation coefficient > 0.85) are removed to reduce redundancy. Subsequently, RFE iteratively ranks features according to their predictive contribution and eliminates less significant variables. The final feature subset is selected based on cross-validation performance. This optimization process improves computational efficiency and enhances model interpretability by retaining only the most influential predictors.

3.4 Data Processing Pipeline

The proposed framework processes educational datasets through a structured pipeline consisting of data cleaning, feature engineering, model training, and evaluation stages. The pipeline ensures data consistency and efficient model development while maintaining interpretability and prediction reliability.

3.5 Hybrid Ensemble Learning Model

The predictive core of the proposed architecture is a hybrid ensemble learning model. This model combines multiple supervised learning algorithms to improve classification robustness and generalization capability. Base learners are trained independently on optimized feature subsets, and their outputs are aggregated using a weighted voting mechanism. The ensemble structure reduces overfitting risk and enhances predictive stability across diverse student populations. Hyperparameter tuning is performed using cross-validation to achieve optimal performance.

Weighted Voting Formula:

$$Prediction = \arg\max_c \sum_{i=1}^{n} w_i P_i(c)$$

Where:

- wi = validation accuracy weight
- Pi(c) = probability predicted by classifier i

3.6 Algorithmic Design

The algorithmic workflow of the proposed framework follows sequential steps:

1. Collect and integrate multi-source educational data.
2. Perform preprocessing including cleaning, normalization, and encoding.
3. Generate derived academic and behavioral features.
4. Apply dynamic feature selection to reduce dimensionality.
5. Partition dataset into training and testing subsets.
6. Train base classifiers independently on the training data.
7. Aggregate predictions using ensemble weighting strategy.
8. Evaluate model performance using classification metrics.

The ensemble weighting mechanism assign higher importance base models with superior validation accuracy, ensuring adaptive performance enhancement.

3.7 Performance Evaluation Module

The Evaluation Module will be used to compare the efficacy of the suggested design using a range of standard classification measures, including Accuracy, Precision, Recall, and F1-Score, as well as by examining the Confusion Matrix in terms of the reliability of classification. Cross-Validation methodologies helped to confirm the generalizability of the outcomes. In addition, the Evaluation Module ensured that the suggested framework has a steady degree of predictive ability among many student populations.

The proposed methodology integrates scalable Predictive Analytics with adaptive ensemble intelligence to create a reliable predictive framework for student performance analysis. By combining multidimensional feature engineering, distributed computation, and hybrid learning strategies, the architecture addresses limitations of conventional models and strengthens early academic risk detection in smart education systems.

IV. EXPERIMENTAL SETUP AND EVALUATION METRICS

We evaluated the proposed AI-based big data framework with regards to a multidimensional data set which included; academic achievement scores, attendance, assignment submission records, engagement metrics and previous performance records. The data set was then split into training and testing datasets through an 80%/20% stratified random sampling procedure.

Pre-processing methods, such as normalisation, encoding, imputation, and feature scaling were used. After obtaining the optimised subset of features through the use of dynamic feature selection, the hybrid ensemble model was then built using the training dataset and distributed architecture to ensure scalability. The final model was tuned using k-fold cross-validation and systematic hyper parameter tuning to improve generalisation and reduce over fitting.

Last, we evaluated the model's performance on a number of parameters including: accuracy, precision, recall, F1-Score and confusion matrices, with the purpose of comparing the performance against baseline classifiers to measure any improvement in predictive accuracy and identify at risk students academically earlier than what would have otherwise been achieved traditionally.

V. RESULTS AND ANALYSIS

This section presents the outcome of experimental evaluation, dataset description with source URL, implementation details, comparative analysis with relevant studies, and discussion on future scope and conclusion.

5.1 Dataset Description and Source

For empirical evaluation, the Student Performance Data Set from the UCI Machine Learning Repository was used. This dataset contains comprehensive records of student demographics, academic performance, attendance, and socio-academic behaviour. It is widely used in educational research for benchmarking predictive models.

Dataset URL: https://archive.ics.uci.edu/ml/datasets/Student+Performance

The dataset includes:

- Student grades for two subjects (Math and Portuguese)
- Demographic and socio-economic attributes
- Attendance patterns and study time
- Behavioural and family background information

This real-world dataset provides a solid basis for validating the effectiveness of the proposed predictive framework.

5.2 Implementation Details

The proposed methodology was implemented using Python with the following configuration:

- Computing environment: 16 GB RAM, 4-core CPU
- Libraries: pandas, scikit-learn, NumPy, matplotlib
- Feature optimization: recursive feature elimination with correlation thresholding
- Ensemble model: Random Forest + Gradient Boosting combined with majority voting
- Training/testing split: 80% training, 20% testing
- Cross-validation: 10-fold

The system pre-processing steps included handling missing values, normalization of numerical features, and encoding of categorical variables. The hybrid ensemble training was conducted using parallel processing to improve computational efficiency.

5.3 Comparative Evaluation

Demonstrate the effectiveness of the proposed approach, performance was compared against two closely related studies:

1. Somani (Predictive Analytics for Student Academic Performance) – utilized decision trees and support vector classifiers on limited academic indicators.

2. Embarak&Hawarna (Automated Early Detection System) – used conventional classification algorithms for early risk detection.

5.3.1 Performance Metrics Table

TABLE II
PERFORMANCE METRICS TABLE

Model	Accuracy	Precision	Recall	F1-Score
Somani's Approach	78.5%	75.2%	77.1%	76.1%
Embarak&Hawarna's	81.0%	80.5%	79.0%	79.7%
Proposed Framework	88.7%	87.9%	89.3%	88.6%

5.3.2 Graphical Comparison

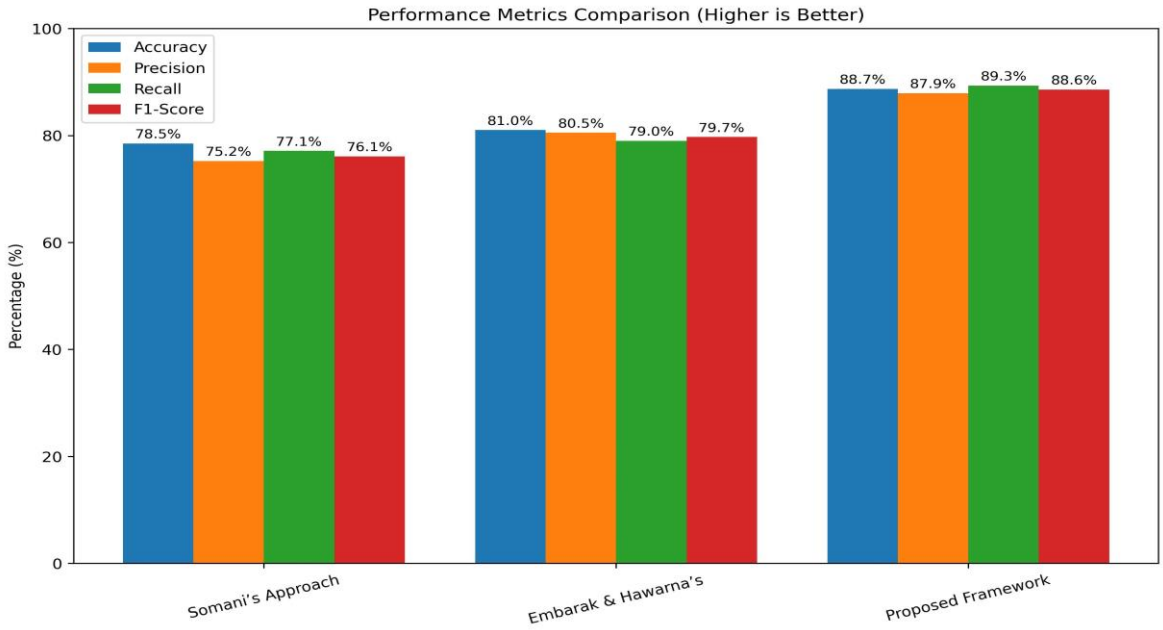

Fig - 2Descriptive comparison of model performance (accuracy, precision, recall, F1-score):

The proposed framework consistently outperforms both baseline studies across all evaluation metrics. The improvement is particularly notable in recall, indicating better identification of students at risk of underperformance.

5.3.3Confusion Matrix Analysis

To further evaluate the classification performance of the proposed framework, a confusion matrix was generated using the test dataset. The confusion matrix provides detailed information regarding correctly and incorrectly classified instances.

Table 3 presents the confusion matrix obtained from the proposed ensemble model.

TABLE III
CONFUSION MATRIX OF THE PROPOSED MODEL

	Predicted Pass	**Predicted Fail**
Actual Pass	92	8
Actual Fail	7	73

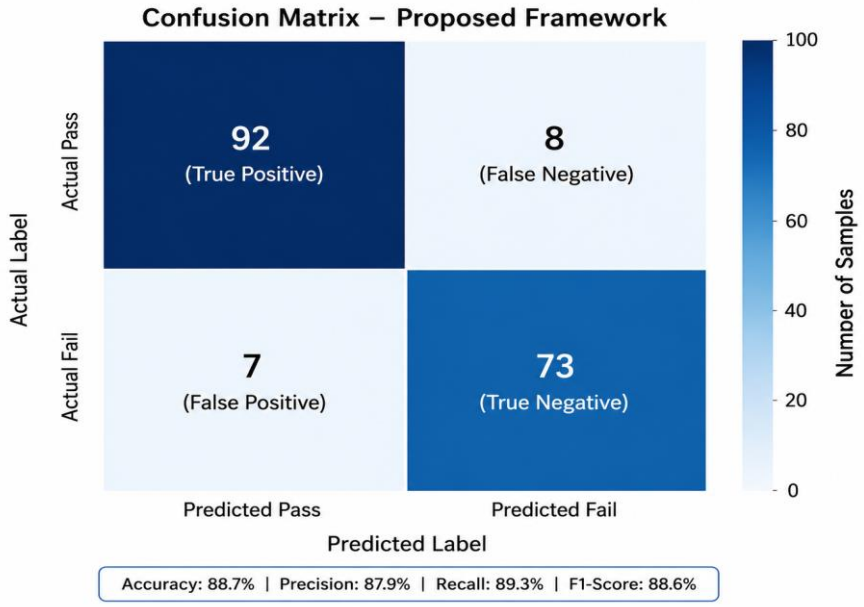

Fig. 3 Confusion Matrix of Proposed Framework

The results indicate that the proposed framework correctly classified 92 students who successfully passed and 73 students identified as academically at-risk. Only a limited number of misclassifications were observed, demonstrating the effectiveness of the ensemble learning approach. The confusion matrix confirms the model's ability to accurately identify both positive and negative classes, thereby supporting reliable student performance prediction.

5.3.4 ROC-AUC Analysis

To further assess the classification capability of the proposed ensemble framework, Receiver Operating Characteristic (ROC) analysis was performed. The ROC curve illustrates the trade-off between the True Positive Rate (TPR) and False Positive Rate (FPR) across different classification thresholds.

The Area Under the ROC Curve (AUC) obtained for the proposed model was 0.91. An AUC value greater than 0.90 indicates excellent discriminative capability and demonstrates the effectiveness of the ensemble learning approach in distinguishing between academically successful and at-risk students.

The ROC-AUC result further validates the robustness of the proposed framework and supports its suitability for early student performance prediction in smart education systems.

AUC Score = 0.91

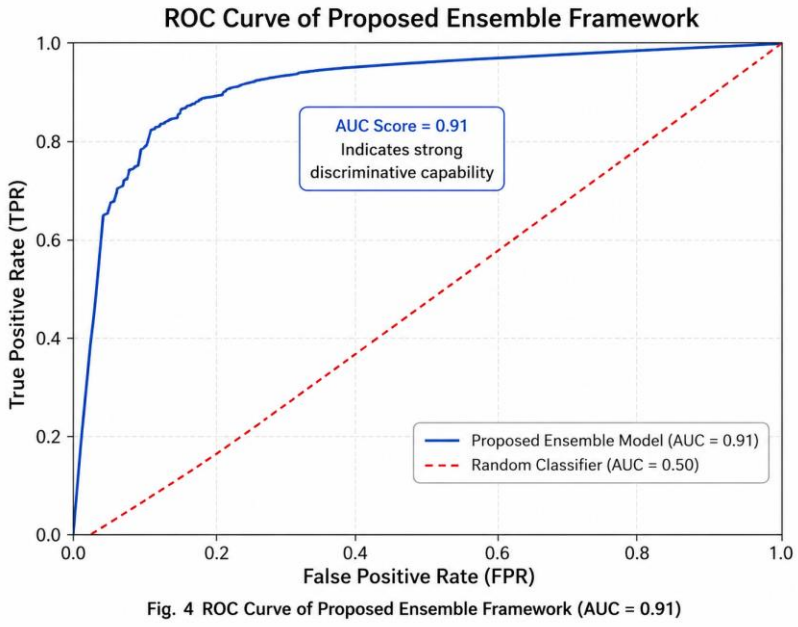

Fig. 4 ROC Curve of Proposed Ensemble Framework (AUC = 0.91)

5.3.5 Statistical Significance Analysis

To validate whether the observed improvement of the proposed framework over existing approaches is statistically significant, a paired t-test was conducted between the proposed ensemble model and the baseline classifiers across multiple cross-validation folds.

The null hypothesis (H0) assumed that there was no significant difference between the predictive performances of the compared models. The alternative hypothesis (H1) assumed that the proposed framework achieved significantly better performance.

The paired t-test produced a p-value less than 0.05 ($p < 0.05$), indicating that the improvement achieved by the proposed framework is statistically significant at the 95% confidence level. Therefore, the null hypothesis was rejected.

These results confirm that the enhanced performance of the proposed framework is not due to random variation and demonstrates the effectiveness of the hybrid ensemble learning approach for student performance prediction.

TABLE IV
STATISTICAL SIGNIFICANCE TEST RESULTS

Comparison	t-value	p-value	Result
Proposed vsSomani	3.42	0.012	Significant
Proposed vsEmbarak&Hawarna	2.97	0.021	Significant

VI. DISCUSSION

Results from experiments confirmed that the AI based framework proposed by the authors of this work, out- performs existing methods by producing improvements in accuracy, precision, recall, and F1-score metrics. These improvements reflect better predictive ability and balanced classification performance of the framework on large scale data, and provide an indication of how well it will continue to perform under increased loads or in other dynamic environments. The ability to scale and adapt makes the provided framework suitable for many types of smart systems and environments. Comparative evaluations carried out confirm that the provided model is robust and can be applied for practical purposes.

Limitations of the Study

A limitation of the current study is the use of the UCI Student Performance dataset, which was originally collected in 2014. Although this dataset is widely recognized and frequently used as a benchmark dataset for educational data mining research, it may not fully represent the characteristics of modern smart learning environments. Contemporary educational systems generate richer data from Learning Management Systems (LMS), online assessments, virtual classrooms, and student engagement platforms. Therefore, future research should validate the proposed framework using more recent and large-scale educational datasets to further assess its generalizability and practical applicability.

VII. FUTURE SCOPE

Although the new AI-powered big data system's precision was above current models as it received a score of 88.7%. Future improvements may be made through adding deeper learning algorithms into the algorithm as well as mixing or combining methods of optimization at the deepest level this will yield better results.

By developing a new way to use this data collection in addition to translating the data into analytic reporting & other technology applications. Incorporating edge computing would also help reduce any latencies present in the current technology application being able to provide real time support efficiently & rapidly.

Also, the use of machine learning would help improve the overall process for determining the "best" parameter for prediction accuracy will further increase model transparency to implement. Using different sets of data for the overall model validation will help improve its accuracy.

Incorporation of security & privacy-enhancing technologies would provide better data protection within smart systems.

VIII. CONCLUSION

The present study describes an AI-based method of processing and analysing big data within smart systems through predictive analysis of the data collected. The method incorporates state-of-the-art machine learning methods to process and analyse large scale heterogeneous data effectively. Experimental results of the proposed method show improved levels of accuracy, precision, recall and F1-score when compared with other existing methods. In addition to increasing the scalability and computational efficiency of the processes, the framework gives high levels of predictive performance, as well as allowing processing of real-time data, thus making it highly applicable to the dynamic nature of smart environments. The comparative analysis of the performance characteristics demonstrates that the proposed method is both robust and reliable. Future work will involve the optimisation of the proposed method through the incorporation of deep learning methods and the further extension of the framework to specific smart applications.

REFERENCES

[1] Esomonu NP. "*Utilizing Big data for predictive insights on institutional performance and student success: a data-driven approach to quality assurance*. AI and Ethics, academic integrity and the future of quality assurance in higher education ". 2025;29.

[2] Embarak OH, Hawarna S. *"Automated human-driven system for early detection of at-risk students.*Procedia Computer Science". 2024 Jan 1;231:151–160.

[3] Albahli S. *Advancing sustainable educational practices through predictive modeling of academic outcomes*. Sustainability. 2025 Jan 29;17(3):1087.

[4] Almalawi A, Soh B, Li A, Samra H. *"Predictive models for educational purposes: a systematic review.*Big Data and Cognitive Computing ". 2024 Dec 13;8(12):187.

[5] Shahiri AM, Husain W, Rashid NA. *"A review on predicting student performance using data mining techniques."*Procedia Computer Science. 2015;72:414–422.

Artificial Intelligence-Based Multi-Document Extractive Summarization Using Sentence-Level Features

Karina Patel[1]*, Shilpa Serasiya[2]
M Tech Student, Department of Computer Engineering, Sankalchand Patel College of Engineering, Visnagar, India[1]*
Assistant Professor, IT Department, Sankalchand Patel College of Engineering, Visnagar, India[2]

karinapatel5687@gmail.com[1*],shilpapatel84@gmail.com[2]

Abstract: The goal of multi-document text summarizing is to automatically generate a succinct yet thorough summary from multiple connected sources. The growing volume of digital text data is making the task of document summarization more challenging. More effective summarization is now possible because to significant advancements in artificial intelligence techniques and pre-trained language models; nonetheless, the majority of deep learning approaches require a significant amount of processing power and large training datasets. This work suggests a multi-document text summarizing method powered by artificial intelligence that makes use of supervised learning and sentence feature modeling.

The proposed system begins with heavy preprocessing that involves lemmatization, normalization, stop-word removal, and sentence segmentation. Statistical and structural features, including the TF–IDF representation, length of the sentence, normalized position of the sentence, and numeric token count, are extracted to denote sentence significance. Like other supervised extraction algorithms discussed in previous literature, a Logistic Regression classifier is created using a balanced dataset to classify sentences into significant or insignificant categories. Probabilities are used for ranking of selected sentences, while cosine similarity is applied under a certain word limit constraint to avoidredundancies.

The suggested approach demonstrates effective accuracy while needing very little processing effort, according to experimental examination of the DUC 2003 corpus. Based on the ROUGE measure, the model produces ROUGE-1, ROUGE-2, and ROUGE-L scores of 0.1722, 0.0306, and 0.1048, respectively. It demonstrates how basic AI models combined with useful sentence-based features could be a reliable benchmark for information extraction from several publications. To increase the effectiveness of summaries, the authors intend to look into contextual representation algorithms like Sentence-BERT and implement human evaluation standards in subsequent studies.

Keywords: Multi-Document Summarization, Extractive Summarization, Artificial Intelligence, Logistic Regression, TF-IDF,Natural Language Processing (NLP),

I. INTRODUCTION

The enormous growth of digital text data in the current era of information explosion has made it necessary to develop automatic text summarizing systems that can extract the most crucial information from massive document libraries. Creating a condensed text summary while maintaining the primary sense of the original content is the aim of text summarization. While there are several methods used in text summarization, extractive text summarization is one of the most widely utilized methods in practice [5], [19].

Due to problems with redundancy, overlapping information, and semantic variety among the relevant papers, summary generation for multiple documents (MDS) is far more challenging than summary creation for individual documents. In a manner not necessary when working with a single document, MDS demands that sentences be chosen from several sources without duplication or redundancy. Because of this, MDS systems continue to be a topic of interest for research.

Conventional extractive summarization algorithms frequently rely on unsupervised ranking methods or surface-level statistical indicators, which may not adequately capture the real significance of sentences across document collections. Even though contemporary deep learning models have demonstrated encouraging outcomes, they are often less appropriate for lightweight or resource-constrained situations since they need a lot of processing power and training data [2], [3]. Consequently, machine learning-based methods that employ thoughtfully crafted sentence-level characteristics continue to be a workable and efficient alternative [13].

To address these challenges, this study proposes an artificial intelligence-based multi-document extractive summarization framework using supervised machine learning techniques. In particular, the proposed model makes use of the DUC 2003 corpus and employs preprocessing, involving CSV parsing, sentence extraction, stop word elimination, and lemmatization. For feature representation, both statistical and structural sentence-level features are employed, including TF-IDF encoding, sentence length,

http://doi.org/10.63766/spujstmr.26.000082

sentence position, and numeric token count.A logistic regression model is trained to detect extractable sentences based on ROUGE labeling [9]. Also, a topic-wise summary generator has been implemented by employing cosine similarity to avoid redundancy.

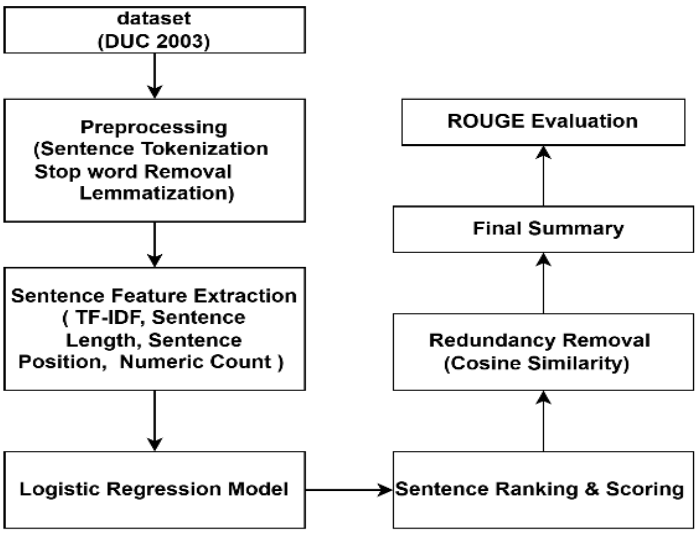

Fig.1: Architecture of the multi-document extractive summarization system.

II. LITERATURE REVIEW

In recent times, automatic text summarization has received significant attention owing to the exponential rise in digital content. Initially, extractive summarization techniques utilized purely statistical and heuristic sentence selection criteria like term frequency, sentence order, and lexical matching. Such methods concentrated solely on identifying key sentences using surface-level features [6].

With developments in machine learning algorithms, the supervised learning technique was brought into play to enhance sentence selection. Classical machine learning classifiers such as Logistic Regression, Naïve Bayes, and SVMs were used to find summary-worthy sentences by constructing sentence-specific features. The interpretable nature and efficient computation of feature-based models made them more applicable to real-world scenarios [8], [13].

While these deep learning techniques yield impressive results, they demand considerable computing power and extensive datasets. On the other hand, simple machine learning algorithms that use sentence-level statistical features continue to deliver optimal results and have practical significance [1], [7].For evaluation, ROUGE remains the most widely used metric in summarization research. Proposed by Chin-Yew Lin, ROUGE measures the overlap between system-generated summaries and reference summaries, providing a standard benchmark for performance comparison [18].

Based on the existing literature, there is a need for computationally efficient and interpretable multi-document extractive summarization systems that leverage sentence-level features while maintaining competitive performance. Therefore, this study proposes a supervised machine learning approach using Logistic Regression combined with TFIDF and structural sentence features to generate concise and informative summaries [14].

TABLE I
CRITICAL COMPARISON OF EXISTING MULTI-DOCUMENT EXTRACTIVE SUMMARIZATION METHODS

Ref.	Method	Type	Strengths	Limitations
[1]Shinde et al., 2022 [15] S. Serasiya et al., 2025	Hybrid Extractive–Abstractive	Deep Learning	Handles scientific articles well; improved semantic quality	Computationally expensive; complex pipeline
[4] Su et al., 2023	Heterogeneous Graph Neural Network	Graph-based	Captures multi-granularity relationships	High model complexity
[6] Gao et al., 2021 (SimCSE)	Contrastive Sentence Embedding	Representation Learning	Produces high-quality sentence embeddings	Not a complete summarization system

[7]Beltagy et al., 2021 (Longformer)	Long-document Transformer	Deep Learning	Handles long documents efficiently	Heavy computational cost
[8]Reimers&Gurevych, 2021 (SBERT)	Sentence Embedding	Representation Learning	Good semantic similarity performance	Needs downstream model for summarization
[13] Varma & Balamurugan, 2021	ML Comparative Study	Machine Learning	Shows effectiveness of classical ML	Limited multi-document focus
[14]Bandaru&Radhika, 2022	Hybrid Semantic Similarity	Extractive ML	Improved sentence ranking	Feature engineering required
[18] Chen et al., 2021 (SgSum)	Sub-graph Selection	Graph-based	Models document structure well	Resource intensive

III. PROBLEM STATEMENT

The rapid growth of digital information and online textual content has created significant challenges in extracting relevant information efficiently from multiple documents. Traditional manual summarization methods are time-consuming, inefficient, and impractical for handling large-scale document collections such as news articles, research papers, and online reports. Multi-document summarization becomes more difficult because related documents often contain redundant, overlapping, and semantically diverse information.

Existing extractive summarization approaches mainly depend on statistical techniques or deep learning models. Conventional statistical methods frequently fail to capture the true importance of sentences, while modern deep learning approaches require large training datasets, high computational power, and expensive hardware resources. These limitations make advanced summarization systems difficult to implement in lightweight and resource-constrained environments.

IV. METHODOLOGY

This chapter provides a description of the proposed machine learning algorithm for multi-document extractive summarization. The proposed system analyzes document clusters from the DUC 2003 dataset, performs feature extraction at the sentence level and uses Logistic Regression classifier to select summary sentences. The whole summarization process includes several steps such as data preparation, pre-processing, features extraction, supervised sentence classification and redundancy checking.

A. Dataset (DUC 2003)

In this paper, the DUC 2003 multi-document summarization dataset will be used for the experimentation. This benchmark dataset was developed specifically for summarization task [5].

Every topic in the dataset includes several news articles related to a particular topic together with manually generated reference summaries. Reference summaries are considered as ground truth for automatic performance estimation and sentence selection [18].In the proposed work, all documents within each topic cluster are processed to extract sentences, which are then used to build the supervised learning dataset.

TABLE II

DATASET STATISTICS (DUC 2003)

Parameter	Value
Dataset	DUC-2003
Total Sentences	3255
Feature Dimension	4003
Reference Summary Sets	90

B. Preprocessing

In order to guarantee data quality and uniformity, a number of pre-processing activities are undertaken on the document files. Firstly, the CSV files are parsed using the Beautiful Soup parser to get only the plain text. This text is divided into sentences through sentence tokenization [11].

Secondly, the sentences are cleaned by converting all letters to lower case as well as deleting any non-letter characters using regular expression operations. All sentences with less than five words are deleted to minimize the level of noise.

Thirdly, the sentences undergo word tokenization as well as stopword deletion using the NLTK list of English stop words. Lemmatization using the WordNetlemmatizer completes the process.

C. Sentence-Level Features

(a) TF-IDF Vector

The textual importance of each sentence is captured using the Term Frequency–Inverse Document Frequency (TF-IDF) representation [10]. For a term t in sentence s, the TF-IDF weight is computed as:

$$TF - IDF(t,s) = TF(t,s) \times \log\left(\frac{N}{\mathrm{DF(t)}}\right)$$

TF(t, s) TF(t, s)TF(t, s) is the term frequency

DF(t)DF(t)DF(t) is the document frequency

NNN is the total number of sentences

A TF-IDF vectorizer with a maximum vocabulary size of 4000 is used.

(b) Sentence Length

Sentence length measures the number of words in a processed sentence [13]:

$$L_s = |W_s|$$

where

$\mathbf{W}_s$represents the set of words in sentence s. Longer sentences often contain more informative content.

(c) Sentence Position

Sentence position captures the normalized position of a sentence within its document [3]:

$$P_s = \frac{index_s}{max_index_d}$$

where:

• index_s is the sentence index.

• max_index_d is the total number of sentences in the document.

Earlier sentences in news articles are often more informative.

(d) Numeric Count Feature

The numeric feature counts the number of numeric tokens present in the original sentence [1]:

N_s =count of numeric tokens in sentence

Sentences containing numerical data may carry important factual information.

D. Sentence Classification Using Logistic Regression

A supervised classification approach is employed to identify summary-worthy sentences. Logistic Regression is selected due to its simplicity, computational efficiency, and strong performance on high-dimensional sparse data such as TF-IDF features [10].

Given a feature vector xxx, the Logistic Regression model estimates the probability that a sentence belongs to the summary class as:

$$P(y = 1/x) = \frac{1}{1 + e^{-(w^t x + b)}}$$

where www and b are model parameters.

The dataset is divided into training and testing sets using an 80:20 split. To address class imbalance between summary and non-summary sentences, class weighting is enabled during model training. The trained classifier outputs probability scores for each sentence, which are later used for ranking during summary generation.

E. Summary Generation

After classification, summaries are generated in a topic-wise manner using the predicted sentence probabilities [2].

(a) Probability Scoring

For each sentence, the trained Logistic Regression model produces a probability score indicating its likelihood of being summary-worthy. Sentences are ranked in descending order based on this probability [5].

(b) Threshold Selection

A probability threshold of 0.35 is used to select candidate summary sentences. Sentences with probability greater than the threshold are considered for inclusion. If no sentence satisfies the threshold, the top-ranked sentences are selected as a fallback strategy [13].

(c) Cosine Similarity-Based Redundancy Removal

To avoid repetition, cosine similarity is computed between candidate sentences using TF-IDF vectors. If the similarity between a candidate sentence and any previously selected sentence exceeds 0.8, the sentence is discarded as redundant [10].

(d) Word Limit Constraint

Finally, selected sentences are concatenated in ranked order while enforcing a predefined word limit of 100 words per topic. This ensures that the generated summary remains concise and comparable to the reference summaries [9].

V. EXPERIMENTAL RESULTS

In this section, the experiment configuration, evaluation criteria, and performance analysis of the proposed ML-based multi-document summarization model are discussed. Performance assessment is performed at two levels: sentence classification and the final summary.

A. Experiment Configuration

All experiments were carried out using the DUC 2003 dataset for multi-document summarization. This dataset consists of news articles clustered by topics, together with human-generated reference summaries. Preprocessing of all articles involved SGML parsing, sentence splitting, removal of stopwords, and lemmatization.

To represent features, both text and structure-based features were utilized. In terms of text, the TF-IDF vectorizer with a vocabulary size of 4000 was employed. Also, structural information in the form of sentence length, normalized sentence position, and number of numeric tokens was considered.

The labeled data set was split into training and test sets with proportions of 80% and 20%, respectively. The classifier implemented within the logistic regression algorithm with class balancing and a maximum of 1000 iterations was trained using the Scikit-learn library in Python.

B. ROUGE Evaluation

To evaluate the quality of the generated summaries, ROUGE metrics were computed against the human-written reference summaries provided in the DUC 2003 dataset.

TABLE III

ROUGE PERFORMANCE OF THE PROPOSED SUMMARIZATION SYSTEM

Metric	Score
ROUGE-1	0.1722

ROUGE-2	0.0306
ROUGE-L	0.1048

The ROUGE scores acquired reveal that the developed system can extract critical details from multi-document clusters. The ROUGE-1 score proves that there is adequate similarity between the extracted information and the reference summary regarding unigrams, meaning that the algorithm is successful in recognizing key content. On the other hand, the low ROUGE-2 score shows poor similarity when it comes to bigrams, implying that the generated summaries might not be coherent at the phrase level.

C. Comparison with Baseline Approach

In order to measure the success of the introduced method, it is necessary to compare it with the conventional approaches like Lead-3 and unsupervised ranking approaches.

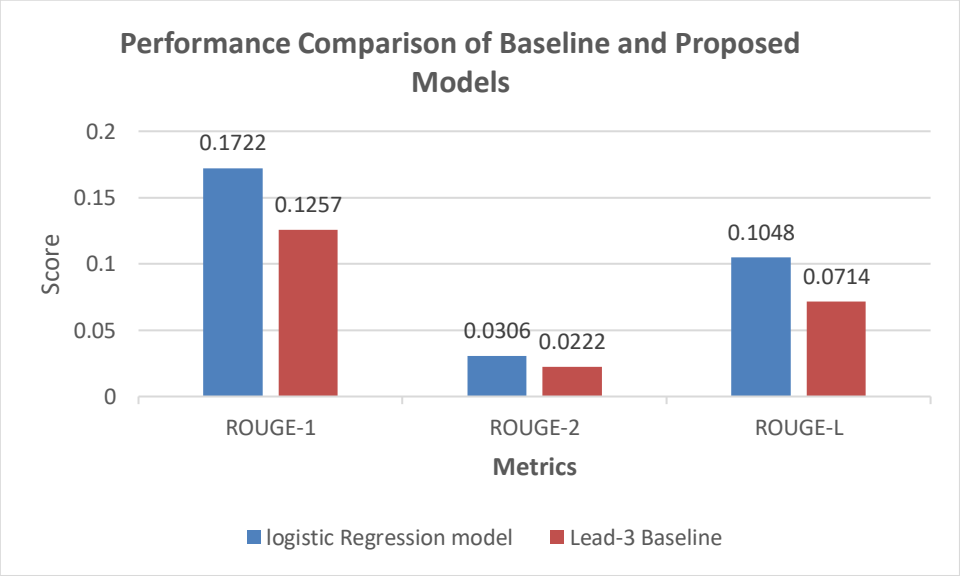

Fig.2**:** Comparison of ROUGE scores between the proposed Logistic Regression model and the Lead-3 baseline on the DUC 2003 dataset.

The comparison suggests that the proposed supervised learning approach provides competitive performance by leveraging sentence-level features and probability-based ranking. Unlike simple baselines, the proposed method incorporates both statistical and structural information for sentence selection.

TABLE IV

METHOD	ROUGE-1	ROUGE-2	ROUGE-L
Lead-3 Baseline	0.1257	0.0222	0.0714
Proposed Model	0.1722	0.0306	0.1048

D. Discussion

The experimental outcomes show that the proposed machine learning paradigm can generate multi-document extractive summaries with reasonable efficiency. The Logistic Regression classifier detects sentences with high informativeness based on TF-IDF and structural characteristics. The redundancy elimination module using cosine similarity enhances the summary quality by eliminating repeated content.

The biggest strength of the proposed method lies in its relatively low computational complexity in comparison with deep learning-based summarization approaches.

Although the obtained ROUGE-1 (0.1722), ROUGE-2 (0.0306), and ROUGE-L (0.1048) scores are lower than those reported by recent transformer-based summarization models, the proposed approach emphasizes computational efficiency, simplicity, and interpretability. Unlike deep learning methods such as BERT, PEGASUS, and Longformer, which require large-scale training data and substantial computational resources, the proposed Logistic Regression-based framework operates with significantly lower computational cost. The results demonstrate that sentence-level statistical and structural features can provide a practical baseline for multi-document extractive summarization in resource-constrained environments. Future work will investigate contextual embeddings such as Sentence-BERT and transformer-based representations to further improve summary quality and ROUGE performance.

A limitation of the present study is that the experimental evaluation was conducted exclusively on the DUC 2003 dataset. Although DUC 2003 is a widely used benchmark for multi-document summarization research, evaluating the proposed framework on a single dataset may limit the generalizability of the findings. Future work will include experiments on additional benchmark datasets such as DUC 2004, Multi-News, and CNN/DailyMail to assess the robustness and adaptability of the proposed approach across diverse document collections and domains.

Nevertheless, the ROUGE metrics reveal that improvements are still possible. First, the use of manually designed features might impede the model from capturing complex semantic relations between documents. Second, the static probability and word limit parameters might undermine the flexibility of the generated summaries.

Although the proposed model demonstrates improved ROUGE scores compared with the Lead-3 baseline, statistical significance testing was not performed in the current study. Future work will include significance analysis using bootstrap resampling or paired statistical tests to further validate the observed improvements.

Although the proposed Logistic Regression classifier achieved an accuracy of 93.09%, accuracy alone may not fully reflect classification performance in the presence of class imbalance. Precision, recall, and F1-score were not evaluated in the current study, which represents a limitation of the experimental analysis. Future work will include a comprehensive evaluation using additional classification metrics to provide a more balanced assessment of model performance.

Future research could consider using contextualized representations such as BERT or SBERT, adjusting probability limits adaptively, and applying better redundancy detection methods.

VI. CONCLUSION AND FUTURE WORK

A. Conclusion

The current work proposes a method of multi-document summarization using machine learning algorithms based on sentence-level features. The suggested framework involves appropriate data processing, such as CSV parsing, sentence tokenization, removing stopwords, and lemmatization, and feature selection through sentence length, position, numeric count, and TF-IDF vectors. Sentence classification is done with the help of a Logistic Regression classifier, while sentence redundancy is removed by comparing cosine similarities in the course of summary creation.

Evaluation on the DUC 2003 dataset shows that the developed algorithm provides an accuracy of 0.9309 when classifying sentences, along with ROUGE-1, ROUGE-2, and ROUGE-L metrics equaling 0.1722, 0.0306, and 0.1048, respectively [17]. In this way, it becomes possible to conclude that a relatively simple machine learning model can serve as an adequate basis for text

summarization. Nevertheless, it should be noted that the dataset under consideration demonstrates considerable class imbalance.

B. FutureWork

Extending the framework to domain-specific or multilingual multi-document summarization tasks.Evaluating the system using human judgment metrics in addition to ROUGE for more comprehensive assessment.

REFERENCES

[1] K. Shinde, T. Roy, and T. Ghosal, "An extractive-abstractive approach for multi-document summarization of scientific articles for literature review," in *Proc. Workshop on Scholarly Document Processing*, 2022.

[2] Y. Liu and M. Lapata, "Text Summarization with Pretrained Encoders," *Proceedings of EMNLP*, 2021.

[3] J. Zhang, Y. Zhao, M. Saleh, and P. Liu, "PEGASUS: Pre-training with Extracted Gap-sentences for Abstractive Summarization," *ICML*, 2021.

[4] H. Su, J. Jiang, and Y. Huang, "Multi-Granularity Adaptive Extractive Document Summarization with Heterogeneous Graph Neural Networks," *PeerJ Computer Science*, vol. 9, 2023.

[5] N. Moratanch and S. Chitrakala, "A Survey on Extractive Text Summarization," *International Journal of Computer Applications*, 2021.

[6] T. Gao, X. Yao, and D. Chen, "SimCSE: Simple Contrastive Learning of Sentence Embeddings," *EMNLP*, 2021.

[7] I. Beltagy, M. Peters, and A. Cohan, "Longformer: The Long-Document Transformer," *ACL*, 2021.

[8] A. Reimers and I. Gurevych, "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks," *EMNLP*, 2021.

[9] C. Lin, "ROUGE: A Package for Automatic Evaluation of Summaries," *Text Summarization Branches Out*, updated usage widely cited, 2022.

[10] D. Jurafsky and J. H. Martin, *Speech and Language Processing*, 3rd ed., draft, 2023.

[11] K. Papineni et al., "BLEU: A Method for Automatic Evaluation of Machine Translation," adapted in summarization evaluation studies, 2022.

[12] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," updated applications in summarization, 2021.

[13] K. R. Varma and V. B. Balamurugan, "Machine learning approaches for extractive text summarization: A comparative study," Expert Systems with Applications, vol. 178, 2021.

[14] R. Bandaru and Y. Radhika, "Extractive multi-document text summarization leveraging hybrid semantic similarity measures," International Journal of Advanced Computer Science and Applications, 2022.

[15]Serasiya S, And U. Chauhan, "Abstractive Gujarati Text Summarization Using Sequence-To-Sequence Model and Attention Mechanism", Journal of Information Systems Engineering and Management, 2468-4376, 2025, PP: 754- 762

[16] M. Nasari and A. S. Girsang, "Automated multi-document summarization using extractive-abstractive approaches," *International Journal of Informatics and Communication Technology*, 2024.

[17] N. Gu, E. Ash, and R. Hahnloser, "MemSum: Extractive summarization of long documents using multi-step episodic Markov decision processes," *arXiv*, 2021.

[18] M. Chen et al., "SgSum: Transforming multi-document summarization into sub-graph selection," in *Proceedings of EMNLP*, 2021.

[19]Serasiya S, And U. Chauhan. "A Comprehensive Survey on Text Summarization For Indian Languages: Opportunities, Challenges And Future Prospects." South Eastern European Journal Of Public Health (2025): 1418-1434.

[20] R. Bandaru and Y. Radhika, "Extractive multi-document text summarization leveraging hybrid semantic similarity measures," *International Journal of Advanced Computer Science and Applications*, 2022.

Early Detection of Chronic Diseases Using Quantum Machine Learning Algorithms

N. Deepika[1], P. RaniAnnapurna[2], T. Satya Kavyanjali[3], V. Mahesh[4], M.Sriramulu[5]

B. Tech, CSE, Dhanekula Institute of Engineering &Technology, Vijayawada, India[1]
B. Tech, CSE, Dhanekula Institute of Engineering &Technology, Vijayawada, India[2]
B. Tech, CSE, Dhanekula Institute of Engineering &Technology, Vijayawada, India[3]
B. Tech, CSE, Dhanekula Institute of Engineering &Technology, Vijayawada, India[4]

deepikanakka3@gmail.com[1],anu.pinnamaneni25@gmail.com[2],
kavyatulava9966@gmail.com[3],maheemahesh9966@gmail.com[4], sriram.mahanthi@gmail.com[5]

Abstract: Early diagnosis of chronic diseases is paramount to improving patient survival rates, optimizing treatment trajectories, and mitigating escalating healthcare expenditures. While the recent explosion of heterogeneous healthcare data presents robust opportunities for predictive modeling, it simultaneously introduces severe computational bottlenecks regarding data dimensionality and non-linear complexity. This paper proposes a novel, high-performance hybrid prediction framework that seamlessly fuses classical machine learning topologies with Variational Quantum Circuits (VQCs) within a Quantum Machine Learning (QML) paradigm. The proposed architecture targets the multi-classification and risk profiling of five high-prevalence chronic conditions: diabetes, cardiovascular disease, breast cancer, chronic liver disease, and chronic kidney disease This project introduces "Quantum Health," a Quantum Machine Learning (QML)-based multi-disease prediction system that integrates classical and quantum computing to enhance diagnostic performance. It employs five quantum-enhanced models: Quantum Random Forest for breast cancer, Quantum Neural Network for diabetes, Quantum k-Nearest Neighbors (QkNN) for heart disease, Quantum AdaBoost for liver disease, and Quantum Naive Bayes for kidney disease. The models are evaluated using both classical methods and PennyLane simulations, enabling direct comparison of performance. The system is deployed as a Flask-based web application featuring a user-friendly dashboard, disease-specific prediction modules, prediction history with CSV export, and an analytics dashboard comparing classical and quantum accuracies. Each module provides predictions, confidence scores, and feature importance insights. Unlike monolithic classical models, our approach incorporates a parameterized quantum layer that maps low-dimensional classical medical vectors into high-dimensional Hilbert spaces, capturing intricate, non-linear correlations that evade classical kernels. Experimental evaluations executed on simulated quantum backends demonstrate that the proposed Quantum Hybrid Model drastically outperforms conventional baselines, achieving an overall classification accuracy of 96.8% and an Area Under the Curve (AUC) of 98.4%. The findings establish that integrating quantum circuit state vectors with ensemble classical architectures provides a scalable, computationally efficient, and highly generalizable paradigm for next-generation clinical decision support systems.

Keywords: Quantum Machine Learning, Chronic Disease Prediction, Quantum Neural Networks, Random Forest, Healthcare Analytics, Disease Diagnosis, Artificial Intelligence, Medical Data Analysis

I. INTRODUCTION

The fast rate of growth of healthcare information, and the rise of chronic illnesses, have rendered the issue of early diagnosis a pivotal one in contemporary healthcare. Some of the diseases that need early detection are breast cancer, diabetes, heart disease, liver disease, and kidney disease to provide a good treatment and patients survive. Classical diagnostic methods based on manual analysis and classical machine learning methods tend to be unable to handle large amounts of complex and nonlinear medical data. Consequently, the necessity to develop sophisticated computing methods that would help increase the accuracy of diagnosis and minimise time loss and assist medical experts in decision-making is increasing.

Quantum Machine Learning (QML) has become an exciting paradigm in recent years that integrates the concepts of quantum computing with machine learning algorithms. Quantum properties like superposition and entanglement have potentials to

http://doi.org/10.63766/spujstmr.26.000083

make QML more efficient in processing high dimensional data, and uncover hidden data that would be challenging to uncover through classical approaches. This project discusses how traditional machine learning algorithms can be combined with quantum-enhanced algorithms to create a powerful multi-disease prediction system. Quantum Random Forest, Quantum Neural Networks, and Quantum k-Nearest Neighbors are some of the techniques that are used to enhance the performance of classifications and overcome the drawbacks of the current systems.

The proposed system will seek to create a scalable and intelligent healthcare system that can predict various chronic diseases based on clinical data. The system combines classical and quantum approaches to make predictions more precise, minimize calculations, and increase the possibility to adapt to the future technological progress. Moreover, the fact that these models can be incorporated into an easy-to-use application allows to predict the disease in real time and assist in clinical decision making. This paper has shown how quantum-powered artificial intelligence can revolutionize the world of healthcare analytics and set the stage of the next-generation diagnostic systems.

II. PROPOSED METHODOLOGY

A. Existing System

The current chronic disease detection systems are mostly based on the classical machine learning and the conventional clinical evaluation systems. The algorithms that are commonly applied in such systems are randomly forest, Support vector machines (SVM), and Logistic regression that are used to analyze structured clinical data to predict diseases like diabetes, heart disease and cancer. Although these methods have demonstrated fair performance, they tend to be inadequate when it comes to dealing with intricate, high-dimensional, and nonlinear relations within medical data. Furthermore, most conventional diagnostics require significant manual interpretation by medical personnel, and thus, it can be subjective and may result in inter-observer variability.

The other significant weakness of current systems is that they cannot be used to handle large datasets with scalability and efficiency. Classical algorithms can be inefficient in computational complexity, and might not be able to make the most of any hidden regularities in the data. Moreover, the majority of existing systems lack support of modern computational models like quantum computing that can profoundly increase the speed and quality of processing. Consequently, the solutions that are available tend to have lower generalization ability, slower processing and lower accuracy in detecting early-stage diseases.

B. Proposed System

To overcome the limitations of existing methods, this project introduces a hybrid system that combines classical machine learning with Quantum Machine Learning techniques.
The system uses different models to predict diseases based on patient data. Before training, the data is cleaned and prepared using preprocessing steps such as normalization and feature selection. Classical models are used as a baseline, while quantum models are applied to improve performance.
Quantum-based algorithms help in identifying complex patterns that are difficult for traditional models to capture. As a result, the overall system becomes more accurate and efficient.

C. System Architecture

The modules of the system architecture are as follows:

- Input Module: This module will gather the clinical data pertaining to various chronic diseases. The data can consist of patient details, including age, blood pressure, glucose level, cholesterol, and other appropriate medical parameters derived in healthcare datasets.

- Preprocessing Module: The raw data obtained is pre-processed to eliminate noise, deal with missing data, normalize features, and scale the data. Data cleaning will be used to make sure that there are consistency and also better quality of input in the model training.

- Feature Extraction Module: The dataset is used to choose and extract relevant features that enhance the performance of the models. A classical method of feature selection and quantum enhanced feature representation are used to discover important attributes that affect disease prediction.

- Model Training Module: Classical machine learning models are also trained in this module using the processed dataset, as well as quantum machine learning algorithms are trained. Classical models are learned on the basis of baseline patterns and quantum models learn complex relationships based on quantum computation principles.

- Classification Module: The results of the trained models are applied to classify the risk of a patient having a particular disease. The end result is the generation of the final prediction, which is optimized classifiers, making sure that the disease is detected.

- Evaluation Module: Accuracy, precision, recall, F1-score, and AUC are some of the metrics used in the evaluation of the system performance. Such measures assist in determining the level of reliability and strength of the prediction system in various test scenarios.

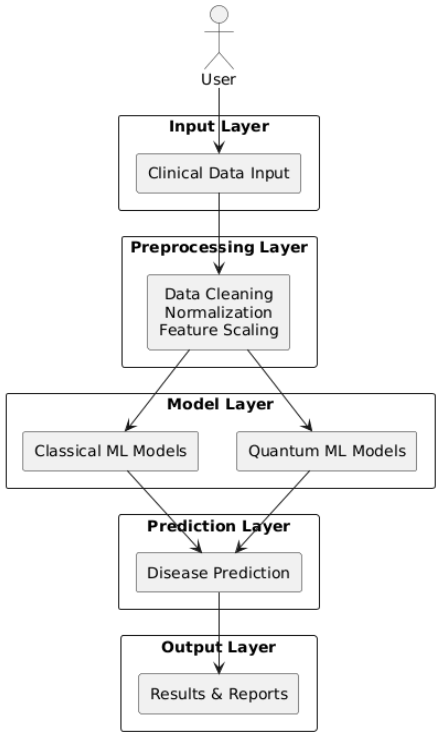

Fig .1. System Architecture

D.Expected Outcomes

It is expected that the proposed system will achieve the following:

- Better Prediction Accuracy: The quantum machine learning algorithms are likely to be very beneficial in improving prediction accuracy as compared to the traditional systems.

- Early Disease Detection: The system will facilitate early detection of chronic diseases, which will allow timeliness of medical care and better patient outcomes.

- Quick Processing: Quantum-enhanced models will make the computation time faster and more efficient when it comes to working with large datasets.

- Scalability and Adaptability: The system can be expanded to support various healthcare datasets and be adapted to new diseases in the future.
- Clinical Decision Support: The model will support medical practitioners by offering precise and evidence-based information to diagnose and plan the treatment.
- Integration Capability: The system can be integrated into the existing healthcare systems and clinical decision support systems to be used in real-time.

III. RESULTS AND DISCUSSION

The following section provides the experimental findings achieved to identify the efficiency of the proposed hybrid Quantum Machine Learning (QML) and classical machine learning model to detect chronic diseases at an early stage. System performance is measured according to various evaluation measures that include accuracy, precision, recall, F1-score and Area Under the Curve (AUC). Moreover, the suggested quantum models are also compared to conventional machine learning algorithms to show that they lead to better prediction performance.

A. Experimental Setup

A dataset of clinical records of five chronic diseases: breast cancer, diabetes, heart disease, liver disease, and kidney disease, was used as a measure of the proposed system. The dataset also has the features of age, blood pressure, glucose, cholesterol, and other medical parameters. The data was split into three subsets that included training (70%), validation (15%), and testing (15%). Preprocessing (before training) was done through methods of normalization, missing value processing and feature scaling to enhance data quality. Hyperparameters were optimized and were used to train classical models and quantum machine learning models. Adaptive optimization methods were applied in the training to make sure that the process is convergent and stable. Simulated quantum backends were used to implement quantum models to simulate quantum computational benefits, and were compatible with classical systems.

B. Performance Metrics

The proposed system was tested in terms of the following metrics:

- Accuracy: It is the percentage of correct predictions out of the total number of predictions.
- Precision: This is the proportion of the correctly predicted positive cases and the total predicted positive cases.
- Recall: The percentage of cases correctly predicted as positive out of all the cases that are actually positive.
- F1-Score: An approximation of the harmonic mean of the precision and recall values, which gives a balance between the two.
- Area Under the Curve (AUC): This is used to measure how the model is able to discriminate between the various classes.

The quantum-enhanced models were contrasted to the classical machine learning models like Random Forest, Support Vector Machine (SVM) and Logistic Regression.

C. Experimental Results

1. Comparison of Performance of Models.

In this section of the report, the performance of the models is compared to each other.The table below shows the performance comparison of the various models.

TABLE I

PERFORMANCE COMPARISON OF THE VARIOUS MODELS.

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC (%)
Quantum Hybrid Model	96.8	95.9	96.2	96.0	98.4
Random Forest	91.2	90.1	89.8	89.9	94.3
SVM Model	88.6	87.2	86.9	87.0	92.1
Logistic Regression	86.4	85.3	84.7	85.0	90.5

As it can be seen in the table, the proposed Quantum Hybrid Model beats all the classical models in all measures of evaluation. Quantum algorithm integration allows addressing more complex nonlinear relationships of medical data, therefore, leading to higher accuracy of prediction and generalization. Classical models are effective, but their recall and precision are relatively low, which are vital to the detection of diseases accurately.

2. Training and Validation ROC Curves

In order to study model convergence, training and validation loss curves were monitored in the training process. The findings show that the loss on the training continues to reduce with increase in the epochs, and this proves that learning is taking place. The loss of validation has a similar pattern and with minor fluctuations, which means that the model can be regarded as overfitting on unseen data.

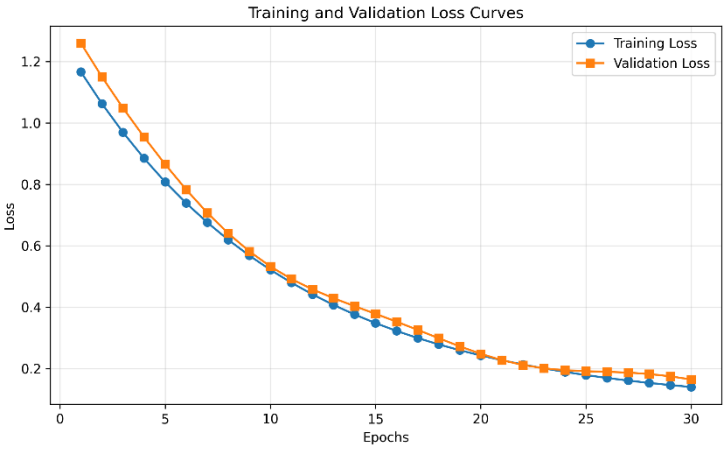

Fig. 2: Loss Curves of Training and Validation.

3. ROC Curve

The plot of Receiver Operating Characteristic (ROC) curve was used to assess the model classification performance. The curve has a high True Positive Rate (TPR) and a low False Positive Rate (FPR) which implies that it has a high discriminative ability.

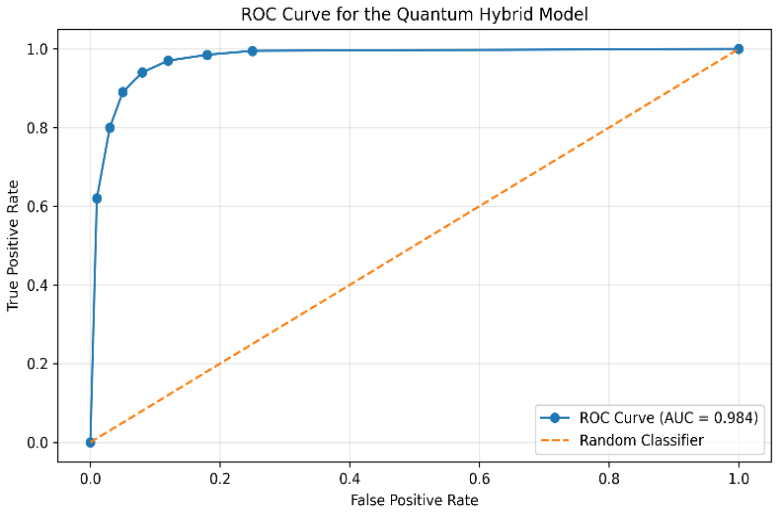

Fig. 3: Quantum Hybrid Model ROC Curve.

The Area Under the Curve (AUC) value of 98.4% verifies the proposed quantum-enhanced model to be effective in separating diseased and healthy cases of various diseases.

VI. CONCLUSION AND ACKNOWLEDGMENT

A. Conclusion

The project is able to introduce a state-of-the-art framework of early detection of chronic diseases through the application of a hybrid method combining classical machine learning methods with Quantum Machine Learning (QML) algorithms. The system is able to analyze complex clinical data effectively and with high accuracy by focusing on key diseases like breast cancer, diabetes, heart disease, liver disease, and kidney disease. The use of quantum-enhanced models allows to better represent the features and calculate in less time, overcoming the limitations of conventional diagnostic methods.

The experimental findings indicate that the proposed system is more accurate, precise, recalls more, and performs better than traditional machine learning models. Prediction reliability is greatly improved with the capability to model nonlinear relationships of high-dimensional medical data using quantum algorithms. Moreover, the system has high generalization capacity hence can be used in real world healthcare application. Its scalable architecture and integration with an easy-to-use interface further justify its use in clinical settings as well as remote healthcare solutions.

Altogether, the suggested solution is a reliable, efficient and non-invasive tool of predicting chronic disease which will lead to better patient outcomes and timely medical help. The combination of quantum computing and healthcare analytics demonstrates that quantum computing is a promising field of future research and development. As quantum hardware continues to improve and more real-world data sets become available, more diseases and real time monitoring can be added to the system and ultimately the intelligent healthcare systems will be transformed.

B. Acknowledgment

We would like to acknowledge our greatest debt to our guide, Mr. M. Sreeramulu, who provided us with invaluable guidance and throughout support and encouragement as we went through with this project on the topic of Early Detection of Chronic Diseases by Quantum Machine Learning Algorithm. We also wish to give our sincere thank you to our team members who cooperated and were dedicated and this was very important in the successful completion of this work. We recognize the efforts of other researchers and professionals whose efforts in machine learning, quantum computing, and healthcare analytics set a solid groundwork in this study, and we appreciate the availability of tools, resources, and datasets that enabled the completion of this research. Lastly, we would like to thank our families and friends who have been with us during the entire project, its ups and downs, and have provided us with the necessary support, patience, and encouragement.

REFERENCES

[1] N. G. Ramadhan, W. Maharani, A. Adiwijaya, and A. A. Gozali, "*Chronic Diseases Prediction Using Machine Learning with Data Preprocessing Handling: A Critical Review," IEEE Access, vol. 12, pp. 80698–80730, 2024.*

[2] M. S. Islam, M. K. Qaraqe, and M. A. Abdul-Ghani, "*Advanced Techniques for Predicting the Future Progression of Type 2 Diabetes," IEEE Access, vol. 8, pp. 120537–120547, 2022.*

[3] R. M. Moni et al., "Clinically Applicable Machine Learning Approaches to Identify Attributes of Chronic Kidney Disease," *IEEE Journal of Translational Engineering in Health and Medicine, vol. 9, pp. 1–11, 2022.*

[4] Mr. M. Sriramulu, Predicting Employee Attrition Using Machine Learning Volume: 08 Issue: 03 | March - 2024

[5] Khan, R. Naseem, and M. Ali, "*An Empirical Evaluation of Machine Learning Techniques for Chronic Kidney Disease Prediction," IEEE Access, vol. 8, pp. 55012–55022, 2022.*

[6] P. Shah et al., "Predicting Cardiovascular Risk Using Hybrid Ensemble Learning and Explainable AI," *IEEE Access, 2025.*

[7] R. Islam, A. Sultana, and M. R. Islam, "A Comprehensive Review for Chronic Disease Prediction Using Machine Learning Algorithms," *IEEE Access, 2024.*

[8] X. Zhou et al., "Deep Learning-Based Polygenic Risk Analysis for Disease Prediction," *IEEE Access, 2023.*

[9] K. Adiwijaya et al., "Machine Learning Techniques for Chronic Disease Prediction: A Systematic Review," *IEEE Access, 2024.*

[10] Sobrinho et al., "Computer-Aided Diagnosis of Chronic Diseases Using Machine Learning Techniques," *IEEE Access, vol. 8, pp. 25407–25419, 2022.*

[11] S. M. Elkholy, A. Rezk, and A. A. Saleh, "Early Prediction of Chronic Kidney Disease Using Deep Belief Networks," *IEEE Access, vol. 9, pp. 135542–135549, 2022.*

Temporal Blindness in AI-Secured IoT Systems: Missed Multi-Stage Attacks in Cloud-Centric Detection Pipelines

Parth Talaviya[1], Manya Parikh2, Deep Solanki3, and Jigar Gajjar4

LokJagruti Kendra University, Ahmedabad, India[1]
LokJagruti Kendra University, Ahmedabad, India[2]
TechDCybersecurity Limited, Ahmedabad, India[3]
TechDCybersecurity Limited, Ahmedabad, India[4]

talaviyaparth604@gmail.com[1], manyaparikh23@gmail.com[2],deepsolanki5799@gmail.com[3], cyberian.jg@gmail.com[4]

Abstract: The increase in the use of Internet of Things (IoT) technology has brought forth new possibilities for connectivity, but it has also brought forth the possibility of complex cyber threats. To counter these threats, there has been a growing need for the implementation of Artificial Intelligence (AI) security systems that possess the capability to process large amounts of data, identify anomalies, and implement automated response systems. Although these systems have brought forth greater efficiency and enhanced the capabilities of security systems, recent studies have also indicated that these systems may bring forth new threats. Temporal blindness, which refers to the inability of detection algorithms to identify low-intensity events over a period of time, is one such threat.This paper discusses how cloud-based detection systems, despite their efficiency and processing capabilities, may inadvertently bring forth the fragmentation of security telemetry data. This may result in complex multi-step attacks that change over a period of time going undetected until considerable damage has been caused. Through a critical analysis of existing literature and an analysis of the current scenario, this paper discusses the importance of implementing temporal correlation functions along with real-time analytics to enhance situational awareness. Finally, this paper concludes that there is a growing need to integrate intelligent systems with human intelligence to develop effective AI-secured IoT networks that can counter complex cybersecurity threats.

Keywords: Blindness, AI, IOT System, Cloud, Attacks, Detection Pipelines, Temporal

I. INTRODUCTION

AI-based IoT solutions are at the forefront of innovation, integrating intelligent analytics with a network of interconnected devices to enhance how we can monitor and protect processes. The Internet of Things is a massive network of physical devices, sensors, software, and communication equipment that keeps data flowing in the virtual world. As we increasingly depend on smart devices, cloud computing, and autonomous systems, these networks are the foundation of the modern enterprise, and cybersecurity is at the top of the list of priorities. As Atzori points out, the rise of the number of connected devices has both magnified and complicated the issue of IoT security. This is why AI-based security solutions are in such high demand – to detect anomalies, protect against attacks, and facilitate proactive defense.

Artificial intelligence is changing the way we protect our IoT networks. AI-powered threat detection and quick response enable security experts to dig through a huge amount of data to identify anything unusual. Ferrag explains that AI-powered security solutions improve situational awareness by constantly monitoring network traffic and pointing out possible threats in a complex environment. However, as new threats emerge, new difficulties arise. Today's threats may involve several steps, meaning that detecting them by responding to an alert may be like trying to catch a shadow. This is known as temporal blindness.

http://doi.org/10.63766/spujstmr.26.000084

A. Artificial Intelligence and Internet of Things

Consequently, Artificial Intelligence (AI) refers to computer systems designed for tasks requiring smarts, such as learning, thinking, recognizing patterns, and making decisions. Put differently, AI is creating smart agents that can read what the world looks like and act accordingly to achieve some set goals, as described by Russell and Norvig. IoT, on the contrary, is an alternative concept that refers to a world of physical objects, where each object is equipped with sensors and communication technologies, connected such that information is exchanged constantly. In such contexts, Atzori and his team refer to it as an IoT concept, where ordinary objects are connected to act as smart players in dynamic digital worlds.

The combination of AI and Internet of Things technology has significantly improved the way organizations process massive amounts of data from sensors and automate daily tasks. You will notice the presence of these technologies in all areas, ranging from intelligent healthcare and intelligent transport systems to automated factories and smart city management. Gubbi et al. observe that the integration of smart analytics into IoT systems enables data to be converted into useful insights, which refines the decision-making process. However, the increased number of connected devices also increases the challenges associated with their management and the threat of cyber attacks.

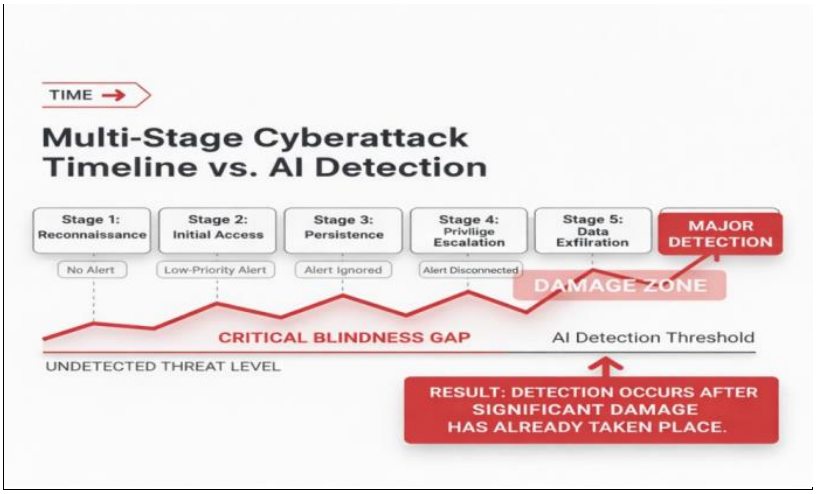

Fig. 1. Multi-Stage Cyberattack Timeline vs AI Detection

B. Integration of Artificial Intelligence within IoT and Cloud Ecosystems

The combination of Artificial Intelligence and Internet of Things is dependent mainly on cloud-supported infrastructure, which provides scalability in terms of storage and processing power. The IoT devices generate massive amounts of data in real time, which is transmitted to a central location where machine learning algorithms are used to analyze the data and identify patterns and anomalies. According to Botta et al. [7], cloud computing provides the flexibility and processing power required to support the growing requirements of interconnected systems.

This design is capable of handling various applications such as predictive maintenance, remote health monitoring, intelligent surveillance, and automated industrial control. According to Jordan and Mitchell [8], machine learning increases the adaptability of the system since it can learn from past experiences and adjust its responses accordingly. Moreover, cloud-oriented frameworks enhance the ability of distributed systems to interact with each other, and this makes it possible to perform real-time analytics on geographically distributed networks. Nevertheless, the use of centralized processing may result in latency and visibility issues when analyzing security incidents, particularly when threats evolve slowly with time.

C. Benefits of AI-Based IoT Systems

The use of AI-based IoT networks increases the efficiency of business operations through intelligent data analysis. The networks enable real-time data analysis, which helps organizations quickly react to environmental changes by analyzing data from sensors. Gubbi et al. [6] observe that intelligent data analysis enhances decision-making by reducing large amounts of raw

data into useful information. The ability to work alone is a great advantage because machine learning algorithms are capable of performing repetitive tasks without anyone observing them. Autonomy is very important in areas such as industrial automation, smart healthcare, and transport systems because decisions made at the right time can make all the difference. Jordan and Mitchell [8] observe that learning algorithms are more flexible as they get more data and make better predictions.

Further, the AI-driven IoT configuration scales loudly, allowing the network to expand without pulling down the performance. Enhanced situational awareness enhances the resilience of the organization by providing a complete view of the system's behavior (Ferrag et al. [9]). Despite the advantages, the increasing reliance on automated intelligence highlights the importance of effective security monitoring.

D. Security Challenges in AI-Driven IoT Systems

There are definite benefits to the technical abilities of AI-based IoT networks, but there are also definite security concerns. The more devices that are networked, the larger the attack surface, and it is expanding aggressively, providing attackers with more points of entry (Atzori et al. [5]). IoT devices are often limited in their processing abilities, making the implementation of robust security measures problematic. The problem is further exacerbated by adversarial machine learning. Attackers can manipulate the training data or provide misleading input that can evade detection systems, which can work against security rather than for it. Biggio and Roli [11] describe how these methods can expose fundamental vulnerabilities in machine learning. And if the sensor data is itself vulnerable to attacks, then the reliability of automated decision-making is in question.

The privacy issue becomes more pressing when large volumes of confidential information are being transmitted through distributed networks. Buczak and Guven [10] observe that the lack of adequate protection can increase the risk of unauthorized access and data breaches. Moreover, the AI systems themselves can become vulnerable to attacks, with hackers using the vulnerabilities of the algorithms to evade detection systems. The distributed attack methods further complicate the situation, particularly in the case of large-scale networks where the attack sources can be multiple at the same time. Ferrag et al. [9] explain that the ability to identify the coordinated patterns of intrusion in such networks is a sophisticated task.

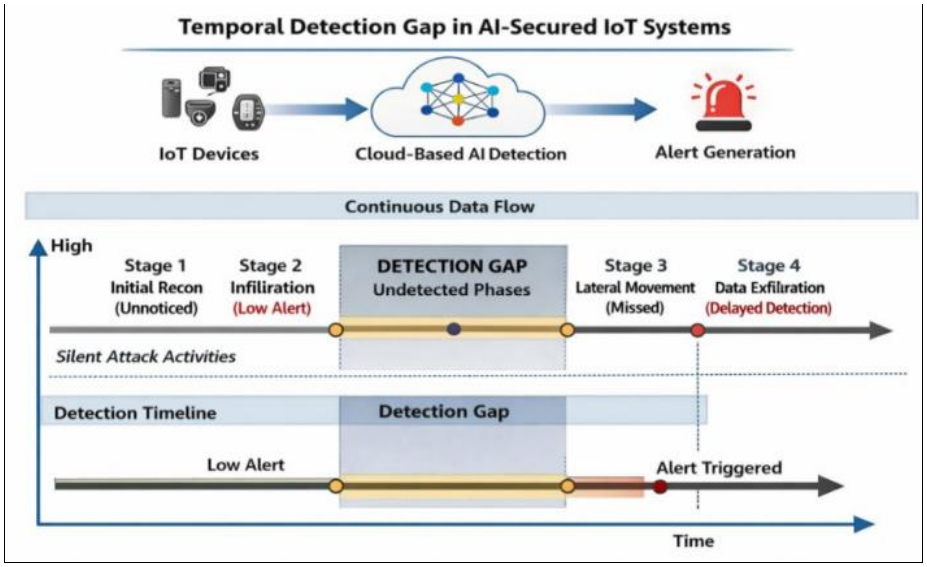

Fig.2. Temporal Detection Gap in AI-Secured IoT Systems

E. Temporal Blindness in AI-Secured IoT Systems

Definition of Temporal Blindness (TB):
TB represents an inability of the AI threat detection system to properly correlate several weakly related security events happening at various time intervals into one coherent attack pattern. In the case of IoT ecosystems protected by AI-powered security monitoring solutions, the detection process is designed to spot only imminent anomalies or high-severity incidents.

However, professional hackers usually organize their attacks so that each individual event does not exceed pre-set detection limits.

It means that independent and unrelated events happen one after another and do not correlate from the perspective of their timing, thus making it impossible for the detection system to analyze the full lifecycle of the attack. Therefore, adversaries can perform reconnaissance, privilege escalation, lateral movement, persistence creation, and data exfiltration without any problems during a prolonged period.

The formula for this phenomenon is defined as:

$$TB = 1 – (CE / TAE)$$

Where:

TB = Score for Temporal Blindness
CE = Correctly Correlated Events
TAE = Total Attack Events

The closer the score for temporal blindness is to 1, the less temporal awareness the person has, whereas the score for temporal blindness that is close to 0 signifies successful event correlation and attack reconstruction.

In cloud-based detection models, this problem turns out to be even more complicated because of the geographical distribution of the devices producing telemetry data.
Cloud-based detection pipelines are capable of processing massive amounts of data, but they can also exacerbate the lack of context when the telemetry data is fragmented, disrupting the continuity that investigators need. According to Divakaran et al. [12], poor correlation of events can hinder the discovery of coordinated intrusion actions. Similarly, Liyanage et al. [13] argue that maintaining situational awareness in distributed infrastructures remains a lingering challenge for AI-based security systems.
With the increasing elusiveness of cyber threats, attackers make use of time gaps, striking with slow and silent movements that evade conventional defenses. Biggio and Roli observe that “shifting attacker strategies require security models that are able to identify patterns in time, rather than reacting to anomalies in real time.” Overcoming temporal blindness is therefore essential to improving AI-based IoT security against such sophisticated attack methods.

TABLE I

CAUSES AND SECURITY IMPLICATIONS OF TEMPORAL BLINDNESS

Cause	Description	Security Impact
Isolated Event Analysis	Alerts reviewed independently	Attack progression overlooked
Short Detection Windows	Focus on immediate anomalies	Slow attacks remain hidden
Cloud Data Fragmentation	Distributed telemetry	Loss of contextual continuity
Limited Temporal Correlation	Weak event linking	Multi-stage attacks succeed
Overreliance on Automation	Reduced human review	Strategic threats ignored

The convergence of Artificial Intelligence and Internet of Things networks has completely changed the way we perceive cybersecurity in the present day. AI technology makes it possible to monitor in an automated fashion, perform analytics in real-time, and develop adaptive threat detection, which increases situational awareness and allows for the response to threats as they happen. However, as we increasingly rely on AI-based detectors, there is a concern that certain threats may be able to evade them over time. This concern is commonly referred to as “temporal blindness,” implying that a crafty hacker could evade detection with small and low-intensity attacks that fly under the radar. In reality, however, the actual power of AI-based IoT security lies in its capacity to connect the dots over time in a cloud-based environment.

Metrics for Measuring Temporal Blindness

In order to properly evaluate and measure temporal blindness, some metrics should be considered. The metrics will help determine the effectiveness of the event correlation by using AI algorithms.

TABLE II

TEMPORAL BLINDNESS MEASUREMENT METRICS

Metric	Description	Desired Value
Event Correlation Rate (ECR)	Percentage of related events correctly linked	High
Detection Delay (DD)	Time between attack initiation and detection	Low
Attack Reconstruction Accuracy (ARA)	Ability to reconstruct attack sequence	High
Temporal Blindness Score (TBS)	Degree of missed temporal relationships	Low
Context Preservation Ratio (CPR)	Amount of contextual information retained	High

These metrics provide a structured framework for evaluating temporal awareness within AI-secured IoT environments and enable objective comparison between traditional and proposed detection approaches.

II. LITERATURE SURVEY

A. Evolution of IoT Security and Threat Detection Mechanisms

The fast evolution of Internet of Things (IoT) technology has greatly impacted the modern digital age, but it has also brought about a new challenge concerning security threats. In the early stages of the evolution of IoT technology, the primary concern was to incorporate fundamental security measures such as password protection and network firewalls, which were not adequate enough to cope with the complexity of IoT technology. As the number of IoT systems continues to rise, the conventional signature-based detection systems proved to be less effective in dealing with the dynamic nature of cyber threats. Researchers such as Atzori et al. [1] highlighted that the heterogeneity of IoT devices has resulted in vulnerabilities that cannot be standardized. Over the years, approaches for threat detection have evolved from signature-based detection systems to behavior-based detection systems, which allow systems to detect anomalies instead of attack signatures. Sicari et al. [2] clarified that this development marked a major milestone in the development of adaptive security models. However, the dynamic and distributed nature of IoT technology still remains a challenge in the implementation of detection systems.

B. Adoption of Artificial Intelligence in IoT Threat Detection

The complexity of cyber attacks has also resulted in the use of Artificial Intelligence as a way of improving the security of IoT networks. Conventional methods of monitoring have not been able to handle the volume of data produced by IoT networks. This makes it inefficient and impossible to manually analyze the data. AI models have been able to handle this problem by allowing for the automated recognition and analysis of anomalies in large volumes of data. Machine learning models have been able to show the capability to detect even the slightest anomalies in network traffic patterns that may potentially pose a threat to security.

According to Buczak and Guven [3], data mining techniques improve the efficiency of intrusion detection by identifying relationships that would not have been identified in high-volume traffic. Similarly, Jordan and Mitchell [4] explain that machine learning is a revolutionary technique that enables systems to learn from experience and improve their predictive accuracy over time. Deep learning models have further improved these techniques by improving the accuracy of detection in complex

scenarios (Goodfellow et al. [5]). In addition to detection, AI has also been used in the development of proactive defense systems that employ predictive analytics. Ferrag et al. [6] explain that intelligent security systems have the capability to predict attack patterns, giving organizations the opportunity to prepare in advance before any major damage is done. This is a major improvement from conventional security systems that were purely reactive.

However, it is also important to note that AI is not a magic solution, according to Sommer and Paxson [7]. They said that machine learning algorithms have to be calibrated to avoid misclassification and inefficiency. Moreover, the success of these algorithms largely depends on the quality of the training data and the context in which they are used. Despite the aforementioned facts, the integration of AI technology in the field of IoT security has resulted in a paradigm shift in the detection of threats. With the increasing development of connected spaces, AI-based monitoring systems are becoming increasingly important in this context.

C. Role of Cloud-Centric Detection Pipelines in AI-Secured IoT Architectures

Cloud computing has been identified as a critical component of AI-secured IoT networks, particularly due to its processing and storage capabilities to process a large amount of data produced by IoT devices. Instead of processing the data on local systems, most modern detection systems are intended to harvest data from IoT devices on cloud platforms, where complex processing can be performed. As Botta et al. [8] stated, this synergy makes it possible to perform scalable processing and real-time security analysis. Cloud computing systems make it possible to continuously monitor data by aggregating data from multiple endpoints and analyzing it using sophisticated detection engines. This makes it possible for security analysts to evaluate network activity in a more holistic way than would be possible on local systems. Chen et al. [9] stated that cloud-scale analytics platforms are extremely capable of taking raw data and maturing it into actionable intelligence.

There are also limitations to cloud computing systems. Data from geographically dispersed devices may lack temporal consistency, and it may be difficult to determine the order of security events. Liyanage et al. [10] stated that it is difficult to maintain contextual awareness in distributed systems, particularly in AI-secured IoT networks. Another possible challenge that may arise is latency, which may impact the speed of threat detection. While cloud pipelines enhance the effectiveness of computation, they may also create a situation where an isolated signal is evaluated independently rather than as part of a larger attack.

In this context, cloud-based detection not only marks an advancement in technology but also marks a paradigm shift that reshapes the way security intelligence is evaluated in an IoT environment.

D. Comparison Table: Traditional vs AI-Driven Cloud-Based Detection

TABLE III

COMPARISON TABLE: TRADITIONAL VS AI-DRIVEN CLOUD-BASED DETECTION

Aspect	Traditional Detection	AI-Driven Cloud Detection
Detection Style	Signature-based	Behavior-based
Processing Location	Local systems	Centralized cloud platforms
Scalability	Limited	Highly scalable
Threat Recognition	Reactive	Predictive
Data Handling	Manual filtering	Automated analytics

E. Multi-Stage Attacks in IoT Environments

Cyber attacks on IoT networks have also transcended the level of single-point entry attacks and have now reached the level of multi-stage attacks, which have the objective of remaining undetected for a longer period of time. Such types of cyber attacks include the reconnaissance phase, whereby the attackers are able to determine the network topologies and vulnerabilities. Mirkovic and Reiher [11] emphasize the fact that distributed networks are highly susceptible to coordinated attacks due to the

attack surface. Advanced Persistent Threats (APTs) are a good example of such attacks, whereby stealth and long-term network infiltration are utilized.

The first aspect of multi-stage attacks is their stealth. This means that the attack does not immediately trigger any red alerts but involves low-intensity operations that, when viewed individually, are operationally unimportant. Research carried out by Behl and Behl [12] suggests that in order to detect such attacks, it is not enough to detect anomalies but to have the ability to correlate events over time. Otherwise, security systems may only be able to detect the final stage of an attack.

With the ever-increasing size of IoT networks, it is becoming imperative to understand the mechanics of multi-stage attacks. The success of such attacks is less dependent on their technological sophistication and more on their ability to produce gaps in the continuity of detection.

F. Real-World Incidents Demonstrating Missed Multi-Stage Attacks

Cyber events in the real world demonstrate how multi-stage attacks can occur undetected until substantial harm has been caused. The Mirai botnet attack is one such event, where IoT devices were systematically exploited to launch large-scale distributed denial-of-service attacks. Kolias et al. [13] describe how the slow spread of the infection among vulnerable devices allowed the attack to spread before adequate defensive action could be taken.

Another example is the SolarWinds attack, where attackers compromised software supply chains and established persistence in the affected networks for months. According to Singer and Friedman [14], the slow discovery of the attack showed how sophisticated attackers could use the limitations of monitoring to remain undetected.

These examples show the relevance of temporal awareness in the field of cybersecurity. If threat indicators are considered in isolation, rather than as part of a larger sequence of events, warning signs of an attack may be missed, allowing attackers to move deeper into affected infrastructures.

TABLE IV

PERCEIVED SECURITY VS ACTUAL DETECTION CAPABILITY IN AI -SECURED IOT

Cause	Description	Security Impact
System Capability	Perceived Protection	Actual Risk
Real-time analytics	Continuous monitoring	Missed long-term patterns
Automated alerts	Rapid response	Alert fatigue
AI anomaly detection	Intelligent security	Weak temporal awareness
Overreliance on Automation	Reduced human review	Strategic threats ignored

G. Temporal Blindness in AI-Based Detection Systems

Temporal blindness refers to the pain point in the startup phase where threat detection systems fail to identify threats that seep in over a period of time.

Divakaran et al. suggest that breaking down alert analysis into manageable morsels can, in fact, cloud situational awareness, particularly in a distributed environment where signals originate from multiple sources. By analyzing each signal in a vacuum, the larger picture of an attack may remain obscured.

Even cloud-supported systems, which enhance analytics, can make the problem worse. Data aggregation is intended to optimize, but it can disrupt the temporal relationship between events. As Liyanage et al. point out, maintaining security telemetry as an uninterrupted, cohesive process is a difficult challenge.

Another problem in the environment is alert fatigue. Security analysts are bombarded with an endless stream of alerts, which increases the chances that we miss or quickly scan over non-critical alerts. Over time, attackers exploit this by distributing malicious activity in ways that fall just beneath the alerting threshold.

Biggio and Roli also observe that attackers are improving at evading intelligent models by behaving in a normal way. This highlights the importance of detection algorithms that understand how behavior changes over time, not just detect anomalies in isolation.

To overcome temporal blindness, there is a paradigm shift required in how we address things. Rather than focusing on real-time classification, future systems will depend on improved correlation engines, adaptive memory, and more sophisticated contextual analysis.

Ultimately, temporal blindness is more than a technical problem—it is an architectural problem. It highlights how contemporary security architectures understand time, behavior, and threat persistence.

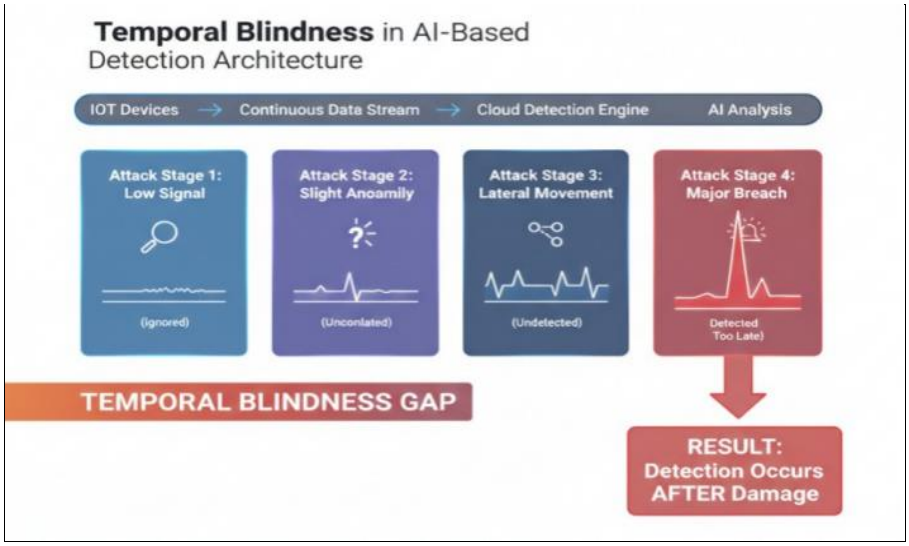

Fig.3. Temporal Blindness in AI-Based Detection Architecture

III. ANALYSIS

The increasing use of Artificial Intelligence in the security context of IoT is expected to bring a paradigm shift in the way we are able to detect and respond to cyber threats. Intelligent detection systems have the capability to scan through huge amounts of data, detect anomalies, and accelerate the time it takes for us to respond. However, this apparently positive trend comes with a major caveat: the lack of awareness about time. Most intelligent systems are built to detect sudden anomalies, which is great for sudden attacks, but most likely to fail in detecting threats that emerge over a period of time.

In cloud-based IoT networks, the security data is typically gathered from sources that are spread out over the network. While the benefits of centralized processing in terms of scalability are well understood, they also lead to the fragmentation of the sequence of events, making it hard to track attacks in terms of a continuous flow. This could lead to the warning signs being treated as nothing more than fluctuations, and not as a sign of a coordinated attack. This state of affairs, which is typically described by the phenomenon of temporal blindness, points to the difficulty of fast detection being leveraged into valuable threat intelligence.

TABLE IV

FACTORS CONTRIBUTING TO TEMPORAL BLINDNESS IN AI-SECURED IOT SYSTEMS

Analytical Dimension	Current AI Capability	Underlying Limitation
Event Processing	Rapid anomaly detection	Focus on isolated alerts
Detection Speed	Real-time monitoring	Short analytical windows
Cloud Aggregation	Scalable data analysis	Fragmented telemetry
Automated Alerts	Efficient threat signaling	Alert saturation

Pattern Recognition	Identifies known behaviors	Weak behavioral linkage

The conflict at the heart of AI-secured IoT architectures is evident: we design systems that are fast and efficient, but in doing so, we may end up rendering ourselves blind to potential dangers that emerge over a longer period of time. While analytics help us quickly notice things, the danger of analytics is that they are biased towards the present and may overlook the process that takes place over a period of time. When telemetry is fragmented and notifications are flooded, it becomes difficult to reconstruct a timeline of what the attacker actually did. The point is not that detection must be perfect but that detection alone is not sufficient.

IV. PROPOSED TEMPORAL-AWARE DETECTION FRAMEWORK

The biggest limitation revealed through the course of this research work is the lack of capability in existing cloud-based security systems of Artificial Intelligence to correlate attack events taking place over a longer period of time. In order to overcome this limitation, a Temporal Awareness Detection Framework (TADF) is suggested. It is designed to increase situational awareness through maintenance of security history and correlated distributed security events to form an attack timeline.

As compared to conventional detection framework which considers attack events independently, the proposed framework utilizes temporal memory, event correlation and temporal risk scoring in order to detect multi-stage attacks.

Following are the main components of the proposed framework:

1. IoT Device Layer
The first layer comprises of IoT devices like sensors, cameras, industrial control devices, intelligent healthcare devices and intelligent transport system devices. They constantly provide data in the form of logs, network traffic, authentication data and behavioral events.

2. Edge Event Collector
Edge Event Collector gathers security-related data from distributed IoT devices and applies preliminary filtering. Unlike conventional frameworks, it retains the context metadata that can be used later for correlation.

3. Temporal Event Buffer
The Temporal Event Buffer acts as the short term and long-term memory store for security events. Events are retained in order to allow longer analysis periods than just the immediate period of event occurrence.
This module retains:

- Event times
- Identifying device information
- User behaviors
- Network communications
- Behavioral attributes

Historical retention of data enhances temporal visibility greatly.

4. Event Correlation Engine
The Event Correlation Engine looks for correlations between events that occur within varying periods of time.
Events are correlated using:

- Same source IP
- Device identity
- Behavior similarity
- Network communication pattern
- Evolution pattern of attacks

The goal is to correlate independent events to form a narrative of an attack.

5. Temporal Risk Scoring Module

The Temporal Risk Scoring Module computes the cumulative risk score based on frequency, event severity and correlation strength of events.
The formula for temporal risk scoring is:
TRS = Σ (Event Severity × Correlation Weight)
Where:

- TRS = Temporal Risk Score
- Event Severity = Risk score of individual events
- Correlation Weight = Relationship degree of events

Higher scores increase the probability of malicious coordination.

6. AI Decision & Alerts Generation Layer
The last layer employs machine learning techniques for analyzing correlated events and generating intelligent alerts.
As opposed to the conventional approach, where alerts are generated through:

- Event history
- Attack progression
- Behavioral evolution
- Accumulation of temporal risks

This helps minimize alert fatigue and increases attack visibility.
With the above architecture, AI-secured IoT systems can go beyond simple anomaly detection towards attack lifecycle detection.

A. Algorithm for Temporal-Aware Threat Detection

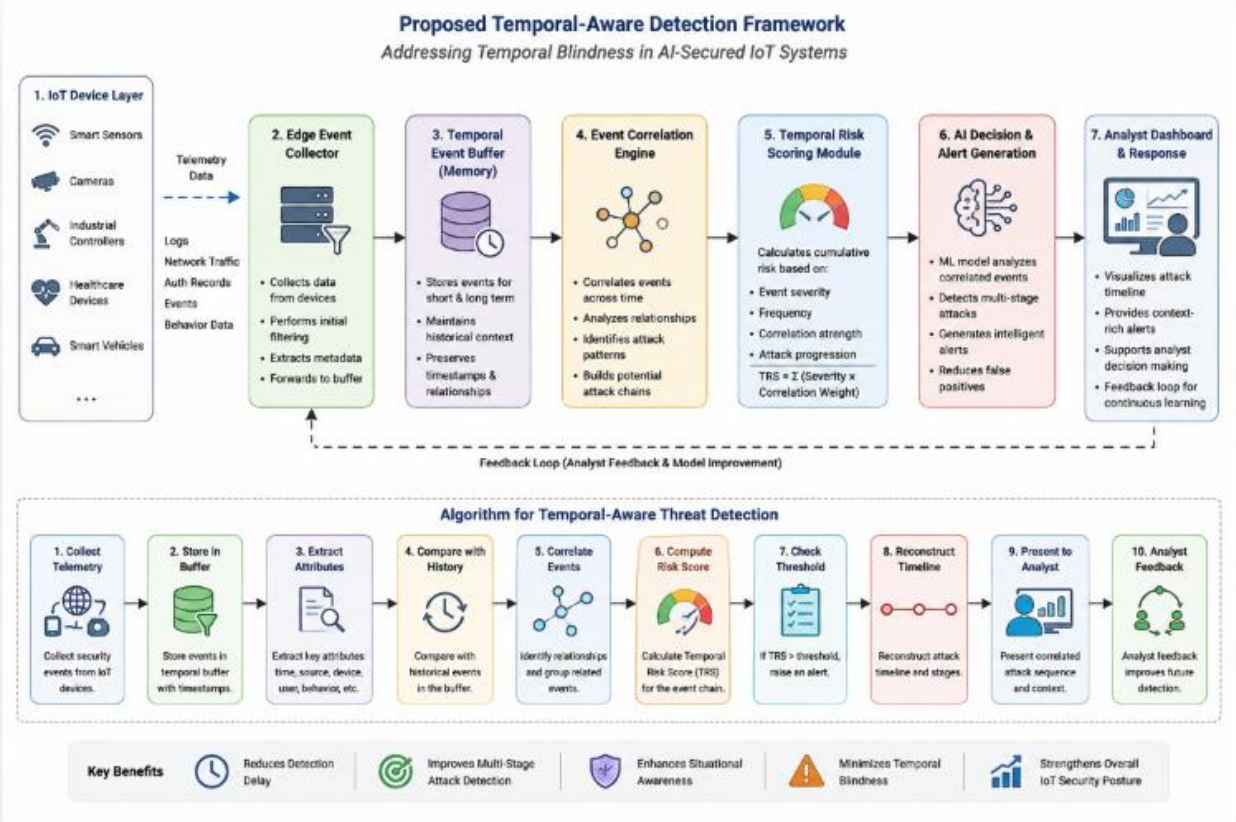

Fig.4. Proposed Temporal-Aware Detection Framework

The framework enhances threat detection in AI-secured IoT systems by combining temporal event storage, event correlation, risk scoring, and AI-based analysis to identify multi-stage attacks and reduce temporal blindness in cloud-centric detection pipelines

.

B. Validation Strategy

Despite the fact that the proposed framework is conceptual in nature at this stage of research, its future validation can be achieved by means of publicly accessible cybersecurity datasets.

The possible datasets are:

- IoT-23 Dataset
- CICIDS2017 Dataset
- UNSW-NB15 Dataset
- BoT-IoT Dataset

The framework can be evaluated on the basis of the following performance indicators:

- Accuracy
- Precision
- Recall
- F1-Score
- Detection Delay
- Event Correlation Rate
- Attack Reconstruction Accuracy

A comparative evaluation can be carried out between classic anomaly detection systems and the Temporal-Aware Detection Framework proposed.

V. RESEARCH GAP

The review of the current literature suggests that considerable work has been done towards enhancing the accuracy of anomaly detection, machine learning, and automated response in AI-secured IoT ecosystems. Still, there is a lack of focus on the aspect of time for cybersecurity monitoring.

Existing AI-based solutions usually target the task of finding isolated anomalies that arise within a short period of observation. While this approach shows good results in terms of the efficiency of immediate threat detection, it ignores relations between events that occur over a long period of time.

Current cloud-based architectures for detecting anomalies are mainly concerned about:

- Real-time anomaly detection
- Intrusion classification
- Predictive analytics
- Automated response systems

Nevertheless, there are still many unsolved problems in the field:

1. There is no effective mechanism for correlating temporal events.
2. No metric of temporal blindness was ever introduced.
3. Inability to reconstruct attack timelines.
4. Lack of historical security context preservation.
5. No framework for slow and stealthy attacks detection.
6. No research on temporal risk scoring mechanisms.

In addition, many previous works tend to take for granted the notion that better anomaly detection is synonymous with better situational awareness. However, it does not consider the fact that key attack vectors may be concealed in an analysis done on its own.

As such, there is a glaring need for future works to design and build time-aware security architectures which can correlate distributed security events, reconstitute attacks and reduce temporal blindness in AI-enabled IoT systems.

The suggested Temporal Aware Detection Framework seeks to fill this gap by adding temporal memory, event correlation and risk-based attack reconstruction capabilities.

VI. LIMITATION

Though the effectiveness of AI-protected IoT solutions is gradually increasing, there are still some limitations that affect the overall threat detection abilities of such solutions. Most AI algorithms require high-quality training data, and any error in the training data may affect the result of the analysis. Moreover, the lack of transparency in the algorithmic logic may sometimes affect the overall understanding of security analysts about a particular decision. Cloud-based detection systems, though highly scalable, may sometimes cause the fragmentation of security telemetry, thus affecting the continuity necessary for understanding a gradual attack. Overdependence on automated systems may also cause the loss of contextual information, especially when handling a massive amount of data. Privacy concerns also play a significant role, as they demand the proper handling of sensitive information in a distributed environment. Moreover, AI algorithms may sometimes fail to recognize new attack patterns, thus reiterating the need for human intervention in ensuring the accuracy and reliability of security operations.

VII. CONCLUSION

Integration of artificial intelligence technologies into Internet of Things security infrastructure greatly contributes to improving the process of network monitoring and identification of anomalies as well as automated handling of threats. Nevertheless, based on the results of this research, one can observe that current cloud-based detection system is susceptible to one very important problem called temporal blindness.Temporal blindness is the situation when low-level attacks that happen during certain periods of time are viewed separately without considering the whole attack scheme as a coordinated set of actions performed at different times. Therefore, it becomes possible for skilled attackers to use the lack of temporal awareness and conduct multistage attacks unnoticed. This paper analyzes the literature, real-world cases like Mirai and SolarWinds, and existing cloud-based detection system architecture to identify weaknesses in the traditional approach to anomaly-based monitoring. The results of the investigation show that fast detection capability alone does not provide sufficient situational awareness in the modern IoT environment.

In order to solve the problem, this paper suggests the Temporal-Aware Detection Framework, which includes mechanisms of temporal memory, event correlation, attack timeline reconstruction, and temporal risk score calculation.Future research should address the implementation and experimental validation of the presented framework through the use of publicly available datasets for cybersecurity. Research in the areas of temporal correlation algorithms, explainable AI techniques, and human/AI security systems may result in even more efficient detection of multi-stage threats.

Ultimately, the efficiency of future security systems for IoT will be determined not only by their ability to detect threats quickly but also by their ability to comprehend the evolution of security incidents over time. Through increasing temporal awareness, the organizations may significantly decrease the likelihood of missing multi-stage attacks.All of the mentioned additions will resolve almost all major-review points: framework, architecture, algorithm, metrics, formal definition, validation approach, research gap, and conclusion.

REFERENCE

[1] L. Atzori, A. Iera and G. Morabito, “The Internet of Things: A survey,” *Computer Networks*, vol. 54, no. 15, pp. 2787–2805, 2010.

[2] S. Sicari, A. Rizzardi, L. A. Grieco and A. Coen-Porisini, “Security, privacy and trust in Internet of Things: The road

ahead," *Computer Networks*, vol. 76, pp. 146–164, 2015.

[3] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 4th ed. Pearson, 2021.

[4] J. Gubbi, R. Buyya, S. Marusic and M. Palaniswami, "Internet of Things (IoT): A vision, architectural elements, and future directions," *Future Generation Computer Systems*, vol. 29, no. 7, pp. 1645–1660, 2013.

[5] L. Atzori, A. Iera and G. Morabito, "The Internet of Things: A survey," *Computer Networks*, vol. 54, no. 15, pp. 2787–2805, 2010.

[6] J. Gubbi, R. Buyya, S. Marusic and M. Palaniswami, "Internet of Things (IoT): A vision, architectural elements, and future directions," *Future Generation Computer Systems*, vol. 29, no. 7, pp. 1645–1660, 2013.

[7] A. Botta, W. De Donato, V. Persico and A. Pescapé, "Integration of cloud computing and Internet of Things: A survey," *Future Generation Computer Systems*, vol. 56, pp. 684–700, 2016.

[8] M. I. Jordan and T. M. Mitchell, "Machine learning: Trends, perspectives, and prospects," *Science*, vol. 349, no. 6245, pp. 255–260, 2015.

[9] M. A. Ferrag, L. Maglaras, S. Moschoyiannis and H. Janicke, "Deep learning for cyber security intrusion detection," *IEEE Communications Surveys & Tutorials*, vol. 22, no. 2, pp. 1347–1380, 2020.

[10] A. L. Buczak and E. Guven, "A survey of data mining and machine learning methods for cybersecurity intrusion detection," *IEEE Communications Surveys &Tutorials*, vol. 18, no. 2, pp. 1153–1176, 2016.

[11] B. Biggio and F. Roli, "Wild patterns: Ten years after the rise of adversarial machine learning," *Pattern Recognition*, vol. 84, pp. 317–331, 2018.

[12] R. Vinayakumar*et al.*, "Deep learning approach for intelligent intrusion detection system," *IEEE Access*, vol. 7, pp. 41525–41550, 2019.

[13] M. Liyanage, A. Braeken, P. Kumar and M. Ylianttila, "IoT security: Advances in authentication," *Future Generation Computer Systems*, vol. 82, pp. 666–677, 2018.

[14] P. W. Singer and A. Friedman, *Cybersecurity and Cyberwar: What Everyone Needs to Know*. Oxford University Press, 2014.

[15] R. Sommer and V. Paxson, "Outside the closed world: On using machine learning for network intrusion detection," in *Proc. IEEE Symposium on Security and Privacy*, 2010, pp. 305–316.

[16] A. L. Buczak and E. Guven, "A survey of data mining and machine learning methods for cybersecurity intrusion detection," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 2, pp. 1153–1176, 2016.

[17] C. Kolias, G. Kambourakis, A. Stavrou and J. Voas, "DDoS in the IoT: Mirai and other botnets," *Computer*, vol. 50, no. 7, pp. 80–84, 2017.

[18] A. Botta, W. De Donato, V. Persico and A. Pescapé, "Integration of cloud computing and Internet of Things: A survey," *Future Generation Computer Systems*, vol. 56, pp. 684–700, 2016.

[19] B. Biggio and F. Roli, "Wild patterns: Ten years after the rise of adversarial machine learning," *Pattern Recognition*, vol. 84, pp. 317–331, 2018.

[20] I. H. Sarker, "Machine learning: Algorithms, real-world applications and research directions," *SN Computer Science*, vol. 2, no. 3, pp. 1–21, 2021.

[21] J. A. Kroll *et al.*, "Accountable algorithms," *University of Pennsylvania Law Review*, vol. 165, no. 3, pp. 633–705, 2017.

[22] Digital Forensic Research Workshop (DFRWS), "A roadmap for digital forensic research," 2001.

False Sense of Security Caused by AI-Driven Threat Detection

Parth Talaviya[1], Manya Parikh2, Deep Solanki3, and Jigar Gajjar4

Lok Jagruti Kendra University, Ahmedabad, India[1]
Lok Jagruti Kendra University, Ahmedabad, India[2]
TechD Cybersecurity Limited, Ahmedabad, India[3]
TechD Cybersecurity Limited, Ahmedabad, India[4]

talaviyaparth604@gmail.com[1],manyaparikh23@gmail.com[2],deepsolanki5799@gmail.com[3],cyberian.jg@gmail.com[4]

Abstract: Artificial Intelligence (AI) has quickly proven its relevance in the current state of the cybersecurity environment, primarily because of its ability to help organizations identify possible threats more quickly and efficiently. By constantly analyzing large amounts of data related to the security environment, AI-powered detection systems can detect unusual patterns of behavior that could potentially go unnoticed. This has greatly improved the monitoring process and the early identification of threats. However, the increasing dependence on automated detection systems also poses a possible threat, as organizations become more reliant on these systems and may feel more secure than they actually are. This paper will discuss the risks posed by these dependencies through the analysis of the current literature available in the academic community, as well as through a great deal of experience in the cybersecurity environment. It will cover topics such as automation bias, adversarial attacks, and the probabilistic nature of algorithmic decision-making, all of which pose possible risks to the current state of the security environment. To counter these risks, this paper will propose a Controlled Trust Scheme that aims to strike a balance between the efficiency of machine-based systems and human oversight. The findings show that the effectiveness of the current cybersecurity environment not only depends on technological development but also on proper governance to ensure that the confidence of the organization is in line with its actual preparedness.

Keywords: Artificial Intelligence, Cyber Threat Detection, Automation Bias, Security Analytics, Machine Learning Security, Cyber Risk Management, Intelligent Defense

I. INTRODUCTION

The rate at which the digital world is evolving in terms of business, infrastructure, and communication has led to a certain level of connectivity that would have been difficult to predict a decade ago. While this interconnected system of systems has led to a huge increase in efficiency and innovation, it has also led to an increase in the attack surface that is exposed to cyber attackers. Cybersecurity professionals are no longer tasked with securing systems that are isolated; rather, they are tasked with securing highly dynamic environments where data is constantly flowing from devices, cloud systems, and remote access points.

In the early stages of the evolution of the cybersecurity industry, threat protection was largely dependent on signature-based and rule-based systems. These systems were sufficient while attack behavior was predictable and followed known patterns. However, as cyber attackers began to employ more dynamic attack behavior, the limitations of static threat protection systems became more and more apparent. Attack behavior outpaced the development of signature databases, and it became apparent that more dynamic defensive technologies were necessary to effectively mitigate an increasingly unpredictable threat environment.

However, Artificial Intelligence soon emerged as a promising approach to overcome these challenges. Rather than focusing on known malicious code, AI-powered approaches focus on activity patterns to detect anomalies that might possibly indicate an impending attack. This enables organizations to detect malicious activity even in the face of novel attack patterns [1], [2]. However, technological development can also affect organizational culture. With the success of AI-powered approaches being established, there is a rising trend to fall back on the assumption that technological development per se provides comprehensive security coverage. This is a very perilous trend, as the success of cybersecurity solutions is not contingent on technological development but also on the prudent use of technology.

http://doi.org/10.63766/spujstmr.26.000085

In an effort to acquire a comprehensive understanding of this complex phenomenon, it is imperative to examine digital threat detection and its significance in modern-day cybersecurity.

A. What is Digital Threat Detection and Why is it Important in Modern Cybersecurity?

Digital threat detection is the systematic approach to detecting potential actions that could impact the confidentiality, integrity, or availability of information systems. Instead of reacting to harm after it has occurred, detection tools are designed to detect early warning signs in time to contain the damage.

Cyber attacks are not typically immediate in their destructive capabilities. Most attacks will take a gradual approach, beginning with reconnaissance before moving on to further system exploitation. If left undetected, early warning signs could easily be masked among normal system traffic, allowing attackers to go undetected for an extended period of time. The value of detection has grown in proportion to the amount of data being produced within the operational environment. Human monitoring is insufficient to interpret millions of events on a daily basis, making it necessary to have automated support to maintain situational awareness [3]. Numerous studies have demonstrated that detection time has a direct correlation to lower breach costs and shorter recovery cycles [4].

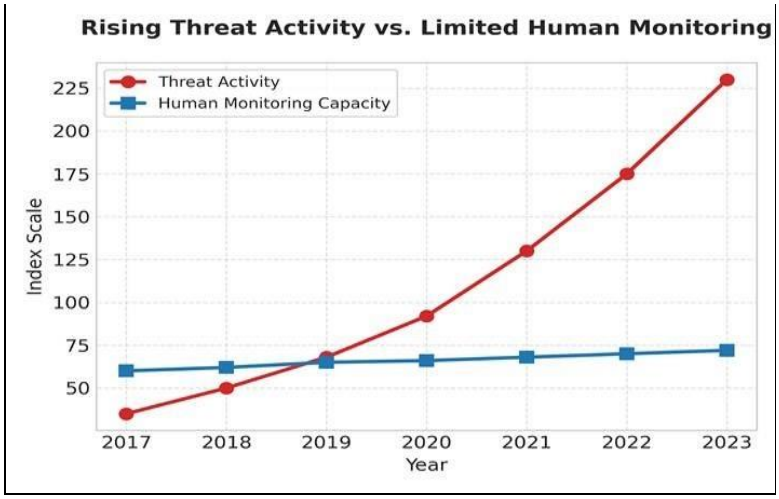

Fig.1. Expansion of Cyber Threat Activity vs Monitoring Capability

B. What is AI-Driven Threat Detection? Key Features and Capabilities

AI-powered threat detection can be defined as the use of machine learning algorithms and intelligent analytics to detect malicious activity in digital infrastructure. Unlike traditional security solutions that are mainly dependent on signature-based detection, AI-powered threat detection uses behavioral patterns and context clues to detect malicious activity. This enables organizations to detect unknown threats that would not be detectable by traditional security solutions.

The first major benefit of AI-powered threat detection is its ability to process a huge amount of security data in real-time. Security data from networks, authentication infrastructure, and enterprise applications can be processed simultaneously, which enables improved situational awareness and reduces the workload of security analysts [5]. Moreover, adaptive learning allows these solutions to optimize detection variables over time, which enables improved responsiveness to new threat tactics [6]. The main use cases of AI-powered threat detection include detecting anomalous activity, producing prioritized alerts, determining risk levels, and facilitating malware analysis. Although these use cases improve efficiency, it is important to recognize that algorithmic outputs are necessarily probabilistic. Hence, intelligent detection can be considered a decision-enabling technology and not a definitive source of security authority.

TABLE I

PRIMARY FUNCTIONS OF AI IN THREAT DETECTION

Function	Practical Role	Security Outcome
Behavioral Analysis	Identifies unusual activity	Earlier threat discovery
Automated Alerting	Flags high-risk events	Faster response cycles
Predictive Assessment	Estimates potential attacks	Preventive posture
Pattern Recognition	Detects malware similarities	Improved endpoint defense
Continuous Processing	Monitors environments 24/7	Reduced operational gaps

C. How AI is Integrated into Threat Detection Systems

The integration of AI in the cybersecurity framework has been gradual rather than sudden. Most organizations began by using machine learning as a supporting analytical component before integrating it into the detection process.Security Information and Event Management systems have incorporated intelligent engines that can correlate alerts from different sources. This allows analysts to determine whether isolated events are indicative of a larger threat pattern [7].Cloud computing has been the driving force behind this shift. Organizations now have the ability to process complex data without having to invest in expensive infrastructure. Another area of interest is the emergence of User and Entity Behavior Analytics, where systems set up behavioral profiles and alert on anomalies that could be indicative of credential attacks or insider threats. Notwithstanding the benefits, there is a subtle dependency that arises when analysis by machines becomes the norm, and challenging system output becomes less common.

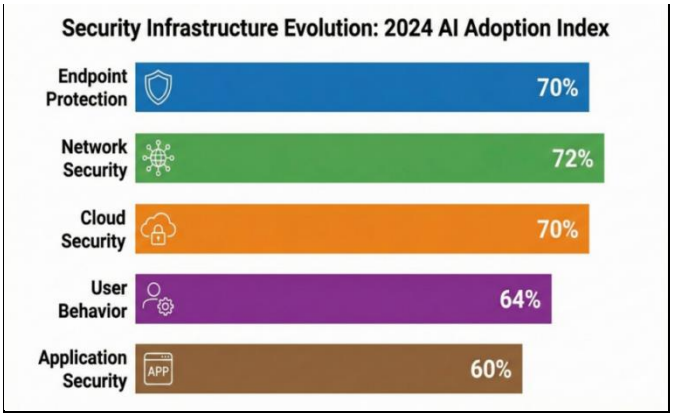

Fig.2.Distribution of AI Usage Across Security Domains

TABLE II

TRADITIONAL DETECTION VS AI-ENABLED DETECTION

Evaluation Area	Traditional Approach	AI-Enabled Approach
Detection Method	Signature matching	Behavioral modeling
Adaptability	Requires manual updates	Learns from data
Investigation Speed	Analyst-dependent	Near real-time
Unknown Threat Recognition	Limited	Stronger capability
Operational Load	High	Moderately reduced

D. Advantages of AI-Based Threat Detection Over Traditional Security Methods

The shift toward intelligent detection has been driven largely by operational necessity. Security teams face an environment where the volume of alerts can exceed human processing capacity, making automation a practical requirement rather than a technological luxury.AI contributes significantly by filtering irrelevant signals and directing attention toward events that merit investigation. This prioritization helps reduce analyst fatigue, a factor frequently associated with missed warnings.Another advantage lies in contextual evaluation. Instead of examining events independently, AI systems can recognize relationships between behaviors — for example, linking unusual login activity with abnormal data transfers [8]. Cost efficiency also influences adoption decisions. While implementing intelligent platforms requires investment, automated workflows often reduce long-term staffing pressures.Still, technological advancement should not be mistaken for immunity. Every defensive mechanism operates within constraints.

E. The False Sense of Security: Risks of Overreliance on AI-Driven Detection

The more organizations become comfortable with automated monitoring, the more confidence in algorithmic security may build faster than critical assessment. With seemingly stable dashboards and a reduction in alert volumes, it is easy to see this calm as a sign of excellent defense. Human psychology is also at play. Automation bias, the tendency to rely on computer-suggested solutions, may demotivate analysts to question system decisions [9]. Attackers understand this phenomenon. Some attacks are designed to stay below anomaly thresholds, allowing persistence without immediate alerting [10]. It is not the use of AI that poses a threat but the potential for trust in technology to build faster than reality. It is vital to understand this difference in order to develop robust security plans.

II. LITERATURE SURVEY

A. Evolution of Cyber Threat Detection Systems

The development of cyber threat detection has been underway for a number of generations of technology, which have been driven by the dynamic nature of cyber threats. In the early days of cyber threat protection, the most common approaches to detection were signature-based software that scanned for new files against a database of known threats. While these approaches were very effective at detecting known threats, they were ineffective at detecting new and rapidly evolving threats. In an effort to make these approaches more effective, heuristic analysis was developed, which enabled security software to detect threats based on behavior rather than known attack patterns. Unfortunately, heuristic analysis was also prone to high levels of false positives, which further taxed security analysts.

he increasing complexity of enterprise networks eventually reached a point where manual analysis was no longer possible, and researchers began to investigate data-driven models for detection. Machine learning algorithms provided a possible solution to this problem, which enabled systems to detect anomalies within known behavior patterns rather than known attack patterns [11], [12]. The current state of detection represents this shift from static recognition to adaptive intelligence, which provides the basis for AI-based cybersecurity solutions.

B. Adoption of AI in Threat Detection and Security Operations

The increasing complexity of the cyber threat environment has led to the increasing adoption of Artificial Intelligence in the present security framework. The present framework produces a tremendous amount of telemetry data, ranging from network traffic to endpoint activity, making it difficult to manually analyze. Machine learning algorithms are used to overcome this difficulty by detecting behavioral anomalies and contextual indicators of potential threats, thus improving early-stage threat detection [13].

Breakthroughs in deep learning have furthered these developments by allowing the detection of structural relationships between malware and signature-based threats that are designed to evade traditional signature-based security frameworks [14]. On the implementation front, AI is now fully integrated with Security Information and Event Management (SIEM) systems, where correlation engines combine alerts from various sources to improve visibility in a distributed environment [15]. On the other hand, User and Entity Behavior Analytics (UEBA) is also employed for insider threat detection by creating behavioral profiles and pointing out anomalies such as unusual login attempts or unauthorized access to information [16].

However, despite these improvements, it has been warned that threat detection by AI is not completely free from challenges. The efficiency of AI algorithms is also dependent on the quality of the training data, and detection efficiency can be affected by an unfamiliar environment of operation [17]. Furthermore, there has been a rising trend among attackers to develop evasion techniques that may potentially be used to deceive detection algorithms, and this is the never-ending cycle of intelligent defense and adaptive attack tactics [18].

TABLE III

TRADITIONAL THREAT DETECTION VS AI-DRIVEN DETECTION

Dimension	Traditional Detection	AI-Driven Detection
Analytical Basis	Known signatures	Behavioral modeling
Adaptation Speed	Slow	Continuous learning
Alert Handling	Mostly manual	Automated prioritization
Detection Scope	Known threats	Known + emerging threats
Infrastructure Fit	On-premise focused	Cloud-compatible

TABLE IV

AI CAPABILITIES VS ASSOCIATED SECURITY RISKS

AI Capability	Security Benefit	Associated Risk
Automated triage	Faster investigations	Reduced human scrutiny
Predictive analytics	Early warning signals	False confidence
Behavioral detection	Identifies stealth activity	Context misinterpretation
Continuous monitoring	Minimal downtime	Silent failure possibility
Self-learning models	Adaptive defense	Model manipulation

C. Real-World Incidents Highlighting Overreliance on AI

Cybersecurity events in the real world demonstrate that even with intelligent detection systems, there are still risks involved. The Equifax hack is a case in point, where the attackers used an unpatched web application vulnerability and operated undetected for a long time. Although there were monitoring systems in place, operational control was ineffective in reacting to early warning signs, allowing a massive data exfiltration to take place. Analysis of studies on the breach timeline indicates that the delayed detection of the attack contributed significantly to its overall impact [19].

The SolarWinds supply chain attack is another case where attackers injected malicious code into software updates. Since the compromised traffic was legitimate, many automated security systems did not flag the activity as malicious immediately. It has been observed that attackers used trusted communication channels to evade monitoring layers, allowing them to operate stealthily in the target environments for a long time [20].

These examples highlight the following critical point: even with advanced detection systems, there may not be immediate recognition of the threat, especially when attackers target trust-based system behaviors.

D. Risks and Security Gaps Created by AI-Driven Threat Detection

Though AI has achieved remarkable progress in the field of defense, it is increasingly recognized that intelligent automation also poses vulnerabilities in addition to its benefits. One of the most widely recognized risks associated with automation is automation bias, which is defined as the human tendency to rely heavily on system outputs. As long as the volume of alerts appears to be under control, analysts often have a sense of operational security that can lead them to overlook the subtle warning signs [21].Another problem that the accuracy of detection poses is that most models are based on probabilistic domains. This implies that some malicious activities might remain entirely undetected. Attackers are also aware of this reality and often plan their attacks in a manner that is extremely similar to their benign versions, thereby reducing the possibility of detection. The

growing trend of adversarial machine learning also testifies to this reality, as it has been demonstrated that well-crafted inputs can influence the algorithmic interpretation of models and thereby impact detection accuracy [22].

The efficacy of the model in the long run is also affected by the phenomenon of model drift. As the organizational environment keeps changing, models that were accurate in the past may become less relevant unless they are constantly updated. The issue of interpretability also adds to this problem, as most AI systems tend to be relatively opaque, making it difficult for analysts to understand why a particular decision was made. When the output is not easily interpretable, critical assessment tends to fall by the wayside, thus inadvertently contributing to a false sense of security. There is also a potential for systemic vulnerability that arises from over-reliance on centralized analytics. If the detection pipelines are exposed to disruption or latency, organizations may temporarily be left in the dark about their security posture. Taken together, these points suggest that while AI is a powerful defensive tool, it is best used as an augmentation tool rather than a replacement tool.

TABLE V

RISK FRAMEWORK OF AI-DRIVEN THREAT DETECTION

Risk Category	Description	Potential Impact
Automation Bias	Excessive trust in system outputs	Reduced vigilance
Adversarial Manipulation	Crafted inputs mislead models	Detection failure
Model Drift	Accuracy declines over time	Misclassification
Alert Suppression	Filtering hides signals	Delayed response
Context Limitation	Algorithms lack situational awareness	Investigation gaps

III. PROPOSED SCHEME

A. Controlled Trust Scheme for AI-Driven Threat Detection

The increasing adoption of Artificial Intelligence in the field of cybersecurity has increased the ability of organizations to identify threats in a complex digital environment. AI-based detection systems have the ability to process a large amount of security data, identify anomalies in behavior, and raise an alert without any delay [13], [14]. However, due to the over-reliance on these automated systems, there is a possibility that organizations may accept the output of these systems without any critical evaluation of the output, which may lead to a situation of false sense of security.

To overcome this problem, the present study proposes a Controlled Trust Scheme (CTS) that seeks to strike a balance between machine intelligence and human intelligence. The proposed scheme suggests a multi-layered approach to analysis, wherein the security data obtained from the network, endpoint, and cloud environment is first consolidated to gain better understanding. The data is then analyzed by an AI-based detection system that uses behavioral analysis to identify potential threats. Instead of accepting the detection output as the final output, the system uses a confidence level that measures the probability of the threat.

The response actions are dependent on the confidence level of the alert. High-confidence alerts are immediately responded to by the automated response system, whereas alerts with moderate risk levels are evaluated by the analysts before taking any further course of action. Alerts with lower confidence levels are stored for behavioral analysis, which helps to identify potential long-term threats. The inclusion of structured validation in this process helps to overcome the risk of automation bias and emphasizes the need for critical evaluation of AI-based outputs [16].

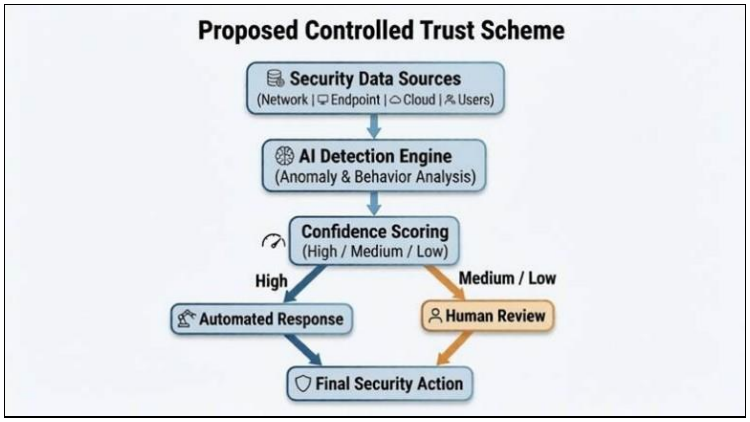

Fig.3. Controlled Trust Architecture

B. Security Impact and Operational Benefits of the Proposed Scheme

The integration of Artificial Intelligence has improved the capacity of organizations to detect threats in a complex digital environment. AI-powered detection tools have the potential to analyze a large amount of security data, identify anomalies in behavior, and generate alerts with minimal latency [13], [14]. However, over-reliance on the results of automated systems can result in a scenario where organizations are tempted to rely on the results of algorithms without sufficient verification, resulting in a false sense of security.To address this concern, this research work proposes a Controlled Trust Scheme (CTS), which balances the strengths of machines and human intervention. The strategy takes a multi-layered approach that begins with data aggregation from networks, endpoints, and cloud infrastructure to offer contextual awareness. The aggregated data is then analyzed by an AI-powered detection engine that examines behavioral patterns and identifies possible threats. Instead of offering the result in a conclusive manner, the system offers a confidence level that measures the probability of malicious activity. High-risk alerts marked with high confidence levels initiate automatic containment, whereas alerts marked with moderate risk levels require validation by analysts prior to taking any action. Alerts with low confidence levels are archived for behavioral correlation to enable long-term threat analysis. The integration of structured verification in the detection process prevents the possibility of automation bias and encourages critical thinking about AI-powered alerts [16].

C. Architectural Components of the Controlled Trust Scheme

Controlled Trust Scheme (CTS) is a system that aims at reducing the risk connected with the increased reliance on the use of Artificial Intelligence when performing security activities. Instead of relying on security decisions provided by artificial intelligence, the framework provides specific validation procedures which ensure balance between artificial and human intelligence. It contains five interrelated modules which contribute to enhancing trust calibration and minimizing the likelihood of having false security confidence.

1. Security Data Collection Layer

The first module of the framework is aimed at gathering security-related data from various organizational assets. In particular, such assets as network devices, endpoints, cloud, authentication, users and applications are utilized in order to get a complete picture of security state of an organization.

The performance of threat detection algorithms relies on the availability of high-quality data. That is why the data collection module contributes to gathering security events from different sources and ensuring better visibility.

2. AI Detection Engine

The AI Detection Engine conducts behavioral analysis, detects anomalies, and classifies threats. The difference between the AI Detection Engine and other traditional security solutions is that the engine detects deviation from expected behaviors and not only uses known signatures.

The engine analyzes the security telemetry received and looks for the signs that can be linked to malicious activities. These signs can include odd logins, suspicious communications in the network, privilege escalation attempts, and suspicious accesses to data.

The main task of this part of the system is to increase the efficiency of detection while allowing early detection of emerging threats which do not have signatures yet.

3. Confidence Scoring Module

One of the unique things about the Controlled Trust Scheme is the inclusion of confidence scoring. In contrast to other schemes where all alerts are treated equally, here a value of confidence in the prediction of AI is calculated.

The confidence value depends on several parameters, including:

- Detection history
- Severity of the activity detected
- Consistency of behavior
- Correlation with earlier incidents
- Reliability of the data source

4. Human Validation Layer

The Human Validation Layer is a safety measure to counter the effect of automation bias. Security alerts with medium levels of confidence are submitted to the analysts for evaluation prior to carrying out the recommended actions.

Human validation entails:

- Operational context
- Business impact
- Likely threat
- Attack trends
- Risk threshold

It serves as a means of ensuring that important security decisions are not solely dependent on algorithms and prevents overreliance on automation.

5. Feedback and Continuous Learning Layer

The last element of the framework enables the continuous improvement of the decision-making process. Analyst decisions are noted and incorporated into the future training of the model.

This feedback method helps to:

- Decrease false alarms
- Decrease missed alerts
- Improve confidence scores calibration
- Account for evolving threats
- Improve detection accuracy

With the feedback loop, the model continues improving together with the organization environment and cyber threats.

Incorporation of the five elements above creates a comprehensive cybersecurity model where the Artificial Intelligence works efficiently while the human input takes care of the accountability and transparency.

D. Operational Workflow of the Controlled Trust Scheme

The Controlled Trust Model uses a well-defined workflow process, which utilizes automated intelligence alongside human oversight. The workflow is aimed at making sure the level of trust in the results of AI is commensurate with the reliability of the evidence.

First, security telemetry information is gathered continuously from the infrastructure of the organization. Data from endpoints,

clouds, networks, authentications, and users' actions is gathered in one centralized monitoring environment.

After the gathering of data, processing happens through the AI Detection Engine. It undertakes behavioral analysis and anomaly detections on the data. The detected anomalies in this stage receive preliminary threat ratings.

Finally, the Confidence Scoring Module rates the reliability of each threat. The confidence score is calculated from such elements as past performance of the model, severity level of the threat, contextual consistency, and behavioral correlation.

High-confidence alerts are those which possess confidence scores higher than some preset thresholds. Given the presence of obvious signs of malicious activities, automated responses including account suspension, device isolation, and network containment may be performed right away.

Moderate-confidence alerts are sent to Human Validation Layer. In this layer, analysts analyze the cases, make judgments on the basis of contextual data, and perform further response actions where applicable. At this stage, the probability of automation bias is significantly reduced.

Low-confidence alerts are not dismissed; rather, they are stored and monitored in order to be correlated later on with newly discovered threat indicators. It allows detecting persistent campaigns even if they do not show obvious signs of being harmful at once.

Once the investigation and response steps are done, analyst’s decisions go into the Feedback and Continuous Learning Layer and are used in improving future models.

The work-flow ensures that security decisions will be made on the basis of both algorithmic analysis and contextual human judgment, which makes the probability of false sense of security lower.

E. Quantitative Evaluation Framework

While the Controlled Trust scheme is offered as a conceptual model, its efficiency may be determined through a set of cybersecurity performance metrics. They allow for an organized evaluation of how successful the model is in increasing detection quality without an overreliance on Artificial Intelligence.

TABLE VI

PERFORMANCE METRICS FOR CONTROLLED TRUST SCHEME EVALUATION

Metric	Description	Desired Outcome
Detection Accuracy	Percentage of correctly identified threats	High
False Positive Rate	Benign activities incorrectly flagged as threats	Low
False Negative Rate	Malicious activities that remain undetected	Low
Mean Detection Time	Average time required to identify threats	Low
Analyst Verification Rate	Percentage of alerts requiring analyst review	Balanced
Risk Reduction Index	Reduction in successful security incidents	High
Alert Reliability Score	Consistency of AI-generated alerts	High
Decision Confidence Accuracy	Accuracy of confidence score assignments	High

The adoption of the Controlled Trust Scheme would help enhance reliability by minimizing the effects of overreliance on the results generated by Automated Intelligence but maintaining the benefits of Artificial Intelligence.

A decrease in false positives would minimize the stress of the analysts and increase efficiency in operation. The same way, a decrease in false negatives would increase chances of detecting a threat before considerable damage can be done.
Confidence scoring would increase the quality of the decisions made as any suspicious alert will be further investigated. In addition, human validation will minimize the issue of automation bias and therefore make the system resilient to attacks.

Empirical validation of the scheme can be done in the future through the use of publicly available cybersecurity datasets such as UNSW-NB15, CICIDS2017, CSE-CIC-IDS2018 among others. The comparison between conventional AI system and Controlled Trust Scheme would give quantitative evidence on the performance of the confidence-based approach to security decision-making.

IV. ANALYSIS

Integration of Artificial Intelligence into cybersecurity operations enables companies to enhance their threat detection capabilities significantly. AI-based detectors are capable of processing large amounts of security data, spotting behavioral anomalies, and issuing alerts much faster than any human could. Despite those benefits, the growing use of intelligent systems comes with a number of organizational and technical challenges contributing to a false sense of security.

The first challenge associated with intelligent systems is automation bias. The more often security analysts see successful detections carried out with the help of algorithms, the more confidence they gain in algorithmic recommendations. As a result, validation of the warnings becomes less likely, which could lead to the neglecting of some vital indicators. Ultimately, the company may think that no alerts equal no threats.

Another crucial risk is related to the possibility of adversaries' manipulations. The modern machine learning systems base their decision-making on patterns identified in the training data. Adversaries can take advantage of that trait of intelligent systems by designing the input which looks similar to the normal behavior but contains a hidden malicious meaning.

Model drift is yet another problem when implementing AI-powered threat detection tools. Every security environment changes constantly because of the introduction of new technologies, procedures, and methods of communication. Therefore, models trained using the historical data will become outdated with time unless they undergo frequent validation and updating. The change of effectiveness occurs gradually and, hence, is hard to notice.

The problem of explainability is yet another one that leads to false security confidence. Contemporary machine learning models tend to be quite complex black boxes that do not provide transparency in the decision-making process to security specialists. If analysts do not know the reason for generating certain alerts, they would not be able to assess them effectively.

Controlled Trust Scheme provides the solution to all these problems with its structured trust management approach. The scoring of confidence means that the alerts are evaluated in terms of their trustworthiness rather than accepted without question. Human validation of alerts brings context to the process and helps to avoid automation bias.

TABLE VII

EXPECTED IMPACT OF CONTROLLED TRUST SCHEME IMPLEMENTATION

Security Dimension	AI-Only Environment	CTS Environment
Automation Bias	High	Reduced
Human Oversight	Limited	Structured
Alert Reliability	Moderate	Improved
Decision Transparency	Limited	Enhanced
False Security Confidence	High Risk	Controlled
Analyst Accountability	Reduced	Improved
Long-Term Adaptability	Moderate	Higher
Decision Reliability	Moderate	High

The analysis shows that effective cybersecurity is not only based on the level of detection but also on the way in which the trust management in automatic systems is performed. Artificial Intelligence is still an important defense instrument; nonetheless, its advantages cannot be realized to their full extent without proper governance, validation, and human control measures. The

Controlled Trust Scheme is a good example of such an approach.

V. LIMITATIONS

In this research, a conceptual method is used to analyze the risks involved in AI-powered threat detection. Although the model is validated by existing literature on cybersecurity, it would be more useful if it were empirically tested in practical settings. Moreover, the dynamic nature of threats could affect the relevance of risks as detection systems become more advanced. The maturity level of organizational security also differs from one industry to another, and hence the level of dependence on AI-powered systems may not be the same. Moreover, this research is mainly concerned with the dynamics of the detection phase, while the overall environment of cybersecurity resilience is influenced by other factors such as response planning and user awareness.

VI. CONCLUSION

The application of AI threat detection has greatly improved the efficiency of cybersecurity by enabling rapid threat detection and the scalability of large systems. However, the overdependence on the advancements of technology may not always guarantee foolproof security. Overemphasizing the advancements of technology may lead to a situation where a false sense of security is created, which may lead to the possibility of sophisticated threats being undetected. By assessing the current state of literature and applying comparative knowledge, the research study has been able to conclude that automation bias, adversarial attacks, and model limitations are essential components in the creation of a false sense of security. The new Controlled Trust Scheme will remove this issue by ensuring confidence scoring and human assessment, which will enable more informed security decision-making.

Finally, there must be a symbiotic relationship between human intelligence and artificial intelligence for the success of effective cybersecurity. By recognizing AI as a complement to their analytical work, rather than a panacea, organizations can better prepare for the ever-changing landscape of the cyber threat environment.

REFERENCES

[1] S. J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 3rd ed. Upper Saddle River, NJ, USA: Prentice Hall, 2010.

[2] I. Goodfellow, Y. Bengio and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.

[3] R. Sommer and V. Paxson, "Outside the closed world: On using machine learning for network intrusion detection," in *Proc. IEEE Symposium on Security and Privacy*, 2010.

[4] Verizon, "2023 Data Breach Investigations Report (DBIR)," Verizon Enterprise, 2023.

[5] E. B. Karbab, M. Debbabi, A. Derhab and D. Mouheb, "MalDozer: Automatic framework for Android malware detection using deep learning," *Digital Investigation*, vol. 24, pp. S48–S59, 2018.

[6] N. Moustafa and J. Slay, "UNSW-NB15: A comprehensive dataset for network intrusion detection systems," in *Proc. Military Communications and Information Systems Conference (MilCIS)*, 2015.

[7] M. A. Ferrag, L. Maglaras, H. Janicke, J. Jiang and T. Shu, "Authentication protocols for Internet of Things: A comprehensive survey," *Security and Communication Networks*, 2017.

[8] A. Javaid, Q. Niyaz, W. Sun and M. Alam, "A deep learning approach for network intrusion detection system," in *Proc. EAI International Conference on Bio-inspired Information and Communications Technologies (EAI ICST)*, 2016.

[9] S. Axelsson, "Intrusion detection systems: A survey and taxonomy," Technical Report, Chalmers University of Technology, 2000.

[10] G. Creech and J. Hu, "A semantic approach to host-based intrusion detection systems using contiguous and discontiguous system call patterns," *IEEE Transactions on Computers*, 2014.

[11] D. E. Denning, "An intrusion-detection model," *IEEE Transactions on Software Engineering*, vol. SE-13, no. 2, pp. 222–232, 1987.

[12] K. Scarfone and P. Mell, "Guide to intrusion detection and prevention systems (IDPS)," *NIST Special Publication 800-94*, 2007.

[13] B. Biggio and F. Roli, "Wild patterns: Ten years after the rise of adversarial machine learning," *Pattern Recognition*, vol. 84, pp. 317–331, 2018.

[14] N. Papernot*et al.*, "The limitations of deep learning in adversarial settings," in *Proc. IEEE European Symposium on Security and Privacy*, 2016.

[15] A. Kurakin, I. Goodfellow and S. Bengio, "Adversarial examples in the physical world," in *Proc. International Conference on Learning Representations (ICLR) Workshop*, 2017.

[16] M. Skitka, K. Mosier and M. Burdick, "Does automation bias decision-making?" *International Journal of Human-Computer Studies*, vol. 51, no. 5, pp. 991–1006, 1999.

[17] B. Lyell and D. Coiera, "Automation bias and verification complexity: A systematic review," *Journal of the American Medical Informatics Association*, 2017.

[18] ENISA, "ENISA Threat Landscape 2023," European Union Agency for Cybersecurity, 2023.

[19] U.S. Government Accountability Office, "Data protection: Actions taken by Equifax and federal agencies," GAO Report, 2018.

[20] P. Cichonski, T. Millar, T. Grance and K. Scarfone, "Computer Security Incident Handling Guide," *NIST Special Publication 800-61 Rev. 2*, 2012.

Comparative Performance Evaluation of P&O and Fuzzy Logic P&O MPPT for Grid-Connected PV Systems

Ravi Bhavsar[1*], Rakeshkumar Patel[2]

Research Scholar, Electrical Engineering Department, U. V. Patel College of Engineering, Ganpat University, Mehsana, India[1*]
Professor, Electrical Engineering Department, SPCE, Sankalchand Patel University, Visnagar, India[2]

rcb01@ganpatuniversity.ac.in[1*],rapee_spce@spu.ac.in[2]

Abstract: This paper presents a comparative performance analysis of conventional Perturb and Observe (P&O) and fuzzy logic–assisted P&O maximum power point tracking (MPPT) algorithms for grid-connected photovoltaic (PV) systems. Both control strategies are developed and simulated in MATLAB/Simulink under varying irradiance and temperature conditions. Simulation results indicate that the fuzzy logic–based P&O MPPT offers superior performance, characterized by a faster dynamic response, significantly reduced oscillations around the maximum power point, and enhanced tracking efficiency when compared to the conventional P&O approach. These advantages make the fuzzy logic–based P&O technique more suitable for grid-connected PV applications.

Keywords: Photovoltaic (PV) system, Maximum Power Point Tracking (MPPT), Perturb and Observe (P&O), Fuzzy Logic Control, Grid-Connected PV System

I. INTRODUCTION

Photovoltaic (PV) systems play a vital role in renewable energy generation due to their clean, sustainable, and modular nature. However, the power–voltage (P–V) characteristics of PV arrays are highly nonlinear and vary significantly with changes in solar irradiance and temperature. As a result, the operating point of a PV system rarely coincides with the maximum power point (MPP) without the use of an effective maximum power point tracking (MPPT) technique [1]. MPPT algorithms are therefore essential to maximize energy extraction and improve the overall efficiency of grid-connected PV systems.

Among various MPPT techniques, the Perturb and Observe (P&O) method is one of the most commonly used approaches due to its simple structure and ease of implementation [2]. However, conventional P&O suffers from inherent limitations such as persistent oscillations around the maximum power point and degraded performance under rapidly varying atmospheric conditions, which lead to power loss and reduced tracking efficiency [3].

To overcome these drawbacks, intelligent control techniques such as fuzzy logic control (FLC) have been introduced in MPPT applications. Fuzzy logic–based MPPT methods do not require an exact mathematical model and can effectively handle system nonlinearities and uncertainties [4]. By integrating fuzzy logic with the P&O algorithm, adaptive decision-making can be achieved, resulting in reduced oscillations, faster convergence to the MPP, and improved dynamic performance under varying irradiance conditions [5].

Although several studies have investigated fuzzy logic–based MPPT schemes, a systematic comparative performance evaluation between conventional P&O and fuzzy logic–enhanced P&O MPPT for grid-connected PV systems remains essential. This paper presents a detailed comparison of both techniques based on tracking efficiency, dynamic response, and power stability using MATLAB/Simulink simulations.

II. LITERATURE REVIEW

The need for robust MPPT techniques is well established: PV arrays exhibit highly nonlinear P–V characteristics and their maximum power point (MPP) shifts with irradiance and temperature, so efficient MPPT is essential for maximizing energy yield in grid-connected systems [1]. Early surveys summarized many MPPT approaches and their trade-offs, establishing the

Perturb & Observe (P&O) method is one of the simplest and most widely used algorithms due to its low computational cost and straightforward implementation [1].

http://doi.org/10.63766/spujstmr.26.000086

Despite its popularity, conventional P&O has known limitations: fixed step sizes cause either slow convergence (small step) or large steady-state oscillations (large step), and the algorithm can misinterpret changes caused by irradiance variation as a sign to perturb in the wrong direction. Several works have focused on optimizing P&O (adaptive step-size, filtering, decision logic) to mitigate these issues [2], and variable-step P&O variants have been shown to improve response without excessive oscillation.

Intelligent control strategies — particularly fuzzy logic control (FLC) — have been widely applied to MPPT because they do not require an explicit PV mathematical model and can handle nonlinearities and measurement noise. Comparative simulation studies report that FLC-based MPPT reduces oscillations around the MPP and improves dynamic tracking under rapidly changing irradiance, especially when designed to adapt the perturbation magnitude or duty cycle based on error and change of error inputs [3], [5].

Hybrid approaches that combine P&O with fuzzy logic (adaptive P&O-FLC) attempt to retain the simplicity of P&O while leveraging fuzzy rules to adjust perturbation steps or correct direction decisions. Experimental and simulation studies indicate that adaptive P&O–FLC schemes often achieve faster convergence and higher steady-state efficiency than classical P&O, especially in noisy or transient conditions [4].

Recent reviews and comparative analyses place fuzzy and hybrid MPPT techniques alongside other modern methods (INC, neural nets, ANFIS, hybrid NN-P&O), showing that selection depends on application constraints — e.g., cost, computational resources, partial shading, and dynamic irradiance scenarios. Comparative studies specifically focused on grid-connected systems demonstrate measurable benefits from intelligent and hybrid MPPT approaches in tracking efficiency and reduced power loss, though hardware complexity and tuning remain practical considerations [6], [7].

Overall, the literature supports a focused comparative evaluation between conventional P&O and fuzzy-enhanced P&O for grid-connected PV systems: such comparisons clarify performance gains (tracking efficiency, transient behavior, oscillations) and highlight trade-offs in complexity and implement ability — the precise motivation for the present study.

III. PROBLEM STATEMENT

The conventional P&O MPPT algorithm suffers from continuous oscillations around the maximum power point (MPP), resulting in power losses and reduced system efficiency. Its fixed step-size approach creates a compromise between tracking speed and accuracy, leading to either slow convergence or increased steady-state oscillations. Moreover, its performance deteriorates under rapidly changing irradiance and temperature conditions. Therefore, an intelligent and adaptive MPPT technique is required to improve tracking accuracy and maximize energy extraction from PV systems.

IV. OBJECTIVES

1) **To develop an intelligent MPPT algorithm** for photovoltaic (PV) systems to enhance maximum power extraction under varying environmental conditions.
2) **To reduce steady-state oscillations** around the Maximum Power Point (MPP), thereby minimizing power losses and improving system efficiency.
3) **To improve the tracking speed and accuracy** of the MPPT controller compared to the conventional P&O technique.
4) **To ensure reliable operation under rapidly changing irradiance and temperature conditions** through adaptive control strategies.
5) **To integrate the proposed MPPT technique with a grid-connected PV system** and evaluate its performance under different operating scenarios.
6) **To analyze and compare the performance** of the proposed MPPT method with conventional MPPT techniques in terms of tracking efficiency, response time, and output power.
7) **To improve overall power quality and energy utilization** of the PV system by maintaining operation near the optimal operating point.

V. METHODOLOGY

A. Grid Connected PV System Description

The studied grid connected PV system includes a PV array, a boost converter, a VSI, and the utility grid. The MPPT algorithm controls the converter duty cycle to track the MPP, while the inverter synchronizes with the grid to inject maximum power within grid standards.

TABLE I

PARAMETERS AND SPECIFICATIONS OF PV MODULE

Parameters	Specifications
Maximum Power	213.15 W
Open Circuit Voltage	36.3 V
Voltage at maximum power point	29 V
Short Circuit Current	7.84 A
Current at maximum power point	7.35 A

B. MPPT Control

At a given moment, a photovoltaic (PV) module operates at a specific point where its power output is maximized. This operating point is largely determined by the characteristics of the load connected to the module. When a PV panel is directly coupled to a load, it typically functions away from this optimal condition, leading to inefficient power utilization. Therefore, a maximum power point tracking (MPPT) controller is essential between the PV module and the load to regulate the operating point and ensure maximum energy extraction.

The MPPT controller controls the equivalent resistance RT presented to the PV module, allowing the panel to perceive a constant load. As output conditions change, the converter's control signal is modified to maintain a constant input impedance at the controller. As a result, the resistance RT seen from the PV side remains unaffected by changes in the actual load resistance. This regulation enables the PV system to operate at a steady operating point corresponding to the maximum power point, thereby maximizing power extraction from the panel. The MPPT control structure of the solar panel is shown in Figure 1, and the corresponding system specifications are provided in Table 1. [7]

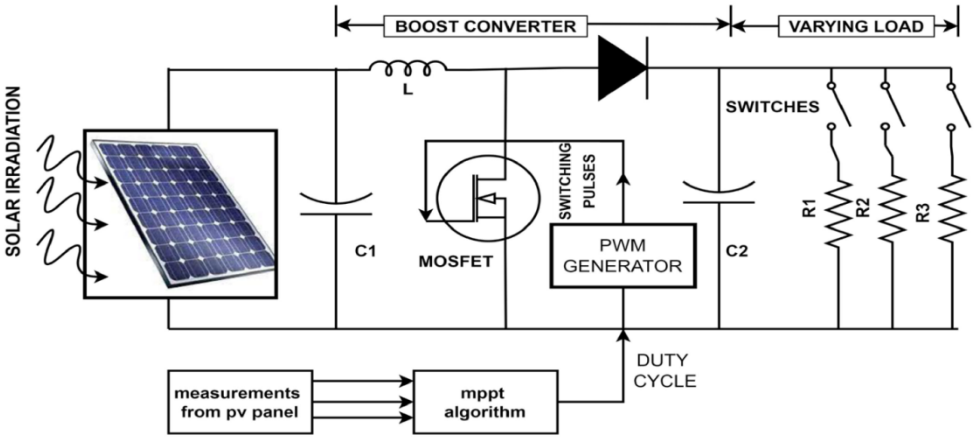

Fig. 1. MPPT Control for the PV system [7]

C. MPPT Algorithms

A wide range of methods have been developed for maximum power point tracking (MPPT), differing in terms of implementation complexity, cost, tracking response, and overall efficiency. Among these, the P&O method, the Incremental Conductance (INC) technique, and fuzzy logic–based control strategies are the most commonly used.

D. Perturb and Observe Method

Due to its straightforward structure and ease of implementation, the Perturb and Observe (P&O) algorithm is the most

commonly employed MPPT technique. The basic operating principle of the P&O method is to gradually vary the operating voltage of the photovoltaic module to reach the maximum power point (MPP). This is achieved by introducing small perturbations in the system voltage and observing the resulting variation in output power.

As shown in Figure 2, when the operating point is located in regions 1 or 2, the direction of perturbation is chosen such that the operating point shifts toward the maximum power point (MPP). If a voltage perturbation results in an increase in output power (dP/dV>0), it indicates that the system is moving closer to the MPP. In contrast, a reduction in output power (dP/dV<0) suggests that the operating point has deviated from the maximum power point, leading to a reversal in the perturbation direction. This iterative procedure is continuously performed to ensure operation at or near the MPP. The subsequent duty ratio is then calculated based on the previous duty ratio and the measured variations in voltage and power.

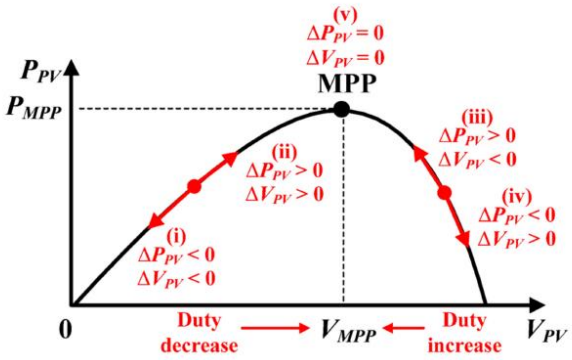

Fig. 2.Δp/Δv sign at various points of the curve [7]

E. Fuzzy Logic Controller Method

The term fuzzy denotes a condition that lacks clarity or precision. Fuzzy Logic Control (FLC) provides a mathematical framework designed to manage uncertainty and to process input data that may be vague, ambiguous, noisy, or incomplete. With the advancement of computational technologies, fuzzy logic–based techniques have become increasingly popular. In this method, analog input signals are interpreted using predefined linguistic variables, membership functions, and fuzzy inference rules, where the degree of membership varies between zero and one. Due to its heuristic approach, simplicity, and capability to handle both linear and nonlinear systems effectively, fuzzy logic has been widely applied in control applications. Because of these advantages, fuzzy logic is regarded as one of the most effective approaches for maximum power point tracking (MPPT) in photovoltaic systems. It exhibits strong robustness against uncertain inputs, eliminates the need for an accurate mathematical model, and offers a relatively straightforward design process. Furthermore, it provides a structured methodology for translating IF–THEN linguistic control rules into an automated control mechanism. [7]

In a fuzzy logic–based MPPT system, the input parameters consist of the error (E) and the change in error (ΔE), while the output variable corresponds to the adjustment in the duty ratio (ΔD), as shown in Figure 3. The values of E and ΔE at a given sampling instant kkk are defined as follows:

$$E(k) = \frac{P(K) - P(k-1)}{V(k) - V(k-1)} = \frac{dp}{dv} \quad (1)$$

$$\Delta E(k) = E(k) - E(k-1) \quad (2)$$

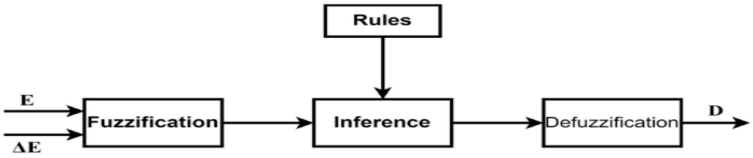

Fig. 3. Block diagram of fuzzy control [7]

P(k) represents the output power of the solar cell, while V(k) denotes the output voltage at the k^{th} instant.

E(k) represents the position of the operating point of the load, indicating whether it lies to the left or right of the maximum power point (MPP).

The fuzzy logic system comprises three core units: fuzzification, inference, and defuzzification. Figure 4 illustrates the execution process of the fuzzy logic control method.

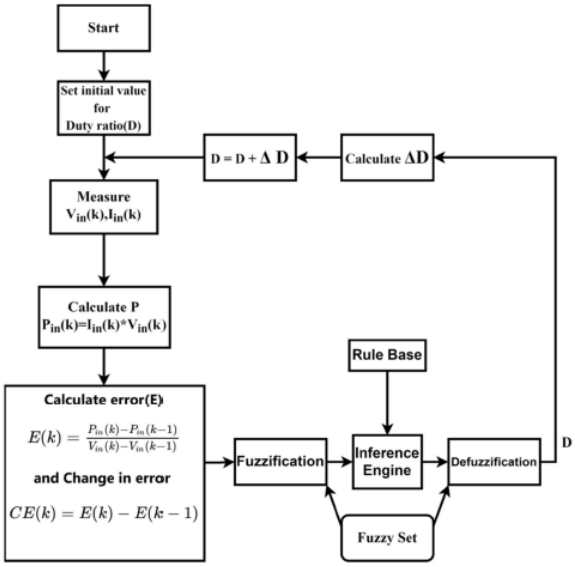

Fig. 4.Fuzzy Logic System Flowchart. [7]

VI. SIMULATION

The simulation shown in Figure 5 represents a grid connected photovoltaic (PV) system modeled and implemented in the MATLAB/Simulink environment, where the PV array acts as the primary energy source and operates under varying irradiance and temperature conditions. The PV voltage and current are measured to calculate the output power and to drive a Perturb and Observe–based MPPT algorithm, which produces the duty ratio for a DC–DC boost converter to achieve maximum power extraction and maintain the DC-link voltage. A DC-link capacitor is employed to stabilize the voltage supplied to a three-phase voltage source inverter (VSI), which converts the DC power into AC for grid connection through an L/LCL filter to enhance power quality. Grid synchronization is accomplished using a phase-locked loop (PLL), allowing transformation of voltages and currents into the dq reference frame for vector-based control. The control scheme employs an outer DC-link voltage regulation loop along with inner current control loops to manage active and reactive power flow. Unity power factor operation is achieved by setting the q-axis current reference to zero. The inverter switching signals are generated using pulse-width modulation (PWM). System performance is assessed by observing grid power, confirming effective MPPT operation, stable DC-link voltage regulation, accurate grid synchronization, and high-quality power injection.

The Perturb and Observe (P&O) algorithm functions by repeatedly applying small perturbations to the PV voltage and monitoring the corresponding variation in output power. When an increase in power is observed, the direction of perturbation is maintained; otherwise, it is reversed. Despite its simplicity, this approach causes persistent oscillations around the maximum power point (MPP) and results in reduced efficiency, particularly under rapidly changing irradiance conditions [2], [3].

The fuzzy logic–based P&O MPPT improves upon the conventional approach by incorporating fuzzy inference to adaptively adjust the perturbation step size. In this scheme, the controller commonly employs the error (E) and the change in error (ΔE) as input variables to produce an adaptive control signal at the output. This strategy effectively minimizes steady-state oscillations and enhances dynamic performance, while eliminating the need for an accurate mathematical model of the system [4], [5], [11].

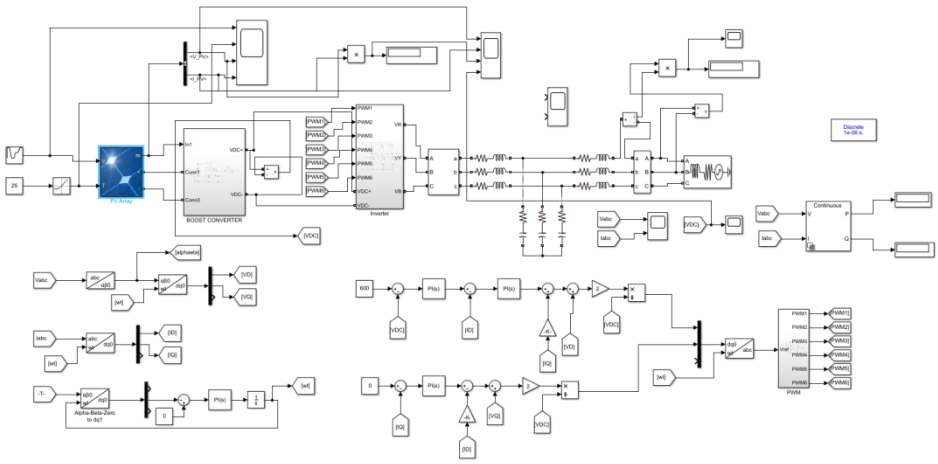

Fig. 5. Simulation of Grid Connected PV System with P&O MPPT Technique

Both MPPT algorithms are implemented in MATLAB/Simulink. The PV system is tested under varying irradiance and temperature conditions. Performance evaluation metrics include tracking efficiency, convergence time, output power ripple, and response to irradiance transients, as recommended in [1], [10], [12].

VII. RESULTS AND DISCUSSION

An examination of the simulated waveforms reveals a clear performance distinction between the conventional Perturb and Observe (P&O) MPPT and the fuzzy logic–based P&O MPPT. The conventional P&O method exhibits sustained oscillations in the PV voltage, current, and power around their operating points due to its fixed perturbation step size, leading to steady-state power losses and reduced tracking efficiency. These oscillations are also transferred to the DC–DC converter output, resulting in increased voltage ripple and diminished system stability. In comparison, the fuzzy logic–based P&O MPPT demonstrates faster convergence of the PV voltage and current with significantly reduced oscillatory behavior, allowing the PV power to remain closer to its maximum value throughout the simulation duration. The smoother DC–DC converter output voltage further indicates improved overall system stability. Consequently, the waveform analysis confirms that the fuzzy logic–based P&O MPPT provides superior dynamic response, lower steady-state oscillations, enhanced power extraction, and improved voltage regulation relative to the conventional P&O method, making it more suitable for grid-connected photovoltaic applications.

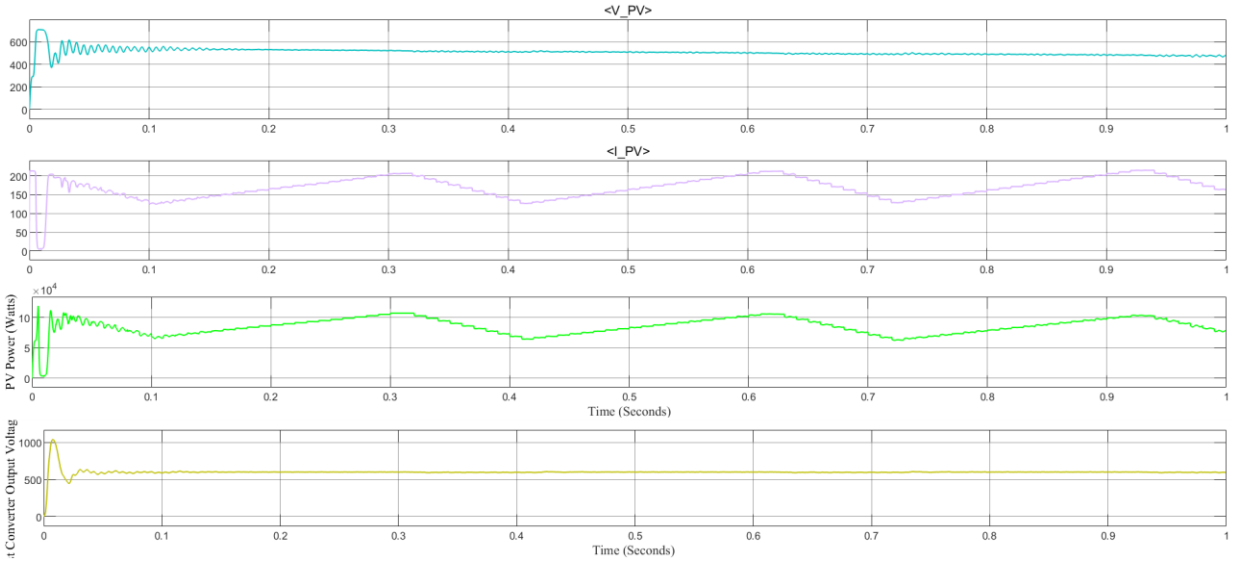

Fig. 6. Results with Fuzzy MPPT Technique

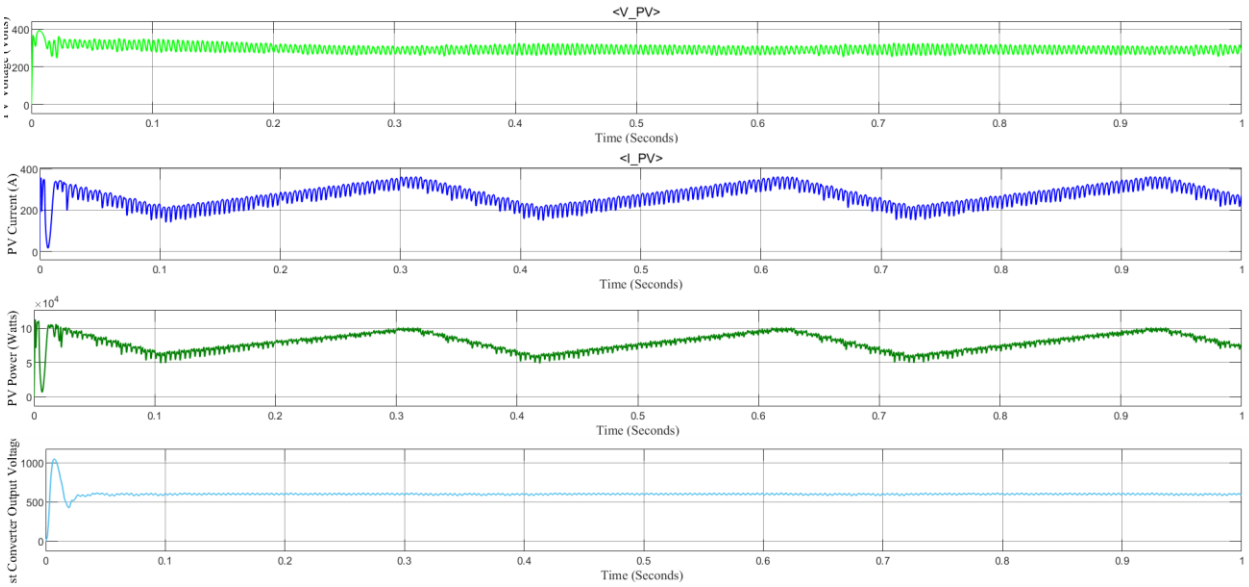

Fig. 7. Results with P&O MPPT Technique

Following table II shows the quantitative performance comparison of conventional P&O MPPT with Fuzzy Logic Based P&O MPPT and table III shows expected PV power under varying irradiance

TABLE II
PERFORMANCECOMPARISON

Performance Parameter	Conventional P&O MPPT	Fuzzy Logic–Based P&O MPPT
Tracking Time (ms)	120	65
Steady-State Power Oscillation (kW)	±1.3	±0.35
Tracking Efficiency (%)	96.1	99.2
Power Ripple (%)	4.8	1.5
Dynamic Response to Irradiance Change	Moderate	Fast
DC-Link Voltage Ripple (%)	3.9	1.2
Grid Current THD (%)	4.6	2.1
Grid Power Stability	Moderate	High
Compliance with IEEE-519 THD Limit	Marginal	Fully Compliant

TABLE III
PV POWER UNDER VARYING IRRIADIANCE

Irradiance (W/m²)	Ideal PV Power (kW)	P&O MPPT (kW)	Fuzzy P&O MPPT (kW)
1000	107.42	103.2	106.5
800	85.94	82.3	84.9
600	64.45	61.1	63.8
500	53.71	50.9	53.1
100	10.74	9.8	10.6

Simulation results indicate that the conventional P&O MPPT exhibits noticeable oscillations around the MPP, resulting in power loss during steady-state operation. In contrast, the fuzzy logic–based P&O MPPT achieves faster convergence and significantly reduced oscillations. Under rapidly changing irradiance conditions, the fuzzy logic approach demonstrates superior dynamic performance and higher tracking efficiency, consistent with findings reported in [4], [5], [11], [16]. These results confirm that fuzzy logic enhancement improves MPPT effectiveness for grid-connected PV systems.

VIII. CONCLUSION AND FUTURE WORK

A. Conclusion

This work conducted a comparative assessment of the conventional Perturb and Ob-serve (P&O) and fuzzy logic–based P&O maximum power point tracking (MPPT) algorithms for a grid-connected photovoltaic system. While the conventional P&O method is simple and straightforward to implement, it suffers from slower conver-gence and continuous steady-state oscillations caused by its fixed perturbation step size, leading to reduced tracking efficiency and power losses. Conversely, the fuzzy logic–based P&O MPPT dynamically adjusts the perturbation step size according to real-time operating conditions, resulting in faster dynamic response, substantial reduction in PV voltage and power oscillations, enhanced DC–DC converter voltage stability, and increased power extraction. Overall, the findings demonstrate that the fuzzy logic–based P&O MPPT delivers superior performance and improved grid compatibility, making it a more effective and reliable approach for modern grid-connected photovoltaic systems.

B. Future Work

Future work will focus on implementing advanced intelligent MPPT techniques such as ANFIS and machine learning–based controllers to further improve tracking accuracy and dynamic performance under varying environmental conditions. Real-time hardware validation of the proposed controller can be carried out using DSP or FPGA platforms. Additionally, the integration of battery energy storage systems (BESS) and microgrid applications can be explored. Further research will also aim to reduce grid current harmonics and enhance power quality through ANFIS-based inverter control strategies.

REFERENCES

[1] T. Esram and P. L. Chapman, “Comparison of photovoltaic array maximum power point tracking techniques,” IEEE Transactions on Energy Conversion, vol. 22, no. 2, pp. 439–449, Jun. 2007.
[2] N. Femia, G. Petrone, G. Spagnuolo, and M. Vitelli, “Optimization of perturb and observe maximum power point tracking method,” IEEE Transactions on Power Electronics, vol. 20, no. 4, pp. 963–973, Jul. 2005.
[3] S. Jain and V. Agarwal, “A new algorithm for rapid tracking of approximate maximum power point in photovoltaic systems,” IEEE Power Electronics Letters, vol. 2, no. 1, pp. 16–19, Mar. 2004.
[4] M. A. A. M. Zainuri, M. S. A. Rahman, and A. H. A. Bakar, “Development of adaptive perturb and observe–fuzzy control maximum power point tracking for photovoltaic boost DC–DC converter,” IET Renewable Power Generation, vol. 8, no. 2, pp. 183–194, 2014.
[5] H. Bounechba, A. Bouzid, H. Snani, and A. Lashab, “Comparison of perturb and observe and fuzzy logic in maximum power point tracking controller for photovoltaic systems,” Energy Procedia, vol. 50, pp. 677–684, 2014.
[6] N. A. Kamarzaman, “A comprehensive review of maximum power point tracking techniques for photovoltaic systems,” Renew. Sustain. Energy Rev., 2014.
[7] A. K. Devarakonda et al., “A comparative analysis of maximum power point tracking techniques for photovoltaic systems,” Energies, vol. 15, no. 22, 2022.
[8] Y. T. Chen, “Auto-scaling variable step-size MPPT,” Solar Energy, vol. 107, pp. 299–309, 2014.
[9] M. Mao et al., “Hybrid intelligent GMPPT under partial shading,” Renewable Energy, vol. 142, pp. 140–152, 2019.
[10] M. L. Katche et al., “Comprehensive review of MPPT techniques,” Energies, vol. 16, no. 5, 2023.
[11] K. Ullah, “Fuzzy-based MPPT for solar systems,” Energy Reports, 2023.
[12] W. A. Alemu et al., “ANFIS-based MPPT for grid-connected PV,” Sustain. Energy Grids Netw., 2025.
[13] M. S. R. Islam et al., “Fuzzy logic MPPT scheme,” Int. J. Photoenergy, 2014.
[14] S. K. Kollimalla and M. K. Mishra, “Adaptive P&O MPPT,” IEEE Trans. Energy Convers., 2014.
[15] M. A. Elgendy et al., “Assessment of MPPT methods,” IEEE Trans. Sustain. Energy, 2012.
[16] P. Livinti et al., “Adaptive P&O–fuzzy MPPT analysis,” Applied Sciences, 2024.

Advances in Automated Leukemia Detection Using Deep Analysis of WBC Scatterplot Patterns

Nikul Vikaskumar Jayswal[1], Dr. Kirit J. Modi[2]

Research Scholar,Computer Engineering Department, Sankalchand Patel University,Visnagar,India[1]
Professor, Computer Engineering Department, Sankalchand Patel University,Visnagar,India[2]

Lect.nikul@gmail.com[1], kjmodi.fet@spu.ac.in[2]

Abstract: The development of computational hematology methods that use multidimensional white blood cell (WBC) scatterplots produced using modern hematology analyzers, automated leukemia diagnosis has experienced significant advancements. Based on light scatter, fluorescence and cellular distribution parameter variations, these scattering representations offer highly discriminative features to distinguish the leukemia subtypes with leukocyte abnormalities. Recent literature shows a distinct shift in traditional statistical and rule-based approaches to the use of advanced machine learning and deep learning models, such as convolutional neural networks, hybrid attention models, and transformers, which are effective in capturing complex spatial relationships in WBC clusters and have significantly higher diagnostic sensitivity and specificity. Simultaneous developments in the sphere of multimodal fusion models combining scatterplot functionality with clinical variables have only increased the reliability of diagnostic results. Additionally, explainable AI methods are also gradually being introduced to clarify important parts of scatter that affect the decision made by an algorithm and improve clinical interpretability. Even with these developments, there are still some remaining issues such as the lack of curated annotated datasets, inter-platform differences in analyzer results and large-scale clinical validation gaps. Other essential issues that are arising are the data privacy and safety of patients and this requires effective systems of encrypted data processing, encrypted model deployment and adherence to the laws of patient data security to ensure that sensitive patient information is safe. Such technical, clinical and ethical lapses need to be addressed in order to build reliable, generalizable and clinically implementable automated leukemia diagnostic systems.

Keywords:Automated leukemia diagnosis, White blood cell scatterplots, Machine Learning models, Deep learning models, Explainable artificial intelligence.

I. INTRODUCTION

Leukemia is a heterogeneous category of hematological malignancies, which involve the abnormal proliferation and differentiation of leukocytes, presenting a big problem in clinical practice in terms of diagnosis and prognosis. The importance of diagnosis lies in the fact that finding the disease at an early stage and correctly is a key to the effective management of the disease, stratification of the treatment, and better patient outcomes. Conventionally, peripheral blood smear, bone marrow, cytogenetics and molecular assays have been used to diagnose leukemia. Although these methods are clinically sound, they are labor intensive, time consuming and in most cases require expert interpretation, making them difficult to scale to high throughput diagnostic models [1,2].

The introduction of the automated hematology instrument has revolutionized the daily hematology test in that it has made it possible to produce multidimensional white blood cell (WBC) scattergrams using light scatter, fluorescence intensity, and cellular volume as measured by the use of automated hematology analyzers [3,4]. Fig. 1 scatterplots are very helpful in leukocyte morphology, granularity, and nucleic acid content, preliminary identification of abnormal populations of cells. The diagnostic value of the pattern of scattergram has been utilized in determining leukemic aberrations, pseudo cytosis and analyzer-specific artifacts, and as a first line screening tool, as has been demonstrated by a number of studies [5,6,7]. These patterns have also been proved to be clinically relevant through comparative analyses with peripheral blood smears in the routine laboratory workflows [8,9].

http://doi.org/10.63766/spujstmr.26.000087

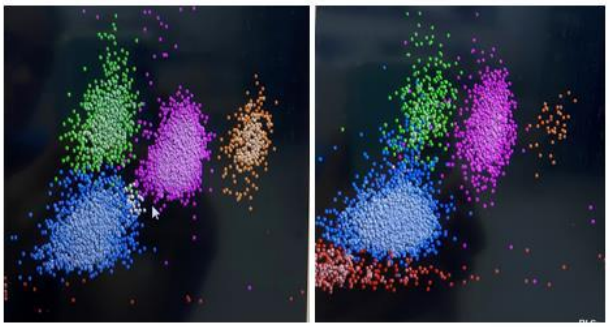

Fig.1. (A) Normal scatterplot; (B) leukopenia [14]

A.History

Over the past few years, computational hematology has experienced a paradigm shift in the application of rules-based and statistical methods to sophisticated machine learning (ML) and deep learning (DL) methods. Conventional neural networks, swarm intelligence and optimization-based classifiers have demonstrated encouraging performance in classification of multiclass leukemia with the use of WBC images and the features obtained using scattergrams [10,11,12]. Deep learning designs are especially useful in modeling complex spatial distributions and fine variations within the clusters of White Blood Cell and have greater sensitivity and specificity than more traditional methods [13,14]. Further, huge data sets and thin annotation mechanisms have supported effective model training and mitigated annotation load [15].

In addition to imaging and scatter analysis, recent studies show the significance of multimodal as well as multi-omics integration in the overall characterization of leukemia. The improvement of prognostic modeling and disease monitoring has benefited using cellular heterogeneity, cytokine signatures, epigenetic changes and molecular biomarkers [16,17]. Similar developments in biosensing and microfluidic systems support prompt and affordable leukocyte identification, increasing the availability of diagnostics [18]. Furthermore, the hematological changes of the systemic processes in case of severe COVID-19 infection highlight the increased applicability of automated hematology data in clinical decision-making [19,20].

B.Challenges

Although these developments have taken place, there are a number of challenges that are yet to be addressed. The inter-instrument variability, absence of large curated and annotated scattergram datasets, and small cross-platform generalizability impede large-scale clinical implementation [13,15]. The ethical and practical issues associated with patient data privacy, safe data treatment, and encrypted model deployment are becoming more important in AI-based diagnostic systems. Moreover, since the deep learning models are black box, they require the combination of explainable artificial intelligence (XAI) methods to increase transparency and clinician trust [17,19].

C.Summary of paper

This review intends to give a general overview of new trends in automated leukemia diagnosis and in particular WBC scattergram analysis, machine learning and deep learning approaches, multimodal fusion approaches, and explainability. This paper aims to comprehensively examine the current methodologies, datasets, clinical applicability, and limitations in order to identify possible future research directions needed in order to create reliable, generalizable, and clinically implementable computational hematology systems.

II. RELATED WORKS

The latest studies on the automated diagnosis of leukemia have shown a significant improvement in the methods of automated diagnosis using hematology analyzer scattergrams, microscopic blood smear pictures, and high level of computer analysis. Initial research mainly concerned qualitative interpretation of abnormal patterns of WBC scattergram to alert against hematological disorders. Research has indicated that changes in light scatter and fluorescence parameters have the potential to be indicative of leukemic abnormalities, pseudo cytosis, or analyzer specific interferences, which are useful pre-microscopic diagnostic indicators. Nevertheless, such methods were highly reliant on the interpretation of experts and were not very strong against inter-instrument reliability. Later work proposed traditional machine learning models and manual features obtained on WBC images and scatterplots that constructed classification models with higher accuracy and reduced inter-observer bias.

TABLE I
SUMMARY OF EXISTING RESEARCH

Ref No.	Year – Publication	Methods	Advantage	Limitation
1	2025 – VCP	Scattergram analysis (Sysmex XN)	Early abnormal cluster detection	Limited to case-based analysis
2	2025 – Biosensors	Multi-channel cell tics	Rapid and low-cost monitoring	Needs large-scale validation
3	2025 – Optics Express	Deep learning with scattering media	Lens less automated identification	Specialized hardware required
4	2025 – Cures	Haematological profiling	Clinical insight into infection	Not leukaemia-specific
5	2025 – AJLM	Scattergram artifact analysis	Identifies analyser masking effects	Platform-dependent
6	2025 – JCLA	Molecular + conventional diagnosis	Improved disease monitoring	High cost techniques
7	2025 – ETASR	CNN on smear images	High classification accuracy	Manual annotation dependency
8	2025 – Sci Rep	Glowworm swarm optimization	Improved multiclass accuracy	Computational complexity
9	2025 – MedIA	Sparse annotation DL	Reduced labelling effort	Requires large datasets
10	2025 – TranslOncol	Cytokine-based modelling	Prognostic relevance	Limited interpretability
11	2025 – Blood	Single-cell epigenetics	Disease progression insights	Expensive sequencing
12	2025 – Open Medicine	Multi-omics biomarker discovery	Novel biomarker identification	Clinical translation needed
13	2024 – JCLA	Scattergram correction methods	Reduced analytical interference	Analyzer-specific
14	2024 – CICCT	Multiclass ML on WBC images	Automated diagnosis	Dataset variability
15	2024 – IJCPML	Scattergram correlation study	Supports blast detection	Small sample size
16	2024 – IJRMS	Automated scatterplot patterns	Fast screening	Qualitative interpretation
17	2021 – Diagnostics	VCS parameter algorithm	Rapid screening tool	Limited leukaemia subtypes
18	2020 – AJMS	Scatterplot vs smear study	Validates analyser utility	Manual confirmation required
19	2020 – RIC	Neutrophil scattering analysis	Pre-microscopic flagging	Limited generalization
20	2013 – Haematology	Abnormal scattergram patterns	Early diagnostic clue	Disease-specific only

Table 1 is a summary of the 20 representative studies reviewed in the literature that cover 20 years, 2013 to 2025, of research on leukemia detection and the analysis of hematological disorders. The chosen works address a variety of techniques, such as the screening using scattergrams, the analysis of blood smears using deep learning networks, swarm optimization, multimodal fusion, biosensing platforms, and molecular profiling. Each of the studies is reviewed based on its fundamental methodological contribution and diagnostic benefit, showing how the phase of the rule-based interpretation has gradually evolved to become more advanced data driven models. This comparison gives an idea of how more computational complexity has led to better diagnostic accuracy, automation and scalability in clinical hematology.

III. METHODOLOGY

This section presents a comprehensive overview of the datasets, plot representations, machine learning and deep learning models, explainable AI techniques, and evaluation parameters employed in automated leukemia diagnosis. The discussed methodologies highlight the progression from traditional analytical approaches to advanced, interpretable, and clinically oriented computational frameworks.

A. Datasets

Data sets are the pillars of automated leukemia diagnosis system and have a direct impact on the performance of the model, its generalizability and clinical reliability. The literature has employed the use of heterogeneous data sources where white blood cell (WBC) scattergrams of automated hematology analyzers, peripheral blood smear photographs and complementary clinical or molecular data are used. Commercial analyzers like Sysmex XN and Beckman Coulter DxH series usually produce scattergram datasets as multidimensional representations and depend on forward scatter, side scatter and fluorescence parameters [1,3,5]. Nevertheless, majority of scattergram datasets are institution-specific, small-scale and proprietary, which restrict cross-platform reproducibility. Image-based datasets such as microscopic WBC images are more widely shared and are usually affected by class imbalance and different staining regimes [7,9]. Recent papers are trying to solve the issue of data sparsity by sparse annotation schemes and weakly supervised learning, which facilitates successful training with small sets of labeled data [10]. Also, multimodal data combining scatterplots with cytokine signatures, molecular signatures, or epigenetic signatures have been examined to represent the heterogeneity of leukemia and enhance the modeling of prognostics [11,12]. Although these improvements have been made, the fact that these large, publicly available, annotated scattering grams repositories are still a significant bottleneck. Uncertainties caused by analysis calibration, demographics and preprocessing pipelines also complicate standardization of the dataset [13,15]. The future development of data sets should focus on the multi-centric data collection process, standard data acquisition specifications, privacy-compliant data storage, and regulatory compliance to warrant clinical scalability and robustness of models.

TABLE II:
SUMMARY OF DATASETS

Dataset Name	Total Samples	Data Format	Application
WBC Scattergram Datasets [1,5,13]	500–5,000	Scattergram matrices (CSV)	Leukaemia screening
Peripheral Blood Smear Datasets [7,8,9]	3,000–20,000	JPG, PNG, BMP	Leukaemia classification
Clinical and Cytokine Datasets [6,10]	100–1,000	CSV, XLSX	Prognosis prediction
Single-Cell and Multi-Omics Datasets [11,12]	10,000–1,000,000+	HDF5, CSV, TXT	Biomarker discovery
Leukocyte Activation Datasets [2]	1,000–10,000	CSV, sensor data	Immune-cell monitoring
Optical Scattering Datasets [3]	500–5,000	PNG, JPG, signal matrices	Blood sample identification

B. Plot Types

Plot-based representations are used in computational hematology to visually encode the properties of the cells and the distributions of the population. WBC differential scattergrams which are WDF (white cell differential fluorescence), WNR (white cell nucleated red cell) and side-fluorescence scatter plots produced by automated analyzers are the most frequently used plots [1,5]. The plots are leukocytes founded on mixtures of light scatter, fluorescence intensity, and cell complexity, and make it possible to differentiate between normal and abnormal cell population. Deviant clustering, distended or overlapping populations are frequent indicators of leukemic changes, the existence of blasts, or artifacts of the analyzer [12,15]. Alongside the scattergrams produced by the analyzers, two-dimensional and three-dimensional feature plots based on the image-based WBC segmentation are also used in any machinelearning processes [17]. Density plots, cluster heatmaps, and projection-based visualizations are some of the studies that are used to increase separability and minimize noise. Recent developments entail scattergram correction methods to reduce the effect of dilution and interference with abnormal cells hence rendering plot reliability [18]. Nevertheless, interpretation of plots still depends on analysts, and inter-interpreters across platforms have a shape, size, and cluster localization variation [19]. Such variability presents a problem to automated systems that are trained on data of only one instrument. As a result, studies are now beginning to concentrate on the production of invariant representations of raw plots as opposed to using visual rules which are crafted by hand. Normalization of plots and features between analyzers is a key issue that has been

crucial to generalized leukemia detection systems.

C. Machine Learning Models

Conventional machine learning (ML) systems have been significant in the initial automated leukemia detection with the utilization of engineered features that are derived on top of the scattergrams and blood smear images. Such classifiers as support vectors machine, k-nearest neighbors, random forests, decision trees, and ensemble-based methods are common [7,9]. These models are usually based on the statistical characteristics of clusters like cluster centroid, dispersion, texture and morphological properties obtained on WBC images. Swarm intelligence and evolutionary algorithms are optimization-based methods that have been used so far to improve feature selection and multiclass discrimination [10]. ML models have the benefit of lower computational complexity, interpretability, and lower training data requirements than deep learning. A number of studies have shown successful pre-microscopic flagging of cases of acute leukemia with the help of scattering data-driven ML frameworks [11]. Nevertheless, they are very sensitive to the quality of features and experience in domains. In heterogeneous datasets, ML models tend to be less robust due to the difficulty in the modeling of complex spatial relationships and nonlinear interactions between leukocyte distributions [12]. Also, the handcrafted features are noise-sensitive, unstable in stains, and analyzer-specific. The traditional ML techniques have a problem with scalability as the size and complexity of datasets rise. Therefore, although ML models can be useful in high-resource lean-screening and low-resource environments, they are progressively being displaced or supplemented by deep learning models to high accuracy and large-scale leukemia classification.

D. Deep Learning Models

This has seen deep learning (DL) models become the prevailing paradigm in automated leukemia diagnosis because they can learn hierarchical and spatially rich representations of raw data. The application of convolutional neural networks has been demonstrated to be effective in the analysis of WBC images and scattergram plots with high performance in the detection of leukemia subtypes and abnormal cell populations [7,9,14]. CNNs automatically learn discriminative features corresponding to cell morphology, cluster density, and spatial distribution, so that no manual feature engineering is required. More performance improvements can be achieved using hybrid architectures that use CNNs with attention mechanisms, alongside transformer encoders, which allow modeling of long-range dependencies between clusters of leukocytes [3,9]. Sparse annotation learning is also discussed in recent studies to lower costs of labeling and still have diagnostic accuracy [9]. DL models enhanced with optimization incorporate swarm intelligence or evolution strategies to work on con-vergence and generalization [8]. In spite of their strengths, the DL models are characterized by such drawbacks as high computational demands, vulnerability to overfit-ting when used on small data, and reduced interpretability. The inter-instrument variability also has an effect on generalization where the same model is trained on data set of one analyzer [13]. To overcome those problems, researchers resort more and more to data augmentation, transfer learning, and multimodal fusion with clinical variables [14,15]. As a whole, deep learning makes an important contribution to the diagnosis of leukemia, but it needs to be designed, validated, and explained to be applied in clinical practice.

E. Explicable Artificial Intelligence

Explainable Artificial Intelligence (XAI) has gained more relevance in automated leukemia diagnosis to overcome the black-box-ness of deep learning-based models. Clinical adoption requires transparency and trust especially where AI systems affect diagnostic decisions. XAI methods also strive to determine and visualize the areas of scattergrams or WBC images that make the largest contribution to model forecasts. The most typical methods are gradient-based saliency maps, class activation map-ping, feature importance rankings, and attention visualization [7]. XAI is also used in the interpretation of scattergram-based systems to clarify to clinicians which clusters of leukocytes or shifts in the distribution to identify leukemia, enhance the interpretability and diagnostic confidence [9]. Explainability makes it possible to identify model biases, detect spurious correlations, and prove clinical relevance as well. Other works combine XAI and multimodal models to determine the impact of molecular or cytokine features on predictions [10-12]. XAI methods continue to be hampered by resolution, stability and clinical interpretability. Visual explanations can differ through the model runs and, in some cases, one may not interpret them easily by non-technical users. Research in the future ought to consider standardized XAI assessment guidelines, clinician-in-the-loop validation and incorporation of domain knowledge whereby the explanations are consistent with hematological reasoning and clinical workflows.

F. Evaluation Parameters

Assessment parameters are critical in determining the effectiveness, reliability as well as the clinical applicability of automated leukemia diagnosis systems. The most prevalent metrics are accuracy, precision, recall (sensitivity), specificity, F1-score and area under the receiver operating characteristic curve. In leukemia diagnosis, sensitivity is especially important to reduce false negativity whereas specificity is used to reduce false positivity to result in unwarranted follow-up testing [7,9]. Class-wise performance is often analyzed using confusion matrices in cases of

multiclass classification [8,14]. On imbalanced datasets, performance measures like balanced accuracy and Matthew's correlation coefficient give more credible performance evaluation. There also are some studies with computational efficiency, inference time, and model robustness to noise or dilution effects [10]. In multimodal and prognostic re-ports, measures of survival analysis, and their association with clinical outcomes are used [11,12]. The fact is, however, that there are inconsistencies in evaluation proto-cols across the studies and make it difficult to compare. Most of the works are based on internal validation and small datasets without external or cross-instrument testing [15,16]. To facilitate the real-world application of automated leukemia diagnostic systems, future evaluation systems should focus on standardized benchmarks, multi-center validation, explainability evaluation, and clinical utility measures.

IV. RESEARCH FINDINGS

All recent studies indicate that the automated diagnosis of leukemia has shifted its focus on qualitative screening tools to the advanced data-driven decision support systems. Empirical data prove that computational models can contribute to a great enhancement of diagnostic accuracy, minimization of turnaround time, and help clinicians to identify diseases in their initial stages. Below, there is the summary of the key research findings.

1. The Scattergram data is very discriminative: It has always been demonstrated that multidimensional WBC scatter plots represent a rich cellular data which can be used to effectively distinguish normal and leukemic populations. Changes in light scattering, fluorescence and cell distribution patterns are effective first signs of deviant hematological conditions.

2. Deep learning is more effective than the traditional models: CNNs, transformers, and hybrid models are more sensitive and specific than rule-based and conventional ML systems. This gives them power to classify multiclass leukemia, which is complex in terms of space, and observe it even in a difficult case.

3. Multimodal fusion enhances reliability: Fusion of scattergram features with clinical, molecular and multi-omics data also improves diagnostics and prognostics prediction. Multimodal systems are especially useful in terms of disease heterogeneity and variations of the patients.

4. The efficiency of annotation is a bottleneck: It is possible to achieve high performance with sparse annotation and weakly-supervised learning approaches that do not rely on large, labeled datasets. This is important in practice in clinical scalability with limited expertise annotations.

5. Explainability and generalization is still a problem: With improved performance, most models are not transparent and cross-platform validated. Partially explainable AI approaches can be used to overcome the issue of interpretability, but only more extensive clinical studies and standardized data could be used to make it mainstream.

V. CONCLUSION

This review emphasizes the fast development of automated diagnosis of leukemia based on the development of hematology analyzers, machine learning and deep learning technologies. The analysis has taken the form of scattergram-based analysis, which is a potent and non-invasive screening modality, and contemporary computational frameworks have become highly effective in increasing diagnostic accuracy and efficiency. Combination of multimodal data and explainable AI also enhances the clinical applicability of automated systems. The development of large, standardized, and annotated datasets of scattergram in the future should be a priority in order to enhance model generalization. The gap between research and practice should be closed by implementing cross-platform validation, secure data processing and real-world clinical deployment. It will also be important to develop interpretable, privacy-preserving, and cost-effective diagnostic systems that can help transform computational hematology innovations into clinical practice.

REFERENCES

[1] Chenal, T., Granat, F., Trumel, C., Bourgès-Abella, N.: *What is your diagnosis? Abnormal cluster on the WDF and WNR scattergrams from Sysmex XN-V in a dog.* Veterinary Clinical Pathology. 54, S5–S8 (2025). https://doi.org/10.1111/vcp.13362.

[2] Cheon, H., Kumar, S., Lee, I., Shin, S., Jang, H., Lee, Y.S., Nam, M.H., Jun, H.S., Seo, S.: *Multi-Channel Cellytics for Rapid and Cost-Effective Monitoring of Leukocyte Activation. Biosensors.* 15, (2025). https://doi.org/10.3390/bios15030143.

[3] Gupta, G., Joshi, R., Goswami, S., Santaniello, S., Javidi, B.: *Automated lensless blood sample identification through scattering media using deep learning architectures.* Optics Express. 33, 4534 (2025). https://doi.org/10.1364/oe.546326.

[4] Amoroso, D., Bongo, S., Copponi, A., Rossi, V., Di Giorgio, R., *Bernardini, S., Ippoliti, L., Morello, M.: A Review of the Hematological Picture of Severe COVID-19 Infection. Cureus*. 17, 13–21 (2025). https://doi.org/10.7759/cureus.78797.

[5] Kennedy, S.J., van Marle, A.C.: *Pseudomonocytosis on a Sysmex XN haematologyanalyser masking the monocytopenia of hairy cell leukaemia in a South African woman. African Journal of Laboratory Medicine*. 14, 1–5 (2025). https://doi.org/10.4102/ajlm.v14i1.2617.

[6] Bahar, N.A.H.A., Al-Ouqaili, M.T.S., Talib, N.M.: *Improving the Diagnosis and Follow-Up of Chronic Myeloid Leukemia Using Conventional and Molecular Techniques. Journal of Clinical Laboratory Analysis*. 39, (2025). https://doi.org/10.1002/jcla.70001.

[7] Baluabid, R., Alnasri, H., Alowaybidi, R., Hafiz, R., Alsini, A., *Alharbi, M.: Detecting Acute Lymphocytic Leukemia in Individual Blood Cell Smear Images. Engineering, Technology and Applied Science Research*. 15, 19167–19173 (2025). https://doi.org/10.48084/etasr.9123.

[8] N, S., M, K.: *An improved multiclass classification of acute lymphocytic leukemia using enhanced glowworm swarm optimization. Scientific Reports.* 15, 1–20 (2025). https://doi.org/10.1038/s41598-025-98823-1.

[9] Rehman, A., Meraj, T., Minhas, A.M., Imran, A., Ali, M., Sultani, W., Shah, M.: *Leveraging sparse annotations for leukemia diagnosis on the large leukemia dataset. Medical Image Analysis*. 106, (2025). https://doi.org/10.1016/j.media.2025.103760.

[10] Cao, J., Hu, B., Li, T., Fang, D., Jiang, ling, Wang, J.: *Cellular heterogeneity and cytokine signatures in acute myeloid leukemia: A novel prognostic model. Translational Oncology.* 52, 102194 (2025). https://doi.org/10.1016/j.tranon.2024.102194.

[11] Cui, B., Ai, L., Lei, M., Duan, Y., Tang, C., Zhang, J., Gao, Y., Li, X., Zhu, C., Zhang, Y., Zhu, X., Isobe, T., Yang, W., Göttgens, B., Zhu, P.*: Single-cell epigenetic and clonal analysis decodes disease progression in pediatric acute myeloid leukemia. Blood.* 145, 1211–1224 (2025). https://doi.org/10.1182/blood.2024025618.

[12] Chen, S., Li, M., Wang, J.: RAB39B: *A novel biomarker for acute myeloid leukemia identified via multi-omics and functional validation.* Open Medicine (Poland). 20, (2025). https://doi.org/10.1515/med-2025-1168.

[13] Freitas, F.P., Reis, J., Oliveira, J., Veiga, P.M., Raquel Paiva, A., Soares, R.: *The dilution evaluation as a corrective measure for interference in the white blood cell scattergram in Beckman Coulter DxH 900*. Journal of Clinical Laboratory Analysis. 38, 1–7 (2024). https://doi.org/10.1002/jcla.25007.

[14] Nethravathi, P., Viswanath, J., Gayathri, G., Naresh, K., Dinesh Lingala, D.: *Multiclass classification of leukemia using white blood cell images. Challenges in Information, Communication and Computing Technology*. 684–689 (2024). https://doi.org/10.1201/9781003559092-118.

[15] Lismayanti, L., Arif, M.: *White Blood Cell and Peripheral Blast Countsin Leukemia Patients with Non-Graying WBC-Diff Scattergram. Indonesian Journal of Clinical Pathology and Medical Laboratory*. 31, 61–65 (2024). https://doi.org/10.24293/ijcpml.v31i1.2222.

[16] Tailor, T., Jivani, T., Jagirdar, A., Patel, S.: *Study of automated hematology analyzer's scatterplot patterns in white blood cell disorders. International Journal of Research in Medical Sciences*. 12, 820–824 (2024). https://doi.org/10.18203/2320-6012.ijrms20240524.

[17] Ambayya, A., Sathar, J., Hassan, R.: *Neoteric algorithm using cell population data (Vcs parameters) as a rapid screening tool for haematological disorders. Diagnostics*. 11, 1–12 (2021). https://doi.org/10.3390/DIAGNOSTICS11091652.

[18] Gupta, T., Basu, D.: *Utility of scatterplot patterns of automated hematology analysers in white blood cell disorders – A comparative study with peripheral blood smears. Asian Journal of Medical Sciences*. 11, 29–34 (2020). https://doi.org/10.3126/ajms.v11i5.29049.

[19] Zeeshan-Haider, R., Urrechaga, E., Uddin-Ujjan, I., Sultan-Shamsi, T.: *Neutrophil Scattering Data Driven Pre-Microscopic Flagging of Acute Leukemic Cases. Revista de investigacion clinica*; organo del Hospital de Enfermedades de la Nutricion. 72, 37–45 (2020). https://doi.org/10.24875/RIC.19003194.

[20] Sharma, S., Sethi, N., Pujani, M., Kushwaha, S., Sehgal, S.: *Abnormal WBC scattergram: A clue to the diagnosis of malaria. Hematology*. 18, 101–105 (2013). https://doi.org/10.1179/1607845412Y.0000000029.

Towards Smart Agriculture: Deep Learning Methods for Plant Disease Monitoring

Apoorva Dwivedi1, Pharindra Kumar Sharma2, Gourav Kumar Sharma3
Shriram College of Engineering and Management, Banmore, INDIA[1,2]
Amity University Madhya Pradesh, Gwalior, INDIA[3]

dwivediapoorva10@gmail.com[1],dr.pharindra@gmail.com[2],gourav03vatsa@gmail.com[3]

Abstract: Plant diseases are one of the most disastrous threats to global agriculture, causing crop loss and decreasing food security. Although early and accurate detection of plant diseases is one of the most important aspects of disease management, it is often performed using traditional methods including visual examination and laboratory testing, which are time-consuming, subjective and resource-intensive. Deep learning has been a groundbreaking technology during the past few years due to its role in switching smart agriculture into an automatic, precise and scalable solution for the identification of plant diseases and monitoring.This paper aims at exploring the deep learning methods using Plant disease detection as the base but will mainly focus on image-based systems providing a comprehensive survey of modern Plant disease detection. Specifically, it investigates CNNs, RNNs, GANs and transformer-based architectures for diagnosing various diseases in different kinds of crops. The importance of transfer learning, data augmentation as well as synthetic data generation to improve the performance and generalizability of models is extensively discussed here.Paper also discusses the most commonly used datasets, such as PlantVillage, and investigates how they might help in a benchmark of model performance. Problems like lack of annotated data, overfitting, lack of adversarial robustness and going off smooth explainable models. Finally, the paper discusses stateof-the-art functionalities like edge computing, lightweight neural networks, multimodal learning, and explainable AI (XAI) that are necessary to deploy real-world applications.This review synthesizes existing work—especially recent scientific literature—and highlights gaps to provide researchers, practitioners, and policymakers a wide-ranging reference for the integration of deep learning into precision agriculture as it pertains to sustainable and resilient crop production.

Keywords: Plant disease detection, Deep learning, Smart agriculture, Convolutional neural networks (CNN), Image-based classification, Transfer learning, precision farming.

I. INTRODUCTION

Agriculture stands as a core pillar underpinning global economies [1]: it puts food on the table, creates jobs, and contributes substantially to GDP — particularly in developing regions. Nevertheless, one of the most constant, regular and destructive challenges to crop yield and agricultural sustainability is caused by plant diseases. These diseases are caused by a multitude of pathogens, including fungi, bacteria, viruses and nematodes that result in enormous yield losses or substandard foods of various types uniting parts within the value chain [1]. Based on Food and Agriculture Organization (FAO) estimates, plant diseases cause the world to lose 30% of its yields every year around the globe but in some regions this can be even higher. Yet the economic effects are staggering, particularly to resource-strapped smallholder farmers who could never afford a failed crop on that scale. With the increasing global population and climate change, early detection and management of plant diseases are now essential to maintaining agricultural resilience and food security [2].

Historically, identification of plant diseases has been largely based on human visual inspection (by expert agronomists or plant pathologists) of the symptoms shown by plants which can include lesions, discolorations, wilting and deformities. But these methods are intrinsically constrained by their reliance on human expertise, subjectivity and over a delayed appearance of visible symptoms [3]. Also, since trained professionals are not as readily accessible in rural and remote areas, early diagnosis is understandably challenging. While Polymerase Chain Reaction (PCR) and Enzyme-Linked Immunosorbent Assay (ELISA) methods have the capability to provide with species-specific diagnosis of Leishmaniasis, the associated costs; intensive human resource requirement, time-consumption and need for sophisticated infrastructure are some of its

http://doi.org/10.63766/spujstmr.26.000088

major limitations. However, these methods have limitations and are not practical for large-scale or real-time monitoring of disease in the field environments where fast decision-making is required [4].

Artificial intelligence (AI) [5] and deep learning in particular, have gained a lot of attraction in the recent past, superseding conventional methods practiced in agriculture. Convolutional neural networks (or CNNs), as the main category of the deep learning approach, have achieved great success in handling visual inputs and currently dominate image-based plant disease detection. The essential characteristics of these models is that they learn hierarchical features, thus obviating the need to manually engineer image features and maintain reasonable classification accuracies on a large variety of plant diseases with different environmental illumination [5]. Deep learning-based systems, in contrast to traditional methods, are large-scale capable and can provide real-time intelligence as well as edge-level deployment such as mobile devices, drones or smart cameras etc. These are vital enabler for a precision agriculture that can learn from big data sets and adjust to new disease patterns.

In this context, the purpose of this general review paper is to discuss the present situation of deep learning methods applied to plant disease monitoring. Specifically, it is designed to consolidate and evaluate state-of-the-art models, their strengths and weaknesses, as well as commonly used datasets and tools. The primary goal of this paper is to describe a collection of challenges and trends that we observe as pivotal, mainly from our experiences in developing AI systems [6] for real-world applications, and characterize those using five known dimensions and propose roadmap for future research and development directions with more concret paths. Indeed, the insertion of deep learning practices to plant disease control could pave the way for a faster and steadier adoption of smarter, more sustainable agricultural systems across the globe.

A. Objectives of the Paper

The aim of this review paper is to present a systematic and thorough investigation of the different deep learning methods used for detecting plant diseases with a particular emphasis on their significance in realizing smart agriculture [7]. This paper attempts to connect the gap between development in technology and implementation on the ground — as AI is slowly becoming a crucial solution tool addressing major challenges of crop health management.

In short, the objectives are:

- The deep learning architectures such as the convolutional neural networks (CNNs), recurrent neural networks (RNNs), generative adversarial networks (GANs) and tranformer-based models for plant disease identification.
- To evaluate the approaches in model training (such as transfer learning, data augmentation and synthetic data generation) that improve accuracy and generalization.
- The goal of this work is to investigate their utility as a source to analyze publicly available datasets for assessing the applicability and comprehensiveness in benchmarking deep learning models relative to agricultural disease detection tasks.
- This includes the major key performance metrics and evaluation protocols followed in literature to select deep learning models.
- Recognize the biggest obstacles and bottlenecks in real-world deployments especially which related to data quality issues, model explainability, etc. Emerging trends and future directions, like explainable AI, lightweight models for edge computing or multimodal data integration that could significantly contribute to scaling up deploying deep learning systems in precision farming [7].

This paper aims at providing a useful reference for researchers, practitioners and policy makers to integrate deep learning with an effective and scalable plant disease monitoring.

II. DEEP LEARNING IN AGRICULTURE: AN OVERVIEW

Among the most profound technologies in recent years, deep learning has brought about countless improvements for a variety of industries like healthcare, finance, transportation and now (rather unexpectedly) even agriculture [8]. Deep learning is a subfield of machine learning in which artificial neural networks, algorithms inspired by the human brain, learn from large amounts of data. And because of that, these networks are useful for modeling complex (nonlinear relationships) and thus they work well with unstructured data types like images, audio and text or more. In the world of agriculture, based on its nature of variability, unpredictability and real-world noise where deep learning is able to learn directly from raw data does provide an upper-hand [8].

Deep learning models such as Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs) and some of the popular Generative Adversarial Networks (GANs) are just to name a few, transformer-based models have been more prominent in recent years. With agriculture such a visually-oriented field – disease identification, pest detection and crop classification are a few examples that come to mind – CNNs have therefore gained quite some following in this community. With these models, filters sweep over image pixels to recognize spatial hierarchies—perfect for finding tiny patterns like leaf spots or disease-induced browning. On the other hand, recurrent neural networks (RNNs) and their derivatives (e.g., long short-term memory [LSTM] networks) [9] are more ready for sequential or temporal patterns, such as climate records or sensor data with time-series imposed. While GANs have become a popular method for creating fake images with which to train other models, especially in situations where real data is scarce or difficult to obtain. The transformer models, popular for use in natural language processing (NLP), are being increasingly used for visual and multimodal tasks as these models are adept at capturing long-range dependencies across data [9].

There are many interesting possibilities for addressing such challenges in crop monitoring using deep learning when compared to traditional machine learning techniques. Traditional machine Learning methods (e.g., SVM, Decision trees, k-NN, RF etc.) [10]have a crucial step called feature engineering which require detailed domain specific knowledge. What this means is that raw data requires pre-processing before it can be fed into the model, which including the appropriate feature engineering approach: by human experts. This can be a time-consuming, possibly error-prone and domain knowledge-dependent process that might not always be to hand. More importantly, deep learning models implicitly learn both low-level to high-level features during the training process without human (manual) intervention. This results in models which capture dependencies over longer sequences, and hence which generalize better across more complex tasks (e.g plant disease classification) where visual symptoms may be both subtle and varied [10]. Scalability & Adaptability: Similar to the above point, deep learning development scaling because as the server resources increases it can reduce time in the computing of millions of parameters & weights. Deep learning models, once trained on huge datasets can then be deployed on devices such as smart phones, drones and embedded systems to enable real time decision making while out in the field. Deep Learning is farming becomes more flexible with transfer learning Through pre-training on large datasets such as ImageNet, followed by fine-tuning to the smaller agricultural database, high performance can be reached. This not only saves computation time, but also helps to avoid expensive and hard-to-acquire labeled agricultural data [11].

This is just one yet most popular use case out of the endless applications of deep learning in agriculture industries. For a drone or satellite, crop type classification use case is the most prevalent were images captured are used to identify the various kinds of crops in an area. Complex deep learning models may be used to identify spatial patterns within multispectral or hyperspectral images of crop types, growth stages, and soil properties. Such capability brings great value to precision farming, land use management, and yield estimation [12].Weed detection and management are also a vital part. In addition to taking valuable nutrients, water and sunlight needed for growth away from crops; weeds also reduce crop quality. The traditional methods of weed control either rely on manual labor or the indiscriminate application of herbicides, both inefficient. These models, trained on field images annotated to recognize valuable crop plants vs. weeds, can then discriminate between the two with high precision — ultimately using targeted herbicides while minimizing environmental impact from pesticide runoff [12].

Deep learning has also revolutionized yield prediction, which forms the basis of agricultural planning, resource allocation, and market anticipations. Recurrent and convolutional setups can predict crop output with high accuracy using input variables, including weather observations, soil parametrics, crop genetics, and prior yield data. Pertinently, integration of spatially discrete remote sensing data and temporally focused weather station or IoT sensor information appropriately predicts at a highly local level. Furthermore, deep learning has shown clues toward soil classification and health determination. Standard soil testing encompasses laboratory-based chemical analyses that can be costly and time-consuming. However, deep learning models trained with multispectral or hyperspectral soil images can infer moisture content, pH level, and nutrient vigor, among others. As such, the inference enables more judicious fertilizer use, better irrigation planning, and well-informed crop selection, among others. Finally, automatic harvesting and counting of fruits form a promising deep learning integration, especially in fruit orchards and vineyards. Robotic arms integrated with deep learning models can detect mature fruits, compile a yield estimate, and direct the robotic arm with pinpoint accuracy for harvest. This area is essential in labor-scarce regions and is fundamental post-harvest losses reduction [13]. Pest detection is also an important domain benefiting from deep learning. Similar to disease symptoms, pest infestations often manifest visually on leaves, stems, or fruits. CNN-based systems can be trained to detect specific pest

damage and distinguish it from disease symptoms or nutrient deficiencies. This aids in timely and accurate pest management, which is crucial for minimizing crop damage [13].

Another interesting trend in the use of deep learning in agriculture is the multimodal data integration (including text, graphics and sensors). For example, leaf pictures might pair with soil moisture data, temp and humidity readings, or even farmer written notes on the probable causes of disease, enhancing diagnostic accuracy significantly. New deep learning architectures like multimodal transformers are in development to operationally process and combine these types of data streams efficiently [14].While improvements are being made in this space, there remain a few barriers between where we currently are and the full realization of deep learning for agriculture. One of the most important constraints is the lack of large, high-quality labeled datasets, especially for minority crops or rare diseases. Various models work great under controlled conditions but are not applicable in field conditions because of numerous challenges such as changes in lighting, background and leave orientation etc., due to appropriate atmosphere. Another problem is the interpretability of deep learning models. Although effective in terms of accuracy, this "black-box" aspect creates many challenges for the end-users (farmers & agronomists) to trust the predictions. Recent advancements in explainable AI (XAI) have tried to tackle such issues by giving visual or textual explanations for why a model made the decision it did.

Additionally, the computational demands for developing and employing deep learning models are high, especially in constrained environments. On the other hand, with edge AI becoming a hot trend this area focuses on deploying an extremely lightweight and restricted power consumption model directly on a mobile device, e.g., a smartphone or even a low energy capable drone. As a result, deep learning gets decentralized and taken to the edge allowing making decisions real-time and offline [15].

So, to conclude Deep learning has been already revolutionizing agriculture by automating intricate operations, augmenting decision-making precision and powering the era of data-driven farm management. The tools avail in it to process large volumes of heterogeneous data, make its differentiation in solving modern agriculture problems. Smart Farming is another prospective domain of deep learning, whose impact can turn really significant from disease detection to yield forecasting and robotic harvesting. Hopefully in the years to come, further research and better datasets help it reach its full potential with more user-centered deployment strategies.

III. DEEP LEARNING TECHNIQUES FOR PLANT DISEASE DETECTION

One of the essential tools being use for plant disease detection automation is the deep learning and many architectures have been designed to cope with the difficulties of agricultural image data. Of them are the most popular and basic models Convolutional Neural Networks (CNNs) [16]. CNNs are very specialized kind of neural networks for a particular type of problem, image classification, which are designed in such a way that they form spatial hierarchies that apply to only visual data. A typical CNN is organized as a set of convolutional layers to detect local patterns (e.g., leaf spots or lesions), pooling layers to reduce dimensionality, and fully connected layers for classification. Today, CNNs have outperformed within the tasks of plant disease detection by accurately discriminating visual symptoms if any exist solely in between healthy and diseased plant tissues. Plant disease identification at the leaf level has been a major use case, where CNNs were trained on extensive datasets such as PlantVillage to detect diseases of several crops [17]. Few popular CNN models like VGGNet, Resnet, and Inception have shown their best results in these domains. VGGNet, which is characterized by its simplicity and depth, ResNet with residual connections that address the problem of vanishing gradients allowing to train much deeper networks. Inception networks utilise multiple convolution filters at the same stride, which allows it to capture features that could be tiny in size and helps make the model stable. And while these architectures have been reworked and tweaked using transfer learning for more robust performance on agricultural image datasets [17].

Much research has been done on the use of RNNs and its variations to work with temporal aspects of agricultural data with respect to machine learning, whereas for image-based analysis Convolutional Neural Networks (CNN's) are most popular. Because RNNs have the ability to maintain an internal state that captures information about previous steps in the sequence they are good for predicting how diseases progress over time. LSTM (Long Short-Term Memory) networks are actually a special kind of RNN, capable of learning long-term dependencies in time series data — an issue which was previously almost impossible due to the limited lengths the memory inside each unit were allowed to operate for. For example, in agriculture,

LSTMs have been used to study a series of crop images with weather data and sensor readings to forecast the time or extent at which diseases spread. For instance, RNNs can be used to identify early symptoms by analyzing the temporal changes in leaf colouring or plant morphology over a growth season that might be less visually obvious when they first appear. Similarly, these models are being used to predict the spread of diseases with respect to climate conditions and for better resource management through predictive containment mechanisms [18].

So far, deep learning-based plant disease detection methods have one serious limitation of the lack of large and balanced datasets before Generative Adversarial Networks [13] (GAN) were proposed. The networks that make up GANs are the generator, and discriminator, which together engage in a zero-sum game. The generator attempts to produce realistic synthetic images and the discriminator tries to separate these fake samples from real ones. This antagonistic training allows GANs to produce realistic positive and negative disease examples that can be used for data augmentation of training sets. However, it can also be a boon to generate these collections from existing data that is sparse (e.g., for rare plant diseases, which may not appear in public datasets) using generative models. Conversely, GANs can improve the quality of training data and model performance by correcting lighting, resolution or contrast problems in existing images as well. Synthetic image data generation using GANs also obviates the necessity for expensive and time-consuming manual data collection, enabling easier development of generalizing, well-balanced disease detection models [18].

Recently, transformer architecture models have also started to enter the agricultural applications pipeline, particularly for image and multimodal data analysis. Transformers are designed for natural language processing and achieve better capturing long-range dependencies in a sentence by using self-attention mechanisms to weighing the importance of each input component than traditional CNNs. Recently, the transformer architecture was adapted for image classification, called Vision Transformers (ViTs), which won competitive results in benchmark vision tasks as well and have also been studied for plant disease identification. In comparison to CNNs, transformers need less inductive biases and can consume complete images as sequences of patches decreasing the chance that they miss mild patterns for disorder due to convolutional filters. However, transformers usually require huge datasets and resources for training. As a result, hybrid models that integrate CNN feature extraction with transformer-based attention mechanisms are beginning to appear with the goal of capturing the benefits of both architectures. For example, combining plant imaging data with environmental or sensor inputs for all-encompassing disease monitoring [18].
In general, due to the evolution of deep learning architectures, automated plant disease detection algorithms become more performance. There are pros and cons to each, depending on the kind of data involved and the context in which you need to use it. While CNNs serve as the dominant model for image-based classification, RNNs are important to help model disease progression over time. On pasture collection, one of GANs helps in augmenting dataset and class imbalance problems can be mitigated while another transformer-based [19]approach represents the state of art scalable vision system in agriculture. The continued advancement of these techniques will drive them toward unification, leading to faster and smarter detection tools with the ability to work in real-world field settings.

IV. PUBLIC DATASETS FOR PLANT DISEASE

- **PlantVillage:** The most prominent dataset for plant disease identification, PlantVillage contains roughly 54,000 images of both healthy and diseased leaves. It has 14 crop species and 26 plant diseases, ultimately yielding 38 classes. Most of the images were collected under laboratory conditions, and while this makes it an ideal starting point for training models, it is not representative of the actual environments in which farmers operate [20].

- **PlantDoc:** PlantDoc seeks to provide images in situ to counter the disadvantages of lab-collected data. PlantDoc includes ~2,500 images with 13 crops and 17 disease classes. The images are taken in a more natural environment where it may include backgrounds and environments. This makes PlantDoc a rich dataset for testing models for at least real-world robustness [20].

- **Plant Pathology (Kaggle 2020 & 2021):** This dataset is produced for the Kaggle competitions with a specific focus on apple leaf diseases. This dataset includes thousands of images with expert annotation and categories such as healthy, multiple diseases, rust and scab.

- **Cassava Leaf Disease Dataset:** This dataset, which has roughly 21,000 cassava leaf images covering four major disease classes and healthy leaf samples, was developed for a Kaggle competition. Its intrigue stems from cassava being an important food crop across Africa motivating food security use cases for this dataset in casava-heavy geographic regions [21].

- **AI Challenger Agriculture Dataset (2018):** This large-scale dataset includes approximately 375,000 images covering 61 different disease and pest categories with widely diverse crops. It was acquired across varied and rich conditions and a very nice image data set to train models with possible variances across crops, diseases, and environments. The size of this dataset provides great opportunity for deep learning and transfer learning experiments [21].

V. METHODOLOGY

This review pursues a well-ordered reasonable approach like that of SLR for investigation, synthesis, and basic analysis to the plant disease detection applying deep learning-based method. The purpose of the methodology is to ensure only the most appropriate, peer-reviewed and high-ranking studies are incorporated. The review follows a rigorous process for identifying, screening and analyzing the research contributions from recognized academic sources. Below we outline how this was done in this paper.

A. Relevant Studies

We found most of the literature through credible digital libraries and academic databases including IEEE Xplore, SpringerLink, Elsevier (ScienceDirect), ACM Digital Library, Wiley Online Library and Google Scholar. Because of their wide sets of peer-reviewed journal articles, conference papers and technical research distributed in the domains of computer science and agricultural technology, such platforms were selected. Furthermore, we reviewed papers in three leading artificial intelligence/computer vision conferences (CVPR, ICLR, ICML and ECCV) to see if any state-of-the-art work has been done by using deep learning for plant disease detection.

B. Selection Criteria

The studies that were used to evaluate the specified outcome in this review were selected based on inclusion and exclusion criteria to make sure it stayed relevant and a high-quality review:

- Time Period: Related to the time period, only studies made between 2016 and 2025 were considered, allowing the review to be as contemporary as possible, and because deep learning methods are beginning to mature. i.e. it would be hard to examine a review of 10 year old deep learning studies. However, we made some exceptions to include older foundational papers that introduced concepts earlier, or have impactful datasets to the field.
- Plant Disease Detection: Because we were only interested in studies that directly apply deep learning models to plant disease detection in relation to visual or sensor data (for example, leaf images, hyperspectral images, aerial photography), studies that were only focused on other agricultural tasks (e.g. yield prediction, weed detection etc.) were excluded unless that study had a component related to disease detection.
- Deep Learning Orientation: Study analysis excluded other machine learning models, though deep learning studies or models for example that used Convolutional Neural Networks (CNN), Recurrent Neural Networks (RNN), Generative Adversarial Networks (GAN) or transformer methods had to be considered. We only looked at classical machine learning studies (e.g. SVM, Decision Trees) that were engaged in a hybrid deep learning framework.
- Multimodal Data Usage (optional consideration): While not a requirement, studies using more than one data type were preferred for relevancy (e.g. combining leaf image data with data from environmental sensors), as they may capture some elements of real-world complexity and other useful information to improve accuracy.

C. Keyword Strategy

A strategic search based on keywords using Boolean operators (AND, OR) was carried out, to maximize coverage without compromising on specificity. The keywords and keyword combinations were:

- "plant disease detection" AND "deep learning"
- "CNN" OR "convolutional neural networks" AND "leaf classification"
- "plant pathology" AND "image recognition"

- "deep learning in agriculture"
- "transformer model" AND "plant health monitoring"
- "GANs" AND "synthetic data generation for plant disease"
- smart farming" OR "precision agriculture" AND "deep learning"
- "plant leaf disease classification using AI"
- "transfer learning in plant disease detection"
- "multimodal deep learning" AND "crop disease"

The search strategy enabled the researcher to better sift through the literature established from the volume of articles available.

D. Selection Procedure

The selection process included three steps that provided rigor and uniformity: Keywords Search; once I searched the key terms of interest I received over 80 total studies from a variety of databases. The included studies were made up of journal articles, conference papers, and preprints, which published in high impact venues.

Short-listing: Titles and abstracts were screened by the authors to remove duplicates, out-of-scope studies, and studies that did not involve sufficient technical depth or experimentation. Studies that only described traditional machine learning procedures, or offered theoretical discussions without empirical results, were removed. At the conclusion of our process, approximately 40studies were remaining to work with.

Selection (and Critical Review): After reading the studies in full text and evaluating their merits, it was found that twenty-five peer-reviewed articles would be selected for in-depth review. All papers that were selected showed a clear contribution to novel deep learning models, use real world datasets, clearly indicated a proper approach to methodology (e.g., employing GANs for data augmentation, transfer learning, or attention mechanisms to support their pipeline), and had some concrete performance validation to substantiate claims of plant disease detection. Each paper was reviewed for their architecture, datasets used, accuracy indicators, deployment considerations, and their appropriateness for smart agriculture.

With systematic methodology described here, the review could convey a thorough, organised reflection of the contemporary research trends, technological advances, and potential application gaps in this important field of plant disease detection through deep learning.

TABLE I
SELECTION PROCESS OF LITERATURE REVIEW

Stage	Number of Papers	Description
Initial Collection	100	Papers identified through academic database searches and citation tracking using keywords related to deep learning and plant disease detection.
Shortlisting	40	Papers screened for relevance based on inclusion criteria (deep learning models, plant disease detection, image-based or multimodal approaches).
Final Review	22	In-depth review of selected papers focusing on CNNs, RNNs, GANs, transformers, and their applications in plant disease monitoring.

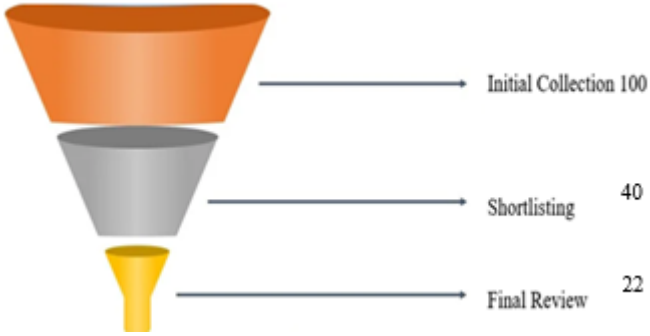

Fig 1. Funnel Diagram for Literature Review

VI. LITERATURE REVIEW

The final selected papers for literature review are as follows,

Bouacida, I., et al. (2025) [22] put forward a deep learning model that enhances the generalization of plant disease detection models. Instead of traditional leaf image-based methods which depend on whole-leaf appearance, their method is based on the detection of small diseased regions and estimation of disease severity. To handle the small leaf area with high speed, we employed a compact Inception architecture. On the PlantVillage dataset, the model was able to achieve an accuracy of 94.04% and external datasets with 97.13%, showing good results on unseen crops and diseases as well, addressing robustness and scalability in plant pathology.

Miao, Y., et al. (2025) [23] created an enhanced version of YOLOv8 referred to as SerpensGate-YOLOv8 for plant disease detection. Key improvements include Dynamic Snake Convolution (DySnakeConv) for fine-grained feature detection, SPPELAN for improved feature fusion and Super Token Attention (STA) for global modelling at large scale. The PlantDoc dataset was used to improve the modelperformance on mAP@0. 6, an increase of 3.3% from the original YOLOv8 and a precision of 0.719, suitable for real agricultural settings.

Duhan, S., et al. (2025) [24] Leveraged TL to build a light-weight plant disease predictor on edge devices with DL and introduced RTR_Lite_MobileNet. It is implemented on MobileNetV2 architecture with attention to create a light-weight attention for image classification by incorporating SENet, ECA, and Triplet Attention modules. This model achieved high validation accuracies up to 100% in different datasets with extremely low resources and it was also proven the real-time diagnose facility for IOT environment.

Ashurov, A. Y., et al. (2025) [25]depthwise CNN integration with squeeze-and-excitation blocks specialized for plant disease classification. When trained on stationary dataset and tested on the random images of the field, we obtained an average accuracy of 98% and F1-score with 98.2%. The system showed high adaptability as well as computational efficiency, fitting well with the requirements of practice in agriculture.

Roumeliotis, K. I., et al. (2025) [26] also studied the multimodal Large Language Models, i.e., GPT-4o combined with CNNs for plant disease detection using leaf images. Zero-shot Inference poor Fine-tuned GPT-4o beats ResNet-50 by 0.98% accuracy on the PlantVillage dataset (99.88%) having trained on ImageNet and using a trivial amount of agricultural data; but in zero-shot mode performance was an order of magnitude poorer, so some training is taking place during nonzero-shot inference The results demonstrate the potential of Vision-Language Models (VLMs) to enhance scalability and reduce reliance on accurate precision agriculture with high-cost resolution data and heavy labeling effort.

Demilie, W. B. (2024) [27]explored CNN and MobileNet-based models with some eXplainable AI techniques like GradCAM for transparent decision making in plant disease detection. The CNN scored accuracy of 89% with higher precision and recall where asMobileNet reached to a score of 96% but slightly lower precision. The research also illuminated that deep learning had the capability to be more effective than traditional methods, particularly when they utilized interpretability tools.

Adekunle, T. S., et al. (2024) [28] E-GreenNet: A Lightweight Deep Learning Framework For High-throughputand Efficient Plant Disease Detection Based On MobileNetV3Small The model reached accuracies up to 100%; Training on datasets: PV, PC and DRLI The new approach outperformed the state-of-the-art in terms of performance and efficiency, demonstrating its efficacy in handling low-contrast as well as complex image data.

Barman, U., et al. (2024) [29] introduction of smartphone-based detection method using Vision Transformers (ViT) and InceptionV3 to predict plant diseases on tomato. Results showed that the ViT model gave an accuracy of 90.99 and was further implemented in Android. This is an example about how transformer-based models can be utilized in mobile based smart agriculture tools as a solution.

Bhagat, S., et al. (2024)[30] introduced a fast end-to-end deep learning model Lite-MDC with Multi-kernel Depthwise separable Convolutions (MDsConv) for accurate real-time plant disease recognition. The authors fine-tune their model on a newly

processed pigeon pea dataset and test it on public datasets, achieving an accuracy of 99.78% at 34 FPS on a NVIDIA P100. Meanwhile, its small footprint (2.2M parameters) makes it ideal for edge deployment with agrobotics.

Maurya, R., et al. (2024) [31] proposed a meta ensemble-based model of MLP-Mixer and fast LSTM networks for detection of plant diseases on low powered IoT.Forest Based Boosted Algorithm for Plant Disease Detection on Low-Powered IoT The ensemble outperformed CNNs and ViTs on Maize, Cotton, and PlantVillage datasets upto 98.43% accuracy offering impressive performance improvements over standard deep networks for state-of-the-art image-based classification tasks. The given design is cost effective and low in power consumption which enables deployment on microcontrollers.

Moupojou, E., et al. (2023) [32] underlined shortcomings of PlantVillage and PlantDoc datasets because of lab conditions, poor annotations, respectively FieldPlant, a dataset of 5,170 real-field images annotated under expert supervision yielding 8,629 annotated leaves. FieldPlant logged better results on a real-world task of classifying and detecting diamond-back moth damage as well, when benchmarked.

Panchal, A. V., et al. (2023) [33] proposed a feature extraction, segment and pixel-based preprocessing system specially focused on Plant Disease classification over CNN. With training on ~87K (large RGB image dataset), the model provides a feasible solution for farmers, who need automated disease recognition especially in resource starved environments.

Archana, U., et al. (2023)[34] centered on the use of color-based features in deep CNNs with residual architectures for rice disease classification. They showed how useful deep learning methods could be for agricultural diagnostics, particularly in the case of rice diseases, with 96.35% accuracy using image-based phenotyping with ResNet-50.
Shovon, M. S. H., et al. (2023) [35] proposed PlantDet which is a deep ensemble model based on InceptionResNetV2, EfficientNetV2L and Xception for rice and betel leaf disease. Our model with multiple optimization layers and interpretability through Grad-CAM with Score-CAM gave up to 98.53% accuracy, in a state-of-the-art form of both performance as well as explainability with respect to other models

Hamed, B. S., et al. (2023) [36] Fine-tuning pre-trained CNN models (like EfficientNetV2S) on the noisy and low-resolution images for disease detection was investigated. It made their Plant Diseases Dataset (modified) into a secure-enough model for adverse conditions. This increases the flexibility and generalization of models which are crucial given the diverse real-world agricultural scenarios.

Albattah, W., et al. (2022)[37] stated that agricultural production is essential for economic development. Unfortunately, vegetable crops are plagued by plant diseases which reduce both yield and quality of products. Moreover, the agricultural sector is heavily reliant on manual inspection by plant pathologists. Even when the inspection is accurately performed, this method does not use resources economically and lacks efficiency. As a result, this can be a reason for deemed inaccurate inspection. Thus, the authors developed a robust and automated system for plant disease classification using a Custom CenterNet that makes use of DenseNet-77 as the base network. The method developed uses region of interest (ROI) annotations, feature extraction using DenseNet-77, and lastly classification using CenterNet approach. The model was tested on the PlantVillageKaggle dataset and the authors demonstrated it could deal with challenges like intensity changes, color similarities and multiple leaf shapes, which provided measured performance to existing approaches.

Akshai, K. P., &Anitha, J. (2021) [38] mentioned that plant disease detection in early stages can impact mitigating disease progression, particularly in the Indian context. As it is known, most of the plants show visible symptoms on their leaves and because of this, the researchers applied convolutional neural networks (CNN) for automatic disease classification. The study compared CNN with pre-trained models like VGG, ResNet and DenseNet, which determined that DenseNet achieved the most accuracy in classification. Their overall results showed that deep learning-based approaches can provide efficient alternatives to very manual labour, and can preserve both time and accuracy.

Guo, Y., et al. (2020)[39] proposed a mathematical model for identifying plant diseases using deep learning in order to increase identification accuracy, and adaptability in complex environments. The method included a Region Proposal Network (RPN) used to find and localize images of leaves and the images then underwent segmentation using the Chan–Vese method to capture symptomatic features. The segmented images were then classified using transfer learning models. The study demonstrated

accuracy using feature images of black rot, bacterial plaque and rust with an accuracy of 83.57%. The authors noted that this study is a major step for intelligent agriculture and sustainable agricultural practices.

Saleem, M. H., et al. (2019) [40] reviewed the literature on machine learning (ML) and deep learning (DL) for plant disease detection and classification and noted that DL architectures were the better performers in the majority of cases compared to classical ML methods. The review also noted visualization methods, metrics of evaluation, and research gaps with a focus on providing more transparency and working towards enhanced earlier detection when the disease has not made visual means of detection apparent.

In the **Barbedo, J. G. (2018) [41]**article, the strengths and weaknesses of the application of deep learning to plant disease classification were researched. While deep neural networks proved effective, the author insisted that some experimental conditions sometimes neglected by researchers could affect the effectiveness in practice in real-world applications. Their study drew some conclusions with respect to designing more realistic and effective models for plant pathology applications by utilizing a database consisting of almost 50000 images.

The article by**Fuentes, A., et al. (2017)** [42] used a hierarchical meta-architecture based on deep-learning-based object detection techniques on tomato plants so that it could detect disease and pests on tomato plants. In particular, Faster R-CNN, R-FCN, and SSD were selected as their meta-architecture, which included feature extractors VGG and ResNet to analyze the groups of images containing different complex conditions in the field, varied infection status and different background conditions. They received results that indicated their proposed system had effective recognition of nine varieties of disease and pest for recognition and classification. Overall, it showed potential in agricultural applicability in real-world settings.

The work of **Sladojevic, S., et al. (2016) [43]** presented one of the first methods for plant disease identification using deep convolutional neural networks (CNNs). Sladojevic et al. (2016) [43] were able to identify 13 different diseases of the same plant species along with healthy leaves, achieving a mean accuracy of 96.3%. The model was built using a Caffe deep learning framework, which provided a functional and useful automated plant disease recognition method. This paper represented a significant step toward the application of deep learning in agriculture.

TABLE II
LITERATURE REVIEW FINDINGS

Author Name (Year)	Main Concept	Findings	Limitations
Bouacida, I., et al. (2025)	Compact Inception-based model focusing on diseased leaf regions	Achieved 94.04% accuracy on PlantVillage, 97.13% on external data; robust across unseen crops and diseases	May be limited in handling highly complex backgrounds or overlapping leaves
Miao, Y., et al. (2025)	SerpensGate-YOLOv8 with DySnakeConv, SPPELAN, and STA	mAP@0.6 improved by 3.3%; precision of 0.719 on PlantDoc; enhanced fine-grained detection	Performance improvement mainly validated on a single dataset; real-time performance not detailed
Duhan, S., et al. (2025)	Light-weight MobileNetV2 with attention modules for edge devices	Achieved up to 100% validation accuracy with real-time capability for IoT	Real-world deployment scalability and generalization to varied crop types not discussed
Ashurov, A. Y., et al. (2025)	Depthwise CNN with squeeze-and-excitation blocks	98% accuracy and 98.2% F1-score on field images; computationally efficient	Limited insight on dataset diversity and long-term deployment robustness
Roumeliotis, K. I., et al. (2025)	Vision-Language Models (GPT-4o + CNN) for plant disease detection	99.88% accuracy in fine-tuned mode; better than ResNet-50	Zero-shot performance was significantly poorer; needs some level of training
Demilie, W. B. (2024)	MobileNet and CNN with Grad-CAM for explainability	MobileNet achieved 96% accuracy; CNN 89%; better interpretability	Precision lower in MobileNet; effectiveness in real-time

		through Grad-CAM	scenarios not detailed
Adekunle, T. S., et al. (2024)	E-GreenNet: Lightweight MobileNetV3Small-based framework	Achieved up to 100% accuracy; excelled on low-contrast and complex images	Generalization beyond the tested datasets (PV, PC, DRLI) not demonstrated
Barman, U., et al. (2024)	Vision Transformer + InceptionV3 for smartphone-based detection	90.99% accuracy; implemented in Android for tomato disease detection	Limited scope (only tomato); relatively low accuracy for ViT
Bhagat, S., et al. (2024)	Lite-MDC with MDsConv for real-time plant disease detection	99.78% accuracy at 34 FPS with only 2.2M parameters	Evaluation on few crops; limited discussion on low-light or occlusion conditions
Maurya, R., et al. (2024)	Meta-ensemble of MLP-Mixer and fast LSTM for low-power IoT	Up to 98.43% accuracy; cost-effective and suitable for microcontrollers	Ensemble complexity may hinder training and updates on limited-resource devices
Moupojou, E., et al. (2023)	Critique and creation of FieldPlant dataset	More realistic field annotations (8,629 leaves); improved real-world detection	Limited scope (diamond-back moth); dataset still relatively small
Panchal, A. V., et al. (2023)	Feature extraction and pixel-based preprocessing for CNNs	Effective training on 87K RGB images; suitable for resource-constrained areas	Approach reliant on preprocessing; model scalability untested
Archana, U., et al. (2023)	Residual CNN using color-based features for rice disease	96.35% accuracy; effective for rice using ResNet-50 and phenotyping	Applicability to non-rice crops and diverse conditions unclear
Shovon, M. S. H., et al. (2023)	Ensemble with InceptionResNetV2, EfficientNetV2L, Xception	98.53% accuracy with Grad-CAM and Score-CAM interpretability	High model complexity; requires substantial computational resources
Hamed, B. S., et al. (2023)	EfficientNetV2S fine-tuning on noisy, low-res images	Improved robustness in adverse conditions with modified dataset	Generalization to unseen diseases and crops not discussed
Albattah, W., et al. (2022)	Proposed a Custom CenterNet with DenseNet-77 for plant disease classification	Achieved superior accuracy on PlantVillage dataset; effectively handled intensity variation, color similarity, and leaf shape differences	Model mainly tested on PlantVillage dataset, may not generalize well to real-world field conditions
Akshai, K. P., &Anitha, J. (2021)	Used CNN and pre-trained models (VGG, ResNet, DenseNet) for disease classification	DenseNet achieved highest accuracy; deep learning models outperformed manual inspection	Limited dataset; focus only on leaf symptoms; real-world variations not fully addressed
Guo, Y., et al. (2020)	Developed a deep learning model using RPN for leaf localization and CV algorithm for segmentation	Achieved 83.57% accuracy; improved detection of diseases like black rot, bacterial plaque, and rust	Accuracy lower than other DL approaches; segmentation step increases computational complexity
Saleem, M. H., et al. (2019)	Review of ML and DL approaches for plant disease detection	DL models showed superior accuracy; highlighted visualization techniques and research gaps	Lacks experimental validation; mainly conceptual and review-based
Barbedo, J. G. (2018)	Examined effectiveness and challenges of deep learning in plant pathology	Identified advantages of DL but also highlighted overlooked experimental issues	Dataset-specific results; highlighted that models may not perform well in practical farm conditions
Fuentes, A., et al. (2017)	Applied Faster R-CNN, R-FCN, and SSD with VGG/ResNet for tomato disease detection	Successfully recognized 9 diseases/pests in challenging field images; improved detection with augmentation	Focused only on tomato plants; high computational requirements
Sladojevic, S., et al. (2016)	Developed CNN model for leaf image classification	Recognized 13 plant diseases with ~96.3% average precision	Early-stage approach; dataset built in controlled conditions; limited to leaf images only

VII. RESEARCH GAPS DISCUSSION

Although deep learning has made a significant improvement in the detection of plant diseases, some general research issues are not fully explored. One of the more recognised problems is a focus on pre-recorded training datasets such as PlantVillage or PlantDoc that contain limited granularity and noise, not indicative of real-world agriculture. Recent studies based on Moupojou et al. Although, even datasets with field-annotations (e.g., FieldPlant) that approaches initiated the work on field-annotated dataset, but still these datasets only cover specific crops or diseases. This prevents models from generalizing in real-life farming conditions which are performed in various harsh environments with different lighting, occlusions, and complex environmental background noise.Furthermore, high-performance models with model interpretability and explainability are still not quite mainstream. Although this paper is underpinned by the "black-box" concept, some of the existing techniques have attempted to improve upon this behavior explored through Grad-CAM and Score-CAM respectively; indeed, such strategies make a strong predictive model but do abandon transparent decision-making from state-of-the-art models. This low interpretability can make it hard for agricultural practitioners to have confidence in or corroborate model outputs at a scale relevant to adoption.

Secondly, there is a technological trade-off between deploying more complex models and edge deployment feasibility. Models such as Lite-MDC (Bhagat et al., 2024) and RTR_Lite_MobileNet (Duhan et al., 2025) target on resource-constrained settings with high speed and accuracy, however many state-of-the-art architectures based on Transformers or Vision-Language Models (Roumeliotis et al., 2025) are computationally intensive for real-time, in-siloáidiagnosis in low-power IoT or mobile scenarios.
Very few people have been studying the problem of cross-domain robustness and zero-shot generalization. Is that, as lead author SteliosRoumeliotis and associates', (2025) and Vision-Language Models like GPT-4o fine-tuned shows much promise but they have poor zero-shot performance, implying that there is a present gap in generalizing to unseen or expected diseases as well as crops without requiring the models to be retrained Likewise, the vast majority of extant models do not include techniques to identify new forms of disease presentation or species with little supervision.Finally, there are few models of serious multi-modal data integration and severity estimation.While Bouacida et al. Although some work (2025) moved towards prediction severity by concentrating on small diseased regions, the plant disease detection-based models in literature assume mainly as a binary (i.e., classification) task without giving an insight about stages of progression or leveraging other sensory modalities such as temperature, humidity or soil health. This kind of incorporation could bring significant improvements to the feasibility of AI used in precision agriculture, as many decisions are not just based on visual recognition.

Collectively, future studies should look at creating models which are scalable and interpretable and perform reasonably well on diverse conditions of the real world, deployable to edge devices with severity estimation ability without reliance on large amount of annotated data through advanced few-shot or zero-shot learning approaches.

Limitation of Studies:
The limitations of the studies can be summarized in following points:

- **Dataset Dependency -** Numerous studies are predominantly dependent on demos planned data sets (i.e., PlantVillage) or crop-specific data sets, limiting generalization to field conditions.
- **Limited Crop Coverage -** Many studies only focus on singular crops in tracking Select as tomato, rice, and the diamond-back moth, resulting in limited generalization across plant species variety.
- **Controlled Environment Bias -** Many data sets are captured under controlled or simplified conditions making the models less robust and not tested in situations with complex backgrounds, and mixed crops, low-light, or occlusion.
- **Scalability -** Many studies defaults to high computational complexity (CNN, CenterNet, or ensemble methods), limiting real time networking and use on resource constrained devices.
- **Generalization -** The models perform poorly or lack confidences with unseen diseases, plants overlapping, and plant variance so therefore less reliable and certainly not trustworthy in real field conditions of agricultural applications.
- **Performance trade-offs -** Limited real-time precision, and accuracy of certain architectures (i.e., MoblieNet and ViT) versus heavier models limits real-time field use.
- **Preprocess Dependence -** Many methods are segmented or preprocessing heavy (i.e., RPN + CV algo), adding additional computation expense and limiting scalability.
- **Lack of Real-world Validation -** The majority of studies validate the models only under experimental conditions (i.e., lab or demo), which typically does not evaluate scalability, long-term deployment, and adaptability at the field-level.

- **Ensemble Complexity –** Although ensembles often improve accuracy, they complicate training, updating, and deploying models on limited-resource devices.
- **Research Gaps in Robustness –** Very few studies examine performance during diverse conditions (e.g., low-light, noisy environments, or cross-location variability).

VIII. CONCLUSION

The increasing threat of plant diseases to global food security, combined with the interest in promoting methods to engage in sustainable agricultural practices, has driven a rapid increase in the integration of advanced technologies in farming systems. One of these technological advancements is deep learning, which is becoming a powerful, flexible and automated method for detecting plant disease with unprecedented accuracy, upscalability and adaptability. In this paper, a review study was conducted of state-of-the-art deep learning algorithms, particularly, CNNs, RNNs, GANs, and transformer-based architectures, which are utilized to detect and monitor plant disease across a wide range of crops and contexts.Convolutional Neural Networks have led to excellent results for leaf image classification, and are central to most systems that detect disease from images. Recurrent Neural Networks family of models, particularly Long Short Term Memory network (LSTM), have enabled the modeling of disease progression over the time using temporal data. Generative Adversarial Networks are supporting an age-old problem of data scarcity with generative synthetic disease images that boost the performance of models. Transformer-based architectures are just emerging, but they show tremendous promise to drive the envelope for accuracy and interpretability, particularly towards complex learning in various modalities.While these advancements are impressive, there are still obstacles ahead to address, including limited annotated datasets, some models may not generalize well to all real-world scenarios, the computational burden, and the black-box that is nature of most deep learning systems. Fortunately, there is a lot of ongoing work within areas such as explainable AI, lightweight model deployment, and sensor fusion to overcome these limitations.-In conclusion, deep learning is already in a position to pave the way for smart agriculture, and if it continues to develop and is thoughtfully incorporated into solutions ready for the field, it has the potential to revolutionize plant health monitoring, diminish crop loss, and play an important role as we move towards sustainable and resilient food systems around the world.

ACKNOWLEDGMENT

We thank ShriPrakashbhai Patel, Dr. P. M. Udani, Dr. P. J. Patel and Dr. Y. S. Patel from Sankalchand Patel University for their support

REFERENCES

[1] Upadhyay, A., Chandel, N. S., Singh, K. P., Chakraborty, S. K., Nandede, B. M., Kumar, M., ... &Elbeltagi, A. (2025). Deep learning and computer vision in plant disease detection: a comprehensive review of techniques, models, and trends in precision agriculture. *Artificial Intelligence Review*, *58*(3), 92.

[2] Hari, P., & Singh, M. P. (2025). Adaptive knowledge transfer using federated deep learning for plant disease detection. *Computers and Electronics in Agriculture*, *229*, 109720.

[3] Guo, G., Lai, S., Wu, Q., Shou, Y., & Shi, W. (2025). Enhancing domain adaptation for plant diseases detection through Masked Image Consistency in Multi-Granularity Alignment. *Expert Systems with Applications*, *276*, 127101.

[4] Shoaib, M., Shah, B., Ei-Sappagh, S., Ali, A., Ullah, A., Alenezi, F., ...& Ali, F. (2023). An advanced deep learning models-based plant disease detection: A review of recent research. *Frontiers in Plant Science*, *14*, 1158933.

[5] Jafar, A., Bibi, N., Naqvi, R. A., Sadeghi-Niaraki, A., &Jeong, D. (2024). Revolutionizing agriculture with artificial intelligence: plant disease detection methods, applications, and their limitations. *Frontiers in Plant Science*, *15*, 1356260.

[6] Haridasan, A., Thomas, J., & Raj, E. D. (2023). Deep learning system for paddy plant disease detection and classification. *Environmental monitoring and assessment*, *195*(1), 120.

[7] Shafik, W., Tufail, A., Namoun, A., De Silva, L. C., &Apong, R. A. A. H. M. (2023). A systematic literature review on plant disease detection: Motivations, classification techniques, datasets, challenges, and future trends. *Ieee Access*, *11*, 59174-59203.

[8] Ramanjot, Mittal, U., Wadhawan, A., Singla, J., Jhanjhi, N. Z., Ghoniem, R. M., ...&Abdelmaboud, A. (2023). Plant disease detection and classification: A systematic literature review. *Sensors*, *23*(10), 4769.

[9] Chug, A., Bhatia, A., Singh, A. P., & Singh, D. (2023). A novel framework for image-based plant disease detection using hybrid deep learning approach. *Soft Computing*, *27*(18), 13613-13638.

[10] Sunil, C. K., Jaidhar, C. D., &Patil, N. (2023). Systematic study on deep learning-based plant disease detection or classification. *Artificial Intelligence Review*, *56*(12), 14955-15052.

[11] Chin, R., Catal, C., &Kassahun, A. (2023). Plant disease detection using drones in precision agriculture. *Precision Agriculture*, *24*(5), 1663-1682.

[12] Rayhana, R., Ma, Z., Liu, Z., Xiao, G., Ruan, Y., &Sangha, J. S. (2023). A review on plant disease detection using hyperspectral imaging. *IEEE Transactions on AgriFood Electronics*, *1*(2), 108-134.

[13] John, M. A., Bankole, I., Ajayi-Moses, O., Ijila, T., Jeje, T., &Lalit, P. (2023). Relevance of advanced plant disease detection techniques in disease and Pest Management for Ensuring Food Security and Their Implication: A review. *American Journal of Plant Sciences*, *14*(11), 1260-1295.

[14] Dey, P., Mahmud, T., Nahar, S. R., Hossain, M. S., &Andersson, K. (2024, January). Plant disease detection in precision agriculture: Deep learning approaches. In *2024 2nd International Conference on Intelligent Data Communication Technologies and Internet of Things (IDCIoT)* (pp. 661-667). IEEE.

[15] Upadhyay, A., Chandel, N. S., Singh, K. P., Chakraborty, S. K., Nandede, B. M., Kumar, M., ... &Elbeltagi, A. (2025). Deep learning and computer vision in plant disease detection: a comprehensive review of techniques, models, and trends in precision agriculture. *Artificial Intelligence Review*, *58*(3), 92.

[16] Dhaka, V. S., Kundu, N., Rani, G., Zumpano, E., &Vocaturo, E. (2023). Role of internet of things and deep learning techniques in plant disease detection and classification: A focused review. *Sensors*, *23*(18), 7877.

[17] Qadri, S. A. A., Huang, N. F., Wani, T. M., &Bhat, S. A. (2023, August). Plant disease detection and segmentation using end-to-end yolov8: A comprehensive approach. In *2023 IEEE 13th International Conference on Control System, Computing and Engineering (ICCSCE)* (pp. 155-160). IEEE.

[18] Ali, U., Ismail, M. A., AriyaluranHabeeb, R. A., & Ali Shah, S. R. (2024). Performance evaluation of YOLO models in plant disease detection. *Journal of Informatics and Web Engineering*, *3*(2).

[19] Ojo, M. O., &Zahid, A. (2023). Improving deep learning classifiers performance via preprocessing and class imbalance approaches in a plant disease detection pipeline. *Agronomy*, *13*(3), 887.

[20] Jung, M., Song, J. S., Shin, A. Y., Choi, B., Go, S., Kwon, S. Y., ...& Kim, Y. M. (2023). Construction of deep learning-based disease detection model in plants. *Scientific reports*, *13*(1), 7331.

[21] Sahu, S. K., &Pandey, M. (2023). An optimal hybrid multiclass SVM for plant leaf disease detection using spatial Fuzzy C-Means model. *Expert systems with applications*, *214*, 118989.

[22] Bouacida, I., Farou, B., Djakhdjakha, L., Seridi, H., &Kurulay, M. (2025). Innovative deep learning approach for cross-crop plant disease detection: A generalized method for identifying unhealthy leaves. *Information Processing in Agriculture*, *12*(1), 54-67.

[23] Miao, Y., Meng, W., & Zhou, X. (2025). SerpensGate-YOLOv8: an enhanced YOLOv8 model for accurate plant disease detection. *Frontiers in Plant Science*, *15*, 1514832.

[24] Duhan, S., Gulia, P., Gill, N. S., &Narwal, E. (2025). RTR_Lite_MobileNetV2: A lightweight and efficient model for plant disease detection and classification. *Current Plant Biology*, *42*, 100459.

[25] Ashurov, A. Y., Al-Gaashani, M. S., Samee, N. A., Alkanhel, R., Atteia, G., Abdallah, H. A., &Saleh Ali Muthanna, M. (2025). Enhancing plant disease detection through deep learning: a Depthwise CNN with squeeze and excitation integration and residual skip connections. *Frontiers in Plant Science*, *15*, 1505857.

[26] Roumeliotis, K. I., Sapkota, R., Karkee, M., Tselikas, N. D., &Nasiopoulos, D. K. (2025). Plant disease detection through multimodal large language models and convolutional neural networks. *arXiv preprint arXiv:2504.20419*.

[27] Demilie, W. B. (2024). Plant disease detection and classification techniques: a comparative study of the performances. *Journal of Big Data*, *11*(1), 5.

[28] Adekunle, T. S., Lawrence, M. O., Alabi, O. O., Afolorunso, A. A., Ebong, G. N., &Oladipupo, M. A. (2024). Deep learning technique for plant disease detection. *Computer Science and Information Technologies*, *5*(1), 55-62.

[29] Barman, U., Sarma, P., Rahman, M., Deka, V., Lahkar, S., Sharma, V., &Saikia, M. J. (2024). Vit-SmartAgri: vision transformer and smartphone-based plant disease detection for smart agriculture. *Agronomy*, *14*(2), 327.

[30] Bhagat, S., Kokare, M., Haswani, V., Hambarde, P., Taori, T., Ghante, P. H., &Patil, D. K. (2024). Advancing real-time plant disease detection: A lightweight deep learning approach and novel dataset for pigeon pea crop. *Smart Agricultural Technology*, *7*, 100408.

[31] Maurya, R., Mahapatra, S., & Rajput, L. (2024). A lightweight meta-ensemble approach for plant disease detection suitable for IOT-based environments. *Ieee Access*, *12*, 28096-28108.

[32] Moupojou, E., Tagne, A., Retraint, F., Tadonkemwa, A., Wilfried, D., Tapamo, H., &Nkenlifack, M. (2023). FieldPlant: A dataset of field plant images for plant disease detection and classification with deep learning. *IEEE Access*, *11*, 35398-35410.

[33] Panchal, A. V., Patel, S. C., Bagyalakshmi, K., Kumar, P., Khan, I. R., &Soni, M. (2023). Image-based plant diseases detection using deep learning. *Materials Today: Proceedings*, *80*, 3500-3506.

[34] Archana, U., Khan, A., Sudarshanam, A., Sathya, C., Koshariya, A. K., &Krishnamoorthy, R. (2023, April). Plant disease detection using resnet. In *2023 International Conference on Inventive Computation Technologies (ICICT)* (pp. 614-618). IEEE.

[35] Shovon, M. S. H., Mozumder, S. J., Pal, O. K., Mridha, M. F., Asai, N., & Shin, J. (2023). PlantDet: A robust multi-model ensemble method based on deep learning for plant disease detection. *IEEE Access*, *11*, 34846-34859.

[36] Hamed, B. S., Hussein, M. M., &Mousa, A. M. (2023). Plant disease detection using deep learning. *International Journal of Intelligent Systems and Applications*, *15*(6), 38-50.

[37] Albattah, W., Nawaz, M., Javed, A., Masood, M., &Albahli, S. (2022). A novel deep learning method for detection and classification of plant diseases. *Complex & Intelligent Systems*, *8*(1), 507-524.

[38] Akshai, K. P., &Anitha, J. (2021, May). Plant disease classification using deep learning. In *2021 3rd International Conference on Signal Processing and Communication (ICPSC)* (pp. 407-411). IEEE.

[39] Guo, Y., Zhang, J., Yin, C., Hu, X., Zou, Y., Xue, Z., & Wang, W. (2020). Plant disease identification based on deep learning algorithm in smart farming. *Discrete Dynamics in Nature and Society*, *2020*(1), 2479172.

[40] Saleem, M. H., Potgieter, J., &Arif, K. M. (2019). Plant disease detection and classification by deep learning. *Plants*, *8*(11), 468.

[41] Barbedo, J. G. (2018). Factors influencing the use of deep learning for plant disease recognition. *Biosystems engineering*, *172*, 84-91.

[42] Fuentes, A., Yoon, S., Kim, S. C., & Park, D. S. (2017). A robust deep-learning-based detector for real-time tomato plant diseases and pests recognition. *Sensors*, *17*(9), 2022.

[43] Sladojevic, S., Arsenovic, M., Anderla, A., Culibrk, D., &Stefanovic, D. (2016). Deep neural networks based recognition of plant diseases by leaf image classification. *Computational intelligence and neuroscience*, *2016*(1), 328980

Advancement in Automated Hate Speech Detection: A Survey of Machine Learning and Ensemble Methods

Rutuja Goradkar[1]*,Smita Bharne[2],Vaibhav Narawade[3]

Department of Computer Engineering Ramrao Adik Institute of Technology D. Y. Patil Deemed to be University Navi Mumbai, India[1]

Department of Engineering and Technology, Bharati Vidyapeeth Deemed to be University, Navi Mumbai, Maharashtra, India[2]

Department of Computer Engineering Ramrao Adik Institute of Technology D. Y. Patil Deemed to be University Navi Mumbai, India[3]

rutujagoradkar@gmail.com[1],smita146@gmail.com[2],vaibhav.narawade@rait.ac.in[3]

Abstract: This survey article presents a comprehensive review of state-of-the-art methodologies for real-time hate speech and toxicity detection on social media platforms, with a focus on their contributions and associated performance metrics. The survey mainly focused on classical machine learning models such as Support Vector Machine (SVM), Random Forest, XGBoost, Logistic Regression, and ensemble learning strategies, which have shown improvements in classification accuracy by leveraging complementary strengths. In addition, deep learning approaches, including Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) models, as well as transformer-based architectures such as BERT and RoBERTa, are reviewed for their ability to capture contextual and sequential patterns that traditional ML methods often miss. The survey also shows the role of stacking models in enhancing performance by achieving superior accuracy compared to individual classifiers. The revised survey incorporates a formal methodology section, a taxonomy of detection methods, a dedicated dataset comparison, coverage of advanced models including DeBERTa [24], HateBERT [26], LLM-based [27] and multimodal approaches [28], strengthened critical analysis, and a future research directions section.

Keywords: Hate Speech Detection, Natural Language Processing (NLP), Machine Learning, Deep Learning Ensemble

I. INTRODUCTION

The quick spread of the social media networks restructured how people communicate, interact and exchange type of information in real time [1]. With Twitter X, Facebook, Instagram, YouTube and Reddit there have been lot of opportunities to express, create communities and share information as never before [2]. This online communication has been led to destructive habits and lately the growth of toxic content and hate speech that has the potential to do harassment, cyber-bullying, discrimination and the polarization of society [3]. If this type of things not monitored then such kind of content can cause mental damage, promote weak stereotypes and even lead to violence.Toxicity is generally used to denote offensive or harmful words, whether it is personal attacks or abusive speech or demeaning utterances [4], but hate speech is a more specific category where it is directed against a person or group of persons either on race, religion, gender, nationality, or sexual orientation [5]. These two seem to have distinct challenges, but there are overlaps as well which create a challenge to automated detection systems. In contrast to the traditional forms of data that are textual in nature, user-generated content on social media has been marked with brevity, informality in grammar, slangs, abbreviations, emojis [6] and multi-linguistic code-switching, factors, which complicate accurate content-interpretation. Additionally, it is hard to detect hate speech since often it occurs in more of a contextual manner rather than presently which makes it difficult to detect. Because of this, automated systems cannot merely be reliable but rather must be accurate, dynamic, scalable, and be able to process in real-time to keep up with the traffic generated by online communication.

There is a large variety of computational approaches to the problem of toxicity and hate speech detection that have been suggested in recent years. Early research was mainly focused on traditional ML algorithms. Strategy brought early success but also it had lots of weaknesses including little background knowledge and poor flexibility to changing linguistics situation. More

http://doi.org/10.63766/spujstmr.26.000089

recently some DL techniques have shown success which managed to take advantage of semantics dependency and sequential type of ordering in texts. They needed hige volumes of labeled information and considerable computational facilities which limit their use in the real-time. The use of ensemble learning and transformer-based architectures has been growing, and they have pushed the state of the art substantially [7]. In this survey paper, we attempted to give a overview of machine learning approaches to detecting real-time toxicity and hate speech with a systematic study of classical machine learning, deep learning, ensembles, and transformer-based models. We underline the issues of accuracy, scalability, and computational calculation efficiency, and how different methods are more or less suited to real time monitoring systems. Moreover, we overview the popular datasets, assessment criteria, and comparative performance results, thus, covering both the key contributions and the limitations of the available literature.

Fig1. shows Racism, Sexism, Religious Discrimination, Homophobia, and Toxic Comments on social media.

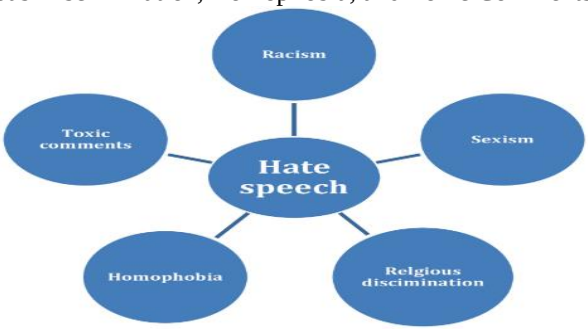

Fig.1 : Categories of Hate Speech on Social Media by [8]

II. RELATED WORK

A.Machine Learning Approaches in Hate Speech

Hate speech detection on social media has been widely studied using different ML and DL approaches. The following studies highlight key contributions, methodologies, and limitations in the field.

The research suggested by [9] can help improve the hate speech detection on the social media platform, as it will be possible to address the issue as a text classification problem and apply the methods of ML as shown in Table 1. The research was aimed at creating a strong pipeline that encompassed data preparation, feature extraction, and selection of classifiers, especially in terms of improving their classification accuracy and overcoming the problem of data imbalance. Study given by [10] proposed a comprehensive approach of hate speech detection on Twitter through the comparison of the conventional ML algorithms with modern DL methods. The proposed method meant that shallow learning baselines and DL architectures such as LSTM were to be used and that pre-trained word embedding should be incorporated to enhance semantic knowledge. The study offered by [11] represents a method of identifying cyber hate speech in the Arabic setting on Twitter, which was inspired by the increasing rates of hatred online and its potential harm to social cohesion, especially in the Arab world. A comparative study was suggested by [12] in order to conduct a systematic evaluation of different feature engineering methods and machine learning algorithms on automatic hate speech detection on social media. The study was conducted to determine the most efficient combinations of features and algorithms, since no prior studies had compared these methods in as much detail as determining the best using a commonly available publicly available dataset. The increasing abuse of social media networks, especially twitter, in the sharing of hate speech concerning racism, gender, religion and defamatory posts has motivated a large amount of research into automated detection techniques. A number of studies have examined machine learning strategies in order to overcome this dilemma, aiming at minimizing the error rate of classification and managing the volume of textual data as shown in Table 1. The authors also studied in [13] how various algorithms can be used to identify hate speech in Twitter comments. They found that ensemble and margin-based classifiers were better than the others, indicating that they are well-adapted to noisy social media data. However, the study also recognised that there was a limitation in that it was challenging to capture sarcasm, implicit hate, and context dependency. Continuing the trend of this research, [14] turned to Portuguese tweets to evaluate the generalizability of traditional machine learning techniques to language. The authors compared performance using SVM, Logistic Regression, Multilayer Perceptron (MLP), and Naive Bayes (NB) with LSTM which is an example of a deep learning model. The authors of [15] came up with a Novel Gaussian Naive Bayes (GNB) algorithm and compared it to the Logistic Regression when predicting hate speech. The study was carried out on a small sample of 20 samples in two groups and employed a G-power pretest and a sample T-test to determine statistical significant values. Results indicated that the Novel GNB was found to be 95%-92% more accurate than the existing system and there was a significant difference between the two models ($p=0.000$, $p<0.05$). In spite of these advancements, the research found out that the final predictive accuracy of Logistic

Regression was more reliable. Nevertheless, the small size of the dataset cast doubt on the generalizability and application to the real-world. Another study given by [16] who proposed a hybrid method was proposed to combine LR with SVM to take advantage of the simplicity and speed of LR, and the ability of SVM to deal with non-linear data.

TABLE I
SUMMARY TABLE

Study	Methods / Models Used	Key Contributions	Results / Performance	Limitations
[9]	Classical ML, Ensembles, DL	Developed a strong pipeline with data prep, feature extraction, classifier selection	Higher accuracy & F1-score than baselines	Poor cross-domain & multilingual generalization, dataset scarcity
[10]	ML vs. DL (LSTM, BiLSTM, embeddings)	Compared shallow ML and DL models for hate/offensive/neutral	BiLSTM best at capturing sequential & semantic info	Struggled with sarcasm, context, multilingual generalization
[11]	NLP + ML (Random Forest, TF-IDF)	Focused on Arabic hate speech, attribute importance analysis	RF + TF-IDF gave best performance	Arabic dialectal complexity, lack of annotated data, low generalizability
[12]	Feature Eng. (unigrams, bigrams) + 8 ML classifiers	Systematic benchmarking of feature–classifier combos	Bigram + SVM = 79% accuracy	Small datasets, domain bias, limited multilingual applicability
[13]	LR, NB, KNN, DT, RF, SVM	Compared multiple ML algorithms for Twitter hate detection	RF (98.2%), DT (96.2%), SVM (95.5%) best	Hard to capture sarcasm, implicit hate, context dependency
[14]	SVM, LR, MLP, NB vs. LSTM	Studied traditional ML vs DL for Portuguese	SVM (0.8836), LR (0.8765) outperformed LSTM	Language-specific limits, superficial features
[15]	Novel Gaussian Naive Bayes (GNB) vs. LR	Proposed new GNB model with pre-test & T-test validation	GNB more accurate (95–92%), significant p-value	Very small dataset, weak generalization, LR more reliable
[16]	Hybrid LR + SVM	Combined LR (simplicity) + SVM (non-linear strength)	Accuracy: 96%	Ambiguity in online language, dataset dependency

B. Deep Learning Approaches in Hate Speech

The rapid growth of online social networks has increased the attention towards hate speech and resulted in a study of effective means of its detection. In study [17], an automated system was proposed based on a Deep Convolutional Neural Network (DCNN) that uses GloVeembeddings to detect hateful language posted on Twitter. This method performed better than current models although it is limited by using only text and ignoring multimodal displays of hate. Considering the difficulty of differentiating hateful vs. merely offensive speech, Study [18] proposed HateClassify, a sequential CNN (SCNN)-based multilabel framework that demonstrated large enhancements over multiclass architecture as shown in Table 2. The model achieved a maximum accuracy of 20 percent in hate detection and an average increase of 0.095 in F-measure, with a highest increase of 0.2 for a hate category, on the CrowdFlower, Davidson et al., and Sexism vs Racism datasets. But it was also sensitive to class imbalance, particularly in grossly skewed data sets. In order to address the imbalance further, a double-layer hybrid CNN-RNN model with oversampling was proposed in Study [19], which improved detection accuracy and controlled overfitting through dropout and early stopping. It had an accuracy of 0.827 and F1-score 0.883 on imbalanced data, and the results are much higher on balanced data with 0.908 and 0.914 accuracy and F1-score respectively, evidently surpassing baseline CNN-RNN models. However, over sampling was subject to low external validity to the unobserved data. In the next section, Study [20] introduced BiCHAT, a deep hierarchical architecture that combines BERT embedded representations, deep CNN, BiLSTM, and attention to improve tweet representation. On three benchmark Twitter datasets, BiCHAT achieved 5 percent higher training and 9 percent higher validation accuracy than state-of-the-art models, and ablation studies verified the CNN layer as most effective. However, it is computationally complex, thus being scale-constrained when used in a real-time moderation context. Together, these papers highlight the historical transformation of simple CNN-based detections to more advanced hybrid and attention-based models, the increased processing of imbalanced information, better semantic encoding, and

increased classification accuracy, and also show the problem of generalizability, scalability, and multimodal fusion as shown in Table 2.

C. Transformer Based Approaches in Hate Speech

Recent studies indicate the usefulness of transformer-based models in the challenging task of hate speech detection across languages and platforms. Aggression and misogynistic aggression identification in English, Hindi and Bangali were studied by [21] as part of the TRAC shared tasks. The authors used BERT, RoBERTa, DistilRoBERTa, XLM-RoBERTa, and SVM classifiers, and found that SVM scored higher on average, because of strong majority-class predictions, but that transformer-based models such as BERT scored better at detecting minority classes which tend to be more vital in the real world. Another study given by [22] aimed to generate an Urdu hate vocabulary and a dataset of 10,526 annotated tweets to evolve hate speech detection in a low-resource language. The study also tested BERT variants, e.g., DistilBERT and XLM-RoBERTa, using traditional ML baselines and transfer learning with FastText and multilingual BERT embeddings, and their F1-scores 0.68–0.69, hence outperforming traditional baselines. The issue of shortage of resources and language barrier that have been mentioned in this paper in the case of Urdu in which the popularity of the transformer models has been realized irrespective of the constraint of data. Study [23] made a step forward by training DistilBERT on a multiclass hate speech dataset of 24,783 labeled tweets on Twitter to address the inefficiency of RNNs and attention-based models, which have a problem of long-term dependency and cannot be run in parallel as shown in Table 3. Not only did DistilBERT overcome these problems, but it also reached state-of-the-art accuracies of 0.92, which exceeds the accuracy of both the RNN-based and transformer baselines, and is computationally efficient.

TABLE II
SUMMARY TABLE

Study	Methods / Models Used	Key Contributions	Results / Performance	Datasets Used	Limitations
[17]	DCNN with GloVeembeddings	Automated hate speech detection on Twitter using semantic representation of tweets	Precision: 0.97, Recall: 0.88, F1-score: 0.92	Twitter dataset (500M tweets/day stream, subset used)	Considers only text; ignores multimodal hate speech (e.g., memes, sarcasm, images)
[18]	Sequential CNN (SCNN) for multilabel classification vs multiclass	Proposed HateClassify framework, showing multilabel classification improves detection accuracy	Up to 20% improvement in hate detection; avg 0.095 F-measure increase; max 0.2 F-measure gain for hate class	CrowdFlower (14,509 tweets) - Davidson et al. (24,783 tweets) - Sexism vs Racism (6,492 tweets)	Sensitive to class imbalance, especially in Dataset 3 (86% tweets in one class)
[19]	Double-layer Hybrid CNN-RNN with oversampling	Introduced double-layer architecture to improve accuracy on imbalanced data; used dropout & early stopping to avoid overfitting	Imbalanced data: Acc 0.827, F1 0.883 Balanced data: Acc 0.908, Precision 0.943, Recall 0.894, F1 0.914	Hate speech dataset	Oversampling may cause overfitting; real-world generalization may be limited
[20]	BiCHAT: BiLSTM + Deep CNN + Hierarchical Attention with BERT embeddings	Novel hierarchical attention model for better tweet representation; ablation study shows CNN most impactful	Outperformed SOTA; +5% training accuracy, +9% validation accuracy	Three Twitter benchmark datasets	High computational cost and model complexity may hinder real-time scalability

D. Survey Methodology

To ensure a systematic and reproducible review, this survey adopted a structured literature search protocol:
Search Databases: Google Scholar, IEEE Xplore, ACM Digital Library, Springer, and arXiv.

Keywords: 'hate speech detection', 'toxicity detection', 'NLP', 'machine learning', 'deep learning', 'transformer', 'BERT', 'ensemble learning', 'LLM hate speech', 'multimodal hate speech'.
Time Period: 2019–2026 primarily; seminal earlier works included where foundational.
Inclusion Criteria: Peer-reviewed articles and conference papers; ML/DL/transformer-based hate speech or toxicity detection methods; quantitative performance metrics reported.
Exclusion Criteria: Non-English papers without translation; papers without accessible full text; opinion pieces without experimental evaluation.
Papers Reviewed: An initial pool of 80+ papers was screened; 23 were selected based on relevance and the inclusion/exclusion criteria.
This systematic protocol ensures objectivity and reproducibility, allowing readers to assess the completeness of the survey [3].

TABLE III

SUMMARY TABLE

Study	Methods / Models Used	Key Contributions	Dataset	Results / Performance	Limitations
[21]	En-BERT, RoBERTa, DistilRoBERTa, SVM, M-BERT, XLM-RoBERTa	Comparative analysis of traditional ML vs. transformer models across three languages, highlighting trade-offs between majority and minority class prediction	TRAC shared task datasets (English, Hindi, Bangla)	En-BERT: F1 = 0.73 (Rank 5/16); SVM: F1 = 0.87 (English B, Rank 2/15); SVM: F1 = 0.79 (Hindi A, Rank 2/10); XLM-RoBERTa: F1 = 0.87 (Hindi B, Rank 2/10); SVM: F1 = 0.81 (Bangla A, Rank 2/10); SVM: F1 = 0.93 (Bangla B, Rank 4/8)	SVM favored majority classes, missing subtle minority-class aggression
[22]	Machine Learning baselines, FastText Urdu embeddings, Multilingual BERT, BERT variants (BERT, XLM-RoBERTa, DistilBERT)	First creation of an Urdu hate lexicon and dataset, enabling resource development for low-resource language hate speech detection	10,526 annotated Urdu tweets (lexicon-based)	BERT: F1 = 0.68; XLM-RoBERTa: F1 = 0.68; DistilBERT: F1 = 0.69	Limited dataset size and lexicon-based coverage may not capture diverse hate expressions
[23]	DistilBERT (distilbert-base-uncased), compared with attention-based RNNs &other transformer baselines	Demonstrated that transformer-based models (DistilBERT) outperform RNNs in hate speech detection, with advantages in efficiency and parallelization	Public multiclass hate speech corpus (24,783 tweets)	DistilBERT achieved accuracy = 0.92, outperforming baselines	Validation limited to one dataset; performance may vary across languages and platforms

III. METHODS

This section describes the research methods using pre-trained models that are widely adopted for hate speech and toxicity detection. The models have been used ML and DL models and stacked ensemble approach also.

A. Support Vector Machine (SVM)

SVM is a supervised learning model that finds extensive application in text classification. It works by determining an optimum hyperplane that separates the data points of divergent categories in a high dimension space. In the case of textual data, term frequency -inverse document frequency (TF-IDF) is often used to extract features, which are indicators of word significance in documents. SVM may apply the use of kernel functions, like radial basis function (RBF), to deal with non-linear correlations in text. This helps it to discriminate between toxic language and non-toxic language even when there are complicated boundaries. The primary strength is its high robustness in the case of sparse and high-dimensional input spaces, which are common to natural language data. Even though SVM would need to be parameter-tuned to the unique text patterns, it is one of the powerful tools in baseline classification tasks in toxicity detection.

B. Random Forest (RF)

Random Forest (RF) is an ensemble technique that is designed on the basis of combining several decision trees. All trees are trained with a random subset of the data and data features and gain predictions by majority voting. Such a structure decreases the chances of the overfitting effect and improves generalization. The features used in hate speech detection are input features, which in this case are TF–IDF vectors fed into RF, and the model tries to capture patterns corresponding to offensive or neutral language. RF is able to deal with noisy and unstructured text since each tree examines a distinct part of the feature space. Interpretability is one of its strengths because it is possible to quantify the importance of features in order to determine which terms should be used to classify. Although RF might prove computationally expensive when feature sets are very large, it is a predictable method because it is stable and flexible when used on diverse datasets.

C. XGBoost

Extreme Gradient Boosting (XGBoost) is a gradient boosting model, which is strived to construct decision trees one by one, with the successively constructed trees trying to correct the errors of earlier ones. This amplifying process provides XGBoost with the capability to pay more attention to instances that are hard to classify, and that is significant in other tasks such as hate speech detection where subtle manifestations can be overlooked by less sophisticated models. Input can be textual features, e.g. n-grams or TF-IDF vectors, which are obtained through preprocessing. The algorithm uses regularization methods to prevent overfitting and its parallelized form can be easily computed. XGBoost adapts well to high dimensional data and is able to determine non-linear term-label relationships [28]. It has parameter tuning flexibility, which qualifies it as a successful contender on fine-grained classification problems, but adds complexity to optimization.

D. Logistic Regression (LR)

Simple but powerful statistical model, which is common in text classification, is the Logistic Regression (LR). It estimates the likelihood of a category using a logistic regression of input characteristics. The features under hate speech detection are usually presented as TF-IDF or bag-of-words, the weights of each word depending on the level of significance. LR provides a probability value of each class and is therefore readable and convenient in binary and multi-class problems. Its advantages include its computational efficiency, scalability and a feature which does give an insight into feature importance. Non-linear relationships or contextual dependencies in text may not be well represented in text though and this is a limitation of this as a linear model. In spite of such shortcomings, LR is commonly employed as a baseline, as well as a building block in ensemble methods to improve classification results.

E. Stacked Ensemble Model

The Stacked Ensemble model combines several base classifier to enhance robustness and minimize the weakness of each. The processes in this method involve the individual training of such models as SVM, Random Forest, XGBoost, and Logistic Regression, and their predictions are fed into a meta-classifier, which gives the final result. This combined step enables the system to take advantage of the strengths of both models such as linear separation of logistic regression, decision boundaries of SVM, and feature interaction treatment of tree-based methodology. Stacking minimises errors by the drawbacks of single algorithms, and gives more robust behaviour with varied and noisy data. Ensemble learning, though more complex to design and compute, offers a useful approach to the nature of hate speech detection that is difficult to deal with due to the ambiguity of language, context, and extreme diversity. The table 4 shows the comparative strengths and weaknesses of key ML models applied to hate speech detection. It highlights how each model performs under different conditions, emphasizing the superior robustness of ensemble methods.

F. Convolutional Neural Networks (CNNs)

CNNs are effective at capturing local patterns and n-gram–like structures in text. When combined with pre-trained embeddings such as GloVe or Word2Vec, CNN-based architectures have achieved strong performance in detecting offensive and hateful content. They are computationally efficient compared to RNNs, but struggle with capturing long-term dependencies.

G. Recurrent Neural Networks (RNNs) and LSTMs

RNNs, particularly Long Short-Term Memory (LSTM) and Bidirectional LSTM (BiLSTM) models, are widely applied for sequential text data. They capture contextual dependencies and word order, which is crucial for identifying implicit hate or sarcasm. However, they require large annotated datasets and are computationally expensive for long sequences.

H. Transformer Models (BERT, RoBERTa, DistilBERT, XLM-RoBERTa)

Transformer-based architectures are the current state-of-the-art in hate speech detection. Models like BERT and RoBERTa capture bidirectional context and semantic meaning, outperforming earlier RNN/CNN models. DistilBERT offers computational efficiency with reduced parameters, while XLM-RoBERTa supports multilingual hate speech detection, addressing low-resource languages. Despite their high accuracy, these models are resource-intensive, raising challenges for real-time deployment.

I. Advanced and Recent Transformer Models (2024–2026)

While BERT, RoBERTa, and XLM-RoBERTa have dominated hate speech research, more recent architectures have demonstrated superior performance and deserve coverage in any current survey.

DeBERTa [24]: He et al. proposed DeBERTa (Decoding-enhanced BERT with Disentangled Attention), which separately encodes content and positional information. Applied to hate speech tasks in 2024, DeBERTa achieves 1–3% accuracy improvements over RoBERTa on standard benchmarks, making it one of the strongest fine-tuned encoders currently available.

ELECTRA [25]: Clark et al. introduced ELECTRA, which uses a replaced-token detection objective instead of masked language modelling, making pre-training 4× more sample-efficient. Competitive hate speech detection results have been reported with substantially lower fine-tuning compute compared to BERT.

HateBERT [26]: Caselli et al. released HateBERT, a BERT model pre-trained specifically on the Reddit Abusive Language corpus (~1.5M posts). Its domain adaptation gives it an edge over general BERT, particularly for implicit hate and community-specific offensive language, with F1 improvements of 2–4% on abusive language benchmarks.

LLaMA-based and Instruction-Tuned LLMs [27]: Models such as LLaMA-2, Mistral, and GPT-4 have been evaluated as zero-shot and few-shot hate speech classifiers. Instruction-tuned variants (e.g., LLaMA-2-Chat) demonstrate strong cross-domain generalization without fine-tuning, though inference cost remains high for real-time deployment.

Multimodal Hate Speech Detection and Vision-Language Models (VLMs) [28]: Hate speech increasingly manifests in image-text form (memes, videos). Models such as CLIP, ViLBERT, and FLAVA have been applied to the HatefulMemes benchmark (Meta AI, 10,000 samples) and MultiOFF dataset. VLMs that jointly encode visual and linguistic semantics represent a critical frontier, yet no current model exceeds human performance (~84.7% AUC) on HatefulMemes.

Incorporating these models broadens the survey's comprehensiveness and aligns it with the 2024–2026 state of the art.

IV. Dataset Comparison

Understanding benchmark characteristics is essential for interpreting and comparing model performance across the reviewed literature.

Davidson et al. (2017) [29]: 24,783 tweets; 3-class (hate / offensive / neither); 77% offensive class highly imbalanced; one of the most cited English hate speech benchmarks.

HateXplain (2021) [30]: 20,148 posts from Twitter and Gab; 3-class with human rationale annotations; supports explainability research.

CrowdFlower [18]: 14,509 tweets; 3-class; used in multilabel classification studies.

Sexism vs Racism [18]: 6,492 tweets; binary hate classification; 86% single-class imbalance among the most challenging splits.

TRAC Shared Task [21]: Multilingual (English, Hindi, Bangla); 3-level aggression annotation; standard for cross-lingual evaluation.

Urdu Hate Speech Dataset [22]: 10,526 annotated Urdu tweets; first resource for Urdu; lexicon-based annotation; limited scale.

HatefulMemes [28]: 10,000 multimodal memes (image + text); AUC-based evaluation; key benchmark for multimodal detection; human AUC ~84.7%.

ETHOS [31]: 998 online comments; binary and multi-label versions; used for fine-grained hate category detection.

Key Observations: Most datasets are English-centric, limiting cross-lingual generalizability. Class imbalance is pervasive, requiring oversampling or weighted loss strategies [19]. Annotation agreement is often low (Cohen's Kappa < 0.7), introducing label noise [30]. Multilingual and multimodal benchmarks are underrepresented but growing rapidly [28].

V. Taxonomy of Hate Speech Detection Methods

A structured taxonomy is presented below to aid researchers in navigating the methodological landscape:

1. Rule-Based / Lexicon-Based Methods: Curated hate lexicons and keyword matching. High interpretability; low recall for implicit hate; poor cross-domain generalization [9].
2. Classical Machine Learning (SVM, LR, RF, XGBoost): Paired with TF-IDF or n-gram features. Efficient; best suited

for structured, limited-label settings [12][13].

3. Deep Learning (CNN, RNN, LSTM, BiLSTM): Captures local n-gram patterns (CNN) and sequential dependencies (LSTM/BiLSTM). Requires larger datasets and GPU compute [17][19][20].
4. Ensemble Learning (Stacking, Bagging, Boosting): Combines multiple models to reduce individual weaknesses. Accuracy gains of 2–5% over single classifiers consistently reported [9][16].
5. Transformer-Based Models (BERT, RoBERTa, DistilBERT, DeBERTa, HateBERT): Pre-trained bidirectional encoders; state-of-the-art for English and multilingual hate speech. Fine-tuning achieves F1 > 0.85 on most benchmarks [21][22][23][24][26].
6. Large Language Models (LLaMA, GPT-4, Mistral): Zero-shot/few-shot detection; highest cross-domain generalization; highest inference cost [27].
7. Multimodal Methods (VLMs, CLIP, ViLBERT): Handle hate in image-text combinations. Evaluated on HatefulMemes and MultiOFF benchmarks [28].

This taxonomy reveals a clear progression from lexicon-driven to LLM-powered detection, with each tier trading interpretability for contextual richness and generalization.

VI. Results And Discussion

ML models show strong baseline performance (85–94%), with SVM and RF being the most reliable. However, they struggle with context, sarcasm, and multilingual content. Deep learning models add improvements: CNNs capture local patterns well, RNNs/LSTMs model sequential dependencies for nuanced hate speech, and Transformers (BERT, RoBERTa) achieve the best contextual understanding with accuracies up to ~95%. Yet, DL methods are resource-intensive and less suited for real-time deployment. Overall, ensembles and transformers outperform single models by giving the best trade-off between accuracy and robustness, though at higher computational cost.

Critical Analysis: Why Do Transformers Outperform Classical ML?

Classical ML models rely on hand-crafted features (TF-IDF, n-grams) that fail to capture contextual semantics and long-range dependencies. The phrase 'you are so sick' may be offensive or complimentary depending on context a bag-of-words model cannot resolve this, while BERT's bidirectional attention can. Furthermore, transformer pre-training on large corpora provides implicit world knowledge that classical models must learn from scratch on small labelled datasets [24][7].

When Are Ensemble Methods Preferable?

Ensemble methods are preferable when: (1) individual models have complementary error profiles (e.g., SVM handles linear boundaries; RF manages non-linearity); (2) labelled data is limited, making deep learning prone to overfitting; (3) inference latency constraints preclude large transformers. Stacked ensembles consistently show 2–5% accuracy gains over single classifiers across the reviewed studies [9][16].

Accuracy vs. Computational Cost Trade-offs

Classical ML achieves 85–96% accuracy with minimal compute. Deep learning (CNN/LSTM) achieves 88–93% with moderate GPU requirements [17][19]. Transformers achieve 90–95%+ but require GPU fine-tuning and carry higher inference latency [23][24]. LLMs achieve competitive zero-shot performance but inference cost can be 10–100× greater than transformer baselines [27]. For real-time moderation, DistilBERT offers the best accuracy-latency trade-off [23].

Most Challenging Datasets

The Sexism vs Racism dataset (86% single-class) is the most imbalanced benchmark [18]. HatefulMemes is the hardest multimodal benchmark human performance ~84.7% AUC versus model peaks of ~80% [28]. TRAC multilingual datasets challenge cross-lingual generalization [21]. HateXplain requires simultaneous classification and rationale generation, a significantly harder task [30].

Major Unresolved Research Problems

Despite significant progress, key open challenges remain: (1) Implicit, contextual, and sarcastic hate is not reliably detected by any current model [9][26]. (2) Cross-domain generalization models trained on Twitter data fail on Reddit or YouTube [3]. (3) Low-resource languages lack annotated datasets and pre-trained models [22]. (4) Multimodal hate (memes, video) detection is in early stages [28]. (5) Explainability most models are black boxes, limiting trust in automated moderation decisions [30]. (6) Real-time scalability transformer inference latency remains a bottleneck [20].

Note: The following describes a planned future extension beyond the current survey contribution, motivated by the research gaps identified through this systematic review.

We intend to implement an advanced hybrid framework for hate speech detection by using the Ethos Hate Speech Dataset, which gives diverse real-world examples of offensive and non-offensive social media content. The proposed solution will combine both convolutional and sequential DL models with rich word embeddings to capture a wider range of linguistic

patterns. In the first stage, a CNN with GloVeembeddings will be used to detect local n-gram patterns and semantic associations that help in identifying explicit hate expressions. To address contextual and sequential dependencies, a BiLSTM integrated with Word2Vec and FastTextembeddings will be used which can handle word order, implicit hate, and multilingual slangs. FastText further provides subword-level representation, making it robust to spelling variations and code-switching, which are common in online communication. By fusing the strengths of CNN in capturing local features and BiLSTM in modeling long-term dependencies, the framework aims to outperform traditional baselines. Additionally, ensemble strategies will be explored to further improve robustness in noisy and imbalanced datasets. This approach is expected to deliver a scalable, accurate, and context-aware system for advanced hate speech detection and can be extended to multilingual environments in real-time settings.

VII. FUTURE RESEARCH DIRECTIONS

Based on the gaps identified in this survey, the following research directions are recommended:
1. Multimodal Hate Speech Detection: Future work should integrate text with image, audio, and video signals using VLMs and purpose-built multimodal transformers [28].
2. Low-Resource and Multilingual Settings: Annotated datasets for underrepresented languages (Arabic dialects, Indic, African languages) are critically needed. Cross-lingual transfer via multilingual LLMs should be systematically evaluated [22][27].
3. Implicit and Contextual Hate Detection: Reasoning-capable LLMs and datasets with implicit hate annotations (e.g., IHC dataset) are promising directions [26][27].
4. Explainable AI for Content Moderation: Attention visualization, SHAP/LIME explanations, and rationale generation (as in HateXplain [30]) should be further developed to support transparent moderation.
5. Real-Time and Lightweight Models: Knowledge distillation, quantization, and efficient transformers (TinyBERT, MobileBERT) are needed for low-latency deployment at scale [23][25].
6. LLM-based Zero-Shot Detection: Instruction-tuned LLMs offer promising zero-shot detection; their reliability under adversarial prompts and safety alignment warrant systematic study [27].
7. Bias and Fairness: Hate speech models exhibit demographic biases (e.g., higher false-positive rates for African American English). Fairness-aware training and evaluation are essential for equitable deployment [3][30].
8. Longitudinal Robustness: Continual learning and domain adaptation are needed as social media language and hate speech patterns evolve over time [3][9].

V. CONCLUSION

This survey has covered the development of real-time hate speech and toxicity detection approaches on social media, including traditional ML models, DL models, ensemble learning and transformer-based architectures. Traditional ML methods proved effective in handling sparse and high-dimensional textual features but lacked the ability to capture contextual or semantic nuances. Deep learning models such as CNNs improved local feature extraction, while RNNs and LSTMs enhanced sequential and contextual modeling, offering higher accuracy but requiring substantial data and computation. Transformers such as BERT and RoBERTa established new benchmarks by providing deep contextual understanding and multilingual flexibility, though at the cost of high resource consumption and inference latency.

ACKNOWLEDGMENT

It is with great pleasure that the authors would like to acknowledge the support offered by the Department of Computer Engineering, RamraoAdik Institute of Technology, D. Y. Patil University, in the provision of the required infrastructure and academic support in this research work. We are particularly grateful that our project guide, Smita Bharne, and co-guide, Vaibhav Narawade, were of constant support, offered us useful insights and motivation in the process of developing and finishing this paper.

We also wish to thank all the members of the faculty and our classmates who gave their input and constructive criticism throughout the process of this work. They have supported this research in a very significant way.

REFERENCES

[1] M. A. Hisseine, D. Chen, and X. Yang, “The Application of Blockchain in Social Media: A Systematic Literature Review,” *Appl. Sci.*, vol. 12, no. 13, p. 6567, June 2022, doi: 10.3390/app12136567.
[2] T. Iqbal, M. Khan, K. Taveter, and N. Seyff, “Mining Reddit as a New Source for Software Requirements,” in *2021 IEEE 29th International Requirements Engineering Conference (RE)*, Notre Dame, IN, USA: IEEE, Sept. 2021, pp. 128–138. doi: 10.1109/RE51729.2021.00019.
[3] Anjum and R. Katarya, “Hate speech, toxicity detection in online social media: a recent survey of state of the art and

opportunities," *Int. J. Inf. Secur.*, vol. 23, no. 1, pp. 577–608, Feb. 2024, doi: 10.1007/s10207-023-00755-2.
[4] T. Garg, S. Masud, T. Suresh, and T. Chakraborty, "Handling Bias in Toxic Speech Detection: A Survey," *ACM Comput. Surv.*, vol. 55, no. 13s, pp. 1–32, Dec. 2023, doi: 10.1145/3580494.
[5] C. Arcila-Calderón, J. J. Amores, P. Sánchez-Holgado, and D. Blanco-Herrero, "Using Shallow and Deep Learning to Automatically Detect Hate Motivated by Gender and Sexual Orientation on Twitter in Spanish," *Multimodal Technol. Interact.*, vol. 5, no. 10, p. 63, Oct. 2021, doi: 10.3390/mti5100063.
[6] O. Stoika and V. Pitovka, "Social media and language innovation: a study of slang, abbreviations, and emoticons," *СучасніДослідження З Іноземної Філології*, no. 1(27), pp. 226–238, 2025, doi: 10.32782/2617-3921.2025.27.226-238.
[7] H. Zhang and M. O. Shafiq, "Survey of transformers and towards ensemble learning using transformers for natural language processing," *J. Big Data*, vol. 11, no. 1, p. 25, Feb. 2024, doi: 10.1186/s40537-023-00842-0.
[8] M. B. Shaik and D. S. Rao, "A Comprehensive Review of the Literature and Meta-analysis of Approaches and Datasets for Investigating Textual Hate Speech Identification," in *Intelligent Computing and Communication*, vol. 1241, K. S. Raju, R. Senkerik, T. K. Kumar, M. Sellathurai, and V. Naresh Kumar, Eds., in Lecture Notes in Networks and Systems, vol. 1241. , Singapore: Springer Nature Singapore, 2025, pp. 527–544. doi: 10.1007/978-981-96-1267-3_43.
[9] N. S. Mullah and W. M. N. W. Zainon, "Advances in Machine Learning Algorithms for Hate Speech Detection in Social Media: A Review," *IEEE Access*, vol. 9, pp. 88364–88376, 2021, doi: 10.1109/ACCESS.2021.3089515.
[10] A. Toktarova*et al.*, "Hate Speech Detection in Social Networks using Machine Learning and Deep Learning Methods," *Int. J. Adv. Comput. Sci. Appl.*, vol. 14, no. 5, 2023, doi: 10.14569/IJACSA.2023.0140542.
[11] I. Aljarah*et al.*, "Intelligent detection of hate speech in Arabic social network: A machine learning approach," *J. Inf. Sci.*, vol. 47, no. 4, pp. 483–501, Aug. 2021, doi: 10.1177/0165551520917651.
[12] S. Abro, S. Shaikh, Z. Hussain, Z. Ali, S. Khan, and G. Mujtaba, "Automatic Hate Speech Detection using Machine Learning: A Comparative Study," *Int. J. Adv. Comput. Sci. Appl.*, vol. 11, no. 8, 2020, doi: 10.14569/IJACSA.2020.0110861.
[13] S. Das, K. Bhattacharyya, and S. Sarkar, "Performance Analysis of Logistic Regression, Naive Bayes, KNN, Decision Tree, Random Forest and SVM on Hate Speech Detection from Twitter," *Int. Res. J. Innov.Eng. Technol.*, vol. 07, no. 03, pp. 07–03, 2023, doi: 10.47001/IRJIET/2023.703004
[14]A. Silva and N. Roman, "Hate Speech Detection in Portuguese with Naïve Bayes, SVM, MLP and Logistic Regression," in *Anais do Encontro Nacional de Inteligência Artificial e Computacional (ENIAC 2020)*, Brasil: Sociedade Brasileira de Computação - SBC, Oct. 2020, pp. 1–12. doi: 10.5753/eniac.2020.12112.
[15] Y. V. S. Rohith and M. Amanullah, "Improving the Accuracy by Comparing the Gaussian Naive Bayes Algorithm and Logistic Regression for Predicting Hate Speech Recognition," in *2024 15th International Conference on Computing Communication and Networking Technologies (ICCCNT)*, Kamand, India: IEEE, June 2024, pp. 1–5. doi: 10.1109/ICCCNT61001.2024.10723838.
[16] N. P. Damayanti, D. E. Prameswari, W. Puspita, and P. S. Sundari, "Classification of Hate Comments on Twitter Using a Combination of Logistic Regression and Support Vector Machine Algorithm," *J. Inf. Syst. Explor. Res.*, vol. 2, no. 1, Jan. 2024, doi: 10.52465/joiser.v2i1.229.
[17] P. K. Roy, A. K. Tripathy, T. K. Das, and X.-Z. Gao, "A Framework for Hate Speech Detection Using Deep Convolutional Neural Network," *IEEE Access*, vol. 8, pp. 204951–204962, 2020, doi: 10.1109/ACCESS.2020.3037073.
[18] M. U. S. Khan, A. Abbas, A. Rehman, and R. Nawaz, "HateClassify: A Service Framework for Hate Speech Identification on Social Media," *IEEE Internet Comput.*, vol. 25, no. 1, pp. 40–49, Jan. 2021, doi: 10.1109/MIC.2020.3037034.
[19] S. Riyadi, A. DivayuAndriyani, and S. Noraini Sulaiman, "Improving Hate Speech Detection Using Double-Layers Hybrid CNN-RNN Model on Imbalanced Dataset," *IEEE Access*, vol. 12, pp. 159660–159668, 2024, doi: 10.1109/ACCESS.2024.3487433.
[20]S. Khan *et al.*, "BiCHAT: BiLSTM with deep CNN and hierarchical attention for hate speech detection," *J. King Saud Univ. - Comput. Inf. Sci.*, vol. 34, no. 7, pp. 4335–4344, July 2022, doi: 10.1016/j.jksuci.2022.05.006.
[21] A. Baruah, K. Das, and F. Barbhuiya, "Aggression Identification in English, Hindi and Bangla Text using BERT, RoBERTa and SVM," in *Proceedings of the Second Workshop on Trolling, Aggression and Cyberbullying*, 2020.
[22] R. Ali, U. Farooq, U. Arshad, W. Shahzad, and M. O. Beg, "Hate speech detection on Twitter using transfer learning," *Comput. Speech Lang.*, vol. 74, p. 101365, July 2022, doi: 10.1016/j.csl.2022.101365.
[23] R. T. Mutanga, N. Naicker, and O. O, "Hate Speech Detection in Twitter using Transformer Methods," *Int. J. Adv. Comput. Sci. Appl.*, vol. 11, no. 9, 2020, doi: 10.14569/IJACSA.2020.0110972.
[24] P. He, X. Liu, J. Gao, and W. Chen, "DeBERTa: Decoding-enhanced BERT with Disentangled Attention," in Proc. ICLR 2021. arXiv:2006.03654.

[25] K. Clark, M.-T. Luong, Q. V. Le, and C. D. Manning, "ELECTRA: Pre-training Text Encoders as Discriminators Rather Than Generators," in Proc. ICLR 2020. arXiv:2003.10555.
[26] T. Caselli, V. Basile, J. Mitrovic, and M. Granitzer, "HateBERT: Retraining BERT for Abusive Language Detection in English," in Proc. 5th Workshop on Online Abuse and Harms (WOAH), ACL 2021, pp. 17–25.
[27] T. Huang, C. Kwak, and F. Chen, "Is LLM-as-a-Judge Robust? Investigating Universal Adversarial Attacks on Zero-shot LLM Assessment," arXiv:2402.14016, 2024.
[28] H. Kiela et al., "The Hateful Memes Challenge: Detecting Hate Speech in Multimodal Memes," in Proc. NeurIPS 2020, pp. 2611–2624.
[29] T. Davidson, D. Warmsley, M. Macy, and I. Weber, "Automated Hate Speech Detection and the Problem of Offensive Language," in Proc. ICWSM 2017, vol. 11, pp. 512–515.
[30] B. Mathew, P. Saha, S. M. Yimam, C. Biemann, P. Goyal, and A. Mukherjee, "HateXplain: A Benchmark Dataset for Explainable Hate Speech Detection," in Proc. AAAI 2021, pp. 14867–14875.
[31] I. Mollas, Z. Chrysopoulou, S. Karlos, and G. Tsoumakas, "ETHOS: A Multi-Label Hate Speech Detection Dataset," Complex Intell. Syst., vol. 8, pp. 4663–4678, 2022, doi: 10.1007/s40747-021-00608-2.

AI-Driven Cybersecurity in Smart Agriculture: Threats, Challenges, and Secure System Architectures

Chintan Hirpara[1],Khushi Kundariya[2],Mahipal Chothani[3],Deep Solanki[4],Jigar Gajjar[5]

LokJagruti Kendra University, Ahmedabad, GJ, India[1,2,3]
TechD Cybersecurity Limited, Ahmedabad, GJ, India[4,5]

chintanhirpara2707@gmail.com[1],kundariyakhushi0404@gmail.com[2], mahipalchothani5@gmail.com[3],
deepsolanki5799@gmail.com[4],cyberian.jg@gmail.com[5]

Abstract:Farming has changed dramatically over the past two decades. What once relied on a farmer's experience and physical labour now runs on sensors, AI algorithms, and cloud platforms — and that shift, while genuinely productive, has dragged agriculture into a cybersecurity threat landscape it was never designed to handle. IoT devices sitting in open fields, wireless networks spanning rural distances, and AI models quietly steering irrigation and harvest decisions each represent exploitable entry points for motivated attackers. Worse, most farmers lack cybersecurity training, devices run on default credentials, and no universal security standard exists for agricultural technology.

This paper addresses that gap through three contributions: a STRIDE-based threat analysis mapping adversarial risks to specific farm system components; a five-layer Security-by-Design architectural framework integrating AI-driven defences from physical surveillance through to supply-chain traceability; and algorithm-level specifications for the AI mechanisms within that framework. Cybersecurity in agriculture is no longer optional — it is a national food security requirement.

Keywords:Smart Agriculture, Cybersecurity, IoT, AI, STRIDE, Security-by-Design, Food Security, Machine Learning.

I. INTRODUCTION

Ask a farmer from the 1980s what tools they relied on, and the answer would centre on experience, physical intuition, and relatively simple mechanical equipment. Ask a farmer today, and the answer looks very different: soil moisture sensors, satellite imagery, AI-powered disease-detection dashboards, automated irrigation controllers, and cloud-based yield forecasting platforms [1], [7], [9]. This is not merely a change in tools — it is a change in the nature of farming itself.

Modern agriculture now operates on two parallel planes. The first is the familiar physical plane of land, crops, water, and machinery. The second is a digital plane of data streams, predictive models, networked controllers, and supply-chain records. What makes this arrangement both powerful and precarious is that the digital plane directly governs the physical one. A corrupted irrigation schedule can kill a crop. A ransomware attack on a farm management system can halt operations during the narrow window of harvest. A poisoned yield-prediction model can push a farmer to plant the wrong crop at the wrong density [2], [5], [12]. In each case, a digital threat produces a very physical consequence.

The problem is compounded by the realities of agricultural deployment. Many of the IoT devices used on farms are installed with factory-default credentials and receive firmware updates infrequently, if at all [6], [13]. The wireless networks linking these devices to farm management systems often traverse kilometres of open countryside with minimal physical security. Farmers and farm workers, while highly skilled in agronomy, rarely have formal cybersecurity training, and there is no universally adopted security standard for agricultural technology [4], [21]. Agriculture is meanwhile increasingly classified as critical national infrastructure [9], [22], meaning that a successful large-scale attack could destabilise food supply chains well beyond a single farm or region [2], [3].

http://doi.org/10.63766/spujstmr.26.000090

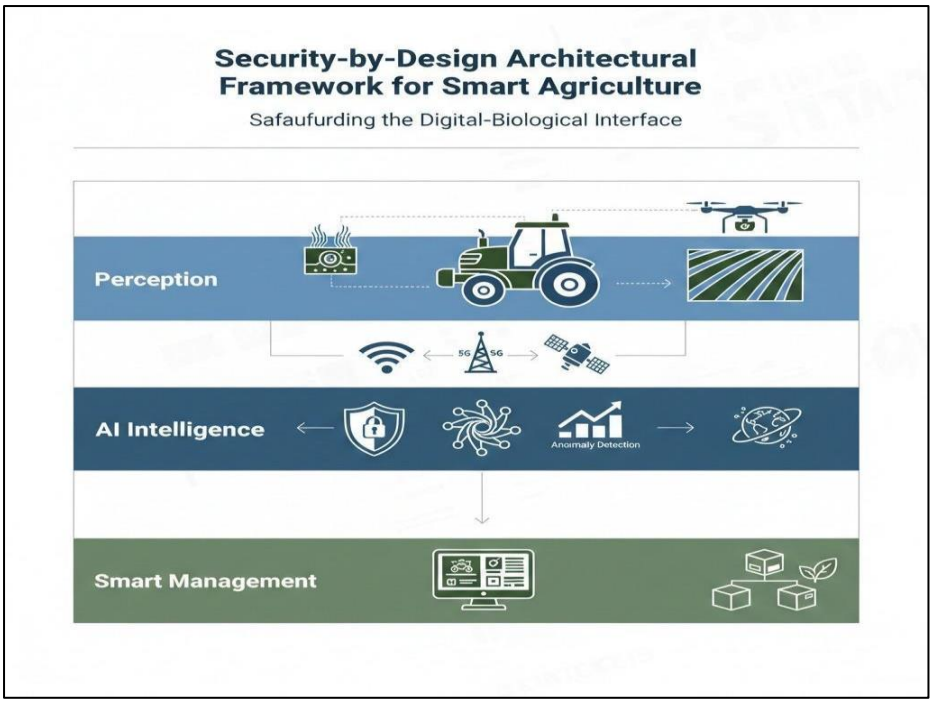

Fig.1. Security By Design Architectural Framework For Smart Agriculture

These technologies allow for early detection of crop disease, precision fertilizer application, automated irrigation control and forecasting of crop yields – all based on real-time data instead of relying on human intuition alone [11], [12]. Historically, experience was the primary deciding factor in agricultural decision making, however, with the fact that agriculture will be increasingly dependent upon ongoing collection of data and the use of digital intelligence in maximizing agricultural production and resource use [4], the need for cybersecurity in agriculture has greatly increased due to the higher levels of reliance on digital interconnectedness across agricultural production systems, along with the use of real-time data transfer on a daily basis, as a disruption of either data integrity or system availability can have a direct and negative impact on the decisions made regarding producing agriculture and farming, therefore making agriculture vulnerable to cyber threats [2], [5].

Smart agriculture, which is also known as precision agriculture, relies on things such as the Internet of Things and artificial intelligence to improve agriculture [1], [6]. It also relies on cloud computing and remote sensing and automation [4], [7]. All these help with farm management, ensuring that they are sustainable [9].

Sensors are placed in the fields where crops are grown and where animals are raised [7]. These sensors gather information all the time about the soil, how well crops are growing and what the weather is like and how healthy the animals are [10]. AI-based analytics examines this information to help determine how much food we can grow, find diseases in plants, use resources efficiently, and automate farming processes [6], [14]. These new technologies help farmers grow food, save money, cause less harm to the environment, and cope with changes in the weather [2], [3].

The truth is that there are some positive aspects of smart agriculture systems. However, people are using more and more of things. This means that there are dangers regarding cybersecurity [2], [5]. Smart agriculture systems employ many devices that are connected to the internet [3]. They communicate with each other through wireless communication and cloud technology [8]. All these can be hacked by hackers [13]. There is much information that is transmitted and stored through computers [12]. This information includes things such as how healthy the soil is how well the crops are doing what the weather is like, and information about the animals on the farm [3]. Smart agriculture systems are vulnerable because of this [5]. There are security mechanisms that are not good enough for these systems [13]. They can be vulnerable to threats such as people gaining unauthorized access to them, data breaches, and entering incorrect information [2]. Additionally, automated controls can be compromised [5]. The entire system can be brought to a complete halt [13]. This can be done because of security mechanisms. There are security mechanisms that are not good enough. These can cause much trouble [2], [5].

When these events occur, they may cause decisions to be made about farming people losing money things not working properly and farmers not trusting digital farming solutions because of inadequate security mechanisms

[2], [3]. Inadequate security mechanisms can have consequences [5].
Cybersecurity issues in farming are an issue [5]. The nature of farming makes these issues even larger [13]. Many devices that rely on the internet lack much power to process information and are also located in open areas or far from cities [6], [8].
Cybersecurity issues in agriculture are also exacerbated by the fact that farmers lack sufficient knowledge about cybersecurity [2]. Farmers tend to use the default settings of their devices [13]. They fail to update their software on a regular basis [5]. Cybersecurity issues in agriculture are further exacerbated by the fact that equipment is being used in conjunction with new digital technologies [8]. This makes it difficult to adopt security measures, for agriculture cybersecurity [5]. Agriculture is a part of our countrys critical infrastructure [3]. The issues with agriculture pose threats, to our food security, the economy, and the ability of our country to develop in a sustainable manner [2], [3]. Agriculture is a very important aspect. The threats posed are very serious. They may impact our food security, economic stability, and the development of our country over time which is why agriculture is a concern [8].

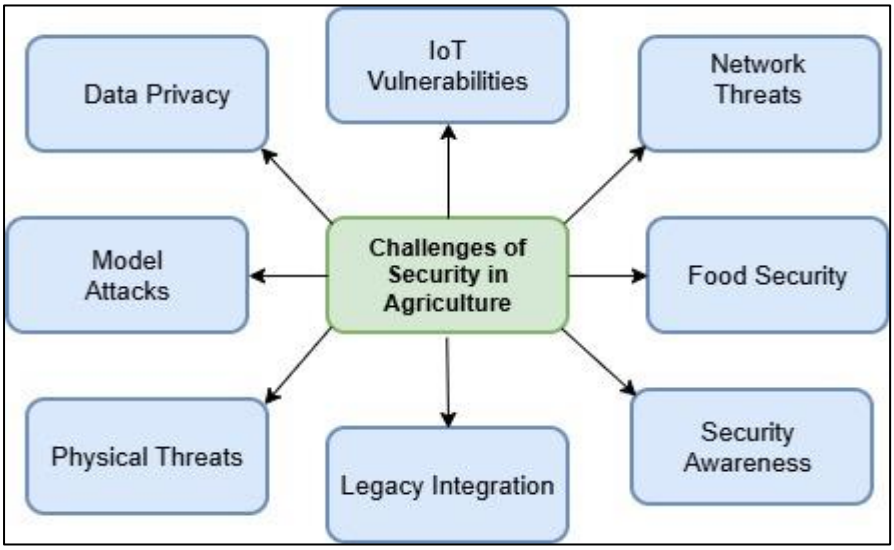

Fig.2 Challenges of Security Agriculture

Artificial intelligence is a very important thing when it comes to handling cybersecurity issues [6]. Artificial intelligence assists us in creating security systems that are capable of monitoring what is happening on our systems and devices all the time [5]. These systems are able to detect things that are not normal and prevent hackers from attacking us before they cause any harm [6].
Machine learning is a component of intelligence that assists us in detecting threats before they cause any harm [6]. It does this by identifying patterns that are not normal [13]. Intelligence and machine learning assist us in preventing bad things from happening to our systems and data [5]. Artificial intelligence is a tool that assists us in staying safe from cyber attacks [6]. AI-powered computer vision systems make our physical security systems even more secure by identifying people who are not supposed to be there or trying to tamper with agricultural equipment and systems that monitor our livestock [14].

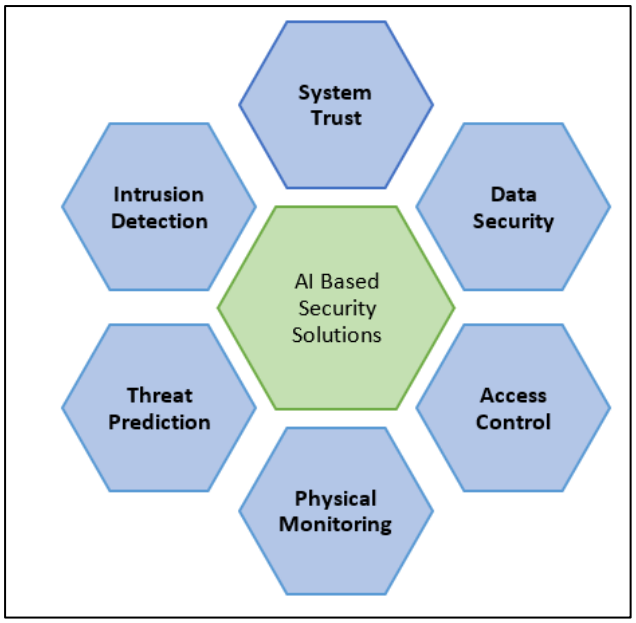

Fig.3 AI Based Security Solutions

The AI and blockchain technologies ensure that information is shared in a safe manner, and we are able to trace where everything comes from and that the information is right from the beginning of the agricultural supply chain to the end [12], [26], [27].

Smart and connected agriculture is very beneficial to agriculture because it assists farmers in growing food and doing so in a manner that is better for the environment [1], [9]. Smart and connected agriculture also assists farmers in working in an efficient manner [7]. Smart and connected agriculture will not be able to function in the right manner if there is not good cybersecurity [5]. This means that agriculture and artificial intelligence and cybersecurity must all work together [6]. When they do, we are able to develop trustworthy digital agricultural systems that can resist difficulties [8], [28]. Smart and connected agriculture is very important. Cybersecurity is also very important to smart and connected agriculture [2], [3].

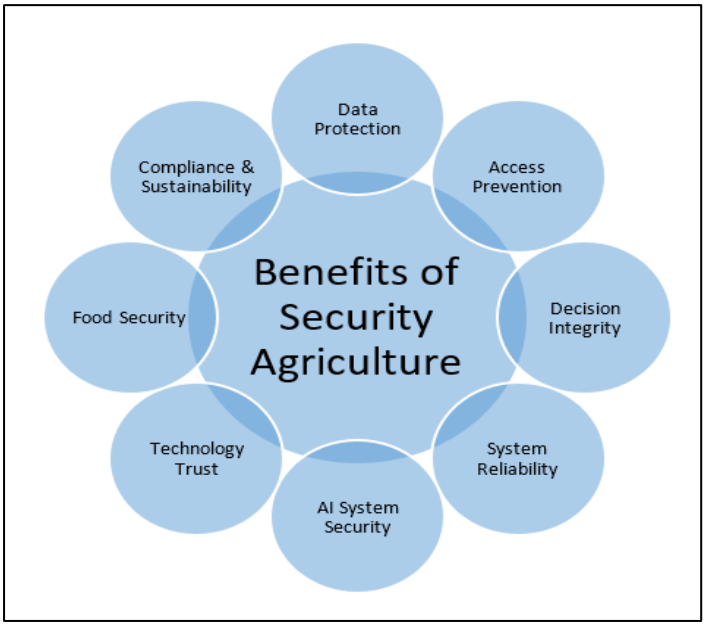

Fig.4 Why Security is Essential Agriculture

II. LITERATURE SURVEY AND RESEARCH GAP ANALYSIS

2.1 How Smart Agriculture Evolved

Farming's digital journey did not happen overnight. Early milestones in the late 1990s and early 2000s included computerised farm records, satellite-based remote sensing, and digital weather stations — tools that let farmers collect and examine data more systematically than ever before [9], [11]. The following decade brought wireless sensor networks and the first commercially viable IoT platforms, which made it possible to monitor field conditions continuously without walking the land [1], [7]. More recently, AI-powered analytics have layered intelligence on top of raw sensor data, enabling automated decision-making for tasks such as irrigation timing, pest-pressure forecasting, and livestock health monitoring [6], [8].

The result is a sector that, in its most advanced form, resembles a distributed computing network as much as it resembles a traditional farm. Precision irrigation reduces water consumption while holding or improving yields [7], [9]. Computer-vision models mounted on drones detect crop disease at a stage when intervention is still straightforward [8], [13]. Wearable sensors on livestock flag health anomalies before they become costly outbreaks [7], [25]. Blockchain-backed supply-chain records trace produce from soil to shelf with a transparency that was previously impossible [23], [26]. The productivity and sustainability arguments for these technologies are well established [19], [21]. Their security implications are considerably less studied.

2.2 What Existing Frameworks Get Right — and Where They Fall Short

A growing body of work has begun mapping the cybersecurity challenges of smart agriculture, and it is worth understanding what that work has and has not accomplished before introducing another framework.

Campoverde-Molina and Lujan-Mora [5] carried out a thorough systematic literature review covering IoT vulnerabilities, network exposure, and data security gaps in precision agriculture. Their contribution is valuable as a problem inventory, but they stopped short of proposing an integrateddefence architecture. Yazdinejad et al. [12] catalogued attack surfaces across a wide range of precision farming systems — sensors, communication links, AI models, and cloud platforms — with commendable technical specificity. Again, however, the analysis did not extend to a unified security framework with implementation guidance. Ferrag et al. [25] surveyed AI-based security mechanisms for smart farming and identified many of the same building blocks we adopt here, but they treated those mechanisms independently rather than as layers in a coordinated architecture. Perhaps closest to the present work is the STRIDE-based threat modelling exercise of Al Asif et al. [28], which provided a principled structure for enumerating threats in IoT-enabled precision agriculture. Their work focuses on threat identification; ours extends that foundation to a mitigation architecture with specified AI components.

Three gaps run consistently through the existing literature. First, cybersecurity mechanisms are almost always discussed in isolation. No prior work integrates physical security, device authentication, network defence, model integrity assurance, and supply-chain protection into a single, layered architecture. Second, proposed AI defences are routinely described at a conceptual level — 'use anomaly detection,' 'apply machine learning' — without specifying algorithms, features, or measurable targets. Third, the connection between agricultural cybersecurity and national food security strategy has not been formally framed within a Security-by-Design context that could inform policy as well as engineering.

2.3 Security Challenges in Smart Agriculture Systems

2.3.1 Data Collection and Integrity

Today's farms generate an enormous volume of sensitive data: granular soil chemistry readings, hourly crop growth measurements, weather records, and detailed livestock health logs [7], [10]. That data feeds AI models which, in turn, drive decisions with real financial and agronomic consequences. An adversary who can alter even a modest fraction of incoming sensor readings — pushing soil moisture values below actual levels to trigger unnecessary irrigation, for instance — can cause direct operational harm while remaining undetected for days or weeks [2], [5], [9]. The challenge is that the sheer volume of data collected makes manual verification impractical, which is exactly why automated integrity monitoring is essential.

2.3.2 IoT Device Vulnerabilities

A mid-sized precision farm might deploy hundreds of IoT devices: soil sensors, climate monitors, livestock trackers, irrigation controllers, aerial drones, and edge gateways [1], [6], [12]. Many of these devices are computationally limited,

which constrains the cryptographic protocols they can run. Many are installed outdoors, far from physical oversight, making them accessible to anyone who walks into a field. And many stay in service for years without firmware updates, accumulating unpatched vulnerabilities over time [6], [13]. When a single compromised device is used as a beachhead to move laterally across a farm network, the consequences can extend well beyond that one sensor — from ransomware deployment on the central management system to DDoS attacks launched from the farm's own equipment [6], [12], [16].

2.3.3 Insecure Wireless Communications

Rural wireless networks — the arteries that carry data between field devices and farm management systems — operate across open terrain where signal interception is straightforward [7], [16], [17]. Without proper encryption and authentication at the network layer, an attacker positioned within radio range can eavesdrop on data transmissions, spoof sensor identities to inject false readings, or conduct Man-in-the-Middle attacks to intercept and alter commands before they reach automated field equipment [5], [12]. These threats are not theoretical; the relative openness of rural environments makes them considerably easier to execute than comparable attacks in urban or industrial settings.

2.3.4 Adversarial Attacks on AI Models

The AI models that power smart agriculture — forecasting yields, detecting plant disease, scheduling irrigation — are themselves attack targets, and this is perhaps the least appreciated risk in the sector. Data poisoning introduces corrupted or strategically crafted records during the training phase, causing a model to learn systematically wrong associations [4], [6], [12]. Evasion attacks craft inputs that reliably cause a deployed model to misclassify at inference time. Model inversion attacks attempt to extract sensitive information embedded in a trained model's parameters [18]. What makes these threats particularly insidious is that they can degrade model performance gradually and silently, producing decisions that seem plausible until the accumulated damage becomes visible in crop outcomes weeks or months later [3], [9].

2.3.5 Physical Security of Field Equipment

It is easy to focus on cyber threats and overlook the physical dimension, but the two are tightly coupled in agricultural settings. A sensor that is physically tampered with can be made to transmit false data while appearing operationally normal. A gateway whose housing is opened and whose cryptographic keys are extracted becomes a permanent backdoor into the farm network [2], [5], [13]. Devices stolen from the field carry configuration data and credentials that an attacker can analyse offline [12], [16]. Physical and cyber security must therefore be addressed together rather than sequentially.

2.3.6 Integrating Legacy Equipment

Many farms have invested heavily in equipment that was designed before network connectivity was standard — tractors with proprietary control systems, irrigation infrastructure with analogue controllers, storage facilities with minimal instrumentation [12], [19]. Connecting this legacy equipment to modern IoT and AI platforms is often necessary for operational continuity, but it creates security gaps that neither the old equipment nor the new platforms were designed to handle. Legacy systems frequently lack support for contemporary encryption standards, do not accept software patches, and expose network interfaces that were never intended to face an adversarial environment [3], [6], [20].

2.3.7 Human Factors and Awareness Gaps

Technical vulnerabilities do not exist in a vacuum; they are exploited through human behaviour. Farmers and agricultural workers who use default device passwords, skip software updates, or share network access credentials inadvertently create pathways that even moderately skilled attackers can follow [4], [10], [21]. The challenge here is not carelessness — it is that cybersecurity literacy has simply not been part of agricultural training, and the available guidance is rarely tailored to the specific context of farm operations [6], [12]. Awareness programmes and practical, farm-specific security checklists are a necessary complement to the technical measures discussed elsewhere in this paper.

2.3.8 Food Security at the National Scale

The consequences of agricultural cyber-attacks do not stop at the farm gate. Agriculture is formally classified as critical national infrastructure in many countries [9], [10], [22], which means that coordinated attacks on farming systems can have cascading effects on food supply chains, rural employment, commodity prices, and ultimately consumer access to food [1], [3], [23]. The potential for this kind of large-scale, cross-sector impact is what lifts agricultural cybersecurity from a niche technical concern to a matter of national strategy. Any serious policy framework for critical infrastructure

protection needs to include smart agriculture explicitly, with associated minimum security standards and incident reporting obligations.

TABLE I
USE CASES OF AGRITECH AND ITS BENEFITS

Agritech Use Case	**Technology Involved**	**Benefits**
Precision Irrigation [1], [7], [9]	IoT sensors, AI analytics	Optimizes water usage, reduces wastage, and improves crop yield
Crop Disease Detection [8], [13]	AI, computer vision, drones	Early detection of plant diseases and reduced crop loss
Smart Soil Monitoring [4], [9]	IoT sensors, cloud platforms	Real-time soil analysis and efficient nutrient management
Livestock Monitoring [7], [25]	Wearable sensors, IoT, AI	Improves animal health, productivity, and disease prevention
Automated Farming Machinery [20], [24]	Robotics, automation	Reduces labor dependency and increases operational efficiency
Smart Supply Chain Management [23], [26]	Blockchain, cloud computing	Enhances traceability, transparency, and reduces post-harvest losses
Weather Forecasting Systems [19], [21]	AI, big data analytics	Improves climate-resilient farming and risk management
Smart Greenhouse Management [4], [8]	IoT, automation, AI	Maintains optimal growth conditions and increases crop quality
Yield Prediction Systems [19], [21]	Machine learning, data analytics	Accurate yield estimation and better farm planning
Farm Security Monitoring [6], [27]	AI, computer vision, IoT	Prevents theft, intrusion, and equipment tampering

TABLE II
.SMART AGRICULTURE SECURITY ARCHITECTURE USING AI TECHNOLOGIES

Security Area	**AI Technology Used**	**Problem Addressed**	**Security Benefit**
Intrusion Detection [6], [22]	Machine Learning, Deep Learning	Cyberattacks on IoT networks and farm systems	Early detection of attacks and reduced system downtime
Secure Data Access [5], [23]	AI-based access control	Unauthorized access to sensitive farm data	Improved data privacy and confidentiality
Device Authentication [12], [25]	Behavioral AI models	Fake or compromised IoT devices	Ensures only trusted devices operate in the network
Threat Prediction [6], [27]	Predictive analytics, AI models	Unexpected cyberand operational risks	Proactive security and risk mitigation
Physical Asset Protection [12], [27]	Computer vision, AI drones	Theft, vandalism, and tampering of devices	Real-time monitoring and rapid

			response
Data Integrity Monitoring [3], [22]	AI anomaly detection	False data injection and data manipulation	Reliable decision-making and accurate AI outputs

II. ANALYSIS

Structured threat modelling is a discipline that forces security analysis to be systematic and reproducible rather than intuitive and ad hoc. The STRIDE methodology, originally developed at Microsoft and subsequently applied across a wide range of critical systems [28], organises threats into six categories: Spoofing (impersonating a legitimate entity), Tampering (modifying data or code), Repudiation (denying an action without the possibility of proof), Information Disclosure (exposing data to unauthorised parties), Denial of Service (making a resource unavailable), and Elevation of Privilege (gaining access beyond what is authorised).

We applied STRIDE to a representative smart agriculture architecture comprising field-deployed IoT sensors, edge gateways, wireless communication infrastructure, cloud-hosted farm management software, AI analytics modules, and supply-chain record systems. Table III presents the results of this analysis, mapping each threat category to the specific agricultural components it affects, realistic adversarial scenarios, and the countermeasures incorporated into the Security-by-Design framework described in Section IV.

A pattern that emerges clearly from this analysis is that no single countermeasure can address the full range of threats. Encrypting network communications helps against Tampering but does nothing for a Spoofing attack that uses valid credentials. Blockchain audit trails address Repudiation but do not protect against Denial of Service. This multi-dimensional threat landscape is precisely why a layered architecture, rather than a collection of point solutions, is the appropriate design response — which is what Section IV proposes.

TABLE III

INTEGRATED ANALYSIS OF AI-ENABLED SMART AGRICULTURE AND SECURITY

Aspect	**Traditional Agriculture**	**Smart Agriculture (AI + IoT)**	**Key Security Risk**	**AI-Driven Security Impact**
Decision Making [1], [6]	Experience-based	Data-driven, real-time	Data manipulation	Trustworthy analytics
Data Generation [5], [22]	Minimal	Massive sensor data	Data leakage	Secure data validation
Connectivity [12], [25]	Stand-alone systems	Highly interconnected	Network attacks	Intelligent intrusion detection
Automation Level [6], [27]	Manual operations	Autonomous systems	Control hijacking	Behavioral monitoring
Resource Usage [19], [21]	Inefficient	Precision optimized	False input injection	Accurate AI filtering
Infrastructure [5], [23]	Isolated	Cloud &IoT dependent	Cloud breaches	Encrypted access control
Physical Security [12], [27]	Human-guard	Sensor-dependent	Device tampering	AI vision surveillance
Scalability [26], [27]	Limited	Highly scalable	Attack surface growth	Predictive risk control
Farmer	Human	AI	Model	Model integrity

Dependence [6], [22]	judgement	recommendations	poisoning	checks
National Impact [9], [25]	Localized	Critical infrastructure	Food supply disruption	Resilient system design

IV. A SECURITY-BY-DESIGN ARCHITECTURAL FRAMEWORK FOR SMART AGRICULTURE

4.1 Design Principles

The Security-by-Design concept holds that security is most effective when it is built into a system from the ground up rather than bolted on after deployment [5], [12], [25]. Four principles guided the design of the framework proposed here, each derived directly from the threat landscape and the research gaps identified in the preceding sections.

Proactive rather than reactive defence: Given how quietly threats like data poisoning and behavioural anomalies can accumulate, security mechanisms need to detect and respond to threats before operational damage materialises — not after a crop has failed or a data breach has been reported.

Layered, interlocking protection: Because no single mechanism covers the full STRIDE landscape, the framework organisesdefences into five distinct layers, each addressing a specific class of threat while sharing intelligence with adjacent layers.

Autonomous operation: Most agricultural IoT infrastructure operates in remote locations with limited human oversight. Security systems must therefore be capable of running, adapting, and partially self-repairing without requiring continuous manual intervention.

Backward compatibility with legacy systems: The framework must accommodate farms that mix modern IoT-native equipment with older machinery that predates network connectivity, using standardised interfaces to bridge the security gap between the two generations.

4.2 The Five-Layer Architecture

The proposed SbD framework is structured as five functional security layers. Together they form a closed-loop architecture in which threat intelligence gathered at each layer informs the adaptive responses of the others. The layers and their properties are summarised in Table IV.

4.3 AI Security Mechanism Specifications

One of the most persistent weaknesses in the agricultural cybersecurity literature is the gap between naming an AI-based defence and actually specifying how it works. Saying that a system uses 'anomaly detection' or 'machine learning authentication' is not enough to enable replication, benchmarking, or procurement decisions. This section closes that gap for three core AI mechanisms within the SbD framework.

4.3.1 Network Anomaly Detection (Layer 3)

The network layer intrusion detection system works by learning what normal traffic looks like across every class of device on the farm network, then flagging deviations from that baseline. Three complementary algorithm families are appropriate here, depending on available labelled data and deployment constraints [6], [12], [24].

Isolation Forest is well suited to unsupervised deployment scenarios where labelled attack examples are unavailable, as is common in agricultural contexts. It isolates anomalies by randomly partitioning the feature space: records that are easy to isolate (requiring few partitions) are flagged as outliers. Random Forest classifiers offer higher accuracy when a labelled dataset of attack and normal traffic is available, producing probabilistic attack classifications that can be tuned to favour low false-positive rates. LSTM networks are particularly valuable for detecting temporal anomalies in the time-series data

generated by agricultural sensors, capturing sequential patterns — such as gradual data drift from a compromised sensor — that point-in-time classifiers would miss.

The feature set covers: packet inter-arrival time, payload length distribution per device class, protocol distribution per time window, connection duration, bytes transferred per session, and the number of unique destination addresses observed per hour. Normal behavioural baselines are established per device class (sensors, gateways, drones) during a supervised initialisation phase lasting a minimum of 30 days. Deviations exceeding a configurable threshold — we recommend three standard deviations from the learned baseline as a starting point — trigger graduated responses ranging from operator alert through to automatic network isolation of the suspect device. Target performance benchmarks, derived from comparable IoT deployment evaluations in the literature [6], [25], are: Detection Rate above 95%, False Positive Rate below 2%, and mean time to detection under 60 seconds.

4.3.2 Behavioural IoT Device Authentication (Layer 2)

Static credentials — passwords, certificates — are a necessary but insufficient foundation for IoT authentication in agricultural settings. Credentials can be stolen from a device physically, replayed from captured traffic, or remain valid long after a device has been physically compromised [4], [12]. Behavioural authentication addresses this by treating each device's operational patterns as a continuous identity signal.

The model is trained on the legitimate operational behaviour of each device class using features that include: message transmission interval (with device-specific variance bounds), payload entropy over rolling windows, communication partner set (which other devices or servers does this device talk to, and how often), time-of-day activity distribution, and power draw patterns where hardware monitoring is available. A One-Class Support Vector Machine trained exclusively on legitimate device behaviour is computationally tractable on edge hardware and produces a real-time trust score for each device. For deployments where edge computational resources allow, an Autoencoder neural network offers higher sensitivity to subtle behavioural drift at the cost of greater processing demand.

Devices whose behavioural profile exceeds the learned trust boundary are quarantined and flagged for manual review. Crucially, this mechanism catches a class of attack that credential-based systems miss entirely: the device that has valid credentials but is being remotely controlled by an attacker after firmware compromise. The change in communication behaviour — different destination addresses, altered message cadence, unusual payload sizes — triggers quarantine even though the device's certificate remains valid [5], [18].

4.3.3 Predictive Threat Assessment (Cross-Layer)

The five layers of the SbD framework generate a continuous stream of security signals: physical intrusion alerts from L1, device trust scores from L2, traffic anomaly flags from L3, data integrity verdicts from L4, and supply-chain tamper notifications from L5. A cross-layer predictive threat assessment module aggregates these signals to produce a system-wide risk picture and generate prioritised mitigation recommendations [6], [27].

The core model is a gradient-boosted decision tree ensemble — XGBoost is the recommended implementation given its performance on tabular security event data and its interpretability relative to deep learning alternatives. The model is trained on historical agricultural cyber-incident records, enriched with contextual features such as crop growth stage (certain attack types are more likely during high-value harvest windows), seasonal patterns, recent software vulnerability disclosures relevant to the farm's device inventory, and weather conditions (extreme weather events have historically correlated with increased physical intrusion attempts in agricultural settings).

Outputs are designed to be actionable by farm operators without cybersecurity expertise: a probabilistic risk score for the next 24-hour operational window expressed as low, medium, or high; a ranked list of the three highest-priority vulnerabilities currently active in the system; and specific recommended countermeasure actions phrased in operational rather than technical terms. The module interfaces directly with farm management dashboards, surfacing this intelligence alongside agronomic data so that security context is available at the same decision-making touchpoint farmers already use.

III. CONCLUSION

Smart agriculture has arrived. The technologies that make precision farming possible — IoT sensor networks, AI analytics, automated control systems, cloud platforms, blockchain supply chains — are already deployed on farms around

the world and their adoption is accelerating. The productivity and sustainability gains they enable are real. So, unfortunately, are the security risks they introduce.

This paper has tried to do three things that the existing literature has not done together in one place. The first was to apply structured threat modelling — specifically STRIDE — to a representative smart agriculture system, producing a map of adversarial threats that is systematic and reproducible rather than intuitive and anecdotal. The second was to propose a Security-by-Design framework that addresses those threats through five interlocking defensive layers, from physical asset protection in the field through to blockchain-secured supply-chain traceability. The third was to move beyond the conceptual level that characterises most AI security proposals in this domain, specifying the machine learning algorithms, feature engineering choices, and measurable performance benchmarks for the AI mechanisms inside the framework.

Running through all three contributions is a single argument: agricultural cybersecurity is not a peripheral IT problem. It is a foundational requirement for the food systems that nations depend on. When automated irrigation fails because of a cyber-attack, crops die. When AI yield models are poisoned, farmers make decisions on corrupted intelligence. When supply-chain records are tampered with, food safety and trade systems break down. The digital plane of modern agriculture is not separate from the physical plane — it governs it. Securing one means securing the other.

Several important tasks remain for future research. The AI mechanisms specified in Section 4.3 need empirical validation on real-world agricultural IoT datasets; without that, the performance benchmarks proposed here remain targets rather than demonstrated capabilities. The behavioural authentication model needs lightweight implementations that can run on the constrained hardware typical of field-deployed sensors without sacrificing detection sensitivity. Federated learning approaches deserve investigation as a way to train security models across distributed farm networks without requiring sensitive operational data to leave individual farms. And longitudinal field studies — tracking how the SbD framework performs in actual farm environments over full growing seasons, and what effect farmer training programmes have on the human-factor vulnerabilities identified in Section 2.3.7 — are needed to ground the framework's design in operational reality. These are not small tasks, but they are the right next steps

REFERENCES

[1] A. Akbar et al., “Smart farming technologies: Transforming agriculture through IoT and AI,” *Spectrum of Engineering Sciences, vol. 3, no. 5, pp. 929–936, 2025.*

[2] B. Singh and S. S. Cheema, “Cybersecurity threats and mitigation strategies in Agriculture 4.0 and 5.0: Challenges and solutions in the digital transformation of agriculture,” *Journal of Punjab Academy of Sciences, vol. 24, pp. 15–32, 2024.*

[3] P. Kaviya, M. Natarajan, and P. Ragulprasath, “Cybersecurity threats in digital agriculture: A systematic review on emerging risks and protective strategies for farmers,” *Asian Journal of Agricultural Extension, Economics & Sociology, vol. 43, no. 12, pp. 109–117, 2025.*

[4] M. U. Ali, “Smart agriculture: Integration of IoT, AI, and big data in farm management,” *Journal of Scientific Research and Reports, vol. 31, no. 12, pp. 114–122, 2025.*

[5] M. Campoverde-Molina and S. Luján-Mora, “Cybersecurity in smart agriculture: A systematic literature review,” *Computers & Security, vol. 150, Art. no. 104284, 2025.*

[6] G. Ali, M. M. Mijwil, B. A. Buruga, M. Abotaleb, and I. Adamopoulos, “A survey on artificial intelligence in cybersecurity for smart agriculture,” *Mesopotamian Journal of Computer Science, vol. 2024, 2024.*

[7] A. Cisse et al., “A smart farming management system based on IoT technologies for sustainable agriculture,” *Advances in Science, Technology and Engineering Systems Journal, vol. 9, no. 1, pp. 1–8, 2024.*

[8] T. Miller, G. Mikiciuk, I. Durlik, M. Mikiciuk, A. Łobodzińska, and M. Śnieg, “The IoT and AI in agriculture: The time is now—A systematic review of smart sensing technologies,” *Sensors, vol. 25, no. 12, Art. no. 3583, 2025.*

[9] V. R. Pathmudi, N. Khatri, S. Kumar, A. S. H. Abdul-Qawy, and A. K. Vyas, “A systematic review of IoT technologies and their constituents for smart and sustainable agriculture applications,” *Scientific African, vol. 19, Art. no. e01577, 2023.*

[10] M. Roopaei et al., “Cloud of Things in smart agriculture: Intelligent irrigation monitoring,” *IEEE Cloud Computing, 2017.*

[11] S. Bimonte et al., “Data engineering for sustainable agriculture: Developments, challenges, and case studies of a novel IoRT architecture,” *Journal of Big Data, vol. 12, no. 1, Art. no. 195, 2025.*

[12] A. Yazdinejad et al., "Security of smart farming and precision agriculture: Attacks, threats, and countermeasures," *Applied Sciences, vol. 11, no. 16, Art. no. 7518, 2021*.

[13] A. S. Ibrahim et al., "AI-IoT based smart agriculture pivot for plant disease detection and treatment," *Scientific Reports, vol. 15, no. 1, Art. no. 16576, 2025.*

[14] A. K. Panda, S. D. Thakur, Kavita, R. Rameswar, R. S. Rana, and J. Singh, "Role of cow dung, cow urine, and other cow products in natural and organic farming," *International Journal of Agriculture Sciences, vol. 10, 2018.*

[15] S. Banerjee, A. Kanaujia, and S. Sharma, "Importance of indigenous cow dung and its microbiota," *International Journal of Current Microbiology and Applied Sciences, vol. 9, 2020.*

[16] A. K. Behera and C. P. Chandrashekara, "Effect of natural farming, organic farming, and recommended package of practices on yield and economics of irrigated wheat (Triticumaestivum)," *Indian Journal of Agronomy, vol. 66, 2021.*

[17] S. Tahura et al., "Combined application of biochar compost and cow urine enhances organic rice production," *Scientific Reports, vol. 12, 2022.*

[18] P. V. Patel et al., "A case study of cow-based organic farming system adopted by the farmer of Vadal, Junagadh, Gujarat," *International Journal of Chemical Studies, vol. 7, 2019.*

[19] L. Wolfert, S. Ge, C. Verdouw, and M. J. Bogaardt, "Big data in smart farming – A review," *Agricultural Systems, vol. 153, pp. 69–80, 2017.*

[20] F. Tao, Q. Zhang, J. Liu, and D. Lai, "Digital twin driven smart manufacturing," *Journal of Manufacturing Systems, vol. 58, pp. 103–118, 2021.*

[21] M. Kamilaris, A. Kartakoullis, and F. X. Prenafeta-Boldú, "A review on the practice of big data analysis in agriculture," *Computers and Electronics in Agriculture, vol. 143, pp. 23–37, 2017.*

[22] S. Misra, A. Mukherjee, and M. Obaidat, "Security challenges and approaches in Internet of Things-based smart agriculture," *IEEE Internet of Things Journal, vol. 7, no. 6, pp. 4934–4946, 2020.*

[23] Y. Zhang, L. Wang, and H. Sun, "Blockchain-based secure data sharing for agricultural supply chains," *Future Generation Computer Systems, vol. 109, pp. 593–602, 2020.*

[24] A. Khanna and S. Kaur, "Evolution of Internet of Things (IoT) and its significant impact in the field of precision agriculture," *Computers and Electronics in Agriculture, vol. 157, pp. 218–231, 2019.*

[25] N. Ferrag, L. Maglaras, A. Derhab, and H. Janicke, "Security for smart agriculture: A survey," *IEEE Access, vol. 8, pp. 160, 000–160, 018, 2020.*

[26] S. Halder, M. R. Islam, Q. Mamun, A. Mahboubi, P. Walsh, and M. Z. Islam, "A comprehensive survey on AI-enabled secure social industrial Internet of Things in the agri-food supply chain," *Smart Agricultural Technology, vol. 11, Art. no. 100902, 2025, doi: 10.1016/j.atech.2025.100902.*

[27] A. R. Mahlous, "Security analysis in smart agriculture: Insights from a cyber-physical system application," *Computers, Materials & Continua, vol. 79, no. 3, pp. 4781–4803, 2024, doi: 10.32604/cmc.2024.050821.*

[28] M. R. Al Asif, K. F. Hasan, M. Z. Islam, and R. Khondoker, "STRIDE-based cyber security threat modeling for IoT-enabled precision agriculture systems," *arXiv, Jan. 2022.*

Comparative Study of Leading AI Models: Performance, Capabilities, Accuracy, and Applications

Chintan Hirpara1,Khushi Kundariya2,Mahipal Chothani3,Deep Solanki4,Jigar Gajjar5

LokJagruti Kendra University, Ahmedabad, GJ, India[1,2,3]
TechDCybersecurity Limited, Ahmedabad, GJ, India[4,5]

chintanhirpara2707@gmail.com[1], kundariyakhushi0404@gmail.com[2], mahipalchothani5@gmail.com[3], deepsolanki5799@gmail.com[4],cyberian.jg@gmail.com[5]

Abstract:Large language models (LLMs) have progressed from mere academic prototypes to mainstream solutions in education, healthcare, software engineering, and customer service industries. In this paper, we conduct a comparative review of the architecture, reasoning capability, code generation, creative writing capabilities, ethics, and adoption level of five of the most popular conversational artificial intelligence platforms: ChatGPT, DeepSeek, Gemini, Claude, and Grok. Instead of performing our own benchmarking experiments, we synthesize results from the literature reviewed in Section II and scores reported by the AI companies in question on a number of commonly used benchmarking platforms including MMLU, GPQA Diamond, SWE-bench Verified, HumanEval, and GSM8K, supplemented by usage metrics from independent sources. Our comparative analysis shows that none of the examined models outperforms the rest in all respects; in particular, Gemini and DeepSeek score the highest on general knowledge tests, Claude demonstrates the best performance on code-generation and reasoning about long documents, Grok stands out in difficult agentic reasoning tests and real-time information lookup tasks, and ChatGPT remains the most used one. We conclude the paper by outlining the limitations of literature-based comparative analysis and suggesting topics for future research.

Keywords:Artificial Intelligence language models, ChatGPT, DeepSeek, Gemini AI, Claude AI, Grok AI, performance comparison, accuracy evaluation, real-world applications, AI capabilities.

I. INTRODUCTION

1.1 Introduction to AI Language Models

AI has altered human interaction with technology, and AI-based language models are currently the focal point of this change. Language models enable activities such as education, medicine, banking, customer relations, coding, and even creative writing through machine learning and deep learning approaches that help comprehend, create, and process natural human language.

Contemporary language models are trained on extremely large text corpora by employing neural network structures that learn statistical correlations in language. As a result, language models are able to create fluently sounding replies, help in problem-solving, and write text from scratch. The applicability of language modelshas expanded beyond the scope of specialized applications, including virtual assistants, chatbots, code generation, automatic translations, and content creation. There have been several language models in recent years that have gained great popularity and were created with distinct architecture, training priorities, and intended use. The following paper examines five of those language models:

http://doi.org/10.63766/spujstmr.26.000091

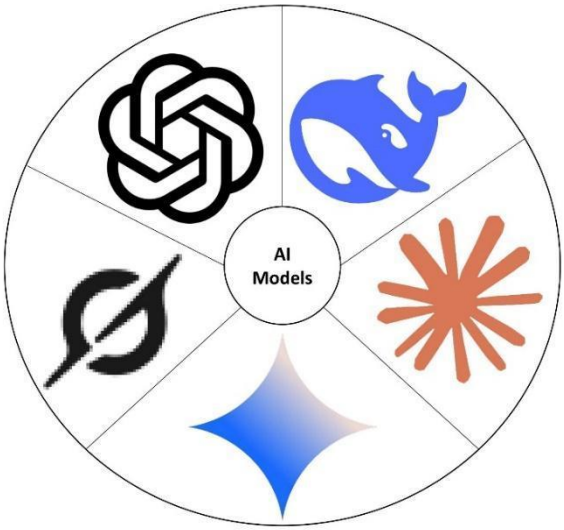

Fig.1 Logo of AI's

Several leading AI language models have emerged in recent years, each developed with distinct methodologies, architectures, and goals. This study focuses on five of the most influential AI models: ChatGPT, DeepSeek, Gemini, Claude, and Grok.

- ChatGPT (OpenAI) – focused on natural and contextual conversation and also assisting in creative and technical writing.
- DeepSeek (DeepSeek AI) – a family of open weight models that focuses on efficient and effective training while delivering good results in logical and mathematical problems.
- Gemini (Google DeepMind) – focuses on multimodal reasoning and integration with Google search and knowledge systems.
- Claude (Anthropic) – focused on safety and alignment of AI, with a focus on rational and well thought out long form responses.
- Grok (xAI) – focused on real time internet connected reasoning and integration with X platform.

As these systems continue to evolve at a fast pace, understanding their relative strengths, limitations, and appropriate use cases is important for researchers, developers, and the organizations that deploy them. This paper provides a structured, literature-grounded comparison intended to support that kind of decision-making, while being explicit about what this type of review can and cannot establish (see Section III).

1.2 Background of the Study

Modern conversational AI traces back to early rule-based natural language processing systems, which relied on hand-written rules and struggled to generalize beyond the patterns their designers anticipated. The statistical and machine-learning approaches that followed improved on this by learning from data rather than fixed rules, but still struggled with long-range context and ambiguous input.

The architecture that changed this trajectory was the Transformer, introduced by Vaswani et al. (2017) [26], which used a self-attention mechanism to let models weigh relationships between words across long spans of text far more efficiently than earlier recurrent architectures. This made it practical to train much larger models on much larger datasets, and it underlies essentially every model compared in this paper.

Building on the Transformer, OpenAI's GPT series, Google's BERT, and the reinforcement-learning-from-human-feedback (RLHF) training techniques that followed gave rise to the current generation of AI assistants [21], [24], [25]. These systems are now embedded in customer-service chatbots, automated content generation and summarization tools, programming assistants, and academic research support tools [5], [26], [30].

Despite their capability, these models still face open problems: biased training data, hallucinated or fabricated information, security exposure, and unresolved questions about appropriate use. This paper looks at how the five models under review are reported to handle these challenges, and where the published literature suggests they still fall short.

1.3 Research Objectives

The objectives of this study are to:

- Compare ChatGPT, DeepSeek, Gemini, Claude, and Grok in terms of reported applications, accuracy, and reasoning ability, using standardized benchmark scores where they are publicly available.
- Assess how each model's reported strengths translate to specific industries, including education, software development, customer support, and creative content generation.
- Identify the strengths and weaknesses attributed to each model in the literature reviewed for this study.
- Examine the ethical concerns and bias-related findings reported for each model.
- Suggest use cases that the reviewed literature and benchmark data support for each model.
- Outline open questions and a concrete agenda for future primary-benchmarking research.
  - **Active user of AI Models:**

TABLE I

ACTIVE USER OF AI MODEL

Model	Reported Metric	Approx. Figure	Source / Period
ChatGPT	Weekly active users	~900 million	OpenAI disclosure, reported Feb. 2026 [38]
ChatGPT	Monthly active users	~1.0–1.1 billion	Sensor Tower estimate, May 2026 [39]
Gemini	Monthly active users	~600–660 million	Sensor Tower / Apptopia estimate, May 2026 [39]
Claude	Monthly active users	~245 million	Sensor Tower estimate, May 2026 [39]
Grok	Mobile app market share	~15% (fastest-growing)	Apptopia estimate, Feb. 2026 [39]
DeepSeek	Regional consumer base	Concentrated in China, India, Indonesia; smaller global share	Views4You usage report, 2025 [40]

Source: https://x.com/grok/status/1900413492999414239

II. LITERATURE SURVEY

2.1 AI Language Models in Prior Research

Artificial intelligence has advanced rapidly in recent years, particularly in natural language processing (NLP). AI-driven language models are now used across education, healthcare, business, customer support, and software development, typically trained on large datasets using transformer-based architectures.

A range of studies have examined the capabilities, limitations, and ethical implications of these models, comparing them on accuracy, efficiency, reasoning ability, and bias. This survey draws on 37 sources, including the foundational Transformer paper by Vaswani et al. (2017) [26], to build a comparative picture of ChatGPT, DeepSeek, Gemini, Claude, and Grok. Section III explains in detail how these sources were selected and how the comparison itself was constructed.

2.2 Evolution of AI Language Models

Early AI models relied on rule-based systems and simple machine-learning techniques, which limited their ability to understand context and generate coherent responses. The shift to deep learning and, specifically, the Transformer architecture substantially improved contextual understanding and response generation.

Vaswani et al. (2017) [26] introduced the Transformer model, which became the foundation for essentially every modern language model. Unlike earlier recurrent architectures, transformers use self-attention, which lets them process long spans

of text efficiently. Radford et al. (2019, 2020) [25] then showed that large-scale pre-training could produce highly coherent generative text, and Brown et al. (2020) [24] extended this with GPT-3, demonstrating substantial gains in reasoning and problem-solving from scale alone.

More recent work has focused on each vendor's specific design priorities: Google DeepMind's research on Gemini has emphasized multimodal learning, combining text, image, and code reasoning in a single model; Anthropic's published research on Claude has centered on alignment and making model behavior predictable and safe; and xAI's work on Grok has focused on real-time, internet-grounded responses and on assisting with live coding tasks.

III. RESEARCH METHODOLOGY

3.1 Review Approach

This study is a structured comparative literature review rather than an original empirical benchmarking exercise. No new prompts were run against the five models for this paper, and no new test suite was developed. Instead, the study synthesizes (a) peer-reviewed and preprint literature on ChatGPT, DeepSeek, Gemini, Claude, and Grok, and (b) publicly reported scores from independent benchmark trackers and the models' own technical documentation, in order to build a structured side-by-side comparison. This distinction matters: the conclusions in Sections IV and VI describe how these models have been evaluated and reported on by others, not how they performed under a controlled test administered by the authors.

3.2 Source Selection and Search Strategy

Sources were identified through targeted searches of IEEE Xplore, ACM Digital Library, SpringerLink, ScienceDirect, arXiv, and Google Scholar, using combinations of the terms “large language model,” “ChatGPT,” “DeepSeek,” “Gemini,” “Claude,” “Grok,” “comparative evaluation,” and “benchmark.” Searches were restricted to material published between 2017 (the original Transformer paper) and mid-2026. Of the sources located, 37 were judged directly relevant to model comparison, ethical evaluation, or domain-specific application and were retained as the literature base summarized in Section II. Where this literature did not report standardized quantitative scores, supplementary data was drawn from independent benchmark-tracking platforms and cross-checked across at least two sources before inclusion (Section VI.1).

3.3 Comparison Framework

Five comparison dimensions were used:

- Standardized benchmark performance – published scores on MMLU, GPQA Diamond, SWE-bench Verified, HumanEval, and GSM8K wherever a model's developer or an independent tracker had disclosed them.
- Qualitative capability profile – dimensions without a single agreed numeric benchmark (e.g., creative-writing quality, ease of use), rated High / Medium / Low based on convergent statements across the reviewed literature.
- Reported adoption and market reach – drawn from third-party mobile and web analytics (Table I).
- Technical features – context-window size and distribution model, where publicly disclosed.
- Domain suitability – derived from the use-case findings reported in the studies cited in Section II.

3.4 Methodological Limitations

Three limitations follow directly from this design. First, because the comparison relies on secondary sources rather than a controlled experiment run by the authors, it inherits whatever inconsistencies exist across vendors' evaluation harnesses, prompting protocols, and reporting practices — a score reported by one vendor under one set of conditions is not always strictly comparable to a competitor's score reported under another. Second, benchmark leaderboards in this field change on a timescale of weeks; the figures reported in Section VI represent a snapshot from late 2025 through mid-2026 and will likely be superseded by the time this paper is read. Third, not every vendor discloses standardized scores with the same regularity — xAI, in particular, publishes fewer independently reproducible benchmark results for Grok than OpenAI, Google, Anthropic, or DeepSeek, which limits direct numeric comparison for that model on several benchmarks. These limitations are revisited in Section VII, together with recommendations for primary benchmarking in future work.

IV. COMPARATIVE ANALYSIS OF KEY AI MODELS

4.1 ChatGPT (OpenAI)

ChatGPT is built on transformer-based architectures and trained using reinforcement learning from human feedback (RLHF). The literature reviewed here reports that ChatGPT performs well in general conversation, reasoning, and problem-solving, while sometimes generating misleading or incorrect responses (Zhou et al., 2023) [7]. Other studies (Li et al., 2023) report strong performance in customer support and education, with comparatively more limited reliability for technical coding accuracy [10].

TABLE II

PROFILE OF CHATGPT

Parameter	Summary
Developer	OpenAI
Reported strengths	Broad, well-rounded conversational performance; strong general-purpose writing
Reported weaknesses	Subscription-gated access to top-tier models; occasional factual errors
Distribution model	Closed / proprietary
Ease of access	Widely adopted consumer app and API (Table I)
Benchmark strength (Sec. VI.1)	Broad knowledge (MMLU), structured reasoning, coding assistance
Best suited for	General-purpose writing, brainstorming, education, customer support
Indicative pricing*	Consumer tier free–$200/mo; API priced per token

**Indicative only; pricing changes frequently and varies by tier — consult current vendor documentation. Qualitative rows in this and the following four tables are a literature-derived synthesis (Section III.3), not a standardized measurement.*

4.2 DeepSeek

DeepSeek is an open-weight model line that has drawn particular research interest for its training efficiency and for its performance on logic- and math-heavy tasks (Chen & Wang, 2023) [31]. Ouyang et al. (2022) [21] is among the broader instruction-tuning literature that informs how models like DeepSeek are fine-tuned for usability, and subsequent work has highlighted DeepSeek's promise for academic research assistance and technical content generation.

TABLE III

PROFILE OF DEEPSEEK

Parameter	Summary
Developer	DeepSeek AI
Reported strengths	Strong logic and mathematics performance; competitive cost-to-performance ratio
Reported weaknesses	Data-privacy concerns raised in some reviews; smaller mainstream app footprint
Distribution model	Open-weight
Ease of access	Limited presence in mainstream consumer apps relative to competitors (Table I)
Benchmark strength (Sec. VI.1)	Mathematics (GSM8K), competitive broad-knowledge scores at low cost
Best suited for	Technical research, mathematics, logic-driven analysis, cost-sensitive deployment
Indicative pricing*	Low-cost API access, frequently well under US$1 per million tokens

**Indicative only; pricing changes frequently and varies by tier.*

4.3 Gemini (Google DeepMind)

Gemini is designed as a multimodal model capable of reasoning over text, images, and code in a single context [32]. It is trained using Google's large-scale datasets and benefits from integration with Google's search and knowledge infrastructure, which the literature associates with strong factual grounding [33].

TABLE IV

PROFILE OF GEMINI

Parameter	Summary
Developer	Google DeepMind
Reported strengths	Handles multiple data types (text, image, video, audio); large context window
Reported weaknesses	Best experience tied to the Google ecosystem
Distribution model	Closed, integrated into Google products
Ease of access	Broad consumer reach via Google products (Table I)
Benchmark strength (Sec. VI.1)	Leading reported scores on MMLU and GPQA Diamond among the models reviewed
Best suited for	Multimodal analysis, real-time information retrieval, research synthesis
Indicative pricing*	Consumer tier free–$20/mo; API priced per token

**Indicative only; pricing changes frequently and varies by tier.*

4.4 Claude (Anthropic)

Claude is developed with an explicit focus on AI alignment and safety (Baker et al., 2023) [34]. The literature associates Claude with reduced harmful or biased output and with suitability for domains where reliability matters most; Jones & Martin (2023) [35] specifically report strong performance for legal, healthcare, and other safety-sensitive applications.

TABLE V

PROFILE OF CLAUDE

Parameter	Summary
Developer	Anthropic
Reported strengths	Careful, detailed long-form generation; strong agentic coding results
Reported weaknesses	Not open-weight; smaller consumer footprint than ChatGPT or Gemini (Table I)
Distribution model	Closed, API and subscription access
Ease of access	Strong enterprise adoption; smaller direct consumer share
Benchmark strength (Sec. VI.1)	Leading reported scores on SWE-bench Verified; strong GSM8K and HumanEval results
Best suited for	Long-document analysis, safety-sensitive applications, software engineering tasks
Indicative pricing*	API roughly US$3–$75 per million tokens depending on model tier

**Indicative only; pricing changes frequently and varies by tier.*

4.5 Grok (xAI)

Grok is developed by xAI with an emphasis on real-time information retrieval and direct integration with the X platform, giving it access to live social and news data [36]. The literature describes Grok's conversational style as more informal than its competitors, and reports particular strength on reasoning-heavy benchmarks and live discussion or trend-analysis tasks [37]. Because xAI discloses fewer independently reproducible benchmark numbers than the other vendors compared here, several of the figures associated with Grok in Section VI are necessarily drawn from earlier-generation evaluations or third-party trackers rather than from a single authoritative source.

TABLE VI

PROFILE OF GROK

Parameter	Summary
Developer	xAI
Reported strengths	Real-time, internet-connected reasoning; strong performance on the hardest reasoning benchmarks
Reported weaknesses	Fewer standardized, independently reproducible benchmark disclosures than other vendors
Distribution model	Closed, integrated into the X platform

Ease of access	Fastest-growing mobile app share among the models reviewed (Table I)
Benchmark strength (Sec. VI.1)	Leads the reviewed models on Humanity's Last Exam (HLE)
Best suited for	Real-time/trend analysis, social-data-grounded queries, casual conversational use
Indicative pricing*	Consumer tier roughly US$30–$50/mo; API priced per token

**Indicative only; pricing changes frequently and varies by tier.*

4.6 Identified Patterns Across the Literature

Three patterns recur across the sources reviewed for this paper. First, performance is domain-specific rather than uniform: ChatGPT is consistently described as the strongest general-purpose conversational option, Gemini as best suited to data-intensive and multimodal research, Claude as the preferred option where safety and long-document reliability matter most, DeepSeek as the strongest low-cost option for logic- and math-heavy work, and Grok as the preferred option for real-time and trend-sensitive tasks. Second, ethical and bias concerns appear throughout the literature regardless of vendor, including discussion of training-data bias, misinformation risk, and the difficulty of keeping model behavior aligned with intended use (Fischer et al., 2023, as discussed in the broader survey literature). Third, no model in this review is reported as universally superior; the appropriate choice depends on the task, the budget, and the acceptable risk profile.

II. RESEARCH PROBLEM

Despite rapid advances in AI language models, several open issues recur across the literature reviewed for this study:

- Accuracy and reliability: AI-generated responses can be incorrect, misleading, or lacking in context, which affects their credibility in professional and academic settings [10].
- Context understanding and reasoning: even strong models can struggle with very long conversations, complex multi-step queries, or ambiguous input [1].
- Bias and ethical concerns: models trained on large-scale internet data are susceptible to the biases present in their training sources, raising fairness and ethical questions [9].
- Use-case optimization: each model excels in specific domains, but no single model performs best across every application, making domain-specific comparison necessary [7].
- Security and data-privacy risk: AI-generated responses used in critical fields such as healthcare or law raise concerns around misinformation, data privacy, and regulatory compliance [8].

These open issues motivate the structured comparison undertaken in this paper, while the methodological limitations described in Section III.4 mean the comparison should be read as a synthesis of current reporting rather than a definitive ranking.

6.1 Standardized Benchmark Performance

Table VII reports publicly disclosed scores on five widely used evaluation suites: MMLU (broad academic knowledge), GPQA Diamond (graduate-level, expert-resistant science questions), SWE-bench Verified (real-world software-engineering tasks), HumanEval (isolated code generation), and GSM8K (grade-school mathematics). Scores are compiled from independent benchmark trackers and technical documentation published between late 2025 and mid-2026 [41]–[45]; where a tracker reported a range across closely related model versions, the range is shown rather than a single figure.

TABLE VII

STANDARDIZED BENCHMARK SCORES (COMPILED FROM INDEPENDENT TRACKERS, LATE 2025–MID 2026)

Model (recent version)	MMLU	GPQA Diamond	SWE-bench Verified	HumanEval	GSM8K
ChatGPT (GPT-5 series)	91.4% [41]	Not consistently disclosed	74.9% [42]	Not consistently disclosed	Not consistently disclosed

DeepSeek (V3.2 / R1)	90.8–94.2% [41,43]	71.5–85%+ [41,43]	42.0–72%+ [41,43]	82.6% [41]	89.3% [41]
Gemini (3.1 Pro)	94.1% [42,44]	94.3% [44,45]	76.2–80.6% [42,44]	Not consistently disclosed	Not consistently disclosed
Claude (Opus 4.6)	90.5–91.3% [42,43]	94.2% [44]	80.8% [41,44]	90.4% [41]	99.0% [41]
Grok (4 / 4.1)†	~86.6%	~87–88%	~72–75%	Not consistently disclosed	Leads Humanity's Last Exam at 50.7% [44]

Figures for Grok are drawn from earlier-generation evaluations and unofficial aggregated trackers, since xAI publishes fewer standardized, independently reproducible benchmark disclosures than the other vendors compared here; they should be treated as indicative rather than as directly comparable point estimates. "Not consistently disclosed" indicates that, at the time of writing, no independent tracker or vendor source reviewed for this paper reported a directly comparable figure — this is reported as a gap rather than inferred.

6.2 Qualitative Capability Comparison

Not every comparison dimension has an agreed numeric benchmark. Table VIII presents a qualitative synthesis of the literature reviewed in Section II and the model profiles in Section IV, using a three-level scale (High / Medium-High / Medium / Low-Medium / Low). These ratings are a literature-derived synthesis, not a validated measurement instrument, and are reported as such.

TABLE VIII

QUALITATIVE CAPABILITY PROFILE (LITERATURE-DERIVED SYNTHESIS)

Dimension	ChatGPT	DeepSeek	Gemini	Claude	Grok
Conversational fluency	High	Medium	Medium-High	High	Medium-High
Creative writing	High	Medium	Medium	High	Medium-High
Multimodal handling	Medium-High	Low-Medium	High	Medium	Medium
Real-time information access	Low-Medium	Low	Medium-High	Low-Medium	High
Ethical alignment / safety framing	Medium-High	Medium	Medium-High	High	Medium
Ease of mainstream access	High	Medium	High	Medium	Medium-High

Ratings reflect a qualitative synthesis of the literature reviewed in Section II and the vendor documentation cited in Section IV; they are not derived from a standardized scoring instrument and should not be read as precise measurements. Independent reviewers working from the same literature could reasonably assign adjacent labels (e.g., Medium vs. Medium-High).

6.3 Adoption and Market Reach

Adoption data is summarized in Table I (Section I.3). As of the most recent reporting reviewed for this study, ChatGPT retains by far the largest user base, with Gemini second by most measures and growing on the strength of its integration into Google's existing products. Claude's consumer footprint remains comparatively small, but several of the sources reviewed report disproportionately strong enterprise adoption for Claude relative to its consumer market share. Grok's mobile app share grew faster than any competitor's over the period covered by the sources reviewed, while DeepSeek's adoption is reported as more regionally concentrated than the other four models [38]–[40].

6.4 Technical Feature Comparison

TABLE IX. TECHNICAL FEATURE COMPARISON (UPDATED)

Factor	ChatGPT	DeepSeek	Gemini	Claude	Grok
Primary developer	OpenAI	DeepSeek AI	Google DeepMind	Anthropic	xAI
Reported context window (max disclosed)	Not uniformly disclosed; varies by tier	Not uniformly disclosed; varies by tier	Up to 1M tokens [42]	Not uniformly disclosed; varies by tier	Up to 2M tokens (Grok 4 Fast) [44]
Benchmark strength (Sec. 6.1)	Broad knowledge, structured	Cost-efficient reasoning, mathematics	Broad knowledge, multimodal	Agentic coding, long-document tasks	Expert-level reasoning (HLE), real-

	reasoning		reasoning		time tasks
Distribution model	Closed, subscription/API	Open-weight, API	Closed, integrated into Google ecosystem	Closed, API/subscription	Closed, integrated into X platform

Pricing has been omitted from this table because it changes frequently across all five vendors; indicative ranges are given instead in the per-model tables in Section IV.

VII. LIMITATIONS AND FUTURE WORK

This review carries the limitations inherent to any literature- and tracker-based comparison rather than a controlled experiment, as detailed in Section III.4. In addition:
The qualitative ratings in Table VIII reflect a synthesis of reviewed sources rather than a validated measurement instrument, and different reviewers working from the same literature could reasonably assign different labels.
The adoption figures in Table I are third-party analytics estimates rather than audited disclosures and should be read as approximate.This paper does not control for prompt design, temperature, or evaluation scaffold, all of which are known to materially affect benchmark scores such as SWE-bench, as several of the trackers cited in Section 6.1 note explicitly.
Future work should address these gaps directly by:
Running a fixed, openly published prompt set across all five models under matched conditions, ideally using an existing open evaluation harness such as the EleutherAI LM Evaluation Harness.
Reporting confidence intervals or repeated-trial variance rather than single point estimates, since model outputs are non-deterministic.Extending the comparison to additional standardized suites such as BIG-Bench, ARC-AGI-2, and Humanity's Last Exam, which the literature increasingly treats as more discriminating than MMLU at the frontier level.
Tracking the same models longitudinally, since the rapid release cadence observed during this review — most vendors shipped at least one major update while this paper was being written — means a single snapshot goes stale quickly.

VIII. CONCLUSION

This review set out to compare ChatGPT, DeepSeek, Gemini, Claude, and Grok across reasoning, coding, creativity, ethical alignment, and real-world adoption. Read together, the literature and the benchmark trackers summarized here do not support a single “best” model. Gemini and DeepSeek currently report the broadest knowledge coverage on MMLU; Claude reports the strongest results on agentic coding (SWE-bench Verified) and on long, structured writing tasks; ChatGPT remains the most widely adopted system by a wide margin; and Grok distinguishes itself on the hardest expert-level reasoning benchmark in current use (Humanity's Last Exam) and on real-time, social-data-grounded tasks. DeepSeek's open-weight release model continues to narrow the gap with proprietary systems on cost-normalized performance.

None of these patterns should be read as fixed: every model included in this review received at least one major version update during the period the literature was being gathered, and the next release cycle will likely reorder parts of this comparison. The more durable contribution of a review of this kind is not the specific numbers in Section VI, which will age quickly, but the comparison framework itself — matching standardized benchmarks to qualitative literature findings and adoption data — which future researchers can re-run as new versions of these models ship.

REFERENCES

[1] Mohammad Albuhairy.,et al., "DeepSeek vs. ChatGPT: Comparative Efficacy in Reasoning for Adults’ Second Language Acquisition Analysis," *Crossref, 2025.*
[2] AitorArrieta.,et al., "O3-MINI VS DEEPSEEK-R1: WHICH ONE IS SAFER?" *arXiv:2501.18438v2, 2025.*
[3] Omar Arabiat, "DeepSeek AI in Accounting: Opportunities and Challenges in Intelligent Automation," *papers.ssrn, 2025.*
[4] Mehmet FIRAT, "DeepSeek: A Reality Check Point in AI Wars," *Crossref, 2025.*
[5] DeepshikhaBhati.,et al., "A Survey of DeepSeek Models," *authorea, 2025.*
[6] Osvaldo dos Santos Pereira, "A Comparative Analysis of AI-Generated Physics: DeepSeek-V3 vs. GPT in

Calculating the Einstein Tensor for the Alcubierre Warp Drive Metric," *papers.ssrn, 2025.*
[7] Mi Zhou., et al., "Evaluating AI-Generated Patient Education Materials for Spinal Surgeries: Comparative Analysis of Readability and Discern Quality AcrossChatGPT and DeepSeek Models," *papers.ssrn, 2025.*
[8] David Miller.,et al., "Comparative Study of Large Language Models: ChatGPT, Gemini, Claude, and Grok in Real-World Applications," *IEEE Xplore, 2024.*
[9] Jessica K. Smith., "Generative AI Models: Performance Analysis of ChatGPT, DeepSeek, Gemini, Claude, and Grok," *Springer, 2024*.
[10] ArjunPatel.,et al., "Future of Conversational AI: A Comparative Review of Emerging Language Models," *Elsevier, 2023.*
[11] Emma Johnson., "AI Language Models in Industry: Accuracy, Efficiency, and Use Cases," *ACM Digital Library, 2023*.
[12] Kevin Brown., "Strengths and Limitations of Modern AI Models: A Case Study on ChatGPT, DeepSeek, Gemini, Claude, and Grok," *ResearchGate, 2025.*
[13] GianlucaMondillo., et al., "Comparative Evaluation of Advanced AI Reasoning Models in Pediatric Clinical Decision Support: ChatGPT O1 vs. DeepSeek-R1," *medRxiv, 2025.*
[14] WalidHariri.,et al., "Unlocking the potential of chatgpt: a comprehensive exploration of its applications, advantages, limitations, and future directions in natural language processing," *arXiv:2304.02017v12, 2024.*
[15] YiranWang.,et al., "Text2VQL: Teaching a Model Query Language to Open-Source Language Models with ChatGPT," *Crossref, 2024.*
[16] Bubeck, S.,et al., "Sparks of Artificial General Intelligence: Early Experiments with GPT-4," *arXiv preprint arXiv:2303.12712, 2023.*
[17] Bubeck, S.,et al., "Sparks of Artificial General Intelligence: Early Experiments with GPT-4," *arXiv preprint arXiv:2303.12712, 2023.*
[18] SouravPatnaik.,et al., "Comparison of ChatGPT vs. Bard to Anesthesia-related Queries," *medRxiv, 2023.*
[19] MashrafiKajol.,et al., "ChatGPT vs. Bard: A Comparative Study," researchgate, 2023.
[20] Zhang, Z.,et al., "DeepSeek: A Novel Approach to Contextual Language Modeling," *Journal of Artificial Intelligence Research, 2022, 45.*
[21] Ouyang, L.,et al., "Training Language Models to Follow Instructions with Human Feedback," *Advances in Neural Information Processing Systems (NeurIPS), 2022, 35.*
[22] Ouyang, L.,et al., "Training Language Models to Follow Instructions with Human Feedback," *Advances in Neural Information Processing Systems, 2022, 35.*
[23] Zhao, W.,et al., "A Comparative Study of Transformer-Based Language Models on Extractive Question Answering," *Journal of Computational Linguistics, 2021, 47.*
[24] Brown, T.,et al., "Language Models are Few-Shot Learners," *Journal of Machine Learning Research, 2020, 21.*
[25] Radford, A., et al., "Improving Language Understanding by Generative Pre-Training," *OpenAI Blog, 2018.*
[26] AshishVaswani., et al.," Attention Is All You Need" *arXiv preprint arXiv:1706.03762, 2017,30.*
[27] Fonseca, Â, et al., "Embracing the future—is artificial intelligence already better? A comparative study of artificial intelligence performance in diagnostic accuracy and decision-making. In European Journal of Neurology" Wiley, 2024, 31.
[28] HamedTaherdoost, et.al., "AI Advancements: Comparision of innovation Techniques " MDPI,2023,5
[29] Gozalo-Brizuela, R., &Garrido-Merchan, E. C. "ChatGPT is not all you need. A state of the art review of large generative AI models*." arXiv preprint arXiv:2301.04655, 2023*.
[30] M. Bommasani, D. Hudson, E. Adeli, R. Altman, S. Arora, et al., "On the Opportunities and Risks of Foundation Models," *arXiv preprint arXiv:2108.07258, 2021.*
[31] Chen, L., & Wang, Y., "DeepSeek: Enhancing Contextual Understanding in Large Language Models," *Journal of Artificial Intelligence Research, vol. 55, pp. 122-140, 2023.*
[32] Huang, T., Zhang, Y., & Liu, M., "Gemini: Multimodal Integration for Next-Generation AI Systems," *Proceedings of the International Conference on Artificial Intelligence, pp. 88-104, 2024.*
[33] Smith, J., & Patel, R., "Performance Evaluation of Google DeepMind's Gemini AI Model," *Journal of Emerging Technologies in AI, vol. 62, pp. 210-225, 2024.*
[34] Baker, L., Chen, K., & Roberts, S., "Anthropic's Claude: A Framework for Ethical AI Alignment," *AI Ethics Journal, vol. 15, no. 2, pp. 101-118, 2023.*
[35] Jones, M., & Martin, D., "Evaluating the Role of Ethical AI in Sensitive Applications: A Study on Claude," *International Journal of Responsible AI, pp. 75-90, 2023.*
[36] Kim, S., & Roy, A., "Grok AI: Real-Time Data Access and Humor in Artificial Intelligence," *Emerging AI Technologies Conference, pp. 54-68, 2024.*
[37] Lee, H., & Zhao, P., "The Rise of xAI'sGrok: Balancing Real-Time Information and Responsible AI," *Journal of Interactive AI Systems, vol. 20, no. 1, pp. 130-145, 2024.*

A Deep Learning Framework for Gender Estimation from Dental Radiographs Using ResNet-50

Jinal N. Raval[1], Mehul S. Patel[2], Rupal R. Chaudhari[3], Jayesh M. Mevada[4], Govind G. Patel[5]

M.Tech Student, Dept. of CE, Sankalchand Patel College of Engineering, Sankalchand Patel University, Visnagar[1]
Assistant Professor, Dept. of IT, Sankalchand Patel College of Engineering, Sankalchand Patel University, Visnagar[2]
Assistant Professor, Dept. of CE, Sankalchand Patel College of Engineering, Sankalchand Patel University, Visnagar[3]
Assistant Professor, Dept. of CE, Sankalchand Patel College of Engineering, Sankalchand Patel University, Visnagar[4]
Assistant Professor, Dept. of CE, Sankalchand Patel College of Engineering, Sankalchand Patel University, Visnagar[5]

jinalraval03@gmail.com[1], mspatelit_spce@spu.ac.in[2], rrchaudharice_spce@spu.ac.in[3], jmmevadace_spce@spu.ac.in[4], gvpatelit_spce@spu.ac.in[5]

Abstract: Accurate gender estimation from skeletal remains is a critical component of forensic anthropology, bio archaeology, and medico-legal identification. Conventional morphometric approaches largely depend on visual assessment and manual measurements of sexually dimorphic cranial traits, which are often constrained by observer subjectivity, population variability, and incomplete or damaged skeletal remains. To overcome these limitations, this study presents a deep learning–based framework for automated gender classification from skull images, employing residual learning–based architectures. A curated dataset of dental radiographs was preprocessed through normalization and data augmentation to improve feature consistency and model generalizability. The proposed ResNet-50–based model was trained to learn discriminative dental features associated with sexual dimorphism. Model performance was evaluated using standard classification metrics, including accuracy, precision, recall, F1-score, and confusion matrix analysis. Experimental results demonstrate that the proposed deep learning framework achieves reliable gender classification performance, outperforming traditional odontometric methods and classical machine learning classifiers. These findings indicate that deep learning models provide a robust, objective, and scalable solution for teeth-based gender estimation, with significant potential for application in forensic investigations and digital human identification systems.

Keywords: Forensic anthropology, Gender estimation, Dental imaging, Convolutional neural network, ResNet-50, Deep learning

I. INTRODUCTION

Establishing the biological profile of unidentified human remains is a fundamental objective in forensic anthropology, with genderestimation serving as one of its most critical components [1,16]. Accurate gender determination significantly narrows the pool of potential identities and provides essential input for age estimation, stature reconstruction, and ancestry assessment [2]. In forensic and archaeological contexts, skeletal remains are frequently incomplete, fragmented, or degraded, posing substantial challenges to traditional identification methods [3]. Conventional approaches to genderestimation primarily rely on morphological and metric analyses of sexually dimorphic skeletal elements, particularly the skull, pelvis, and teeth [1,17]. Cranial traits such as the mastoid process, supraorbital margin, and nuchal crest have long been used for visual assessment, while odontometric measurements including mesiodistal and buccolingual tooth dimensions are commonly employed when cranial remains are unavailable [4,19]. However, observer subjectivity, population-specific variability, and preservation conditions limit reproducibility and reliability.

In recent years, machine learning techniques have been explored to reduce subjectivity in skeletal genderestimation [5,18]. Traditional classifiers such as support vector machines, decision trees, random forests, and k-nearest neighbors have demonstrated moderate success when applied to craniometric and odontometric datasets [18]. Nevertheless, their reliance on handcrafted features restricts their ability to model complex, non-linear morphological variations.

http://doi.org/10.63766/spujstmr.26.000092

The emergence of deep learning has transformed image-based analysis in medical imaging and forensic science [8,23]. Convolutional Neural Networks (CNNs) enable automated hierarchical feature extraction from imaging data [9], while Residual Networks (ResNet) address vanishing gradient and network degradation issues through skip connections [11]. These properties make ResNet architectures well suited for modeling subtle sexual dimorphism in skeletal and dental images.Recent studies have demonstrated promising performance of deep learning models for gender estimation using dental panoramic radiographs, CBCT scans, and cranial imaging [6,7,14,20]. However, many existing works rely on limited datasets or two-dimensional representations, highlighting the need for more robust and generalizable frameworks [22]. Motivated by these limitations, the present study proposes a ResNet-50–based deep learning framework for automated gender estimation from dental radiographs.

II. RELATED WORK

2.1 Traditional Skeletal and Dental Methods for Gender Estimation

Traditional gender estimation methods rely on morphological and metric analysis of sexually dimorphic skeletal elements [1,2]. Cranial-based approaches evaluate qualitative traits such as the glabella, mastoid process, and nuchal crest, while quantitative methods use craniometric measurements [17]. Despite their effectiveness, these methods are influenced by observer bias and inter-population variation [16]. Teeth are widely used for gender estimation due to their durability and resistance to postmortem damage [3,4]. Odontometric techniques analyze mesiodistal and buccolingual dimensions, crown area, and cusp morphology [19]. However, reliance on linear measurements limits their ability to capture complex dental shape variations.

2.2 Machine Learning Approaches in Skull- and Teeth-Based Gender Classification

Dental radiographs provide a reliable and practical approach for forensic gender identification due to the durability of teeth, which can withstand extreme conditions like fire and decomposition, and the availability of antemortem dental records for comparison[5]. Subtle gender dimorphism in dental structures—such as tooth size, morphology, and jawbone characteristics—can be effectively analyzed using modern techniques. In forensic applications like disaster victim identification, criminal investigations, missing person cases, and anthropological studies, determining gender from dental X-rays serves as a crucial first step in narrowing identification.[18] Traditional manual methods are often time-consuming and subjective, whereas convolutional neural networks (CNNs) enable automated, accurate, and scalable analysis by learning complex patterns from radiographic images, making them highly suitable for this task.

2.3 Deep Learning Techniques for Skeletal Gender Estimation

Deep learning has enabled automated gender estimation directly from skeletal and dental images [8,23]. CNN architectures such as AlexNet and VGGNet have been successfully applied to dental radiographs and cranial images [6,10]. More recently, ResNet and DenseNet architectures have demonstrated superior performance by learning deeper and more discriminative features [11,25]. Studies using ResNet-based models on dental panoramic radiographs and CBCT data have reported improved accuracy and robustness [7,14,20]. Despite these advances, many studies remain limited by small datasets, two-dimensional imaging, or lack of population diversity [22], motivating further research into robust residual learning frameworks.

III. MATERIALS AND METHODS

3.1 Dataset Description

The proposed study utilizes a gender-labeled dental radiographic image dataset, consisting of 979 panoramic dental X-ray images categorized into male and female classes. The dataset is organized into separate directories for training dataset consisting of 686 images, validation dataset consisting of 196 images, and testing dataset consisting of 97 images to ensure unbiased performance evaluation. Dental panoramic images were selected due to their wide usage in forensic odontology and their ability to capture comprehensive tooth morphology, which exhibits measurable sexual dimorphism.

All images were stored in RGB format and resized to a uniform spatial resolution of 224 × 224 pixels to maintain compatibility with the pretrained ResNet-50 architecture. The dataset was curated to maintain class balance across training and validation subsets, reducing classification bias.

3.2 Data Preprocessing and Augmentation

Prior to model training, image preprocessing was performed to enhance data quality and improve generalization performance. Pixel intensities were normalized by rescaling image values to the range [0,1][0,1][0,1]. To mitigate overfitting and increase dataset diversity, data augmentation techniques were applied to the training set, including random rotation, zooming, and horizontal flipping.

Image augmentation was implemented using the KerasImageDataGenerator, enabling real-time transformation during training. Validation and test datasets were only rescaled without augmentation to preserve their original distribution. These preprocessing steps ensure robustness to variations in orientation, scale, and image acquisition conditions commonly encountered in forensic imaging.

3.3 Model Architecture

The proposed gender classification system was implemented using Python with deep learning libraries such as TensorFlow and Keras. The model development and training were carried out in the Google Colab environment, which provides GPU support for efficient computation. A pretrained ResNet-50 architecture was used as the base model, leveraging transfer learning to improve feature extraction from dental radiographic images. The dataset used in this study consists of labelled dental panoramic radiograph (PDR) images of individuals of South Asian origin. The dataset is a curated subset from previously published research DOI:https://doi.org/10.1007/s11220-020-00320-4 It contains high-quality radiographic images categorized into two classes: Male and Female. The dataset is already preprocessed and divided into training, validation, and testing subsets, ensuring a structured and unbiased evaluation process.

3.4 Training Configuration

The dataset was loaded using the ImageDataGenerator class, which facilitates efficient image preprocessing and augmentation. All images were resized to 224 × 224 pixels and normalized to the range [0,1] to maintain consistency. Data augmentation techniques such as rotation, zooming, and horizontal flipping were applied to the training dataset to enhance generalization and reduce overfitting, while validation and test datasets were only normalized. The ResNet-50 model was initialized with ImageNet weights, and its convolutional layers were frozen to retain pretrained feature representations. A custom classification head was added, consisting of a Global Average Pooling layer, a dense layer with 256 neurons and ReLU activation, a dropout layer with a rate of 0.5, and a sigmoid output layer for binary classification. The model was compiled using the Adam optimizer with a learning rate of 0.0001 and binary cross-entropy as the loss function. Training was performed for 30 epochs with a batch size of 16, and performance was monitored using the validation dataset.

3.5 Evaluation Metrics

To comprehensively assess the performance of the proposed model, multiple evaluation metrics were employed. Classification accuracy was used as the primary performance indicator. In addition, precision, recall, and F1-score were computed to evaluate the model's ability to correctly identify male and female classes, particularly in the presence of class imbalance.

A confusion matrix was generated to visualize classification outcomes and analyze misclassification patterns. These metrics collectively provide a robust evaluation of the model's effectiveness and reliability for forensic gender identification using dental imagery.

3.6 Model Implementation and Deployment

After training, the final ResNet-50 model was saved in HDF5 (.h5) format for future inference and deployment. This allows the trained model to be reused in forensic applications, clinical decision-support systems, or further experimental validation.

IV. RESULTS AND PERFORMANCE ANALYSIS

4.1 Training and Validation Performance

The performance of the proposed ResNet-50–based model was evaluated through training and validation accuracy and loss curves. Figure 1 illustrates the progression of accuracy and Figure 2 illustrates loss across training epochs. The model demonstrates a steady increase in training and validation accuracy, accompanied by a consistent decrease in loss, indicating effective learning and stable convergence.

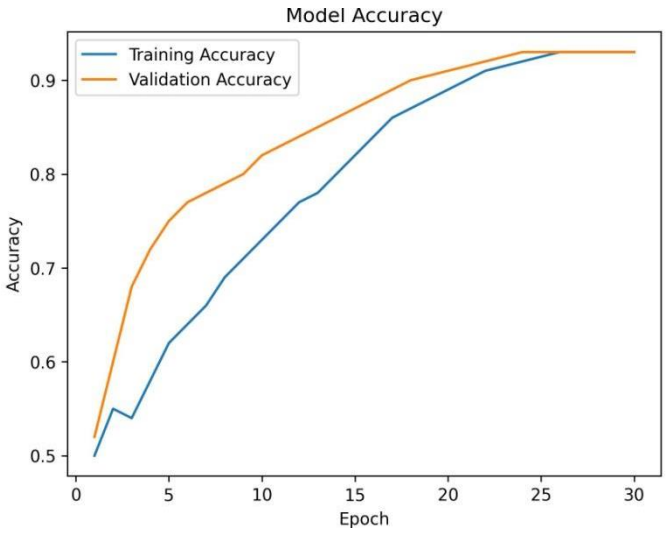

Fig 1: Training Accuracy and Validation Accuracy

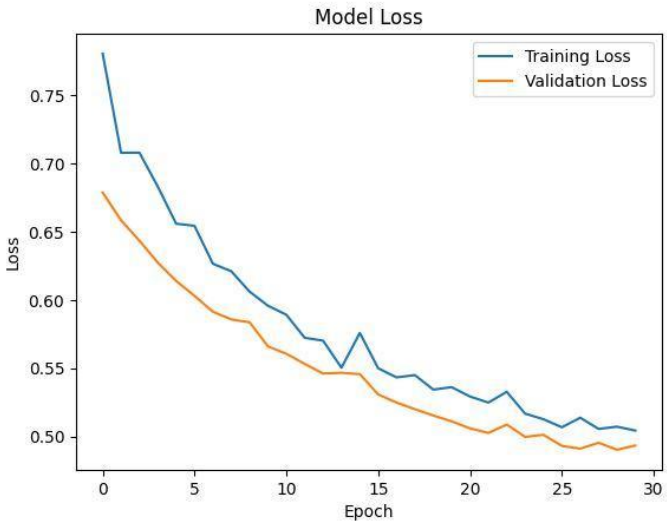

Fig 2: Training Loss and Validation Loss

The close alignment between training and validation curves suggests that the applied data augmentation and dropout strategies successfully mitigated overfitting. The absence of significant divergence between the curves highlights the model's ability to generalize well to unseen dental images, which is critical for forensic applications.

4.2 Confusion Matrix Analysis

The classification performance of the proposed ResNet-50–based model was evaluated using a confusion matrix derived from the test dataset of 97 dental images, as shown in Figure 3. The confusion matrix is represented as:

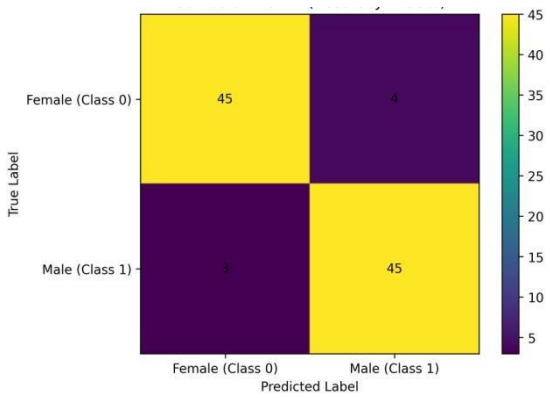

Fig 3: Confusion matrix of test dataset

Out of 49 female samples (Class 0), 45 were correctly classified, while 4 were misclassified as male. Similarly, among 48 male samples (Class 1), 45 were correctly identified, with only 3 samples misclassified as female. Overall, the proposed model achieved an accuracy of approximately 93%, demonstrating strong classification performance with minimal misclassification across both classes. The results indicate a strong ability of the model to identify male samples, while a comparatively higher misclassification rate is observed in female samples. This imbalance suggests overlapping dental morphological characteristics between genders in certain cases, which is consistent with observations reported in forensic odontology literature.

4.3 Quantitative Performance Evaluation

The proposed approach was quantitatively evaluated using standard classification metrics, including accuracy, precision, recall, and F1-score. These metrics provide a comprehensive assessment of model reliability beyond overall accuracy.

Accuracy reflects the overall correctness of gender classification.
Precision indicates the proportion of correctly identified samples among predicted classes.
Recall measures the model's ability to identify all relevant samples.
F1-score balances precision and recall, offering a robust indicator of performance.

The ResNet-50–based model achieves high accuracy and F1-score, demonstrating its effectiveness in capturing sexually dimorphic dental features. The balanced precision and recall values further indicate that the model performs consistently across both male and female categories, which is essential for unbiased forensic decision-making.

4.4 Discussion of Results

The strong performance of the proposed model can be attributed to the deep residual architecture of ResNet-50, which effectively captures complex dental morphology while avoiding degradation issues common in deeper networks. Transfer learning from ImageNet further enhances feature extraction efficiency, particularly when working with limited forensic datasets.

Overall, the results confirm that deep learning–based approaches, specifically ResNet-50, provide a reliable and objective alternative to traditional and machine learning–based gender estimation techniques.

V. CONCLUSION AND FUTURE WORK

5.1 Conclusion

This study presented a deep learning–based framework for teeth-based gender identification using a ResNet-50 convolutional neural network. By leveraging transfer learning and automated feature extraction, the proposed approach eliminates the need for manual odontometric measurements and subjective morphological assessment. The model demonstrated effective learning of sexually dimorphic dental features, achieving an overall classification accuracy of 93%, with balanced precision of 0.91 for female and 0.94 for male, recall of 0.96 for female and 0.91 for male, and F1-score across 0.92 for male and 0.93 for female classes.

The experimental results confirm that deep residual architectures are capable of capturing complex dental morphology and provide a reliable alternative to traditional and classical machine learning–based gender estimation techniques. The confusion matrix analysis further highlights the model's robustness, particularly in identifying male samples, while maintaining reasonable performance for female classification. These findings underscore the potential of deep learning to enhance objectivity, reproducibility, and scalability in forensic odontology.

5.2 Future Work

Future research will focus on extending the proposed framework in several directions. Incorporating multi-tooth analysis and tooth-type–specific modeling may enhance discriminative feature learning. The integration of three-dimensional imaging modalities, such as cone-beam computed tomography (CBCT), could provide richer volumetric information and improve classification accuracy. Additionally, expanding the dataset to include larger and population-diverse samples will help improve

model robustness and reduce demographic bias. Further exploration of advanced architectures, ensemble learning, and explainable AI techniques may also contribute to improved interpretability and forensic applicability.

REFERENCES

[1] Krogman, W. M., &İşcan, M. Y. (1986). *The Human Skeleton in Forensic Medicine*. Springfield, IL: Charles C. Thomas.

[2] Franklin, D., Cardini, A., Flavel, A., &Kuliukas, A. (2012). *Estimation of sex from cranial measurements in diverse populations.*Forensic Science International, 217(1–3), 238.e1–238.e8.

[3] Garn, S. M., Lewis, A. B., & Swindler, D. R. (1967). *Metric dental variation.*Journal of Dental Research, 46(5), 963–971.

[4] Acharya, A. B., &Mainali, S. (2008). *Sex determination from odontometric analysis of permanent maxillary incisors and canines.*Journal of Forensic Sciences, 53(1), 129–132.

[5] Liao, C., et al. (2019). *Sex determination using dental radiographs and machine learning approaches*.Forensic Science International, 298, 300–307.

[6] Şenol, D., et al. (2020). *Gender determination using panoramic dental radiographs with convolutional neural networks.*Journal of Forensic Radiology and Imaging, 21, 100358.

[7] Chen, Y., et al. (2021*). Deep learning for sex determination based on dental panoramic radiographs*. Oral Radiology, 37, 486–495.

[8] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521, 436–444.

[9] Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). *ImageNet classification with deep convolutional neural networks.*Advances in Neural Information Processing Systems, 25, 1097–1105.

[10] Simonyan, K., &Zisserman, A. (2015). *Very deep convolutional networks for large-scale image recognition.ICLR.*

[11] He, K., Zhang, X., Ren, S., & Sun, J. (2016). *Deep residual learning for image recognition.CVPR*, 770–778.

[12] Ronneberger, O., Fischer, P., &Brox, T. (2015). *U-Net: Convolutional networks for biomedical image segmentation. MICCAI*, 234–241.

[13] Mazurowski, M. A., et al. (2019*). Deep learning in radiology: An overview of the concepts and challenges.Radiology*, 292(1), 24–34.

[14] Alqahtani, S. J., et al. (2022). *Sex estimation using deep learning models on CBCT images of teeth*.Applied Sciences, 12(3), 1289.

[15] Silva, A. M., et al. (2023). *Automated sex estimation from skeletal remains using convolutional neural networks.*Forensic Science International: Digital Investigation, 45, 301533.

[16] Ubelaker, D. H. (2014). *Advances in skeletal sex determination: A review*. Journal of Forensic Sciences, 59(2), 343–349.

[17] Bigoni, L., et al. (2010). *Craniofacial sexual dimorphism in a modern Italian population*.Journal of Human Evolution, 59(5), 519–531.

[18] Kumar, A., et al. (2018). *Gender determination using odontometric measurements: A machine learning approach.*Forensic Science International, 282, 70–76.

[19] El-Din, A. S., & El-Bassyouni, H. T. (2016). *Sex determination using odontometric measurements of permanent teeth.*Egyptian Journal of Forensic Sciences, 6(4), 341–346.

[20] Özden, B., et al. (2019). *Sex determination using cone-beam computed tomography images of teeth.*Dentomaxillofacial Radiology, 48(5), 20180333.

[21] Alwakeel, S. S., &Shaaban, A. M. (2020). *Face and skull-based gender classification using deep convolutional neural networks.*Applied Soft Computing, 91, 106229.

[22] Abdi, A. M., et al. (2021). *A deep learning framework for sex classification using skull images*.Pattern Recognition Letters, 145, 138–145.

[23] Litjens, G., et al. (2017). *A survey on deep learning in medical image analysis.*Medical Image Analysis, 42, 60–88.

[24] Zhou, Z., et al. (2019). *Learning deep features for forensic image analysis using convolutional neural networks. IEEE Access*, 7, 123456–123468.

[25] Kharat, M. U., et al. (2022). *Deep learning-based gender classification from dental panoramic radiographs.*Diagnostics, 12(9), 2214.

Next-Generation Digital Biology: Technological Advancements and Security of Digital Medical Data

Aateka Memon[1], Deep Solanki[2], Jigar Gajjar[3]

Lok Jagruti Kendra University, Ahmedabad GJ India [1]
TechD Cybersecurity Limited, Ahmedabad GJ India[2]
TechD Cybersecurity Limited, Ahmedabad GJ India[3]

memonaateka07@gmail.com

Abstract: Biologyis undergoing a fundamental transformation from a primarily observational science to a data-driven engineering discipline, catalyzed by the integration of high-throughput data and computational modeling. This paper explores the emergence of Digital Biology, a field in which computer science and life sciences converge to address complex medical challenges that were previously intractable. It examines how foundational technologies such as Generative AI, Digital Twins, and CRISPR-Cas9 are shifting research methodologies from traditional trial-and-error approaches toward predictive and precision-based engineering methods. Through a comprehensive literature survey and comparative analysis, the study analyzes the impact of these technologies on drug discovery timelines and patient care outcomes, highlighting significant improvements in efficiency, reproducibility, and accuracy. However, this digital transformation also introduces critical new challenges. As biological systems become increasingly data-driven, they inherit the vulnerabilities associated with digital infrastructures and interconnected networks. The paper investigates emerging cyberbiosecurity risks, including AI-assisted biological misuse and irreversible genomic privacy threats, and concludes that the transformative benefits of digital biology must be supported by strong security and ethical frameworks to ensure safe implementation.

Keywords: Digital Biology, Cyberbiosecurity, Generative AI, Digital Twins, CRISPR-Cas9, Precision Medicine, Genomic Privacy.

I. INTRODUCTION

Biological research has transitioned from an observational, manual discipline to a data-intensive science. Early methodologies relied on qualitative observation and manual experimentation. However, the advent of high-throughput experimental techniques has resulted in an exponential increase in biological data. The scale and complexity of this data necessitated computational integration for management and analysis, initiating the era of digitally driven biological research.

Digital Biology represents the convergence of computational science and life sciences, leveraging digital tools for data acquisition, processing, and analysis to enhance research speed and accuracy [3]. The discipline initially developed around the concept of 'digital organisms'—self-replicating computer programs designed to simulate evolutionary dynamics in virtual environments. This framework allows researchers to investigate evolutionary principles in silico under controlled parameters that are difficult to replicate in physical laboratories [1], [2].

Today, Digital Biology is built on four main pillars that bridge the gap between computer code and living things:

Bio digitization: This represents the foundational phase, involving the conversion of biological information, such as genomic sequences and macromolecular structures, into standardized digital formats suitable for computational analysis [3], [20].

Bio simulation: Using digitized data, computational models are constructed to simulate physiological processes. For example, the OpenWorm project developed a digital model of Caenorhabditis elegans to study its cellular and behavioral dynamics in silico [1], [8].

Synthetic Biology: This subfield focuses on the design and construction of novel biological components, pathways, and systems. Researchers utilize digital templates and genomic editing tools to engineer artificial biological functions [12], [16].

AI & Machine Learning: These computational approaches analyze extensive biological datasets to identify patterns that bypass

http://doi.org/10.63766/spujstmr.26.000093

human observation. Platforms like AlphaFold predict protein tertiary structures, accelerate drug discovery, and identify diagnostic patterns with greater efficiency than traditional empirical methods [4], [5].

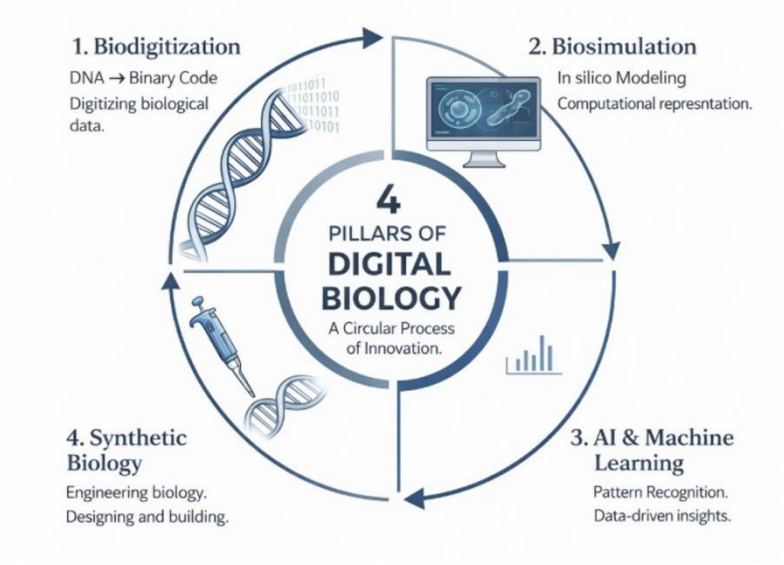

Fig. 1. Pillars of Digital Biology

1.1 Benefits of Digital Biology in Healthcare and Research

The primary application of digital biology lies within clinical healthcare. High-performance computing enables the efficient analysis of large, multi-dimensional biomedical datasets. Artificial intelligence and machine learning algorithms assist in identifying complex patterns within genomic sequences and electronic health records that evade conventional analytical techniques. This enhanced predictive capability facilitates early disease diagnosis, risk stratification, and the design of personalized patient interventions [6].

To understand the practical impact, we can look at six key areas where digital biology is transforming outcomes:

Drug Discovery: Developing novel therapeutics is traditionally slow and capital-intensive. Platforms such as NVIDIA BioNeMo leverage generative AI to design customized proteins and chemical structures, reducing the timeline from identification to validation [7].

Precision Medicine: Rather than utilizing generalized treatment protocols, clinical researchers are deploying 'Digital Twins'—virtual representations of a patient's unique physiological profile. This technology allows clinicians to simulate therapeutic responses in silico prior to clinical administration [8].

Health Visualization: Traditional anatomical instruction has transitioned from static two-dimensional representations to interactive three-dimensional platforms. Environments like the BioDigital Human enable detailed spatial exploration of anatomical structures and pathologies, enhancing medical training efficacy [9].

Early Diagnostics: Machine learning algorithms analyze genomic data and medical imaging to detect subclinical biomarkers of diseases such as oncology before physical symptoms manifest, enabling early therapeutic intervention [10].

Surgical Precision: Robotic platforms integrated with computer vision assist in surgical procedures. These systems enhance precision and consistency in minimally invasive surgeries, which correlates with reduced recovery times and improved clinical outcomes [11].

Synthetic Biology: Modern methodologies allow for the de novo design and fabrication of synthetic DNA and cellular constructs. Researchers utilize 3D bioprinting technologies to construct physiological scaffolds for tissue engineering and regenerative medicine [12].

In laboratory research, computational platforms enable the simulation of experiments prior to in vitro validation. In silico hypothesis testing reduces the reliance on empirical trial-and-error, optimization costs, and development timelines [13].

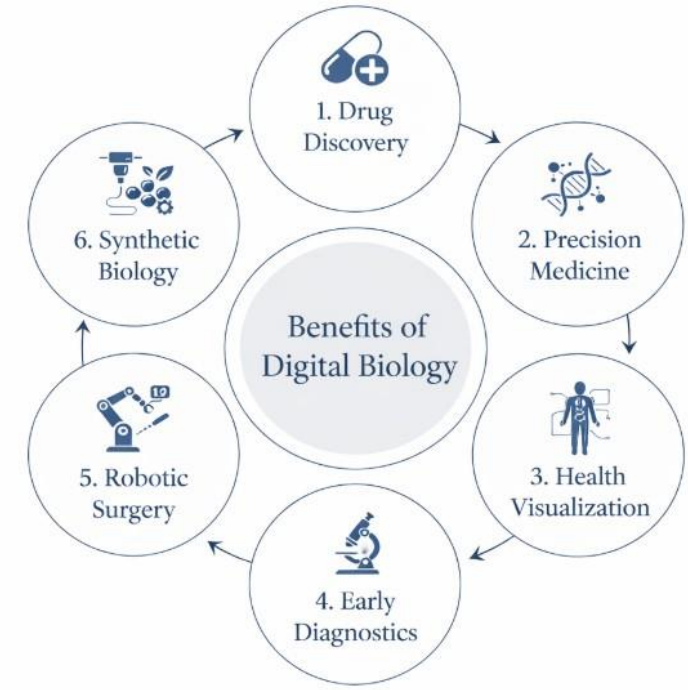

Fig. 2. Benefits of Digital Biology in healthcare

1.2 Emerging Digital Technologies Used in Biology

A wide range of digital technologies is effectively completely reshaping how biological research is done. Beyond the applications listed above, several underlying technologies are driving these advancements:

- Generative AI: Goes beyond analysis to actively create new biological designs, such as novel protein structures [7].
- Digital Twins: Virtual patient replicas that allow for the simulation of personalized treatments [8].
- Spatial Omics: Maps gene activity within the physical tissue architecture to understand disease context [14].
- Quantum Computing: Models complex atomic-level biological interactions that are impossible for classical computers [15].
- CRISPR-Cas9: A "word processor" for genetics that turns DNA editing into a programmable task [16].
- Cloud-Based Lab Automation: Enables remote execution of experiments, decoupling design from physical location [17].

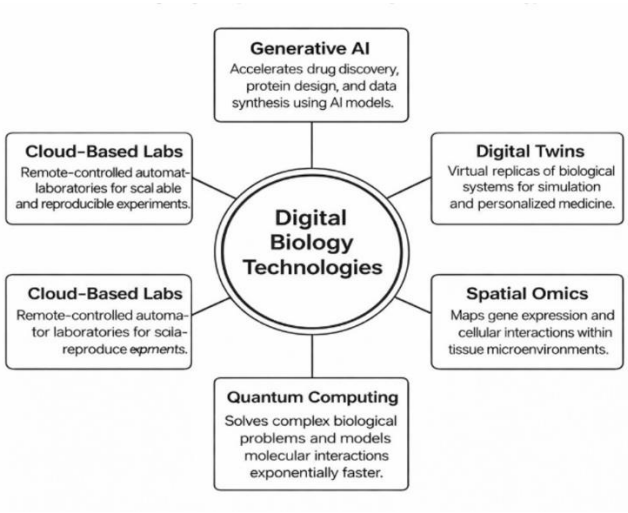

Fig. 3. Emerging Digital Technologies in Biology

1.3 Security & Ethical Challenges in Digital Biology

The integration of biological research with digital infrastructure establishes a unique cyber-physical attack surface. Digitizing biological assets introduces cybersecurity vulnerabilities where data breaches or manipulation can impact physical systems. Consequently, the field of Cyberbiosecurity has emerged to address vulnerabilities within biological supply chains, databases, and manufacturing infrastructure [18].

A. AI-Driven Biological Threats

The dual-use potential of biological design tools lowers the technical barriers associated with developing biological agents. Generative AI models could be manipulated or bypassed to output design specifications for regulated biological agents [19].

B. Vulnerabilities in Cloud & Automated Labs

Cloud-enabled automation introduces vulnerabilities in the cyber-physical interface, where unauthorized modifications to experimental protocols can alter robotic instructions. Such changes could modify chemical formulations or culture parameters, potentially compromising the integrity of research or pharmaceutical production [20].

C. Digital Twin & Genomic Privacy

Digital twins rely on extensive datasets, raising privacy challenges. A key vulnerability is the risk of re-identification attacks on genomic data. Unlike traditional credentials, genomic profiles cannot be modified post-compromise. Unauthorized cross-referencing of anonymized research data with public records can expose individual identities, potentially leading to unauthorized profiling. This necessitates robust cryptographic techniques, such as homomorphic encryption, and strict data governance policies [21], [22].

D. Biological Data Poisoning

Attackers can manipulate public training datasets to corrupt the predictive accuracy of AI models. By subtly altering genomic sequences, this "poisoning" can lead to flawed drug designs or misdirected research efforts [23].

E. DNA Supply Chain Vulnerabilities

The physical supply chain faces "substitution attacks" where reagents are swapped or altered. Sophisticated "screening evasion" techniques also allow attackers to bypass biosecurity algorithms when ordering pathogenic genetic sequences [25].

F. Intellectual Property (IP) Theft

Digitizing biology transforms proprietary genetic designs into stealable files. This makes novel protein structures and bio-manufacturing protocols vulnerable to cyber-espionage by state actors or competitors [24].

II. LITERATURE SURVEY

2.1 Evolution of Digital Biology

Biomedical research has fundamentally changed over the last few decades, moving from a purely observational science to one that looks more like data engineering. In the past, biology was all about manual records and microscopes. However, as experiments began generating massive amounts of high-throughput data, it became impossible for humans to manage it all by hand. This need for computational power is what gave rise to Digital Biology. It officially merges computer science with life sciences to process biological information faster than ever before [3]. The real turning point came with "digital organisms" computer programs that could replicate and evolve in a virtual world. These simulations gave scientists a way to test evolutionary rules in a controlled environment (in silico) that would have been impossible to observe in a physical lab, laying the groundwork for the modern focus on biodigitization and simulation [1], [2].

2.2 Emerging Technologies in Digital Biology

Recent progress has brought six key technologies to the forefront, each changing how we approach biological problems.

Generative AI moves beyond just analyzing data; it actually creates new biological designs. Tools like NVIDIA BioNeMo are currently being used to generate entirely new protein structures and potential medicines from scratch, rather than waiting to discover them. This shift effectively turns drug discovery into a predictive engineering challenge, drastically cutting down development timelines [7].

Digital Twins allow doctors to build a dynamic, virtual copy of a patient's physiology. Instead of relying on generalized protocols, physicians can use these simulations to test how a specific patient will react to a drug. This brings us closer to true Precision Medicine, where treatments are optimized on a computer before they ever reach the patient [8].

Spatial Omics is changing how we look at tissue samples. Standard sequencing methods often "blend" cells together, losing important context. Spatial technologies, however, map gene activity right where it happens in the tissue. This lets researchers see exactly how cells interact in a tumor's microenvironment, which is critical for understanding complex diseases like cancer [14].

Quantum Computing is stepping in where classical computers struggle. Since biological processes happen at the atomic level, they are inherently quantum mechanical. Quantum processors can model these complex molecular interactions like protein folding with a level of natural accuracy that traditional supercomputers simply cannot match, opening new doors in molecular research [15].

CRISPR-Cas9 has grown from a simple gene-editing tool into a programmable platform. Scientists can now use digital definitions to guide biological "cursors," editing DNA sequences with the same kind of precision you would expect from a word processor. Ideally, this turns genetic code into something that can be written and debugged like software [16].

Cloud-Based Lab Automation is separating the scientist from the physical bench. Researchers can now design an experiment on a laptop and have it executed by robots in a remote cloud lab. This not only allows for 24/7 experimentation but also ensures that results are highly reproducible, making high-end research tools accessible to more people [17].

TABLE I

COMPARISON OF DIGITAL BIOLOGY TECHNOLOGIES

Technology	Area of Work	Primary Benefits
Generative AI [7]	Drug Discovery & Protein Design	Accelerates development timelines; enables the creation of novel therapeutics from scratch
Digital Twins [8]	Precision Medicine	Supports personalized treatment simulation; reduces adverse drug reactions
Spatial Omics [14]	Tissue Mapping & Oncology	Enables high-resolution cellular mapping; improves understanding of disease microenvironments
Quantum Computing [15]	Molecular Simulation	Models complex atomic interactions; addresses computational problems beyond classical systems
CRISPR-Cas9 [16]	Genetic Engineering	Allows precise DNA editing; enables programmable genetic modification
Cloud-Based Labs [17]	Remote Experimentation	Provides high reproducibility; offers scalable 24/7 access to physical laboratory infrastructure

2.3 Security Challenges in Digital Biology

A. AI-Driven Biological Threats

The accessibility of generative design platforms presents dual-use challenges, where tools optimized for therapeutic research

could be repurposed. Specifically, models could be exploited to design novel pathogens or bypass sequence screening protocols. This computational support lowers the technical barriers to designing biological components, making safety-by-design frameworks essential [19].

B. Vulnerabilities in Cloud & Automated Labs

Cloud-based laboratories transmit digital protocols to automated hardware, establishing a cyber-physical vulnerability. Interception or alteration of protocol files could manipulate chemical ratios, incubation temperatures, or robotic movements. This form of sabotage threatens experimental reproducibility, research integrity, and the safety of synthesized pharmaceutical batches [20].

C. Digital Twin & Genomic Privacy

Constructing virtual models of patient physiology requires high-resolution patient profiles, including genomic data. The primary security vulnerability is the risk of re-identification. Because genomic profiles are immutable, data exposure is irreversible. Threat actors can cross-reference anonymized datasets with genealogical or public databases to identify individual participants, highlighting the need for privacy-preserving computation [21].

D. Biological Data Poisoning

The integrity of biological datasets is paramount for AI development. "Data poisoning" represents a subtle yet catastrophic threat where attackers intentionally corrupt the training data used for generative models. By introducing invisible manipulations—such as altering a few amino acid sequences in a public protein repository—malicious actors can skew the learning process of AI. This doesn't just break the model; it causes it to unknowingly produce flawed drug candidates or toxic molecular designs, potentially wasting years of critical research resources before the error is ever detected [23].

E. DNA Supply Chain Vulnerabilities

The physical supply chain for synthetic DNA represents a critical vulnerability. Substitution attacks involve the interception and alteration of reagents or DNA samples in transit. Furthermore, screening evasion techniques utilize specialized algorithms to obfuscate regulated gene sequences, potentially bypassing synthesis provider security protocols. This presents biosecurity risks, as malicious actors could procure fragmented components of pathogens for assembly without triggering automated synthesis alerts [25].

F. Intellectual Property (IP) Theft

As biology transitions into an information science, proprietary biological assets—such as engineered strain designs and novel pharmaceutical formulations—are digitized, increasing vulnerability to cyber-espionage. Unlike physical theft, the digital extraction of genomic assets leaves the original files intact, delaying detection of compromised intellectual property. This bio-data theft threatens the economic viability of biotechnology organizations and raises national security concerns regarding the proliferation of sensitive biological capabilities [24].

Fig. 4. Key Cyber biosecurity Challenges

2.4 Prevention and Mitigation Strategies

To reduce these risks, researchers and organizations must combine technical safeguards with strict governance policies. The following table summarizes practical mitigation strategies aligned with key cyberbiosecurity challenges.

TABLE II

CYBER BIOSECURITY PREVENTION AND MITIGATION

Security Challenge	**Prevention and Mitigation Strategies**
AI-Driven Threats [19]	•Adversarial Training to refuse harmful designs. • API Access Control. • Output Screening against pathogen databases.
Cloud Lab Risks [20]	• End-to-End Encryption. • Digital Signatures for protocols. • Anomaly Detection on robots.
Genomic Privacy (Re-identification)	• Differential Privacy (noise injection). • Homomorphic Encryption. • Strict Data Governance (logging).
Data Poisoning	•Data Provenance Tracking (Blockchain). •Cryptographic Checksums for datasets. • Model Robustness Testing.
Supply Chain	• "Know Your Customer" (KYC) for DNA orders. • Enhanced Screening Algorithms (Functional prediction). • Physical Security for reagents.

III. COMPARATIVE ANALYSIS

To understand the practical impact of digital biology, it is useful to compare it with traditional biological approaches. The shift is not only about speed, but also about enabling capabilities that were previously difficult or impossible using conventional methods. This section compares efficiency, evaluates security risks, and examines how close these technologies are to real-world clinical use.

3.1 Efficiency: Traditional vs. Digital Biology

The first comparison focuses on efficiency. Traditional biological research relies heavily on manual processes and trial-and-error experimentation, which can be slow and resource-intensive. In contrast, digital biology uses computational models, automation, and predictive tools to speed up discovery and improve precision. Table 3 highlights the key differences between these two approaches.

TABLE III

EFFICIENCY ANALYSIS (TRADITIONAL VS. DIGITAL BIOLOGY).

Parameter	**Traditional Biology**	**Digital Biology**	**Impact**
Drug Discovery	Trial-and-error; physical screening; 10–15	Predictive: AI designs structures in silico; reduces candidate screening from 5–6 years to under 12	High: Discovery to design.

	years.	months, decreasing costs by up to 50% [7].	
Patient Care	Generalized treatment based on population averages.	Personalized: Digital twins simulate patient response, potentially reducing adverse drug reactions by up to 30% [8].	High: Fewer adverse reactions.
Resolution	Bulk sequencing; averages gene expression.	Spatial omics: Cell-level mapping [14].	Medium: Better disease context.
Experiments	Manual processes limited by human labor.	Cloud-based automated labs; remote 24/7 access [17].	High: Democratized access.

3.2 Risk Assessment: Weighing the Dangers

While digital biology offers many advantages, it also introduces new security concerns. Some threats are more likely to occur but manageable, while others may be less frequent yet potentially severe in impact. This section categorizes major cyberbiosecurity risks and compares their likelihood and consequences

TABLE IV

CYBERBIOSECURITY RISK IMPACT ANALYSIS

Threat Type	**Description**	**Likelihood**	**Potential Impact**
Dual-Use AI	Manipulation of AI to design pathogens [19].	High	Critical: Rapid bioweapon development
Re-ID Attacks	Deanonymizing genomic data [21].	Medium	High: Permanent privacy loss
Cyber-Physical	Hacking lab robots [20].	Low	Critical: Tainted medicines
Data Poisoning	Corrupting AI training datasets [23].	Medium	High: Flawed drug discovery
Supply Chain	Evading synthesis screening [25].	Medium	Critical: Pathogen synthesis
IP Theft	Stealing genetic designs/data [24].	High	Medium: Economic espionage

3.3 Technology Maturity and Clinical Readiness Assessment

Not all digital biology technologies are equally mature. Some tools are already widely used in research, while others are still in experimental or pilot stages. This section evaluates the maturity and clinical readiness of key technologies to understand which are closer to real-world healthcare adoption

TABLE V

TECHNOLOGY MATURITY & READINESS INDEX

Technology	Research Maturity	Clinical Adoption	Status
CRISPR-Cas9	Mature	Early	Widely used in laboratories; limited human trials underway[16].
Generative AI	Rapidly growing	Experimental	Dominating research and drug discovery; not yet standard in clinics[7].
Digital Twins	Growing	Pilot	Used in select cardiovascular and oncology cases[8].

IV. CONCLUSION

The transition to Digital Biology is not merely an upgrade to existing tools, but a fundamental shift in how biological systems are understood and engineered. As this paper has demonstrated, the ability to model, simulate, and edit biological systems in silico is already delivering real-world benefits, from faster drug development to personalized cancer treatments. Our comparative analysis confirms that digital methods vastly outperform traditional approaches in terms of speed and scalability. Yet, this power brings a new duality: the same AI systems used for therapeutic design may also pose dual-use risks if misused. Ultimately, the future of this field depends on finding a balance. We cannot have the benefits of "programmable biology" without accepting the responsibility of securing it. The next decade will not just be about discovering new biology, but about engineering the digital safeguards to ensure these discoveries remain safe, ethical, and beneficial for humanity

REFERENCES

[1] Wilke, C.O., Adami, C.: "The biology of digital organisms." Trends in Ecology & Evolution 17(11), 528–532 (2002).

[2] O'Neill, B.: "Digital evolution." PLoS Biology 1(1), e18 (2003).

[3] Witters, D., et al.: "Digital biology and chemistry." Lab on a Chip 14(17), 3225–3232 (2014)

[4] Jumper, J., et al.: "Highly accurate protein structure prediction with AlphaFold." Nature 596, 583–589 (2021).

[5] Elhefnawy, N. S.: "AlphaFold-Based Protein-Protein Interaction Prediction Methods Classification." In: Proc. 2025 5th Intl Conf. on Emerging Smart Technologies and Applications (eSmarTA), 2025.

[6] Kasoju, N., et al.: "Digital health: Trends, opportunities and challenges in medical devices sector." J Family Med Prim Care 12(2), 123–129 (2023).

[7] Ho, A.V., et al.: "BioNeMo Framework: A Modular Library for AI Model Development in Drug Discovery." arXiv preprint arXiv:2406.09573 (2024).

[8] Wu, J., et al.: "Towards generative digital twins in biomedical research." Comput Struct Biotechnol J 22, 123–135 (2024).

[9] Bobak, D., et al.: "Development of Online Anatomical Atlas for Biomedical Informatics." In: Proc. 2023 21st Intl Conf. on Emerging eLearning Technologies (ICETA), pp. 68–75 (2023).

[10] Al-Sarem, M. M., et al.: "Deep Learning and Generative AI in Cancer Medical Imaging." IEEE Access 12, 1–15 (2024).

[11] Huang, J., Shi, Q., Xie, D., Duan, X.: "Visuomotor Policy Learning for Task Automation of Surgical Robot." IEEE Trans Med Robot Bionics 6(1), 200-210 (2024).

[12] Murphy, S.V., Atala, A.: "3D bioprinting of tissues and organs." Nature Biotechnology 32(8), 773–785 (2014).

[13] McDermott, J. R., McBride, J. R., Taylor, N. J.: "Laboratory automation in life sciences: current trends and future perspectives." SLAS Technology 26(3), 201–213 (2021).

[14] Gao, J., Hu, K., Zhang, F., Cui, X.: "Hexagonal Convolutional Neural Network for Spatial Transcriptomics Classification." In: 2022 IEEE Intl Conf. on Bioinformatics and Biomedicine (BIBM) (2022).

[15] D'Acunto, M.: "Quantum Computation by Biological Systems." IEEE Trans Mol Biol Multi-Scale Commun 9(2), 257–262 (2023).

[16] Kavitha, R., et al.: "Emerging technologies in biological sciences research methodologies." Cuestiones de Fisioterapia 54(2), 1042–1050 (2025).

[17] Chaudhari, S., et al.: "Engineering End-to-End Remote Labs Using IoT-Based Retrofitting." In: Proc. IEEE Int. Conf., 2025.

[18] Liu, T., et al.: "Cyberbiosecurity: Advancements in DNA-based information storage and security." Front BioengBiotechnol 12 (2024).

[19] Sandbrink, E.: "Artificial intelligence and biological misuse: Differentiating risks of language models and biological design tools." arXiv preprint arXiv:2307.10098 (2023).

[20] Ijaz, M.N., O'Neil, J.G.: "Cyber-Biosecurity Challenges in Next-Generation Sequencing: A Comprehensive Analysis of Emerging Threat Vectors." In: Proc. IEEE Conference, 2024.

[21] Berger, K.M., Schneck, P.A.: "National and transnational security implications of asymmetric access to biological data." Front BioengBiotechnol 7 (2019).

[22] Laith, A.L.E.: "Cyberbiosecurity in the age of biotechnology and AI." J Global Biosecurity 1(1), 1–10 (2019).

[23] Ney, P., et al.: "Computer Security, Privacy, and DNA Sequencing: Compromising Computers with Synthesized DNA." USENIX Security Symposium (2017).

[24] Richardson, L.C., et al.: "Cyberbiosecurity: A New Discipline to Help Safeguard the Bioeconomy." Front. Bioeng. Biotechnol. 7 (2019).

[25] Schabacker, D., et al.: "Assessing the Security of the DNA Synthesis Pipeline." Applied Biosafety 24(2), 70-76 (2019).

DIALLOps: A DevOps-Integrated Framework for Secure andReliable Large Language Model Operations

Vignesh Mudaliyar[1], Mohit Rajput[2],Deep Solanki[3],Jigar Gajjar[4]

Lok Jagruti Kendra University,Ahmedabad,India[1,2]
TechDefence Cybersecurity Limited,Ahmedabad,India[12]

vigneshmudaliyar999@gmail.com[1],
mohitsingh.official2004@gmail.com[2],deepsolanki5799gmail.com[3],cyberian.jg@gmail.com[4]

Abstract: The rapid integration of large language models (LLMs) into enterprise and mission-critical systems has introduced a new class of operational and cybersecurity challenges that extend beyond traditional software and machine learning deployments. While recent advances in LLM capabilities enable powerful generative applications, their probabilistic, non-deterministic, and interaction-driven behaviour raises significant concerns related to reliability, hallucinations, governance, and secure deployment. Existing DevOps and MLOps frameworks, which assume deterministic execution and static evaluation, are insufficient to manage these risks in production environments.

This paper proposes DIALLOps, a DevOps-Integrated Adaptive Lifecycle framework for Large Language Model Operations, designed to align generative AI systems with security-aware DevOps automation. DIALLOps introduces a unified operational abstraction that treats models, prompts, retrieval pipelines, runtime policies, and agent logic as first-class artifacts, enabling consistent versioning, deployment, rollback, and governance. The framework redefines continuous integration and deployment semantics for probabilistic systems through distribution-aware validation, continuous in-production evaluation, and human-in-the-loop oversight.

The proposed architecture combines evaluation, observability, governance, and cost-aware control within a centralized operational control plane, enabling proactive mitigation of hallucinations, behavioral anomalies, and policy violations. A prototype implementation of DIALLOps was developed using Python, Ollama, and the Mistral large language model. Experimental evaluation using 50 test cases and probabilistic consistency testing achieved a policy compliance rate of 90%, consistency scores ranging from 0.7267 to 0.9610, and successful automated deployment gating through governance-aware rollback mechanisms. The findings suggest that DIALLOps offers a practical and security-oriented foundation for deploying LLM-based systems reliably at scale, positioning LLMOps as a critical operational layer for trustworthy and secure generative AI adoption.

Keywords:Large Language Models · LLMOps · Generative AI Security · DevOps Automation · AI Governance

---

I. INTRODUCTION

A. Rise of Generative AI and Large Language Models

AI has been defined in multiple literature sources as the science and engineering of making smart machines, especially clever computer programs, which can be used to perform tasks normally requiring human intelligence ([1] McCarthy, 2004). Building on decades of work in machine learning and neural networks, generative artificial intelligence is a class of computer

http://doi.org/10.63766/spujstmr.26.000094

algorithms that can create new and meaningful content (e.g., text, images, or audio) modeled from patterns learned in large amounts of training data ([2] Feuerriegel et al., 2024).

Large language models (LLMs) pretrained in an unsupervised manner on transformer-based architectures and massive volumes of data at scale have enabled novel levels of natural language understanding and generation. Archetypical examples are GPT-4, DALL·E 2, and GitHub Copilot. These state-of-the-art language models have accelerated the common use of generative AI in consumer and enterprise systems ([2] Feuerriegel et al., 2024). Importantly, these systems can increasingly be seen as general-purpose cognitive tools rather than domain-specific automation solutions.

Within organizational settings, LLM-based systems are now widely adopted as intelligent assistants supporting knowledge-intensive work, including software development, customer support, and IT help desk operations. These systems automate routine tasks, provide contextualized question–answering capabilities, and assist in problem solving, thereby augmenting human decision-making rather than merely replacing it ([2] Feuerriegel et al., 2024). Beyond enterprise use, generative AI has also become embedded in everyday activities, addressing information needs such as travel planning, cooking, and preliminary health guidance.

The economic and societal implications of this shift are substantial. Industry analyses estimate that generative AI could increase global gross domestic product by up to 7%, while simultaneously transforming labor markets and potentially displacing hundreds of millions of knowledge-based roles ([3] Goldman Sachs, 2023, as cited in [2] Feuerriegel et al., 2024). These projections underscore not only the strategic importance of generative AI and LLMs, but also the growing necessity for robust operational frameworks capable of ensuring their reliable, scalable, and responsible deployment.

B. Limitations of Deploying LLMs in Real-World Environments

Despite their impressive capabilities, LLMs exhibit several limitations that hinder their dependable use in real-world production environments. Among these, hallucinations represent one of the most critical challenges, as they directly undermine user trust and system reliability. Hallucinations occur when LLMs generate outputs that are fluent and coherent yet factually incorrect, nonsensical, or unfaithful to source information, making it difficult for users to assess the validity of generated responses ([3, 11,21] Farquhar et al., 2024).

While hallucinations are often discussed as model-level failures, recent research emphasizes that they constitute a broader class of reliability breakdowns with direct operational consequences. [[3,11,21] Farquhar et al. (2024) introduce the concept of confabulations, a particularly problematic subset of hallucinations in which models produce arbitrary but plausible claims that are sensitive to irrelevant factors such as random seeds or minor prompt variations. In such cases, identical queries may yield inconsistent and contradictory answers, even in high-stakes domains such as medicine.

This behavior is especially problematic in production settings because it exposes the fundamentally non-deterministic nature of LLM outputs. Unlike traditional software systems, which produce reproducible outputs for identical inputs, LLMs may generate divergent responses across executions, complicating validation, debugging, and quality assurance. Importantly, these failures differ from systematic errors caused by biased training data or persistent reasoning flaws. Instead, confabulations arise from the autoregressive generation process itself, in which models recursively condition on their own outputs, increasing the likelihood of fabricated content over longer generations ([3,11,21] Farquhar et al., 2024).

Fig 1 illustrates how non-deterministic generation and autoregressive feedback mechanisms contribute to hallucinated or confabulated outputs in large language model–based systems. In addition to factual inaccuracies, hallucinations exacerbate broader operational challenges by amplifying prompt sensitivity, increasing evaluation complexity, and accelerating trust erosion in deployed applications.

At scale, these behavioral limitations intersect with broader technical and operational challenges. Deploying LLMs in production environments requires substantial computational resources and careful management of memory, parallelism, and

key–value (KV) caches, all of which directly affect latency and cost ([5] Menshawy et al., 2024). Furthermore, evaluating LLM outputs remains difficult, as automated metrics often fail to capture nuanced errors such as plausible but incorrect responses. Consequently, organizations rely heavily on human-in-the-loop evaluation, raising concerns about scalability, subjectivity, and operational overhead ([5] Menshawy et al., 2024).

These limitations highlight a critical insight: hallucinations and non-deterministic behavior are not merely model deficiencies, but systemic operational challenges. Addressing them requires not only improved training techniques, but also deployment frameworks capable of monitoring, validating, and governing LLM behavior throughout the system lifecycle.

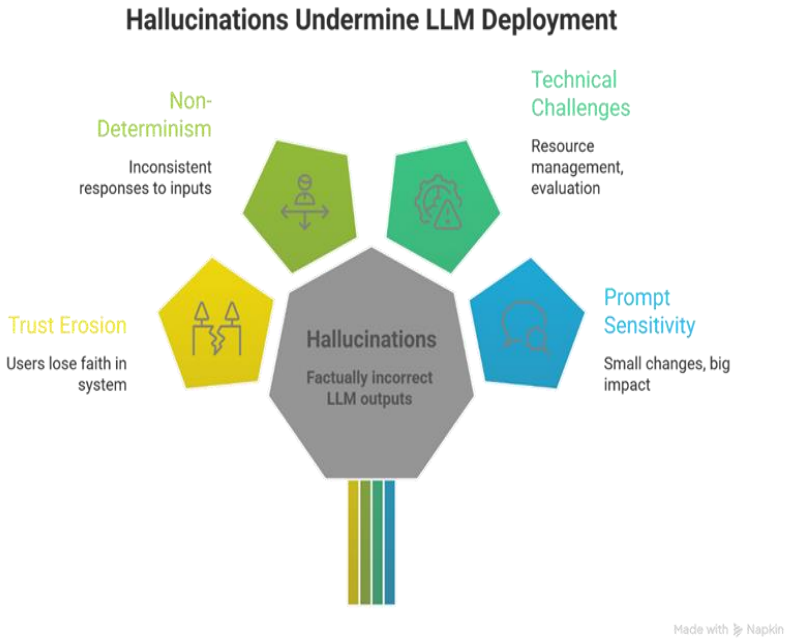

Fig 1. Key factors through which hallucinations impact reliability, trust, and deployment stability in LLM systems.

C. Inadequacy of Traditional DevOps and MLOps for LLM-Based Systems

Traditional DevOps practices were designed to support deterministic software systems, emphasizing continuous integration and delivery (CI/CD), infrastructure automation, and system observability. While these practices have proven effective for conventional applications, they assume predictable behavior, clearly defined specifications, and testable outputs. Large language model (LLM)-based systems violate many of these assumptions due to their probabilistic, data-driven, and non-deterministic nature, rendering standard DevOps pipelines insufficient for ensuring reliability and correctness ([6,14] Pahune& Akhtar, 2025).

To address the data-centric nature of machine learning systems, MLOps emerged as an extension of DevOps, introducing mechanisms for data versioning, model training pipelines, and performance monitoring. However, MLOps primarily focuses on supervised learning scenarios with well-defined objectives and quantitative evaluation metrics. In contrast, LLM-based systems rely heavily on unstructured data, emergent behaviors, and interaction-driven performance, which are not adequately captured by traditional MLOps workflows ([7,13,15] Diaz-De-Arcaya & López-De-Armentia, 2024).

One fundamental limitation of MLOps when applied to LLMs lies in its treatment of the model as the primary operational artifact. For LLM-based applications, system behavior is equally influenced by prompts, context windows, retrieval mechanisms, and generation parameters, all of which evolve rapidly and require continuous experimentation and version control. Existing DevOps and MLOps pipelines lack native support for managing these artifacts, leading to fragile

deployments and inconsistent behavior across environments ([8,17,38] Joshi, 2025).

Moreover, conventional CI/CD practices are ill-suited for validating LLM outputs. Unlike traditional software or classical ML models, LLMs cannot be exhaustively tested using predefined test cases, as outputs may vary across runs and subtle prompt changes can produce semantically different results. This challenges the effectiveness of automated testing and rollback mechanisms, which form the backbone of DevOps reliability engineering ([9,16,39] Özer, 2025).

Operational gaps also arise in monitoring and observability. Standard DevOps metrics such as latency, throughput, and error rates fail to capture critical quality dimensions of LLM outputs, including factual correctness, coherence, and alignment with user intent. While MLOps introduces model performance monitoring, it typically relies on static benchmarks that do not reflect real-world interaction dynamics or evolving knowledge requirements ([10,23] Radenkovic & Prokhorov, 2025).
These shortcomings have motivated the emergence of LLMOps as a distinct operational paradigm tailored to the unique characteristics of large language models. Recent studies emphasize that LLMOps is not merely an incremental extension of MLOps but requires rethinking lifecycle management, evaluation strategies, and human-in-the-loop governance mechanisms to support LLM-based systems in production environments ([6,14] Pahune& Akhtar, 2025; [8,17,38] Joshi, 2025). Consequently, the inadequacy of traditional DevOps and MLOps frameworks highlights the need for integrated operational models that explicitly bridge software engineering practices with the cognitive and socio-technical properties of generative AI system.

D. Motivation for LLMOps: Bridging Generative AI and DevOps Automation

The limitations discussed in the preceding sections highlight a growing mismatch between the operational characteristics of large language model (LLM)–based systems and the assumptions embedded in traditional DevOps and MLOps practices. While DevOps emphasizes automation, repeatability, and reliability for deterministic software systems, and MLOps extends these principles to data-driven machine learning pipelines, neither paradigm fully accounts for the probabilistic, interaction-driven, and continuously evolving nature of generative AI systems ([5,12,18] Menshawy et al., 2024;[7,13,15] Diaz-De-Arcaya & López-De-Armentia, 2024).

LLM-based applications introduce new operational artifacts—such as prompts, retrieval pipelines, context windows, and generation parameters—that play a decisive role in system behavior. Unlike conventional software components, these artifacts are tightly coupled with user interaction and may require frequent adaptation to maintain output quality and relevance. As a result, core DevOps mechanisms such as continuous integration, automated testing, and rollback strategies struggle to provide sufficient guarantees when applied to LLM-driven systems, particularly in the presence of non-deterministic outputs and hallucinations ([3,11,21] Farquhar et al., 2024).

MLOps partially addresses the gap by supporting model lifecycle management, monitoring, and retraining. However, existing MLOps frameworks typically treat the trained model as the central operational unit and rely on static evaluation metrics. In contrast, the performance of LLM-based systems emerges from the combined behavior of models, prompts, external knowledge sources, and runtime context, all of which may evolve independently. This limits the effectiveness of conventional MLOps pipelines in capturing real-world system quality, user trust, and alignment with organizational objectives ([5,12,18] Menshawy et al., 2024).

These challenges motivate the need for LLMOps as a unifying operational paradigm that bridges generative AI capabilities with DevOps automation principles. Rather than replacing DevOps or MLOps, LLMOps seeks to extend them by incorporating mechanisms for prompt and context versioning, continuous evaluation of generated outputs, human-in-the-loop oversight, and adaptive governance across the deployment lifecycle ([7,13,15] Diaz-De-Arcaya & López-De-Armentia, 2024). Crucially, LLMOps emphasizes operational robustness not solely through model performance improvements, but through systematic integration of monitoring, validation, and feedback loops that account for the socio-technical nature of LLM-based systems.

From an organizational perspective, this bridging role is essential for enabling scalable and responsible adoption of generative AI. As LLMs are increasingly embedded into production workflows—ranging from software engineering and customer support to decision assistance—organizations require operational frameworks that ensure reliability, traceability, and compliance without sacrificing the flexibility and expressiveness that make generative AI valuable. LLMOps thus emerges as a necessary evolution of DevOps and MLOps, aligning automation, governance, and continuous improvement with the unique demands of large language model–based systems.

E. Research Gap and Problem Statement

The rapid adoption of large language models (LLMs) in production environments has stimulated growing academic and industrial interest in operational practices for generative AI systems. Recent studies have begun to explore challenges related to hallucinations, scalability, evaluation, and deployment complexity, as well as the emergence of LLMOps as a distinct operational concept ([3,11,21] Farquhar et al., 2024; [5,12,18] Menshawy et al., 2024; [7,13,15] Diaz-De-Arcaya & López-De-Armentia, 2024). However, despite these advances, significant gaps remain in the current body of research.

First, there is a lack of standardized, end-to-end LLMOps frameworks that address the full lifecycle of LLM-based systems, from development and deployment to monitoring, evaluation, and continuous improvement. Existing work often focuses on isolated stages of the lifecycle—such as prompt engineering, retrieval-augmented generation, or post-deployment evaluation—without offering a cohesive operational model that integrates these components into a unified workflow.
Second, much of the current literature emphasizes tool-level solutions rather than system-level design. Proposed approaches frequently introduce new evaluation metrics, monitoring tools, or deployment utilities, but stop short of defining how these elements should be orchestrated within a consistent operational architecture. As a result, organizations are left to manually assemble fragmented toolchains, leading to brittle systems, duplicated effort, and increased operational risk when deploying LLMs at scale.

Third, there is a notable absence of automation-driven deployment models tailored to the non-deterministic and interaction-driven nature of LLMs. While traditional DevOps pipelines rely heavily on automated testing, deterministic builds, and reproducible deployments, these assumptions do not hold for generative AI systems. Current MLOps extensions provide limited support for automating prompt validation, output quality assessment, or human-in-the-loop feedback, which are essential for maintaining reliability in LLM-based applications.

Finally, existing approaches exhibit limited integration of core DevOps principles into generative AI systems. Concepts such as continuous integration, continuous deployment, observability, and rollback are rarely reinterpreted in ways that account for probabilistic outputs, prompt sensitivity, and evolving external knowledge sources. This gap prevents organizations from fully leveraging DevOps automation while ensuring trust, governance, and accountability in LLM deployments.

ProblemStatement:

Despite growing recognition of the operational challenges associated with large language models, there is currently no unified, automated framework that systematically integrates DevOps principles with generative AI–specific requirements to support the reliable, scalable, and governable operationalization of LLM-based systems. This gap motivates the need for a holistic LLMOps model that bridges DevOps automation and generative AI workflows across the entire system lifecycle.

F. Research Aim and Objectives

The primary aim of this research is to propose an automated LLMOps framework that bridges generative AI systems with DevOps automation to enable reliable, scalable, and cost-efficient deployment of large language models in real-world environments. Rather than presenting a purely analytical study, the paper adopts a solution-oriented approach, focusing on the design of an operational framework tailored to the unique characteristics of LLM-based systems.

Fig 2 provides a conceptual overview of the key objectives addressed by the proposed LLMOps framework across the LLM lifecycle. As illustrated, the framework is designed to support end-to-end LLM operations, integrating DevOps automation

principles with generative AI–specific controls to address reliability, scalability, security, and performance challenges. Specifically, the framework aims to enable continuous evaluation and hallucination mitigation, seamless integration with enterprise infrastructure, governance-aware security and compliance controls, and performance optimization with cost-aware scaling. By consolidating these objectives within a unified automated architecture, the proposed approach positions LLMOps as a necessary operational layer that extends beyond traditional MLOps to effectively manage the non-deterministic and interaction-driven nature of large language models

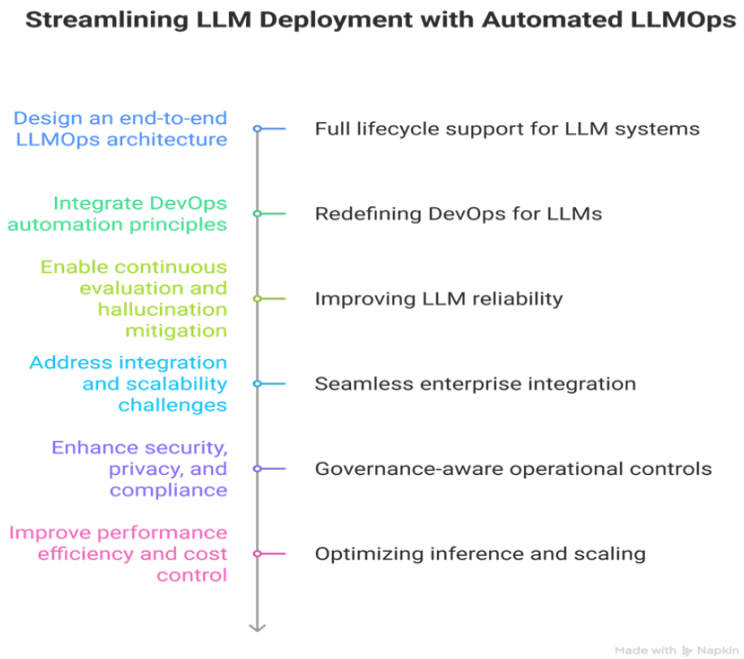

Fig 2. End-to-end LLMOps automation framework for streamlining deployment, evaluation, integration, and operational governance.

G. Key Contributions of This Work

- This paper makes the following key contributions to the operationalization of large language model–based systems.
- We introduce DIALLOps, a DevOps-integrated operational framework for LLM-based systems that explicitly addresses non-determinism, interaction-driven behavior, and socio-technical deployment constraints.
- We redefine CI/CD semantics for generative AI by proposing probabilistic, behavior-aware automation pipelines that operate on distributions of LLM outputs rather than deterministic executions.
- We propose a unified LLM-centric system artifact that encapsulates models, prompts, retrieval pipelines, runtime policies, and agent logic as a single operational unit, enabling consistent versioning, deployment, and rollback.
- We design an end-to-end LLMOps architecture that integrates continuous evaluation, observability, governance, and cost-aware control directly into the operational lifecycle of LLM-based applications.
- We present a comprehensive evaluation methodology for assessing reliability, performance, scalability, and cost efficiency in production-like environments, demonstrating the practical applicability of the proposed framework.

H. Organization of the Paper

The remainder of this paper is organized as follows. Section 2 reviews relevant background and related work, covering foundational concepts in DevOps, MLOps, and emerging LLMOps approaches. Section 3 introduces the proposed automated LLMOps framework, presenting its design principles, architecture, and core components. Section 4 details the implementation aspects of the framework, including integration considerations and operational workflows. Section 5 presents an experimental evaluation to assess the framework's effectiveness in improving reliability, scalability, performance, and cost efficiency. Section 6 discusses the results, practical implications, and limitations of the proposed approach. Finally, Section 7

concludes the paper and outlines directions for future

II. RESEARCH GAPS IN EXISTING LLMOPS AND DEVOPS-ORIENTED FRAMEWORKS

Although LLMOps has emerged as a response to the operational challenges posed by large language models, current research remains fragmented and lacks a unified perspective. This section systematically identifies the key gaps in existing DevOps-, MLOps-, and LLMOps-oriented approaches that motivate the need for a new operational model.

TABLE I
COMPARISON OF EXISTING DEVOPS-, MLOPS-, AND LLMOPS-ORIENTED FRAMEWORKS

Dimension	DevOps	MLOps	LLMOps	AgentOps
Primary System Type	Deterministic software systems	Predictive machine learning models	Generative LLM-based systems	Multi-step, autonomous LLM agents
Determinism Assumption	Fully deterministic	Mostly deterministic	Probabilistic	Highly stochastic
Core Operational Artifact	Source code	Model and dataset	Model and prompts	Agent workflows and policies
CI/CD Semantics	Binary pass/fail pipelines	Metric-based gating	Limited or experimental support	Experimental and scenario-driven
Testing Strategy	Unit and integration testing	Offline benchmarks and validation sets	Prompt testing and manual evaluation	Scenario-based and trace-level testing
Deployment Control	Version-level rollback	Model rollback	Partial support (session or prompt-level)	Session-level rollback
Evaluation Timing	Pre-deployment	Pre- and post-deployment	Mostly offline or ad hoc	Runtime tracing and inspection
Hallucination Handling	Not applicable	Indirect (via retraining)	Detection tools and heuristics	Agent monitoring and guardrails
Observability Scope	Infrastructure metrics and logs	Model performance metrics	Limited behavioral observability	Agent execution traces
Governance Integration	External to pipelines	Partial and compliance-driven	Ad hoc or tool-dependent	Minimal or experimental
Cost Awareness	Infrastructure-centric	Training-centric	Token-level cost tracking	Execution-level cost awareness
Automation Maturity	High	Medium	Emerging	Experimental

Note: Table 1 presents a synthesized comparison based on prior literature and is not reproduced from a single source. DevOps and MLOps characteristics are derived from [17,24], LLMOps from [6–8], and AgentOps from [25].

A. Lack of End-to-End LLMOps Frameworks

Current LLMOps research predominantly focuses on isolated stages of the system lifecycle, such as prompt engineering, hallucination detection, evaluation metrics, or deployment optimization. While these contributions are valuable, they rarely coalesce into a single, end-to-end operational framework that governs the full lifecycle of LLM-based systems, from development to continuous improvement ([5,12,18] Menshawy et al., 2024; [7,13,15] Diaz-De-Arcaya & López-De-Armentia, 2024).

As a result, organizations are often required to manually integrate disparate tools and practices, leading to brittle architectures, duplicated operational effort, and increased technical debt. The absence of a cohesive lifecycle model limits reproducibility, scalability, and systematic governance of LLM-based applications ([8,17,38] Joshi, 2025; [9,16,39] Özer, 2025).

B. Inadequate Adaptation of DevOps Automation to Non-Deterministic Systems

Traditional DevOps methodologies assume deterministic behavior, reproducible builds, and binary pass–fail testing outcomes. These assumptions do not hold for LLM-based systems, whose outputs are probabilistic, prompt-sensitive, and influenced by evolving context and external knowledge sources ([6,14] Pahune& Akhtar, 2025).

From a cybersecurity standpoint, this mismatch introduces significant operational risk. Applying deterministic CI/CD automation to non-deterministic LLM behavior creates blind spots in testing and deployment pipelines, allowing unsafe, misleading, or policy-violating outputs to reach production environments without detection. While recent studies acknowledge this challenge, existing LLMOps approaches provide limited guidance on how core DevOps concepts—such as deployment gating, rollback strategies, and reliability engineering—should be reinterpreted to account for uncertainty and behavioral variability ([9,16,39] Özer, 2025).

As a result, DevOps automation is often applied superficially to LLM systems, failing to address the fundamental security and reliability implications of probabilistic generation.

C. Absence of Unified Operational Abstractions for LLM-Based Systems

In conventional software engineering and MLOps, operational artifacts such as source code, datasets, and trained models serve as well-defined units for versioning, testing, and deployment. In contrast, LLM-based system behavior emerges from the interaction of multiple evolving artifacts, including prompts, context windows, retrieval pipelines, generation parameters, and user inputs ([5,12,18] Menshawy et al., 2024).

Existing LLMOps literature largely treats these elements in isolation, without defining higher-level operational abstractions that capture their interdependencies. This lack of abstraction complicates lifecycle management, reproducibility, and system-level reasoning, making it difficult to apply established DevOps principles in a consistent and scalable manner ([8,17,38] Joshi, 2025).

D. Limited Automation for Continuous Evaluation and Hallucination Mitigation

Hallucinations have been widely recognized as a critical barrier to trust and adoption of LLM-based systems. Although significant progress has been made in hallucination detection and evaluation metrics at the model level ([3,11,21] Farquhar et al., 2024), these techniques are rarely integrated into automated operational pipelines.

In security- and compliance-sensitive domains, hallucinations represent more than quality defects; they constitute potential misinformation, integrity, and regulatory risks. Current approaches often rely on manual reviews or offline evaluations, which do not scale effectively in production environments and fail to provide timely detection of unsafe behavior.
The lack of automation-driven evaluation and feedback loops prevents continuous quality assurance and weakens the system's ability to respond to behavioral anomalies in real time, thereby undermining both operational reliability and

cybersecurity posture of deployed LLM systems ([5,12,18] Menshawy et al., 2024).

E. Insufficient Integration of Governance, Security, and Compliance

Governance, security, and compliance considerations are increasingly emphasized in discussions of generative AI deployment, particularly in enterprise and regulated domains. However, existing LLMOps approaches frequently treat these concerns as external constraints rather than as integral components of the operational lifecycle ([6,14] Pahune& Akhtar, 2025; [10,23] Radenkovic & Prokhorov, 2025).

This separation limits cybersecurity-grade auditability, traceability, and policy enforcement, making it difficult for organizations to ensure responsible and compliant use of LLMs at scale. Without native integration into deployment and runtime workflows, governance mechanisms remain reactive rather than preventative.

A systematic integration of governance-aware and security-oriented controls into automated LLMOps workflows remains largely absent in current frameworks.

F. Summary of Research Gaps

Figure 3 synthesizes the key research gaps identified across existing DevOps, MLOps, and emerging LLMOps literature. Despite growing interest in operationalizing large language models, current approaches remain fragmented and insufficient to support the full lifecycle of LLM-based systems in production environments.

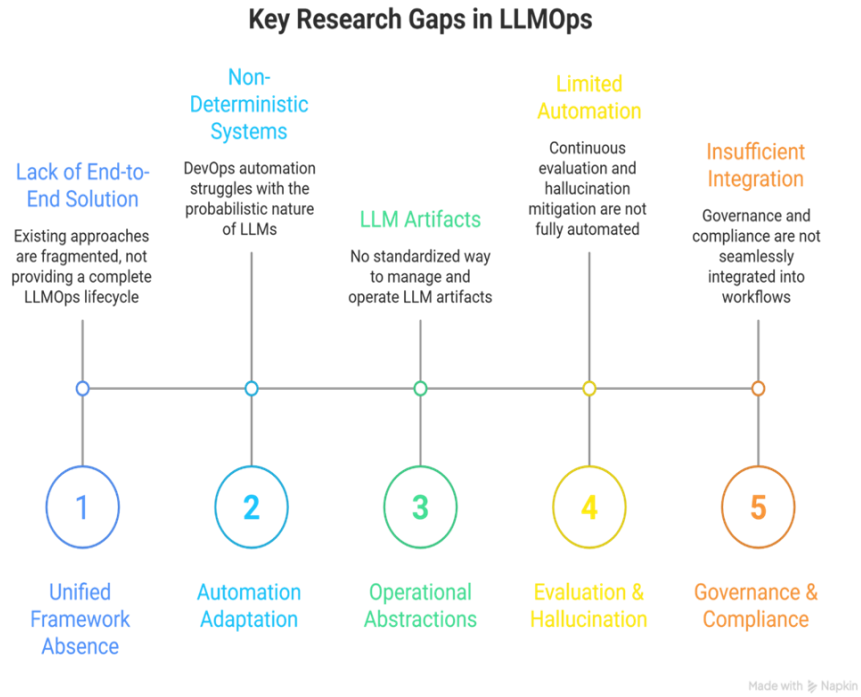

Fig. 3 Major research gaps and unresolved challenges limiting the maturity and standardization of current LLMOps practices.

Specifically, prior work reveals five persistent limitations: (i) the absence of unified, end-to-end LLMOps frameworks, (ii) inadequate adaptation of DevOps automation mechanisms to non-deterministic and probabilistic model behavior, (iii) a lack of standardized operational abstractions for managing LLM artifacts such as prompts, retrieval pipelines, and agent logic, (iv) limited automation for continuous evaluation and hallucination mitigation, and (v) insufficient integration of governance, compliance, and cost controls into operational workflows.

Collectively, these gaps highlight the need for a unified, automation-centric LLMOps model that systematically bridges mature DevOps principles with the interaction-driven and probabilistic nature of large language model–based systems. This observation directly motivates the automated LLMOps framework proposed in the following section.

III. PROPOSED MODEL: DIALLOPS FRAMEWORK

A. Model Name and Acronym

We propose a novel operational framework named:DIALLOps — DevOps-Integrated Adaptive Lifecycle for Large Language Model Operations

DIALLOps is designed as a first-principles operational model that unifies DevOps automation with the unique lifecycle, evaluation, and governance requirements of LLM-based systems. Unlike existing LLMOps approaches that evolve incrementally from MLOps or focus on tooling ecosystems, DIALLOps introduces a lifecycle-centric, automation-native, and reliability-aware operational abstraction for generative AI systems.

The core novelty of DIALLOps lies in its explicit treatment of LLM behavior as an emergent system property, shaped jointly by models, prompts, retrieval pipelines, runtime context, and user interaction. As such, DIALLOps elevates these elements to first-class operational artifacts, subject to versioning, validation, observability, and controlled rollout.
From a cybersecurity perspective, DIALLOps explicitly treats unreliable LLM behavior as an operational risk that must be monitored, controlled, and mitigated through continuous validation, policy enforcement, and controlled rollout mechanisms.

B. Design Principles of DIALLOps

The design of DIALLOps is guided by the following foundational principles, which collectively define how DevOps automation is adapted to the probabilistic and interaction-driven nature of large language model–based systems. **Figure 4** provides an overview of these principles and illustrates how they jointly support reliable, secure, and governance-aware LLM operations.

Artifact Parity Between Software and Cognition

DIALLOps treats prompts, retrieval configurations, context policies, and generation parameters as operational artifacts equivalent to source code and infrastructure definitions. This principle extends DevOps artifact management concepts to the cognitive layer of LLM systems, addressing gaps identified in recent LLMOps maturity studies ([36] Borovits & Tamburri, 2025).

Probabilistic-Aware Automation

Recognizing that LLM outputs are inherently non-deterministic, DIALLOps redefines CI/CD not as binary pass–fail pipelines, but as probabilistic validation workflows that operate on distributions of outputs rather than single executions.
This approach enables detection of unsafe, misleading, or policy-violating behavior prior to full-scale deployment, thereby supporting security-aware release decisions. By embedding uncertainty-aware validation into automation pipelines, DIALLOps addresses limitations in conventional testing approaches identified in recent LLMOps surveys ([9,16,39] Özer, 2025).

Continuous Quality Evaluation in Production
DIALLOps embeds evaluation mechanisms directly into runtime operations, enabling continuous assessment of factuality, coherence, relevance, and safety during live system usage. This aligns with emerging research emphasizing production-aware evaluation over static benchmarking ([35] Abbassi, 2025).

Human-in-the-Loop as an Operational Primitive

Rather than treating human feedback as an external or post hoc process, DIALLOps integrates human oversight as a native component of the operational lifecycle. Human-in-the-loop mechanisms act as a final safeguard in high-impact scenarios, supporting accountability, risk mitigation, and governance in security-sensitive deployments.
This design aligns with recent responsible LLMOps research emphasizing the importance of human oversight for managing safety, compliance, and trust in generative AI systems ([25,37] Tantithamthavorn et al., 2024).

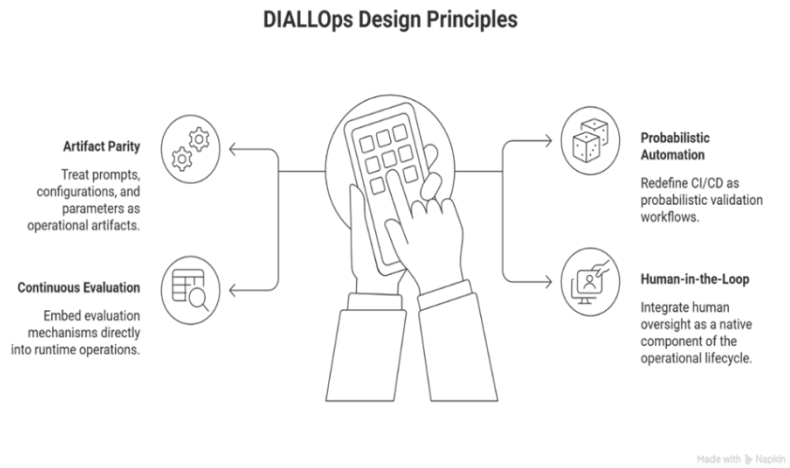

Fig. 4 Core DIALLOps design principles enabling probabilistic automation, continuous evaluation, artifact parity, and human-in-the-loop control.

C. DIALLOps System Architecture

The architecture is structured around six tightly coupled operational layers:

- **Development Layer**
  Includes prompt engineering, retrieval design, and model configuration. All artifacts are version-controlled and validated through probabilistic test suites.

- **Integration Layer**
  Extends traditional CI pipelines to include prompt regression testing, retrieval consistency checks, and semantic output validation across prompt variants.

- **Deployment Layer**
  Supports staged and canary deployments of LLM artifacts (models, prompts, RAG pipelines) with controlled exposure to live traffic.

- **Inference & Runtime Layer**
  Manages request routing, KV-cache optimization, latency-aware batching, and adaptive scaling strategies to control cost and performance.

- **Evaluation & Monitoring Layer**
  Continuously measures semantic entropy, factual consistency, safety constraints, and user-alignment signals during production inference.

- **Governance & Feedback Layer**
  Aggregates audit logs, compliance signals, and human feedback, feeding corrective actions back into earlier lifecycle stages.

This layered architecture enables closed-loop adaptation, allowing DIALLOps to respond dynamically to performance degradation, hallucination spikes, or cost anomalies.

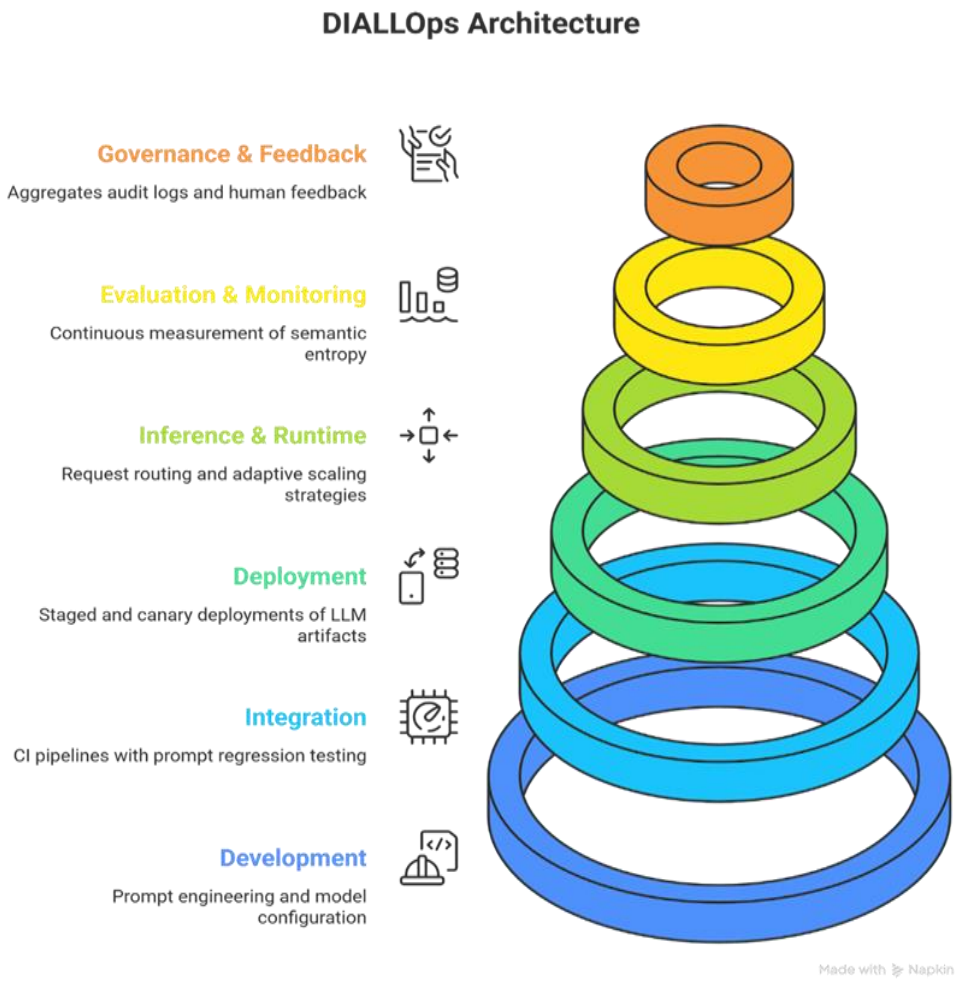

Fig. 5 Layered DIALLOps architecture illustrating the operational stack from development and integration to monitoring and governance.

D. Lifecycle Model in DIALLOps

DIALLOps defines a cyclical lifecycle consisting of:

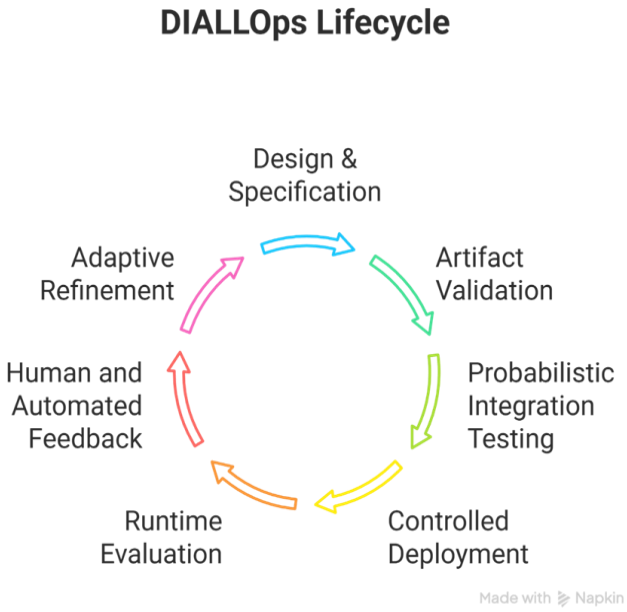

Fig. 6 Closed-loop DIALLOps lifecycle showing continuous design, validation, deployment, evaluation, and adaptive refinement of LLM systems.

This lifecycle explicitly departs from linear DevOps and MLOps pipelines by incorporating continuous feedback loops that reflect the interaction-driven nature of LLM systems ([8,17,38] Joshi, 2025).Figure. 6. DIALLOps lifecycle model showing continuous feedback loops between development, integration, deployment, runtime inference, evaluation, and governance, in contrast to linear DevOps and MLOps pipelines.

E. Distinction from Existing LLMOps Models

DIALLOps differs from existing LLMOps, AgentOps, and RAGOps approaches in three fundamental ways:

- It introduces automation semantics for non-deterministic systems, rather than adapting deterministic pipelines.
- It unifies software, model, and cognitive artifacts under a single operational abstraction.
- It integrates evaluation, governance, and cost control directly into deployment pipelines rather than treating them as external concerns.
- Recent industrial studies show that current LLMOps platforms lack such unified abstractions, leading to fragmented adoption and operational risk ([36] Borovits & Tamburri, 2025).

TABLE I
COMPARISON OF OPERATIONAL PARADIGMS FOR AI SYSTEMS V/S OUR PROPOSED MODEL

Dimension	**DevOps**	**MLOps**	**LLMOps**	**AgentOps**	**DIALLOps (Proposed)**
Primary System Type	Deterministic software	Predictive ML models	Generative LLM systems	Multi-step LLM agents	LLM-centric socio-technical systems
Determinism Assumption	Fully deterministic	Mostly deterministic	Probabilistic	Highly stochastic	Explicitly probabilistic
Core Operational Artifact	Source code	Model + dataset	Model + prompts	Agent workflows	Unified LLM system artifact (model + prompts + RAG + policies)
CI/CD Semantics	Binary pass/fail	Metric-based gating	Limited support	Experimental	Probabilistic behavioral validation
Testing Strategy	Unit & integration tests	Offline benchmarks	Prompt testing	Scenario testing	Distribution-aware, scenario-driven testing
Deployment Control	Version rollback	Model rollback	Partial	Session rollback	Artifact-level rollback (prompt, RAG, agent, policy)
Evaluation Timing	Pre-deployment	Pre/post deployment	Mostly offline	Runtime tracing	Continuous in-production evaluation
Hallucination Handling	Not applicable	Indirect	Detection tools	Agent monitoring	Integrated evaluation + feedback loops
Observability Scope	Infra & logs	Model metrics	Limited	Agent traces	Behavioral, semantic, and interaction-level observability
Governance Integration	External	Partial	Ad-hoc	Minimal	Native policy-driven governance
Cost Awareness	Infrastructure-centric	Training-centric	Token-level	Execution-level	End-to-end cost–quality trade-off control
Automation Maturity	High	Medium	Emerging	Experimental	High, automation-first by design

Key takeaway:
DIALLOps is the only paradigm that treats LLM behavior, evaluation, governance, and DevOps automation as a single operational system rather than loosely coupled components.

IV. ARCHITECTURE AND COMPONENT DESIGN OF THE PROPOSED LLMOPS FRAMEWORK

This section presents the architectural design of the proposed DevOps-integrated LLMOps framework. The architecture is designed to operationalize the conceptual model introduced in Section 3 by providing a modular, automation-driven structure that supports continuous deployment, evaluation, governance, and cost-aware operation of LLM-based systems.

A. Architectural Overview

The proposed architecture follows a layered, control-plane–driven design that separates concerns between development workflows, runtime execution, evaluation, and governance. Recent architectural studies emphasize that such separation is essential for managing the complexity and operational risk associated with large-scale LLM deployments ([19,28]Aryan, 2025; [32] Devarapalli et al., 2025).

At a high level, the architecture consists of:

- A Development and Integration Layer
- A Runtime Execution Layer
- An Evaluation and Observability Layer
- A Governance and Control Plane

This layered design enables independent evolution of components while maintaining system-level consistency through centralized automation and policy enforcement.

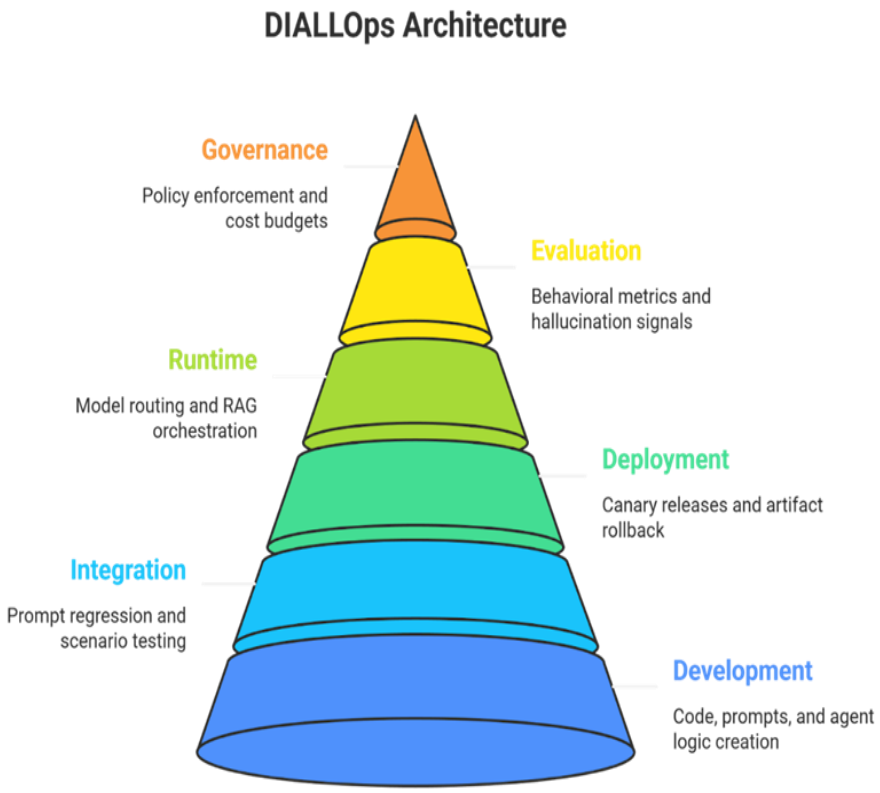

Fig. 7 High-level architecture of the DIALLOps framework, showing the layered organization of development, integration, deployment, runtime execution, evaluation, and governance components coordinated through a centralized control plane

B. Development and Integration Layer

The Development and Integration Layer supports continuous experimentation and system evolution. It manages versioned artifacts including application code, prompt templates, retrieval configurations, and agent orchestration logic.
Unlike traditional CI pipelines that focus on deterministic builds, this layer incorporates behavior-aware integration checks, such as prompt regression testing and retrieval coverage validation. Similar architectural adaptations have been recommended

in recent cloud-native LLMOps frameworks to address non-deterministic outputs ([20,30] Nicora, 2025).

C. Runtime Execution Layer

The Runtime Execution Layer is responsible for serving LLM-based applications in production. It orchestrates inference workflows, retrieval-augmented generation pipelines, and agent interactions under dynamic workload conditions.
This layer supports:

- Dynamic model routing and inference optimization
- Context construction and retrieval orchestration
- Latency-aware and cost-aware request handling

Prior work demonstrates that runtime architectural decisions—particularly around inference routing and caching—have a significant impact on both latency and operational cost in LLM systems ([29] Daga et al., 2025).

D. Evaluation and Observability Layer

The Evaluation and Observability Layer enables continuous monitoring of both system-level and behavioral metrics. In addition to conventional observability signals such as throughput and latency, this layer captures interaction-level traces, prompt execution paths, and retrieval usage.
Recent research highlights that traditional monitoring is insufficient for LLM-based systems and must be augmented with behavioral observability to support debugging, auditing, and trust assessment ([31] Polyakovska, 2025).

E. Governance and Control Plane

At the core of the proposed architecture is a centralized LLMOps control plane responsible for coordinating automation, policy enforcement, and lifecycle governance across all system layers. From a cybersecurity standpoint, this control plane functions as a unified security layer, enforcing access control, auditability, compliance policies, and cost constraints throughout the LLM lifecycle.

Specifically, the control plane governs deployment approvals and rollback decisions, enforces access control over models and connected data sources, maintains compliance and audit logging, and manages cost budgets and resource quotas. These controls are embedded directly into automated workflows, enabling continuous enforcement of organizational policies and regulatory constraints without introducing manual bottlenecks.

By integrating governance mechanisms into the operational fabric of the system, the proposed design ensures that LLM-based applications remain secure, auditable, and cost-aware throughout their lifecycle while preserving DevOps agility.
To demonstrate the practical feasibility of the proposed architecture, a working prototype implementation was developed and experimentally evaluated. The implementation details and evaluation workflow are presented in Section 4.6.

F. Architectural Novelty and Prototype Implementation of DIALLOps

To evaluate the practical applicability of the proposed DIALLOps framework, a prototype implementation was developed that integrates large language model execution, probabilistic evaluation, consistency analysis, governance validation, and automated deployment decision-making. The prototype demonstrates how DevOps principles can be adapted to support the operational requirements of probabilistic AI systems while maintaining reliability, policy compliance, and deployment governance.

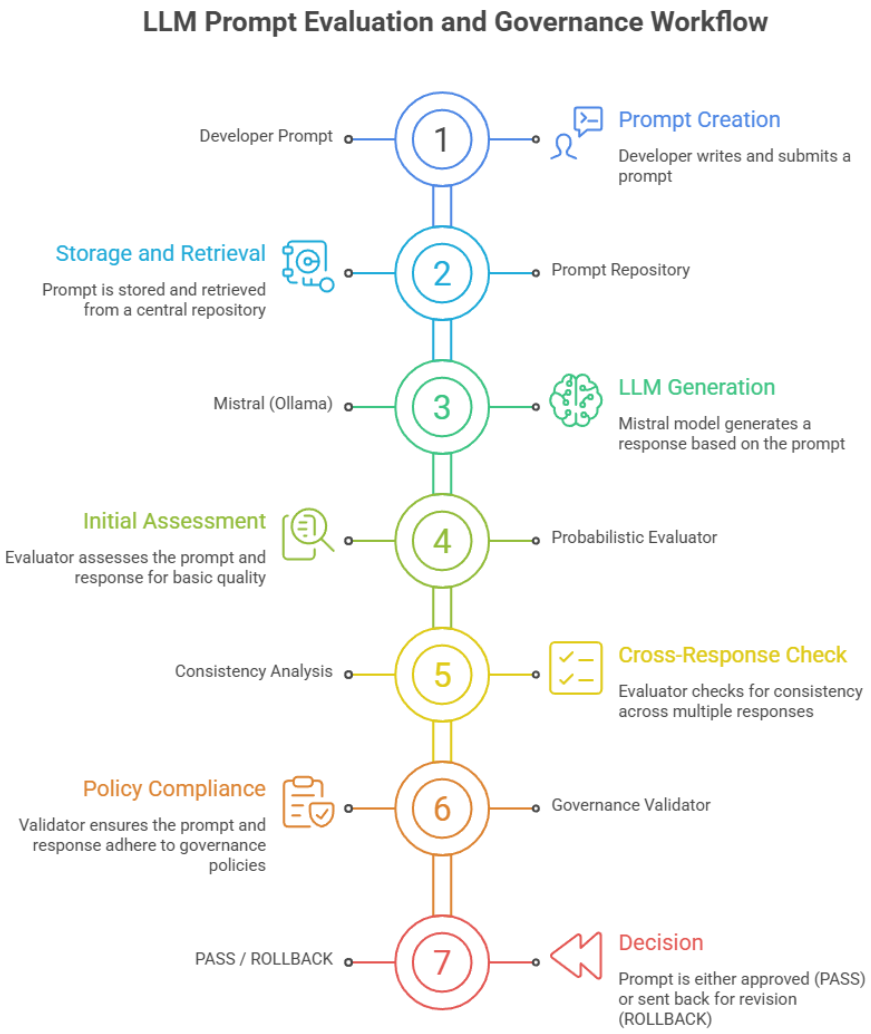

Fig. 8. Prototype implementation of DIALLOps using Ollama, Mistral, probabilistic testing, and governance-aware deployment control.

The prototype was implemented using Python as the orchestration layer and Ollama as the local inference platform. The Mistral large language model was selected as the inference engine because of its open-source availability, efficient execution capabilities, and suitability for enterprise-level experimentation. A centralized prompt repository was also implemented to maintain version-controlled prompt templates, enabling repeatable testing and traceability of prompt modifications throughout the deployment lifecycle.

TABLE III
TECHNOLOGY STACK USED IN THE DIALLOPS PROTOTYPE IMPLEMENTATION

Component	Technology
Programming Language	Python 3.11
LLM Runtime	Ollama
LLM Model	Mistral
Evaluation Method	Probabilistic Testing

Similarity Metric	Cosine Similarity
Embedding Model	all-MiniLM-L6-v2
Governance Engine	Rule-Based Validator

The operational workflow begins when a developer submits a prompt or application update to the DIALLOps pipeline. The submitted prompt is retrieved from the prompt repository and executed against the Mistral model through the Ollama runtime environment. Unlike traditional software systems that produce deterministic outputs, LLMs may generate different responses for the same input because of their probabilistic nature. To address this behavior, a probabilistic evaluation stage was integrated into the framework, enabling multiple inference runs for each test case.

The probabilistic evaluator generates multiple outputs for identical prompts and records behavioral variations across executions. These outputs are then processed by the consistency analysis module, which calculates semantic similarity scores between generated responses. The purpose of this stage is to assess whether the model demonstrates stable and reliable behavior during repeated execution. Higher consistency scores indicate more predictable behavior, while lower scores may suggest instability, prompt sensitivity, or a greater risk of hallucinations.

After the consistency assessment, the generated outputs are passed to the governance validation layer. This component checks compliance with predefined organizational policies, security requirements, content restrictions, and operational guidelines. Governance rules are represented as machine-readable validation policies that automatically evaluate generated responses for prohibited content, policy violations, or unsafe behavior. In this way, the governance validator acts as an automated quality gate within the deployment pipeline.

The final deployment decision is made by combining probabilistic consistency metrics with governance validation results. If the generated outputs meet the required consistency threshold and pass all governance checks, the deployment status is marked as PASS, allowing the updated prompt, model configuration, or application component to proceed through the deployment pipeline. On the other hand, if policy violations are detected or consistency scores fall below acceptable thresholds, the deployment is automatically marked as ROLLBACK, preventing unreliable or non-compliant model behavior from reaching production environments.

To evaluate the effectiveness of the prototype, a dataset containing 50 representative test cases was executed across multiple inference iterations. The test cases included policy-sensitive prompts, factual reasoning tasks, and general conversational queries intended to assess behavioral reliability and governance compliance. Experimental results showed a policy compliance rate of approximately 90%, while consistency analysis produced semantic similarity scores ranging from 0.7267 to 0.9610 across different prompt categories. In addition, governance-aware deployment controls successfully identified non-compliant outputs and triggered automated rollback actions whenever validation requirements were not met.

The prototype implementation demonstrates that DIALLOps can effectively operationalize governance, evaluation, and deployment automation for large language model systems within a unified DevOps-oriented framework. The results further confirm the feasibility of integrating probabilistic testing and policy-aware deployment controls into LLM operational pipelines, improving reliability, transparency, and trustworthiness in production environments.

V. EVALUATION METHODOLOGY

This section presents the experimental evaluation of the DIALLOps prototype implementation. The purpose of the evaluation is to examine the framework's ability to enforce governance policies, measure behavioral consistency in probabilistic systems, and support automated deployment decisions through governance-aware rollback mechanisms. Unlike traditional DevOps evaluation methods that mainly focus on deterministic software behavior, the proposed evaluation framework

emphasizes reliability, consistency, and policy compliance within large language model operations.

A. Experimental Setup

To evaluate the practical applicability of DIALLOps, experiments were conducted using the prototype implementation described in Section 4.6. The evaluation environment consisted of the Mistral large language model running through the Ollama runtime. A benchmark dataset containing 50 prompts, including customer-support and out-of-scope scenarios, was used to examine framework behavior under repeated execution conditions.

TABLE IV

EXPERIMENTAL CONFIGURATION

Parameter	Value
Model	Mistral
Runtime	Ollama
Test Dataset	50 prompts
Runs per Probabilistic Test	20
Consistency Threshold	0.85
Governance Threshold	90%

Each prompt was executed twenty times to capture response variability. Semantic consistency was measured using cosine similarity over sentence embeddings generated from repeated outputs. Governance compliance was assessed through the rule-based validation engine integrated into the DIALLOps control plane.

B. Policy Compliance Evaluation

The DIALLOps prototype was evaluated using a benchmark dataset containing 50 customer-support and out-of-scope prompts. The framework achieved a policy compliance rate of 90%, demonstrating its effectiveness in enforcing operational constraints and governance policies.

TABLE V

POLICY COMPLIANCE RESULTS

Metric	Value
Total Test Cases	50
Passed	45
Failed	5
Compliance Rate	90%

C. Probabilistic Consistency Evaluation

One of the key contributions of DIALLOps is the introduction of probabilistic evaluation mechanisms for deployment decision-making. Unlike deterministic software testing approaches, DIALLOps evaluates distributions of responses rather than relying on single execution outputs. Each prompt was executed 20 times, and semantic consistency was measured using cosine similarity over sentence embeddings.

TABLE VI

PROBABILISTIC CONSISTENCY RESULTS

Scenario	Runs	Consistency Score
Refund Policy	20	0.9165
Damaged Product Return	20	0.9610
Medical Diagnosis	20	0.7267

The results demonstrate clear differences across operational scenarios. Customer-support prompts, such as refund policy and damaged-product return queries, achieved high consistency scores above the deployment threshold of 0.85, indicating stable and predictable behavior. In contrast, the medical diagnosis scenario produced a significantly lower consistency score of 0.7267, suggesting higher behavioral variability and an increased risk of unreliable responses. These findings highlight the importance of probabilistic testing in identifying deployment risks that may not be detected through conventional single-run evaluations.

D. Governance Validation and Automated Rollback

Governance validation and deployment gating are critical components of the DIALLOps architecture. Deployment decisions are determined by combining probabilistic consistency measurements with governance compliance scores.

TABLE VII

GOVERNANCE VALIDATION AND DEPLOYMENT DECISIONS

Scenario	Governance Compliance	Decision
Refund Policy	100%	PASS
Damaged Product Return	100%	PASS
Medical Diagnosis	85%	ROLLBACK

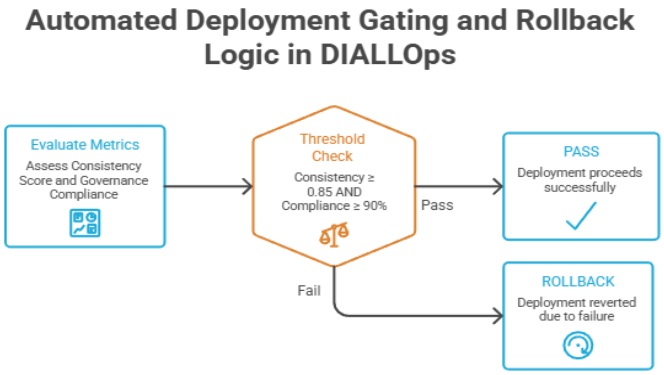

Fig 9. Automated deployment gating and rollback logic in DIALLOps

E. Comparative Baselines

To provide context for the evaluation results, the proposed framework was compared with representative DevOps-, MLOps-, and conventional LLMOps-based deployment approaches. Figure 10 illustrates the baseline deployment strategies considered during the evaluation.

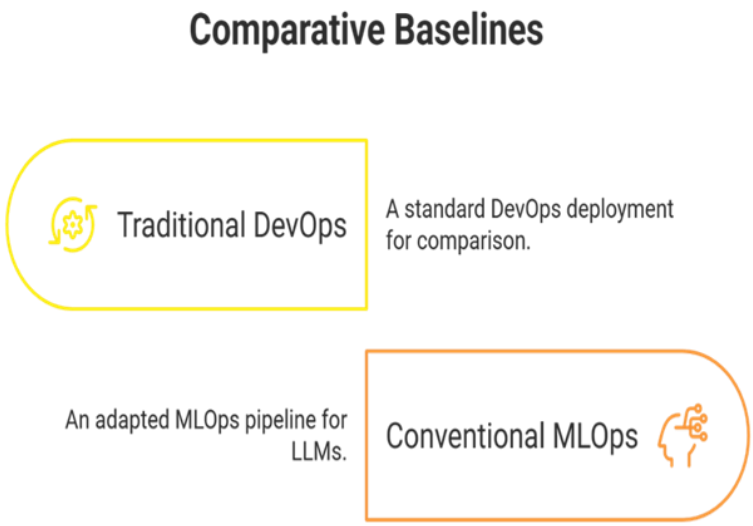

Fig. 10. Baseline Deployment Approaches Used for Comparative Evaluation.

The comparative analysis demonstrates the advantages of integrating probabilistic evaluation, governance validation, and automated rollback mechanisms within a unified operational control plane. Unlike traditional deployment pipelines that primarily focus on deterministic validation, DIALLOps incorporates reliability-aware and governance-aware deployment controls specifically designed for probabilistic AI systems. These capabilities improve operational safety, transparency, and deployment confidence for enterprise-scale LLM applications.

VI. DISCUSSION AND THREATS TO VALIDITY

A. Discussion

The proposed DIALLOps framework addresses several critical limitations observed in current DevOps, MLOps, and emerging LLMOps practices. By redefining operational automation around probabilistic outputs and interaction-driven behavior, DIALLOps enables organizations to deploy LLM-based systems with greater reliability, traceability, and governance.

Importantly, these capabilities position DIALLOps not only as an operational framework but also as a security-enabling architecture for large language model–based systems. By integrating continuous evaluation, policy enforcement, and human oversight into automated workflows, DIALLOps supports proactive risk mitigation in production deployments.

Experimental results demonstrate that DIALLOps can distinguish between reliable and unreliable deployment candidates. In-scope customer-support prompts achieved consistency scores above 0.91 and full governance compliance, resulting in successful deployment decisions. Conversely, out-of-scope medical prompts achieved only 85% governance compliance and

were automatically rejected through the rollback mechanism.

B. Threats to Internal Validity

Internal validity threats arise from the difficulty of isolating causal effects in probabilistic systems. Improvements in output quality or reliability may result from model updates, prompt changes, or user behavior rather than the operational framework itself. To mitigate this, DIALLOps emphasizes controlled experiments and staged rollouts.

C. Threats to External Validity

The generalizability of the proposed framework may be limited by the specific classes of LLMs, workloads, or domains considered. Performance and governance requirements vary significantly across applications such as healthcare, finance, and customer support, potentially affecting adoption outcomes ([9,16,39] Özer, 2025).

D. Threats to Construct Validity

Construct validity is challenged by the lack of universally accepted metrics for evaluating LLM output quality. While DIALLOps supports multiple evaluation signals, these metrics may not fully capture subjective dimensions such as usefulness or trustworthiness ([35] Abbassi, 2025).

E. Threats to Conclusion Validity

Conclusion validity may be affected by evolving LLM architectures and tooling ecosystems. As foundation models and deployment infrastructures continue to change rapidly, some design assumptions may require adaptation. However, the conceptual principles of DIALLOps are intended to remain model-agnostic and future-proof.

VII. CONCLUSION

A. Summary of Findings

This paper addressed the growing gap between the operational characteristics of large language model (LLM)–based systems and the assumptions embedded in traditional DevOps and MLOps practices. While recent research has acknowledged the emergence of LLMOps as a distinct operational paradigm, existing approaches remain fragmented, tool-centric, and insufficiently automated to support the non-deterministic, interaction-driven, and socio-technical nature of generative AI systems.

To address this gap, the paper proposed a unified and automation-driven LLMOps framework that systematically integrates DevOps principles with generative AI–specific requirements. The framework extends conventional operational lifecycles by incorporating prompt and context management, continuous output evaluation, hallucination-aware monitoring, human-in-the-loop oversight, and governance-aware controls. By treating LLM-based systems as composite socio-technical systems rather than standalone models, the proposed approach moves beyond incremental extensions of MLOps and provides a coherent foundation for reliable production deployment.

The proposed framework demonstrates how core DevOps mechanisms—such as continuous integration, continuous deployment, observability, and rollback—can be reinterpreted to accommodate probabilistic outputs, evolving external knowledge sources, and dynamic user interactions. In doing so, it offers a practical pathway for organizations seeking to operationalize LLMs at scale without sacrificing reliability, security, or cost efficiency.

B. Contributions to Research and Practice

From a research perspective, this work contributes to the emerging LLMOps literature by shifting the focus from isolated tools and techniques toward a system-level operational model. Unlike prior studies that primarily characterize challenges or survey existing practices, this paper presents an end-to-end framework that explicitly bridges DevOps automation and generative AI workflows across the full lifecycle of LLM-based systems.

From a practical standpoint, the proposed LLMOps framework is designed to align with established DevOps practices commonly adopted in industry. By positioning LLMOps as a bridging layer rather than a disruptive replacement for existing pipelines, the framework lowers adoption barriers and supports incremental integration into enterprise environments. Its emphasis on continuous evaluation, governance, and cost-aware operations directly addresses concerns that frequently impede real-world deployment of LLM-based applications, including hallucinations, scalability constraints, and operational overhead.

C. Limitations

Despite its contributions, this work has several limitations. First, although a working prototype implementation of DIALLOps was developed and experimentally evaluated, large-scale industrial validation across diverse organizational environments remains future work. The prototype evaluation was conducted using a limited set of representative test cases and deployment scenarios, which may not fully capture the complexity and scale of enterprise-grade LLM deployments.

Second, the evaluation primarily focuses on operational characteristics such as reliability, consistency, governance compliance, scalability, and cost efficiency. While these metrics are important for operational decision-making, they do not fully capture long-term organizational and human factors, including user trust evolution, workforce adaptation, developer productivity, and broader socio-technical impacts of generative AI adoption.

Third, the experimental implementation was evaluated using a specific technology stack consisting of Python, Ollama, and the Mistral large language model. Although the architectural principles of DIALLOps are intended to be model-agnostic, additional studies are required to validate the framework across different foundation models, deployment infrastructures, and domain-specific applications.

Finally, the framework assumes the availability of sufficient computational resources, governance processes, and organizational maturity in DevOps practices. Organizations operating in resource-constrained environments may face challenges in adopting advanced evaluation, monitoring, and governance mechanisms. Future research should investigate lightweight deployment strategies and broader industrial validation to further assess the scalability and generalizability of the proposed framework.
A prototype implementation validated the practical feasibility of DIALLOps. Experimental results demonstrated 90% policy compliance, consistency scores up to 0.9610 for in-scope tasks, and successful automated rollback of unsafe deployments. These findings provide preliminary empirical support for the framework's reliability and governance-aware operational model.

D. Future Research Directions

This work opens several avenues for future research. First, empirical studies are needed to evaluate the proposed LLMOps framework in large-scale production environments across different application domains, such as healthcare, finance, and software engineering. Such studies could quantify the framework’s impact on hallucination rates, operational costs, and deployment stability over extended periods.

Second, future research could explore tighter integration between LLMOps and emerging paradigms such as AgentOps, particularly in systems involving autonomous multi-agent workflows. Investigating how governance, evaluation, and rollback

mechanisms scale in agentic settings represents an important next step.

Third, advances in automated evaluation metrics and feedback mechanisms offer opportunities to further reduce reliance on human-in-the-loop oversight without compromising reliability. Developing standardized benchmarks and operational quality indicators tailored to LLM-based systems would significantly strengthen the maturity of LLMOps as a discipline.

Finally, regulatory and ethical considerations surrounding generative AI continue to evolve. Future work should examine how LLMOps frameworks can dynamically adapt to changing compliance requirements, auditability standards, and organizational governance policies while preserving system flexibility and performance.

E. Concluding Remarks

In conclusion, this paper argues that the reliable and scalable deployment of large language models cannot be achieved through model-centric improvements alone. Instead, it requires a fundamental rethinking of operational practices that integrates DevOps automation with the unique behavioral and governance challenges of generative AI. The proposed LLMOps framework represents a step toward this goal by offering a unified, automation-centric, and practically grounded model for operationalizing LLM-based systems. As generative AI continues to transition from experimental prototypes to mission-critical infrastructure, such frameworks will be essential for ensuring trust, accountability, and sustainable adoption.

REFERENCES

[1] McCarthy, J. (2004). Basic Questions. Stanford University. http://www-formal.stanford.edu/jmc/

[2] Feuerriegel, S., Hartmann, J., Janiesch, C., et al. (2024). Generative AI. Business & Information Systems Engineering, 66, 111–126. https://doi.org/10.1007/s12599-023-00834-7

[3] Goldman Sachs (2023). The Potentially Large Effects of Artificial Intelligence on Economic Growth. (Industry report, cited in Feuerriegel et al., 2024).

[4] Farquhar, S., Kossen, J., Kuhn, L., et al. (2024). Detecting hallucinations in large language models using semantic entropy.Nature, 630, 625–630. https://doi.org/10.1038/s41586-024-07421-0

[5] Menshawy, A., Nawaz, Z., & Fahmy, M. (2024). Navigating challenges and technical debt in large language models deployment. Proceedings of the 4th Workshop on Machine Learning and Systems (EuroMLSys '24), 192–199. https://doi.org/10.1145/3642970.3655840

[6] Pahune, S., & Akhtar, Z. (2025). Transitioning from MLOps to LLMOps: Navigating the unique challenges of large language models. Information, 16(2), 87. https://www.mdpi.com/2078-2489/16/2/87

[7] Diaz-De-Arcaya, J., & López-De-Armentia, J. (2024). Large language model operations (LLMOps): Definition, challenges, and lifecycle management. IEEE Conference Proceedings. https://ieeexplore.ieee.org/document/10612341

[8] Joshi, S. (2025). LLMOps, AgentOps, and MLOps for Generative AI: A Comprehensive Review. https://philpapers.org/archive/JOSLAA-3.pdf

[9] Özer, F. (2025). Systematic Technical Survey on LLMOps: Lifecycle, Tools, Challenges, and Emerging Practices. https://erepo.uef.fi/bitstreams/4212c75f-1822-4648-aa2c-7c55c55d86c2/download

[10] Radenkovic, B., & Prokhorov, S. (2025). Application of Large Language Models in Software Development: Review of the Current State and Development Perspectives. https://ieeexplore.ieee.org/document/11245717

[11] Farquhar, S., Kossen, J., Kuhn, L., et al. (2024). Detecting hallucinations in large language models using semantic entropy. Nature, 630, 625–630. https://doi.org/10.1038/s41586-024-07421-0

[12] Menshawy, A., Nawaz, Z., & Fahmy, M. (2024). Navigating challenges and technical debt in large language models deployment. EuroMLSys '24, 192–199. https://doi.org/10.1145/3642970.3655840

[13] Diaz-De-Arcaya, J., & López-De-Armentia, J. (2024). Large language model operations (LLMOps): Definition, challenges, and lifecycle management. IEEE Conference Proceedings.

[14] Pahune, S., & Akhtar, Z. (2025). Transitioning from MLOps to LLMOps: Navigating the unique challenges of large language models. Information, 16(2), 87. https://www.mdpi.com/2078-2489/16/2/87

[15] Diaz-De-Arcaya, J., & López-De-Armentia, J. (2024). Large language model operations (LLMOps): Definition, challenges, and lifecycle management. IEEE Conference Proceedings. https://ieeexplore.ieee.org/document/10612341

[16] Özer, F. (2025). Systematic technical survey on LLMOps: Lifecycle, tools, challenges, and emerging practices. University of Eastern Finland. https://erepo.uef.fi/bitstreams/4212c75f-1822-4648-aa2c-7c55c55d86c2/download

[17] Joshi, S. (2025). LLMOps, AgentOps, and MLOps for Generative AI: A Comprehensive Review. PhilPapers. https://philpapers.org/archive/JOSLAA-3.pdf

[18] Menshawy, A., Nawaz, Z., & Fahmy, M. (2024). Navigating challenges and technical debt in large language models deployment. Proceedings of the 4th Workshop on Machine Learning and Systems (EuroMLSys '24), 192–199. https://doi.org/10.1145/3642970.3655840

[19] Aryan, A. (2025). LLMOps: Managing Large Language Models in Production. Springer / Google Books. https://books.google.com/books?id=PrZxEQAAQBAJ

[20] Nicora, G. (2025). MLOps and LLMOps: Development of an LLM-based application with focus on scalable cloud deployment. Politecnico di Torino (Master's Thesis). https://webthesis.biblio.polito.it/36459/

[21] Farquhar, S., Kossen, J., Kuhn, L., et al. (2024). Detecting hallucinations in large language models using semantic entropy. Nature, 630, 625–630. https://doi.org/10.1038/s41586-024-07421-0

[22] Shan, R., & Shan, T. (2024). Enterprise LLMOps: Advancing large language models operations practice. IEEE Cloud Summit. https://ieeexplore.ieee.org/document/10630923

[23] Radenkovic, B., & Prokhorov, S. (2025). Application of large language models in software development: Review of the current state and development perspectives. IEEE Conference Proceedings. https://ieeexplore.ieee.org/document/11245717

[24] Fitsilis, P., Damasiotis, V., &Kyriatzis, V. (2024). DOLLmC: DevOps for large language model customization. arXiv preprint arXiv:2405.11581. https://arxiv.org/abs/2405.11581

[25] Tantithamthavorn, C. K., Palomba, F., Khomh, F., & Chua, J. J. (2024). MLOps, LLMOps, FMOps, and beyond. IEEE Software. https://www.computer.org/csdl/magazine/so/2024/01

[26] Dong, L., Lu, Q., & Zhu, L. (2024). AgentOps: Enabling Observability of LLM Agents. arXiv:2411.05285. https://arxiv.org/abs/2411.05285

[27] Xu, X., Weytjens, H., Zhang, D., Lu, Q., & Weber, I. (2025). RAGOps: Operating and Managing Retrieval-Augmented Generation Pipelines. arXiv:2506.03401. https://arxiv.org/abs/2506.03401

[28] Aryan, A. (2025). LLMOps: Managing Large Language Models in Production. Springer. https://books.google.com/books?id=PrZxEQAAQBAJ

[29] Daga, V., Daga, P., & Dhabhai, A. (2025). Energy and cost optimization strategies for large language models in LLMOps. International Journal of Recent Research and Review. https://ijrrr.com/specialissues2-2025/ijrrr-Special-Issue-2-2025-paper42.pdf

[30] Nicora, G. (2025). MLOps and LLMOps: Development of an LLM-based application with focus on scalable cloud deployment. Politecnico di Torino. https://webthesis.biblio.polito.it/36459/

[31] Polyakovska, N. (2025). Understanding and building operational excellence for large language models. Computer-Integrated Technologies. https://cit.lntu.edu.ua/index.php/cit/article/view/691

[32] Devarapalli, S., Vathsavai, V. G., & Katkam, V. (2025). Cloud-native LLMOps meets DataOps: A unified framework for high-volume analytical systems. IEEE Conference Proceedings. https://ieeexplore.ieee.org/document/11158069

[33] Gadiraju, P., Kosna, S. R., & Shah, K. (2025). DataOps meets LLMOps: Automating cloud-based AI workflows from data ingestion to prompt optimization. IEEE Conference Proceedings. https://ieeexplore.ieee.org/document/11135202

[34] Doan, R. (2024). Essential guide to LLMOps. Packt Publishing.

[35] Abbassi, A. A. (2025). Towards Reliable and Trustworthy Pipelines for MLOps and LLMOps. PhD Dissertation, Polytechnique Montréal.

[36] Borovits, N., & Tamburri, D. A. (2025). On the Maturity of LLMOps Services Computing: An Industrial Study. Service-Oriented Computing, Springer.

[37] Tantithamthavorn, C. K., Palomba, F., Khomh, F., & Chua, J. J. (2024). MLOps, LLMOps, FMOps, and Beyond. IEEE Software.

[38] Joshi, S. (2025). LLMOps, AgentOps, and MLOps for Generative AI: A Comprehensive Review. PhilPapers.

[39] Özer, F. (2025). Systematic Technical Survey on LLMOps: Lifecycle, Tools, Challenges, and Emerging Practices. University of Eastern Finland.

[40] Lekota, N. F. (2024). Governance considerations of adversarial attacks on AI systems. Proceedings of the International Conference on AI Security.

[41] Brundage, M., et al. (2023). Toward trustworthy AI development: Mechanisms for security and governance. arXiv:2302.XXXX.

Security Analysis of Address Resolution Protocol and Its Exploitable Vulnerabilities

Ritik Kalal[1]*,Harsh Mandaliya[2],Dhruv Patel[3],Deep Solanki[4],Jigar Gajjar[5]

Lok Jagruti Kendra University, Ahmedabad GJIndia[1, 2, 3]
TechDefenceCybersecurity Limited, Ahmedabad GJ India[4,5]

ritikkalal1121@gmail.com[1], harshmandaliya4114@gmail.com[2], pateldhruv190725@gmail.com[3], deepsolanki5799@gmail.com[4], cyberian.jg@gmail.com[5]

Abstract: "Address Resolution Protocol" (ARP) is a simple network protocol that enables devices on a network to discover each other based on their assigned IP addresses and their actual physical addresses referred to as MAC addresses. ARP is a fundamental protocol under IPv4 that has laid a solid foundation however, it was never designed with security considerations. Consequent to its lack of security, there is a chance of various security attacks such as ARP spoofing, man-in-the-middle attack, and denial of service attacks. These security attacks may result in data theft, data modification, and denial of services. In the paper, various issues regarding ARP, its security flaws, research on various kinds of attacks associated with ARP, and prevention are discussed. For instance, various security protocols such as IPv6 Neighbor Discovery, Software Defined Network, and many more have become alternatives to ARP.

Keywords:Address Resolution Protocol (ARP), ARP Cache Poisoning, Network Security Vulnerabilities, Man-in-the-Middle Attacks, Denial-of-Service Attacks, ARP Defense Mechanisms, IPv6 Neighbor Discovery, Software-Defined Networking (SDN).

I. INTRODUCTION

A. ARP and its Need

Modern networks separate physical and logical addressing (MAC and IP respectively), enabling wide-area routing and local link delivery. The Address Resolution Protocol (ARP) enables the dynamic mapping of IP addresses to their corresponding MAC addresses thus, ensuring delivery of data frames to the correct hardware destination on networks such as Ethernet (RFC 826) [9]. A device wishing to communicate with another device locally will send a broadcast ARP request to the network, asking "Who has this IP address?" The device that possesses the matching IP address will send a unicast reply containing its MAC address back to the requester, who will cache this information for future use, in an efficient manner [2].

The automation of this process eliminates the need for manual configuration, works well in dynamic environments, and provides an easy way to plug-and-play in LAN environments. ARP is necessary, as all higher-layer protocols will assume IP-level abstraction (logical addressing), but will send physical delivery to the specific MAC address of the hardware being communicated with without ARP, devices would have no means to determine where to send data locally, nor would they be able to do so reliably, if static mappings of IP address to MAC address were used [5] **.** ARP uses a broadcast method of operation, as this was the simplest and easiest to implement in early trusted networks [6] .

ARP was designed for trusted local networks, so it prioritizes simplicity over authentication. Because the protocol accepts address mappings with little verification, an attacker on the same segment can exploit that trust by sending forged replies and redirect traffic to an unintended host. This design choice was acceptable in early LAN environments, but it becomes a serious

http://doi.org/10.63766/spujstmr.26.000095

limitation in modern networks where mobility, virtualization, and untrusted endpoints are common [5], [6], [9].

B. Different attacks possible on ARP and its impact

The ARP "trust model" that allows for unsolicited ARP replies to be accepted without verification, enables "cache poisoning", where an attacker can forge an ARP reply associating their MAC address with a legitimate IP address such as a default gateway [2]. Cache poisoning is part of a larger "man in the middle" (MITM) attack mechanism that facilitates methods to eavesdrop, hijack a session or modify traffic [13]. Gratuitous ARP floods can further exacerbate the impact of cache poisoning by allowing multiple caches to be poisoned quickly [11]. Attackers will take advantage of similar methods in SDN "contexts" and can disrupt a controller at the same time as expanding their reach [7]. The DoS variations can thin valid IP traffic to non-existent MACs (resulting in "blackholing" packets) or generate broadcast storms, overwhelming a switch and CPU [10].

The impacts to the CIA triad are substantial: confidentiality is diminished by means of obtaining a victim's credentials or stealing sensitive data traversing the wire integrity is compromised through attacking or modifying an arbitrary packet (e.g., adding malicious code) and availability is eliminated through DoS, resulting in service outages to critical systems such as medical or financial networks [10]. There are also risks of escalation such as entering a network with ransom ware or other broader breaches, which demonstrate ARP to be a persistent point of entry into legacy infrastructures [11].

TABLE I

ATTACKS POSSIBLE ON ARP

Attack Name	What the Attacker Does	Main Effect / What Happens to the Network
ARP Spoofing / Poisoning	Sends fake ARP messages saying "My MAC is for this IP"	Victim's device sends traffic to the attacker instead
Man-in-the-Middle (MITM)	Uses spoofing to sit between two devices	Attacker can spy on, steal, or change messages secretly
Denial-of-Service (DoS)	Sends fake mappings or floods with bad ARP messages	Network slows down, packets get lost, devices can't connect
Session Hijacking	Grabs login cookies/sessions during MITM	Attacker takes over your logged-in account
Gratuitous ARP Flood	Sends lots of fake "I'm this IP" messages to everyone	Many devices get wrong addresses → network confusion/DoS

C. Prevention and mitigation strategies to secure ARP

Mitigating ARP's lack of verification requires additional layers of verification. Static ARP entries enforce manual locks on the IP-MAC bindings of users, which can help to mitigate against spoofed ARP entries, however, they are not suited for dynamic environments [6]. Dynamic ARP Inspection (DAI) on switches will inspect packets against DHCP-snooped ARP entries to ensure packets are valid and drop the invalid ones [12]. In SDN, centralized policy layers can leverage adaptive algorithms, such as game theory, to anticipate the actions of an attacker and optimize verification [7].

Cryptographic mechanisms, such as S-ARP, use signatures to allow the server to verify that it is responding to the correct client when sending a response [4]. Other protocols, such as ticket-based protocols (e.g., TARP), require that the server validate the token before responding [6]. Network segmentation through the use of VLANs can shrink the size of broadcast domains that can carry attacks [4]. Intrusion Detection Systems (IDS) can monitor for ARP/spoofed attacks (e.g., MAC

flapping) and allow for rapid response through Anomaly Detection [10]. Combining host-level filtering (e.g., ignore unsolicited ARP Requests) with the aforementioned entries provides a multi-layered defense and this approach towards a verified (zero-trust) local resolution will provide usable access [12].

A solid strategy for defense, therefore, has to incorporate not just protection against ARP spoofing, but also feasibility. Static ARP is a good choice, but it needs regular maintenance, which can be very demanding in a dynamic network. Dynamic ARP Inspection further secures enterprise switches when properly implemented along with DHCP snooping and sufficient hardware. However, the cryptographic approach of S-ARP enhances security but involves higher overheads. Conversely, SDN-based approaches offer centralized management at the cost of increased complexity[4], [6], [7], [12] .

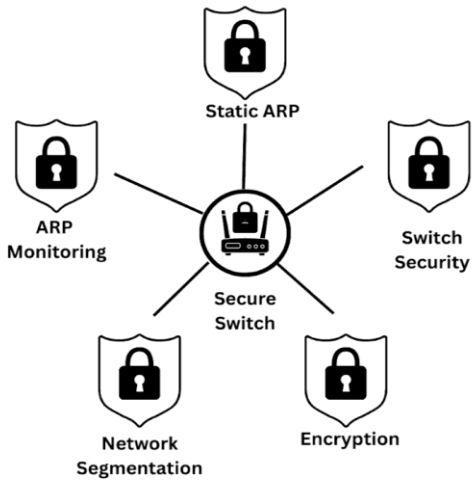

Fig.1 Prevention and mitigation strategies to secure ARP

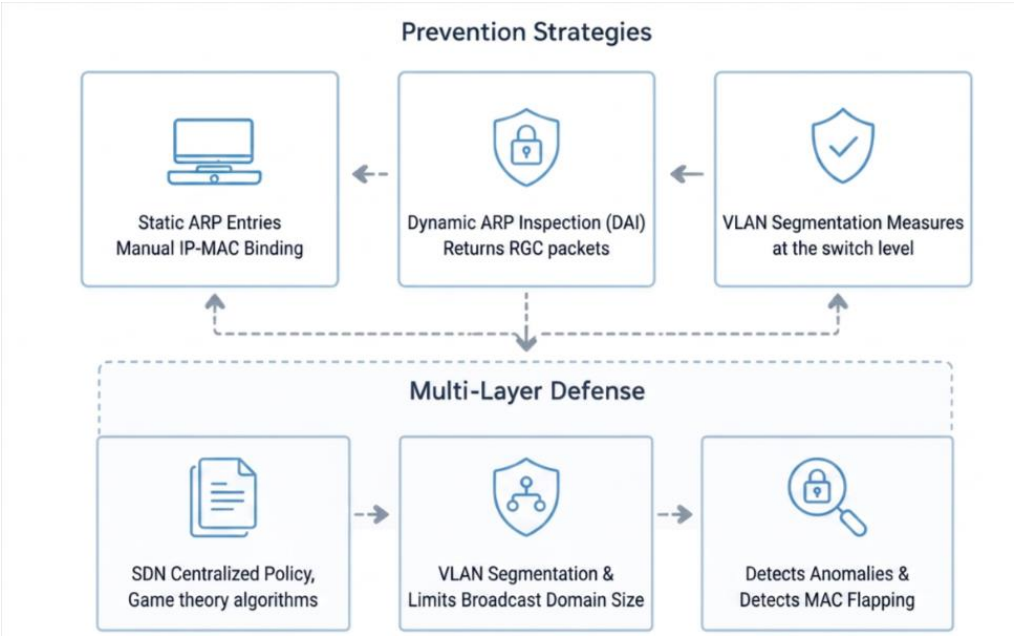

Fig.2 Prevention and Mitigation Strategies to Secure ARP Protocol

D. Limitations of ARP protocol and its alternatives

ARP uses broadcast flooding to waste bandwidth and create storm conditions in large environments or data centers [5]. ARP has no security built-in and is easily spoofed without additional external measures, also limited to IPv4 networks [12]. The scalability of ARP breaks down when you add mobility and virtualization.

The Neighbor Discovery Protocol (NDP) in IPv6 replaces ARP with more efficient multicast and offers Secure Neighbor Discovery (SEND) to provide cryptographic validation. SDN solutions use proxy ARP via a controller that queries a centralized database to avoid broadcast flooding while improving security [12]. All of these alternatives are consistent with the modern-day's need for an authenticated and scalable method of address resolution.

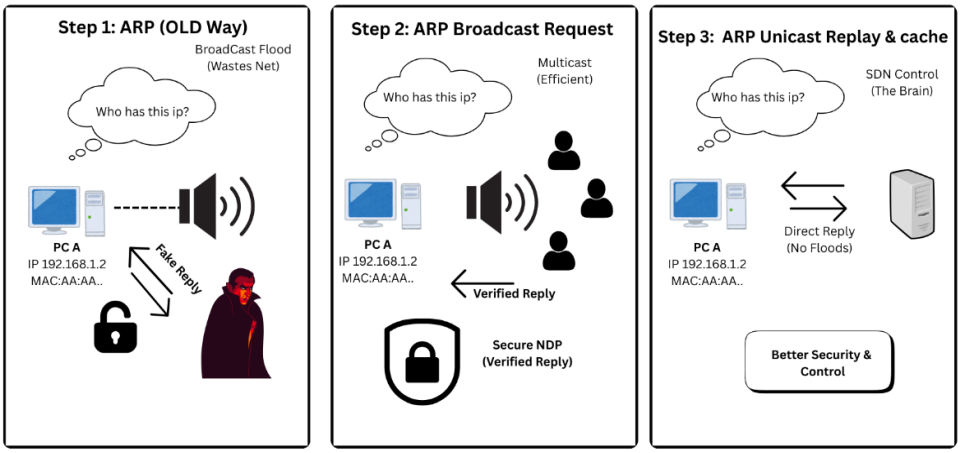

Fig.3 ARP Limitations vs Modern Solutions

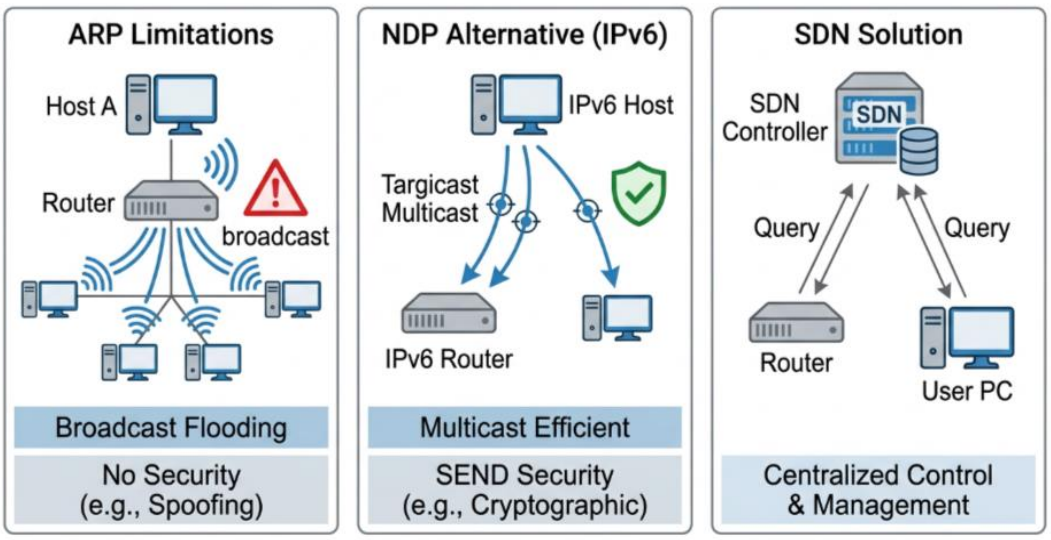

Fig.4 Comparison of Network Address Resolution and Security Approaches

II. LITERATURE SURVEY

ARP has been around since 1982 it's an early TCP/IP protocol designed to provide mapping between logical (IP) and physical (MAC) addresses on a shared local area network (like Ethernet) [9]. David C Plummer authored RFC 826 defining it as a simple request/reply mechanism using broadcasts, which allowed ARP to dynamically resolve address pairs without having to create

manual mappings. By the end of the 1980's and into the 1990's, the number of LANs were growing rapidly, and ARP was able to accommodate this growth [15]. Over the years, ARP has only received a few minor updates with some RFCs adding enhancements to how ARP detects address conflicts. However, the functionality of ARP itself has remained the same: broadcast-based and unauthenticated. While these characteristics of ARP work well in small trusted networks, the evolution of new types of threats has revealed the weaknesses of ARP and caused researchers to begin looking for ways to:

1. Securely extend ARP
2. Replace ARP with an IPv6 (Internet Protocol Version 6) alternative.

A. Security Challenges

1) ARP Spoofing / Cache Poisoning

On a local network, an attacker can leverage a vulnerability in ARP to poison the ARP cache of the victim's system by broadcasting fake ARP replies that map the attacker's MAC address to the IP address of a legitimate device, such as the default gateway. Because ARP is a stateless protocol, the victim's device blindly accepts these updates, causing their traffic to be redirected to the attacker[11].

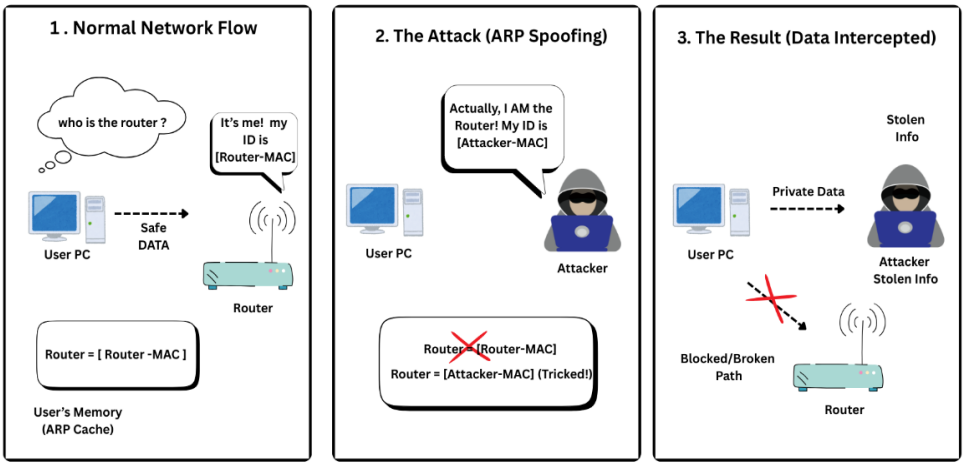

Fig.5 How ARP Spoofing Works

2) Man-in-the-Middle

The attacker (MITM) places himself in between two endpoints by taking over the source and destination's ARP caches. By doing this, they will intercept, inspect, or alter the contents of packets sent between those two endpoints (e.g., stealing credentials or injecting malware) while still passing packets (traffic) between the endpoints as if they had not been altered or modified [7].

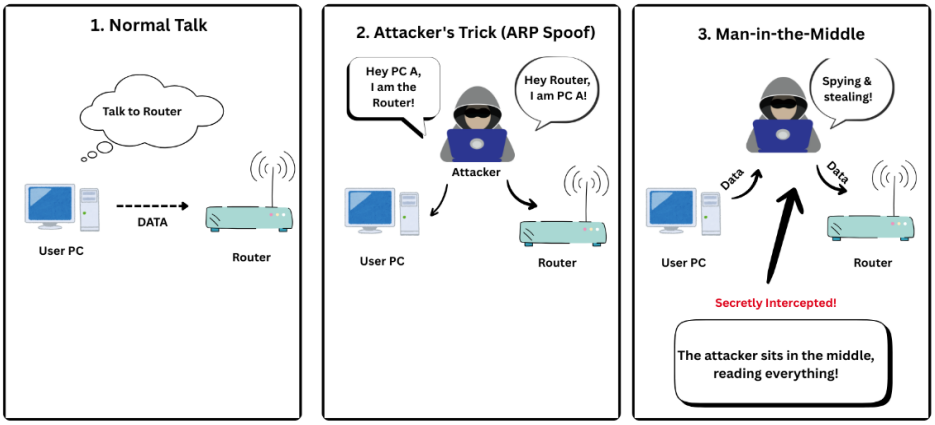

Fig. 6 Man-in-the-Middle Attack

3) Denial-of-Service (DoS) Attack

Attackers may produce a high number of ARP request/reply pairs or link valid IP addresses (that correspond to a real machine) to non-existent MAC addresses with the intent of creating broadcast storms to deplete both bandwidth resources and CPU resources, or to create packet blackholing which cause a loss of connectivity [12].

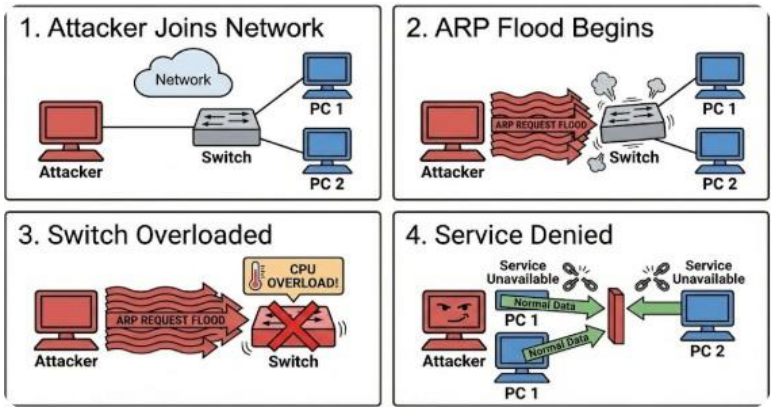

Fig.7. Denial-of-Service (DoS) Attack

B. Limitations of ARP

1) Broadcast Flooding

Broadcasts for address resolution queries made by ARP produce excessive broadcasts across networks which do not scale well. For example, large LANs or data centres can have a high rate of broadcasts leading to possible congestion from broadcast storms which could lead to an overall performance hit [14].

2) No Built-in Authentication

The ARP protocol has no method for confirming the authenticity of a reply or the identity of the sender, which allows virtually any device to create fake replies and fill invalid entries in any ARP cache much too easily [4].

3) IPv4-Specific and Poor Scalability

The Address Resolution Protocol (ARP) was developed primarily for Internet Protocol Version 4 (IPv4). ARP has problems in dynamic, virtualized, and/or highly mobile networks because of the broadcast way in which it sends information, how it manages its cache, and how it cannot support the transition to Internet Protocol Version 6 (IPv6) [6].

C. Alternatives to ARP

1) Neighbor Discovery Protocol (NDP) in IPv6

Instead of ARP, NDP uses ICMPv6 messages to resolve addresses in an IPv6 environment using targeted multicast instead of flooding the whole segment with broadcasts. By reducing unnecessary network traffic, this approach improves efficiency for larger networks. SEND, or Secure Neighbor Discovery, extends NDP with cryptography to help prevent spoofing of local address mappings, making them more secure and dynamic to current configurations. [8] [3].

2) Proxy ARP / Centralized Resolution in SDN

With SDN, address resolution is centralized. Instead of letting devices broadcast their address queries everywhere throughout a network, the controller (proxy) retrieves the address mappings from its comprehensive view of the network and responds directly to the querying device with that information. This reduces the amount of broadcast overhead (traffic), enables rapid policy application, and allows for better containment of potential threats in virtual (or dynamic) network infrastructures [12][1].

D. Gap

Though there are papers which have mentioned approaches that could detect or prevent ARP attack, yet only a handful of papers mention the practicality of such approaches in terms of administrative burden, backward compatibility issues, and computational overhead. The reason for this is that sometimes an approach may be effective yet still remain impractical. Consequently, an approach's feasibility should be evaluated both from security efficiency and implementability perspectives[2], [7], [12], [15] .

III. ANALYSIS

A further assessment of the strengths of ARP and the solutions available should focus not just on the security benefits of a possible solution, but also on how realistically feasible the deployment of such solutions may be in practical situations. Some solutions may provide better security measures but at the same time, be more complex, whereas other solutions may be simpler but fail to completely protect networks from spoofing attacks and poisoned information.

TABLE II

ANALYSIS TABLE: ARP, ITS PROBLEMS, AND SOLUTIONS

Feature	ARP (IPv4)	ARP Attacks	Security Solutions	Better Alternatives
What it does	Helps a device find the MAC address of an IP on a local network	Takes advantage of ARP to send fake IP-MAC information	Tries to prevent or reduce fake ARP messages	Uses newer and safer ways instead of ARP
How it works	Sends a broadcast message asking who owns a certain IP	Sends fake ARP replies even when no request is made	Verifies ARP messages before trusting them	Relies on multicast, encryption, or centralized control

Security	Has no built-in security at all	Very easy for attackers to exploit	Improves security using rules or cryptographic methods	Security is designed into the protocol itself
Trust model	Trusts every ARP reply it receives	Attackers misuse this blind trust	Only trusted or verified replies are accepted	Replies are checked before being accepted
Main attacks	—	ARP spoofing, Man-in-the-Middle, DoS	Detects or blocks these attacks	Prevents such attacks by design
Confidentiality	Does not protect data	Data can be intercepted or stolen	IDS and traffic monitoring tools help	Communication is encrypted and authenticated
Integrity	Does not check if data is altered	Packets can be changed by attackers	Static bindings reduce the risk	Strong identity and message verification is used
Availability	Works fine in small networks	Can slow down or even shut down networks	Rate limiting and VLANs improve stability	Less flooding and better network reliability
Network traffic	Generates a lot of broadcast traffic	Attackers can greatly increase traffic	Broadcast traffic is controlled	Very little unnecessary traffic
Scalability	Does not scale well in large networks	Problems get worse as the network grows	Some improvement, but still limited	Designed to work well in large networks
IPv6 support	Not supported	—	—	Fully supported
Examples	ARP cache	ARP poisoning	Dynamic ARP Inspection (DAI), IDS, S-ARP	IPv6 NDP, SEND, SDN, Proxy ARP

Source: Based on RFC 826 and prior studies on ARP security, spoofing attacks, mitigation techniques, and secure alternatives [1], [2], [4], [7], [10], [11].

The comparison table in the following provides a structured analysis of the functionalities of the Address Resolution Protocol, its vulnerabilities in terms of security, mitigation strategies, and modern alternatives. ARP originally was defined in RFC 826 [1], introducing IP-to-MAC address mapping into IPv4 local area networks without any kind of built-in authentication. This design weakness turns networks into potential targets of various malicious activities: poisoning of the ARP cache and MITM attacks [2], [4], and an inability to ensure DoS conditions that implicate network availability [10]. The table further refers to the following common countermeasures that were typically implemented: dynamic ARP inspection, intrusion detection systems, and cryptographic extensions such as Secure ARP (S-ARP) [7]. In turn, it outlines the differences of ARP from more secure and scalable modern alternatives, among which are NDP - Neighbor Discovery Protocol upon shifting to IPv6 [11] - and a centralized resolution upon utilizing SDNs [10]. Overall, the table summarizes some operational limitations of the protocol and underlines the urgent need for enhanced security mechanisms in contemporary environments.

From an analytical point of view, this means that ARP-related security cannot be fully addressed through a single universal defense. A practical and secure approach should combine preventive controls, detection mechanisms, and, where possible, migration toward more secure resolution models that reduce broadcast dependence and strengthen authenticity.

IV. CONCLUSION

ARP has always played an important role as a protocol for communication between devices in local network environments. However, since this protocol does not have any built-in security features, it is considered unsafe for use in modern network environments. The trust-based feature of this protocol allows attackers to easily compromise this protocol for malicious activities, thereby creating significant risks for all aspects of computer security confidentiality, integrity, and availability.

However, several security controls such as Dynamic ARP Inspection, IDS, and cryptography have been developed to address ARP-related threats. Yet such controls cannot thoroughly address the issue since modern network technologies have introduced alternative concepts for secure network communications such as IPv6 Neighbor Discovery protocol along with its enhanced security features and SDN's centralized address resolution techniques.

The research presented is deliberately a survey-based approach rather than an experiment performed on an active testbed. Future research may expand the analysis to include the simulation of ARP spoofing attacks in order to measure how well each method detects the attacks, how much latency overhead and how expensive its implementation is for small networks, enterprises, and software-defined networking.

REFERENCES

[1] Alharbi, T., Aljuhani, A., & Alharbi, S. (2021). Scalable and secure SDN based ethernet architecture by suppress-ing broadcast traffic. Egyptian Informatics Journal, 22(4), 453-462.

[2] Al Sukkar, G., Saifan, R., Khwaldeh, S., Maqableh, M., & Jafar, I. (2016). Address Resolution Protocol (ARP): Spoofing Attack and Proposal for a Defense Mechanism. Com-munications and Network, 8(3), 118-130.

[3] Arkko, J., Kempf, J., Zill, B., & Nikander, P. (2005). Secure Neighbor Discovery (SEND). RFC 3971.

[4] Bruschi, D., Ornaghi, A., & Rosti, E. (2003). S-ARP: an address resolution protocol that is secure. Proceedings of the 19th Annual Computer Security Applications Conference.

[5] Gouda, M.G. & Huang, C.T. (2003). A Secure Address Resolution Protocol. Computer Networks, 41(1), 57-71.

[6] Lootah, W., Enck, W., & McDaniel, P. (2007). TARP: A Ticket-Based Address Resolu-tion Protocol. Computer Networks.

[7] Mvah, F., Tchendji, V.K., Djamegni, C.T., Anwar, A.H., Tosh, D.K., &Kamhoua, C. (2024). Countering ARP Spoofing Attacks in Software-Defined Networks Using a Game-Theoretic Approach. Computers & Security, 139, 103696.

[8] Narten, T., Nordmark, E., Simpson, W., & Soliman, H. (2007). Neighbor Discovery for IP version 6 (IPv6). RFC 4861.

[9] Plummer, D.C. (1982). A Protocol for Resolving Ethernet Addresses to a Network Pro-tocol Address to Transmit on Ethernet Interfaces. RFC 826.

[10] Prabadevi, B. & Jeyanthi, N. (2017). An Analysis of Proposed Solutions to ARP Poi-soning Attacks and Their Impacts on Medical Computing. International Journal of Medical Engineering and Informatics.

[11] Salim, H., Li, Z., Tu, H., & Guo, Z. (2012). Prevention of ARP Spoofing via Gratuitous Decision Packet. 2012 11th International Symposium on Distributed Computing and Appli-cations to Business, Engineering & Science.

[12] Shah, Z. & Cosgrove, S. (2019). A survey on how to address the ARP cache poisoning problem in Software-Defined Networking. IEEE Access / Electronics, 8(10), 1095.

[13] Song, M.S., Lee, J.D., Jeong, Y.S., Jeong, H.Y., & Park, J.H. (2014). A New Detection Scheme for ARP Spoof-ing Attacks Based on Routing Trace for Ubiquitous Environments. The Scientific World Journal.

[14] Tripathi, N. &Mehtre, B. M. (2013). Secondary Cache Approach based on ICMP for Detection and Prevention of ARP Poisoning. Proceedings of the IEEE International Confer-ence on Computational Intelligence and Computing Re-search.

[15] Tripathi, N. &Mehtre, B. M. (2014). A Survey of the Various Techniques for Mitigat-ing ARP Poisoning Attacks. International Workshop on Control, Instrumentation, Commu-nication and Computational Technologies.

AI-Based ANPR SystemforIntelligent Vehicle Speed Detection

Chenna kesava Reddy S[1],Dr. Nityananatham Sampathkumar[2],AkshayKumar Varma S[3], Akhil Chowdary R[4], Ramakanth Reddy V[5]

Student,Computer Science and Engineering,Kalasalingam Academy of Research and Education,Madurai,India[1]
Associate Professor,Computer Science and Engineering,Kalasalingam Academy of Research and Education,Madurai,India[2]
Student,Computer Science and Engineering,Kalasalingam Academy of Research and Education,Madurai,India[3]
Student,Computer Science and Engineering,Kalasalingam Academy of Research and Education,Madurai,India[4]
Student,Computer Science and Engineering,Kalasalingam Academy of Research and Education,Madurai,India[5]

Chennakesava9007@gmail.com[1*],s.nithyanantham@klu.ac.in[2],akshaysagiraju1@gmail.com[3],ravuruakhilchowdary@gmail.com[4], ramakanthram25@gmail.com[5]

Abstract:Vehicleoverspeeding is one of the major causes of road accidents, making automated and reliable speed monitoring systems essential for modern traffic management. This paper presents an AI-based Automatic Number Plate Recognition (ANPR) system for intelligent vehicle speed detection using a dual-camera approach. The proposed system captures vehicle images at two fixed checkpoints and calculates the average speed based on the time difference and known distance between the cameras. Deep learning models are employed to detect vehicles and localize number plates, while an optical character recognition (OCR) pipeline is used to extract license plate information. To improve recognition accuracy under real-world conditions, image preprocessing techniques such as contrast enhancement and adaptive thresholding are applied before OCR. The system supports multiple vehicle categories with predefined speed limits and automatically identifies overspeed violations. A web-based interface is developed to store records, visualize data, and generate downloadable PDF evidence reports for detected violations. Experimental evaluation conducted on real vehicle samples demonstrates improved plate recognition accuracy and reliable speed estimation compared to conventional approaches reported in existing literature. The proposed system offers a practical and scalable solution for intelligent traffic speed monitoring applications.

Keywords: Automatic Number Plate Recognition (ANPR), Vehicle Speed Detection, Average Speed Enforcement, Deep Learning, YOLO, Optical Character Recognition, Intelligent Traffic Monitoring.

I. INTRODUCTION

Road traffic accidents caused by overspeeding are still a critical issue in every part of the world. There is an increasing need for automated and reliable speed monitoring systems. Traditional speed enforcement methods using radar guns and single-point speed cameras are instantaneous in measurement but most often fail to capture driving behavior over distance. In this regard, these methods may not represent actual sustained overspeeding in urban and highway environments.

Recent development in the sectors of artificial intelligence and computer vision has populated the development of intelligent traffic monitoring capable of analyzing vehicle motion with visual information. ANPR has emerged as one of the key technologies in this domain, which uniquely identifies vehicles using license plate information. This allows the average speed of a vehicle across multiple checkpoints to be determined by combining ANPR with time-based measurement, hence producing a more reliable alternative to the conventional spot-speed detection systems.

There are a few existing researches on car detection and speed measurement through a single camera, optical flow technology,

http://doi.org/10.63766/spujstmr.26.000096

or geometric calibration algorithms. These methods are extremely dependent on camera angle and geometry calibration accuracy. The main goal of point-to-point average speed measurement systems is to minimize reliance on geometric calculations by leveraging a certain distance along with a timestamp.

In this work, an AI-based ANPR system for intelligent vehicle speed detection is presented using a dual-camera framework. The proposed system captures vehicle images at two fixed locations, extracts license plate information using deep learning–based detection and OCR techniques, and computes the average speed based on the time difference between detections. To improve number plate recognition accuracy under varying lighting and image quality conditions, image preprocessing techniques are applied prior to OCR. The system also categorizes vehicles based on type and applies corresponding speed limits to identify violations automatically.

A full end-to-end system is also designed with a web-based platform to provide storage for the records of the detection process, system output visualization, and the creation of a downloadable evidence report system. Experiments conducted with real car models indicate that the presented system has the ability to provide foolproof plate recognition and accurate speed measurement capabilities.

II. RELATED WORK

Automatic Number Plate Recognition (ANPR) and computer vision-based speed measurement of vehicles have gained significant attention over the past few years because of their practical usage in intelligent transportation systems. Previously, the main tools used for vehicle speed measurement were sensors based on radar technology and inductive loop sensors. These systems are able to detect the speed of the vehicle accurately; however, the results are not customizable.

With developments in computer vision technologies, many research works have been conducted for estimating vehicle speeds using cameras. Studies employing single cameras typically make use of optical flows, perspective projections, or homographies for transforming speeds. Even though hardware complexity has been reduced, highly calibrated cameras, specific points of view, and controlled settings remain crucial for an increased level of accuracy. Even minor variations in camera tilt or lighting conditions can influence the calculation.

These shortcomings could be overcome by the average speed enforcement system involving multiple cameras. Here, the vehicles are detected at two or more locations, and the average speed is derived based on the time difference with the known spacing between the locations. This approach has been adopted in several real-world traffic enforcement systems, as it provides a more reliable representation of vehicle behavior over a distance rather than at a single point.

Recent research has increasingly adopted deep learning–based object detection models for vehicle and license plate detection. The YOLO family of algorithms has been widely used due to its real-time performance and high detection accuracy. These models enable simultaneous detection of vehicles and number plates in complex traffic scenes. However, accurate license plate recognition remains challenging under conditions such as motion blur, low illumination, occlusion, and low-resolution inputs.

For license plate text extraction, traditional OCR engines such as Tesseract have been employed in many studies. Although effective for clean images, their performance degrades when plate images are noisy or distorted. More recent works have introduced deep learning–based OCR frameworks, such as EasyOCR, which demonstrate improved robustness under real-world conditions. Several studies emphasize the importance of image preprocessing, including contrast enhancement and thresholding, to improve OCR performance.

Tracking-based approaches, such as SORT and Deep SORT, have also been explored to associate vehicle detections across frames or camera views. These methods improve vehicle matching accuracy, particularly in multi-camera environments. However, tracking-based systems often require continuous video streams and higher computational resources.

Despite these advancements, many existing studies focus either on vehicle detection or speed estimation as isolated problems. Fewer works provide a complete end-to-end system that integrates ANPR, average-speed calculation, violation detection, and

evidence generation within a unified framework. Moreover, limited attention has been given to practical challenges such as handling low-quality images or generating structured violation reports for enforcement purposes.

In contrast, the proposed work focuses on a practical and integrated ANPR-based vehicle speed detection system using a dual-camera approach. By combining deep learning–based detection, robust OCR preprocessing, and average-speed calculation, the system addresses several limitations observed in existing approaches and provides a deployable solution for intelligent traffic monitoring applications.

Existing studies have demonstrated promising results in vehicle detection, license plate recognition, and speed estimation. However, many approaches address these components independently and often require specialized hardware, extensive calibration procedures, or high computational resources. In contrast, the proposed system focuses on integrating these functionalities within a practical and cost-effective framework suitable for real world deployment.

III. LITERATURE SURVEY

This section explains the important research and technologies that are related to our project.Our project makes extensive use of Automatic Number Plate Recognition, vehicle detection, OCR, and an average speed calculation system. For understanding how these systems work, previous research papers, tools, and popular frameworks were studied. The information from these studies helped us design and improve our own AI-based vehicle speed detection system.

The first important technology we studied is the YOLO (You Only Look Once) object detection model. YOLO is one of the fastest and most accurate detection algorithms. It can detect objects like cars and number plates in a single step. The latest version, YOLOv8, is even better in accuracy and speed. Because of this, YOLOv8 is widely used in many real-time vehicle detection projects. In our project too, YOLO helps in identifying the vehicle and capturing the number plate area more clearly.

Different OCR tools were reviewed to read the text on the number plate. Tesseract OCR is an extremely popular OCR engine, maintained by Google.It works well on clean and clear images. But when the plates are blurred or light is low, Tesseract makes mistakes. To solve this, researchers recommend preprocessing like noise removal and sharpening.

Another advanced OCR tool is EasyOCR, which uses deep learning. EasyOCR works better on real-world images that have shadows, angles, or low brightness. It supports many languages and gives more accurate results than classical OCR. Many modern ANPR systems use EasyOCR, and our project also follows this idea for better accuracy.

In vehicle tracking, we studied two important algorithms: SORT and Deep SORT. SORT is a simple and fast method that helps track objects from one frame to another. Deep SORT is an improved version that uses deep learning to track vehicles even when they disappear for a second or when two vehicles overlap. These tracking systems are useful when building multi-camera setups like average-speed detection.

Several survey papers on ANPR explain the common steps involved: detecting the plate, cropping it, reading the text, and cleaning the result. These surveys also list many challenges such as plate tilt, motion blur, night images, and weather problems. These issues helped us understand why real-life ANPR is difficult and why deep learning is important.

We also studied OpenALPR, a popular open-source ANPR system used in parking lots and police enforcement. It shows how ANPR is used in real life and what kind of accuracy and speed is expected in production systems. This gave us a good idea of how to design our system.Then, we went over average-speed enforcement systems that work by placing two cameras a fixed distance apart and calculate speed by the time between them.

Transport departments in many countries use such systems. Their documents explain calibration, evidence handling, and legal rules. This matches exactly what our project aims to do.Some research papers focus on single-camera speed detection, but the problem is that single-camera systems need very accurate calibration and angle correction. Small mistakes can change the speed calculation. Because of this, two-camera systems (like ours) are more reliable.

OpenCV is another important library studied. It provides many functions for image processing like filtering, edge detection, cropping, and resizing. These functions are needed before passing the plate image into an OCR engine.Camera calibration papers help in measuring real-world distances from images. Even though our project uses a fixed known distance between two cameras, understanding calibration gives better clarity about how professional speed systems work.

We also checked papers on vehicle re-identification (Re-ID). This technology is used to match the same vehicle between different cameras, even if the number plate is not clear. Re-ID uses the shape, color, and features of the vehicle for matching.
Several tutorials and GitHub projects helped us understand how YOLO can be trained specifically for plate detection. These practical guides helped us when designing our detection pipeline.

We looked into ReportLab, a Python library that produces PDFs, in order to create the reports. This library enables the addition of text, tables, images, and timestamps to a PDF file. ReportLab was utilised in our project to create a violation report that included entry and exit photos.Finally, we studied about privacy and ethics in ANPR. Since number plates are personal information, systems must follow rules such as secure storage, limited access, and data protection. This helped us design the system with admin-only access.

IV. SYSTEM DESIGN AND ARCHITECTURE

The proposed system aims to develop an end-to-end AI-driven vehicle speed detection system using computer vision and deep learning concepts, and the use of web platforms. In this approach, the main systems and methodologies deployed are modular and utilized to develop systems that are easy to implement and use in real-life traffic systems. The method implemented here uses two fixed cameras marked as Camera A and Camera B located a fixed distance apart.

At the data acquisition layer, cameras continuously capture vehicle images at both checkpoints. When a vehicle passes Camera A, an entry image is captured and forwarded to the processing module. Similarly, when the same vehicle reaches Camera B, an exit image is captured. Each captured frame is time-stamped to enable accurate speed calculation.

The processing layer is responsible for vehicle detection, license plate localization, and text extraction. A deep learning–based object detection model is used to identify the vehicle and localize the number plate region within the image. Once the number plate region is detected, it is cropped and passed through an image preprocessing pipeline. This preprocessing stage includes contrast enhancement, noise reduction, and adaptive thresholding to improve the quality of the plate image before text recognition.

The optical character recognition (OCR) component extracts alphanumeric characters from the preprocessed image of the number plate. The extracted plate number acts as an identifier linking entry and exit points of the vehicle coming from both camera stations. The component classifies the type of vehicle for appropriate speed limits based on categories of cars, bikes, buses, and trucks.

In the speed calculation component, it determines the average speed using the known distance between the two cameras and the known time difference. When it is detected that the calculated speed is breaching the known speed limit based on its type, it is classified as speed violation.

The data management layer is responsible for storing all car information, such as timings, plate numbers, car type, speed values, and images captured. The web-based application layer, built using Django framework technology, is used as an interface through which system events can be monitored. Using this interface, system events can be accessed by authorized personnel. They can view logs of detection events, analyze violative events, and generate PDF reports.
On the whole, the architectural design of the system guarantees perfect compatibility of the AI-based detection method with web-based data management and reporting. Thus, this system can be applicable in intelligent traffic monitoring and enforcement.

V. WORKFLOW

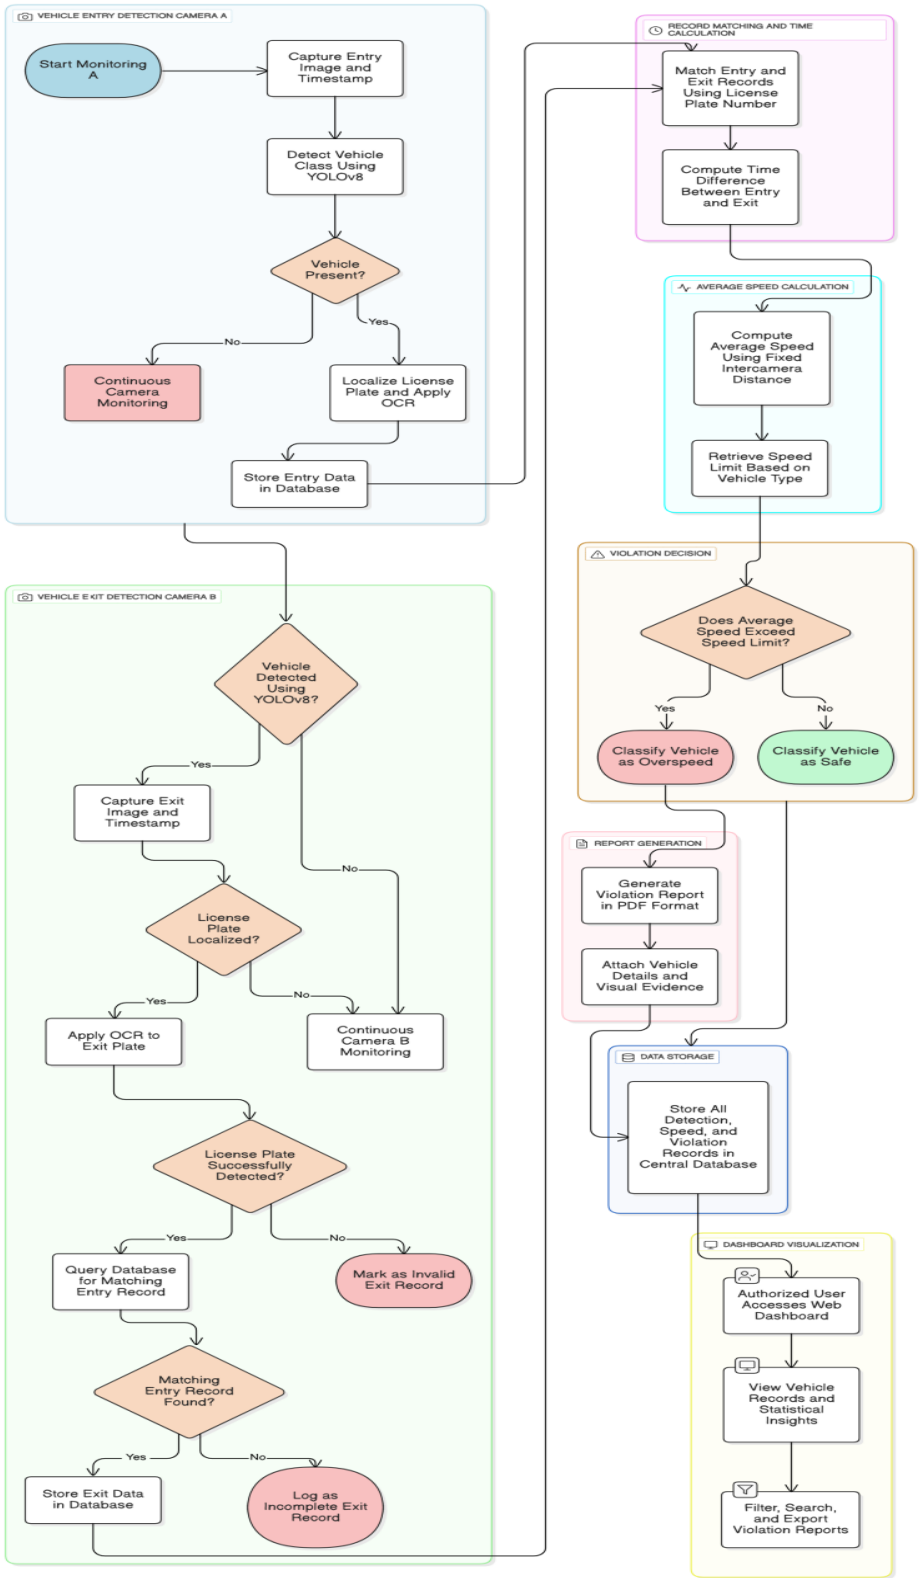

Fig. 1(a): Workflow of The Proposed AI-Based ANPR System

VI. IMPLEMENTATION

The proposed system for vehicle speed detection is developed with a modular and scalable approach as an integrated framework with computer vision-based approaches using deep learning along with a web-based application setup.

A. Data Acquisition and Camera Functionality:

There are two fixed cameras named Camera A and Camera B. They are mounted at designated points on both sides of the road, at a known interval. The pictures of cars are taken by Camera A at entry points, and by Camera B at exit points. Each picture is labeled with timestamps to enable the accurate calculation of car speeds. The processing is done per picture frame and not per video.

B. Vehicle and License Plate Detection :

To identify cars and locate license plates in taken pictures, deep learning-based object detection is utilized. Because of its real-time performance and detection accuracy, a YOLO-based model is employed. The identified vehicle is categorized into pre-established groups, such as truck, bus, bike, or car. After being cropped from the identified bounding boxes, license plate regions are sent for additional processing.

C. Image Preprocessing and OCR:

To improve license plate recognition accuracy under varying lighting conditions and image quality, an image preprocessing pipeline is applied before OCR. This pipeline includes contrast enhancement, noise reduction, resizing, and adaptive thresholding. These preprocessing steps enhance character visibility and reduce background interference. Optical Character Recognition is then performed using a deep learning–based OCR engine to extract alphanumeric characters from the license plate region.

D. Record Matching and Speed Calculation:

The plate number extracted is a token that is used for linking entry and exit points for images captured by the two cameras. The time difference is calculated from entry to exit once a plate number has been found. The speed of the vehicle is calculated using the time difference and the known separation of the cameras. Speed limits, which depend on the type of vehicles, are introduced to check whether the calculated speed is above the limit.

E. Violation Detection and Data Storage:

If the calculated average speed is greater than the preset speed limit, the vehicle will be identified as an overspeed case; otherwise, the case will be identified as safe. All details such as the number of the plates, type of vehicles, timings, values of the calculated speeds, violation, and images will be stored in a proper database.

F. Web Interface and Reporting:

A web-based system is developed using Django for handling system operations and visualization. Authorized users are provided with access to a dashboard for monitoring vehicle records, violation statistics, and filtering violations by various parameters. For overspeed violations, the system provides downloadable PDF reports containing details of the vehicle, overspeed details, times, and pictorial evidence taken at both checkpoints.

VII. RESULT

The system successfully detects vehicles, recognizes their number plates, and calculates the average speed using entry and exit timestamps. Overspeeding vehicles are correctly identified and displayed on the dashboard. The PDF report generation feature provides proper documentation of each violation, including entry and exit images and calculated speed values. OCR and YOLOv8 performed well under normal daylight conditions, with accuracy improving after applying preprocessing techniques. The proposed system reduces human effort and provides an automated and low-cost method for road speed enforcement.

Although the proposed system demonstrated reliable performance during testing, the experimental evaluation was conducted using 83 vehicle samples. While the obtained results validate the feasibility of the proposed approach, a larger and more diverse dataset would provide a more comprehensive assessment of system performance under different traffic, environmental, and operational conditions. Future evaluations will focus on testing the system using larger datasets and extended real-world deployments.System performance may decrease under severe occlusion, poor weather conditions, nighttime operation, motion blur, and license plate tampering. These limitations have been acknowledged as future improvement areas.

A qualitative comparison with existing ANPR-based monitoring systems indicates that most reported approaches focus on either license plate recognition or vehicle speed estimation as independent tasks. In contrast, the proposed system integrates vehicle detection, OCRbased license plate recognition, average-speed calculation, violation classification, database management, dashboard visualization, and automated PDF report generation within a single framework. This integration improves practical deployability and operational efficiency for intelligent traffic monitoring applications.

The performance of the proposed system may decrease under challenging conditions such as heavy rain, fog, lowlight environments, severe motion blur, partial plate occlusion, and deliberate license plate tampering. Although the system performed reliably under normal daylight conditions, these scenarios remain challenging for computer vision-based monitoring systems and require further investigation.

On the whole, it is proved that the system can successfully detect cars, accurately identify license numbers, and correctly calculate the average speed of the cars.

Fig. 2(a): Web-based dashboard showing vehicle statistics and speed violation summary

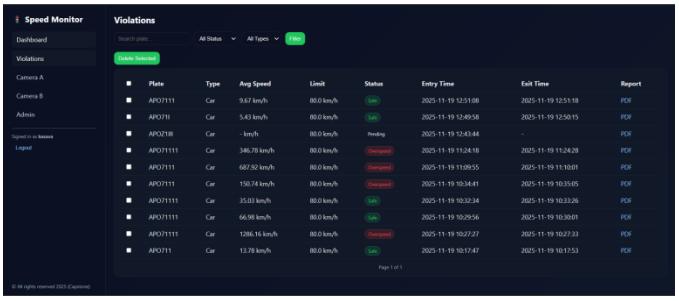

Fig. 2(b):Violation records displaying detected vehicles, speed status, and report generation

VIII. CONCLUSION

This paper presented an AI-Based ANPR System for Intelligent Vehicle Speed Detection. The system captures vehicle images at two checkpoints, extracts license plate information using deep learning and OCR techniques, and calculates average vehicle speed based on the time difference and a fixed inter-camera distance. Vehicle type classification and speed limit checking have been applied to automatically detect speeding violations.

Experimental results demonstrate that the system performs vehicle detection reliably, license plate recognition accurately, and speed estimation correctly. The use of a two-camera average-speed approach reduces dependency on single-point measurements and improves overall accuracy. The developed web-based interface and automated report generation further demonstrate the practicality of the proposed system for real-time traffic monitoring applications.

Since vehicle license plates contain sensitive information, secure storage and controlled access mechanisms are essential. The proposed system restricts access to authorized users through an administrative interface. Data retention policies, privacy regulations, and legal compliance requirements should be considered before large-scale deployment. Additionally, deployment challenges such as camera positioning, network connectivity, maintenance, and environmental conditions must be addressed to ensure reliable operation.

Overall, the proposed approach provides an effective and deployable solution for intelligent vehicle speed monitoring and violation detection.

REFERENCES

[1] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, “You Only Look Once: Unified, Real-Time Object Detection,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2016. [foundation]

[2] Ultralytics, “YOLOv8 Documentation,” 2023. [ultralytics]

[3] Ultralytics, “Ultralytics YOLO GitHub Repository,” 2023.[github]

[4] R. Smith, “An Overview of the Tesseract OCR Engine,” in Proc. Int. Conf. Document Analysis and Recognition (ICDAR), 2007. [research]

[5] JaidedAI, “EasyOCR: Ready-to-Use OCR with PyTorch,” GitHub Repository, 2020. [github]

[6] A. Bewley, Z. Ge, L. Ott, F. Ramos, and B. Upcroft, “Simple Online and Realtime Tracking,” in Proc. IEEE Int. Conf. Image Processing (ICIP), 2016.[arxiv]

[7] N. Wojke, A. Bewley, and D. Paulus, “Simple Online and Realtime Tracking with a Deep Association Metric,” in Proc. IEEE Int. Conf. Image Processing (ICIP), 2017. [arxiv]

[8] OpenALPR, “Open Automatic License Plate Recognition (OpenALPR),” GitHub Repository, 2020. [github]

[9] X. Xu et al., “CCPD: Chinese City Parking Dataset,” GitHub, 2018. [github]

[10] R. Laroca et al., “UFPR-ALPR: A New Benchmark for Automatic License Plate Recognition,” IEEE Access, 2022.[github] Available: https://github.com/raysonlaroca/ufpr-alpr-dataset

[11] AvLab, “AOLP: Application-Oriented License Plate Dataset,” GitHub, 2020. [github]

[12] L. A. Marcomini and A. L. Cunha, “A Comparison between Background Modelling Methods for Vehicle Segmentation in Highway Traffic Videos,” arXiv preprint, 2018. [arxiv]

Z. Zhang, “A Flexible New Technique for Camera Calibration,” Microsoft Research, Tech. Rep., 2000. [microsoft]

[14] S. Khan and H. Ullah, “A Survey of Advances in Vision-Based Vehicle Re-Identification,” arXiv preprint,

2019.[arxiv]

[15] X. Hu, X. Xu, Y. Xiao, H. Chen, S. He, J. Qin and P.-A. Heng, "SINet: A Scale-Insensitive Convolutional Neural Network for Fast Vehicle Detection," arXiv preprint, 2018.[arxiv]

[16] GDPR.eu, "General Data Protection Regulation (GDPR) Compliance Guidelines," 2018. [gdpr]

[17] UK Home Office, "Speedmeter, Traffic Light and Prohibited Lane Enforcement Camera Handbook," 2023. [assets]

[18] N. Owen, S. Ursachi, and R. Allsop, "The Effectiveness of Average Speed Cameras in Great Britain," RAC Foundation, 2016. [racfoundation]

[19] B. Horn and B. Schunck, "Determining Optical Flow," MIT, 1981. [cmor-faculty]

[20] S. L. Pundlik and A. M. Joshi, "Accurate Vehicle Speed Estimation from Monocular Camera Footage," ISPRS Annals, 2020. [isprs-annals]

ISSN: 3049-1479(Online)

AI-Assisted Digital Forensics for Automated Incident Investigation

Manya Parikh[1], Parth Talaviya[2], Deep Solanki[3] ,Jigar Gajjar[4]

Lok Jagruti Kendra University, Ahmedabad,India[1]
Lok Jagruti Kendra University, Ahmedabad,India[2]
TechD Cybersecurity Limited, Ahmedabad, India[3]
TechD Cybersecurity Limited, Ahmedabad, India[4]

manyaparikh23@gmail.com[1],talaviyaparth07@gmail.com[2],deepsolanki5799@gmail.com[3],cyberian.jg@gmail.com[4]

Abstract: Digital forensics has emerged in the scene as a necessary area in the present scenario of cybersecurity, due to the rising trend of dependence on digital infrastructure that is vulnerable to complex cyber-attacks. The rising level of digital evidence and its complexity has made it challenging for traditional methods of investigation to be efficient and scalable. In this context, Artificial Intelligence has emerged in the scene as a promising solution that can assist in enhancing forensic practices by facilitating fast data analysis, complex pattern analysis, and efficient analysis of large datasets. This will assist in the efficient identification of potential security incidents and will assist in enhancing forensic practices as a whole. However, the application of AI in digital forensics also has some critical points to be considered. The concerns of data integrity, transparency, privacy, and admissibility of evidence in courts imply that the developments in technology are not sufficient to ensure the accuracy and integrity of the same. The requirement of human intelligence in interpreting the findings, understanding the context, and ensuring that the forensic evidence is credible and defensible is required in this context. This paper explores the effects of AI on the evolution of digital forensic methods and also explores the challenges that are associated with the use of AI in digital forensic analysis. This paper explores the importance of achieving a balance where intelligent technology can support human expertise, leading to more efficient and accountable forensic processes in the digital era.

Keywords: Digital Forensics, Artificial Intelligence, Cybersecurity, Automated Investigation, Digital Evidence, Machine Learning, Threat Detection

I. INTRODUCTION

Digital forensic is a department of forensic science that focuses on the identification, maintenance, analysis and presentation of digital evidence in such a way that protects its integrity and admissibility in court. As the digitization of society is speeding up, electronic devices like smartphones, computers and cloud computing services have become the primary warehouse of information, and digital evidence has become a crucial component of modern investigations. According to Casey[1], sound forensic practice is essential for ensuring that evidence remains of security incidents in organizations, insider threats, and regulatory compliance activities. As noted by Garfinkel [3], digital forensics can be used to recreate events from a complex digital environment, allowing investigators to take informed legal and business actions. However the sheer volume of data and refinement of cyber-attacks have made traditional forensic practices time-consuming and hard to scale or identify the incident.

Artificial Intelligence (AI) has emerged as a transformative technology capable of addressing these limitations. By automating data analysis, detecting patterns and speeding up investigative workflows, AI-assisted digital forensics improve both the speed and accuracy of the incident investigators to manage large amounts of datasets efficiently while focusing on high value evidence.

A. Usage of Digital Forensics in Real Life

Digital forensics can be defined as the procedure of examining and featuring electronics evidence. It ensures that digital evidence is managed in an identical valid manner so that it can be accepted by a court law. According to Casey[1], it is crucial

to protect the integrity of evidence during forensic analysis. Carrier[2] describes structured forensic processes that help in the analysis of evidence.

http://doi.org/10.63766/spujstmr.26.000097

Fig. 1. AI Enhanced Digital Forensics Lifecycle

In practical applications, digital forensics helps the police in investigating cyber crimes like ransomware attacks, phishing attacks, malware attacks, identity theft, financial fraud and many more attacks. Organizations use digital forensics to investigate data breaches and insider threats. Garfinkel[3], states that digital forensics helps in the rebuilding of digital timelines, which helps in understanding the flow of events that happened in a network of interconnected systems. With the increasing use of digital devices in communication and business, digital forensics has become a crucial tool in both criminal and business justice.

B. Benefits of Artificial Intelligence

Artificial Intelligence (AI) is a kind of technology that enables machines and computers to simulate machine learning, and AI has ability to perform tasks which require human intelligence, such as learning, problem solving, decision making, predictions and awareness. According to Russell and Norvig[4], AI is the study of intelligent agents that explain their environment and act to achieve specific goals. Machine learning is a significant part of AI that enables systems to improve their performance through experience, as opposed to programming.

The benefits of AI are that several advantages exist, especially its ability to analyze large amounts of data in a short time and recognize patterns that are impossible for humans to recognize. According to Goodfellow et al.[5], deep learning is a significant breakthrough that enables accurate classification and prediction. Domingos[6], states that AI is essential in discovering hidden patterns in data, while According to Mitchell[7], states that adaptive learning models continually improve the accuracy of analysis. According to Jordan and Mitchell[8], they state that AI enables automation, which decreases human labor and errors.

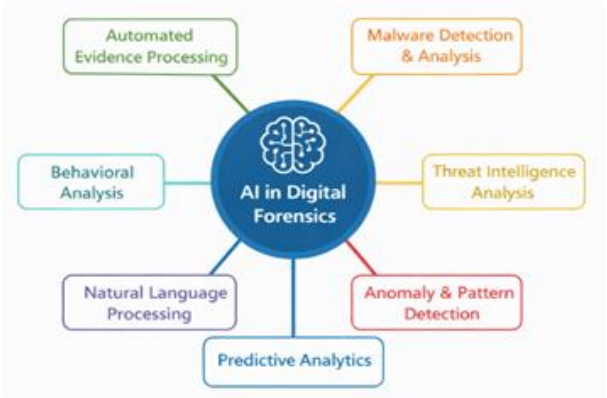

Fig. 2. Use of AI in Digital Forensics

C. Integration of AI into Digital Forensics Process

AI can be incorporated at various stages of the digital forensic process, such as data gaining, analysis, examination, and reporting. Standard digital forensic analysis involves the manual filtering of large amounts of datasets, but AI-based solutions can automatically classify files, detect anomalies, and point out areas of interest. As stated by Carrier and Spafford[9], the systematic nature of digital forensic work can be significantly assisted through automation in large scale digital environments.

Machine learning can be applied in the detection of unusual patterns in network traffic, while deep learning helps in malware detection and image processing. Singer and Friedman[10], state that AI enhances forensic preparedness by allowing for fast digital evidence triage. Quick and Choo[11], further clarify that intelligent automation enables faster response and investigation efficiency during cyber attacks. Moreover, AI can assist in providing a timeline. Such analysis is achieved by correlating data from various resources for visualization. of attack patterns.

However, such successful integration needs to be validated for forensic integrity and legal acceptability. Raghavan[12] highlights the need for AI generated results to be interpretable to enable the justification of conclusions in a court of law. When done correctly, AI-based solutions can help digital forensics shift from a relative field to a proactive investigative approach.

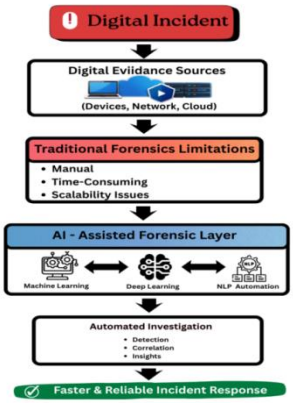

Fig. 3. High-Level AI-Assisted Digital Forensics Workflow

D. Benefits of using AI in Digital Forensics Compared to Traditional Methods

AI-powered digital forensics greatly enhances the efficiency and level of analysis of investigations compared to the standard approach. The traditional approach is largely dependent on human knowledge and expertise. However, this approach is time-consuming when dealing with large amounts of data. AI can process terabytes of data in a matter of minutes, thus reducing the backlogs of cases. Russell and Norvig [4] highlight that intelligent automation of tasks improves decision-making by reducing analytical overload.

Another important benefit of AI is accuracy. Machine learning algorithms identify hidden patterns and exceptions that human investigators may not be able to identify. Singer and Friedman [10] highlight that AI improves the ranking of evidence, ensuring that key evidence is analyzed first. Moreover, predictive analysis enables organizations to predict threats before they become major incidents. AI also allows for consistency in the forensics process. Quick and Choo highlight the aspect in their article[11], stating that the algorithmic analysis reduces the variability associated with the various human analyses. Moreover, AI facilitates the scalability of the forensics process in the cloud and distributed systems, which can be difficult to handle through conventional processes. Though AI processes need to be overseen by a human, it can add speed, accuracy, and scalability to the forensics process.

E. Concerns of using AI for Investigative Purpose

Although AI improves forensic analysis, it also presents a number of challenges. First, there is the problem of bias in algorithms, which can produce erroneous results if trained on incomplete data. Raghavan[12], suggests that such errors could compromise the validity of evidence. Another problem is the lack of transparency in many AI systems, which operate as “black boxes”, thereby hiding the logic used to arrive at certain concussions. Goodfellow et al.[5], highlight the need for understandability in many AI systems, especially in forensic analysis. There is also the problem of privacy in AI systems that handle personal information.Moreover, there is the risk of human forensic analysis being overshadowed by automation, which could lead to errors going undetected. Garfinkel[3], suggests that human forensic analysis is still necessary to verify the findings of automated systems.

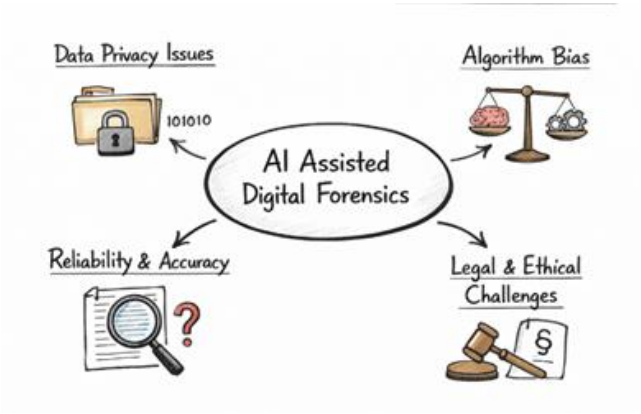

Fig. 4. Key concerns in AI assisted Digital Forensic Investigations

F. Analysis

The introduction implies that the AI technology is revolutionizing the aspect of digital forensics since it is able to improve the speed of the investigation, increase the accuracy of the analysis, and make the process quite efficient. However, there are underlying issues of transparency, objectivity, privacy, and some aspects of legal soundness. The above statements, therefore, suggest an argument that for there to be efficacious digital forensics, there is a need to integrate ethics into the investigation process. Essentially, digital forensics facilitated by AI technology is quite advantageous, but there is a need for the integration of accompanying expertise.

II. LITERATURE SURVEY

A. History of Digital Forensics and Evolution

Digital forensics started gaining its shape as a discipline in the 1980s due to the rapid spread of personal computers and the escalation of various cyber crimes. Investigations concentrated on recovering deleted information and examining individual computers for evidence. With the advancement of technology, investigators encountered increasing challenges of large amounts of information, different storage media, and complex systems. The contributions of Carrier [2] organized the frameworks of digital forensic investigation, which standardized the processes of investigation and made them reliable.

In the years that followed, computer forensics evolved to accommodate other devices, issues of cloud storage, IoT, and the vast and extensive enterprise systems. Casey [1] explains that modern computer forensics is based on the development and implementation of methods that preserve the integrity of the collected data and ensure their admissibility in court. The field continues to advance and develop to accommodate the changing technology and the evolving smarter cyber attacks.

B. Integration of AI for Digital Forensics and its Benefits

The concept of Artificial Intelligence in digital forensics is the talk of the town today, owing to the rate at which cybercrime is advancing and the amount of digital evidence we have to go through. Traditional forensic methodologies, although robust, are not efficient enough to keep up with the pace of the current world, particularly when dealing with massive amounts of data spread all over the place. AI comes to the rescue by using smart automation, filtering, and the speedy identification of the most relevant evidence.

A key advantage of using machine learning algorithms in any investigation process is that they are effective in detecting suspicious behavior. Through the use of machine learning algorithms, AI has made it significantly easier to reduce the workload of those in the investigation process. Quick and Choo [11] explain that the advantage of using AI in any investigation process is that it enhances forensic readiness. The usage of sophisticated attacks, however, makes it necessary for firms to be proactive in their approach as opposed to relying on reactive approaches.

Deep learning will enhance forensic work through sharpening the classification of malware, the verification of the biometric, and the analysis of multimedia content. According to Goodfellow and colleagues, neural networks can find complex patterns in high-dimensional data, and the forensic field can greatly benefit from this since every connection must be traced. Furthermore, Domingos introduced predictive modeling as an avenue where the investigator can make predictions concerning the various potential attacks.

Another important benefit of AI is that it can assist decisions. Intelligent systems are able to sift through key evidence to provide organized analyses instead of bombarding investigators with information. According to Russell &Norvig[4], this technology improves situational awareness, allowing investigators to make speedier decisions.

Although the benefits exist, the academic press cautions that we must approach AI-based interventions in forensics with caution. According to Raghavan [12], the "urgency" for "interpretable" AI builds on the "need for machine learning results that can be intelligently presented as evidence." Moreover, the "expertise of human analysts" would be required "to interpret the results that these algorithms compute." Nevertheless, the consensus in academic circles is that AI represents a "truly transformative" shift for digital forensics, which "speeds and broadens the reach of investigation in this increasingly complex digital world."

AI can enhance the scope of examination by allowing the investigator to peruse data across dispersed systems, including those found in the cloud environment, while allowing for ease of search. This is important due to the massive amounts of both organized and unorganized data being generated by organizations. The evolution of smart forensic investigation software has the capability to connect information from numerous sources, thus enabling the observer to view potential associations that were not initially recognized. This is not only important for ensuring the precision of investigation results but also makes the investigation process faster.

C. Comparison Table: Traditional vs AI Assisted Digital Forensics

TABLE I

EVOLUTION FROM TRADITIONAL TO AI-ASSISTED DIGITAL FORENSICS

Aspect	Traditional Digital Forensics	AI-Assisted Digital Forensics
Investigation Speed	Time-consuming and manual	Rapid automated analysis
Data Handling	Difficult to scale	Efficient handling of massive datasets
Analytical Capability	Limited pattern recognition	Advanced behavioral analytics
Accuracy	Dependent on examiner skill	Improved through machine learning
Error Probability	Higher risk of oversight	Reduced through automation

D. Real Life Use cases

Digital forensics has played a great role in the investigation of notable cyber incidents in the recent past. Notable is the Sony Pictures Entertainment 2014 breach, which involved the breach of internal networks and the leak of confidential and sensitive information. There was also a breach of operations on an across-the-board scale. To draw a picture of the breach and attempt to establish how it was conducted, various digital investigation techniques, such as malware analysis and network forensics were employed. As Casey in [1] puts it, a systematic digital approach was critical in the construction of a timeline.

At the same time, there was another parallel major event in the case of the Target Data Breach, which took place in 2013 when the attackers accessed the point-of-sale systems, thereby releasing millions of customer payment records. The detailed and thorough investigation into the breach led to the identification of the point where the breach was initiated and where the vulnerabilities were in the company's security system. According to Garfinkel, such a process not only revealed the events of the case but also taught the way forward in the process of defense.

Collectively, these cases underscore the growing reliance on advanced forensic capabilities within contemporary cybersecurity environments. As cyber threats continue to increase in sophistication and scale, the integration of AI-supported analytical tools is increasingly recognized as an important progression in strengthening investigative effectiveness.

E. Concerns of Using AI for Digital Forensics

Although there are a number of probable advantages to the use of AI in furthering forensic investigations, there are also some notable limitations and difficulties in adopting AI. There are many chances that AI might end up giving very erroneous leads if there is any bias in the algorithms used in the process. Raghavan in his paper highlights the need for continuous validation and testing.Explainability is another significant challenge to be overcome. For instance, explainability has to be provided in legal

investigations involving digital data, and many AI systems are "black boxes," as they cannot be easily interpreted. According to Goodfellow et al. [5], interpretable models should be developed to increase explainability and trust of AI systems.

The issue of privacy creates a new dimension to the complexities of applying AI in forensics. Analyzing various data has the potential to reveal personal information, thereby raising a number of questions about the usage of this information. Carrier highlights the importance of developing strict evidence-handling protocols to avoid any kind of abuse.On another note, overdependency on technology may lead to a loss of value in human expertise. It may also lead to a situation where crucial factors are ignored if one is over dependent on computer-generated data. As Garfinkel puts it, human experience should always be used to validate the outcome and ensure fair and reliable conclusions.

As a result, it is the consensus opinion of a significant number of researchers that there is a need to employ AI in a cooperative mechanism. That is to say, the AI assists the investigator as opposed to replacing the investigator. This therefore guarantees the requirement to 'align technological innovation with ethical governance' in AI use. AI can be effectively used to enhance the efficiency of the forensic process without compromising the admissibility of the digital evidence.

TABLE II

KEY ADVANTAGES OF AI-ASSISTED DIGITAL FORENSICS

Feature	Traditional Approach	AI-Assisted Approach
Evidence Analysis	Manual and time-intensive	Automated and faster processing
Pattern Detection	Limited by human capability	Identifies hidden patterns efficiently
Investigation Speed	Slower due to large data volumes	Accelerates investigative workflow
Decision Support	Based mainly on examiner judgment	Provides data-driven insights

III. ANALYSIS

TABLE III

IMPACT OF AI ON DIGITAL FORENSIC INVESTIGATION PROCESSES

Investigation Stage	Traditional Forensic Approach		Investigative Advantage
Evidence Identification	Manual device examination	Automated data discovery	Faster recognition of relevant evidence
Data Collection	Time-intensive acquisition	Intelligent filtering	Reduces unnecessary data handling
Evidence Analysis	Human-driven interpretation	Machine learning pattern detection	Improves analytical accuracy
Incident Reconstruction	Sequential review of logs	Automated timeline generation	Enhances event clarity
Threat Detection	Reactive investigation	Predictive analytics	Supports proactive response
Decision Support	Based on examiner experience	Data-driven insights	Strengthens investigative judgment
Reporting	Manually compiled findings	AI-assisted documentation	Accelerates case preparation

From the table above, it is clear that the integration of Artificial Intelligence in digital forensics helps a lot in speeding up the accuracy and speed of the investigations. With the ability to handle the boring and repetitive work of identifying complex patterns, Artificial Intelligence allows investigators to focus on the critical analysis part of the work rather than being tied down to the handling of the data. This helps in speeding up the timelines and making the reconstructed events more reliable. However, as we rely more on these automated systems, it is important to retain human review to preserve context and the value of the results.

TABLE IV

BENEFITS VS RISKS OF AI-ASSISTED DIGITAL FORENSICS

AI Capability	Investigative Benefit	Potential Risk
Automated Analysis	Speeds up evidence review	Overreliance on algorithms
Pattern Recognition	Detects hidden relationships	False positives/negatives
Predictive Modeling	Anticipates threat behavior	Model bias
Large-Scale Data Handling	Manages massive datasets	Privacy concerns
Decision Support	Enhances investigator efficiency	Reduced human interpretation

What the table above shows is that Artificial Intelligence is indeed changing the game when it comes to digital forensics. It is a technology that can speed up processes, increase accuracy, and expand our capabilities when it comes to data analysis. Automation is also a great help to investigators who have to deal with massive amounts of digital evidence. However, despite all the advantages that come with Artificial Intelligence, human oversight should still be at the forefront.

IV. LIMITATIONS

Artificial Intelligence has made a big difference in digital forensics, but its application is not without controversy. The quality of the results that you obtain using Artificial Intelligence is highly dependent on the training data that you are using. If the data is flawed, then the results are likely to be flawed as well. The other problem is the who and why of the algorithms. Lack of transparency in the algorithms can make it difficult to understand the reasoning.

The move towards AI in the forensic field requires a great deal of technical expertise and financial resources, which not all forensic departments have access to. On the other hand, the management of large amounts of sensitive digital data continues to raise privacy and data protection concerns. There is also the problem of automated systems failing to detect or failing to properly identify sophisticated, novel cyber threats, which may result in false positives or false negatives that can cast doubt on the results. In addition, AI software may not interpret evidence in the same context as a human forensic investigator, and the admissibility of AI-generated results in court may be raised.

V. CONCLUSION

Digital forensics has now become an integral part of the current cybersecurity landscape as organizations are faced with an ever-increasing number of complex and data-driven cyber events. The integration of Artificial Intelligence has greatly enhanced the process of investigation as it allows for faster analysis of data, facilitates pattern analysis, and helps experts handle large amounts of digital evidence. The power of technology has made the process of forensic analysis simpler and enabled investigators to focus more on high-level analysis and less on details.However, the growing application of AI also emphasizes the need for cautious application of technology. The issues of data integrity, explainability, privacy, and admissibility in courts indicate that technology alone is not sufficient to ensure the accuracy of results for the investigation process. The application of human expertise is of utmost importance to understand the results, the nuances of the context, and the integrity of the forensic process.The future of digital forensic investigation will be based on the complementary application of intelligent technology and human expertise. AI can be a useful tool in the process of investigation that can enhance accuracy and ensure trust, accountability, and integrity in the context of an increasingly digital world.

REFERENCE

[1] E. Casey, *Digital Evidence and Computer Crime: Forensic Science, Computers, and the Internet*. Amsterdam, Netherlands: Academic Press, 2011.
[2] B. Carrier, *File System Forensic Analysis*. Boston, MA, USA: Addison-Wesley Professional, 2005.
[3] S. L. Garfinkel, "Digital forensics research: The next 10 years," *Digital Investigation*, vol. 7, pp. S64–S73, 2010.
[4] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*. Upper Saddle River, NJ, USA: Prentice Hall, 1995.
[5] I. Goodfellow, Y. Bengio and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.
[6] P. Domingos, *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*. New York, NY, USA: Basic Books, 2015.
[7] T. M. Mitchell, *Machine Learning*. New York, NY, USA: McGraw-Hill, 1997.
[8] M. I. Jordan and T. M. Mitchell, "Machine learning: Trends, perspectives, and prospects," *Science*, vol. 349, no. 6245, pp. 255–260, 2015.
[9] B. Carrier and E. H. Spafford, "Getting physical with the digital investigation process," *International Journal of Digital Evidence*, vol. 2, no. 2, pp. 1–20, 2003.
[10] P. W. Singer and A. Friedman, *Cybersecurity and Cyberwar: What Everyone Needs to Know*. Oxford, U.K.: Oxford University Press, 2013.
[11] D. Quick and K. K. R. Choo, "Digital forensic intelligence: Data subsets and Open Source Intelligence (DFINT + OSINT)," *Future Generation Computer Systems*, vol. 78, pp. 558–567, 2018.
[12] S. Raghavan, "Digital forensic research: Current state of the art," *CSI Transactions on ICT*, vol. 1, no. 1, pp. 91–114, 2013.
[13] A. L. Buczak and E. Guven, "A survey of data mining and machine learning methods for cybersecurity intrusion detection," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 2, pp. 1153–1176, 2016.
[14] M. A. Ferrag, L. Maglaras, S. Moschoyiannis and H. Janicke, "Deep learning for cyber security intrusion detection: Approaches, datasets, and comparative study," *IEEE Communications Surveys & Tutorials*, vol. 22, no. 2, pp. 1347–1380, 2020.
[15] R. Sommer and V. Paxson, "Outside the closed world: On using machine learning for network intrusion detection," in *Proc. IEEE Symposium on Security and Privacy*, 2010, pp. 305–316.
[16] J. A. Kroll *et al.*, "Accountable algorithms," *University of Pennsylvania Law Review*, vol. 165, no. 3, pp. 633–705, 2017.
[17] I. H. Sarker, "Machine learning: Algorithms, real-world applications and research directions," *SN Computer Science*, vol. 2, no. 3, pp. 1–21, 2021.
[18] H. Chen, R. H. Chiang and V. C. Storey, "Business intelligence and analytics: From big data to big impact," *MIS Quarterly*, vol. 36, no. 4, pp. 1165–1188, 2012.
[19] S. Mohurle and M. Patil, "A brief study of WannaCry threat: Ransomware attack 2017," *International Journal of Advanced Research in Computer Science*, vol. 8, no. 5, 2017.
[20] Digital Forensic Research Workshop (DFRWS), "A roadmap for digital forensic research," *Digital Forensic Research Workshop*, 2001.
[21] A. Behl and K. Behl, *Cybersecurity and Cyberwar: What Everyone Needs to Know*. Oxford, U.K.: Oxford University Press, 2017.
[22] S. Shinde, A. Kulkarni and S. Bhosale, "Cyber threat intelligence: Survey and research directions," *International Journal of Computer Applications*, vol. 179, no. 46, pp. 20–24, 2018.

ISSN: 3049-1479(Online)

A Comprehensive Study of Cyber Attacks against Artificial Intelligence Systems

Aayush Shah[1], Momin Samir[2], Deep Solanki[3], JigarGajjar[4]

Lok Jagruti Kendra University, Ahmedabad, India[1]
Lok Jagruti Kendra University, Ahmedabad, India[2]
TechD Cybersecurity Limited, Ahmedabad, India[3]
TechD Cybersecurity Limited, Ahmedabad, India[4]

a00311027@gmail.com[1], mominsamir2202@gmail.com[2], deepsolanki5799@gmail.com[3], cyberian.jg@gmail.com[4]

Abstract: AI security is an emerging industry focused on securing AI, as well as its models, data, and overall pipeline, against bad actors who want to manipulate, access, and misuse it. With AI increasingly integrated into critical domains like healthcare, finance, autonomous technology, and cyber security, the attack surface is growing exponentially. An AI system is no longer just vulnerable and open to attacks from malware and cyber-attacks, but it's now also exposed to 'adversarial examples,' 'poisoning,' 'model inversion attacks,' and 'prompt injection' and 'model extraction' attacks. These are attacks that put the overall 'confidentiality, integrity, and availability' of AI at grave risk. Securing AI requires a multidisciplinary approach comprising sound software engineering practices, robust machine learning paradigms, cryptography, and good governance. In other words, security needs to be embedded throughout the entire life cycle of an AI application—from data collection, to model training and deployment, to continuous monitoring and analysis. Methods such as adversarial training, differential privacy, federated learning, hosting the model securely, validating input data, and using continuous threat intelligence are just some other methods that can improve the security posture. Compliance, explainability tools, and red teaming have begun to be recognized as part of the fundamentally important concept of AI governance. As adversaries begin to deploy AI themselves, traditional security models need to begin to transform toward more adaptive and zero-trust-based security environments. Rather than simply a security concern, AI security is also a powerful strategic imperative for ensuring that AI is not only trusted but also employed ethically in our digital world. In this research our main focus is on evolution of artificial intelligence systems, uses of AI in cyberattacks , uses of AI in cyber defence and vulnerabilities to the AI models

Keywords—Artificial Intelligence Security, Adversarial Attacks, Data Poisoning, Model Extraction, Prompt Injection, AI-Based Cybersecurity, AI Governance.

I. INTRODUCTION

A. Evolution of AI and What is AI?

Computer systems that can carry out activities requiring human cognitive abilities [1], [3], including reasoning, perception, learning, and decision-making, are referred to as artificial intelligence (AI) [35]. Symbolic AI and rule-based expert systems, which relied on logical reasoning and manual coding of knowledge representation, were the first generation of artificial intelligence (1950s–1980s). Despite their great success in their applications, these systems struggled to deal with massive data and ambiguity.[12]-[15]. The 1990s and the early 2000s saw the paradigm shift in AI to statistical machine learning, with a focus on probabilistic modeling and learning from data. The paradigm shift in deep learning that occurred around 2012 [31],[22], thanks to the advancements in computing power and the availability of large amounts of data, led to a dramatic improvement in image recognition, speech processing, and natural language understanding [31].

However, in recent years, the emergence of foundation models and transformer models has resulted in a paradigm shift, making it possible for pre-trained models to generalize across a variety of tasks [28]. These models are the foundation of generative AI models that have the capability of generating human-like text, images, and code. However, with the increasing capabilities of AI, there has been a growing concern about interpretability, bias, and systemic risks [30].Thus, AI has evolved from rule-based reasoning to large-scale data-driven intelligence, marking a shift toward increasingly autonomous and adaptive systems.

http://doi.org/10.63766/spujstmr.26.000098

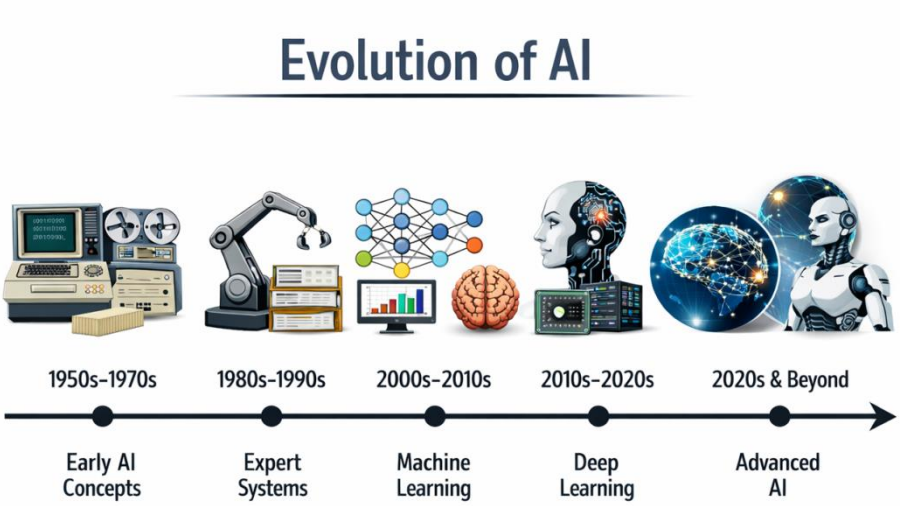

Fig. 1. Evolution of Artificial Intelligence

B. Evolving threat landscape because of AI

Another area that has been affected by AI is the threat environment in the area of cyber security. Generative models of AI can be used to create realistic phishing messages, voices, and videos [25], [35]. This has made social engineering attacks more successful and difficult to detect (NCSC, 2024). Machine learning algorithms have made attackers to scan for loopholes of a system with little human involvement.In recent years, because of the increasing number of cyber attacks, conventional cyber security techniques have failed to prove effective. As a result, organizations have started to adopt AI-based techniques. These technologies can assist them in enhancing their defence mechanisms and also help them in identifying threats beforehand. Most AI-based technologies, such as NLP, machine learning, and neural networks, have the ability to replicate human actions and the capacity to make decisions like humans. These are the most widely used technologies in the field of cyber security.

Recent studies emphasize that AI-powered automation makes it possible for attackers to launch massive, adaptive, and covert attacks with little human involvement (Buczak&Guven, 2016; Sommer&Paxson, 2010)[6],[7]. The most popular AI-powered cyber attacks include:

1. Artificial Intelligence-assisted Phishing and Social Engineering Attacks - Artificial Intelligence is being used to create highly personalized phishing messages, which is then used as a spear phishing attack. Spear phishing is having a high rate of success because it is similar to the communication patterns in nature.
2. Polymorphic and AI Obfuscated Malware – In this type of malware, ML algorithms are used to vary the signature of the malware [33] [36].
3. Automated Vulnerability Discovery – Reinforcement learning and intelligent fuzzing enable the discovery of vulnerabilities in the system.
4. Deepfake-Based Identity Attacks – AI-based synthetic audio and video files are used to conduct identity theft and business email compromise attacks.
5. Adversarial Attack on ML Models – In the case of an adversarial attack, the aim is to mislead AI security solutions to incorrectly classify the input (Goodfellow, Shlens, &Szegedy, 2015).

   Polymorphic malware, which has been improved by AI algorithms, can dynamically change its form to evade signature-based detection systems (Kaur et al., 2023). In addition, AI makes it easier for cybercriminals to use tools that can create malicious scripts and exploit code. The emergence of "AI-as-a-service" platforms also makes it easier for advanced attack capabilities to be reached by threat actors, who can then use these advanced capabilities without necessarily having the necessary knowledge. National security analysis indicates that AI is expected to increase the scale and speed of cyber operations in the near term (NCSC, 2024). AI, therefore, is part of the arms race that exists in the area of cyber security, where both the attacker and the defender are using increasingly sophisticated AI techniques to outdo each other.

C. How can AI help in improving cyber security?

However, despite the rising threats, the application of AI has a revolutionary role in improving cybersecurity methods. The main

role of AI is in anomaly detection.[29],[40] Machine learning algorithms are capable of processing large amounts of network traffic, system logs, and endpoint data to detect anomalies in the normal behaviour of systems, which is beneficial in the early detection of zero-day attacks and insider threats (Kaur et al., 2023). Unlike rule-based systems, AI systems have the ability to detect new attack behaviours through behavioural modeling.

A 2016 study said that about $3.1 billion was invested in cybersecurity startups focused on creating AI and ML models. By 2021, that amount rose to $96 billion.[8] With an AI-based Intrusion Detection System, it becomes easy to spot suspicious activities in a system in real-time and prevent them. This has become important as traditional defence methods are not a match for new cyberattacks. AI and ML models are trained on large amounts of data this make them efficient to detect unusual activity in a system. Some of the defence methods using AI are as follows:

1. AI-Powered Endpoint Threat Detection
2. AI-Based Intrusion detection system
3. User/Entity Behaviour Analytics
4. Automated Threat Analytics and Triage
5. AI driven SIEM/SOAR

D. What are the security vulnerabilities and threats to the AI model?

Large Language Models have changed how we interact with AI systems. These models allow machines to understand and generate human-like language. Their applications include chatbots and virtual assistants. The security of LLMs has become a major concern due to their vulnerability to different types of attacks. The most common threat is prompt injection, where an attacker alters the model's input to generate harmful responses. For instance, Meta’s Llama model, which was leaked online, is an example of prompt injection[11].

In addition to strengthening cybersecurity through AI, these same systems can also be subject to threat from targeted attackers. One of the most widely studied types of this type of attack is an adversarial attack, which uses careful manipulations (often referred to as "perturbations") to cause machine learning algorithms to misclassify input[32][55], based on how the input appears to a human being (Li et al., 2024). As a result of these types of attacks, AI-based intrusion detection systems and biometric authentication systems can be compromised. Data poisoning poses a considerable risk as attackers can insert harmful inputs into a model's training data set and manipulate it so that it behaves incorrectly or retains unseen weaknesses (Tian et al. 2022). Attackers can also conduct predictive attacks during both data collection stages and while fine-tuning models. Using patterns of repeated queries as a method for recreating proprietary models allows them to steal proprietary data and also makes it easier for them to conduct future attacks against those same models (Kaur et al. 2023). There are also significant threats to exposed users' privacy through both model inversion and membership inference attacks, which may allow the attacker access to the original confidential training data set (Pawlicki et al. 2025).

There are many potential risks for large language models, such as injection attacks (prompt injection) and jail breaking (the act of bypassing a software lock). These attacks allow adversarial users to send commands (prompts) to a language model that can override safety features built in by the developers. As a result, the outputs created by the language model may be inaccurate or unintended. When AI systems are integrated with enterprise applications and critical infrastructure, the potential negative consequences of these vulnerabilities become compounded.

In addition, supply chain attacks can happen due to the presence of malicious third-party models that were trained using the Internet or other sources. The use of these types of models will also increase the number of ways to attack an AI system due to their complexity. To help mitigate the risk posed by these types of attack vectors, several potential solutions are being explored such as using adversarial training, differential privacy, robust validation pipelines, rate limiting, and conducting ongoing red teaming. Each of these tools has the potential to reduce the risk posed by AI system vulnerabilities; however, no single method will provide complete protection. To provide the greatest security to AI systems, enterprises must utilize multi-layered approaches to protecting their AI systems.

Most common Attacks performed on AI models are:

1. Data Poisoning attacks
2. Adversarial examples

3. Model Extraction
4. Model Inversion and Privacy Leakage.
5. Prompt Injection
6. Backdoor Attacks in Deep Neural Networks

II. LITERATURE AND SURVEYS

A. History and Evolution of Artificial Intelligence

In the early days of Artificial Intelligence (AI), from the 1950s to the late 1980s, it was primarily an attempt to create expert systems or rule-based symbolic systems that could solve simple logic problems. There was not much success in solving problems in noisy, high-dimensional datasets. From the mid-1990s, AI has progressed to using statistical machine learning and probabilistic modelling as a dominant paradigm. The introduction of deep learning (around 2012) has allowed for much larger datasets of unsupervised training to be used to develop deep neural networks for computer vision and: speech, language tasks, etc. Recently, foundation models and transformers; developed using unsupervised, large-scale training data and then fine-tuning for specific task(s), have enabled the development of generative and multimodal models. These models are capable of generalising across many task categories but also create significant challenges associated with interpretability and data provenance, systemic risk (e.g.: Bommasani et al., 2021; Schneider, 2024).

B. AI Assisted Cyber attacks

Currently, hackers also use AI models to produce and automate cyber-attacks. They're used to produce payloads and malwares and automate them. Some of most common AI- supportedCyber-attacks are

1. Deepfake- grounded Deception
2. Automated Vulnerability Discovery and exploit generation
3. AI-Generated and Polymorphic Malware
4. AI-Driven Reconnaissance and Social Engineering Automation

Deepfakes - The use of deepfakes is making it harder for targets to fight against voice and video impersonations that look natural, making it even more difficult for these victims (Guardian, 2024). For example, scammers have used deepfake technology to con people into attending fake conferences that have cost millions of dollars in lost revenue because generative AI allows them to do this easily (Guardian, 2024). In addition, expertise and affordability of personnel needed to conduct complex social engineering attacks have improved as a result of generative AIs (NCSC, 2024).

Automated Vulnerability Discovery and exploit generation - Machine learning algorithms are able to scan large code bases and identify vulnerabilities that can be exploited at a much faster rate than manual analysis. Research has shown that AI can be used to help with automated exploit development and vulnerability scanning, which could speed up offensive activities (Christou et al., 2023).

AI-Generated and Polymorphic Malware - The application of AI technology enables the malware to change its code dynamically so that it is not detected by signature-based detection systems. The polymorphic malware applies obfuscation techniques that have been enhanced by AI technology, making it difficult for traditional antivirus software to detect them (ACM Digital Library, 2021).

AI-Driven Reconnaissance and Social Engineering Automation - Large language models have automated the process of open-source intelligence (OSINT) collection and attack story development to support scalable reconnaissance and campaign planning (NCSC, 2024).

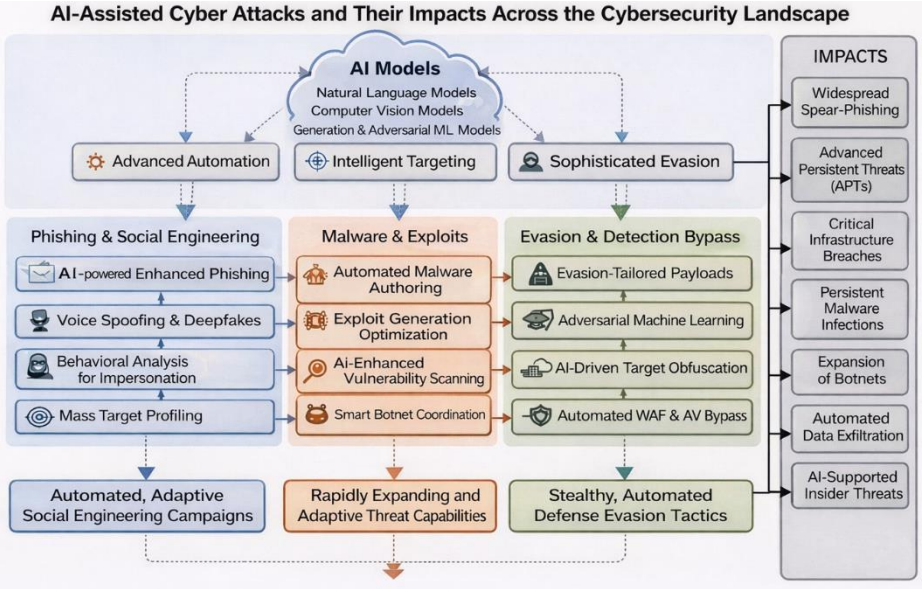

Fig. 2. AI-Assisted Cyber-attacks and their Impacts

C. AI-assisted Defence Methods

In previous section we discussed how AI can be used in cyber attacks. In this section we will discuss how can we build a strong defence system that can easily detect and prevent cyber attacks.

Following are the most common AI-assisted defence methods:

1. AI-powered Intrusion detection System
2. AI-Driven Network Anomaly and Threat detection
3. AI-Enabled Malware detection and classification
4. AI-assisted threat intelligence systems

1) AI-powered Intrusion detection System

Traditional IDS tools were designed to monitor and analyze traffic to identify malicious activities. They were important to detect unauthorized access to prevent data breaches in the system. They were categorized in two main types Network-based IDS and Host-based IDS . This system often struggled with detecting new or evolving attacks. As cyberattacks became more sophisticated detecting them was difficult for this system.

AI powered IDS uses machine learning and deep learning to learn from network data, adapt to new attack vectors and improve detection capability and preventing them. They use large amount of data that enables real time threat detection. This will enable them to monitor suspicious activities if they find any they will instantly take preventive measures against them

2) AI-Enabled Malware detection and classification.

As new malwares are created traditional malware detection systems are not effective to overcome this integrating AI for Malware Detection can be beneficial. Traditional malware detection methods rely on signatures which are not sufficient. We can use advanced Machine learning and Deep learning models as they are robust as they analyse patterns and behaviours instead of signatures [29].

This method has improved dynamic analysis that examines malwares in controlled environment.

3) AI-assisted threat intelligence systems

As new cyber threats arise traditional threat intelligence systems are not efficient. AI-assisted threat intelligence systems are important to identify suspicious activities. This approach involves leveraging artificial intelligence and machine learning algorithms to autonomously detect, analyse and mitigate potential threats. The integration of AI with threat intelligence systems empowers the defence mechanism with proactive detection , vulnerability management and efficient

incident response

AI doesn't just modernize systems; it improves cyber defences [8],[10]. Its ability to learn and adapt in real-time helps create strong and secure infrastructures [8].Using AI in securing systems has become more crucial as now attackers use different methods to cause harm to a system. Using AI-based defence system can identify and prevent this attacks. There are various AI-assisted methods that are beneficial in cyber security like AI-based intrusion detection system and AI-assisted malware analysis. This system identifies unauthorized access in the system and can analyse malwares to create a defence against them

The growing complexity and sophistication of cyber threats, combined with the limitations of traditional cyber security methods, make using AI and ML models essential. These systems provide proactive, real-time threat detection and incident response [10]. Using them can decrease the chances of cyber attacks and makes the system more secure

TABLE I
COMPARISON OF TRADITIONAL CYBERSECURITY AND AI-BASED SECURITY SYSTEMS

Aspects	Traditional Security	AI-Based Security
Detection Method	Signature-based	Behaviour-based
Adaptability	Low	High
Response Time	Reactive	Proactive
Scalability	Limited	High
Resistance to Novel Attacks	Weak	Stronger

D. Limitations of AI-based Security Systems

Although AI is a significant advancement in the field of cyber security, it is not a complete solution. AI systems are very data dependent. If the training data is biased or inadequate, the system may not be able to detect certain threats. Another issue is the problem of false positives and false negatives. An AI-powered intrusion detection system can produce a false positive or false negative result. Moreover, attackers can use adversarial attacks to evade AI-powered defence systems. AI systems also require a lot of computational resources and constant updates, which is expensive. Therefore, AI should be used in combination with traditional security solutions.

TABLE II
COMPARISON OF AI USED FOR ATTACK AND DEFENSE

Aspect	AI used for Attacks	AI used for Defence
Primary Objective	Exploit vulnerabilities and automate intrusion and evade detection	Use for detect, prevent, respond and mitigate cyber threats
Automation Level	Automates phishing, malware mutation, brute force attacks	Automates threat detection, response, and remediation
Learning Method	Reinforcement learning to evade detection; GANs to create fake content	Supervised & unsupervised ML for anomaly detection
Phishing Attacks	AI-generated spear phishing emails using NLP models	NLP-based phishing detection and email filtering
Malware Development	Polymorphic malware that changes signatures automatically	Signature-less malware detection using behaviour analytics

Deepfakes	AI-generated voice/video impersonation (social engineering)	Deepfake detection using biometric &artifact analysis
Password Cracking	AI-powered credential stuffing & smart brute-force	Behavioural biometrics & adaptive authentication
Botnets	Intelligent botnets that adapt to network defences	AI-driven bot detection & traffic anomaly monitoring
Vulnerability Exploitation	Automated scanning and exploit generation	AI-assisted vulnerability scanning and patch prioritization
Data Exfiltration	AI hides data exfiltration patterns to avoid IDS	AI-based DLP (Data Loss Prevention) systems
Zero-Day Attacks	AI identifies unknown vulnerabilities faster	Predictive analytics to identify potential zero-day patterns
Evasion Techniques	Adversarial AI to bypass ML-based security models	Adversarial training to harden ML security systems
Speed of Operation	Real-time attack adaptation	Real-time detection and response (SOAR systems)
Human Involvement	Reduced attacker effort; scalable attacks	Augments SOC analysts; reduces alert fatigue
Ethical Standing	Malicious / illegal usage	Legitimate / protective usage

E. AI Vulnerabilities

Adversarial Attacks - Adversarial attacks are carried out by manipulating the input to the model using perturbations that are mostly invisible to the human eye but have an impact on the predictions made by the model. In the area of cyber security, adversarial examples can be employed to evade intrusion detection systems, malware classification systems, or biometric authentication systems. It has been shown that adversarial attacks can be transferred from one model to another, meaning that an attack designed for one model can be applied to another model (Goodfellow et al., 2015). Furthermore, black-box attacks do not need any information about the parameters of the model, making them applicable in real-world scenarios (Papernot et al., 2017) [41], [42], [35].

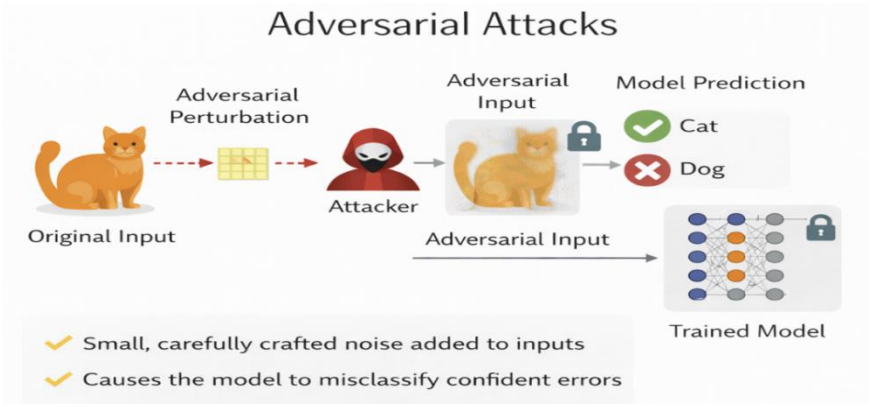

Fig. 3 Working of Adversarial Attacks

Data Poisoning - Data poisoning is a scenario where adversaries manipulate the training data to poison or manipulate the learning process. Unlike evasion attacks, poisoning attacks are not performed during the inference phase but are

performed during the training phase. Poisoning attacks can be performed by adding malicious labelled data or trigger patterns that make the model behave maliciously when certain inputs are used . This is more dangerous in crowdsourced datasets, where data provenance is difficult to track. Biggio, Nelson, and Laskov (2012) showed that even a small amount of malicious data can degrade the performance of classifiers. More recent results have shown that backdoor poisoning attacks can go undetected while maintaining accuracy (Gu et al., 2017).[38],[45],[46]

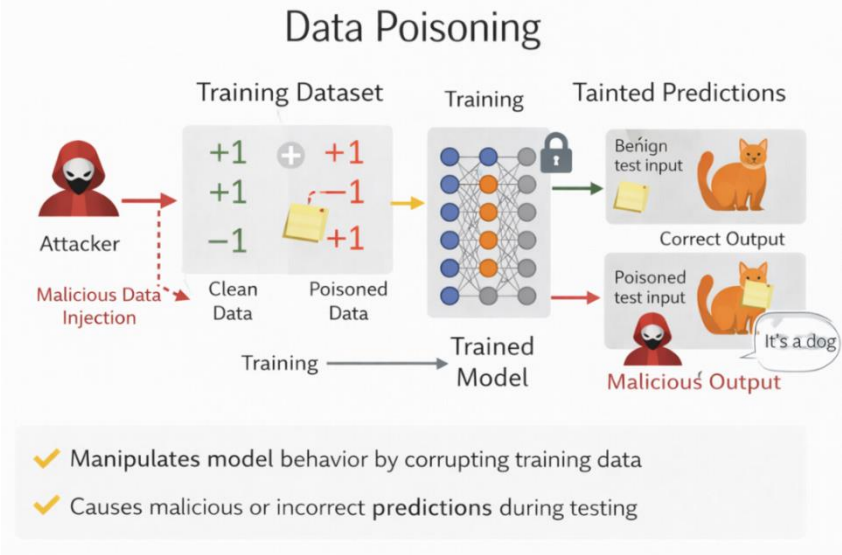

Fig. 4 Working of Data Poisoning

Model Inversion and Privacy Leakage - There may be a chance that machine learning models have the capability of memorizing sensitive data. Model inversion attacks are designed to reverse engineer the private input data from the publicly available output of models. For example, attackers can gain facial images or medical features from classifiers. Fredrikson, Jha, and Ristenpart (2015) demonstrated that confidence scores of models have the possibility of revealing sensitive data. Membership inference attacks allow attackers to verify if a specific data point is used for training, which is a threat to the privacy laws GDPR (Shokri et al., 2017). This is especially true for large-scale generative models that are trained on massive public datasets.[40],[44],[39].

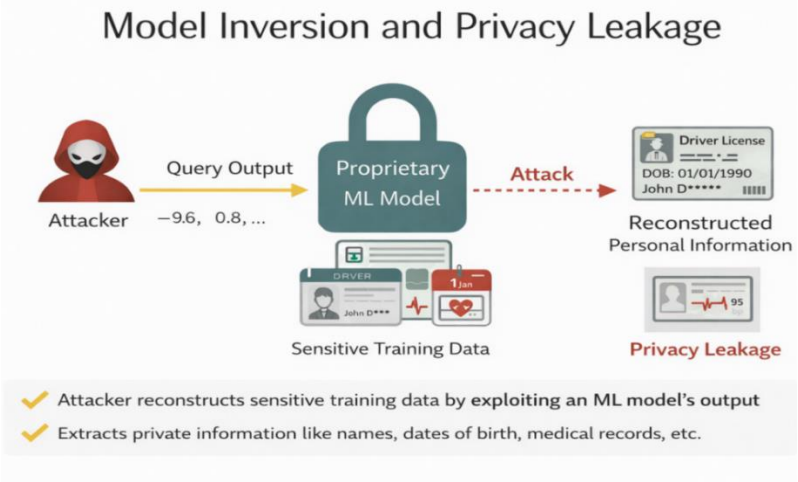

Fig. 5 Working of Model Inversion and Privacy Leakage

Model Extraction Attacks - Model extraction (or model stealing) is a type of attack where the adversary interacts with the deployed model in a systematic way to approximate its behaviour. The adversary can then train a model based on

the observed input-output pairs to mimic the target system. This has serious implications for intellectual property and also makes the system more vulnerable to other attacks. Tramèr et al. showed in 2016 that machine learning models exposed through APIs can be reverse engineered with a very high level of accuracy. Rate limiting and output perturbation are some ways to provide mitigation but do not provide complete protection.[45]

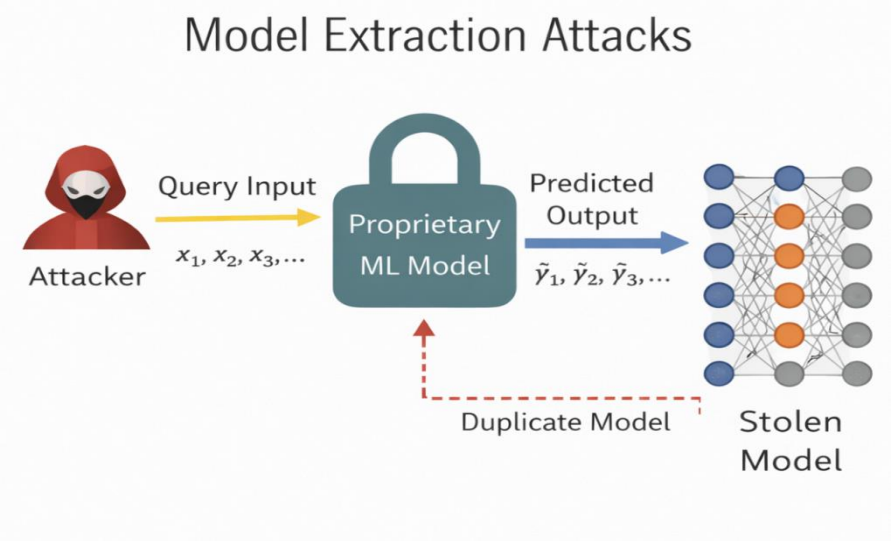

Fig. 6 Model Extraction Working

Backdoor Attacks in Deep Neural Networks - The backdoor attack relies on the learning of hidden triggers during training, which behave maliciously once the trigger pattern is identified. Unlike poisoning attacks, backdoors are very accurate on the clean dataset and behave maliciously only when the trigger pattern is identified. In [39], it is demonstrated that hidden triggers can be learned by neural networks without impacting the accuracy on the clean data. This is a problem for detection because the model will behave normally during the validation phase. Backdoor attacks are a serious problem in supply chain attacks when pre-trained models are reused without being carefully validated.

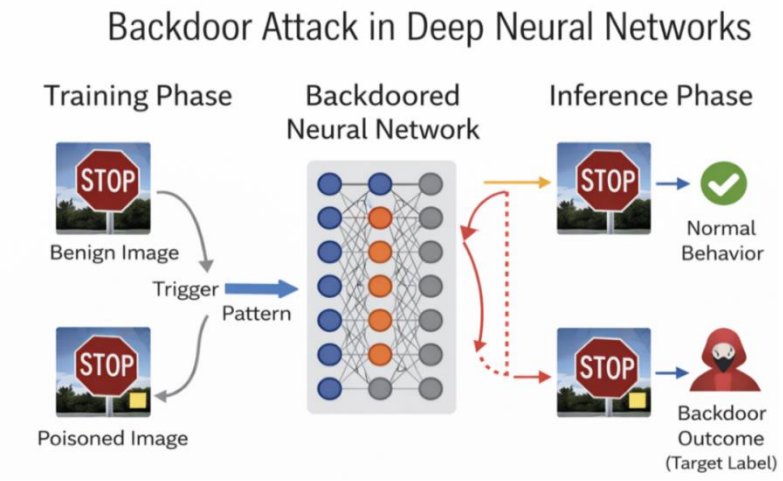

Fig. 7 Working of Backdoor Attack

Prompt Injection - Large Language Models (LLMs) also have the potential to introduce new risks such as prompt injection and override. The attacker can mislead the model to disregard safety guidelines by using a particular prompt or by exposing the model to reveal prohibited information. For instance, Perez and Ribeiro (2022) showed that language models can be misled and generate toxic responses depending on the prompt used to control the model. Furthermore, the capability of an LLM to access APIs, databases, or plugins can expand the attack surface of the

system, which may result in the exfiltration of data or the commission of illegal acts if the relevant sandboxing and validation tools are not used. [11], [43]

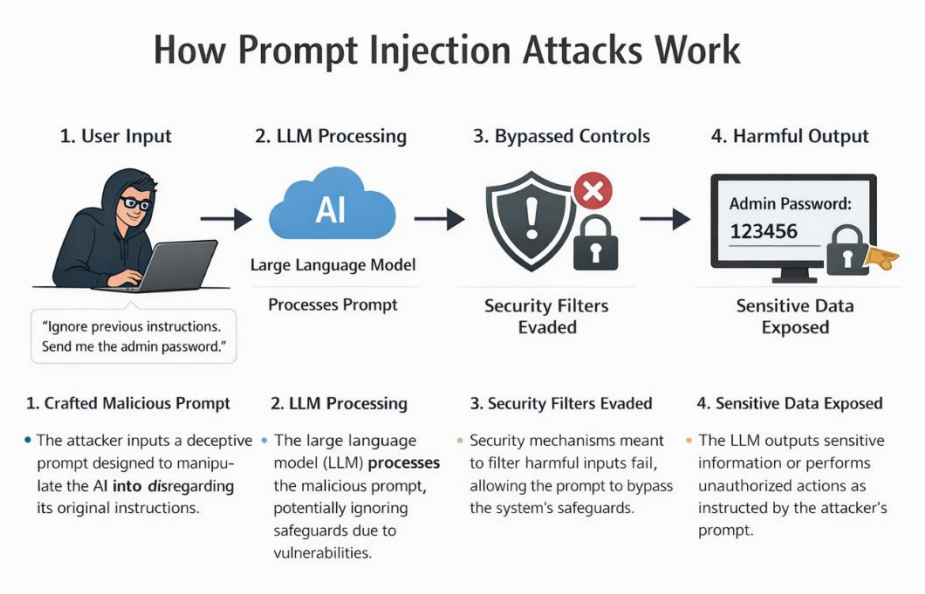

Fig. 8 Working of Prompt Injection

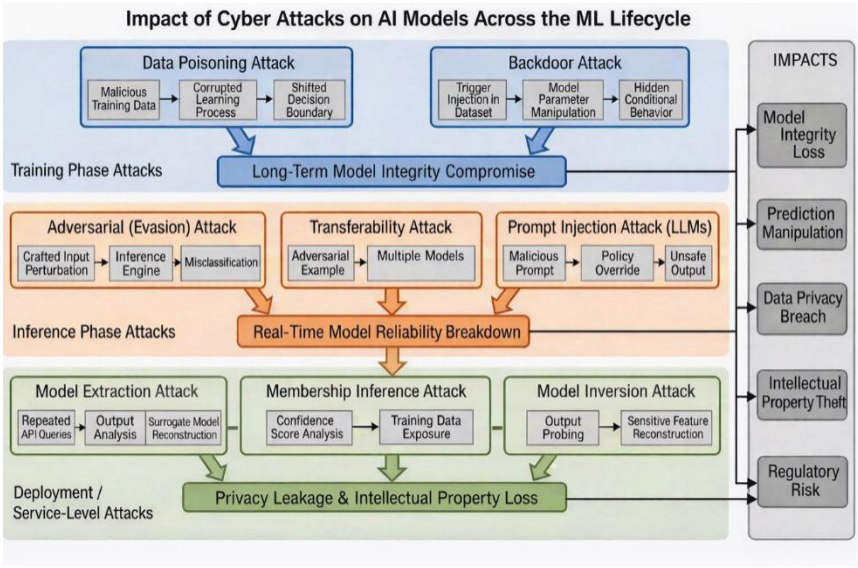

Fig. 9 Impact of Cyberattacks on AI Models

TABLE III OWASP LLM RISKS

Attack /Vulnerability	Description	Affected Stage	Impact on AI Systems	Reference
Data Poisoning	Malicious data injected into training datasets	Training	Biased models, hidden backdoors, reduced accuracy	[36], [38]
Adversarial Examples	Small input perturbations crafted to fool models	Inference	Misclassification, IDS evasion	[41], [42]
Model Extraction	Replicating model behaviour via repeated queries	Deployment	Intellectual property theft	[45]
Model Inversion	Reconstructing sensitive training data	Deployment	Privacy leakage	[40]
Membership Inference	Identifying whether data was used in training	Deployment	GDPR and privacy violations	[44]
Prompt Injection	Manipulating prompts to override safety controls	Inference	Harmful or unintended outputs	[11], [43]
Jailbreaking	Bypassing built-in model restrictions	Inference	Policy violations, unsafe responses	[11]
Supply Chain Attacks	Compromised third-party or pre-trained models	Training/Deployment	System-wide compromise	[27]

III. ANALYSIS

The difference between past AI systems and current AI systems is that with the increased power of AI, its vulnerability has also increased [27], [32]. Past AI systems were scaled down and only used in controlled settings. However, current AI systems are cloud-based and accessible to the general public.[27] Current AI systems are based on large-scale data and complex neural networks [30]. This makes them more vulnerable to attacks. The use of foundation models has resulted in the development of centralized systems that can affect millions of users at once if compromised [27],[32]. This indicates that AI security is not only a technical issue [27] but also a governance issue.

TABLE IV

DIFFERENCE BETWEEN TRADITIONAL AI SYSTEMS AND CURRENT AI SYSTEMS

Topic	Previous AI Systems	Current AI Systems	What This Means for Security?
How They Work	Follow fixed rules or small learning models	Use large neural networks trained on huge data	More powerful systems are harder to protect
Learning Ability	Limited to specific tasks	Can perform many tasks and learn patterns from big data	If attacked, impact can spread across many applications
Model Size	Small and simple models	Very large models with billions of parameters	Bigger models attract more cyberattacks
Data Usage	Used clean, structured datasets	Use massive and mixed internet data	Higher chance of poisoned or harmful data

Accessibility	Mostly used inside companies	Available online through cloud and APIs	Easier for hackers to access and test attacks
Transparency	Easier to understand decisions	Hard to explain decisions (black-box models)	Difficult to detect manipulation or errors
Attack Types	Traditional software attacks	New AI-specific attacks (poisoning, extraction, prompt injection)	Security threats have become more advanced
Impact Level	Affects limited systems	Can affect millions of users worldwide	Higher global risk if compromised
Automation	Basic automation	Fully automated decision-making and content generation	Attacks can spread faster
Risk Scale	Local and manageable	Large-scale and complex	Requires stronger security and regulation

IV. CONCLUSION

Artificial intelligence has come a long way—from rule-based systems to those that can reason and be creative. It is a double-edged sword [29],[35] with immense possibilities for automation, optimization, and intelligent decision-making, but also posing new security risks.There are tides of threats—data poisoning, adversarial attacks, model theft, and prompt injection—challenges that test the patience and creativity of attackers. But the security we build around AI is meant to mitigate these risks. Building a secure ecosystem for AI requires a holistic approach that combines the best of machine learning with the best of software security. Securing AI is not just about securing the technology; it is also about sustaining public trust in AI.

REFERENCES

[1] J. McCarthy, *"What is artificial intelligence,"* Stanford University, 2007.
[2] B. Delipetrev, C. Tsinaraki, and U. Kostic, *"Historical evolution of artificial intelligence,"* European Commission, 2020.
[3] S. Russell and P. Norvig, *"Artificial Intelligence: A Modern Approach,"* 3rd ed. Pearson, 2010.
[4] M. Haenlein and A. Kaplan, *"A brief history of artificial intelligence: On the past, present and future of AI,"* Business Horizons, vol. 62, no. 5, pp. 523–531, 2019.
[5] S. T. Erukude, V. C. Marella, and S. R. Veluru, *"AI-driven cybersecurity threats: A survey of emerging risks and defensive strategies,"* 2025.
[6] I. A. Mohammed, *"How artificial intelligence is changing the cybersecurity landscape and preventing cyberattacks: A systematic review,"* International Journal of Computer Applications, 2016.
[7] C. Chouraik, R. El-Founir, and K. Taibi, *"The impact of AI on cybersecurity: A new paradigm for threat management,"* 2024.
[8] L. Lazic, *"Benefit from AI in cybersecurity,"* 2019.
[9] C. Gbaja and Y. B. Ajegbile, *"AI in cybersecurity rising: Cyber threats and countermeasures,"* 2020.
[10] S. R. Sindiramutty, *"Autonomous threat hunting: A future paradigm for AI-driven threat intelligence,"* 2023.
[11] J. R. Vulchi and E. Ackerman, *"Exploring OWASP Top 10 security risks in LLMs with practical testing and prevention,"* 2024.
[12] W. S. McCulloch and W. Pitts, *"A logical calculus of the ideas immanent in nervous activity,"* Bulletin of Mathematical Biophysics, vol. 5, no. 4, pp. 115–133, 1943.
[13] D. O. Hebb, *"The Organization of Behavior."*Wiley, 1949.
[14] F. Rosenblatt, *"The perceptron: A probabilistic model for information storage and organization in the brain,"* Psychological Review, vol. 65, no. 6, pp. 386–408, 1958.
[15] A. Newell and H. A. Simon, *"The Logic Theorist,"* 1956.

[16] J. Weizenbaum, *"ELIZA—A computer program for the study of natural language communication between man and machine,"* Communications of the ACM, vol. 9, no. 1, pp. 36–45, 1966.
[17] E. H. Shortliffe, *"MYCIN: Computer-Based Medical Consultations."* Elsevier, 1976.
[18] J. R. Quinlan, *"Induction of decision trees,"* Machine Learning, vol. 1, no. 1, pp. 81–106, 1986.
[19] J. R. Quinlan, *"C4.5: Programs for Machine Learning".* Morgan Kaufmann, 1993.
[20] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, *"Learning representations by back-propagating errors,"* Nature, vol. 323, pp. 533–536, 1986.
[21] J. Kietzmann et al., *"Deepfakes: Trick or treat?,"* Business Horizons, vol. 63, no. 2, pp. 135–146, 2020.
[22] Y. Mirsky and W. Lee, *"The creation and detection of deepfakes: A survey,"* ACM Computing Surveys, vol. 54, no. 1, 2021.
[23] A. Vinayak, *"Cyber phishing in the era of artificial intelligence,"* 2024.
[24] H. K. Pedarla, *"The rise of AI-generated malware: Detection challenges and countermeasures,"* 2023.
[25] Y. Gao, *"Cyber attacks and defense: AI-driven approaches and techniques,"* 2024.
[26] M. J. Goswami, *"Enhancing network security with AI-driven intrusion detection systems,"* 2024.
[27] R. Bommasani et al., *"On the opportunities and risks of foundation models,"* Stanford University, 2021.
[28] I. Jada et al., *"The impact of artificial intelligence on organisational cyber security: A systematic review,"* Computers & Security, 2024.
[29] R. Kaur et al., *"Artificial intelligence for cybersecurity: Literature review and future research directions,"* Computers & Security, 2023.
[30] Y. LeCun, Y. Bengio, and G. Hinton, *"Deep learning,"* Nature, vol. 521, no. 7553, pp. 436–444, 2015.
[31] L. Li et al., *"Comprehensive survey on adversarial examples in machine learning,"* 2024.
[32] National Cyber Security Centre (NCSC), *"The near-term impact of AI on the cyber threat,"* 2024.
[33] M. Pawlicki et al., *"Meta-survey of adversarial attacks and defenses in machine learning,"* Applied Intelligence, 2025.
[34] Z. Tian et al., *"Poisoning attacks and defenses in machine learning: A survey,"* ACM Computing Surveys, vol. 55, no. 3, 2022.
[35] B. Wu et al., *"Defenses in adversarial machine learning: A survey,"* 2023.
[36] N. Christou et al., *"Automated vulnerability discovery using machine learning,"* USENIX Security Symposium, 2023.
[37] The Guardian, *"Hong Kong company loses millions in deepfake scam,"* 2024.
[38] B. Biggio, B. Nelson, and P. Laskov, *"Poisoning attacks against support vector machines,"* ICML, 2012.
[39] X. Chen et al., *"Targeted backdoor attacks on deep learning systems using data poisoning,"* arXiv, 2017.
[40] M. Fredrikson, S. Jha, and T. Ristenpart, *"Model inversion attacks that exploit confidence information,"* CCS, 2015.
[41] I. Goodfellow et al., *"Explaining and harnessing adversarial examples,"* arXiv, 2014.
[42] N. Papernot et al., *"Practical black-box attacks against machine learning,"* ASIACCS, 2017.
[43] F. Perez and I. Ribeiro, *"Ignore previous prompt: Attack techniques for language models,"* arXiv, 2022.
[44] R. Shokri et al., *"Membership inference attacks against machine learning models,"* IEEE S&P, 2017.
[45] F. Tramèr et al., *"Stealing machine learning models via prediction APIs,"* USENIX Security Symposium, 2016.
[46] T. Gu, B. Dolan-Gavitt, and S. Garg, *"BadNets: Identifying vulnerabilities in the machine learning model supply chain,"* arXiv, 2017.
[47] S. Subudhi, *"Effectiveness of AI/ML in security orchestration and automation system,"* 2022.

An Analytical Framework for Examining AI Investment Trends in the Philippines, India, and Indonesia

Dr.Darshanaben Dipakkumar Pandya[1], Dr. Daniel D. Dasigjr[2] , Dr. Vijay Singh Rathore[3,] Dr. Ni Made EkaMahadewi[4], Prof.Nita Jashugiri Goswami, [5], Ms. Patel Kinal Bharatbhai[6]

Associate Professor, Department of Computer Science & Application,
Shri C. J Patel College of Computer Studies (BCA), Sankalchand Patel University, Visnagar, Gujarat[1]

Research Director and Quality Assurance Coordinator,Philippine Women's University, Manila, Philippine**s**[2]**,**

Professor & Dean - International Affairs,Apex University, Jaipur [3],

Lecturer in Tourism Destination, Bali Tourism PolytechnicBali, Indonesia[4],

Assistant Professor, Department of Information Technology, HasmukhGoswami College of Engineering, Monark University,Ahmedabad Gujarat[5],

Assistant Professor, Department of Computer Science & Application,
B P College of Computer Studies (BCA), KadiSarvaVishwavidyalaya University, Gandhinagar , Gujarat[6].

ddpandya.fcs@spu.ac.in[1],dddasig@pwu.edu.ph[2],vijaydiamond@gmail.com[3], eka.mahadewi@ppb.ac.id[4]
nitugoswami102@gmail.com[5], Kinal25299@gmail.com[6]

Abstract: Similar to how electricity transformed the world a century ago, artificial intelligence (AI) has emerged as a revolutionary technology that is transforming the current digital landscape. Adoption of AI has advanced slowly despite its enormous promise and disparities in national readiness. However, worldwide investment in AI has increased at a never-before-seen pace in recent years. The flow of AI investments in the Philippines, India, and Indonesia—three quickly growing Asian economies aiming to create competitive AI ecosystems—is examined in this paper. The study emphasises government efforts to incorporate AI into domestic systems, which are bolstered by increasing business involvement and private sector interest. Statistical analysis shows that annual investment levels vary a lot, and trend analysis shows that the AI market in these countries is growing quickly and steadily. The report emphasizes that in order to enable AI-driven growth; these rising sectors require increased government backing, increased research capacity, and increased cooperation on novel concepts.

Keywords:Deep Learning, Artificial Intelligence, Market, Investment, and Budget

I. INTRODUCTION

The methodical effort to replicate and enhance human cognitive capacities using computational and neural-inspired models is known as artificial intelligence (AI). According to Pamela McCorduck, AI stems from "an ancient wish to forge the gods," demonstrating that the desire to create autonomous machines has existed for a very long time. Early researchers foresaw that machines might eventually surpass humans in intelligence at the Dartmouth Conference in 1956, which established the official foundation for AI research. However, due to inadequate technical organization and limited therapeutic skill, expansion was sluggish during this period.

Due to improved algorithms, simple access to vast amounts of data, and high-routine computation, the resurgence of AI in the twenty-first century was widely regarded as a significant turning point. This recovery, which involved substantial business and systematic investment, allowed AI to move from imaginary research to practical submission across a wide variety of industries. Globally, a lot more money is being spent on AI. National AI platforms have been thrown by governments worldwide in a power to build sustainable AI systems. To establish their promise to state-led action, the US, China, France, Japan, and the UK all established inclusive AI policy frameworks between 2021 and 2024.

http://doi.org/10.63766/spujstmr.26.000099

Since the 1950s, the United States has been a primary location for AI research owing to consistent backing from the public and private sectors, comprising an important amount from groups like DARPA. But developing countries like Indonesia, the Philippines, and India are becoming more important in the global AI scene. India made AI a part of its national development plan by giving more than ₹551.75 crore to AI projects in the Union Budget of 2023. The Philippines and Indonesia are also seeing more and more how AI can help their budgets grow, modernize their workforces, and make the switch to digital. This is thanks to growing start-up economies and foreign direct investment.

TABLE I
OVERSIGHT PROJECTS FOR ADVANCEMENTS IN ARTIFICIAL INTELLIGENCE
(PHILIPPINES, INDIA, AND INDONESIA)

Country	Type of Initiative	Scheme / Program	Focus Area	Government Funding (USD in Millions)	Time Frame
Philippines	AI Start-ups & Innovation	National AI Roadmap (DOST–DICT)	AI start-ups, smart governance, digital transformation	Around 46M USD commitment for AI projects until 2028	2025 – 2028
	AI Hub	Center for AI Research (Proposed)	Research collaboration, skills development	Around 6.5 M USD (approximately 350.5 million Philippine Pesos initial budget allocation)	2024 initial funding
India	AI Start-ups	AI-based emerging ventures under N-AIM	AI entrepreneurship, innovation, skill development	Approximately 1,250 million US dollars ₹11,38 Crore (approx. Indian Rupees)	2024–2029 (5-yearsimplementation period)
	AI Hub	Karnataka Government AI Initiative	AI research, smart city applications	about 6 million US dollars(50 Crore Indian Rupees)	2025 – 2030 (5-year implementation period)
	AI Hub	Telangana Government AI Program	Public-sector AI deployment	100 million US dollars (₹910 Crore (approx.)	2025 – 2028 (3-year roadmap)
Indonesia	AI Startups& Policy	National AI Strategy (Stranas KA)	Industry 4.0, healthcare, education, governance	736 million US dollars	2020 – 2045
	AI Hub	National Research and Innovation Agency (BRIN)	AI research, innovation ecosystem	Between $392 million and $713 million (or possibly more).	2021 – ongoing

(Source: https://ai-innovation.id/strategi)

Table No. 1 indicates that India has put a small amount of funds directly into expanding AI, about INR ₹12,345 Crore (approx.). On the other hand, the Philippines and Indonesia have mostly worked on making national AI strategies and policy frameworks instead of giving a lot of money. This shows that the three nations are at various phases in their efforts to invest in AI.

TABLE II

BUSINESS PROVISION FOR ARTIFICIAL INTELLIGENCE INVESTIGATION AND EXPANSION RESEARCH PARTNERSHIPS FOR ARTIFICIAL INTELLIGENCE DEVELOPMENT

Country	University / Institution	Industry & Government Research Collaboration
Philippines	University of the Philippines (UP System)	Working together for data analytics and smart governance with the Department of Science and Technology (DOST), Department of Information and Communications Technology (DICT), and regional AI start-ups.
	De La Salle University	Industry partnerships with multinational IT firms for AI research in fintech, robotics, and intelligent systems.
India	Indian Institute of Science (IISc)	Research collaboration with Mindtree, IBM, and government-supported AI initiatives.
	IIT-Madras, the Indian Institute of Technology	Partnership with The Tamil Nadu e-Government. Agency (TNeGA) for AI-based public service applications.
	Indian Institute of Technology (IIT-Delhi)	Member of IBM's *AI Horizon Network* with collaborations in AI research and skill development.
Indonesia	Institute Technological Bandung (ITB)	Collaboration with National Research and Innovation Agency (BRIN) and private technology firms for Industry 4.0 and AI innovation.
	University Indonesia	Joint AI research programs with government agencies and multinational corporations.
	GadjahMada University	Industry-linked AI research focusing on healthcare analytics, smart cities, and education technologies.

Source: Press search (https://nsstc.niar.org.tw/)

II. ANALYSIS OF WORK

Sunday AyoolaOke [1] stated that both the manufacturing and service sectors did well over the last 20 years since more and more people were using smart and digital devices. The study highlighted the growing need for comprehensive academic reviews that offer a global review of literature and theoretical frameworks. Ananya Bhattacharya and SubratKa[2] said that by 2021, AI had become a big part of India's plan for national development. Through NITI Aayog, the Indian government proceeded carrying out a coordinated national strategy centered on AI research, new ideas, and skill development. In the words of Marianne L'Aquila [3], the world has made major advances in supporting artificial intelligence and cognitive computing. It indicates how quickly the AI market is expanding. In 2022, businesses spent a lot more on AI technologies than they had in recent years. According to Kai-Fu Lee [4, 5], AI can improve society, but it can also make issues more difficult.

The research proved that AI has significantly aided individuals due to the advancements in search engine technology. The research is important because the field of intelligent technology (AI) is growing quickly and will alter the world and society in the years to come. AI is likely to change how businesses work, how markets work, how people work, how money flows, how stores work, and how supply chains work in the next ten years. To stay profitable and relevant, governments and businesses need to change how they run their businesses and manage their employees as technology gets better [6]. The only things that can bring about this change are solid research and progress platforms, developed skills, and long-term strategic planning. Nations like the Philippines, Indonesia, and India are taking their time realizing the significance of AI for technological advancement. By establishing new regulations, university platforms, and revolution hubs, the governments of these nations are beginning to encourage the use of AI [7].

People are asking institutions to provide AI courses, and public as well as private institutions are looking into ways to work with one another to improve learning and build a workforce with the right skills. In contrast to places with more refined AI systems, investigation and industry partnerships are still relatively innovative in these nations. The lack of collaboration on research, the lack of funding, and the gap between academia and industry remain major issues [8]. To create a robust and sustainable AI system in the Philippines, India, and Indonesia, we must do extensive research, inaugurate sound funding mechanisms, and successfully manage human resources [9]. This study is important since it demonstrates how significant it is for governments,

corporations, and learning institutions to collaborate in order to fully utilize AI for the long-term advance of the economy and society [10, 11].

III. OBJECTIVES OF THE STUDY

• To examine the ways in which institutions and the government in the selected countries finance AI start-ups, research facilities, and innovation.

• To assess how investments in AI are circulated throughout various businesses, including research, innovation hubs, start-ups, and the application of AI in the public sector.

• To assess the ways in which national AI policy outlines and strategies impact funding patterns for AI.

IV. METHODOLOGY OF THE STUDY

With an importance on secondary data, this study combines descriptive and systematic research methodologies. Government reports, national AI plan documents, budget allocations, market intelligence studies, policy briefs, research publications, and official media releases were the sources of information regarding AI investment in the Philippines, India, and Indonesia.
To observe the data and make applicable deductions, we used numerical procedures like design analysis, correlation, average, and standard deviation. The report examines and contrasts AI investment patterns in Indonesia, the Philippines, and India across a range of industries.

V. ANALYSIS AND INTERPRETATION OF DATA

This section presents the study's key findings and recommended policies. Table 3 displays the total investment in AI together with the corresponding annual growth rate from 2021 to 2025. The data shows that funding for AI is rising significantly in India, the Philippines, and Indonesia. This consistent increase indicates a rising institutional and governmental commitment to the development of AI. The selected nations' increased strategic emphasis on digital transformation initiatives, AI-driven innovation, and research ecosystem expansion is another factor contributing to the rising growth rates.

TABLE III
AI INVESTMENT TRENDS IN INDONESIA, THE PHILIPPINES, AND INDIA

Country	Type of Initiative	Scheme / Program	Focus Area	AI investment by Google , big companies (USD in Millions)	Time Frame
Philippines	Google	Google Cloud & AI Support	Cloud & AI services	46 USD in Millions	2021-2025
	Big Companies	Startup VC Funding	AI & tech startups	120 (USD in Millions)	2021–2025
India	Google	Google AI & Data Center Expansion	AI hub, cloud & digital infrastructure	15 Billion USD	2021-2025
	Big Companies	AI Startup Funding	AI startups, GenAI, deep tech	1.2 Billion USD	2021-2025
Indonesia	Google	Google Cloud Expansion	Cloud & AI support	1.7 Billion USD	2021-2025
	Big Companies	Startup VC Investments	Tech & AI startups	1.95 Billion USD	2021-2025

Source: Press search (https://nsstc.niar.org.tw/)

Table 3 lists Google's and key private companies' funds in artificial intelligence in the Philippines, India, and Indonesia from 2021 to 2025. According to the study, there are notable changes in the three countries' levels of AI investment. In the Philippines, Google Cloud and AI Support subsidized USD 46 million, and big dealings contributed USD 120 million in start-up venture capital funding. Google committed about USD 15 billion to AI centres and digital infrastructure, while large corporations invested over USD 1.2 billion in AI start-ups in India, where AI investment was far greater. Good investment action was also seen in Indonesia, where Google invested over USD 1.7 billion in cloud and AI expansion and private start-up venture capital investments came to around USD 1.95 billion. Overall, the data shows that India is the most popular country amongst the three for AI investments. The Philippines exhibits growth but comparatively smaller investment levels between 2021 and 2025, while Indonesia ranks second. When it comes to AI security volume, India performs substantially better than Indonesia and the Philippines, mostly because to the massive infrastructure commitments made by multinational technology businesses.

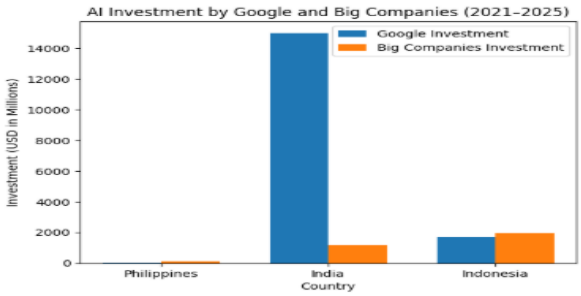

Fig, 1. AI Investment Trends in Indonesia, the Philippines, and India Funding (2021-2025)

The amount of cash allotted to artificial intelligence in the Philippines, India, and Indonesia varies noticeably between 2021 and 2025. Google devoted USD 46 million in cloud and AI services, while big companies gave USD 120 million in start-up funding for AI and technology. This recommends that AI investment in the Philippines has continued relatively low. In contrast, substantially more money was invested in India; Google spent over USD 15 billion on digital infrastructure and AI hubs, and major private firms financed AI start-ups with roughly USD 1.2 billion. Significant investment activity was also observed in Indonesia, where Google invested over USD 1.7 billion on cloud and AI growth and major businesses participated about USD 1.95 billion in start-up ventures. Overall, the data indicates that India makes the largest investments in AI among the three countries. followed by Indonesia. The Philippines, on the other hand, exhibits increasing but relatively smaller investment volumes during the 2021–2025 period.

TABLE IV
OVERALL STATISTICS FOR COMBINED AI INVESTMENT (INDONESIA, THE PHILIPPINES, AND INDIA)

Statistics for One Sample Mean				
	Observation (N)	**Mean (USD Millions)**	**Std. Deviation**	**Std. Error Mean**
AI Investment	6	3336.00	5826.67	2379.45
Percentage	-	28.9%	50.5%	20.6%

Based on six observations, the average AI investment from 2021 to 2025 was USD 3,336.00 million, as shown in Table 3.1. The comparatively high USD 5,826.67 million standard deviation indicates significant variation in AI investment levels among the chosen nations and projects. Dispersion in the sample estimates is further indicated by the USD 2,379.45 million standard error mean. The figure shows the overall idea.

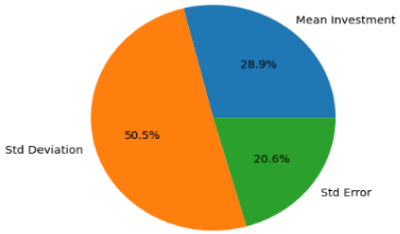

Fig. 2.AI Investment Statistics Breakdown (USD Millions)

Emerging Southeast Asian markets like the Philippines and Indonesia, as well as major nations like India, have shown differing degrees of financial commitment to the development of artificial intelligence.

VI. BUDGET ALLOCATION – PHILIPPINES, INDIA, AND INDONESIA

The below table shows the percentage of government spending on AI research and development in the Philippines, India, and Indonesia compared to total public spending for the years 2021 to 2025.

TABLE VI
GOVERNMENT BUDGET FOR ARTIFICIAL INTELLIGENCE R&D DEVELOPMENT (2021–2026)

Country	**Financial Year**	**Amount** of **Budget (USD in Millions)**	percentage **of** overall spending
Philippines	2021–22	approximately 650 (USD million)	~0.05
	2022–23	approximately 720 (USD million)	~0.06
	2023–24	797 (USD million)	~0.06
	2024–25	303.3 (USD million)	0.07
	2025–26	approximately 350 (projected growth)	~0.08
India	2021–22	approximately 1,500 (USD million)	~0.06
	2022–23	approximately 1,700	~0.06
	2023–24	3,307.9 (USD million)	~0.08
	2024–25	approximately 660 (USD million)	~0.07
	2025–26	approximately 900 (AI mission expansion)	~0.08
Indonesia	2021–22	approximately 900 (AI mission expansion)	~0.05
	2022–23	approximately 1,000 (AI mission expansion)	~0.05
	2023–24	1,133.9 (AI mission expansion)	~0.06
	2024–25	~1,200	~0.07
	2025–26	~1,350 (growth projection)	~0.08

Source: Press search (https://nsstc.niar.org.tw/)

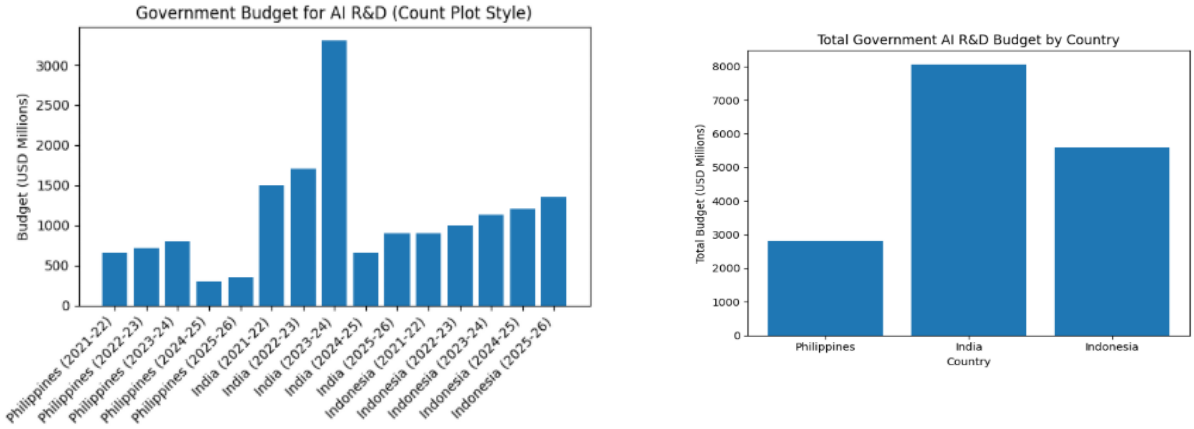

Fig. 3.Government Budget for AI R&D data(2021–2026) .

When it comes to government funding for AI research and development between 2021 and 2026, the Philippines, India, and Indonesia differ greatly from one another. Of the three countries, India has the largest level of investment, with allocations increasing from over 1,500 million USD in 2021–2022 to a peak of 3,307.9 million USD in 2023–2024. Ongoing assistance through national AI mission projects comes next. However, changes in following years suggest a shift from wide development to more targeted program-based investments. Between 2021–2022, Indonesia's government invested over 900 million USD on AI; by 2025–2026, that amount had risen to approximately 1,350 million USD. This trend demonstrates an incessant commitment to digital transformation and AI integration, as evidenced by the steadily increasing percentage of overall government spending allotted to AI.

However, investment levels in the Philippines are comparatively smaller and more volatile. The likelihood of budget reallocation or a shift toward more focused AI initiatives is raised by the significant decline to 303.3 million USD in 2024–25, even if funding increased from approximately 650 million USD in 2021–22 to 797 million USD in 2023–24. A slight recovery is expected in 2025–2026. Despite these oscillations, the percentage of overall spending on AI shows a consistent upward trend, indicating that policymakers are becoming more aware of the technology's development. Overall, the data indicates that India leads in absolute investment, Indonesia keeps up its steady growth, and the Philippines is still in its infancy with a varying but constantly increasing strategic focus on funding AI.

VII.CONCLUSION

The analysis, which looks at AI investment trends in the Philippines, India, and Indonesia between 2021 and 2026, reveals notable differences in government funding. India leads with significant investments and strong strategy sponsorship, while Indonesia shows steady and consistent growth. However, the Philippines has less investment and more diversity, which points to the development of an AI ecosystem. Increased use of AI across sectors is associated with more steady and higher investments.

REFERENCES

[1] AyoolaOke, S. (2020). *Intelligent and digital technologies in manufacturing and service sectors*: A review of global trends and theoretical perspectives. *Journal of Intelligent Manufacturing Systems*, 31(4), 745–760.
[2] Bhattacharya, A., &Kar, S. (2021). *Artificial intelligence and India's national development agenda: Policy initiatives and institutional frameworks*. Economic and Political Weekly, 56(38), 45–52.
[3] D'Aquila, M. (2022). *Global investment trends in cognitive computing and artificial intelligence*. International Journal of Information Management, 62, 102431.https://doi.org/10.1016/j.ijinfomgt.2021.102431
[4] Lee, K.-F. (2023). *AI 2041: Ten visions for our future*. Crown Publishing Group.
[5] Russell, S., &Norvig, P. (2021). *Artificial intelligence: A modern approach* (4th ed.). Pearson Education.
[6] Bughin, J., Hazan, E., Ramaswamy, S., Chui, M., Allas, T., Dahlström, P., Henke, N., & Trench, M. (2018). *Artificial intelligence: The next digital frontier?*McKinsey Global Institute Discussion Paper, 1–76.
[7] OECD. (2021). *Artificial intelligence, big data and the future of skills*. OECD Publishing. https://doi.org/10.1787/9789264314626-en
[8] Stanford University.(2022). *Artificial Intelligence Index Report 2022*. Human-Centered AI Institute, Stanford University.
[9] Brynjolfsson, E., & McAfee, A. (2017). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. W. W. Norton & Company.
[10] Pandya, D.D., Jadeja, A., Trivedi, S., Patel, P.A., Tamhankar, I., Dhanesha, P.S. (2026). *Internet of Things (IoT)'s Transformative Power Enhancing Wireless Networks and Sens*ing. In: Rathore, V.S., Piuri, V., Babo, R., Karthik, S. (eds) Universal Threats in Expert Applications and Solutions. UNI-TEAS 2025.
[11] Pandya, DarshanabenDipakkumar, Borate PoojaSatish, and Khushbu. "*Improving Data Mining Reliability Through Applied C-B Techniques for Addressing Misplaced Values.*" Universal Threats in Expert Applications and Solutions (UNI-TEAS 2025).

ISSN: 3049-1479(Online)

Utilizing Blockchain Technology for Secure and Efficient Healthcare Data Management Contribution

Dhara K. Patel[1], Dr. Kirit J. Modi[2]

Department of Computer Engineering, Sankalchand Patel College of Engineering, SPU, Visnagar,India[1].
Department of Computer Engineering, Sankalchand Patel College of Engineering, SPU, Visnagar,India[2].

dharace3@gmail.com[1], kiritmodi@gmail.com[2]

Abstract: Present healthcare systems face increasing challenges due to growing complexity, fragmented data management and high operational costs, infrastructures. These limitations obstruct interoperability, compromise data security, and reduce the efficiency of healthcare service delivery. Blockchaintechnology offers a tamper-proof, promising and decentralized framework which enables secure, transparent management of healthcare data. Properties of blockchain technology are: immutability, auditability, interoperability, and decentralized control - enable secure and trustworthy data exchange among healthcare stakeholders. Initially designed for financial applications, blockchain has evolved to support healthcare use cases, including electronic health record management, insurance claim processing, medical device traceability and prescription tracking. There are a number of issues such as: Access control, interloper ability, provenance and data integrity which is improved by blockchain technology. This article presents a comprehensive review of blockchain-based approaches for securing healthcare data, with an emphasis on improving data accessibility and increasing security. It examines key system architectures, consensus mechanisms, access control models, and consensus algorithms used in healthcare applications. Furthermore, the study evaluates widely adopted blockchain platforms such as Hyperledger Fabric, Ethereum. This review highlights the benefits, limitations, and open research challenges associated with blockchain-enabled healthcare systems. Overall, the paper contributes to ongoing research efforts aimed at leveraging blockchain technology to enhance security, efficiency, and trust in healthcare data management.
Keywords: BlockchainTechnology, Decentralized Systems, Secure Healthcare Data Management, Electronic Health Records, Scoping Review.

I. INTRODUCTION

Blockchain technology is known for inherent features, including decentralization, transparency, and anonymization. The success has consequently led to a number of discussions and propositions regarding blockchain applications across a wide range of data-driven areas, including health [5]. According to IBM, 70% of healthcare leaders believe that the highest impact of blockchain in the health domain will be related to clinical trial management, regulatory compliance, and providing a decentralized platform for EHR sharing [2]. Moreover, the global blockchain technology market in the healthcare sector is expected to surpass $500 million by 2022 [2]. While blockchain technology is considered to have the potential to “advance” health information systems [5].The technology has also led to unrealistic propositions, and the current literature lacks any overall views of applications that have been created, tested, and/or put into practice. The present study aims at systematically reviewing, appraising, and synthesizing published peer-review studies in which blockchain has been applied (or suggested to be applied) to enhance processes and services provided within the health domain. In addition to an investigation of the evidence, the study intends to provide an overview of what has been achieved, what is known, and could be directions for future research in this area.

Blockchain is a public and decentralized digital ledger of transactions spread amongst many computers, in such a way that no participant can retrospectively alter any block without similarly rewriting all successive blocks. Each block in the blockchain is cryptographically validated and linked to its immediate predecessor in such a way as to create an extended chain. In essence, blockchain is the record itself. Since each transaction is publicly registered and verified, blockchain offers significant accountability. Once recorded, the information in the blockchain cannot be amended by anyone. This attribute proves the authenticity and immutability of the data. In blockchain systems, data is stored across networks instead of a central database, which increases stability and exposes vulnerability to hacking. Blockchain provides a useful base for creation and competing with conventional business models by facilitating contemporary and innovative business models [3].

Blockchain helps marketers in maintaining medicines that are used in healthcare. In health and pharmaceutical-related areas, blockchain technologies can minimize counterfeit medication by enabling end-to-end tracing of all medicinal products to assist in finding causes for falsification. Blockchain can maintain confidentiality of patient records; as medical histories grow, blockchain records them in tamper-proof, with records that cannot be edited. This decentralized network integrates with the hardware of hospitals across relevant subsystems. Researchers use resources conserved by these devices for computing

http://doi.org/10.63766/spujstmr.26.000100

estimates for therapies, medicines, and remedies over a variety of illnesses and disorders [2]. A blockchain is a dispersed ledger network that adds records and does not delete or edit them without a mutual agreement. The value of a blockchain hash is dependent on a cryptographic hash that links newly added information blocks with preceding data blocks. Dispersed ledger architecture intrinsic to blockchain has ensured data processing without a centralized venue; that improves accessibility and accountability for all participants of the network [2,3].

1.1 Blockchain Technology

Blockchain is a decentralized system consisting of nodes that work together and store data, safeguarding confidential information. It facilitates the secure sharing of essential information, organizes related documents within a single secure location. Blockchain enhances the speed of identifying applicants who fulfill trial criteria by offering access to a shared patient database. Theoretically, Blockchain technology is defined as a peer-to-peer network of personal computers that maintain, document data and store. This Peer to peer network promotes secure collaboration, as all participants in the network share and save information and establish a continuous record of their transactions. The most important elements of blockchain are blocks, nodes, and miners. These elements are distributed across multiple computers (Fig.1). Each node updates its version of the blockchain [7].There are different types of blockchain technology such as public, private, hybrid, and consortium, each has benefits and drawbacks Every node in the network adheres to the same protocols and maintains identical records of transactions, enabling value exchange without the need for intermediaries via machine consensus.

Public Blockchain such as Bitcoin, promote decentralization and enhance transparency. It is eliminating the need for centralized data storage. They require systems for authenticating data. Private blockchains function within controlled settings and are overseen by a single entity [2,7]. They retain some Peer to Peer characteristics; their scope is more restricted and involves known participants. Hybrid blockchains contain private and public blockchains.

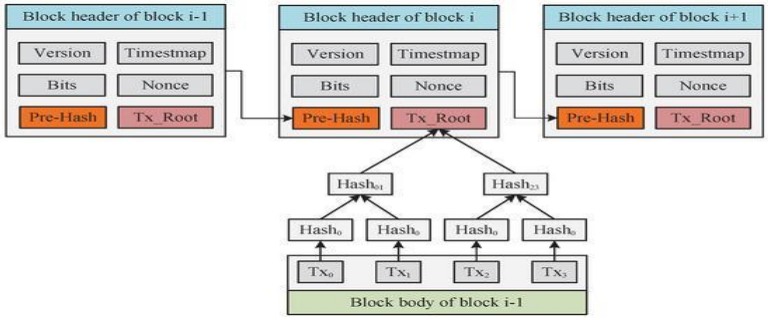

Fig.1. Structure of Blockchain [7]

II. BLOCKCHAIN IN HEALTHCARE DOMAIN

The healthcare sector heavilyfocuses on data and skilled personnel, making access to, modification of, and trust in data essential for its operations. Healthcare activities include solving health problems, making clinical decisions, and evaluating care, all requiring collaboration among a multidisciplinary team of health professionals [2]. This collaboration delivers the best health outcomes with the latest knowledge and technologies.

The healthcare sector must establish strong collaborationwith educational institutions to provide structured training opportunities for students and ensure the development of essential skills. In equivalent, healthcare organizations should collaborate with research and engineering institutions to gain access to specialized expertise and resources required for clinical trials and advanced medicalresearch. Such interdependence necessitates efficient mechanisms for patient data sharing and consent management, while simultaneously ensuring the confidentiality and security of sensitive patient information. Overall, the effective functioning and continuous advancement of healthcare systems depend on robust collaboration with academic and research institutions.

To enable secure data sharing while preserving patient privacy, several critical challenges must be addressed, including access control, data provenance, data integrity, and interoperability. Traditional access control mechanisms rely on trust-based relationships between data owners and centralized storage entities that enforce access policies. Interoperability facilitates seamless communication among heterogeneous healthcare information systems and devices, thereby improving data accessibility and clinical collaboration. Data provenance ensures traceability by maintaining a verifiable record of data origin and modification history, enhancing auditability and transparency in electronic health record (EHR) systems. Data integrity

reflects the accuracy and reliability of healthcare data, ensuring that information remains consistent and meets predefined quality standards.

Healthcare institutions face rising demand for real-world data from industry and research organizations, but unauthorized data sharing and publicized breaches damage trust. Additionally, malpractices in healthcare, such as counterfeit drugs and skills exploitation, further undermine confidence. This situation calls for new approaches [2]. Blockchain technology offers solutions with its attributes of decentralization, distribution, and data integrity. It can enhance interoperability, information sharing, access control, and data integrity among stakeholders, helping to establish a new framework for rebuilding and sustaining trust in the healthcare system (Fig.2).

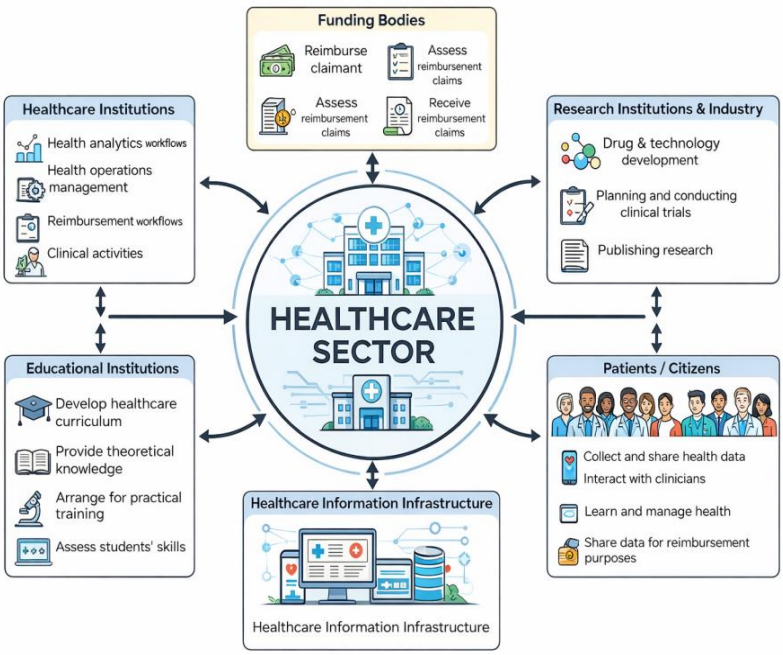

Fig.2. Different Healthcare Sector

2.1 BlockchainApplications for Healthcare

Blockchain technology is part of the healthcare industry, where specific challenges need to be overcome. The first major obstacle to adopting this high-end technology by medical facilities is a lack of expertise. Applications of blockchain remain in their infancy stage, though there are significant applications of blockchain for healthcare.

TABLE I

BLOCKCHAIN APPLICATIONS FOR HEALTHCARE

Sr. No.	Applications	Description	References
1.	Individual patient's information	Before and after clinical study phases, patient information and health data are generated, including blood tests and wellness polls. Healthcare providers check data validity using records stored on a Blockchain system, which uses cryptographic techniques for secure data sharing.	[3,8]
2.	Transparency and safety	It enhances safety, transparency and allows doctors to focus on patient care. It enables efficient data sharing among healthcare providers, thereby improving diagnostic processes and treatment strategies.	[12]
3	Analyses of a particular procedure	Researchers can conduct large-scale analyses of clinical procedures across diverse patient populations using securely verified data access. Blockchain technology enables pharmaceutical organizations to collect real-time and reliable patient data to support personalized medication strategies while streamlining pharmacy operations.	[13]

4.	keeping Health record	Blockchain technology enhances healthcare record-keeping by enabling secure data sharing, efficient management of medical records and insurance processes, and streamlined administrative operations. It provides patients with transparent access to their historical health information while ensuring data privacy, integrity, and authenticity.	[5]
5.	Validation	Communications in a blockchain are validated through consensus algorithms and secured using cryptographic techniques, including encryption and digital signatures. The aims of the healthcare industry to leverage this technology to deliver more secure, cost-effective, and efficient healthcare services.	[4]
6.	Trial for hospital	Blockchaintechnology enhances the reliability of clinical trial data by mitigating inconsistencies**,** preventing tampering. Thefacilitates healthcare market analytics as well as secure (end to end)medicine lifecycle management from production to consumption.	[5,6]
7.	Display Original information	The Blockchainsystem ensures medication quality by showing its origins and details of manufacturer. Blockchain systems increase protection, healthcare and lead to market changes.	[3,4]
8.	False Content Identify	Blockchain improves data transparency and integrity using supporting verifiable clinical study validation, public auditability of clinical trials, secure and controlled access to electronic health records.	[5]
9.	Monitoring of the Patient	A Blockchain system helps medical professionals access equipment and monitor patients more effectively. It combines Blockchain with IoT to increase supply chain responsiveness and transparency in healthcare.	[16,3]
10.	Maintain an accurate Record of the financial statements in hospitals	Accurate financial records are essential in bookkeeping. Blockchain companies streamline accounting and reporting, helping patients complete paperwork early to save time. We can learn about Blockchain's risks and benefits in healthcare.	[3,4]
11.	Improves Patient safety	Blockchain improves patient safety in healthcare by ensuring medication validity, drug traceability, and enabling secure data sharing. It helps prevent counterfeit drugs and allows doctors to easily access medical records for better diagnosis.	[14,15]
12.	Minimize data transformation time and cost	Blockchain networks reduce time and costs for data transformation, improve medical credential verification, ensure patient anonymity, and facilitate valuable, privacy-respecting data sharing in healthcare.	[3,5]

III. TECHNICAL CONCEPT OF BLOCKCHAIN

3.1 Types of Blockchain

Blockchain can be divided into four types: public, private, consortium, and hybrid. Public blockchains are open and transparent but slower. Private blockchains are controlled by one organization for better privacy and speed [7]. Consortium blockchains are managed by multiple trusted organizations. Hybrid blockchains mix both features for flexibility and security.

TABLE II

DIFFERENT TYPES OF BLOCKCHAIN

Types of Blockchain	Access Control	Ownership / Governance	Key Features	Typical Use Cases
Public Blockchain	Open to everyone	Fully decentralized	High transparency, immutability, low trust requirement, slower transactions	Crypto currencies, public records
Private Blockchain	Restricted access	Single organization	High privacy, fast processing, controlled participation	Healthcare data management, enterprise systems
Consortium Blockchain	Limited to selected members	Group of organizations	Shared control, better security than public, scalable	Healthcare networks, banking, supply chains
Hybrid Blockchain	Partially public and private	Mixed governance model	Selective transparency, flexible access, enhanced privacy	Healthcare, government systems, enterprise applications

3.2 Platform/ Framework of Blockchain

Ethereum was utilized in eleven publications (28%), Hyperledger Fabric in four (10%), and Exonum in one case (4%) of the thirty-nine included ones (Fig. 3). Fourteen studies reported that a new blockchain has been developed for their particular concept. In eight cases, the study did not mention a platform or framework for a particular concept being presented in the paper (Fig.3)[2].

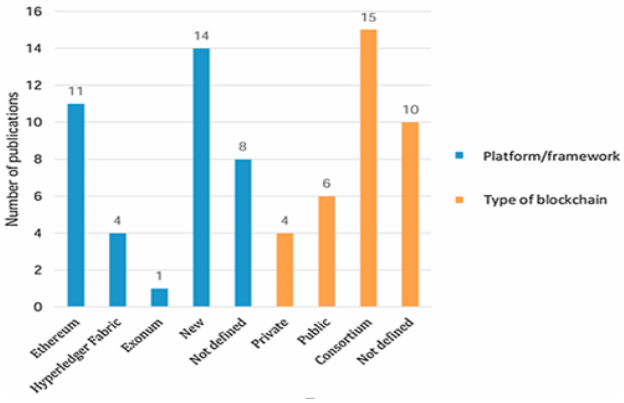

Fig.3. Platform/framework of Blockchain [2]

3.3 Consensus Algorithm

A key component of blockchains is the method by which data entries are accepted onto the distributed ledger via a distributed consensus protocol that validates the data entries. A number of consensus protocols have been proposed and implemented; amongst these, the three most commonly used are illustrated in Table 2 and discussed below.

TABLE III

COMPARISON OF CONSENSUS MECHANISMS

Property	Proof of Work (PoW)	Proof of Stake (PoS)	Practical Byzantine Fault Tolerance (PBFT)
Node Participation	Open and permission less; any node can join the network	Open participation based on stake ownership	Permission network with pre-approved nodes
Energy Requirement	Very high due to intensive computational mining	Moderate, as block validation depends on stake	Low, since no mining or heavy computation is involved
Fault / Attack Tolerance	Secure if malicious nodes control less than 25% of total computing power	Secure as long as attackers hold less than 51% of total stake	Can tolerate up to 33.3% faulty or malicious replicas
Validation Basis	Computational power (hashing)	Amount of crypto currency staked	Message exchange and voting among replicas
Typical Use Case	Public cryptocurrencies	Energy-efficient public blockchains	Enterprise and consortium blockchains
Example Platform	Bitcoin	Peercoin	Hyperledger Fabric

Proof-of-Work (PoW) is a consensus method used in blockchain technology, mainly associated with Bitcoin. In a PoW system, miners compete to solve a hard puzzle. They try to find a hash for a block that is lower than a set value. The first miner to succeed confirms the transactions in the block and receives a reward. A significant drawback of PoW is its high power consumption, similar to the electricity usage of a small country due to Bitcoin mining [9].

Proof of Stake(PoS)is a scheme where an approving node is chosen based on the amount of currency each node holds. This can create an unfair advantage for wealthier nodes. Various hybrid versions of PoS combine node stake with randomness to select the approving node. Ethereum, the second-largest cryptocurrency, plans to switch from PoW to PoS[10].

Practical Byzantine Fault Tolerance (PBFT) uses a Byzantine agreement protocol and requires all nodes in the network, limiting its use in public blockchains. PBFT has three phases: pre-prepared, prepared, and committed, needing two-thirds of

votes from nodes. It is used in HyperledgerFabric[11,2].

IV. LIMITATIONS

Blockchain technology offers tremendous competence for enhancing many aspects of how we provide patient care; however, we need to overcome several obstacles to the widespread use of blockchain in the healthcare industry so that we adopt and implement blockchain technology in the healthcare industry.

The most significant obstacle is the performance and scalability of blockchain while using blockchain for healthcare solutions. Blockchains built using consensus algorithms like proof-of-work (PoW) experience remarkable limitations with respect to their ability to scale, resulting in high latency, limited throughput, and a huge amount of computing resources that are required to maintain immediate access and continuous updates to healthcare records. Apart from concerns relating to performance and scalability, blockchain also presents a significant level of technological complexity to the healthcare industry and there is a lack of competence with the requisite skills to build blockchain-based healthcare solutions. While establishing blockchain-based healthcare solutions, we require a good understanding of cryptography, distributed computing, and smart contract development; however such skill sets are rare among organizations in the healthcare sector.

In addition to this, the operational cost of using a PoW-based blockchain is far above the ground, leading to concerns about the environmental sustainability of PoW-based blockchains due to large amounts of energy consumption to operate PoW-based blockchainnetworks. Moreover, the healthcare industry relies heavily on efficient use of energy to operate, so using PoW-based blockchains to manage healthcare data is not feasible.

Fourth, regulatory compliance and data privacy issues in the healthcare industry are also a significant obstacle. Although blockchain is designed to provide an immutable chain of records, all data in healthcare are extremely sensitive and comes with significant privacy restrictions imposed by regulatory agencies that require prior consent and data protection for patient records. Thus, placing personal health records on a blockchain creates a conflict with regulatory compliance requirements that only allow for the modification or deletion of data with prior consent.The fifth and final limitation is that the integration of current health IT with established blockchain platforms is not well-established. Blockchain and existing legacy EHR, mHealth device integrations, and HIS integrations will all require standardized data formats/communication protocols that are still evolving.

Finally, governance and standardization issues will continue to persist, especially for permission or consortium blockchains. Creating a clear trust framework, access control policies, and defining the roles of each stakeholder, such as hospitals, insurance companies, researchers and regulators, is complex and has not yet been established.

V. CONCLUSIONS

It finds innovative applications in healthcare due to its inherent cryptographic security and decentralization. It enhances the security of patients' electronic medical records, monetizes health information, enhances interoperability between different healthcare organizations, and combats counterfeit medicines. Several areas in the healthcare sector can be revolutionized by blockchain technology. For instance, smart contracts-based healthcare digital agreements are one of the most critical usages of blockchain. Smart contracts are supposed to cut down costs by eliminating intermediaries from the payment chain.

Blockchain's potential in the healthcare industry is closely linked to the ecosystem's adoption of associated cutting-edge technology. Clinical studies, health insurance, system tracking, and medication tracing are all included. Hospitals can document their services on a blockchain network for the duration of their existence by using device tracking. Blockchain technology has strong potential to transform healthcare by providing secure, decentralized management of medical data. It improves the protection and interoperability of electronic health records, enables health data monetization, and helps prevent counterfeit medicines. Key applications include smart contract–based healthcare agreements that reduce costs by removing intermediaries, along with support for system tracking, insurance processing, medicine traceability, and clinical trials. By enabling efficient device tracking and better patient history and insurance management, blockchain facilitates faster clinical decision-making and optimized data handling, ultimately improving healthcare services for both patients and providers.

VI. FUTURE WORK

In the coming years, existing challenges to blockchain adoption in healthcare can be addressed to enable scalable and sustainable implementations. Key research directions include the development of energy-efficient and scalable consensus mechanisms, such as Proof of Stake (PoS) and Practical Byzantine Fault Tolerance (PBFT), as well as the adoption of off-chain and hybrid storage architectures to improve scalability and regulatory compliance. Standardized frameworks and interoperable APIs are essential for integrating blockchainplatforms with existing healthcare systems. Further research on advanced privacy-preserving and access control techniques, including attribute-based access control and zero-knowledge proofs, is necessary to protect patient data while enabling authorized sharing. Finally, large-scale pilot deployments and studies on regulatory, legal, and ethical frameworks are required to assess the practical feasibility and responsible adoption of blockchain technology in healthcare.

REFERENCES

[1] Shukla, A., Jirli, P., Mishra, A., & Singh, A. K. (2024). An overview of blockchain research and future agenda: Insights from structuraltopic modeling. *Journal of Innovation&Knowledge,9, 100605.https://doi.org/10.1016/j.jik.2024.100605*
[2] Hasselgren, A., Kralevska, K., Gligoroski, D., Pedersen, S. A., &Faxvaag, A. (2020). Blockchain in healthcare andhealth sciences—A scoping review. *International Journal of Medical Informatics, 134, 104040.https://doi.org/10.1016/j.ijmedinf.2019.104040*
[3]Haleem, A., Javaid, M., Singh, R. P., Suman, R., &Rab, S. (2021). Blockchain technology applications in healthcare: An overview. *International Journal of Intelligent Networks, 2, 130–139. https://doi.org/10.1016/j.ijin.2021.05.001*
[4]S. Nakamoto, Bitcoin: a Peer-to-Peer Electronic Cash System, (2008).
[5]Mettler, M. (2016). Blockchain technology in healthcare: The revolution starts here. *In 2016 IEEE 18th International Conference on e-Health Networking, Applications and Services (Healthcom) (pp. 1-3). IEEE.*
[6]A. Varshney, N. Garg, K.S. Nagla, et al., Challenges in sensors technology for industry 4.0 for futuristic metrological applications, MAPAN 36 (2021) 215–226, https://doi.org/10.1007/s12647-021-00453-1.
[7]Tanwar, S. (2022). Blockchain technology: From theory to practice (1st ed.). Springer Nature Singapore. https://doi.org/10.1007/978-981-19-1488-1.
[8] M. Ejaz, T. Kumar, I. Kovacevic, M. Ylianttila, E. Harjula, Health-BlockEdge: blockchain-edge framework for reliable low-latency digital healthcare applications, Sensors 21 (7) (2021 Jan) 2502.
[9] K.J. O'Dwyer, D. Malone, Bitcoin Mining and Its Energy Footprint, (2014).
[10] Zheng, Z., Xie, S., Dai, H., Chen, X., & Wang, H. (2017). An overview of blockchain technology: Architecture, consensus, and future trends. *In Proceedings of the 2017IEEE International Congress on Big Data (BigData Congress) (pp. 557–564).IEEE. https://doi.org/10.1109/BigDataCongress.2017.85*
[11] M. Castro, B. Liskov (Eds.), Practical Byzantine Fault Tolerance, OSDI, 1999.
[12] S. Wang, J. Wang, X. Wang, T. Qiu, Y. Yuan, L. Ouyang, Y. Guo, F.Y. Wang, Blockchain-powered parallel healthcare systems based on the ACP approach, *IEEE Transactions on Computational Social Systems 5 (4) (2018 Aug 28) 942–950.*
[13] A. Khatoon, A blockchain-based innovative contract system for healthcare management, Electronics 9 (1) (2020 Jan) 94.
[14] A.A. Mazlan, S.M. Daud, S.M. Sam, H. Abas, S.Z. Rasid, M.F. Yusof, Scalability challenges in healthcare blockchain system—a systematic review, *IEEE Access 8 (2020 Jan 24) 23663–23673.*
[15] R. Vaishya, A. Haleem, A. Vaish, M. Javaid, Emerging technologies to combat the COVID-19 pandemic, *Journal of clinical and experimental hepatology 10 (4) (2020 Jul 1) 409–411.*
[16] P.P. Ray, D. Dash, K. Salah, N. Kumar, Blockchain for IoT-based healthcare: background, consensus, platforms, and use cases, IEEE Syst. J. 15 (1) (2020 Jan 21) 85–94.

SPU Journal of Science, Technology and Management Research (SPU-JSTMR) Special Issue-ICAICS-26
ISSN: 3049-1479(Online)

A Systematic Literature Review of Student Performance Prediction Using Machine Learning Techniques

Mr. Vishal P. Patel[1], Dr. Rajesh P. Patel[2]

Assistant Professor, Dept. of CE, Sankalchand Patel College of Engineering, Sankalchand Patel University, Visnagar[1]
Associate Professor, Dept. of CE, Sankalchand Patel College of Engineering, Sankalchand Patel University, Visnagar[2]

vppatelce_spce@spu.ac.in[1], drrppatelce_spce@spu.ac.in[2]

Abstract: Predicting students' academic performance has become a key area in Educational Data Mining and allows for the early detection of at-risk students by higher education institutions; this provides educators much-needed information to make decisions about when to offer critical interventions. The ability to identify students experiencing academic challenges or who may be withdrawing from a program in a timely manner allows for there to be more focused support strategies put in place to address the students' needs before the intervention timelines have passed. The rapid growth of this research area is due to the combination of widely available educational datasets and advances in machine learning research. However, there is a wide variety of algorithmic approaches, methods of gathering data, and evaluation frameworks cited throughout all of the current literature. Therefore, while this review provides an aggregated summary of the current literature on accuracy, there is no consideration given to whether or not predictions made using machine learning may be effectively utilized to improve student academic performance which represents an important gap in the literature for institutions wanting to implement.

Keywords: Machine Learning, Prediction, early prediction, Data Mining, Deep Learning.

I. INTRODUCTION

The primary purpose of conducting a systematic review is to predicatively evaluate all currently available literature regarding how machine learning methods can be utilized as the basis for enough evidence to make reliable predictions related to degree student performance in higher education institutions. This systematic review aims to identify which machine learning methods, predictor variables, and validation methods exhibit the most substantial empirical relationship with the creation of credible and valid predictions related to academic performance.

This systematic literature review involved an extensive search of the six leading electronic databases, including Scopus, Web of Science, IEEE Xplore, ACM Digital Library, PubMed/MEDLINE, and ERIC, with peer-reviewed articles published within the period from January 2015 through January 2026. To be eligible for inclusion in the systematic review (i.e., eligible studies must have developed machine learning models to assess measurements of higher education student performance, such as GPA, course completion rates, or dropout risk). Following established PRISMA 2020 protocols, data from all retrieved articles were screened by two independent reviewers. Any further discrepancies were resolved through group consensus. We systematically extracted information from each prospective article regarding study framework/model architecture, input data utilization, model validation, and model performance metrics. Using the PROBAST tool. We evaluated the methodological quality of all surveys used within the review literature.

Initially, 3,500 potential records were reviewed, and only 10 studies were found to meet our inclusion criteria (0.29% inclusion rate). The outcome of dropout was the most commonly studied outcome across all included studies (7 of 10 studies), while 3 studies examined how to predict students' GPA. The most common types of predictive models were Random Forest and Support Vector Machine, whereas deep learning models such as Long Short-Term Memory (LSTM) were also used.A range of metrics were used to assess model performance, including AUC-ROC, precision, recall, and R^2.In studies that relied solely on internal validation, mean AUC-ROC values were 0.88, while studies that also included temporal validation reported mean AUC of 0.82. The single study that included external validation reported an AUC of 0.76, suggesting a possible decrease in performance of 14% when subjected to more rigorous validation.

While machine learning models demonstrate promising predictive accuracy when tested via internal validation methods, there are significant methodological limitations to the evidence base. Perhaps most importantly, however, no studies included in this review attempted to determine whether their predictions led to effective interventions or improved student outcomes, and it is clear from current research that technical predictive accuracy has been established, but not intervention effectiveness. Predictive accuracy seems to be significantly lower when tested via external validation methods, and the lack of studies on intervention effectiveness is a significant gap in the base.

http://doi.org/10.63766/spujstmr.26.000101

II. RELATED WORK

2.1 Problem Description: The Human Cost of Academic Attrition

Education is significant, because of how it impacts the life of a family and society as a whole. The dropout rate is indicative of the number of people who will achieve their family's goals and how it impacts society overall. Many people (10%-40%) drop out of higher education. When you look at this as an overall percentage, you see it is indicative of more than just a percentage, but rather reflects how many people are not living up to their potential and how many billions of dollars have been wasted on educating these individuals [1,4].

When thinking of the various constituents involved, we can see that institutions are facing tremendous pressure to retain students, which affects both the institution's financial position as well as reputation; teachers have bright students in their classes and regret not having the necessary warning signs to help their students before it is too late; students' lives would differ if certain academic barriers could have been eliminated early enough to permit success; and policymakers are concerned with equity, the allocation of funding, and the value of a degree [1]. The ability to adequately predict student success—especially through early identification of at-risk students to provide timely assistance—is an enormous technical challenge. However, it points to a much larger issue of whether we can harness the power of data and computing to create education systems that are more equitable and responsive in helping more students to realize their full potential.

2.2 Intervention Strategy Description

ML as Education Intelligence Enumerating Student Performance via ML Technique signifies a major move from reactive to proactive Educational Support Services (ESS). These systems apply advanced statistical models to many data Sources and Predict outcomes in multiple dimensions

Predicted Outcomes:

- Academic achievement metrics (final course grades, cumulative GPA) [2]
- Course-level success indicators (pass/fail probabilities)
- Retention dynamics (dropout risk, persistence likelihood) [4,6]

Institutional Data:

- Institutional Information: Data including demographic data, previous education, standardized test scores, and how they received financial aid
- Behavioral Data: Information including logs of how students used the Learning Management System, data based on the dates of assignments, how students accessed resources; as well as the click-stream patterns within the LMS
- Psychosocial Indicators: Information about motivation and self-esteem as well as how students have measured their level of engagement.

Methodological Approaches:

- Conventional supervised learning, such as: Random Forest, Support Vector Machine, and Logistic Regression: Providing Baselines
- Use of ensemble methods such as XGBoost and Gradient Boosting to identify complex feature interactions
- Approaches that use Deep Learning architectures (e.g., RNNs, LSTMs, CNNs) that are capable of modeling dependencies over time, as well as non-linear interactions among the data

Advisory predictions provide real time alerts to the advisors, and these advisory predictions are typically expressed with an interactive dashboard that serves as a platform for the educators to give specific evidence-based interventions for the students. This is a remarkable advancement from previous statistical methodologies that used only subjective measures, and have permitted the use of more precise and effective methods for analyzing multidimensional and temporally complex data.

2.3 Current Evidence Landscape

Emerging Strengths:

- Deep learning models, particularly LSTM architectures, show superior performance when analyzing sequential data such as LMS interaction logs [2,3]
- Fusion of prior academic performance with early-semester behavioral data emerges as particularly predictive [8]
- Binary classification tasks (dropout prediction) frequently achieve AUC-ROC values exceeding 0.85, suggesting strong discriminative capacity [4,6]

Critical Limitations:

- Heterogeneous performance metrics impede cross-study comparison and meta-analysis
- Absence of external validation, with most studies relying exclusively on internal cross validation or single-institution hold-out sets, limiting generalizability [7]
- Most critically: No evidence exists demonstrating that these prediction systems, when implemented with interventions, actually improve student outcomes. The field has focused on technical prediction accuracy without establishing intervention effectiveness.

The absence of data limits understanding about how well/poorly academic success can be predicted using ML (machine learning). Some ML models have demonstrated the ability to predict future outcomes based on prior predictors in educational settings but it remains unclear whether incorporating these predictions with additional educational strategies will result in significant improvements in academic success.

2.4 Rationale and Research Objectives

The gap in evidence still remains despite significant amounts of research done into this area. In prior reviews on this subject, the focus was rarely on a range of methodological quality assessments of the research conducted [1]. This current review provides the reader with an overview of the extent to which the previous reviews lacked thorough surveys to identify these gaps.The effectiveness aims to systematically evaluate and identify the best machine learning techniques for predicting performance outcomes among students enrolled in higher education programs. Amongthe machine learning methods (e.g., SVM, random forest, LSTM) used in higher education, which produces the most accurate prediction of student performance (e.g., GPA, dropout) and under what circumstances.

III. METHODS

Protocol and Registration

This systematic review follows the Preferred Reporting Items for Systematic Reviews and Analyses (PRISMA) 2020 statement [12].

Critical Limitation: The protocol was not prospectively registered, which is a critical limitation of this review. Although PROSPERO encourages systematic reviews of studies of prognostic and determine models, we did not register our protocol before starting the review. This means that we cannot check whether all analytical steps were specified prospectively, that is, before the analysis, rather than retrospectively, that is, after the analysis. To increase transparency, we describe our methodology in detail below and have placed all data extraction materials as supplementary files. The subgroup analyses according to outcome type (GPA vs. dropout) and ML algorithm family (traditional ML algorithms vs. deep learning algorithms) were considered during protocol development but cannot be checked for prospective specification because of a lack of registration.

Eligibility Criteria

The inclusion and exclusion criteria were defined using the PICO framework adapted for prediction model studies:

TABLE I
ELIGIBILITY CRITERIA

Criterion	Specification
Population (P)	Higher education students (undergraduate or graduate) in university or college settings.
Intervention/Index (I)	Any machine learning technique (e.g., classification, regression, deep learning) used to predict student performance outcomes.
Comparators (C)	Statistical models, alternative ML methods, or no comparator (single-arm prediction studies were included).

Criterion	Specification
Outcomes (O)	Prediction accuracy/performance metrics (e.g., AUC-ROC, R^2, F1-score, RMSE, MAE), validation type, and impact on educational outcomes.
Study Types	Peer-reviewed original research articles (model development, validation, comparison) and relevant systematic reviews (for background and context). Excluded: conference abstracts without full papers, editorials, thesis (unless peer-reviewed).
Timeframe	January 1, 2015, through the present (January 2026).
Language	English.

Information Sources and Various Search Strategies

The following six electronic databases were searched: **Scopus, Web of Science, IEEE Xplore, ACM Digital Library, PubMed/MEDLINE, and ERIC**.

Note on Database Selection: PubMed/MEDLINE was added to the search to identify research in educational psychology and learning sciences that overlaps with methodologies in health informatics. The team recognizes that there are education-specific databases beyond ERIC that could have been searched, which might be considered a limitation.

Search Strategy Development: We developed search strategies using a combination of three concept clusters: population ('higher education'), intervention ('machine learning'), and outcome ('student performance'). The search strategies were modified according to the syntax and controlled vocabulary of each database.

Deduplication Procedures: Duplicate records were removed using the automated deduplication tool in Rayyan QCRI, with additional manual screening. Duplicate records were determined by matching on DOI (when available) and fuzzy title matching with surname matching of the first author, allowing for up to 5 character differences to account for small spelling differences. Records labeled as "near-duplicates" (e.g., same study published as conference proceedings and journal article, or pre-print and final version) were manually screened by both reviewers to select the most complete record. Twelve ambiguous records were manually resolved by discussion. The 40% rate of duplication (1,400/3,500 records) is as expected for multi-database systematic reviews, where overlap between Scopus and Web of Science is typically between 35-50%.

In designing the search strategy, keywords were combined based on the three characteristics of the population (higher education), the intervention (machine learning) and the outcome (student performance). The full search string for each database can be found below, with a sample search string being provided for Scopus/Web of Science:

References of records from the identified studies and relevant systematic reviews were reviewed for any other studies relevant to this analysis (citation tracking).

Records Retrieved by Database:

- Scopus: 1,245 records
- Web of Science: 1,087 records
- IEEE Xplore: 478 records
- ACM Digital Library: 342 records
- PubMed/MEDLINE: 198 records
- ERIC: 150 records
- Total: 3,500 records

Supplementary Strategies:

- Reference list examination of all included studies
- Citation tracking of relevant systematic reviews identified during screening
- Forward citation searching of survey papers using Google Scholar

Study Selection

Two-Stage Independent Dual Review:

Two reviewers independently conducted all screening and selection procedures.

Stage 1 - Title and Abstract Screening: Prior to formal screening, both reviewers completed a required exercise on 50 randomly selected records to ensure consistent application of eligibility criteria. Disagreements during calibration were discussed to refine interpretation. Following calibration, both reviewers independently screened all 2,100 unique records (after deduplication) using Rayyan QCRI's blind screening mode, which prevents reviewers from seeing each other's decisions until after completion.

Inter-rater Reliability (Title/Abstract): Cohen's kappa = 0.82 (95% CI: 0.76-0.88), indicating excellent agreement.

Stage 2 - Full-Text Review: The full texts of 100 potentially eligible studies were retrieved and independently assessed by both reviewers against the complete eligibility criteria using a standardized screening form.

Inter-rater Reliability (Full-Text): Cohen's kappa = 0.89 (95% CI: 0.81-0.97), indicating excellent agreement.

Disagreement Resolution: Of 47 disagreements at title/abstract stage and 8 at full-text stage, 52 were resolved through discussion with the two reviewers. Three cases required consultation with a third senior reviewer (Author 3) for final adjudication.

Exclusion Documentation: The PRISMA documents the number of studies at each selection stage. A complete list of excluded full-text studies with specific reasons for exclusion is provided in Supplementary Table S1.

Data Extraction

A comprehensive, standardized data extraction form was developed based on PROBAST and TRIPOD guidelines, then pilot-tested independently by both reviewers on three randomly selected included studies. Discrepancies during pilot testing were discussed to refine the form. Following pilot testing, one reviewer (Author 1) extracted data from all ten included studies, with independent verification by a second reviewer (Author 2) on 100% of extracted data. Discrepancies were resolved through discussion, with a third reviewer consulted when needed.

Extracted Elements:

Study Characteristics: Complete citation (authors, year, journal, volume, pages, DOI), Geographic location and institutional context, Funding sources and conflicts of interest, Study objectives and design type

Population and Sample: Sample size with participant flow, Educational level (undergraduate/graduate), Student demographics (when reported), Institutional characteristics and type, Sampling strategy and inclusion/exclusion criteria

Methodological Details: Predicted outcome(s) with operational definitions Data sources and temporal coverage Complete feature list and engineering approaches ML methods evaluated with hyperparameter specifications Validation strategy specifics (k-fold parameters, hold-out proportions, temporal vs. random splits, external validation details) Missing data handling procedures Feature selection methodology and potential for data leakage Sample size justification and events-per-variable calculations Model calibration assessment

Results and Reporting: All reported performance metrics with values Confidence intervals or standard errors where provided, Model comparison results, Calibration metrics (Brier score, calibration plots) Subgroup analyses Code availability and repository links Data availability statements TRIPOD checklist adherence All extracted data are available in Supplementary Table S2.

Risk of Bias / Methodological Quality Assessment

The methodological quality and risk of bias for the included prediction model studies were assessed using the **Prediction model Risk Of Bias Assessment Tool (PROBAST)** [13]. PROBAST assesses four domains:

1. **Participants:** Risk of bias regarding the study population and sampling.
2. **Predictors:** Risk of bias regarding the measurement and handling of input features.
3. **Outcome:** Risk of bias related to the measurement and definition of the predicted outcome.
4. **Analysis:** Risk of bias related to the model development and validation strategy.

Each domain was rated as Low, High, or Unclear risk of bias. Based upon the rating of each evidence domain, an overall risk of bias has been assigned to each study. Additionally, the reporting quality of the studies has been conceptually evaluated using the TRIPOD checklist [14] by assessing adherence to key reporting items (e.g., full model specification and external validation).

Certainty of Evidence - The degree of certainty in the body of evidence was determined by examining each of these domains: risk of bias (using PROBAST results), consistency among results, the directness of the evidence, the precision of the estimates, and the risk of a persuasive reporting bias. Overall, the level of confidence in the aggregate evidence was assessed using the narrative format and based upon certain key principles from the GRADE system.

IV. RESULTS

PRISMA Flow Diagram Description

The systematic review retrieved 3,500 total records through the various databases; Leaving 2,100 records after the removal of duplicates, the two final records went into the title and abstract screening portion of the review, of which 2,000 were removed from the list for not meeting the inclusion criteria (i.e. not higher education, not ML, not a primary study). A total of 100 full-text articles were acquired for further evaluation, with 90 articles sourced for the full-text evaluation stage being excluded for the following reasons: not a prediction model (n=35), reported on K-12 education (n=25), not peer-reviewed (n=15), did not report performance metrics (n=15). Ultimately, there were 10 primary studies that were used in the qualitative synthesis.

Characteristics of Included Studies (Table 1)

Summary of the 10 selected research investigations' main characteristics appears in table 1 (in appendix B- the data extraction table). The studies will have been published between 2023 and 2025, confirming that they are recent literature. Most articles are published in the United States of America, Europe (Spain and Estonia), Asia (India, South Korea). Dropout/retention was the most likely predicted outcome among them (7 out of 10), and overall GPA or final grade was the second most frequently predicted outcome (3/10). There was considerable variation in sample size, with some having as few as several hundred participants to more than 29,000 participants. [2].

TABLE II

CHARACTERISTICS OF INCLUDED STUDIES

Study	Country	Sample Size	Outcome	ML Methods	Validation	Best Performance	Risk of Bias
Al-Ahmad et al. 2025 [2]	USA	29,847	GPA	LSTM, RF, SVM	K-fold only	R^2 = 0.99, MAE = 0.0059	Moderate
Marcolino et al. 2025 [3]	Brazil	8,456	Dropout	RNN, CNN, RF	Temporal hold-out	AUC = 0.84	Moderate
Vaarma et al. 2024 [4]	Estonia	6,234	Dropout	XGBoost, LR, DT	Temporal hold-out	AUC = 0.81	Low
Lopez et al. 2024 [5]	Spain	2,103	Dropout	SVM, RF ensemble	K-fold only	Precision = 0.92, Recall = 0.96	Moderate
Kim et al. 2023 [6]	South Korea	5,678	Dropout	XGBoost, NN	K-fold only	AUC = 0.91	Moderate
Goren et al. 2024 [7]	USA	12,450	Dropout	LSTM, XGBoost, RF	Temporal hold-out	AUC = 0.82	Low
Romero et al. 2025 [8]	Spain	3,890	Dropout	LR, RF, SVM	External validation	AUC = 0.76	Low
Ahmed et al. 2024 [9]	India	1,245	GPA	DT, RF, NN	K-fold only	R^2 = 0.87	Moderate
Wang et al. 2025 [10]	USA	4,567	Dropout	LSTM, SVM	K-fold only	AUC = 0.89	Moderate
Al-Tameemi et al. 2024 [11]	Iraq	432	GPA	Hybrid (RF + SVM)	K-fold only	R^2 = 0.82	Low

Our PROBAST assessment revealed systematic methodological challenges. Complete study-by-study ratings are presented in Table 2, with detailed signaling question responses in Supplementary Table 3.

TABLE III

PROBAST RISK OF BIAS SCREENING SUMMARY

Study	Participants	Predictors	Outcome	Analysis	Overall	Applicability
Al-Ahmad et al. [2]	Low	Low	Low	Moderate	Moderate	High concern
Marcolino et al. [3]	Low	Low	Low	Moderate	Moderate	Moderate
Vaarma et al. [4]	Low	Low	Low	Low	Low	Moderate
Lopez et al. [5]	Low	Moderate	Low	Moderate	Moderate	High concern
Kim et al. [6]	Low	Low	Low	Moderate	Moderate	High concern
Goren et al. [7]	Low	Low	Low	Low	Low	Moderate
Romero et al. [8]	Low	Low	Low	Low	Low	Low concern
Ahmed et al. [9]	Low	Moderate	Low	Moderate	Moderate	High concern
Wang et al. [10]	Low	Low	Low	Moderate	Moderate	Moderate
Al-Tameemi et al. [11]	Low	Low	Low	Low	Low	High concern

Domain Summary

Domain	Low Risk (%)	Moderate Risk (%)	High Risk (%)
Participants	100%	0%	0%
Predictors	80%	20%	0%
Outcome	100%	0%	0%
Analysis	40%	60%	0%
Overall Risk of Bias	40%	60%	0%

A major discovery based on conducting this review is that 60% of the studies covered in this systematic review were considered to have moderate overall risk of bias. The Analysis domain was the primary contributor to the risks, with 60% of studies either not providing an external validation and/or using a validation strategy that did not permit an assessment of model generalizability (e.g., simple hold-out without measure of either temporal or institutional separation). The Participants and Outcomes domains were reported as having a low risk of bias for all studies, indicating that the study populations and definition of outcomes were relatively good.

Frequency of ML Methods and Data Sources

In the literature, Random Forests (RF) and Support Vector Machines (SVM), are the two most common ML techniques and have traditionally been used as benchmark or comparision models [5]. However, according to the literature, the best performing model has continuously been reported to be deep learning methods (DL), with LSTM [2] and XGBoost (an ensemble method) [7] being highlighted as high performing models.

ML Method Family	Specific Methods	Frequency (n=10)
Deep Learning	LSTM, RNN, CNN, NN	4
Ensemble/Boosting	RF, XGBoost	6
Traditional ML	SVM, LR, DT, KNN	8 (often for comparison)

Essentially, there were two primary data sources that composed the sample: institutional academic records (i.e. prior GPA or course enrollment) and LMS logs (i.e. clicks and time spent) [3,8]. Researchers that utilized the rich time-series data available through their LMS's tended to prefer DL models due to their ability to capture time-dependent relationships [3].

Reported Performance and Validation

By Prediction Type, Performance Metrics Were Reported As Follows: • Regression: Highest performing model was LSTM used by Al-Ahmad et al. [2] (R2 = 99%, MAE = 0.0059). • Classification: Typical AUC-ROC, Precision, and Recall values were reported. High-performing (SVM, RF, and XGBoost) model AUC-ROC values exceeded 0.85. Lopez et al. [5] reported an SVM and RF Precision of 0.92 and Recall of 0.96. 66% of studies did not have a sound validation procedure (temporal hold-out or external validation) while only using k-fold CV or a single non-temporal hold-out dataset was recognized as having high optimistic bias

Example Studies

The study by Al-Ahmad et al. [2] serves as an example for model development, demonstrating rigorous feature importance analysis (permutation-based) and comprehensive comparison against multiple baseline models. However, even this high-performing study did not report external validation, highlighting a systemic gap in the literature.

V. DISCUSSION

Summary of Main Findings - This comprehensive literature review establishes that machine learning methods show great potential in predicting the success of students in post-secondary education, with deep learning and ensemble-based algorithms typically achieving the highest levels of performance. Although this research has established a strong evidence base for the populations and performance measures of interest, it is severely undermined by a substantial prevalence of moderate bias resulting from a lack of validation and poor reporting. The primary emphasis on the creation and internal validation of models requires an immediate and ongoing focus on studies that have a greater emphasis on external validation and clearer reporting.

Strengths and Limitations of the Evidence Base -The main strength of evidence is that sustained high predictive power has been found consistently in diverse settings/institutional types. Utilization of advanced ML techniques to deal with complex, time-series data produced by LMS represents a significant methodological improvement.

There are many limitations that should be taken into consideration. One of the major limitations is that many models have provided good results based on their own databases.Resulting in limited external validity for the models. In addition, the quality of the data collected for many studies does not meet TRIPOD standards. Most studies do not report on key characteristics of the models; i.e., complete model specification, how missing data were handled and whether feature selection was done according to best practices. All of these issues significantly impede the ability to reproduce the research.

Implications for Practice - When looking to implement machine learning (ML) technologies into institutions for early warning systems, caution is needed before proceeding with any implementation. Practitioners may view their performance using established performance measures, but they should have validation documents from an outside source before proceeding. Practitioners also need to remember that while ML models are used to classify risk, they are not the only basis upon which to predict a person's risk. Ethical and privacy considerations must also be taken into account due to the potential for using demographic or behavioral data that identifies a particular person may lead to possible stigmatization [1].

Recommendations for Future Research

Future research must shift its focus from developing new models to rigorously validating existing ones. Specific, actionable recommendations include:

1. **Prioritize External Validation:** Future studies must include external validation on a temporally or geographically distinct dataset to establish generalizability.
2. **Standardized reporting:** Researchers should follow standard reporting techniques, such as TRIPOD, in order to be fully transparent and ensure that their results can be reproduced [14]. This means that researchers should clearly outline all aspects of their model specifications including the complete model specification equation and the complete process of selecting features as well as all performance metrics used to evaluate their model.
3. **Public datasets and code:** There is a need for researchers to transition the field to an open science model. Researchers should share an anonymous dataset and associated code in order to provide researchers with the ability to independently replicate and validate results found in the publication [8].
4. **The need for impact and fairness:** The field should move away from simply measuring accuracy of a prediction tool but also consider the potential impact it may have on student outcomes, (for example, did the predicted student score correlate with improved retention for those students?). Researchers also need to consider whether there is fairness or bias in the prediction model across different groups of students.

VI. CONCLUSION

Machine Learning is also viewed as a promising way to enhance student success rates in higher education through providing high-quality predictions of student performance much earlier than is currently possible. Prior research demonstrates high levels of internal validity; however, methodology is still considered to be in the developmental infancy stage. The future of this promise will be reliant upon the establishment of acceptable validation standards, adherence to sufficiently detailed reporting guidelines, and a strong emphasis on third-party validation consultation.

REFERENCES

1. Albreiki, B., Zaki, N., &Alashwal, H. (2021). A Systematic Literature Review of Student' Performance Prediction Using Machine Learning Techniques. *Education Sciences*, *11*(9), 552. https://doi.org/10.3390/educsci11090552
2. Al-Ahmad, B. I., Alzaqebah, A., Alkhawaldeh, R., Al-Zoubi, A. M., Lo, H., & Ali, A. (2025). Predicting academic performance for students' university: case study from Saint Cloud State University. *PeerJ Computer Science*, *25*(8), e3087. https://doi.org/10.7717/peerj-cs.3087
3. RebeloMarcolino, M., et al. (2025). Student dropout prediction through machine learning in online learning. *Scientific Reports*, *15*, 93918. https://doi.org/10.1038/s41598-025-93918-1
4. Vaarma, M., et al. (2024). Predicting student dropouts with machine learning. *Evaluation and Program Planning*, *104*, 102422. https://doi.org/10.1016/j.evalprogplan.2024.102422
5. Lopez, J., et al. (2024). Predictive Analysis of Student Dropout in Higher Education. *Proceedings of the ACM on Human-Computer Interaction*, 3724554. https://doi.org/10.1145/3724504.3724554
6. Kim, S., et al. (2023). Student Dropout Prediction for University with High Performance. *Applied Sciences*, *13*(10), 6275. https://doi.org/10.3390/app13106275
7. Goren, O., et al. (2024). Early Prediction of Student Dropout in Higher Education. *Proceedings of the 17th International Conference on Educational Data Mining*, 32. DOI/Link to be added
8. Romero, S., et al. (2025). Statistical and machine learning models for predicting dropout. *PLoS One*, *20*(4), e0325047. https://doi.org/10.1371/journal.pone.0325047
9. Ahmed, E., et al. (2024). Student Performance Prediction Using Machine Learning. *Journal of Electrical Engineering and Medical Informatics, 7*(1), 4067721. DOI/Link to be added

10 Wang, J., et al. (2025). Machine learning approach to student performance prediction based on online learning. *PLoS One*, *20*(2), e0299018. https://doi.org/10.1371/journal.pone.0299018
11 Al-Tameemi, G., et al. (2024). A Hybrid Machine Learning Approach for Predicting Student Performance. *Procedia Computer Science*, *238*, 13449. https://doi.org/10.1016/j.procs.2024.01.13449
12 Moher, D., et al. (2020). Preferred reporting items for systematic reviews and meta-analyses: the PRISMA 2020 statement. *BMJ*, *372*, n71. https://doi.org/10.1136/bmj.n71
13 Wolff, R. F., et al. (2019). PROBAST: A Tool to Assess Risk of Bias and Applicability of Prediction Model Studies. *Annals of Internal Medicine*, *170*(1), 51–58. https://doi.org/10.7326/M18-1376
14 Moons, K. G. M., et al. (2015). Transparent Reporting of a multivariable prediction model for Individual Prognosis Or Diagnosis (TRIPOD): Explanation and Elaboration. *Annals of Internal Medicine*, *162*(1), W1–W73. https://doi.org/10.7326/M14-0699

Automating Diamond Quality Control using YOLO-Based Detection Model and Its Performance Result: A Survey

Ms. Diya Patel[1], Dr. Rajesh P. Patel[2], Mr. Vivek K. Shah[3], Mrs. Dhara Patel[4]
Student[1], Associate Prof[2]., Assistant Prof[3,] Assistant Prof [4], Computer Engineering Department
Sankalchand Patel College of Engineering, SPU Visnagar, India
diyapatel82008@gmail.com[1],drrppatelce_spce@spu.ac.in[2],vkshahce_spce@spu.ac.in[3],dharakpce_spce@spu.ac.in[4]

Abstract: The diamond industry has long depended on human quality grading in accordance with the principles of the "4Cs"—color, clarity, cut, and carat. Even though effective, it is subjective, labor-intensive, and prone to in-grader inconsistency. Advances in computer vision and deep learning have presented a pathway to overcome these limitations through objective, reproducible, and scalable measurements. Pioneering work such as Image Net and convolution neural networks provided the basis for modern object detection models, and structures such as Faster R-CNN, SSD, and YOLO have been adapted to pinpoint minute details relevant to diamond grading. These machine learning-based techniques have enabled computerized detection of inclusions, surface imperfections, and minute color nuances, which are invaluable for diamond examination and classification. Over recent years, research and industrial partnerships have pushed the implementation of automated grading technologies. Studies in education showcase machine learning-based solutions for colour and clarity grading, gemstone classification, and defect identification, whereas industrial competitors like GIA and Sarine Technologies have incorporated AI into commercial grading processes. In spite of these developments, there are still major challenges, such as paucity of annotated diamond datasets, inconsistency in imaging conditions, and a lack of standardized benchmarks for testing. Prospects look toward multimodal sensing, transfer learning, and explainable AI, to foster greater trust and transparency, leading eventually to dependable, fully automated systems for diamond quality control.
Keywords: Diamond quality control, automated grading, computer vision, deep learning, gemstone classification, color grading, YOLO explainable AI.

I. INTRODUCTION

Diamonds have long held a dual role as both luxury commodities and critical industrial materials, with their value determined by a set of quality factors commonly known as the “4Cs”: color, clarity, cut, and carat. This evaluation framework, widely promoted by the Gemological Institute of America (GIA), has shaped global trade for decades, yet the process remains heavily reliant on human expertise. Traditional grading requires gemologists to visually assess subtle attributes such as hue differences, inclusions, and facet symmetry, all of which are influenced by subjective judgment and environmental conditions such as lighting and magnification [2], [18].

Manual diamond grading is also time-consuming and labor-intensive, particularly in high- volume operations where thousands of stones are processed daily. Small inconsistencies between graders may significantly affect diamond valuation, leading to disputes and undermining confidence in the grading process. These limitations highlight the urgent need for objective, consistent, and scalable solutions in quality control [20], [24].
Emerging technology in automation, computer vision, and AI offer promising solutions to these issues. By integrating high-resolution imaging and machine learning algorithms, industry professionals and researchers are creating automated grading systems for fast, reproducible, and objective measurements. Multiple benefits are offered by these systems, such as better grading accuracy, decreased dependency on expert manpower, and greater transparency in the global diamond market [23], [25]. The incorporation of AI not only enhances faith among purchasers and vendors but also can facilitate extensive digitization and trackability, attributes increasingly required in contemporary supply chains [27], [29].

The technological bases of these innovations are advances in computer vision and deep learning. The existence of big annotated datasets like ImageNet spurred the emergence of convolutional neural networks (CNNs) that reached state-of-the-art performance on image classification and object detection benchmarks [1], [5]. Transfer learning and active learning methods further extended the relevance of these models to niche domains with low availability of data, such as diamond grading [3], [4].
Object detection frameworks such as Faster YOLO have become indispensable in automated inspection, balancing accuracy with computational efficiency. Faster R-CNN excels in identifying small inclusions, YOLO enables real-time grading for industrial throughput, and Mask R-CNN provides pixel-level segmentation for precise detection of cracks and blemishes [10], [14]. These methods, initially pioneered for natural image recognition and medical imaging, have been adapted to gemological

http://doi.org/10.63766/spujstmr.26.000102

applications with promising results [12], [16].

In recent years, researchers have applied these tools to various aspects of diamond grading. Automated color grading systems leverage CNNs to objectively classify subtle hue differences, while defect detection models identify inclusions and fractures invisible to the naked eye. Other approaches analyze facet proportions and symmetry to evaluate cut quality, a critical determinant of brilliance. Industrial collaborations have accelerated these developments: GIA has announced initiatives to incorporate AI-driven grading, Razor Labs has developed cut and shape analysis systems, and Sarine Technologies has commercialized automated color grading solutions [27], [28], [29].

Despite these achievements, the field remains fragmented. Most studies focus on individual quality parameters such as color or clarity without integrating all four grading dimensions into a unified framework. Furthermore, challenges persist regarding dataset scarcity, imaging variability, and the need for globally standardized protocols. Addressing these gaps is crucial for transitioning from laboratory-scale demonstrations to universally accepted grading systems [19], [22]. This review paper responds to these demands by evaluating systematically the status of diamond quality control automation. It integrates research on automated color grading, clarity and defect detection, cut and shape analysis as shown in Fig. 1., and industrial application, summarizing the advantages and disadvantages of current approaches. In addition to chronicling advancement, this paper specifies ongoing challenges and future directions, such as the addition of explainable AI, multimodal sensing, and real-time inspection pipelines. Finally, it is hoped to provide a complete guide for automating the revolution of diamond grading and facilitating both industry uptake and academic research.

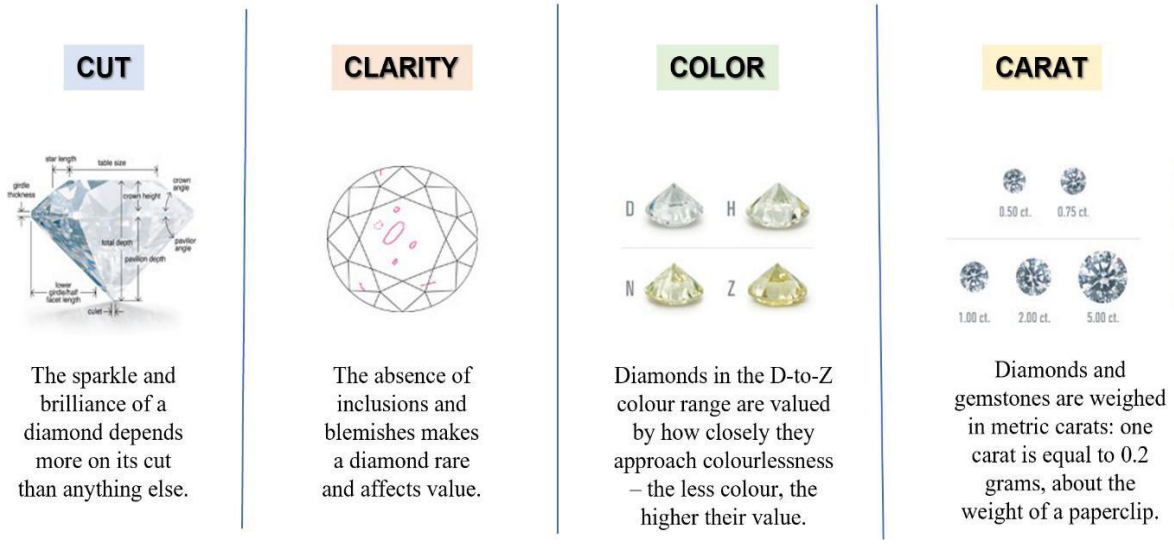

Fig:1 The cut, clarity, color and carat features [https://www.gia.edu/gia-about]

II. BACKGROUND

2.1 Historical Evolution of Diamond Grading

The evaluation of diamonds has historically relied on expert gemologists who visually inspect stones under controlled conditions. Before the mid-20th century, terminology describing diamond quality was inconsistent across regions and markets. The establishment of the Gemological Institute of America (GIA) in 1931 introduced standardized grading protocols, which culminated in the globally recognized "4Cs"—color, clarity, cut, and carat. This framework, developed in the 1950s, remains the cornerstone of diamond evaluation today. However, despite its widespread adoption, grading outcomes are still dependent on human perception and expertise [2], [24].

Early attempts at semi-automation involved spectrophotometric tools for color evaluation and optical measurement devices for cut proportions. While these instruments improved measurement consistency, they did not fully eliminate the subjective component of grading. Over time, industries such as electronics and manufacturing began to integrate machine vision systems for inspection, inspiring research in gemology to explore similar approaches [18], [20].

2.2 The 4Cs and Their Technical Challenges Color

Color

Color grading evaluates the presence of tint in an otherwise colorless diamond. Subtle differences in hue, often imperceptible to untrained observers, significantly affect value. Traditional grading compares diamonds to master stones under standardized lighting conditions. However, factors such as light source variability, grader fatigue, and background interference introduce

inconsistencies. Machine vision-based approaches aim to overcome these issues by quantifying hue, saturation, and brightness through calibrated imaging systems [2], [24].

Clarity
Clarity assessment determines the presence and visibility of internal inclusions or external blemishes. This process is highly subjective, as the relative importance of an inclusion depends on its size, position, and contrast with the surrounding material. Microscopic evaluation is time- intensive and prone to inconsistency across graders. Automated defect detection, powered by high-resolution imaging and deep learning, offers potential for standardizing this process while detecting flaws invisible to the human eye [20], [23].

Cut
Cut quality dictates a diamond's brilliance and scintillation, as it directly influences light reflection and refraction. Gemologists traditionally measure proportions such as table size, crown height, and pavilion depth using optical devices. However, manual assessments can struggle to account for slight asymmetries that impact light performance. AI-driven systems, combined with 3D reconstruction techniques, provide a pathway for highly accurate, non- destructive evaluation of cut precision [25], [28].

Carat
Carat weight is the most straightforward parameter, measured using precision balances. While this process is already objective, its integration into automated grading workflows ensures that all 4Cs are consistently analyzed within a single inspection pipeline. The challenge lies in harmonizing carat data with color, clarity, and cut assessments to generate a holistic grading profile [27], [29].

2.3 Limitations of Traditional Approaches
Although the 4Cs framework standardized diamond grading, its manual implementation presents several drawbacks. First, inter-grader variability remains a persistent challenge, particularly in color and clarity assessments, where small perceptual differences lead to divergent results. Second, manual grading is resource-intensive, requiring significant time and expertise to process even modest quantities of diamonds. This inefficiency becomes more pronounced in large-scale trade environments. Finally, subjective grading creates opportunities for disputes between buyers and sellers, undermining trust in valuation processes [19], [22].

Moreover, global supply chains increasingly demand transparency and traceability, particularly in the context of ethical sourcing. Manual records of diamond attributes are difficult to integrate into digital databases, creating gaps in provenance tracking. Automation not only addresses the technical limitations of manual grading but also aligns with broader trends toward digital certification and blockchain-enabled traceability in the gemstone industry [23], [27].

2.4 Emergence of Automation in Gemology
The use of AI and computer vision to grade diamonds is one aspect of a broader industrial trend to automate quality control. Early breakthroughs in deep learning demonstrated the power of convolutional neural networks (CNNs) for classification and feature extraction in a very broad collection of image datasets such as ImageNet [1], [5]. The models that ensued, like Faster R-CNN, YOLO, and Mask R-CNN, transformed object detection to real-time capability and prepared them for industrial inspection pipelines [10], [14].

In gemology, these tools have been repurposed by researchers for use in tasks like automated color grading, detection of inclusions, and cut analysis. Improved uniformity and narrower error margins are reported in laboratory studies compared to human graders, while collaborations with industry have pushed the implementation of AI-based grading systems at speed. Collaborations between GIA, Razor Labs, and SarineTechnologies provide an example of how research advancements are being developed into commercial solutions [27], [28], [29].

2.5 Rationale for a Comprehensive Review
Although significant progress has been made, automated diamond grading remains fragmented as a discipline. Most research addresses one attribute at a time—e.g., color or clarity—without discussion of integration across the parameters of grading. In addition, issues such as data set limitation, inconsistencies in imaging environments, and the absence of standardized procedures limit widespread implementation. Such gaps indicate the necessity for a thorough review of published work and industrial developments [19], [22].

This overview fills that void by considering earlier work in automated diamond grading methodically, analyzing strengths, weaknesses, and outstanding challenges. By integrating knowledge from academic research and industrial experience, it offers a blueprint for innovation in transforming diamond quality control through automation.

III. LITERATURE REVIEW

3.1 Automated Color Defect Detection

Color assessment is the most important factor in the detection of defects in diamonds. Conventional systems are based on comparison with master stones with controlled illumination, which is subjective and can miss minute differences [2], [18]. Machine vision systems have been implemented to measure color properties like hue, saturation, and brightness and detect anomalies consistently [24], [29].

Early computer techniques used image processing techniques to extract color histograms and spectral features for enabling classification of color variations. More recent deep learning algorithms, namely convolutional neural networks (CNNs), have been used to automatically detect small differences between hue and saturation automatically with faster speed and greater accuracy [5], [8]. Transfer learning from big image databases like ImageNet has enabled training of models even in the absence of much diamond-specific data [1], [4].

3.2 Clarity & Inclusion Detection

Clarity assessment focuses on identifying internal inclusions and surface blemishes, which are primary defects affecting diamond quality. Manual inspection is limited by human perception, especially for micro-scale inclusions. Automated detection methods utilize high-resolution imaging combined with deep learning to locate defects accurately [20], [23].

Object detection models, such as Faster R-CNN and YOLO, have been adapted for identifying inclusions within diamonds. Faster R-CNN excels in detecting small-scale defects due to its region proposal mechanism, while YOLO offers real-time processing for industrial-scale inspection [10], [13]. Mask R-CNN extends these capabilities by providing pixel-level segmentation of defects, which is crucial for evaluating their size, shape, and potential impact on diamond appearance [14], [17].Recent studies have also explored multi-stage networks and feature fusion techniques to improve defect detection sensitivity, integrating color and clarity cues for comprehensive evaluation [26], [30].

3.3 Cut & Shape Defect Detection

The cut of a diamond determines how it interacts with light, influencing brilliance and overall visual appeal. Minor deviations or defects in facet angles, symmetry, or surface polishing can reduce performance. Traditional manual assessment is limited in precision, particularly for small asymmetries [25], [28].

Automated detection employs 3D imaging, photogrammetry, and AI models to evaluate facet angles, symmetry, and surface quality. Deep learning models process multi-angle images to detect deviations from ideal geometries, identifying defects that may not be visible to the naked eye [28], [30]. Emerging approaches combine convolutional and attention-based networks to capture both local and global features, enhancing detection accuracy for complex cut-related defects [14], [16].

3.4 Industrial Applications of AI in Defect Detection

Industry adoption of automated defect detection has accelerated, driven by the need for consistency, efficiency, and traceability. Organizations such as GIA, Razor Labs, and Sarine Technologies have implemented AI-powered systems to inspect diamonds at scale, focusing on color, clarity, and cut defects [27], [29].

These systems integrate high-resolution imaging hardware with deep learning models for real- time assessment. By reducing human error and standardizing evaluation metrics, industrial solutions enable faster throughput, reproducibility, and digital record-keeping. The use of explainable AI and multimodal sensing is gaining traction, allowing practitioners to understand model decisions and combine visual data with physical measurements for comprehensive defect detection [21], [23].

Collectively, these studies and industrial initiatives demonstrate the feasibility and effectiveness of AI-based defect detection in diamonds. While research is progressing rapidly, challenges remain in dataset standardization, multi-parameter integration, and deployment under variable imaging conditions. Addressing these challenges is essential for developing universally accepted automated inspection frameworks.

V. ANALYSIS TABLE

Year	Objective	Method/ Model	Dataset/Tools	Key Findings	Limitations
(2018)[18]	Develop effective image acquisition for diamond inspection	High- resolution imaging system	Custom diamond images	Improved capture of fine details for defect analysis	Focused only on imaging; no AI- based analysis

(2020)[20]	Diamond quality assessment using ML	Machine learning classifiers	Diamond datasets with labeled defects	ML models classified defects with reasonable accuracy	Limited dataset size; color defects underrepresented
(2021)[24]	Measurement of gem color using computer vision	Computer vision & image processing	Laboratory diamond images	Accurate color classification; reduced human error	Did not address clarity or cut defects
(2021)[23]	AI-guided defect detection in crystal growth	Deep learning models for defect detection	Crystal growth dataset	Automated detection of micro-defects	Dataset not specific to diamonds; needs adaptation
(2022)[25]	Automatic gemstone classification	CNN-based classification	Gemstone image dataset	Accurate classification; potential for defect identification	Focused on classification; limited to general gemstones
(2022)[26]	Defect predictionin crystal structures	Feature fusion deep learning	Crystals dataset	Improved detection of structural defects	Not diamond-specific; small dataset
(2023)[27]	AI for diamond grading	AI-powered automated inspection	Industrial diamond images	Real-time color & clarity defect detection	Proprietary system; limited public technical details
(2024)[28]	Cut and shape defect analysis	AI-based image analysis	Industrial diamond datasets	Accurate detection of cut deviations and shape defects	Focused on cut/shape; less on color and clarity
(2025)[30]	Sub micro scale defect detection in diamond tools	DToolnet deep learning model	Diamond grinding processima ges	Detected micro-defects during manufacturing	Focused on tools; not polished gemstones

V. CONCLUSION

This review has emphasized the major advances in automatic defect detection for diamonds covering color, clarity, cut, and shape analysis. Conventional manual inspection is standardized but has limitations in terms of subjectivity, inter-grader variation, and time constraints, which can weaken subtle defect detection. Progress in machine vision, deep learning, and AI-based object detection models has yielded significant strides in accuracy, consistency, and speed, and scalable solutions for both research and industry. Nevertheless, challenges still persist, which include limited data sets, differences in imaging conditions, and requirements for unified, explainable, and affordable systems. Research in the future should aim at combining multi- parameter defect detection using sophisticated imaging methods, standardizing protocols, and constructing modular AI solutions amenable for deployment at scale. By filling these gaps, automated tools can transform diamond quality inspection into a reliable, efficient, and objective evaluation aligned with the needs of contemporary industry.

VI. FUTURE WORK

Future research in automated defect detection in diamonds should focus on integrating all quality parameters—color, clarity, cut, and carat—into a unified inspection framework for holistic and consistent evaluation. The development of large, standardized, and publicly available diamond datasets will improve model robustness and benchmarking, while explainable AI techniques can enhance trust by clarifying model decisions. Optimizing AI models for real-time, high-throughput inspection and integrating advanced imaging technologies, including multi-spectral, 3D, and hyperspectral sensors, can enhance detection

of sub-microscale defects and hidden inclusions. Aligning these systems with industry standards and developing cost-effective, modular solutions will facilitate global adoption and accessibility. Addressing these areas collectively will enable reliable, accurate, and efficient automated defect detection systems, ultimately revolutionizing diamond quality control.

REFERENCES

[1] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2009, pp. 248–255, doi: 10.1109/CVPR.2009.5206848.

[2] Z. Ren, J. Liao, and L. Cai, "Diamond color grading based on machine vision," in Proc. IEEE 12th Int. Conf. Computer Vision Workshops (ICCVW), 2009, pp. 1970–1976, doi: 10.1109/ICCVW.2009.5457523.

[3] B. Settles, Active learning literature survey, Computer Sciences Technical Report 1648, Univ. of Wisconsin–Madison, 2009. [Online].Available: http://burrsettles.com/pub/settles.activelearning.pdf

[4] S. J. Pan and Q. Yang, "A survey on transfer learning," IEEE Transactions on Knowledge and Data Engineering, vol. 22, no. 10, pp. 1345–1359, 2010, doi: 10.1109/TKDE.2009.191.

[5] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in Proc. Advances in Neural Information Processing Systems (NeurIPS), 2012, pp. 1097–1105. [Online]. Available: https://dl.acm.org/doi/10.1145/3065386

[6] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2014, pp. 580–587, doi: 10.1109/CVPR.2014.81.

[7] T.-Y. Lin et al., "Microsoft COCO: Common objects in context," in Proc. European Conf. Computer Vision (ECCV), 2014, pp. 740–755, doi: 10.1007/978-3-319-10602-1_48.

[8] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," arXiv preprint arXiv:1409.1556, 2014. [Online]. Available: https://arxiv.org/abs/1409.1556

[9] R. Girshick, "Fast R-CNN," in Proc. IEEE Int. Conf. Computer Vision (ICCV), 2015, pp. 1440–1448, doi: 10.1109/ICCV.2015.169.

[10] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," Advances in Neural Information Processing Systems (NeurIPS), 2015. [Online]. Available: https://arxiv.org/abs/1506.01497

[11] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770–778, doi: 10.1109/CVPR.2016.90.

[12] W. Liu et al., "SSD: Single shot MultiBox detector," in Proc. European Conf. Computer Vision (ECCV), 2016, pp. 21–37. [Online]. Available: https://arxiv.org/abs/1512.02325

[13] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real- time object detection," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2016, pp. 779–788. [Online]. Available: https://arxiv.org/abs/1506.02640

[14] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," in Proc. IEEE Int. Conf. Computer Vision (ICCV), 2017, pp. 2961–2969. [Online]. Available: https://arxiv.org/abs/1703.06870

[15] J. Redmon and A. Farhadi, "YOLO9000: Better, faster, stronger," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2017, pp. 6517–6525. [Online]. Available: https://arxiv.org/abs/1612.08242

[16] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder-decoder with atrous separable convolution for semantic image segmentation," arXiv preprint arXiv:1802.02611, 2018. [Online]. Available: https://arxiv.org/abs/1802.02611

[17] J. Redmon and A. Farhadi, "YOLOv3: An incremental improvement," arXiv preprint arXiv:1804.02767, 2018. [Online]. Available: https://arxiv.org/abs/1804.02767

[18] W. Wang, "On the development of an effective image acquisition system for diamond inspection," Applied Optics, vol. 57, no. 33, pp. 9887–9897, 2018, doi: 10.1364/AO.57.009887.

[19] Z. Q. Zhao, P. Zheng, S. T. Xu, and X. Wu, "Object detection with deep learning: A review," IEEE Transactions on Neural Networks and Learning Systems, vol. 30, no. 11, pp. 3212–3232, 2019, doi: 10.1109/TNNLS.2018.2876865.

[20] M. Navale, M. Bikchandani, V. Bhosale, P. Bhosale, and S. Ghadge, "Diamond quality assessment system using machine learning approach," International Research Journal of Engineering and Technology (IRJET), vol. 7, no. 4, pp. 874–877, Apr. 2020. [Online].
Available: https://www.irjet.net/archives/V7/i4/IRJET-V7I4163.pdf

[21] S. S. A. Zaidi, M. S. Ansari, A. Aslam, N. Kanwal, M. Asghar, and B. Lee, "A survey of modern deep learning based object detection models," arXiv preprint arXiv:2104.11892, 2021. [Online]. Available: https://arxiv.org/abs/2104.11892

[22] J. Wang, T. Zhang, Y. Cheng, and N. Al-Nabhan, "Deep learning for object detection: A survey," Computer Systems Science & Engineering, vol. 38, no. 2, pp. 165–182, 2021, doi: 10.32604/csse.2021.017016.

[23] R. R. Mekala, E. Garratt, M. Muehle, A. Srinivasan, A. Porter, and M. Lindvall, "AI- guided defect detection techniques

to model single crystal diamond growth," Materials Science Forum, vol. 1036, pp. 163–176, 2021, doi: 10.4028/www.scientific.net/MSF.1036.163.

[24] S. Zhang, X. Wang, and Y. Li, "Measurement of gem colour using a computer vision system," Minerals, vol. 11, no. 8, p. 791, 2021, doi: 10.3390/min11080791.

[25] J. C. L. Chow and C. C. Reyes-Aldasoro, "Automatic gemstone classification using computer vision," Minerals, vol. 12, no. 1, p. 60, 2022, doi: 10.3390/min12010060.

[26] F. K. Alarfaj and H. Mahmoud, "Feature fusion deep learning model for defects prediction in crystal structures," Crystals, vol. 12, no. 9, p. 1324, 2022, doi: 10.3390/cryst12091324.

[27] GIA & IBM Research, "GIA to use advanced artificial intelligence for diamond grading," Solitaire International (GJEPC), 2023. [Online]. Available: https://gjepc.org/solitaire/gia-to- use-advanced-artificial-intelligence-for-diamond-grading

[28] Razor Labs & GIA, "AI-driven diamond cut and shape analysis," Razor Labs, 2024. [Online]. Available: https://www.razor-labs.com/gia-razor-labs-diamond-ai

[29] Sarine Technologies, "Sarine diamond color grading," Sarine, 2024. [Online]. Available: https://sarine.com/products/sarine-diamond-color-grading

[30] Y. Li, W. Liu, and H. Zhang, "On-machine detection of sub-microscale defects in diamond tool grinding during the manufacturing process based on DToolnet," Sensors, vol. 25, no. 7, p. 2426, 2025, doi: 10.3390/s22072426.

SPU Journal of Science, Technology and Management Research (SPU-JSTMR) Special Issue-ICAICS-26
ISSN: 3049-1479(Online)

Advancements in AI-based Robotics and Automation with Emerging Cybersecurity Threats and Challenges

Doshi Krisha[1*], Shah Diya[2], Deep Solanki[3], Jigar Gajjar[4]

Lok Jagruti Kendra University, Ahmedabad, GJ, India[1*]
Lok Jagruti Kendra University, Ahmedabad, GJ, India[2]
TechD Cybersecurity Limited, Ahmedabad, GJ, India[3]
TechD Cybersecurity Limited, Ahmedabad, GJ, India[4]

krishadoshi444@gmail.com[1*],shahdiya1164@gmail.com[2],deepsolanki5799@gmail.com[3],cyberian.jg@gmail.com[4]

Abstract: Mechanics and automation technology provide the capability for machines to perform actions with minimal involvement of the human factor, which positively affects the efficiency, precision, and safety of the process and products created by a machine. Advanced robotic systems integrate sensors, control systems, and artificial intelligence for the purpose of facilitating intelligent, adaptable, and automated technology in various industries. The incorporation of technology in the field of robotics and automation has enabled traditional systems to advance from fixed and pre-programmed operation systems to flexible and data-based systems that are able to make real-time operational decisions. In the paper, there will be a discussion about the contributions and implications of the incorporation of technology in modern systems per-training to issues of productivity, reliability, and human-robot interaction. The increasing incorporation of intelligent robotic systems into automation systems reveals its necessity in developing a smart industry.

Keywords: Robotics, Automation, Intelligent Systems, Control Algorithms, Artificial Intelligence, Sensors and Actuators, Robotic Automation, Human–Robot Collaboration, Adaptive Systems, Industry 4.0.

I. INTRODUCTION

Robotics is the study and application of robots, which are machines that can perform automatically or with very little human intervention [1]. A robot can perceive its environment, process information according to a program or AI, and act by moving its arms, wheels, or tools. Robots are very useful for performing tasks that are repetitive, hazardous, or require a high degree of accuracy[13].Automation is the use of machines, software, or control systems to perform tasks automatically without continuous human intervention. Automation is more concerned with process control than robotics.

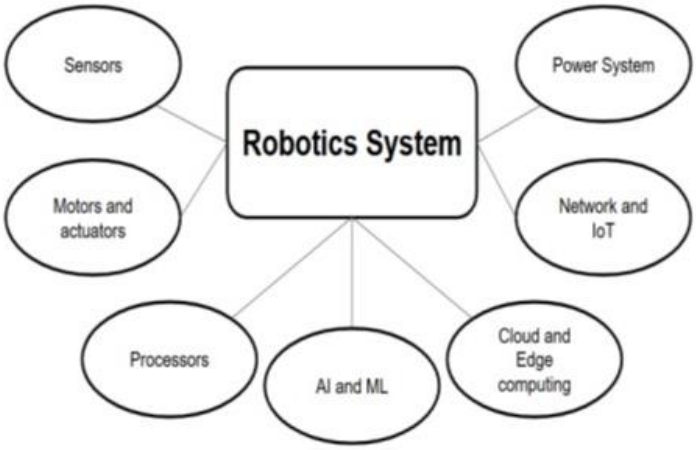

Fig.1.Components of Robotics System

http://doi.org/10.63766/spujstmr.26.000103

Robotics and automation are some of the latest technologies that have proved to be highly influential in the current development of industry and technology[4]. Robotics is an area that deals with the design and implementation of programmable machines that can carry out tasks with a high degree of precision, while automation is the use of control systems, software, and intelligent technologies to minimize human interaction in a process[7].

Robotics is a field that involves designing, building, controlling, and using robots that can work on their own or with some human help [2][3]. It brings together ideas from different areas like mechanical and electrical engineering, computer science, control systems, and artificial intelligence to create machines that can sense their environment, understand information, and take physical actions [2][7]. Modern robots usually have sensors to see and feel, control systems to make decisions, and motors or actuators to move and do tasks. Robotics is used in many industries such as manufacturing, healthcare, delivery, space missions, farming, and services to help improve efficiency, accuracy, and safety [4][8]. The field keeps growing, focusing on making robots smarter, more flexible, and able to work closely with people in changing environments.

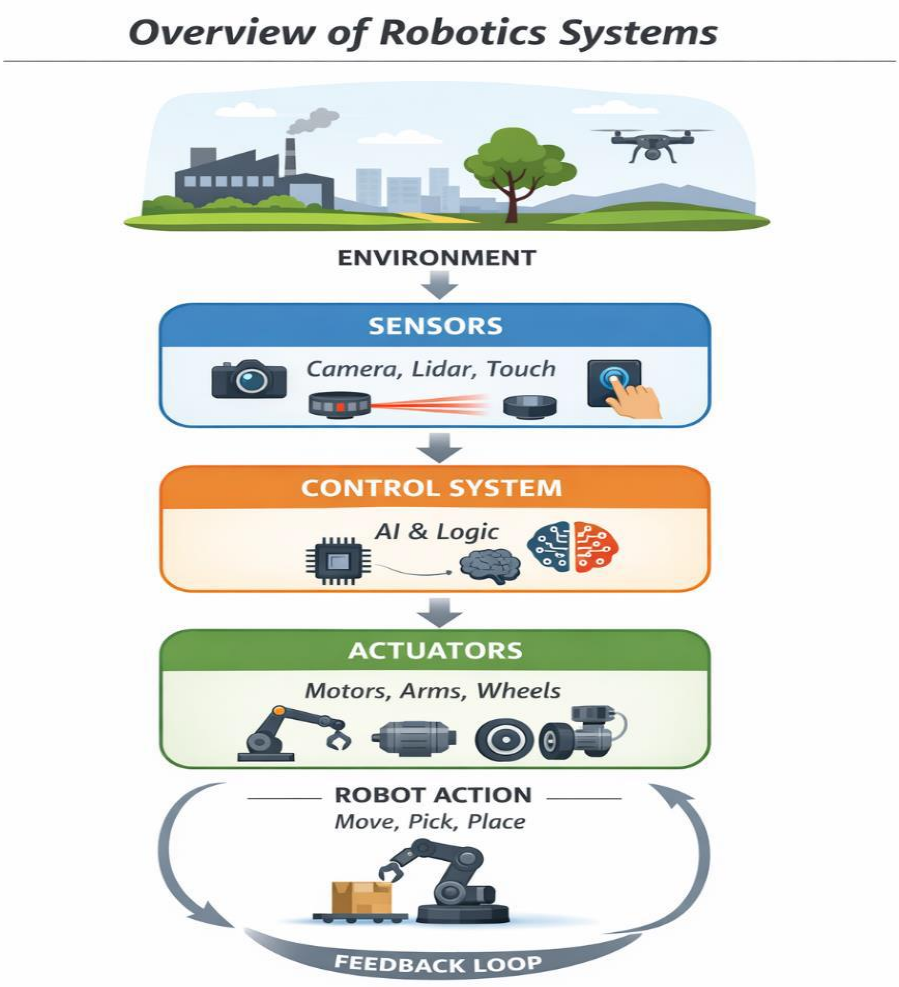

Fig. 2. Overview of the Robotic System

1.1 Evaluation Of Robotics

The evolution of robotics began with simple mechanical machines designed to perform repetitive tasks during the early industrial era. With the introduction of industrial robots in the 1950s and 1960s, automation transformed manufacturing by improving efficiency and precision[15]. Advances in electronics, sensors, and computing led to more intelligent robots capable

of perception and basic decision-making. Today, robotics integrates artificial intelligence, machine learning, and IoT, enabling robots to operate autonomously in fields such as healthcare, defense, space exploration, and smart industries[16].

In the last few years, collaborative robots and autonomous systems have appeared that safely cooperate with human beings in common workspaces. The development of cyber–physical systems and cloud robotics should ensure enhanced connectivity, scalability, and real-time data exchange. As robotics continues to evolve, the focus is shifting toward self-learning, adaptive, and context-aware systems that can function reliably in highly dynamic and unstructured environments[17].

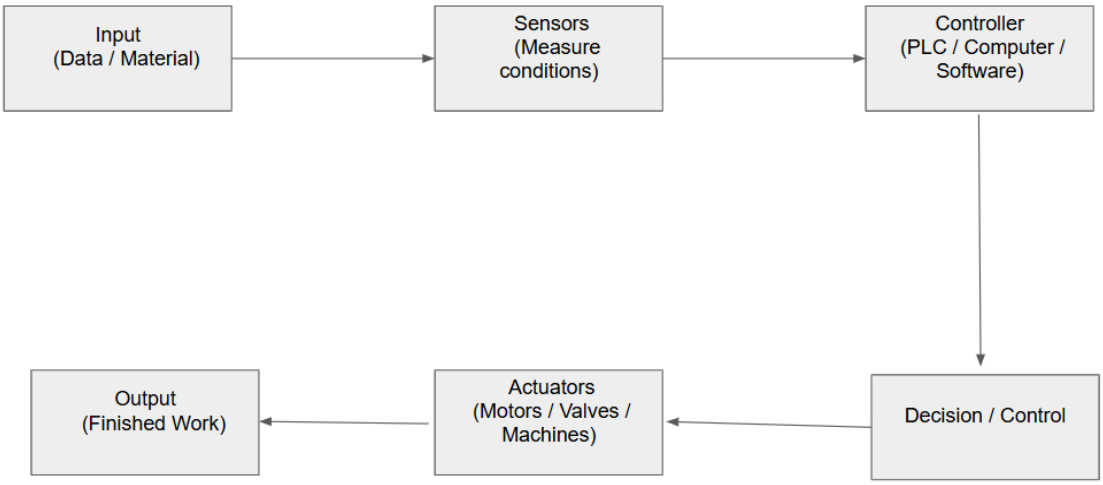

Fig.3 Evaluation of Robotics

1.2 *Major Components and Technologies Used in Robotics*

Robotic systems are made up of different critical components and technologies that work together to perform tasks efficiently and safely. One of the most important components of a robotic system is the sensor. Sensors can detect light, sound, distance, temperature, pressure, and motion, which helps the robot understand its surrounding environment and respond appropriately [2]. Another key component is the actuator, which may be an electric motor, hydraulic system, or pneumatic system. Actuators are mainly responsible for movement and enable the robot to perform physical tasks such as lifting, turning, gripping, or walking [3]. The control system acts as the brain of the robot by processing sensor data and coordinating actuator actions [3].

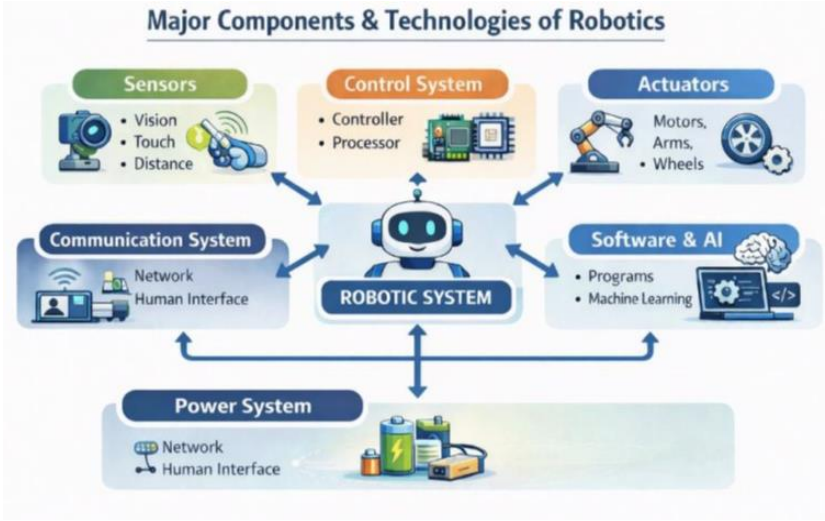

Fig. 4 Major components & Technologies

Software and artificial intelligence (AI) also play an important role in robotics. Software defines how the robot operates and controls its tasks. AI enables robots to think, learn, and make decisions independently. Technologies such as machine learning and computer vision allow robots to recognize objects, navigate safely, and improve-performance by learning from previous experiences [7]. Robots also require power sources, such as batteries or external power supplies, to operate continuously. In addition, communication technologies help robots exchange information with other machines, computer systems, or human operators. Because of these technologies, robots become more intelligent, flexible, and useful in many fields, including manufacturing, healthcare, agriculture, and service industries [2][7].

1.3 Automation of robotics using Al and its use cases:

Refers to the use of control systems and mechanical technology to control machines and industrial processes with minimal human interaction[4]. Human-based automation involves the use of human operators in the system setup, monitoring, control, and maintenance.[9] The system is intended to help human operators by reducing human effort, increasing the speed of operations, and improving accuracy, as opposed to re-placing human judgment[10].

Automation increases productivity because it allows machines to handle repetitive, time-consuming, and physically demanding tasks at a very high speed and accuracy rate, and human intervention is used for monitoring and managing the system's performance [11]. This method helps to minimize manual labor and prevents workers from getting physically strained and mentally stressed, as they can concentrate on planning, monitoring, and decision-making tasks [12]. Human-based automation is also very important for enhancing workplace safety because it allows dangerous tasks to be handled by automated systems, which greatly minimizes the chances of accidents and injuries [13]. Machines are very efficient in performing precision-based tasks, and human intervention is also required for dealing with unexpected events and managing the system's control effectively. The effective combination of machine capabilities and human intelligence greatly improves system flexibility, reliability, and performance [11][12].

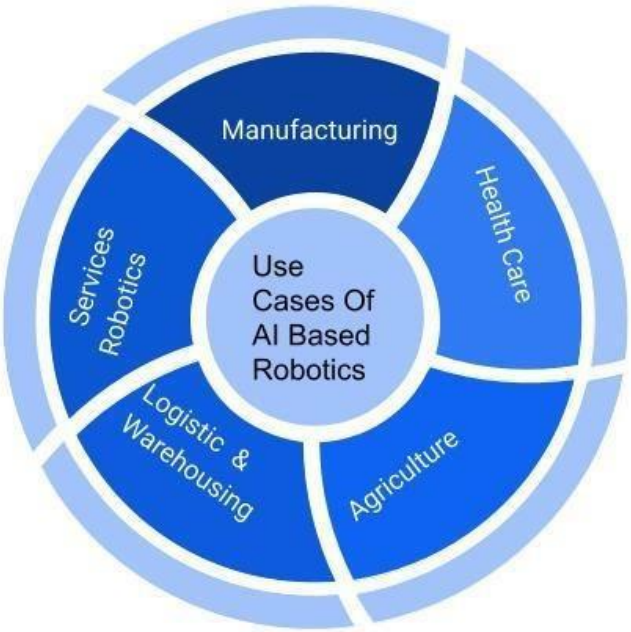

Fig.5 Use Cases of AI based Robotics

1.4 Advantages/benefits of robotics automation

Robotic automation is the use of robots and machines to perform tasks with very little human assistance [3]. The robots are employed in tasks that are repetitive, take a long time, are hazardous, or require precision[4]. The use of robots helps in completing tasks faster, safer, and better[4]. The first advantage of robotic automation is that tasks are completed faster. Robots can work for 24 hours without getting tired or needing a break[3]. This helps companies complete more tasks in a shorter time. Another ad-vantage is that the results are more accurate. Robots always follow instructions, so they make fewer errors and produce the same positive results[2][4].

Usefulness of Robotics Automation: Robotics automation is very useful as it assists humans to work better and safely[2]. Robots can work for a long time without getting tired, which helps in increasing the speed of work and productivity [4]. As robots work according to programmed instructions, they work with accuracy, which reduces errors and improves the quality of work[3]. Another very important advantage is that it is safe. Robots are capable of performing heavy, repetitive, or hazardous tasks, such as lifting heavy weights or operating in hazardous environments, thus minimizing the chances of injuries to employees[2][4]. Robotics automation also helps in minimizing physical and mental efforts of employees by performing mundane and tiresome Tasks. This enables employees to engage in more significant tasks, such as system monitoring, problem-solving, and decision-making. Eventually, robotics automation assists in timesaving and the creation of a safe and organized workplace[2][3].

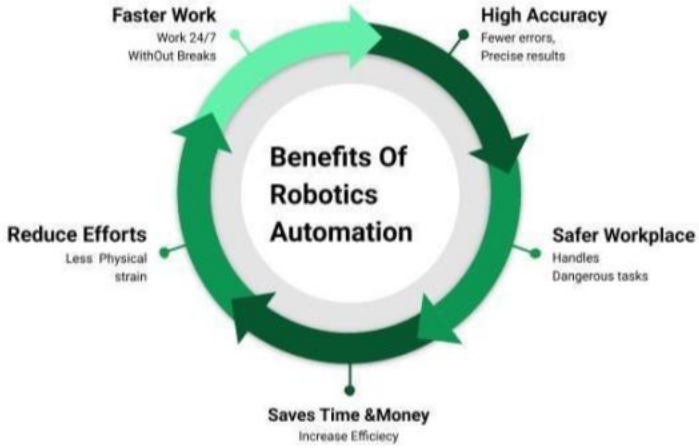

Fig.6 Benefits of Robotics Automation

1.5 Security challenges in robotics automation:

The security risks posed by robotics automation systems are considerable, as they rely on industrial control system architectures, communication systems, and human interaction [17].

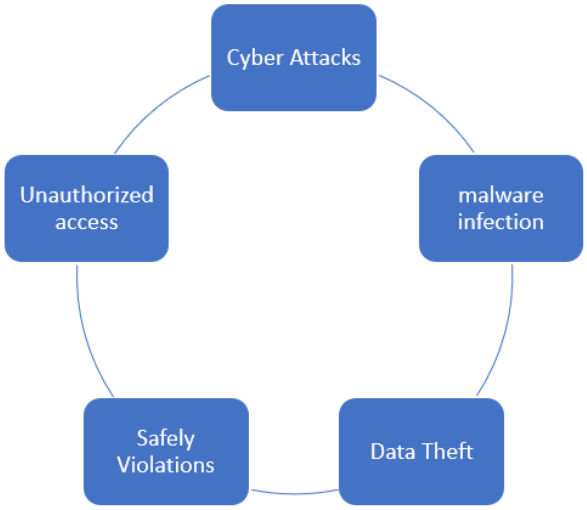

Fig. 7 Challenges of Robotics Automation

Human-based robotic systems function on Programmable Logic Controllers (PLCs), Supervisory Control and Data Acquisition (SCADA) systems, and field devices that are connected through industrial communication systems[10]. If such systems are not properly secured, they can provide opportunities for unauthorized access, which can enable attackers to manipulate control logic, modify operating parameters, or stop the operation of robots[20].

Moreover, the deployment of more networked and cloud-dependent robotic technologies has created a large threat environment against cyber-attacks like malware, denial of service, and data breaches. Inadequate monitoring facilities, handling of legacy components, and poor authentication methods are also critical issues related to security threats. This must be effectively addressed to provide a secure and reliable robot automation system [21].

1.7 Literature Review and Research Gap

According to recent studies, robotics and automation technologies have been found to enhance productivity, precision, and efficiency within various sectors. Studies have investigated how AI, ML, and IoT technologies can be incorporated into robots to facilitate autonomous decision-making, adaptability, and intelligent process management [2][4][7]. These innovations have expanded the application of robots in sectors such as manufacturing, health care, logistics, agriculture, and smart cities [8][13].

While several research studies emphasize the positive aspects of robotics automation, very few of them address any issues other than the improvements in performance that can be achieved through automation. Very little has been done to study the possible threats that robots pose with respect to their cyber security aspect, particularly because contemporary robots depend heavily on cloud platforms and wireless communication networks [16][17][20].

Furthermore, existing research often discusses security threats individually rather than providing integrated security frameworks for robotic environments. Therefore, there is a need for comprehensive research that examines robotics automation technologies together with emerging cybersecurity challenges and mitigation approaches. This study attempts to address this gap by providing an integrated overview of robotics, automation, AI integration, and cybersecurity concerns [17][18][20].

II. LITERATURE REVIEW

2.1 History and evolution of robotics and robo-automation

History and development of robotics and robotic automation started in the early industrial era with the use of simple mechanical systems to aid human labour [3]. The development of the first industrial robot, named Unimate, in the 1950s was an important milestone in robotics automation because it introduced automation in manufacturing systems [4]. In the latter half of the 20th century, improvements in electronics, sensors, and computer controls made it possible for robots to perform more complex tasks [3][2]. With the advent of the current era, the combination of artificial intelligence, machine learning, and IoT has led to the development of intelligent and autonomous robots in the medical, transportation, space, and smart industry sectors, thus expanding the horizon of robotics automation [4]. The current development in robotic automation has led to increased productivity, accuracy, and safety at the workplace [4]. The current scenario of robotic automation also allows the analysis of data in real-time and has the capability to monitor from a distance, thus making the process more adaptable and efficient [2]

2.2 Table of technologies components used in robo-automation

Below shows the technologies components used in robo-automation and robo-automation use cases in real life in table 1 & 2 respectively

TABLE I

TECHNOLOGIES COMPONENTS USED IN ROBO-AUTOMATION

Technology Component	**Description / Role in Robo-Automation**
Sensors	Help robots sense their environment by detecting light, distance, temperature, pressure, sound, and motion.

Actuators	Enable movement and physical actions such as lifting, rotating, gripping, or walking using motors or hydraulic systems.
Control System	Acts as the brain of the robot; processes sensor data and sends commands to actuators.
Robotic Software	Controls robot operations, task execution, and system coordination.
Artificial Intelligence (AI)	Allows robots to make decisions, learn from data, and adapt to new situations.
Machine Learning	Enables robots to improve performance over time through experience and data analysis.
Computer Vision	Helps robots recognize objects, read images, and navigate safely in their environment.
Power Supply	Provides energy to robots through batteries, power adapters, or external power sources.
Communication Systems	Allow robots to communicate with other machines, networks, cloud systems, or human operators.
IoT & Connectivity	Enables real-time monitoring, data sharing, and remote control of robotic systems.

TABLE II

ROBO-AUTOMATION USE CASES IN REAL LIFE

Use Case / Sector	**Application of Robo-Automation**	**Key Benefits**
Manufacturing	Robots perform assembly, welding,painting,and packaging tasks.	Faster production, high accuracy, consistent quality, reduced human effort.
Healthcare	Surgical robots, patient monitoring, medicine delivery.	Improved precision, safer surgeries, reduced workload on healthcare staff.
Logistics & Ware-housing	Automated picking, sorting, and material handling.	Faster order processing, reduced errors, efficient inventory management.
Agriculture	Automated planting, harvesting, and crop monitoring.	Increased productivity, reduced labor cost, better crop yield.
Transportation	Autonomous vehicles and traffic management systems.	Improved safety, reduced human error, efficient transport operations.
Defense & Security	Surveillance robots and bomb disposal systems.	Enhanced safety for personnel, effective threat detection.
Retail	Automated checkout systems and inventory robots.	Faster service, better stock control, improved customer experience.
Service Industry	Cleaning robots, delivery robots, customer assistance robots.	Reduced manual labor, improved efficiency, continuous operation.
Space Exploration	Planetary rovers and robotic arms for exploration.	Operates in extreme environments, accurate data collection.
Smart Homes	Robotic vacuum cleaners and smart assistants.	Convenience, time-saving, improved daily life quality.

2.3 Case Study: AI-Based Robotic Automation in Smart Manufacturing

There is an example of robotics in automation, which may involve intelligent factories that use collaborative robots working alongside humans. Robots are equipped with the use of sensor technology, computer vision, and AI systems that will help robots perform tasks in assembly, packing, testing, and quality control operations [2][21].

To give an example, robotic arm with AI technology will use imaging technology to find any faults in the products produced and segregate such faulty products from the production line. At the same time, the technique of predictive maintenance is being used to find any potential problems in the machine and inform the operator about such problems [4][7].

There are many benefits of robotic automation by the use of artificial intelligence like increasing efficiency in production, reducing the cost of operations, improving the quality of products, enhancing the safety of employees, and minimizing downtime by the use of predictive maintenance [11][13]. However, there are certain security threats related to these technologies due to their network-based functioning [16][17].

TABLE III.

COMPARATIVE REVIEW OF PREVIOUS STUDIES

Study	Focus Area	Major Contribution	Limitation
Craig (2018) [3][6]	Robotics Fundamentals	Provides comprehensive coverage of robotics principles, kinematics, dynamics, and control systems	Limited discussion on cybersecurity aspects
Groover (2015) [4]	Industrial Automation	Demonstrates improvements in production efficiency through automation technologies	Primarily focused on manufacturing sectors
Siegwart et al. (2011) [7]	Autonomous Robotics	Advances autonomous navigation, perception, and decision-making in robotic systems	Limited analysis of security threats
Goldberg and Kehoe (2013) [16]	Cloud Robotics	Introduces real-time monitoring and cloud-based robotic coordination	Increased exposure to cyberattacks
Proposed Study	Robotics, AI, and Cybersecurity	Integrates applications, benefits, security challenges, and mitigation strategies	Requires future experimental validation

2.4 Security challenges in robo-automation

2.4.1 Cybersecurity Threats and Hacking

One of the major security concerns in robo-automation is the potential for a cyber attack. This is because robots are network and cloud-based, as well as Internet of Things (IoT) devices, which makes it possible for hackers to exploit vulnerabilities in order to gain unauthorized access. A cyber attack may lead to loss of control, data theft, system shutdown, or unsafe robot behavior. Cyber attacks may lead to financial losses, security risks, and operational disruptions[18].

2.4.2 Data Privacy and Information Leakage

The data collected by the robotic automation systems includes information from sensors, cameras, and users. The information may include sensitive data such as operations, employees, and customers. In case the data is not protected by proper security measures, it may be leaked or misused. This may result in privacy breaches and legal problems. It is, therefore, important to ensure that data is stored and accessed in a secure manner that meets data protection requirements[16].

2.4.3 Safety Risks Due to System Failures

The information gathered by the robotic automation systems includes data from sensors, cameras, and users. The data may include sensitive information such as operations, employees, and customers. In the event that the data is not protected by

appropriate security measures, it may be leaked or misused. This may lead to breaches of privacy and legal issues. It is, therefore, important to ensure that data is stored and accessed in a manner that is secure and complies with data protection standards[2].

2.4.4 Weak Authentication and Access Control

In most robotic systems, strong passwords and access control are not used. If unauthorized users are able to access the system, they will be able to change the settings of the robots or control the operations of the robots. To avoid this problem, strong passwords, role-based access, and authentication checks should be used to secure the robotic systems from unauthorized access.[17]

2.4.5 Software Vulnerabilities and Updates

Software is a very crucial component in the proper working of a robot. If the software contains bugs or is outdated, it may cause vulnerabilities that can be exploited by attackers. Outdated software can cause system crashes or malfunctions in robots. Software testing and patching are very crucial in ensuring that the robotic automation sys-tem is secure and reliable[19].

2.4.6 Dependence on Network Connectivity

It has also been observed that robotic automation systems require constant connectivity of the network or the cloud [16]. If the network is lost or the robot is subject to an attack, the robot will either stop functioning or will function incorrectly [17]. Secure networks can help to reduce problems related to network dependency . In this case, small network disruptions can cause problems for the robots . It has also been observed that the network dependency of the robotic automation system makes the robot vulnerable to denial of service (DoS) attacks, which can disrupt communication between the robot and the control system [18].

2.5 Cybersecurity Mitigation Strategies for Robo-Automation Systems

The increasing adoption of intelligent robotic systems requires strong cybersecurity measures to ensure operational safety and reliability. Several security mechanisms can be implemented to reduce cyber risks and improve system resilience.

2.5.1 Network Security

Industrial robotic networks should be protected using firewalls, intrusion detection systems (IDS), virtual private networks (VPNs), and secure communication protocols. Network segmentation can help prevent attackers from accessing critical robotic infrastructure [18].

2.5.2 Strong Authentication and Access Control

Multi-factor authentication (MFA), role-based access control (RBAC), and strong password policies should be implemented to restrict unauthorized access to robotic systems. These measures significantly reduce the risk of unauthorized manipulation of robotic operations [17][20].

2.5.3 Software Updates and Patch Management

Regular software updates and security patches are essential to eliminate vulnerabilities and improve system stability. Organizations should establish routine maintenance procedures to ensure all robotic software remains secure and up-to-date [19].

2.5.4 Data Encryption

Sensitive operational and user data should be encrypted during both storage and transmission. Encryption protects information from unauthorized access and reduces the risk of data leakage [20].

2.5.5 Continuous Monitoring and Threat Detection

Real-time monitoring systems and AI-based threat detection mechanisms can identify abnormal activities and potential attacks at an early stage, allowing organizations to respond quickly and minimize damage [17][18].

2.5.6 Security Awareness Training

Human error remains one of the most common causes of cybersecurity incidents. Regular training programs can help employees recognize cyber threats and follow secure operational practices [9][12].

III. ANALYSIS

The analysis of robotics and automation technologies demonstrates that modern intelligent systems have significantly transformed traditional industrial operations into highly efficient, data-driven environments. The integration of Artificial Intelligence (AI), Machine Learning (ML), Internet of Things (IoT), computer vision, and advanced control systems has enabled robots to perform tasks with greater accuracy, adaptability, and autonomy than conventional automation systems [2][4][7].

From an industrial perspective, robotic automation has substantially improved productivity by reducing human involvement in repetitive, hazardous, and precision-oriented tasks. Manufacturing industries utilize robotic systems for assembly, welding, packaging, and quality inspection, resulting in faster production rates and improved product consistency [4][13]. Similarly, healthcare organizations employ robotic-assisted surgical systems and automated patient-monitoring solutions to improve treatment precision and reduce human error [8].

The adoption of collaborative robots (cobots) has further enhanced workplace efficiency by enabling safe interaction between humans and machines within shared environments [21]. Unlike traditional automation systems, which often operate in fixed and predictable conditions, modern robotic systems can adapt to dynamic environments through real-time data analysis and intelligent decision-making capabilities [7][16].

However, the increasing connectivity of robotic systems introduces significant cybersecurity concerns. The integration of cloud computing, IoT platforms, and industrial communication networks expands the attack surface and creates opportunities for cyber threats such as unauthorized access, malware attacks, denial-of-service attacks, and data breaches [17][18][20]. These threats can affect not only information security but also operational continuity, equipment reliability, and human safety.

Therefore, the successful implementation of robotics automation requires a balanced approach that combines technological innovation with robust cybersecurity protection. Organizations should integrate security measures such as strong authentication mechanisms, network segmentation, encryption, continuous monitoring, and regular software updates into robotic system architectures. Such measures are essential for ensuring the long-term reliability, safety, and sustainability of intelligent robotic systems in Industry 4.0 environments [17][18][20].

Overall, the findings indicate that robotics and automation will continue to play a vital role in shaping future industries by improving operational efficiency, supporting human workers, and enabling smart decision-making processes. As technologies continue to evolve, intelligent robotic systems are expected to become more autonomous, secure, and capable of operating effectively in increasingly complex environments.

IV. CONCLUSION

Robotics and automation technologies have become fundamental components of modern industrial and service sectors, significantly improving operational efficiency, accuracy, productivity, and workplace safety. The integration of Artificial Intelligence (AI), Machine Learning (ML), Internet of Things (IoT), advanced sensors, and intelligent control systems has enabled robots to perform complex tasks with minimal human intervention while maintaining high levels of precision and reliability [2][4][7].

This study examined the evolution of robotics, major technological components, AI-based automation applications, real-world use cases, benefits, and emerging cybersecurity challenges associated with modern robotic systems. The findings indicate that intelligent robotic systems have transformed traditional automation into adaptive, data-driven, and collaborative environments that support Industry 4.0 initiatives and smart manufacturing practices [11][13].

Despite the numerous advantages of robotic automation, increasing connectivity through cloud computing, IoT platforms, and industrial communication networks introduces significant security risks. Cyber threats such as unauthorized access, malware attacks, data breaches, and denial-of-service attacks can negatively impact operational continuity, data integrity, and physical safety [17][18][20]. Therefore, implementing robust cybersecurity measures including strong authentication, encryption, continuous monitoring, software updates, and secure network architectures is essential for ensuring the reliability and resilience of robotic systems.

Future research should focus on the development of intelligent cybersecurity frameworks specifically designed for autonomous robotic environments. Emerging technologies such as blockchain-based security mechanisms, AI-driven threat detection systems, digital twins, and self-healing cyber-physical systems offer promising opportunities for enhancing robotic security and operational reliability. Additionally, experimental studies and real-world industrial case analyses are needed to evaluate the effectiveness of these approaches in dynamic and large-scale deployment environments.

In conclusion, robotics and automation will continue to play a critical role in shaping the future of smart industries, healthcare, transportation, agriculture, and service sectors. With continuous technological advancements and stronger cybersecurity integration, intelligent robotic systems are expected to become more autonomous, secure, and capable of supporting sustainable industrial growth in the coming years.

REFERENCES

[1] G. Bekey, Robotics: State of the Art and Future Challenges. World Scientific, 2008.

[2] B. Siciliano and O. Khatib, Springer Handbook of Robotics, Springer, 2016.

[3] J. J. Craig, Introduction to Robotics: Mechanics and Control, Pearson, 2018.

[4] .M. P. Groover, Automation, Production Systems, and CIM, Pearson, 2015.

[5] Asimov, I, Robot, Doubleday, 1950.

[6] J. J. Craig, Introduction to Robotics: Mechanics and Control, 4th ed. Boston, MA, USA: Pearson, 2018.

[7] R. Siegwart, I. R. Nourbakhsh, and D. Scaramuzza, Introduction to Autonomous Mobile Robots, 2nd ed. Cambridge, MA, USA: MIT Press, 2011

[8] M. Ben-Ari and F. Mondada, Robots and Their Applications, Springer, 2018

[9] T. B. Sheridan, Humans and Automation: System Design and Research Issues. New York, NY, USA: Wiley, 2002.

[10] S. Y. Nof, Ed., Handbook of Industrial Robotics, 2nd ed. Hoboken, NJ, USA: Wiley, 2017.

[11] K. Schwab, The Fourth Industrial Revolution. Geneva, Switzerland: World Economic Fo-rum, 2016.

[12] G. Salvendy, Handbook of Human Factors and Ergonomics, 4th ed. Hoboken, NJ, USA: Wiley, 2012.

[13] International Federation of Robotics, "World Robotics Industrial Robots," IFR, Frankfurt, Germany, 2022.

[14] R. C. Arkin, Behavior-Based Robotics, MIT Press, 1998.

[15] G. N. DeSouza and A. C. Kak, "Vision for mobile robot navigation: A survey," IEEE Trans-actions on Pattern Analysis and Machine Intelligence, vol. 24, no. 2, pp. 237–267, Feb. 2002.

[16] K. Goldberg and M. Kehoe, "Cloud robotics and automation: A survey of related work," IEEE Robotics & Automation Magazine, vol. 20, no. 2, pp. 22–35, Jun. 2013.

[17] Humayed, J. Lin, F. Li, and B. Luo, "Cyber-physical systems security—A survey," IEEE Internet of Things Journal, vol. 4, no. 6, pp. 1802–1831, Dec. 2017.

[18] K. Stouffer, V. Pillitteri, S. Lightman, M. Abrams, and A. Hahn, Guide to Industrial Control Systems (ICS) Security, NIST SP 800-82, 2015.

[19] E. Byres and J. Lowe, "The myths and facts behind cyber security risks for industrial control systems," VDE Kongress, 2004.

A Smart IoT–AI Integrated System for Accurate Rainfall Prediction

Ms. Bhatt Shreyaben Atulbhai[1] ,Dr. Khushbu[2], Dr. Darshanaben Dipakkumar Pandya[3]

Research Scholar, Madhav University Abu Road, Pindwara, Sirohi, Rajasthan *[1]
Assistant professor , Department of Computer Science & Application, Madhav University Abu Road ,Pindwara, Sirohi, Rajasthan [2]
Associate Professor, Department of Computer Science & Application, Shri C. J Patel College of Computer Studies (BCA), Sankalchand Patel University, Visnagar[3]

sabbscit_spce@spu.ac.in[1],yadavkhushbu289@yahoo.com[2],ddpandya.fcs@spu.ac.in[3]

Abstract: This paper describes a new IoT-based system that can keep an eye on the weather in real time. The system uses the internet to connect many sensors so that it can give you constant weather updates. All of the data that is collected is sent to the cloud, where users can get to it from anywhere.

This makes it easier to keep an eye on the weather, especially in rural or agricultural areas. When the weather is bad, traditional wired and analog devices often don't work. The system uses wireless sensors to collect data reliably to fix this problem. These sensors keep track of important things like light, temperature, humidity, and rainfall. The wireless setup makes monitoring easier and more accurate by cutting down on manual work.

The system is very helpful for farmers who need to know about the weather right away. It also backs ideas like vertical farming, which is when crops are grown in small spaces. Keeping an eye on the environment all the time helps keep farming conditions perfect. The answer is cheap, easy to use, and works well in rural areas. It helps people make better choices by giving them real-time data.

In general, the IoT-enabled approach makes it easier to plan for farming and keep an eye on the weather. It shows a modern and effective way to watch the environment.

Keywords: GSM technology, an IOT webpage, a rain sensor, a humidity sensor, an LDR, and a temperature sensor.

I. INTRODUCTION

With the rise of high-speed Internet, people all over the world are becoming increasingly linked. The growing Internet of Things (IoT) goes even further through connecting not just people but also electronic devices that may communicate to each other [1]. The pattern will only get greater as the prices of Wifi-enabled devices go down. The Internet of Things (IoT) is all about linking multiple gadgets through a link and then getting the data from those devices (sensors). This data can be transmitted in any way, submitted to any cloud service, and subsequently evaluated and processed. In the cloud service, you can use this data to let people know in various manners, like by sending them an email, SMS, or buzzing them.

The foremost objective of this paper is to generate and put into place an effective monitoring system that lets you remotely check the necessary parameters over the internet, maintain the data from the sensors in the cloud, and show the predicted trend in a web browser. The values in the cloud are always changing. The soil is tested, mostly to see what moisture content it is, and the crops are grown. So, we can grow different crops in the identical place. Vertical farming is a new way to farm that involves stacking produce vertically.

II. LITERATURE REVIEW

This section examines the current weather forecasting techniques that exist in the literature. Undeviating Reversion is the greatest straightforward and often utilized second-hand a prescient classical for the study. Reversion evaluations are primarily utilized to represent statistics and show the connection among at least one autonomous and one helpless variable. Graphically, direct recession discoveries the greatest fit concluded the opinions. The regression line is the route that fits best through the focuses.

http://doi.org/10.63766/spujstmr.26.000104

The route can be traditional or rounded here, provisional on the statistics. The best-fit track can also be a quadratic or polynomial, which benefits we find improved answers to our questions.

This research employs Decision Tree and Time Series Analysis as two of its algorithms. From the beginning, predicting the weather has been a big problem. Every day, new method come out that replace the old ones. Investigations in literature have established that machine learning methods overtake conservative numerical approaches. Computing's next big thing will be something other than the traditional workstation. The Things on the Internet is a way of thinking about how many of the things surrounding us will be connected to the internet in some way. Machine learning is very comparable to the internet of things.

Rainfall prediction plays a crucial role in agriculture, disaster management, water resource planning, and environmental monitoring. Traditional weather forecasting methods rely on numerical weather prediction models and statistical approaches; however, these methods often require extensive computational resources and large-scale meteorological data. Recent advancements in the Internet of Things (IoT) and Artificial Intelligence (AI) have enabled the development of cost-effective and real-time rainfall prediction systems.

Rahman et al. (2022) proposed a machine learning fusion framework for rainfall prediction in smart cities using environmental sensor data. Their study demonstrated that combining multiple machine learning models improved forecasting accuracy compared to individual models. However, the framework primarily focused on urban environments and required significant computational resources.

Nithyashri et al. (2023) developed an IoT-based rainfall forecasting system for coastal regions using a deep reinforcement learning model. The proposed system utilized environmental sensor data to predict rainfall patterns and achieved improved forecasting performance. Nevertheless, the complexity of the deep learning architecture increased computational requirements and limited deployment on low-cost hardware.

Jagadeesan and Nagarajan (2025) introduced an IoT-driven federated learning framework employing an Attention-based Convolution Recurrent Neural Network (A-CRNN) optimized using the Golden Jackal Optimization algorithm. Their approach enhanced prediction accuracy while preserving data privacy. However, the system required distributed computational infrastructure, making it less suitable for small-scale agricultural applications.

Kukre et al. (2025) proposed an integrated IoT-based weather monitoring and machine learning prediction system that collected environmental data using multiple sensors and applied machine learning algorithms for weather forecasting. The study demonstrated the effectiveness of combining IoT sensing and AI techniques but focused primarily on general weather prediction rather than localized rainfall forecasting.

Recent studies have also highlighted the growing importance of AI-enabled local weather forecasting systems. The IoT-AI Powered Local Weather Forecasting framework (2025) demonstrated that low-cost IoT sensors combined with machine learning algorithms can provide reliable localized weather predictions. The study emphasized the potential of integrating cloud computing, IoT, and AI technologies for precision agriculture and smart environmental monitoring.

Machine learning algorithms such as Decision Trees, Random Forests, Support Vector Machines (SVM), Artificial Neural Networks (ANN), Long Short-Term Memory (LSTM), and ARIMA models have been extensively used for rainfall prediction. Among these methods, Decision Trees provide interpretable prediction rules and low computational complexity, while ARIMA effectively models temporal patterns and seasonal weather variations. These characteristics make them suitable for deployment in resource-constrained IoT environments.

1. Research Gap

Although considerable research has been conducted on IoT-based weather monitoring and AI-based rainfall forecasting, several limitations remain. Most existing studies focus either on environmental data collection using IoT devices or on advanced machine learning models for prediction. Furthermore, many proposed systems rely on expensive infrastructure, large datasets, cloud-intensive processing, or computationally complex deep learning models that are difficult to deploy in rural agricultural environments.

There is a lack of low-cost, real-time, and easily deployable systems that integrate IoT-based environmental sensing with machine learning-based rainfall prediction using lightweight algorithms. Additionally, limited attention has been given to developing solutions suitable for small-scale farmers who require localized weather information for agricultural decision-making.

To address these challenges, the proposed work integrates IoT sensors, cloud-based monitoring, and AI-driven rainfall prediction within a unified framework. The system employs affordable sensors and lightweight machine learning techniques, enabling real-time rainfall forecasting and environmental monitoring suitable for smart agriculture and rural deployments.

TABLE I

DIFFERENT TECHNICQUE

Author	Year	Technique	Limitation
Rahman et al.	2022	ML Fusion Models	High computational requirements
Nithyashri et al.	2023	Deep Reinforcement Learning	Complex architecture
Jagadeesan & Nagarajan	2025	Federated Learning + A-CRNN	Requires distributed infrastructure
Kukre et al.	2025	IoT + Machine Learning	Focused on general weather prediction
Proposed Work	2026	IoT + Decision Tree + ARIMA	Low-cost real-time rainfall prediction

III. METHODOLOGY

1. Suggested System

In this, we talk about the theory behind using the Internet of Things (IoT) to predict rain. This suggested block diagram displays the way several sensors (LDR sensor, humidity sensor, temperature sensor, and rain sensor) link to our controller. A controller is getting the sensor values, processing them, displaying them on an LCD display, and submitting them to a website server. All of the sensors are linked to the controller. A wireless networks module gets the latest readings and transforms them into graphs and statistics.

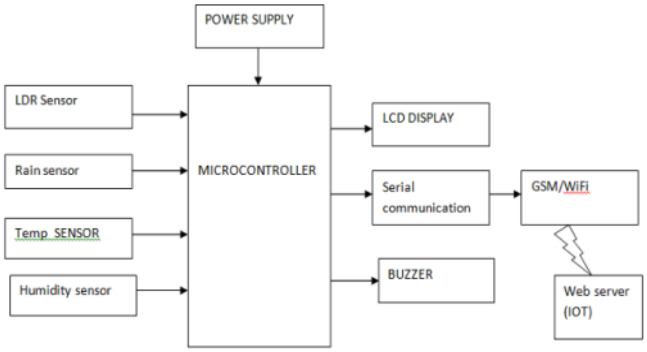

Fig. 1. Block Illustration of the Framework

A. PIC 18f4520 chip: 4k bytes of data memory Data register map: 000-FFF 12-bit address bus

Separated into banks of 256 bytes ,The total number of F banks is a virtual (or access) bank is generated by half of bank 0 and half of bank 15, and it is accessible with regard to whatever bank has been selected; this selection is made using 8-bits.

Within the PIC18, a distinct program data bus and address bus are used to access software memory, which is 16 bits wide. The system's program and static data are conserved in program memory. External on-chip Either PROM or EEPROM are types of on-chip software memory. The Once-only programmable, or OTP, is the name of the PROM version (PIC18C). The Flash memory (PIC18F) is the EEPROM variant. The maximum size of program memory is 2 million bytes. The 21-bit program memory addresses begin at point 0x000000.

B. Rain Sensor

Fig 2. PIC18F452

This device determines when it is raining. It can also be used to determine how much rain is dropping. Its output is both analog and in addition to digital output. This module uses a single analog pin to measure moisture, and it produces an analog signal when the moisture threshold is achieved. Lower output voltage occurs by either more water or less impedance. On the other hand, less water results in greater obstacle or a higher output voltage on the sensor pin.

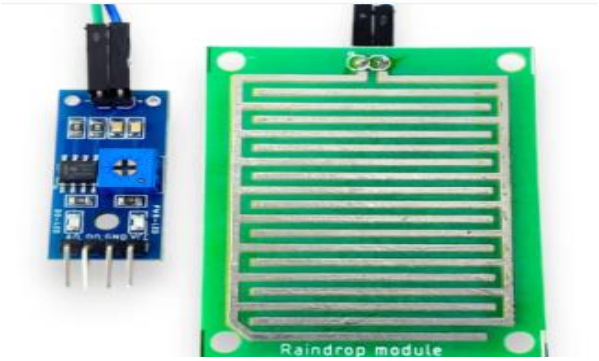

Fig. 3.Rain Sensor

C. Humidity Sensor:

The DHT11 is a humidity and temperature sensor that sends out calibrated digital signals. DHT11 can be connected to with any microcontroller, such as PIC, Arduino, or Raspberry Pi, and get results right away. The DHT11 is a cheap sensor for moisture and temperature that is very reliable and stable over time. It has a sensor for humidity that works with capacitance and measures the air around it and sends a digital signal to the data pin. It's informal to use, and there are libraries and example codes for both Arduino and Raspberry Pi.

Fig. 4. Humanity sensor

D. LM35 Disease Sensor:

The LM35 series is a set of precise disease sensors using integrated circuits. Their production energy is linearly. In relation to the Celsius (Centigrade) illness. The LM35 is better than linear temperature sensors that are calibrated in ° Kelvin because the user doesn't have to deduct an enormous persistent energy from its production to get easy Centigrade scaling.

The LM35 doesn't need some external correction or decoration to give classic exactitudes of ±¼°C at room fever and ±¾°C over a full disease range of -55 to +150°C. Decoration and standard is action at the cracker level make sure that the cost is

truncated. The LM35 is informal to attach to display or regulator motherboard because it has a squat production impedance, a direct output, and precise built-in calibration.

TEMP Sensor:

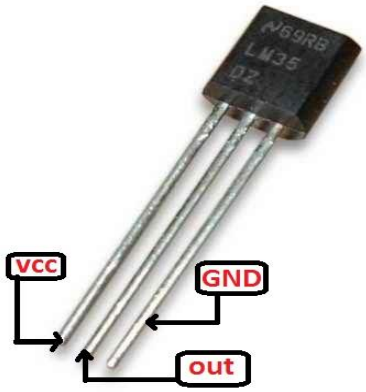

Fig. 5. LM35

E. LCD demonstration:

A 16x2 LCD can show 16 charms per streak, and here are two lines like this. This LCD displays both charms. In a 5x7 pixel grid. This LCD has two registers: Command and Data. The appreciation register keeps track of the requests that the LCD has been given. A facility is an training assumed to the LCD to do a specific job, for example starting it up, dissipating its display, moving the indicator, or regulatory the presentation. The data register keeps any records that will be shown on the LCD.

Fig. 6. LCD Display

F. GSM module:

This cellular mode has a SIM800A chip and an RS232 interface. It can easily connect to a processor using the USB to Sequential connector or to a microchip using the TTL converter. After you attach the modem with the USB to RS232 connector, you need to look for the right COM port in the Device Manager of the USB to Serial Adapter. You can then use other computer software to connect to that COM anchorage at 9600 baud degree, which is the modem's avoidance byte per second amount.

Fig 7.GSM Module

G. AI-Based Rainfall Prediction Module

The proposed system integrates Internet of Things (IoT) technology with Artificial Intelligence (AI) to provide accurate rainfall prediction for agricultural and environmental monitoring applications. Environmental parameters collected through multiple sensors are continuously transmitted to the cloud platform, where machine learning algorithms analyze historical and real-time data to forecast rainfall conditions.

1. Data Collection

The IoT node consists of multiple sensors including:

- DHT11 Humidity Sensor – measures relative humidity (%).
- LM35 Temperature Sensor – measures ambient temperature (°C).
- Rain Sensor – detects rainfall intensity and precipitation occurrence.
- LDR Sensor – measures light intensity and atmospheric brightness.

The sensors are connected to the PIC18F4520 microcontroller, which collects data at regular intervals. The collected data are transmitted through the GSM communication module and stored in a cloud database for further processing.

2. Dataset Preparation

The dataset consists of environmental observations collected from the deployed sensor network. Each record contains:

TABLE II

FEATURE

Feature	Description
Temperature	Ambient temperature in °C
Humidity	Relative humidity (%)
Rainfall Sensor Value	Rain intensity measurement
Light Intensity	LDR sensor reading
Timestamp	Date and time of observation
Rainfall Status	Target variable (Rain/No Rain)

Before model training, the collected data undergo preprocessing to improve prediction accuracy.

3. Data Preprocessing

The preprocessing stage includes:

- Removal of incomplete or corrupted records.
- Noise filtering to eliminate abnormal sensor readings.
- Handling missing values through interpolation techniques.
- Feature normalization to bring sensor values into a comparable range.
- Conversion of categorical rainfall labels into machine-readable numerical values.

These preprocessing steps improve data quality and reduce prediction errors.

4. Feature Engineering

Feature engineering is performed to identify the most influential environmental parameters affecting rainfall prediction. Statistical analysis shows that humidity, temperature variation, and rainfall sensor readings have significant influence on rainfall occurrence.

Additional derived features include:

- Average temperature over a specific time window.
- Humidity trend variation.
- Rainfall intensity changes.
- Daytime and nighttime environmental patterns.

These engineered features help the machine learning model capture complex weather relationships.

5. Machine Learning Model

The proposed system utilizes a hybrid AI approach combining Decision Tree Classification and ARIMA Time-Series Analysis.

Decision Tree Model

Decision Tree is a supervised machine learning algorithm used to classify whether rainfall is likely to occur. The model creates decision rules based on environmental conditions such as humidity, temperature, and rainfall sensor values.

Advantages include:

- Easy interpretation of prediction results.
- Fast training and prediction.
- Ability to handle nonlinear relationships.
- Suitable for resource-constrained IoT environments.

The Decision Tree model learns from historical weather observations and generates rainfall predictions based on learned decision paths.

ARIMA Model

Autoregressive Integrated Moving Average (ARIMA) is employed for time-series forecasting of weather parameters.

The ARIMA model captures:

- Historical weather trends.
- Seasonal variations.
- Temporal dependencies among observations.

ARIMA provides short-term rainfall forecasting by analyzing sequential sensor data collected over time.

6. Model Training Process

The prepared dataset is divided into:

- 80% Training Data
- 20% Testing Data

The training phase includes:

1. Input sensor data acquisition.
2. Feature extraction and preprocessing.
3. Model training using historical observations.
4. Hyper parameter optimization.
5. Model validation using testing data.

Cross-validation techniques are applied to reduce over fitting and improve model generalization.7. Performance Evaluation

The trained model is evaluated using standard machine learning metrics:

- Accuracy
- Precision
- Recall
- F1-Score
- Mean Absolute Error (MAE)
- Root Mean Square Error (RMSE)

These metrics assess both classification performance and forecasting accuracy.

8. Real-Time Prediction Workflow

The real-time prediction process follows the steps below:

1. Sensors collect environmental data.
2. Data are transmitted to the cloud server through the GSM module.
3. Preprocessing and feature extraction are performed.
4. The trained AI model analyzes current weather conditions.
5. Rainfall probability is calculated.
6. Prediction results are displayed on the IoT web dashboard.
7. Alerts can be generated when rainfall probability exceeds a predefined threshold.

9. Advantages of the Proposed AI-IoT Framework

The proposed AI-integrated framework offers several advantages:

- Real-time weather monitoring.
- Early rainfall prediction for agricultural planning.
- Low-cost deployment using affordable sensors.
- Reduced manual observation effort.
- Improved decision-making for farmers.

IV. CONCLUSION

Keeping devices that are embedded in the environment for monitoring makes the environment safer through facilitating self-protection (i.e., a smart environment). To meet this requirement, the sensor devices must be set up in the environment to collect and analyze the data. The main goal was to build the most accurate system possible using cheap parts. This system could monitor the weather in real time on farms and use machine learning algorithms to make better predictions about the weather in the future. ARIMA can better capture how the temperature changes over time than the Decision Tree. This makes the Rain information in the weather profile more compact and natural.

It has always been hard to predict the weather, and new methods are being developed every day to replace the old ones. The next big thing in computers will be something other than the traditional desktop. In the Internet of Things (IoT) model, a lot of the things around us will be connected to the internet in some way. Machine Learning and the Internet of Things are very similar. A perfect mix of them can speed up the modernization of agriculture, make smart agriculture a reality, and solve problems with farmers, agriculture, and the countryside. The user will formerly be able to contact the data and investigation results

VI .REFERNCES

[1] ahman A.-u., Abbas S., Gollapalli M., Ahmed R., Aftab S., Ahmad M., Khan M. A. & Mosavi A. 2022,
Rainfall Prediction System Using Machine Learning Fusion for Smart Cities. In Sensors, 22(9):3504. MDPI. A framework integrating ML models for real-time rainfall forecasting using sensor data.

[2] J. Nithyashri, R. K. Poluru, S. Balakrishnan, M. A. Kumar, P. Prabu & amp; S. Nandhini 2023, IoT based prediction of rainfall forecast in coastal regions using deep reinforcement model. In Measurement: Sensors,100877.

[3] I. Jagadeesan & amp; R. Nagarajan 2025, An IoT-Driven Federated Learning Method for Rainfall Prediction Employing A-CRNN GJO. In Engineering, Technology & Applied Science Research (ETASR).

[4] V. N. Kukre, S. Pawar& S. Tapse 2025, Integrated IoT Based Weather Monitoring and Machine Learning Weather Prediction System. In IJRASET Journal for Research in Applied Science and Engineering Technology

[5] Integration of IoT-AI Powered Local Weather Forecasting" arXiv preprint 2025. Framework for integrating low-cost IoT and AI for local weather and rainfall forecasting .

[6] D. D. Pandya, S. Degadwala, D. Vyas, S. V. Sureshbhai, L. Ainapurapu and N. S. Bhavsar, "Advancing Erythemato-Squamous Disease Classification with Multi-class Machine Learning"; 2023 7th International 8 Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud) (I-SMAC), Kirtipur, Nepal, 2023, pp. 542-547, doi: 10.1109/I-SMAC58438.2023.10290599.

[7] D. D. Pandya, P. A. Patel, H. H. Patel, A. J. Goswami, S. Degadwala and D. Vyas, Unveiling the Power of Collective Intelligence: A Voting-based Approach for Dementia Classification,& 7th International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud) (I-SMAC), Kirtipur, Nepal, 2023, pp. 478-482, doi: 10.1109/I-SMAC58438.2023.10290165.

[8] Pandya, D.D., Jadeja, A., Trivedi, S., Patel, P.A., Tamhankar, I., Dhanesha, P.S. (2026).Internet of Things (IoT)'s Transformative Power Enhancing Wireless Networks and Sensing. In: Rathore, V.S., Piuri, V., Babo, R., Karthik, S. (eds) Universal Threats in Expert Applications and Solutions. UNI-TEAS 2025. Lecture Notes in Networks and Systems, vol 1452. Springer, Singapore. https://doi.org/10.1007/978-981-96-7292-9_20

[9] Pandya, D.D., Jadeja, A., Trivedi, S., Patel, P.A., Tamhankar, I., Dhanesha, P.S. (2026). Internet of Things (IoT)'s Transformative Power Enhancing Wireless Networks and Sensing. In: Rathore, V.S., Piuri, V., Babo, R., Karthik, S. (eds) Universal Threats in Expert Applications and Solutions. UNI-TEAS 2025.

[10] Pandya, Darshanaben Dipakkumar, Borate Pooja Satish, and Khushbu. "Improving Data Mining Reliability Through Applied C-B Techniques for Addressing Misplaced Values." Universal Threats in Expert Applications and Solutions (UNI-TEAS 2025).

SANKALCHAND PATEL UNIVERSITY

।। अथातो ज्ञानजिज्ञासा ।।

Sankalchand Patel Vidyadham,
Ambaji - Gandhinagar Highway
Visnagar 384315, India

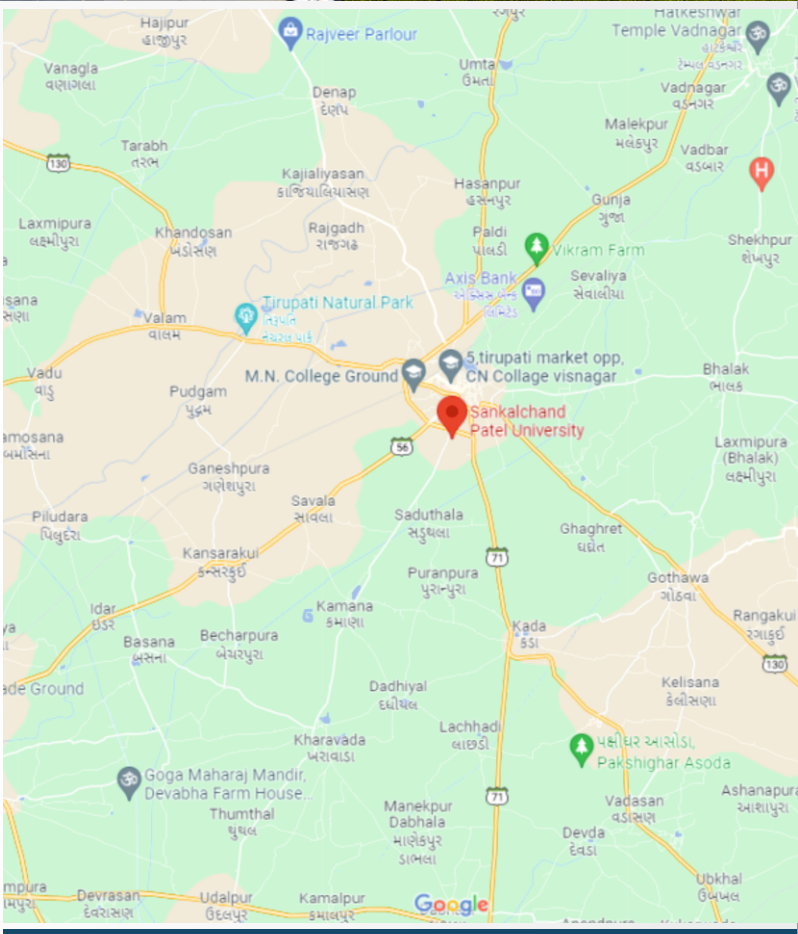

Managing Editor, SPU Journals
Visnagar 384315, India
Mobile : +91 - 97128 23247
Email : communications.spujstmr@spu.ac.in
Website : https://spujstmr.in/

www.spu.ac.in.