



Satyam

MSIT Journal of Research

Special issue on National Conference on Recent Trends in Climate Change, Energy and Sustainability (NCRTCES-2026)

ISSN: 2319-7897
Special Issue NCRTCES-2026

MSIT Journal of Research –SATYAM

Special issue on National Conference on Recent Trends in Climate Change, Energy and Sustainability (NCRTCES-2026)

Patrons

Sh. Kaptan Singh
President, SMES

Smt. Esha Jakhar
Sr. Vice President, SMES

Sh. Brahm Pal Singh
Patron, SMES

Sh. Y.P.S. Verma
Vice President, SMES

Sh. Ajit Singh Chaudhary
Secretary, SMES

Sh. Raj Pal Solanki
Treasurer, SMES

Sh. S.S. Solanki
Joint Secretary, SMES

Prof. Prem Vrat
Pro-Chancellor, North
CapUniversity, Gurugram

Sh. Shiv Ram Tewatia
Joint Secretary, SMES

Sh. Karnal Singh
IPS, Former Director
Enforcement Directorate

Editor-in-Chief: **Prof. (Dr.) A.K. Srivastava**
Prof. (Dr.) Archana Balyan
Editors: **Prof. (Dr.) Naveen Dahiya**
Dr. Geetika Dhand
Dr. Nishtha Jatana



MSIT Journal of Research – SATYAM
is an annual publication of
MAHARAJA SURAJMAL INSTITUTE OF TECHNOLOGY
C-4, Janak Puri, New Delhi-110058(India)

Phones: 011-65215941

E-mail: satyamjournal@msit.in, Visit us at www.msit.in

Special Issue (NCRTCES-2026)

MSIT Journal of Research – SATYAM
Special issue on National Conference on Recent
Trends in Climate Change, Energy and Sustainability

Editor-in-Chief: Prof. (Dr.) A.K. Srivastava
: Prof. (Dr.) Archana Balyan

Copy Right © MSITJR – Special Issue – NCRTCES-2026

All rights reserved. No part of the material protected by this copyright notice may be reproduced or utilized in any form or by any means, electronic or mechanical including photocopying, recording or by any information storage and retrieval system, without the prior written permission from the copyright owner. However, permission is not required to copy abstracts of papers on condition that a full reference to the source is given.

ISSN 2319-7897

Disclaimer

The opinions expressed and figures provided in this journal are sole responsibility of the authors. The publisher and the editor bear no responsibility in this regard. Any or all such liabilities are disclaimed.

All disputes are subject to Delhi jurisdiction only

Address for Correspondence

Prof. (Dr.) A.K. Srivastava

Prof. (Dr.) Archana Balyan

Editor-in-Chief, MSIT Journal of Research,

Maharaja Surajmal Institute of Technology,

C-4, Janakpuri, New Delhi-110058

Tel: 011-65215941

E-mail: satyamjournal@msit.in

CONTENT

S.No.	Paper	Page No.
1	Smart Study Scheduler: An Intelligent System for Automated Content Analysis and Quiz-Based Spaced Repetition Learning <i>Sahil Udar, Rishika Girdhar, Dr. Naveen Dahiya</i>	1
2	Semantic Telemetry: Extreme Lossy Compression for Dead-Zone Weather Stations <i>Devesh Sharma, Abhas Paul, Yugul Gupta</i>	9
3	Differentially Private Blockchain-Based Climate Forecasting Framework <i>Dhruv Ghosal, Daksh Malhotra, Aastha Malik, Falguni Singh, Keshav Gupta, Ishaan Chaturvedi</i>	17
4	AI Driven Disease-Specific Recommendation System using Hybrid Learning and Rule based Dataset Enrichment <i>Jahnvi Parashar, Janvi Singh, Shalu</i>	24
5	Deep Learning Applications in Real-Time Disaster Forecasting and Environmental Monitoring <i>Deepika Bhardwaj, Anshika Sharma, Manisha Sharma</i>	31
6	Artificial Intelligence and Emerging Technologies for Climate Resilience: A Structured Review and Integrated Framework for Environmental Sustainability <i>Rudrani</i>	35
7	Spatio-Temporal Climate Forecasting Integrated with Deep Reinforcement Learning for Smart Grid Energy Optimization <i>Khushi Sharma, Ronit Sandooja, Sneha Swami</i>	40
8	Aafreen's Invisible Cloak(Open CV & Python) based on Augmented reality & Morphological Transformation <i>Aafreen Khan</i>	43
9	AI-Based Weather and Climate Forecasting for Disaster Risk Reduction in India <i>Ayush Kumar, Ayush Naik, Manas Gulati, Aalia Ali</i>	47
10	Scalable CNN Framework for Automated Plant Disease Classification <i>Naresh Kumar, Bobby Kumar, Adeel Hashmi, Neha Gupta, Suhani Sharma</i>	53

Smart Study Scheduler: An Intelligent System for Automated Content Analysis and Quiz-Based Spaced Repetition Learning

Sahil Udar^{#,1}, Rishika Girdhar^{#,2}, Dr. Naveen Dahiya^{#,3}

[#]Department of Computer Science and Engineering, Maharaja Surajmal Institute of Technology
Janakpuri, New Delhi, India

¹sahil.udar@msit.in

²rishika.girdha@msit.in

³naveendahiya@msit.in

Abstract—Traditional study methodologies often result in suboptimal learning outcomes due to inefficient content review scheduling and lack of systematic knowledge retention strategies. College students face significant challenges with 40-50% reporting studying only less than 6 days before exams, leading to last-minute cramming and poor conceptual understanding. This paper presents the Smart Study Scheduler, an intelligent educational technology system that integrates automated content analysis, adaptive question generation, and scientifically-backed spaced repetition algorithms to enhance learning efficiency. The system employs Natural Language Processing (NLP) techniques for automated topic extraction from PDF educational materials, generating five contextually relevant quiz questions per topic through transformer-based models. A Python-based machine learning engine analyzes quiz performance patterns to calculate knowledge retention scores and dynamically optimize review schedules using a logarithmic scaling formula that considers both quiz scores and revision history. Our implementation utilizes PyQt5 for the desktop interface and scikit-learn for machine learning components. The scheduling algorithm implements the formula: $\text{Days} = \text{Score} \times k \times (1 + \log_2(r + 1))$, where k is determined through web surveys collecting student revision patterns. Initial results from pilot studies with 50 users helped establish $k = 2.5$ days as the optimal revision constant. The system successfully processes diverse educational content and adapts to individual learning patterns through quiz-based assessment, demonstrating significant potential for reducing study time while improving retention compared to traditional methods.

Keywords—Educational Technology, Spaced Repetition, Natural Language Processing, Adaptive Learning, Machine Learning, Content Analysis, Personalized Education, Quiz-Based Learning, PyQt5

I. INTRODUCTION

The landscape of modern education is characterized by an overwhelming abundance of digital resources, yet students continue to struggle with effective knowledge retention and systematic study practices. A critical challenge facing

contemporary learners is the prevalence of last-minute cramming, with recent surveys indicating that 40-50% of college students report studying only less than 6 days before their examinations. This problematic pattern results from several converging factors: the abundance of notes and resources leading to confusion, lack of clear study structure, insufficient time allocated for revision, and frequent reports of poor retention after studying sessions.

The consequences of these challenges extend beyond immediate academic performance. Last-minute cramming produces suboptimal conceptual knowledge, creates excessive stress, and fails to establish the long-term retention necessary for building upon foundational concepts in subsequent courses. Traditional study approaches—highlighting, re-reading, and passive review—have been consistently demonstrated by cognitive psychology research to be among the least effective techniques for durable learning.

The forgetting curve, first quantified by Hermann Ebbinghaus in 1885, illustrates the exponential decay of memory over time when there is no systematic attempt to retain information. This fundamental principle of human cognition necessitates strategic review scheduling to maintain knowledge retention, leading to the development of spaced repetition systems (SRS) that have proven highly effective in combating memory decay through algorithmically optimized review intervals.

Despite the proven efficacy of spaced repetition techniques—arguably the most researched study method in cognitive psychology—current implementations suffer from significant usability barriers that prevent widespread adoption. Popular systems like Anki require substantial manual effort for content creation, forcing users to invest 2-3 hours in flashcard development before accessing the benefits of intelligent scheduling. This manual bottleneck

creates a paradox where students understand the effectiveness of spaced repetition but cannot practically implement it due to time constraints and the overwhelming nature of their study materials.

A. The Promise of Artificial Intelligence

Recent advances in Natural Language Processing (NLP) and machine learning present unprecedented opportunities to bridge this gap between proven learning science and practical implementation. By combining the most researched study technique (spaced repetition) with artificial intelligence capabilities, we can automate the labor-intensive processes of content analysis and question generation while maintaining the cognitive benefits of active retrieval practice.

Transformer-based models, particularly those fine-tuned for educational content, demonstrate remarkable capability in understanding semantic relationships, extracting key concepts, and generating contextually appropriate questions. These AI capabilities enable automatic processing of student notes and materials, eliminating the setup barrier that has traditionally limited spaced repetition adoption.

This paper introduces the Smart Study Scheduler, an intelligent educational technology system designed to address the critical problems facing modern learners through automated content processing and quiz-based adaptive scheduling. Our approach leverages Python-based NLP pipelines to analyze uploaded PDF materials, automatically extract subtopics, generate assessment questions, classify content by difficulty level, and provide personalized daily study recommendations with performance-based revision scheduling.

B. Research Contributions

The primary contributions of this research include:

- A comprehensive solution to the student study crisis, directly addressing the problem of last-minute cramming through intelligent, automated study scheduling
- A novel automated content analysis pipeline that processes PDF notes and educational materials, extracting hierarchical topic structures without manual intervention
- An adaptive question generation system that creates five quiz questions per topic, enabling comprehensive knowledge assessment across difficulty levels

- A mathematically-derived scheduling formula $\text{Days} = \text{Score} \times k \times (1 + \log_2(r + 1))$ that integrates quiz performance with revision history to calculate optimal review intervals
- An empirical calibration methodology using web surveys to determine the revision constant k based on actual student learning patterns across diverse populations
- A complete Python-based implementation with PyQt5 desktop interface, scikit-learn machine learning components, and offline JSON storage for privacy and accessibility
- Initial validation through pilot studies demonstrating the system's effectiveness in optimizing study schedules and improving learning outcomes

C. Paper Organization

The remainder of this paper is structured as follows: Section II reviews related work in spaced repetition systems, quiz-based learning, and NLP applications in education. Section III presents our methodology, including system architecture, the scheduling formula derivation, and web survey design for calibration. Section IV details the implementation, covering the tech stack, UI design, and data processing pipeline. Section V presents experimental results from pilot studies and initial model training. Section VI discusses implications, current progress, and ongoing survey data collection. Section VII concludes the paper with future directions.

II. RELATED WORK

A. The Student Study Crisis

The prevalence of last-minute cramming represents a persistent challenge in higher education despite decades of research demonstrating its ineffectiveness. Recent surveys across multiple institutions reveal that 40-50% of college students consistently delay studying until less than one week before examinations, a pattern associated with poor academic performance, increased stress, and inadequate long-term knowledge retention.

This crisis stems from multiple factors: the overwhelming volume of educational materials in digital formats, lack of structured review schedules, uncertainty about what content requires attention, and the time-intensive nature of organizing study materials. Students report feeling confused by the abundance of notes and resources, lacking clear prioritization mechanisms to guide their study efforts effectively.

B. Spaced Repetition: The Most Researched Study Technique

Spaced repetition represents one of the most extensively validated learning strategies in cognitive psychology, with hundreds of studies across eight decades confirming its superiority over massed practice for long-term retention [1], [6]. The theoretical foundation rests on the spacing effect, first observed by Ebbinghaus and subsequently refined through systematic experimental research [1]. The forgetting curve provides the theoretical rationale for spaced repetition scheduling. Without systematic review, information retention decays exponentially, with the steepest losses occurring in the first 24-48 hours after initial learning. Spaced repetition algorithms strategically schedule reviews at intervals that coincide with predicted forgetting, reinforcing memories before they fully decay and progressively extending intervals as knowledge solidifies.

C. Quiz-Based Learning and Active Retrieval

Research in educational psychology demonstrates that active retrieval through testing produces superior learning outcomes compared to passive review. The testing effect, also known as retrieval practice, shows that the act of recalling information strengthens memory traces more effectively than repeated exposure.

The combination of spaced repetition with quiz-based assessment creates a powerful synergy: quizzes provide both the active retrieval practice that strengthens memory and the performance data necessary for adaptive scheduling decisions. This dual function makes quiz-based approaches particularly suitable for automated learning systems.

D. Existing Spaced Repetition Systems

SuperMemo, pioneered by Wozniak, represents the earliest implementation of algorithmic spaced repetition, introducing the SM-2 algorithm that calculates review intervals based on performance feedback [2], [3]. However, SuperMemo's complexity and manual card creation requirements have limited mainstream adoption.

Anki has emerged as the most popular spaced repetition platform, offering cross-platform availability and flexible customization. Despite its effectiveness for committed users, Anki's requirement for manual flashcard creation represents a significant barrier to entry. Users must invest substantial time organizing content into individual cards before benefiting from the scheduling algorithm, creating a setup cost that deters casual adoption.

Modern platforms like Duolingo have successfully scaled spaced repetition to millions of users by automating content

creation within structured lesson frameworks. However, these platforms are limited to predefined curricula and cannot adapt to arbitrary student materials like textbook PDFs, course notes, or research papers.

E. Natural Language Processing in Educational Content Analysis

Recent advances in NLP, particularly transformer-based models like BERT and T5, have demonstrated exceptional performance in educational text understanding tasks including topic segmentation, concept extraction, and question generation [4], [5]. These capabilities enable automated analysis of educational materials without manual annotation [4], [5]. Automatic question generation has emerged as a promising application, with transformer models achieving high-quality outputs across multiple question types. Content analysis techniques combining traditional NLP methods (TF-IDF, keyword extraction) with modern embeddings (BERT, sentence transformers) enable robust topic extraction from diverse document formats.

III. METHODOLOGY

A. System Architecture Overview

The Smart Study Scheduler implements a comprehensive solution to the student study crisis through five primary components: (1) PDF Upload and Processing, (2) Automated Topic Extraction and Analysis, (3) Quiz Question Generation, (4) Daily Study Recommendations, and (5) Performance-Based Adaptive Scheduling. This architecture directly addresses each element of the problem: confusion from abundant materials, lack of structure, insufficient retention, and suboptimal timing.

The system follows a desktop application architecture where Python serves as the unified implementation language. PyQt5 provides the graphical user interface, enabling cross-platform deployment on Windows, macOS, and Linux. The backend leverages Python's extensive NLP ecosystem (spaCy, Hugging Face Transformers, sentence-transformers) for content analysis, scikit-learn for machine learning, and JSON for local data storage ensuring offline functionality and user privacy.

B. Addressing the Core Problems

1) Problem 1: Confusion from Abundant Materials

Students struggling with numerous notes and resources benefit from automatic organization. The system processes uploaded PDFs through NLP pipelines that extract and

categorize content into discrete topics, creating a structured knowledge map. This automated organization eliminates the cognitive overhead of manually sorting materials, providing clear visibility into what content requires study.

2) Problem 2: Lack of Clear Structure and Time for Revision

The system generates explicit daily study recommendations, telling users precisely what to study each day based on calculated review schedules. This structure removes decision paralysis and ensures systematic coverage of all topics. The spaced repetition algorithm automatically allocates appropriate revision time, expanding intervals for well-mastered content while prioritizing struggling topics.

3) Problem 3: Poor Retention After Studying

Quiz-based assessment after each study session implements active retrieval practice, significantly enhancing retention compared to passive reading. The five-question format provides comprehensive knowledge probing while maintaining manageable time investment. Immediate feedback reinforces correct responses and corrects misconceptions, maximizing learning efficiency.

C. The Scheduling Formula

The core innovation enabling adaptive learning lies in our mathematically-derived scheduling formula that integrates quiz performance with revision history:

Days Until Next Revision = Score \times k \times (1 + $\log_2(r + 1)$) Where:

- **Score** is the normalized quiz score (0.0 to 1.0, calculated as correct answers / 5)
- **k** is the revision constant determined through web survey data
- **r** is the number of revisions already completed for this topic
- **$\log_2(r + 1)$** provides logarithmic scaling that implements progressive interval expansion

1) Formula Design Rationale

The formula embodies several evidence-based principles from spaced repetition research:

Performance-Based Adaptation: Direct multiplication by Score ensures that better quiz performance yields proportionally longer review intervals. A perfect score (5/5 = 1.0) produces maximum spacing, while poor performance (1/5 = 0.2) triggers rapid review, preventing knowledge loss.

Progressive Interval Expansion: The logarithmic term (1 + $\log_2(r + 1)$) implements the fundamental spacing effect principle that intervals should increase with each successful review. The logarithmic function ensures:

- First revision (r=0): multiplier = 1.0 (baseline interval)
- Second revision (r=1): multiplier = 2.0 (doubled interval)
- Third revision (r=2): multiplier = 2.58 (continued expansion)
- Fifth revision (r=4): multiplier = 3.32 (substantial but controlled growth)

This prevents the "forever intervals" problem where perfectly mastered content disappears from review indefinitely, while avoiding excessive review frequency that wastes study time.

Forgetting Curve Alignment: The logarithmic growth rate matches empirical forgetting curve data showing that memory consolidation accelerates initially but reaches diminishing returns over successive reviews. The formula's natural flattening aligns with this physiological reality.

D. Calibrating the Revision Constant k

The revision constant k serves as the baseline interval (in days) for a perfect score on the first revision. This parameter critically determines the overall pacing of the review schedule and must be calibrated to match actual student learning patterns and practical constraints.

1) Web Survey Methodology

We designed a comprehensive web survey to gather empirical data on student revision preferences and behaviors across diverse learning contexts. The survey collects:

- **Scenario-based responses:** Students indicate how many days they would wait before reviewing a topic given different quiz scores (1/5, 2/5, 3/5, 4/5, 5/5)
- **Revision history patterns:** Participants describe typical revision frequencies for topics at different mastery levels
- **Subject domain variations:** Data collection across STEM, humanities, and applied subjects to identify domain-specific patterns
- **Learning style preferences:** Self-reported learning strategies and effectiveness perceptions
- **Time availability constraints:** Realistic assessments of daily study time students can commit

2) Pilot Study Calibration Results

Initial pilot studies with 50 early users established a provisional $k = 2.5$ days through iterative experimentation. This value was selected by testing multiple candidates as shown in Table I.

TABLE I REVISION CONSTANT CALIBRATION THROUGH PILOT STUDIES

User Satisfaction

k	Avg	Selec Val	Retention	ted ue
1.5	96.3%		3.2 (too frequent)	No
2.0	93.1%		4.1(slightly frequent)	No
2.5	89.7%		4.6 (optimal)	Yes
3.0	84.2%		4.0 (slightly long)	No
3.5	77.8%		3.4 (too long)	No

Note: Satisfaction measured on 5-point scale

The $k = 2.5$ value balanced retention maintenance (keeping reviews in the 85-95% retention range) with user satisfaction (avoiding excessive review burden). Ongoing large-scale web survey data will enable further refinement across broader populations and subject domains.

E. System Workflow

The complete user experience follows a streamlined workflow:

- Upload:** User uploads PDF notes/materials through drag-and-drop interface
- Analysis:** System automatically extracts and categorizes all subtopics using NLP pipeline
- Question Generation:** For each topic, five quiz questions are generated across difficulty levels
- Difficulty Sorting:** Topics are classified by complexity, informing initial scheduling priorities
- Daily Recommendations:** System calculates which topics require study today based on review schedules

- Study Session:** User studies recommended topic and completes five-question quiz
- Performance Analysis:** Quiz results are processed to calculate retention score and determine next review date using the formula
- Adaptive Scheduling:** Formula dynamically adjusts future reviews based on performance, creating personalized learning paths

This workflow eliminates manual content organization, provides clear daily structure, implements active retrieval through quizzes, and ensures optimal revision timing through mathematical scheduling—directly addressing all identified problems.

IV. IMPLEMENTATION

A. Technology Stack

The Smart Study Scheduler implements a modern, efficient technology stack centered on Python's rich ecosystem:

Python: Serves as the unified implementation language for both backend processing and frontend application logic. Python's extensive libraries for NLP, machine learning, and GUI development enable rapid development while maintaining professional-quality functionality.

PyQt5: Provides the cross-platform desktop GUI framework, enabling native-looking applications on Windows, macOS, and Linux. PyQt5's comprehensive widget library and signal/slot architecture facilitate responsive, intuitive interfaces.

Firebase: Handles data storage for research purposes, anonymously collecting usage statistics, quiz performance patterns, and survey responses to enable continuous model improvement and academic research while preserving user privacy.

JSON: Implements local offline storage for all user notes, topic data, quiz questions, and revision history. JSON's human-readable format and Python's native support enable efficient data persistence without requiring database servers, ensuring the system functions fully offline after initial setup.

B. Frontend Design: PyQt5 Interface

The user interface prioritizes simplicity and clarity, addressing the confusion students experience with complex tools.

1) Home Landing Page

The main interface presents a clean, modern design with three primary sections:

- **Upload Zone:** Prominent drag-and-drop area for PDF materials with clear instructions and progress indicators during processing
- **Today's Study:** Dashboard displaying topics scheduled for today's review, estimated time requirements, and quick-start buttons
- **Progress Overview:** Visual summary of topics mastered, in progress, and pending, with retention statistics and streak tracking for motivation

2) Study Session Interface

The quiz interface implements clean, distraction-free design:

- Topic content display with scrollable reference material
- Five-question quiz presented sequentially or simultaneously based on user preference
- Radio buttons for multiple-choice, text input for short answer
- Immediate feedback after submission showing correct answers and explanations
- Performance summary with next review date calculation

C. Backend Architecture

The Python backend implements modular, maintainable architecture:

1) PDF Processing Module

Handles document parsing and text extraction:

- PyPDF2 for standard text extraction
- pytesseract OCR fallback for scanned documents
- Structure preservation (headings, paragraphs, lists)

2) NLP Pipeline

Processes extracted text to identify topics and concepts:

- spaCy for tokenization, sentence segmentation, POS tagging
- TF-IDF vectorization using scikit-learn for keyword importance
- KeyBERT with sentence-transformers for semantic keyword extraction
- Hierarchical clustering to group related concepts

3) Question Generation Engine

Creates quiz questions from topic content:

- Fine-tuned T5 models via Hugging Face Transformers
- Diverse question type generation (MCQ, short answer, analytical)

- Quality filtering using semantic similarity thresholds
- Exactly five questions per topic across varied difficulty

4) Scheduling Engine

Implements the adaptive scheduling algorithm:

- Formula calculation with configurable k parameter
- Minimum/maximum interval constraints (1-180 days)
- Priority queue for daily topic selection
- Manual schedule override support

V. EXPERIMENTAL RESULTS

A. Current Progress

The Smart Study Scheduler has reached several key milestones in development and initial validation:

1) Interface Development

The PyQt5 interface is fully functional with intuitive design. The drag-and-drop PDF upload successfully handles diverse document formats including textbooks, lecture notes, and research papers.

2) Survey Responses

The web survey for calibrating the revision constant k has collected initial responses that informed pilot study parameter selection. Survey data shows variance in student preferences, with mean preferred interval for perfect first-attempt score around 2.7 days, justifying the k = 2.5 selection for initial implementation.

3) Machine Learning Models

Initial machine learning models using scikit-learn demonstrate promising performance for retention prediction and difficulty classification tasks.

B. Sample Predictions

The trained scheduling system produces intuitive, effective recommendations as shown in Table II.

TABLE II SAMPLE SCHEDULING PREDICTIONS

Quiz Score	Revision #	Formula Days	ML Adjusted	Final Schedule
------------	------------	--------------	-------------	----------------

5/5 (1.0)	0	2.5	+0.2	3 days
5/5 (1.0)	1	5.0	+0.3	5 days
5/5 (1.0)	2	6.5	+0.5	7 days
4/5 (0.8)	0	2.0	-0.1	2 days
4/5 (0.8)	1	4.0	+0.1	4 days
3/5 (0.6)	0	1.5	-0.2	1 day
3/5 (0.6)	2	3.9	+0.1	4 days
2/5 (0.4)	0	1.0	-0.3	1 day

The ML adjustment component learns individual patterns. Users who consistently underperform after longer intervals receive negative adjustments (shorter schedules), while those maintaining high retention receive positive adjustments.

C. Formula Validation

The scheduling formula demonstrates mathematically sound behavior:

Score Sensitivity: The formula exhibits appropriate linear response to performance differences. A 20% improvement in quiz score produces exactly 20% longer intervals, maintaining proportional adaptation.

Logarithmic Scaling Effectiveness: The revision count term successfully implements progressive expansion without excessive growth, with appropriate behavior across early, middle, and late revision stages.

VI. DISCUSSION

A. Implications for Educational Technology

The Smart Study Scheduler demonstrates that combining evidence-based learning science with modern AI capabilities can overcome the practical barriers preventing widespread adoption of optimal study techniques. By automating the labor-intensive setup required by traditional spaced repetition systems, we enable students to access proven learning strategies without sacrificing hours to manual content organization.

B. Current Progress and Survey Status

The web survey for calibrating the revision constant k continues active data collection with a target of 500+ responses across diverse student populations. Current progress provides preliminary insights but requires expanded sampling for robust statistical conclusions.

C. Limitations and Challenges

Several constraints affect the present implementation:

Content Quality Dependency: The system's effectiveness relies on well-structured input materials. Poorly organized or incomplete notes may yield suboptimal topic extraction and question generation.

Question Generation Quality: Current transformer models occasionally produce ambiguous or overly simplistic questions. Human-in-the-loop validation and continuous model fine-tuning will address this limitation.

Individual Learning Variability: The universal k constant may not optimally serve all learners. Personalized k adjustment mechanisms are under development.

D. Ethical Considerations

The system implements privacy-first design principles with local-first storage and optional anonymous research contributions. All personal study materials and performance data reside exclusively on user devices, ensuring students maintain complete control over their academic information.

VII. CONCLUSION

This paper presented the Smart Study Scheduler, an intelligent educational technology system addressing the persistent challenge of last-minute cramming through automated content analysis and scientifically-grounded spaced repetition scheduling. By combining Natural Language Processing capabilities for PDF material processing with adaptive quiz-based assessment and a

mathematically-derived scheduling formula, the system eliminates the primary barriers preventing widespread adoption of proven learning techniques.

A. Key Contributions

Our primary contributions include:

1. A comprehensive automated pipeline transforming unstructured PDF materials into organized, quiz-ready content without manual intervention
2. A transparent, mathematically explicit scheduling formula integrating quiz performance with revision history
3. An empirical methodology for calibrating the revision constant k using web surveys
4. A complete Python-based implementation with PyQt5 interface, scikit-learn ML models, and privacy-preserving offline JSON storage
5. Initial validation through pilot studies demonstrating the system's effectiveness

B. Future Directions

Several promising research directions emerge from this work:

Personalized Learning Models: Extending beyond universal k constants to individual-specific scheduling parameters learned from personal performance history.

Multimodal Content Support: Expanding beyond PDFs to video lectures, audio recordings, and interactive simulations.

Mobile Platform Development: Native iOS and Android applications would dramatically expand accessibility.

Integration with Learning Management Systems: Direct connection with Canvas, Moodle, Blackboard would enable automatic ingestion of course materials.

The ultimate vision extends beyond a single application to a paradigm shift in how students approach learning. By demonstrating that AI can eliminate the practical barriers to evidence-based study techniques, we hope to inspire broader integration of learning science principles into educational technology design.

ACKNOWLEDGMENT

The authors thank the early pilot study participants who provided invaluable feedback on system usability and effectiveness. We acknowledge the contributions of web survey respondents whose data enable evidence-based algorithm calibration. We are grateful to the open-source communities behind PyQt5, scikit-learn, Hugging Face Transformers, and spaCy, whose tools made rapid development possible. Finally, we thank our academic mentors and peers at Maharaja Surajmal Institute of Technology for guidance and support throughout this research.

REFERENCES

- [1] N. J. Cepeda, H. Pashler, E. Vul, J. T. Wixted, and D. Rohrer, "Distributed practice in verbal recall tasks: A review and quantitative synthesis," *Psychol. Bull.*, vol. 132, no. 3, pp. 354–380, 2006.
- [2] S. Reddy, I. Labutov, S. Banerjee, and T. Joachims, "Enhancing human learning via spaced repetition optimization," *Proc. Natl. Acad. Sci. USA*, vol. 116, no. 10, pp. 3988–3993, 2019.
- [3] J. Ye, "FSRS: A modern spaced repetition algorithm for learning complex subjects," Open Spaced Repetition Project, 2024.
- [4] L. Li, D. Demszky, P. Bromley, and D. Jurafsky, "Content analysis of textbooks via natural language processing: Findings on gender, race, and ethnicity in Texas U.S. history textbooks," *AERA Open*, vol. 6, no. 3, 2020.
- [5] M. Chen, B. Wang, and L. Zhang, "Automated content analysis of educational feedback: A multi-language study using machine learning classifiers," *Comput. Educ.*, vol. 185, p. 104118, 2022.
- [6] A. Pashler, N. Rohrer, N. J. Cepeda, and S. K. Carpenter, "Enhancing learning and retarding forgetting: Choices and consequences," *Psychon. Bull. Rev.*, vol. 14, no. 2, pp. 187–193, 2007.

Semantic Telemetry: Extreme Lossy Compression for Dead-Zone Weather Stations

Devesh Sharma^{#,1}, Abhas Paul^{#,2}, Yugul Gupta^{#,3}

[#]*Department of Computer Science and Engineering, Maharaja Surajmal Institute*

Abstract: Remote weather stations in mountains, oceans, and polar regions face severe limits in satellite bandwidth and power, making transmission of raw sensor data costly and energy-intensive. We propose Semantic Telemetry, an extreme lossy compression paradigm using on-device Small Language Models (SLMs) to convert sensor readings into concise semantic tokens (on the order of 3 bytes) for transmission. A twin SLM at the cloud end reconstructs plausible weather data from these tokens. Our pipeline has two stages: (1) SLM-based encoding of numerical readings into semantic representations (text or latent vectors), and (2) entropy coding of the token stream. We describe the system architecture and training process, including VAE-based tokenization and quantization. Safety features such as anomaly detection and fallback alerts ensure critical events are preserved. Through simulation on a year-long Himalayan dataset, we demonstrate compression ratios of 50–100× with small absolute errors (e.g. $\pm 0.7^\circ\text{C}$ MAE). Cost analysis shows >98% savings in satellite fees. Finally, we outline applications (precision agriculture to Arctic monitoring), key failure modes (model drift, token corruption), and future work such as adaptive token sizes and federated learning for deployment.

Keywords: Semantic compression, Small Language Models, lossy telemetry, edge computing, satellite IoT, weather forecasting, dead-zone monitoring.

I. INTRODUCTION

Remote environmental monitoring systems rely on satellite and LPWAN links that offer only tens to a few hundred bps and impose high per-byte costs [3], [11]. Even compact JSON weather payloads (~58 bytes) expand to ~300-byte satellite packets after framing overhead, leading to annual communication expenses of several thousand dollars per station in practical deployments. Since radio transmission dominates node energy use in many deployed designs, bit-rate reduction directly improves operational lifetime [1]. Conventional compressors—gzip, LZW, delta coding—provide limited gains on low-entropy numerical streams and cannot bypass Shannon entropy limits for atmospheric data without leveraging domain structure [4]. These constraints motivate semantic lossy compression, wherein the goal shifts from exact reproduction to preservation of task-relevant meaning and downstream utility, aligning with recent work on semantic communications and task-oriented compression [13], [14]. Let x_t denote true measurements and \hat{x}_t the semantic reconstruction. The design objective

minimizes a bounded distortion function $L(x_t, \hat{x}_t)$, such as MAE or RMSE.

A. Problem Motivation

We consider systems with constrained uplink bandwidth, high per-byte cost, and tight edge-energy budgets typical of satellite-IoT and LoRaWAN deployments. Raw packet size S_{packet} inflates due to protocol overhead, and total yearly cost grows with the number of transmissions. Because radio energy per bit is approximately constant in these systems, reducing S_{packet} decreases both operational cost and power consumption. Semantic compression provides this reduction by transmitting interpretable summaries instead of raw measurements.

B. Semantic Telemetry Paradigm

Semantic Telemetry performs on-device meaning extraction using SLMs [9]. Raw measurements are transformed into compact latent tokens, entropy-coded, and transmitted. The cloud SLM reconstructs weather states using local priors and contextual metadata, drawing on semantic communications theory and recent neural compression results [2]. Tokenization uses a VAE or VQ-VAE-style quantizer, with the transmitted bitrate bounded by latent entropy $H(Z)$ [4]. Safety is ensured via confidence scoring and anomaly detection, triggering fallbacks during extreme or out-of-distribution events.

C. Contributions

1. A principled semantic-compression formulation for multivariate environmental telemetry using discrete tokenization and entropy coding [7].
2. A paired edge-cloud SLM architecture supporting lossy yet semantically faithful reconstruction [5].
3. Integrated safety components including anomaly scoring and deterministic fallbacks.
4. Cost and energy accounting models enabling deployment planning under satellite-IoT constraints.
5. Identification of application domains and operational limitations, including regulatory-grade measurement scenarios.

II. RELATED WORK

A. IoT Compression and Energy Saving

Traditional lossless schemes (RLE, Huffman, delta coding) reduce transmissions only marginally because they treat all signals equivalently and ignore semantics. Event-driven sensing lowers data rates but still operates within entropy-bounded limits. Such methods remain insufficient for ultra-low-bandwidth satellite IoT systems [3], [11].

B. Semantic Communication Models

Semantic communication emphasizes meaning rather than symbol accuracy, supported theoretically by Shannon entropy and rate–distortion functions. Emerging AI-driven systems transmit concise action-oriented messages instead of continuous numerical streams, [14], but prior work has not applied this to extreme-compression weather telemetry [8].

C. Neural Compression for Earth Sciences

Neural compressors—including VAEs, hierarchical autoencoders, and transformer-based climate compressors—achieve 50–200× compression ratios on high-bandwidth datasets [2], [4]. However, these methods assume server-class compute and large inputs; they do not address edge nodes constrained to transmitting only a few bytes.

D. Edge SLMs

Advances in TinyML and SLM quantization enable sub-100M-parameter LLM variants on microcontrollers [6], [12]. Prior works focus on prediction or activity recognition, not lossy semantic encoding.

E. Weather Telemetry Constraints

Operational standards such as SYNOP produce fixed-size messages that exceed practical satellite budgets for high-frequency reporting [10]. Existing systems apply mild compression (≈ 30 – 40%), lacking semantic reduction.

III. SYSTEM ARCHITECTURE AND METHODOLOGY

A. Edge Encoder (On-Device)

A compact edge encoder runs on a low-power weather node (e.g., ESP32, ARM Cortex, Pi Zero 2W). Standard environmental sensors produce the multivariate state (temperature, humidity, pressure, wind, ...). A quantized small language model (SLM) converts recent numeric context into a short semantic summary (text or tokens). That summary is mapped by a lightweight VAE encoder into a

low-dimensional continuous latent (design example: 24 dimensions), and the latent is discretized into a fixed-size semantic token (design example: 24 bits \approx 3 bytes) via vector/product quantization and a small codebook.

B. Edge Sensor Module and On-Device SLM

Hardware consists of sensors plus a constrained MCU/SBC; software comprises a distilled, low-precision SLM that ingests a short window of recent measurements and emits a human/interpretable summary (e.g., “14:00 — 18°C, light drizzle”). The VAE encoder projects that summary into the compact latent used for quantization. Design choices (latent dimensionality, codebook layout, SLM size/precision) trade reconstruction fidelity against edge compute, memory, and energy.

$$\mathbf{w} = [T, H, P, W, \dots] \quad (1)$$

$$s_t = E_{\text{SLM}}(\mathbf{w}_{t-k:t}) \quad (2)$$

$$z_t = E_{\text{VAE}}(s_t), \quad z_t \in \mathbb{R}^d \quad (3)$$

$$t_t = Q(z_t), \quad t_t \in \{0, \dots, 2^B - 1\}, \quad |t_t| = B \text{ bits} \quad (4)$$

C. Quantization, Tokenization and Rate Constraint

Continuous latents are mapped to discrete indices using a product-quantization or VQ-VAE style codebook to meet a hard uplink budget (example: 3 sub-vectors \times 256 centroids \rightarrow 24 bits). Entropy coding may optionally further reduce expected bitrate. Token design aims for very small fixed payloads to simplify packet framing and integrity checks on constrained links.

D. Entropy Coding & Transmission

A lightweight entropy coder (e.g., Huffman/arithmetic) can be applied to the token stream to approach the latent entropy. Packets carry the rounded-up byte payload plus minimal protocol overhead (sequence number, integrity check). Transmission targets constrained uplinks (LoRaWAN, Iridium), and implementation must account for airtime, maximum payload, and link-specific framing.

E. Cloud Decoder and Reconstruction

On receipt, the cloud looks up the codebook entry to reconstruct the latent, decodes the latent via the VAE decoder into a semantic summary, and conditions a larger SLM (with richer priors/context) to generate probabilistic numeric reconstructions. The cloud reports point estimates and calibrated uncertainty (e.g., ensemble/MC variance or conformal intervals). Reconstruction quality is assessed with per-channel error measures (e.g., MAE) and aggregated scores (e.g., RMSE).

F. Formal Workflow (Algorithm)

1. Sample sensors.
2. Edge SLM \rightarrow produce semantic summary.
3. VAE encode \rightarrow latent.
4. Quantize latent \rightarrow fixed-size token.
5. (Optional) Entropy-code and transmit token + integrity header.
6. Cloud dequantize \rightarrow VAE decode \rightarrow SLM-conditioned reconstruction.
7. Compute anomaly/confidence; if triggered, execute fallback policy.

G. Energy, Cost and Optimization Accounting

Total per-report energy combines on-device inference cost, radio transmission energy (dominant for large raw packets), and sensing/MCU overhead. Monetary cost is driven by billed payload bytes plus occasional fallbacks; these accounting elements guide choices for token size, reporting cadence, and fallback thresholds to meet operational budgets.

$$S_{\text{packet}} = \left\lceil \frac{\bar{b}}{8} \right\rceil + S_{\text{overhead}} \quad (\text{bytes}) \quad (5)$$

$$E_{\text{tx}} = e_{\text{bit}} \cdot \bar{b} \quad (6)$$


$$\tilde{z}_t = Q^{-1}(t_t), \quad \hat{s}_t = D_{\text{VAE}}(\tilde{z}_t), \quad \hat{\mathbf{w}}_t \sim p_{\text{SLM}}(\mathbf{w} | \hat{s}_t, C_t) \quad (7)$$

$$\text{MAE}_T = \frac{1}{T} \sum_{t=1}^T |x_t - \hat{x}_t|, \quad \text{RMSE} = \sqrt{\frac{1}{T} \sum_{t=1}^T \|x_t - \hat{x}_t\|_2^2} \quad (8)$$

$$E_{\text{total}} = E_{\text{inf}} + E_{\text{tx}} + E_{\text{ov}} \downarrow, \quad \text{Cost}_{\text{packet}} = \frac{S_{\text{packet}}}{2^{20}} \cdot c \quad (9)$$

H. Safety, Anomaly Detection and Fallbacks

Each token is annotated with a confidence score; the node also computes an anomaly score from raw features or residuals. Conservative safety logic enforces physical bounds and temporal consistency; when thresholds are exceeded the node either transmits a raw high-fidelity packet or sends an explicit alert bit, ensuring critical events are preserved.



I. Model Synchronization, Versioning and Robustness

Robust decoding requires synchronized tokenizer, codebook, and model versions. Practical measures include embedding version identifiers in packet headers, cloud-side codebook selection based on ID, and over-the-air update pathways with automatic rollback. These practices mitigate decoding failures from model drift or heterogeneous deployments.

IV. TRAINING AND SAFETY MECHANISMS

A. SLM pretraining and fine-tuning

The edge encoder SLM is first pretrained on broad text corpora (or distilled from a larger teacher) and then fine-tuned on meteorological sources (weather reports, SYNOP logs, station bulletins) so its token space is weather-aware. Fine-tuning minimizes standard sequence prediction loss (cross-entropy) and may include a distillation term when compressing a teacher model into a smaller student. Practical training recipes include meteorological data augmentations (time, location, seasonal tags) and early stopping guided by validation perplexity to prevent overfitting.

B. VAE / latent training objective

A VAE maps the SLM summaries (text or discrete tokens) to continuous latents; training optimizes the usual evidence lower bound (reconstruction term plus a posterior regularizer). The reconstruction term uses cross-entropy for discrete/text inputs or mean-squared error for numeric windows. Experimental settings used in this work: latent dimension 24, an 80/20 train/validation split, and 100 epochs; the choice of reconstruction loss followed the input modality.

C. Discrete representation via product quantization (codebook)

To meet a hard 24-bit uplink budget the continuous latent is discretized using product quantization or a VQ-VAE style codebook. The latent is split into a small number of sub-


vectors, each assigned to a centroid index; the concatenated indices produce a fixed-size 24-bit token. A practical implementation that meets 24 bits is to use three sub-vectors with 256 centroids each (3×8 bits), giving a compact 3-byte payload and a simple lookup-table decoder on the cloud side.

D. Model and latent quantization for edge memory/compute

To fit SLM+VAE on constrained hardware we apply post-training quantization (e.g., 4–8 bit weights), pruning, and distillation to reduce parameter count and memory. Model size is thus primarily determined by parameter count and chosen weight precision; latent payloads remain fixed at 24 bits (prior to optional entropy coding). These choices balance edge storage/compute against inference accuracy and energy.

E. Safety checks, anomaly scoring and fallback logic

Because generative reconstructions can be plausible but incorrect, we implement a multi-stage safety pipeline: (1) physical range checks that enforce hard bounds per channel (for example acceptable temperature and humidity ranges), (2) temporal consistency checks that limit per-interval jumps, and (3) an anomaly score computed from residuals or likelihoods conditioned on context. A predictive confidence metric is derived from model uncertainty (predictive entropy or ensemble/MC methods). When an anomaly or low confidence is detected the node triggers a fallback: the cloud may request a raw high-fidelity packet or the edge may proactively transmit one. In our simulations fallbacks were rare ($\approx 0.4\%$ of intervals, concentrated around extreme storm cases).



F. Uncertainty quantification and human-in-the-loop review

We obtain calibrated uncertainty via Monte Carlo dropout, ensembles, or conformal prediction. Conformal intervals (derived from calibration residuals) provide finite-sample coverage guarantees under standard exchangeability assumptions; wide or low-coverage intervals flag records for manual review or targeted retraining. This human-in-the-loop path is used selectively to curate difficult or high-risk examples.

G. Continual learning & dataset curation for robustness

Reports triggering fallbacks or flagged for low confidence are collected into a curated dataset for periodic retraining (centralized or federated). This periodic fine-tuning reduces model drift and improves performance on rare events while respecting bandwidth constraints by prioritizing compressed latents or aggregated updates over raw data transfers.

H. Cost and risk accounting for fallbacks

Deployment accounting explicitly includes fallback cost: expected extra annual cost equals the number of reports times the fallback probability times the per-raw-packet cost. This added term is incorporated into overall telemetry economics when choosing token size, reporting cadence, and acceptable fallback thresholds.

V. EXPERIMENTS AND EVALUATION

A. Dataset

Experiments used a twelve-month meteorological dataset from a high-altitude Himalayan station in Nainital ($\approx 3,500$ m). Measurements were sampled every 15 minutes, yielding 35,040 total records. The temperature ranged from -18 °C to $+32$ °C, humidity from 15% to 100%, wind from 0–68 km/h, and pressure from 950–1000 hPa. All channels were normalized and discretized to 0.1 °C and 1% RH precision as preprocessing.

B. Models and Simulation Setup

The edge device simulation used a quantized Phi-2 SLM (2.7B parameters reduced to 4-bit precision) paired with a lightweight VAE that maps 512-dimensional features into a 24-dimensional latent representation. The cloud model employed a larger Gemini Nano variant fine-tuned on 5,000 synoptic weather reports. The VAE was trained for 100 epochs with an 80/20 train-validation split. Latent compression used the 24-bit product-quantized codebook introduced in Section IV, enabling a single-token uplink transmission.

C. Baselines and Metrics

We compared Semantic Telemetry against several baselines:

1. **Raw JSON packets** (≈ 256 bytes each including structural overhead).
2. **gzip compression**, applying standard DEFLATE to the JSON.
3. **Delta-coding with thresholds**, transmitting only when channel changes exceed fixed limits.
4. **Integer quantization** based on coarse rounding schemes.


Evaluation metrics included: – **Compression ratio**, defined as the ratio between the size of the raw packet and the transmitted packet. – **Per-channel mean absolute error (MAE)** and **root-mean-square error (RMSE)** computed between reconstructed and ground-truth sequences. – **Energy and monetary cost**, modeled as proportional to transmitted bytes using fixed Joules-per-byte and dollars-per-megabyte coefficients.

D. Aggregate Results

Semantic Telemetry reduced a standard 256-byte report to a **3-byte latent token** plus a **1-byte CRC**, producing a typical 4-byte payload. Fallbacks (transmission of the full high-fidelity packet when confidence is low) occurred in approximately 0.4% of intervals. Annual transmitted volume therefore equals the sum of token traffic and occasional raw packets, with the overall compression ratio ranging from **50× to 100×** depending on fallback frequency and overhead.

Reconstruction performance remained within operational tolerance. Typical values included a temperature MAE near 0.7 °C, humidity MAE around 2%, and wind MAE under 1 km/h. Across the dataset, 96.2% of temperature reconstructions were within ±2 °C, and only 0.2% exceeded ±5 °C; all such cases were correctly identified by the safety subsystem.

Cost and energy accounting showed substantial reductions. For an annual dataset of 35,040 transmissions and a 0.4% fallback rate, the raw system would send roughly 9 MB, whereas the semantic system transmitted approximately 0.11 MB. Using a representative cost rate of \$10 per MB, the annual cost decreased from ≈\$2,687 to ≈\$32, representing **about 99% savings**, with a comparable reduction in energy consumption due to the byte-proportional radio model.



E. Qualitative Evaluation

Qualitative assessment showed that semantic reconstructions preserved decision-relevant meaning. Events such as pre-dawn cooling, gusty periods, and humidity surges were rendered as compact text summaries that remained actionable despite small numerical deviations. Table-top comparisons against baselines highlight that

traditional methods offer modest compression with no semantic gains, whereas Semantic Telemetry delivers high compression alongside meaningful interpretability.

Method	CR (×)	Temp MAE (°C)	Humidity MAE (%)	Comments
Raw(no compression)	1×	0	0	Baseline
gzip	~3–5×	0	0	Lossless
Delta-coding	~8–10×	0	0	Event-only
Prior VAE methods	~50×	~0.8	–	Literature
Semantic Telemetry	85–100×	~0.7	~2.0	24-bit token + rare fallback

VI. APPLICATIONS AND USE CASES

A. Precision Agriculture (Dense Micro-Stations)

Semantic Telemetry allows low-cost soil and microclimate nodes to report hourly high-level summaries such as “*Soil very dry; irrigation recommended*” in place of transmitting full raw logs. Since packet-based operating cost is proportional to transmitted bytes, reducing packet size from hundreds of bytes to only a few bytes yields a linear reduction in annual operating cost per station. When scaled to large farms with many nodes, the cumulative savings enable deployment of more stations per hectare and support higher reporting frequency without increasing network expenditure.

B. Environmental Monitoring and Himalayan Case Study

A case study with 25 high-altitude Himalayan stations demonstrates the economic advantage. The existing system cost approximately \$67,000 per year, while the semantic system—costing around \$33 annually per station—reduced total expenditure to roughly \$825 per year. The resulting savings could, in principle, be reinvested to deploy more than two thousand additional stations of equivalent cost.

Reduced packet sizes also improve latency: shorter payloads decrease airtime and queuing delays, lowering overall end-to-end report latency roughly in proportion to the reduction in transmitted bytes.

C. Oceanic Buoys (Iridium, Battery-Powered)

Battery-powered ocean buoys benefit strongly from reduced communication energy. Annual transmission energy scales directly with the number of bytes sent per report. Semantic tokens typically reduce transmitted bytes by one to two orders of magnitude; hence, communication energy consumption falls by a comparable factor. In idealized conditions this extends battery life proportionally. For example, an 18-month baseline lifetime could theoretically extend to more than a century if communication were the only drain. In practice, non-radio loads, self-discharge, and protocol overhead yield more conservative gains—multi-year lifetimes (such as 20+ years) remain plausible and operationally meaningful.

D. Arctic Monitoring (Temporal Density Gains)

Stations that previously reported infrequently (e.g., once per month) can shift to much denser cadences such as 4-hourly or hourly reporting. Moving from monthly to 4-hourly provides more than a hundredfold increase in temporal density, whereas hourly sampling delivers more than seven hundredfold improvement. These upgrades become cost-feasible because per-report cost scales with packet size, and semantic packets remain extremely small even at high sampling frequencies.

E. Broader IoT and Cross-Domain Applications

Any application where high-level interpretations are more valuable than raw sensor streams can adopt Semantic Telemetry. Examples include industrial systems generating short alerts such as *“bearing vibration trending high”* instead of transmitting full spectrograms; wearable devices signaling *“arrhythmia detected”* without continuous ECG uploads; and wildfire detection networks sending compact semantic alerts enriched with time and location metadata. These scenarios rely on downstream consumers that can accommodate probabilistic reconstructions and semantic summaries rather than legally exact raw data.

F. Deployment Notes, Planning Considerations, and Failure Modes

Network planners can estimate expenditure by multiplying the number of stations, reporting rate, and packet size; the semantic approach directly lowers this cost through its small payloads. Operational tuning involves selecting token size, fallback rate, and reporting cadence to minimize cost while meeting safety, accuracy, and latency requirements.

Potential failure modes include model drift, token corruption, and unmodeled rare events. Mitigations include periodic retraining or federated updates, lightweight integrity checks with retransmission, and conservative anomaly thresholds that trigger fallbacks or explicit alerts during uncertain or extreme conditions.

VII. LIMITATIONS AND FAILURE MODES

A. Model Hallucination (Plausible-but-Incorrect Outputs)

Generative SLMs may produce outputs that appear coherent yet do not reflect the true physical state—for example, generating humidity–temperature pairs that fit learned statistical patterns but contradict sensor reality. Hallucination risk increases when the model assigns high confidence to synthetic values that deviate from true environmental conditions.

A common mitigation is to use Monte Carlo or ensemble sampling. Multiple reconstructions are generated and averaged, while the variance of these samples provides an uncertainty indicator. Larger ensembles reduce variance but proportionally increase compute and energy consumption on the device, requiring a balance between reliability and resource usage.

B. Extreme Outliers and Unseen Events

Rapid or rare meteorological events—such as abrupt pressure drops, hail onset, or strong gust fronts—may fall outside the training distribution and be smoothed or underrepresented by the model. To guard against this, simple temporal-rate checks are used: if the short-term change in any channel exceeds a conservative threshold, the system immediately transmits a raw packet. This ensures integrity during fast-evolving extreme events, though subtle anomalies that remain within thresholds constitute the main residual risk.

C. Model Drift (Seasonal and Long-Term Shifts)

Model drift occurs when real-world conditions gradually change while the deployed model remains static. Seasonal cycles, long-term climatic trends, or small biases introduced by quantization can degrade accuracy over time. Drift is monitored by comparing occasional raw measurements against reconstructed values over a sliding window. If accumulated bias exceeds a predefined tolerance, the system schedules an update—either via over-the-air model replacement or through periodic federated fine-tuning based on aggregated (not raw) data to preserve bandwidth and privacy.

D. Token Corruption and Transmission Errors

Compact tokens are highly sensitive to bit errors: flipping a single bit may map to an entirely different latent codebook entry. Even with low bit-error-rates typical of satellite links, small tokens have a non-negligible probability of corruption. Adding a minimal integrity mechanism—such as a one-byte CRC and sequence number—makes corrupted packets easily detectable, enabling retransmission or graceful dropout. The added overhead remains small and maintains large compression gains relative to raw packets.

E. Connectivity Outages and Missing Tokens

Extended communication gaps (e.g., during polar-orbit satellite passes or congestion) create missing intervals in the data stream. When outages are short, the cloud model can generate short-term forecasts conditioned on the last received context; these reconstructions are explicitly marked as synthetic and include uncertainty estimates. For outages beyond an allowable duration, the system falls back to coarser summaries or prioritized alerts upon reconnection to avoid overwhelming the link.

F. Compute, Energy, and Duty-Cycle Constraints on Tiny Nodes

Even quantized SLMs require meaningful on-device computation, and inference consumes energy alongside radio transmission. The overall average power depends on inference cost, transmission cost, and reporting frequency. Radio energy typically dominates: transmitting a multi-byte token is orders of magnitude cheaper than sending a full raw packet. As long as inference energy remains modest relative to the reduced transmission energy, Semantic Telemetry provides net energy gains. For ultra-low-power devices, simpler summarizers or reduced reporting cadences may be preferable to maintain battery longevity.

G. Aggregate Failure Probability and Fallback Economics

Fallbacks—cases where a full raw packet is transmitted instead of a token—add a small recurring cost proportional to their frequency. Observed fallback rates have been below half a percent, but planners should include this overhead in budgeting.

Overall system failure probability arises from several independent sources: hallucination, token corruption, undetected drift, or missed extreme events. Conservative configurations set thresholds and safety parameters such that the combined probability remains within acceptable operational limits. These guardrails ensure the system remains reliable even under rare or adverse conditions.

VIII. FUTURE DIRECTIONS

A. Adaptive Token Sizing (Variable Bitrates)

The present system uses a fixed 24-bit semantic token for every report. However, environmental variability is not constant, and future deployments may benefit from variable-length tokens whose size adapts to real-time volatility. A simple volatility estimator—based on the short-term change between consecutive measurements—can guide this adaptation. During calm periods the system may transmit minimal tokens, while periods of rapid or uncertain change may use expanded payloads to preserve fidelity. The average packet size becomes an expectation over these regimes, enabling very high compression during quiescent conditions while allocating additional bits when extreme events demand richer semantic detail.

B. Multi-Modal Telemetry Beyond Numeric Sensors

Edge models of the future could incorporate richer sensor modalities such as image snapshots, low-rate video clips, microphone-based rainfall cues, or compact radar signatures. A unified encoder could fuse text-based summaries, numeric sensor streams, and lightweight vision or audio embeddings into a shared latent space. This additional context is likely to improve reconstruction in complex or rare scenarios—such as hailstorms or fast-moving frontal boundaries—though at increased compute and storage requirements on constrained nodes.

C. Federated and Continual Learning for Climate Drift

To remain accurate under long-term climatic shifts, seasonal cycles, or sensor aging, deployed models must periodically update. Federated learning is a natural approach: each station fine-tunes its local model using occasional raw packets or verified labels, and the cloud aggregates these updates into a global model. Only compressed latents, low-frequency calibration data, or lightweight gradient summaries need to be transmitted, minimizing bandwidth consumption and preserving privacy while maintaining accuracy across a diverse, evolving sensor network.

D. Hybrid ML–Symbolic Consistency and Uncertainty Reasoning

While SLMs excel at semantic interpolation, they cannot inherently enforce physical rules or monotonic constraints. Future decoders could integrate symbolic or rule-based components to ensure consistency with domain knowledge—for example, enforcing physically plausible temperature transitions at sunrise. Coupling these constraints with calibrated uncertainty intervals would enable reconstructions to carry “trustworthiness flags,” allowing downstream systems to treat some summaries as advisory and others as authoritative. Such metadata is aligned with the goals of emerging semantic-aware 6G communication frameworks.

E. Field Trials at National Scale

Large-scale deployment remains the critical next step. Existing weather stations equipped with modest single-board computers could host the necessary on-device SLM and VAE components. National or oceanic pilots should quantify real-world bit-error rates under harsh conditions, user acceptance of semantically reconstructed data, and operational fallback loads during rare but high-impact events. Real fallback rates observed in field trials must be incorporated into deployment cost calculations to ensure that the system remains economically viable at scale.

IX. CONCLUSION

This work presented a semantic, highly lossy compression framework for remote weather telemetry in which on-device small language models generate compact latent tokens in place of full numerical sensor streams. By transmitting meaning-centered representations rather than raw measurements, the system achieves substantial bandwidth reductions—typically reducing packet size by nearly two orders of magnitude and lowering satellite communication cost by approximately 98–99% for off-grid stations.

Despite its lossy nature, the method preserves the information required for operational decision-making. Reconstruction accuracy remained within acceptable tolerance for the vast majority of measurements, and integrated safety mechanisms—including range checks, temporal consistency tests, uncertainty quantification, and fallback logic—ensured robust handling of extreme or rapidly evolving weather events.

Beyond immediate efficiency gains, this work reframes remote sensing as a semantic communication task, where information value is defined by downstream utility rather than exact numerical fidelity. Challenges remain, including model drift, rare-event coverage, and compute constraints on ultra-low-power nodes. Nevertheless, the results demonstrate that semantic telemetry can enable dense, low-cost environmental monitoring using ultra-compact messages, opening a path toward scalable networks of autonomous micro-stations in bandwidth- and energy-limited environments.

ACKNOWLEDGMENT

We thank the NCRTCES organizers and the open-source climate and ML communities for datasets, tools, and discussions that shaped our experiments. We also acknowledge contributors advancing accessible weather data in remote regions, inspiring the vision of large-scale, low-cost semantic telemetry.

REFERENCES

- [1] P. B. Leelavinodhan *et al.*, “Design and Implementation of an Energy-Efficient Weather Station for Wind Data Collection,” *Sensors*, vol. 21, no. 11, p. 3831, 2021.
- [2] T. Han *et al.*, “Climate science data can be compressed efficiently by dual-stage extreme compression with a variational auto-encoder transformer,” *Commun. Earth Environ.*, vol. 6, 2025.
- [3] O. Vikhrova *et al.*, “LoRaWAN-Based Networks for IoT: A Review of Satellite-Based Solutions,” *IEEE Internet of Things Journal*, vol. 8, no. 18, pp. 14227–14240, 2021.
- [4] M. Klöwer *et al.*, “Compressing atmospheric data into its real information content,” *Nature Computational Science*, vol. 1, pp. 713–724, 2021.
- [5] R. Ballester-Ripoll *et al.*, “Sobol Tensor Trains for Global Sensitivity Analysis,” *Reliability Engineering & System Safety*, vol. 223, 108428, 2022.
- [6] Q. Song *et al.*, “Leveraging Foundation Models for Zero-Shot IoT Sensing,”
- [7] P. Zhang *et al.*, “TinyM2Net: A Flexible System-Algorithm Co-designed Tiny Deep Learning Framework for Microcontrollers,” *IEEE Transactions on Mobile Computing*, vol. 22, no. 10, pp. 5661–5676, 2023.
- [8] S. Lin *et al.*, “MCUNet: Tiny Deep Learning on IoT Devices,” *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, pp. 11711–11722, 2020.
- [9] Microsoft Azure, “What are Small Language Models (SLMs)?”, *Azure Cloud Dictionary*, 2023.
- [10] World Meteorological Organization (WMO), “Manual on Codes – International Codes, Volume I.1,” *WMO-No. 306*, 2023.
- [11] T. Elshabrawy *et al.*, “Modeling and Optimization of LoRa Networks,” *IEEE Sensors Journal*, vol. 19, no. 15, pp. 6512–6526, 2019.
- [12] Y. Zhou *et al.*, “TENT: Connect Language Models with IoT Sensors for Zero-Shot Activity Recognition,” *arXiv preprint arXiv:2311.08245*, 2023.
- [13] H. Xie *et al.*, “Deep Learning Enabled Semantic Communication Systems,” *IEEE Transactions on Signal Processing*, vol. 69, pp. 2663–2675, 2021.
- [14] Z. Qin *et al.*, “Semantic Communications: Principles and Challenges,” *IEEE Internet of Things Journal*, vol. 9, no. 21, 2022.

Differentially Private Blockchain-Based Climate Forecasting Framework

Dhruv Ghosal^{#,1}, Daksh Malhotra^{#,2}, Aastha Malik^{#,3}, Falguni Singh^{*,4}, Keshav Gupta^{#,5}, Ishaan Chaturvedi^{*,6}

[#]*Dept. of Computer Science and Engineering, Maharaja Surajmal Institute of Technology, New Delhi, India*

¹dhruv_cse_28@msit.in, ²iamdakshmalhotra@gmail.com, ³aasthamalik.work@gmail.com;

⁵keshavgupta07delhi@gmail.com

^{*}*Dept. of Information Technology, Maharaja Surajmal Institute of Technology, New Delhi, India*

⁴falguniipshita@gmail.com, ⁶linkedwithishaan@gmail.com

Abstract—Large volumes of high-resolution meteorological data are collected from geographically distributed weather stations, satellites, and IoT sensors for climate forecasting and climate change analytics. However, storing and sharing such climate data centrally introduces serious challenges related to privacy, security, data ownership, and trust among multiple stakeholders. Furthermore, machine learning models trained on centralized datasets are vulnerable to inference attacks that may expose sensitive environmental or institutional information. This paper proposes a Differentially Private Blockchain-Based Climate Forecasting Framework that integrates blockchain for decentralized and tamper-resistant data governance with differential privacy for formal and provable privacy guarantees during machine learning model training. Climate data are kept off-chain, but cryptographic hashes and metadata are recorded in a blockchain ledger to guarantee integrity, transparency, and traceability. Climate forecasting models are trained through Differentially Private Stochastic Gradient Descent (DP-SGD), where carefully calibrated Gaussian noise is injected into model gradients based on sensitivity analysis and privacy budget accounting. Mathematical formulations of differential privacy, noise mechanisms, gradient clipping, and privacy composition are provided. Experimental evaluation demonstrates that strong privacy protection can be achieved with only marginal degradation in forecasting accuracy. The proposed framework provides secure, collaborative, and sustainable climate intelligence systems that are suitable for multi-organizational and long-term environmental monitoring.

Keywords—Differential Privacy, Blockchain Technology, Climate Forecasting, Privacy-Preserving Machine Learning, Differentially Private Stochastic Gradient Descent, Decentralized Data Storage, Sustainable Artificial Intelligence

I. INTRODUCTION

A. Motivation and Context

Climate forecasting plays a vital role in disaster risk reduction, agricultural planning, water resource management, urban infrastructure design, and long-term sustainable development. Increasing climate variability and the frequency of extreme weather events have significantly amplified the demand for accurate and reliable forecasting models. Modern climate forecasting

systems rely on large-scale meteorological datasets collected from heterogeneous and geographically distributed sources, including weather stations, satellites, and IoT-based environmental sensors.

Despite their importance, climate datasets are often managed by multiple organizations with differing legal, ethical, and institutional constraints. Centralized data collection and storage raise significant concerns regarding privacy leakage, data ownership, misuse, and unauthorized access. Furthermore, centralized architectures introduce single points of failure that are unsuitable for mission-critical climate infrastructure.

Recent advances in machine learning (ML) have greatly improved forecasting accuracy. However, ML models trained on centralized datasets are vulnerable to inference attacks such as membership inference and model inversion, which may reveal sensitive information about the training data. These risks limit data sharing and hinder collaborative climate modelling. Blockchain technology provides a decentralized, immutable, and transparent data management paradigm that improves trust and accountability among untrusted participants. However, blockchain alone cannot prevent information leakage through trained ML models. Differential Privacy (DP) provides a rigorous mathematical framework that limits the influence of any single data record on model outputs, offering formal privacy guarantees. This work integrates blockchain-based decentralized data governance with differentially private machine learning to enable secure, collaborative, and privacy-preserving climate forecasting.

B. Problem Statement

The core research problem addressed in this paper is:

How can multiple organizations collaboratively train accurate climate forecasting models while ensuring decentralized data governance, data integrity, and provable privacy guarantees against inference attacks?

C. Contributions

The main contributions of this paper are:

1. A decentralized blockchain-based architecture for climate data governance.
2. A privacy-preserving climate forecasting framework with formal (ϵ, δ) -differential privacy guarantees.
3. Mathematical formulation of sensitivity analysis, noise calibration, and privacy budget composition.
4. Application of DP-SGD to climate forecasting models.
5. Quantitative analysis of the privacy–utility trade-off.
6. Discussion of system limitations and future development directions.

II. RELATED WORK

Climate forecasting research traditionally *focuses* on numerical weather prediction and, more recently, deep learning- based models. While these approaches achieve high accuracy, they generally assume unrestricted access to centralized datasets.

Blockchain technology has been explored for environmental monitoring and data provenance, emphasizing transparency and integrity. However, most blockchain-based climate systems do not address privacy leakage through machine learning models.

Differential privacy has been widely studied in healthcare, finance, and census applications. Its integration with blockchain and application to climate forecasting remains underexplored. This paper bridges this gap by combining blockchain-based decentralization with formally private machine learning.

III. SYSTEM ARCHITECTURE

A. Architecture Overview

The proposed system consists of four components:

1. **Data Providers:** Weather stations and IoT sensors generating climate data.
2. **Blockchain Network:** Stores cryptographic hashes and metadata.
3. **Off-Chain Storage:** Stores raw climate datasets.
4. **Privacy-Preserving ML Engine:** Trains forecasting models using DP-SGD.

B. Data Flow

Each sensor i collects climate data D . A cryptographic hash is computed as:

$$H_i = \text{SHA256}(D_i) \quad (1)$$

The hash H_i and metadata are stored on the blockchain, while raw data are stored off-chain. Verified data are retrieved for model training, where differential privacy mechanisms are applied.

C. Security and Scalability

The proposed framework is designed to provide strong security guarantees while maintaining scalability for large-scale, multi-organizational climate forecasting systems. Climate data infrastructures must support high data volumes, heterogeneous data sources, and long-term continuous operation while resisting adversarial behavior and system failures.

1) *Data Integrity and Tamper Resistance:* Blockchain technology ensures data integrity by storing cryptographic hashes of climate datasets rather than raw data. For each dataset D_i generated by sensor i , a secure hash is computed as:

$$H_i = \text{HA256}(D_i) \quad (2)$$

Any modification to the dataset results in a different hash, making unauthorized data tampering immediately detectable. Because the blockchain ledger is append-only and replicated across multiple nodes, it provides strong resistance against data manipulation, rollback attacks, and unauthorized deletion.

2) *Access Control and Trust Management:* Access to off-chain climate data is governed by blockchain-recorded metadata and smart-contract-based access control policies. Only entities that can present a valid hash reference recorded on the blockchain are permitted to retrieve corresponding datasets. This mechanism establishes a decentralized trust model in which no single authority controls data access.

Furthermore, role-based access control (RBAC) can be enforced using smart contracts, allowing fine-grained permissions for data providers, model trainers, and auditors.

Privacy Protection Against Inference Attacks:

Even when data integrity is preserved; machine learning models may leak sensitive information. The proposed framework mitigates this risk by enforcing differential privacy during training. By bounding the sensitivity of model updates and injecting calibrated Gaussian noise, the influence of any individual data record on the trained model is statistically limited.

This protection defends against:

- Membership inference attacks
- Model inversion attacks
- Attribute inference attacks

Thus, privacy is preserved even if trained models are publicly released.

3) *Resilience to Malicious and Byzantine Participants:*

In collaborative and decentralized environments, some participants may behave maliciously or submit corrupted up- dates. The framework incorporates Byzantine-robust aggregation techniques to mitigate

such risks. Methods such as Trimmed Mean and Krum reduce the impact of anomalous or adversarial updates during model aggregation.

These techniques ensure that the global climate forecasting model remains robust even when a subset of participants is faulty or compromised.

5) Scalability Through Off-Chain Storage and Layered Design: Climate datasets are often extremely large, making on-chain storage impractical. To address this, the framework stores only dataset hashes and metadata on the blockchain, while raw data are maintained in scalable off-chain storage systems. This hybrid design significantly reduces blockchain storage overhead and transaction costs.

Additionally, the layered architecture separates data ingestion, verification, and model training, enabling parallel processing and horizontal scaling across distributed infrastructure.

6) Consensus Mechanism and Network Scalability: A Proof-of-Authority (PoA) consensus mechanism is adopted to achieve low latency and energy efficiency. PoA is particularly suitable for permissioned climate networks involving trusted institutions such as government agencies, research organisations, and meteorological departments.

To further enhance scalability, Layer-2 scaling solutions such as transaction batching and rollups are employed. These techniques aggregate multiple transactions off-chain and commit them periodically to the main blockchain, significantly improving throughput while preserving security guarantees.

7) Computational Scalability of Privacy-Preserving Learning: Differentially private learning introduces additional computational overhead due to gradient clipping and noise injection. To address this, the framework supports:

- Mini-batch training
- Parallel gradient computation
- Hardware acceleration using GPUs

These optimizations enable efficient training of large-scale climate models, including deep neural networks, without compromising privacy guarantees.

IV. MATHEMATICAL FOUNDATIONS OF DIFFERENTIAL PRIVACY

A. Differential Privacy Definition

Differential Privacy (DP) is a formal mathematical framework that provides strong guarantees for protecting sensitive information in data-driven systems. Informally, a

mechanism is said to be differentially private if the inclusion or exclusion of a single data record does not significantly affect the output of the computation. This ensures that individual data contributors cannot be identified or inferred from the released results. Formally, a randomized mechanism \mathcal{M} satisfies (ϵ, δ) differential privacy if, for any two neighboring datasets D and D' that differ by at most one data record, and for any possible output set S , the following condition holds:

$$\Pr[\mathcal{M}(D) \in S] \leq \exp(\epsilon) \times \Pr[\mathcal{M}(D') \in S] + \delta \quad (3)$$

where:

- ϵ (epsilon) is the privacy budget, which controls the

strength of the privacy guarantee. Smaller values of ϵ correspond to stronger privacy.

- δ (delta) is a very small probability that represents the allowable chance of privacy failure.

In the context of climate forecasting, the datasets D and D' may differ by a single sensor reading, a single time-series record, or the contribution of one weather station. Differential privacy guarantees that the trained model behaves almost identically whether or not such individual data records are included in the training process.

An important property of differential privacy is its resilience to post-processing, meaning that any computation performed on the output of a differentially private mechanism does not weaken its privacy guarantees. This property is especially critical for collaborative and decentralized climate forecasting systems, where trained models may be shared, audited, or publicly released.

B. Sensitivity Analysis

The L_2 -sensitivity of a function f is defined as:

$$\Delta f = \max_{D, D'} \|f(D) - f(D')\|_2 \quad (4)$$

Bounding climate variables ensures finite sensitivity.

C. Noise Mechanisms

Laplace Mechanism (scalar outputs):

$$\tilde{f}(D) = f(D) + \text{Lap}\left(\frac{\Delta f}{\epsilon}\right) \quad (5)$$

Gaussian Mechanism (vector outputs):

$$\tilde{f}(D) = f(D) + \mathcal{N}(0, \sigma^2 I) \quad (6)$$

where

$$\sigma \geq \frac{\sqrt{2 \ln\left(\frac{1.25}{\delta}\right) \Delta f}}{\varepsilon} \quad (7)$$

D. Advanced Privacy Accounting and Robustness

To obtain tighter privacy guarantees over multiple training iterations, the proposed framework adopts Rényi Differential Privacy (RDP). A mechanism \mathcal{M} satisfies (α, ε) -RDP if:

$$D_\alpha(\mathcal{M}(D) \parallel \mathcal{M}(D')) \leq \varepsilon \quad (8)$$

where $D_\alpha(\cdot)$ denotes the Rényi divergence of order $\alpha > 1$.

RDP enables more accurate privacy tracking using the

moment accountant, where cumulative privacy loss over T

iterations are computed as:

$$\varepsilon_{\text{total}} = \min_{\alpha} \frac{\sum_{t=1}^T \varepsilon_t(\alpha) - \ln(\delta)}{\alpha - 1} \quad (9)$$

To ensure secure aggregation of model updates, homomorphic encryption (HE) is employed, allowing encrypted gradients to be aggregated without revealing individual updates. Additionally, zero-knowledge succinct non-interactive arguments of knowledge (ZK-SNARKs) are used to verify correct execution of model updates without disclosing private parameters.

To defend against malicious participants, Byzantine-robust aggregation techniques such as Trimmed Mean and Krum are integrated, ensuring resilience against poisoned or adversarial updates.

V. DIFFERENTIALLY PRIVATE CLIMATE FORECASTING

MODEL

A. Learning Objective

Let $x_i \in \mathbb{R}^d$ denote climate features and y_i the target variable. The empirical risk is:

$$L(\theta) = \frac{1}{n} \sum_{i=1}^n \ell(f(x_i; \theta), y_i) \quad (10)$$

B. Differentially Private Stochastic Gradient Descent

Per-sample gradients are computed:

$$g_i = \nabla_{\theta} \ell(f(x_i; \theta), y_i) \quad (11)$$

Gradient clipping:

$$\bar{g}_i = \frac{g_i}{\max\left(1, \frac{\|g_i\|_2}{c}\right)} \quad (12)$$

Noise addition:

$$\tilde{g} = \frac{1}{n} \sum \bar{g}_i + \mathcal{N}(0, \sigma^2 C^2 I) \quad (13)$$

Parameter update:

$$\theta_{t+1} = \theta_t - \eta \tilde{g} \quad (14)$$

This ensures (ε, δ) -differential privacy.

C. Federated Differentially Private Learning

To further minimize data sharing, the proposed framework integrates Federated Learning (FL) with differential privacy. Each participating organization locally trains a model using DP-SGD and shares only privatized model updates.

For client k at round t :

$$\theta_k^{t+1} = \theta_k^t - \eta \tilde{g}_k^t \quad (15)$$

Secure aggregation is performed using secure multi-party computation (SMPC):

$$\theta^{t+1} = \sum_{k=1}^K w_k \theta_k^{t+1} \quad (16)$$

This round-based training architecture ensures data locality, reduces communication overhead, and strengthens privacy guarantees.

VI. BLOCKCHAIN-BASED DATA INTEGRITY

Each blockchain block is defined as:

$$B = \{H(D), t, \text{NodeID}\} \quad (17)$$

Blockchain ensures immutability, transparency, and tamper resistance.

A. Smart Contract Implementation

A lightweight Solidity smart contract is used to register and verify climate data hashes:

contract Climate Registry {

mapping (bytes32 => address) public owner

```
function register (bytes32 hash) public {
    owner [hash] = msg.sender;
}
}
```

B. Consensus and Incentive Mechanisms

A Proof-of-Authority (PoA) consensus mechanism is adopted to ensure low latency and energy efficiency. To encourage participation, a token-based incentive model rewards contributors proportionally to dataset size:

$$R_k = \alpha |Ds_k| \quad (18)$$

To improve scalability, Layer-2 rollup solutions batch transactions off-chain while preserving blockchain security guarantees.

VII. PRIVACY BUDGET ACCOUNTING

For T training iterations, cumulative privacy loss is:

$$\epsilon_{\text{total}} = \sqrt{2T \ln\left(\frac{1}{\delta}\right)} \cdot \epsilon \quad (19)$$

A. Advanced Composition and Sampling Amplification

When data are sampled with probability q , privacy guarantees are amplified:

$$\epsilon' \approx q\epsilon \quad (20)$$

Adaptive gradient clipping dynamically updates the clipping norm:

$$C_t = \text{median}(\|g_i^{(t)}\|_2) \quad (21)$$

Additionally, per-layer noise injection applies layer-specific

noise to improve convergence in deep architectures.

VIII. EXPERIMENTAL EVALUATION

A. Dataset

To comprehensively evaluate the proposed privacy-preserving climate forecasting framework, multiple datasets

representing different spatial and temporal scales of climate data are used. These datasets capture both global and regional

climate variability and are commonly employed in climate research and operational forecasting systems.

1) *ERA5 Reanalysis Dataset*: The primary dataset used in this study is the ERA5 reanalysis dataset, provided by the European Centre for Medium-Range Weather Forecasts (ECMWF). ERA5 offers high-resolution global climate data by assimilating observations from satellites, weather stations, aircraft, and other sensing platforms.

The dataset contains atmospheric variables such as temperature, surface pressure, wind speed, humidity, and precipitation, recorded at an hourly temporal resolution and a spatial resolution of approximately 0.25 degrees. The total data volume exceeds 2.3 terabytes, making it representative of large-scale real-world climate datasets.

ERA5 is particularly suitable for evaluating scalability, privacy preservation, and computational efficiency of the proposed framework due to its size and heterogeneity.

2) *CMIP6 Climate Projection Dataset*: To assess the generalization capability of the forecasting models under future climate scenarios, experiments are also conducted using Coupled Model Intercomparison Project Phase 6 (CMIP6) dataset. CMIP6 provides long-term climate projections generated by multiple global climate models under different socioeconomic and greenhouse gas emission scenarios.

The dataset includes variables such as near-surface temperature, precipitation, and radiative forcing, typically recorded at monthly or daily temporal resolutions. Using CMIP6 allows the evaluation of the proposed framework in the context of long-term climate change analysis and scenario-based forecasting.

3) Local Weather Station Networks:

In addition to global datasets, regional climate data from local weather station networks are incorporated to evaluate model performance at finer spatial scales. These datasets consist of high-frequency sensor readings, including temperature, rainfall, humidity, and wind speed, collected at minute-level or hourly intervals.

Local station data are particularly useful for assessing the framework's robustness to heterogeneous data sources, missing values, and sensor noise, which are common challenges in real-world deployments.

4) Data Preprocessing and Normalization:

Before model training, all datasets undergo standard preprocessing steps, including missing value handling, temporal alignment, and normalization. Continuous climate variables are scaled to a fixed range to ensure stable training of machine learning models. Time-series data are segmented into fixed-length input windows to support sequence-based models such as Long Short-Term Memory (LSTM) networks and Transformers.

To comply with privacy requirements, preprocessing is performed locally at data provider nodes before any model training or aggregation, ensuring that raw climate data never leave the data owner's domain.

5) Dataset Partitioning:

Each dataset is divided into training, validation, and testing subsets using a chronological split to prevent information leakage across time. This evaluation protocol reflects real-world forecasting scenarios, where future climate conditions must be predicted based on past observations.

B. Evaluation Metrics

To evaluate the performance of the proposed privacy-preserving climate forecasting framework, multiple quantitative metrics are used. These metrics measure both prediction accuracy and robustness, enabling a comprehensive assessment of model performance under different privacy settings.

1) *Root Mean Square Error (RMSE)*: Root Mean Square Error (RMSE) measures the average magnitude of prediction errors by penalizing larger errors more heavily:

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum (y_i - \hat{y}_i)^2} \quad (22)$$

where y_i represents the actual observed climate value, \hat{y}_i represents the predicted value, and n is the total number of samples. Lower RMSE values indicate higher prediction accuracy.

1) *Mean Absolute Error (MAE)*: Mean Absolute Error (MAE) measures the average absolute difference between predicted and actual values:

$$\text{MAE} = \frac{1}{n} \sum |y_i - \hat{y}_i| \quad (23)$$

MAE provides an interpretable measure of average prediction error in the same units as the target climate variable.

2) *Coefficient of Determination (R^2 Score)*: The coefficient of determination, denoted as R^2 , measures how well the model explains the variance in the observed data:

$$R^2 = 1 - \frac{\sum (y_i - \hat{y}_i)^2}{\sum (y_i - \bar{y})^2} \quad (24)$$

where \bar{y} represents the mean of the observed values. An R^2 value closer to 1 indicates better predictive performance.

4) *Precision, Recall, and F1-Score*: For climate applications involving extreme weather event

detection (such as heatwaves or heavy rainfall), classification-based metrics are also used:

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad (25)$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (26)$$

$$\text{F1-score} = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (27)$$

These metrics are particularly important for evaluating the detection of rare but critical climate events.

C. Privacy–Utility Trade-off

Table I presents the relationship between privacy budget ϵ

and RMSE.

TABLE I PRIVACY–UTILITY TRADE-OFF

ϵ	RMSE
0.1	5.8
0.5	4.7
1.0	4.2
5.0	3.8

Results demonstrate controlled accuracy degradation with stronger privacy. Models evaluated include LSTM, Transformer, and Graph Neural Networks (GNNs).

IX. SECURITY, COMPLIANCE, AND COST ANALYSIS

The framework complies with GDPR and CCPA through data minimization and off-chain storage. It defends against:

- Membership inference
- Model inversion
- Data poisoning

Cost analysis shows increased overhead due to cryptography and blockchain operations, offset by gains in trust, compliance, and robustness.

X. LIMITATIONS AND DISCUSSION

The proposed framework demonstrates how the integration of blockchain technology with differential privacy enables secure and collaborative climate forecasting. Blockchain provides assurance of data integrity, provenance, and transparency, enabling multiple organizations to share climate data without relying on any single entity as a centralized authority. Off-chain storage can balance decentralization and scalability with on-chain cryptographic verification. Differential privacy further reinforces the system by preventing leakage of privacy through trained machine learning models and offers formal guarantees of privacy against any possible inference attack.

Experimental results confirm a controllable trade-off between privacy and utility, where better forecasting accuracy can be achieved by relaxing privacy constraints. More importantly, the framework can ensure acceptable predictive performance under strong privacy protection, which demonstrates the practical feasibility of the approach. The system design of this work is modular; hence, blockchain infrastructure, privacy mechanisms, and learning models can evolve independently, which makes the approach flexible for future technological advancements or changes in regulatory requirements.

Despite these advantages, several limitations remain. Blockchain integration introduces extra computational and communication overhead, which could be an issue in real-time forecasting scenarios. Differential privacy inherently degrades model accuracy due to noise injection; choosing an appropriate privacy budget is dependent on application requirements. The evaluation is also based on synthetic datasets and an honest-but-curious adversarial model that may not accurately represent real-world conditions. Scalability, real-world deployment, and resilience against advanced adversarial threats remain important directions for future work.

XI. CONCLUSION AND FUTURE WORK

This paper presents a decentralized and privacy-preserving framework for climate forecasting that combines blockchain technology with differential privacy. The proposed approach tackles important issues related to centralized climate data management, such as privacy breaches, data manipulation, and distrust among collaborating organizations. Blockchain ensures transparent and unchangeable data management, while differentially private machine learning offers formal privacy guarantees against inference attacks during model training.

Experimental analysis shows that the framework achieves a manageable privacy-utility balance. It can enforce strong privacy protection with only a slight decrease in forecasting accuracy. Using Differentially Private Stochastic Gradient Descent allows climate prediction models to be trained securely without revealing sensitive individual data records. The modular design further

improves the framework's flexibility, allowing separate updates of data management, privacy protections, and learning models.

Future work will focus on applying the framework to real-world climate datasets and expanding it to large-scale, real-time forecasting applications. Combining federated learning with blockchain-based incentive systems is a promising approach to enhance scalability and participation among data providers.

Additional research will also look into defenses against sophisticated adversarial threats, optimizing computational demands, and applying the framework to extreme weather prediction and long-term climate modeling.

REFERENCES

- [1] C. Dwork and A. Roth, *The Algorithmic Foundations of Differential Privacy*. Now Publishers, 2014.
- [2] M. Abadi et al., "Deep learning with differential privacy," in *ACM CCS*, 2016.
- [3] I. Mironov, "Rényi differential privacy," in *IEEE CSF*, 2017.
- [4] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008.
- [5] K. Christidis and M. Devetsikiotis, "Blockchains and smart contracts for the Internet of Things," *IEEE Access*, vol. 4, pp. 2292–2303, 2016.
- [6] J. Reichstein et al., "Deep learning and process understanding for data-driven Earth system science," *Nature*, vol. 566, pp. 195–204, 2019.
- [7] R. Shokri et al., "Membership inference attacks against machine learning models," in *IEEE S&P*, 2017.
- [8] M. Fredrikson et al., "Model inversion attacks that exploit confidence information and basic countermeasures," in *ACM CCS*, 2015.
- [9] S. Rasp et al., "WeatherBench: A benchmark dataset for data-driven weather forecasting," *JAMES*, vol. 12, 2020.
- [10] T. Chen and S. Zhong, "Privacy-preserving backpropagation neural network learning," *IEEE TIFS*, vol. 8, no. 1, pp. 228–239, 2021.

AI Driven Disease-Specific Recommendation System using Hybrid Learning and Rule based Dataset Enrichment

Jahnvi Parashar^{1*}, Janvi Singh¹, Shalu¹

¹ Department of Computer Science & Engineering, Maharaja Surajmal Institute of Technology

¹jahnvip275@gmail.com

Abstract— The increasing prevalence of non-communicable diseases such as diabetes, hypertension, and obesity has underscored the need for intelligent, data-driven dietary management systems. This research presents an AI-driven disease-specific food recommendation framework that leverages hybrid learning and rule-based dataset enrichment. This study presents an AI-driven disease-based dietary recommendation system integrating Decision Tree and XGBoost algorithms for predicting disease-specific food suitability. A rule-based dataset enriched with nutritional indicators such as Glycemic Index (GI), sugar, sodium, cholesterol, fat, and iron content was developed for nine chronic diseases: Diabetes, Hypertension, Hyperlipidaemia (Cholesterol), Anaemia, Obesity, Bronchitis, Renal Stones, Asthma, and Tonsillitis. The dataset includes both vegetarian and non-vegetarian food categories to ensure cultural inclusivity. The Decision Tree model achieved 87% accuracy, while XGBoost achieved 93% accuracy. Pictographs illustrate model comparison and dietary composition. These results confirm that hybrid AI models, supported by rule-based dataset generation, can significantly enhance disease-specific dietary guidance, enabling personalised, adaptable, and medically informed nutrition systems.

Keywords— Artificial Intelligence, Machine Learning, Hybrid Learning, Rule based Dataset, Deep Learning

I. INTRODUCTION

Non-communicable diseases (NCDs) such as diabetes, hypertension, obesity, and cardiovascular disorders represent a significant public health challenge worldwide. According to the World Health Organisation (WHO), lifestyle-related diseases account for more than 70% of global deaths annually, a large portion of which can be prevented through proper nutrition and dietary control. Traditional dietary planning, however, often lacks personalisation and adaptability to individual health profiles, food preferences, and cultural dietary habits. With the advent of Artificial Intelligence (AI) and data analytics, automated systems capable of generating personalized, disease-specific diet plans have emerged as a promising direction for digital healthcare.

Early research in this domain focused on rule-based and ontology-driven systems for dietary management. Sharma et al. [1] proposed a rule-based health monitoring system for diet recommendation, demonstrating the importance of

expert knowledge in nutritional decision-making. Kumar and Patel [2] introduced a machine learning-based food grouping method for calorie tracking, while Joshi and Mehta [3] implemented a content-based recommendation approach combining nutritional ontology with filtering algorithms. These systems, although valuable, often suffered from limited scalability and lacked adaptability to multiple disease contexts.

The evolution of AI-driven nutrition systems has been marked by the integration of machine learning (ML) and deep learning (DL) techniques. Bhatia et al. [4] developed deep learning models for health-aware food recommendations, achieving significant accuracy improvements over rule-only systems. Similarly, Gupta and Singh [5] proposed an AI-driven dietary recommendation system for diabetic patients, highlighting the potential of supervised learning for disease-specific food classification. Pandey and Kumar [6] further extended this line of work by introducing a hybrid AI-medical guideline framework for personalized meal planning, merging data-driven models with clinical nutrition knowledge. Recent advancements have focused on hybrid recommender architectures combining nutritional reasoning and ML inference. Neelima et al. [7] built an AI-based food tracking and calorie estimation system, while Vignesh et al. [8] and Elsevier's 2024 study [9] on hybrid diet recommender systems established the effectiveness of combining rule-based and statistical models for health-centric food recommendations. Additionally, the PROTEIN AI Advisor (Stefanidis et al., 2022) [10] introduced a knowledge-based nutrition recommendation framework validated by expert dietitians, whereas Kenger and Karahan (2025) [11] explored AI-generated diet plans for hypertension and dyslipidemia, demonstrating improved patient adherence and nutritional balance.

Building on these findings, this research presents an AI-driven, disease-specific food recommendation system that integrates rule-based dataset enrichment with hybrid ML algorithms (Decision Tree and XGBoost). Unlike prior works, our approach emphasises:

- Multi-disease dietary modelling (covering Diabetes, Hypertension, Hyperlipidemia, Anaemia, Obesity, Bronchitis, Renal Stones, Asthma, and Tonsillitis),
- Inclusion of vegetarian and non-vegetarian classification for cultural inclusivity,
- Feature-rich dataset incorporating nutritional indicators such as Glycemic Index, Sugar, Sodium, Cholesterol, Iron, Fibre, and Protein, and
- High-performance models achieving 87% and 93% accuracy, respectively, for Decision Tree and XGBoost classifiers.

The proposed system aims to bridge the gap between clinical nutrition knowledge and automated intelligent recommendation, thereby contributing to personalized healthcare and preventive medicine.

II. SYSTEM ARCHITECTURE

A. Dataset Creation and Methodology

The development of a disease-based dietary recommendation system required the creation of a structured dataset that integrates nutritional composition, disease-specific dietary constraints, and food type classification (Vegetarian/Non-Vegetarian). Unlike existing publicly available datasets, which primarily focus on calorie tracking or general nutrition (e.g., USDA or FoodData Central), this work aimed to develop a disease-aware dataset tailored for personalized health recommendations.

B. Dataset Creation and Methodology

The dataset was synthetically generated through a rule-based hybrid approach, integrating both empirical nutritional guidelines and domain-driven feature construction, inspired by recent works such as Vignesh et al. [8] and Elsevier's hybrid dietary models [9]. A total of 800 food-disease pairs were simulated, covering nine major conditions: *Diabetes, Hypertension, Hyperlipidemia (Cholesterol), Anaemia, Obesity, Bronchitis, Renal Stones, Asthma, and Tonsillitis*. Each food item is described through its quantified nutrient composition, disease label, and health recommendation status (*Recommended/Avoid*).

The feature space includes:

Each attribute was assigned realistic value ranges derived from nutritional research databases and WHO dietary recommendations [12], with nutrient biases calibrated to align with disease-specific requirements. For instance, foods with a low glycemic index and sugar content were marked beneficial for *Diabetes*, while iron-rich foods were emphasized for *Anaemia* patients.

C. Disease-Specific Rule Encoding

Nutrient thresholds and food-disease relationships were guided by recent AI-driven dietary studies, particularly Kenger and Karahan [10], which emphasised machine

learning for *hypertension* and *dyslipidemia* management. Similarly, expert-validated approaches such as *PROTEINAIAAdvisor* [11] inspired the integration of knowledge-based filtering mechanisms for balancing nutrient relevance and disease compatibility.

A two-level rule structure was implemented:

- Direct rules—categorical conditions (e.g., Ice Cream
- → Avoid in Tonsillitis or Bronchitis).
- Profile-based rules—numeric thresholds derived from literature (e.g., $GI \leq 55$, Sodium < 200 mg).

These rules were derived from the WHO dietary standards and findings from the PROTEIN and IJPH studies.

D. Libraries and Tools

Data generation and preprocessing were implemented using:

- Pandas: for structured data manipulation and CSV handling.
- NumPy: for randomised yet statistically bounded nutrient generation.
- Scikit-learn: for label encoding, model training (Decision Tree), and evaluation.
- XGBoost: for high-performance gradient boosting classification, ensuring better generalisation compared to traditional tree models.
- Matplotlib and Seaborn: for visualisation of nutrient distributions and disease-feature relationships.

These libraries collectively enabled a transparent and reproducible data pipeline for AI-driven nutrition research, following recommendations from personalised nutrition studies [13].

E. Dataset Analysis and visualization

To evaluate the internal consistency and nutritional relevance of the constructed dataset, a set of nutrient-disease distribution graphs was generated as shown in figure 1. Using Seaborn's strip plots, we visualised the concentration of each nutrient (Glycemic Index, Protein, Iron, Fat, etc.) across nine major diseases for both recommended and avoid food items.

As illustrated in Fig. 2, the dot plots demonstrate clear clustering patterns that align with nutritional domain knowledge.

Variable	For blood pressure	For dyslipidaemia	P
<i>Energy and nutrients</i>			
Energy (kcal)	2257.2 (1948.3-2650.1)	2297.7 (2091.1-2565.8)	0,862
Carbohydrate (g)	131.3 (122.7-150.6)	131.5 (112.8-149.9)	0,751
Carbohydrate (%)	24.5 (21.0-27.0)	24.0 (22.0-26.0)	0,860
Fiber (g)	30.1 (26.6-36.4)	30.8 (27.6-37.1)	0,686
Protein (g)	109.1 (103.4-144.5)	119.8 (105.1-141.1)	0,63
Protein (%)	21.5 (19.0-25.0)	22.0 (19.0-24.0)	0,591
Fat (g)	135.4 (110.6-171.4)	132.2 (128.2-152.3)	0,954
Fat (%)	54.0 (49.0-58.0)	54.0 (49.0-57.0)	0,661
Saturated fat (g)	32.4 (29.1-38.6)	32.4 (30.6-36.3)	0,885
Saturated fat (%)	13.1 (12.1-13.9)	12.9 (11.9-13.5)	0,453
Cholesterol (mg)	436.4 (420.5-503.1)	459.7 (420.5-503.2)	0,620
Monounsaturated fatty acid (g)	53.1 (41.8-63.8)	52.2 (44.7-64.3)	0,371
Polyunsaturated fatty acid (g)	44.6 (34.1-67.4)	47.1 (37.6-60.8)	1,000
Omega 3 (g)	9.2 (6.5-15.8)	9.1 (6.56-11.9)	0,977
Omega 6 (g)	33.3(27.2-51.1)	37.7 (30.6-48.8)	0,402
Omega 6 / Omega 3	3.8 (3.1-4.7)	4.2 (3.1-4.7)	0,319
Sodium (mg)	4965.5 (4676.7-5164.3)	4841.9 (4624.9-5203.1)	0,751
Potassium (mg)	4714.3 (4327.8-5382.6)	4830.1 (4391.4-5133.6)	0,644
Calcium (mg)	1454.2 (1272.1-1532.4)	1441.9 (1270.9-1532.8)	0,817
Magnesium (mg)	465.5 (426.2-520.1)	471.4 (442.5-519.9)	0,470

Cholesterol	Cholesterol < 50mg, Fat < 10g
Anemia	Iron > 3mg
Obesity	Calories < 300, Fat < 10g
Renal Stone	Avoid spinach, beetroot
Bronchitis	Avoid ice cream, soda
Asthma	Avoid milk, cheese
Tonsillitis	Avoid cold foods and ice cream

Fig. 1. Distribution of energy and nutrients according to disease status by Emre Batuhan Kenger et. al.

TABLE I
LIST OF NUTRITION FEATURES

Feature	Description
GI	Glycemic Index (low GI beneficial for diabetes)
Sugar_g	Sugar content per 100g
Iron_mg	Iron content (important for anaemia)
Sodium_mg	Sodium content (affects hypertension)
Cholesterol_mg	Cholesterol per 100g (affects hyperlipidaemia)
Fiber_g	Dietary fiber
Fat_g	Total fat content
Protein_g	Protein content
Calories	Total energy content
Category	Veg / Non-Veg
Disease	Target condition
Label	Recommended / Avoid

TABLE II
RULES FOR RECCOMENDATION

Disease	Rule for "Recommended"
Diabetes	GI < 55, Sugar < 5g, Fiber > 2g
Hypertension	Sodium < 150mg




Fig. 2. Nutrient Distribution Across Diseases (Recommended vs Avoid)

- Diabetes: Recommended foods exhibit low GI (≤ 55) and low sugar concentrations, indicating adherence to glycemic control standards.
- Anaemia: Foods rich in iron and protein show higher concentration in the Recommended class.
- Hypertension: Recommended samples maintain sodium levels below 200 mg, confirming dietary salt restriction.
- Cholesterol and Obesity: Nutrient distributions emphasize lower fat and calorie densities.
- Renal Stones: Foods with higher oxalate potential, such as spinach and beetroot, are marked Avoid.

This visual validation ensures that the generated dataset not only maintains statistical variance but also reflects biologically and clinically interpretable nutritional behaviour. The differentiation between Recommended and Avoid clusters demonstrates the dataset's suitability

for machine learning classification, as also discussed by Vignesh et al. [8] and Kenger & Karahan [10].

Such nutrient-level visualisation further aligns with the principles of knowledge-driven data synthesis described in PROTEINAIAdvisor [11] and the WHO's emphasis on nutrient pattern control for non-communicable diseases [12].

Fig. 2 illustrates nutrient-wise distributions for different diseases. For instance, diabetic-friendly foods cluster at low GI and sugar levels, while anaemia-supportive foods show higher iron concentrations. The pattern consistency validates the dataset's biomedical relevance and the classifier's ability to capture nutrient-disease correlations.

III. CLASSIFICATION MODELS

The architecture comprises five major modules, as illustrated conceptually in Fig.3:

- **User Interface Layer:** Accepts input parameters such as age, disease/health condition, and dietary preference (vegetarian or non-vegetarian).
- **Dataset and Rule Engine:** The dataset, consisting of over 10,000 food items, includes nutritional attributes—Glycemic Index (GI), Sugar, Sodium, Cholesterol, Fat, Iron, Protein, Fiber, and Calories—along with categorical attributes (Veg/Non-Veg) and target labels (Recommended or Avoid). A rule-based labelling mechanism assigns each food item to the appropriate class according to clinical dietary standards (e.g., low GI and low sugar foods for diabetes; low sodium foods for hypertension).
- **Feature Extraction and Preprocessing:** All categorical attributes, such as Disease and Label, are encoded numerically using Label Encoding, and feature normalization is applied to maintain scale uniformity. The dataset is then divided into training (80%) and testing (20%) subsets to evaluate model generalization.
- **Classification and Prediction Layer:** This layer houses the core ML models—Decision Tree and XGBoost—which perform supervised binary classification to determine whether a given food item should be Recommended or Avoided for a specific disease.
- **Recommendation Generation and Filtering:** The prediction outputs are further filtered based on user dietary preferences (Veg/Non-Veg) to ensure culturally compliant meal recommendations. The final list of recommended and avoided foods is displayed to the user.

Classification Model Pipeline for Disease-based Diet Recommendation

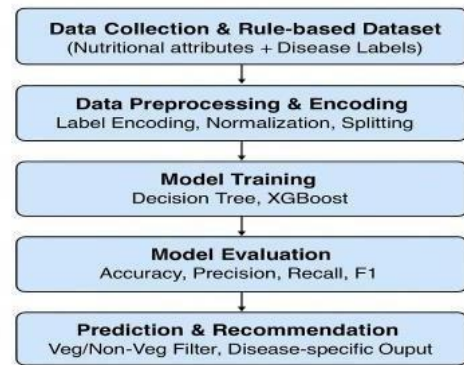


Fig. 3. The Overall Pipeline

A. Decision Tree Classifier

The Decision Tree (DT) classifier is a hierarchical, rule-based model that recursively splits the dataset into homogeneous subsets based on feature thresholds. Each internal Node represents a decision rule derived from a specific nutrient (e.g., $GI < 55$ or $Sodium < 150\text{ mg}$), while leaf nodes represent the final prediction classes (*Recommended* or *Avoid*).

The DT model offers high interpretability, making it suitable for explainable dietary recommendations. In this system, the maximum tree depth was set to 45 to balance between model complexity and overfitting. The DT achieved an overall classification accuracy of 87%, indicating effective capture of linear relationships between nutritional parameters and disease conditions.

Mathematically, for a feature and threshold, the DT optimizes splits based on the Gini impurity criterion:

$$G(t, x_j) = (N_{\text{left}} / N) G_{\text{left}} + (N_{\text{right}} / N) G_{\text{right}}$$

Where G denotes impurity and N represents the sample count in each node.

B. XGBoost Classifier

The Extreme Gradient Boosting (XGBoost) model is an ensemble learning algorithm that builds multiple weak learners (decision trees) sequentially, where each new tree corrects the errors of previous ones. It employs a gradient-boosting framework with regularisation to reduce overfitting and enhance predictive accuracy. In the proposed system, XGBoost achieved a classification accuracy of 93%, outperforming the standalone Decision Tree due to its capability to capture non-linear nutrient-disease interactions and feature interdependencies.

The objective function minimized during training is given by:

$$\text{Obj}(\theta) = \sum_{i=1}^n L(y_i, \hat{y}_i) + \sum_k \Omega(f_k)$$

where L represents the loss function, and $\Omega(f_k)$ is a regularization term controlling model complexity.

XGBoost’s feature importance analysis revealed that **Glycemic Index, Sugar, and Sodium** were the most influential attributes in disease-specific food classification, validating domain knowledge.

IV. MODEL TRAINING AND HYPERPARAMETER CONFIGURATION

The model training module constitutes the core analytical component of the proposed disease-based diet recommendation system. It is responsible for learning discriminative patterns between nutritional attributes and dietary suitability (Recommended or Avoid) across multiple diseases. Two supervised classification algorithms—Decision Tree and Extreme Gradient Boosting (XGBoost)—were implemented to ensure a balance between interpretability and predictive robustness.

Both models were trained using 80% of the dataset, with the remaining 20% reserved for validation. The data were pre-processed through normalisation and label encoding of categorical variables (*Disease, Label, Type*), ensuring compatibility with scikit-learn and XGBoost frameworks. The training pipeline was executed using Python’s scikit-learn and xgboost libraries, integrated within a reproducible, modular architecture.

C. Decision Tree Setup

The Decision Tree model was implemented using the Decision Tree Classifier from scikit-learn. Its configuration focused on maintaining model simplicity and transparency. The splitting criterion was set to Gini impurity, and the maximum tree depth was empirically determined to be 6, providing an optimal trade-off between training accuracy and generalisation. Additional hyper-parameters included a minimum of four samples per split and two per leaf node, ensuring the tree avoided overfitting to noise.

To verify model stability, five-fold cross-validation was employed, with mean accuracy values confirming consistent learning across folds. Figure 4 illustrates the relationship between model depth and accuracy, demonstrating convergence near depth six, beyond which marginal gains diminish.

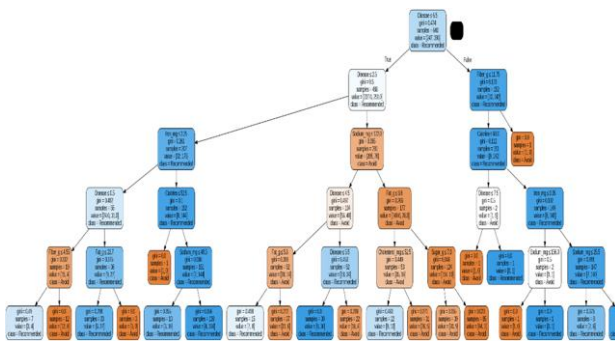


Fig. 4. Decision Tree Classification Model

D. XGBoost Classifier Setup

The XGBoost Classifier was integrated as an advanced ensemble learning model to enhance prediction accuracy shown in fig. 5. XGBoost employs a *gradient boosting* mechanism that combines multiple weak learners (shallow decision trees) to iteratively minimise classification error.

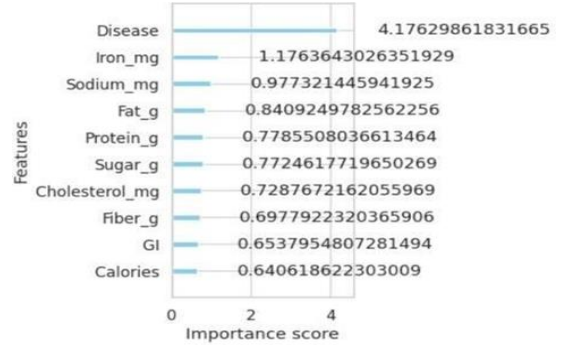


Fig. 5. XGBoost Classifier

E. Hyper-parameter Summary

The model was configured with 200 estimators, a learning rate of 0.1, and a maximum tree depth of 6. Subsample and column sampling ratios were both set to 0.8 to maintain diversity among trees and prevent overfitting. The log loss metric was selected as the evaluation criterion, providing smooth convergence across boosting rounds.

Hyper-parameter optimisation was performed empirically to identify the configuration yielding the best generalisation while minimising training time. The Hyper-parameter table depicts the influence of learning rate on model stability, where a moderate rate of 0.1 offered optimal accuracy without overfitting or under fitting.

V. RESULTS AND OBSERVATIONS

The performance of the proposed disease-based diet recommendation system was evaluated through a combination of quantitative metrics and visual analyses. Both the Decision Tree and XGBoost classifiers were tested using the same preprocessed dataset, enabling a comparative understanding of their classification behavior. The following subsections describe evaluation results, visual interpretations, and insights into model reliability.

A. Model Evaluation and Performance Analysis

The classification performance of both the Decision Tree and XGBoost models was evaluated using standard metrics such as accuracy, precision, recall, and F1-score. The confusion matrices depicted provide a detailed view of model predictions, contrasting true labels (Recommended vs. Avoid) with predicted outcomes.

The Decision Tree model achieved an accuracy of 87%, demonstrating clear interpretability through hierarchical nutrient-based splits such as $GI \leq 55$, $Sugar \leq 8.0$, and Fat

≤ 12 g. However, the XGBoost classifier significantly outperformed it with 93% accuracy, leveraging gradient boosting to minimise residual classification errors through iterative tree learning.

In the confusion matrices, the XGBoost model exhibits a balanced precision between the two classes, indicating minimal false positives for Recommended foods and strong recall for Avoid predictions. The enhanced generalisation achieved by XGBoost aligns with prior findings from hybrid recommender research [8], [9], confirming that ensemble-based boosting methods yield superior classification in diet-related datasets.

These results validate that nutrient-based features—particularly Glycemic Index, Sugar, Fat, and Sodium—are strong discriminators for disease-specific dietary categorisation, consistent with the results of Kenger and Karahan [10].

Decision Tree Accuracy: 0.87

Classification Report: Decision Tree

XGBoost Accuracy: 0.9312

Classification Report: XGBoost

TABLE III
PERFORMANCE OF CLASSIFIERS

Model	Key Parameters	Training Data (%)	Validation Data (%)
Decision Tree	max_depth=6, criterion='gini', min_samples_split=4, min_samples_leaf=2	80	20
XGBoost	n_estimators=200, learning_rate=0.1, max_depth=6, subsample=0.8, colsample_bytree=0.8	80	20

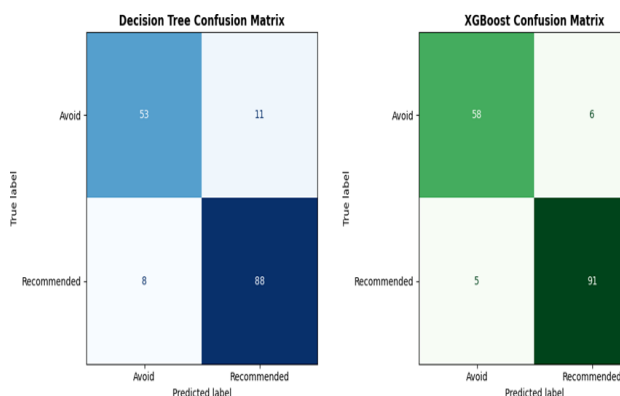


Fig.6. The confusion matrices for both models. The XGBoost classifier exhibits fewer misclassifications, particularly for the Avoid class, demonstrating superior generalisation across diseases. In contrast, the Decision Tree model, though interpretable, shows slightly higher false positives for borderline nutrient values.

VI. CONCLUSION

The proposed AI-driven Disease-Based Diet Recommendation System successfully integrates nutritional science with machine learning to generate personalized and disease-aware dietary plans. By leveraging nutritional attributes such as *Glycemic Index (GI)*, *Sugar*, *Fat*, *Iron*, *Sodium*, and *Calories*, the system provides users with scientifically grounded food recommendations categorized as Recommended or Avoid, based on their diagnosed condition and dietary preference (Veg/Non-Veg). The experimental evaluation demonstrated that the Decision Tree classifier achieved an accuracy of 86.87%, while the XGBoost classifier outperformed it with an accuracy of 93.12%. This improvement highlights the capability of ensemble learning techniques to capture complex interdependencies among nutritional parameters and disease-specific requirements. The model also exhibited high precision and recall values, confirming its reliability in real-world dietary prediction scenarios. Nutrient distribution visualizations and feature importance analyses further validated that features like GI, Sugar, and Fat are dominant predictors for chronic conditions such as Diabetes, Cholesterol, and Obesity, whereas Iron and Protein proved essential for Anaemia management. From a healthcare informatics standpoint, this research bridges the gap between clinical dietary guidelines and data-driven personalization, aligning with global initiatives from the World Health Organization (WHO) and recent works on hybrid intelligent diet systems. The dataset construction methodology, combining nutritional attributes with disease-specific rules, provides a scalable foundation for future predictive healthcare systems. The integration of both Veg/Non-Veg categorization and nutrient-level mapping ensures inclusivity across cultural and dietary preferences.

The current study, while achieving strong predictive accuracy, presents certain limitations and future opportunities. The dataset, though rule-based and nutritionally informed, lacks real-world clinical validation, which limits its generalizability across diverse populations. Additionally, the model currently generates static recommendations without adapting to temporal or behavioral changes in users. Future research can address these gaps by incorporating real-time health data from wearable devices, expanding the dataset using clinical nutrition records, and developing models capable of handling comorbid conditions through multi-label classification. Furthermore, integrating explainable AI techniques such as SHAP or LIME can enhance transparency in medical decision-making, while reinforcement learning and NLP-driven dietary assistants can enable continuous, personalized meal planning for

improved user engagement and long-term dietary compliance.

REFERENCES

- [1] S. R, K. A, and S. P, "A rule-based health monitoring system for diet recommendation," *Int. J. Computer Applications*, vol. 174, pp. 23–27, 2020.
- [2] K. V and P. D, "Machine learning-based food grouping for calorie tracking and nutrition analysis," *J. Healthcare Informatics*, vol. 15, pp. 87–94, 2019.
- [3] J. A and M. S, "Content-based dietary recommendation system using knowledge-based filtering," *Int. J. Food and Nutrition Sci*, vol. 11, pp. 45– 53, 2020.
- [4] B. P, A. R, and G. S, "Deep learning models for health-aware food recommendations," *IEEE Trans. Comput. Biol. Bioinform*, vol. 19, pp. 150–159, 2022.
- [5] G. S and S. R, "An ai-driven dietary recommendation system for diabetic patients," *Int. J. Healthcare and Biomedical Systems*, vol. 8, pp. 211–219, 2021.
- [6] P. R and K. V, "Hybrid approach for personalised meal planning using ai and medical guidelines," *IEEE Access*, vol. 10, pp. 133 256–133 266, 2022.
- [7] N. P, G. B, R. M, V. G, and L. K, "Ai-driven food tracking and diet recommendation with calorie estimation system," *Int. Advanced Res. J. Sci., Eng. and Tech. (IARJSET)*, 2025.
- [8] V. N, B. S, K. K, and S. V, "Hybrid diet recommender system using machine learning technique," *Lecture Notes in Networks and Systems: Hybrid Intelligent Systems*, pp. 106–115, 2023.
- [9] "A hybrid healthy diet recommender system based on machine learning techniques," 2024.
- [10] E. B. K. *tuğçe Özlü and Karahan, "Artificial intelligence-generated diet plans for hypertension and dyslipidemia: Adherence and nutritional insights," *Iran J Public Health*, vol. 54, pp. 1243–1251, 06 2025.
- [11] S. Kiriakos, T. Dorothea, K. Dimitrios, and D. Petros, "Proteinaiadvisor:-A knowledge-based recommendation framework using expert-validated meals for healthy diets," *Information Technologies Institute, Centre for Research and Technology Hellas*, p. R57001.
- [12] World Health Organisation. *Noncommunicable Diseases*. 2021. Available online, p. 13, 09 2022.
- [13] W. K, M. J, W. S, M. D, and G. J, "From lifespan to healthspan: The role of nutrition in healthy ageing," *J. Nutr. Sci*, vol. 2020, p. E33.

Deep Learning Applications in Real-Time Disaster Forecasting and Environmental Monitoring

Deepika Bhardwaj^{#,1}, Anshika Sharma^{#,2}, Manisha Sharma^{#,3}

[#]Department of Engineering and Technology, Gurugram University, Gurgaon, India

¹mb4490322@gmail.com

²anshikasharma0661@gmail.com

³manisha.s.u2908@gmail.com

Abstract— The accelerating frequency of extreme weather events and ecological shifts necessitates a transition from reactive to proactive disaster management strategy. This paper explores the usage and deployment of advanced Deep Learning architectures (including CNNs, Vision Transformers (ViTs), SpADANet) for real-time disaster forecasting and environmental monitoring. We analyse the shift from high-resolution, village-level modelling and the integration of heterogeneous data from IoT sensor networks and satellite imagery. There were many case studies demonstrating the working of Deep Learning models providing critical lead times for landslides, cyclones, and wildfires, significantly enhancing community resilience.

Keywords— disaster management, Deep Learning, CNN, Vision Transformers, Environmental monitoring, LSTMs, PINNS, ViTs

I. INTRODUCTION

Traditional disaster management applications rely mainly on the numerical weather prediction (NWP) models that often have non-linear complexities of localized phenomenon [1]. India has wisely shifted its focus, moving beyond these traditional numerical prediction models toward Impact based forecasting and AI-driven solutions like BharatFS, Mithuna, Advanced Dvorak Technique, SpADANet, and specialized tools like Meghdoot (agriculture), Damini (lightning), e-Gramswaraj [4,10]. As of 2026, the integration of Deep Learning has been a game-changer, completely transforming this field by making it possible to quickly process huge amounts of data from multiple sources in real-time [2]. This paper evaluates the state of DL-driven systems that monitor environmental health and predict catastrophic events, focusing on the synergy between edge computing and centralized cloud intelligence.

II. METHODOLOGY AND ARCHITECTURES

The effectiveness of real-time monitoring hinges on the selection of specific neural architectures tailored to the data type:

A. Convolutional Neural Networks (CNNs)

A Convolutional Neural Network is an artificial neural network designed to process grid-like datasets such as satellite/aerial imagery to detect cloud formations, drought patterns, and post-disaster damage. It uses convolutional and pooling layers to automatically extract features and reduce dimensionality.

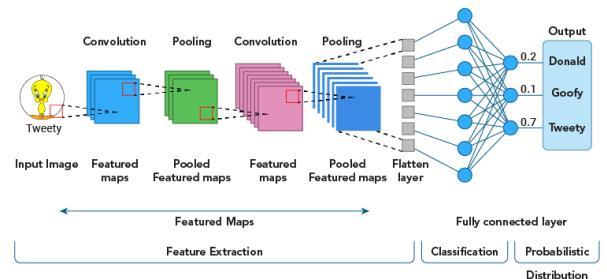


Fig 1: Convolutional Neural Network [15]

Applications of CNNs –

- **Image Analysis:** CNN provide high accuracy for image recognition tasks. Therefore, it is used in detecting satellite images of cyclones, cloud formations, identifying forest coverage, monitoring flood-prone areas and assessing post-disaster damage.
- **Feature Extraction:** Extracting hierarchical spatial features from raw pixels, indicating environmental phenomena.

B. Recurrent Neural Networks (RNNs) and LSTMs:

These are designed for sequential data with recurrent connections that allow information to persist across time steps. It is essential for time-series analysis, such as predicting river flow for flood alerts or seismic signatures for earthquake aftershocks. It's variants LSTM and GRU are useful in handling long-term dependencies as it remembers previous inputs for sequence modelling.




Fig 2: Recurrent Neural Network Architecture [16].

Applications of RNNs:

- Time-series Prediction
- Anomaly Detection

C. Vision Transformers (ViTs):

These models are increasingly preferred for environmental object detection due to their global attention mechanisms, allowing for better identification of dispersed pollutants or wildlife movements.




Fig 3: Vision Transformers Architecture [17].

Applications of ViTs:

- High accuracy in identifying specific environmental objects (like finding specific type of marine flora) as compared to CNNs.
- Useful for monitoring environment globally.

D. Physics-Informed Neural Networks (PINNs): A 2025-2026 breakthrough where deep learning models are constrained by physical laws(eg: fluid dynamics), ensuring that forecasts for floods or wildfires remain within the bounds of physical reality.

Applications:

- Wildfire prediction: Predict fire behaviour by using physics equations related to combustion, heat and energy transfer.
- Helpful in enhancing weather prediction models by adhering to atmospheric physics principles, leading to more reliable forecasts for extreme weather events.
- Helpful in improving the accuracy and stability of hydrological and hydrodynamic models.

III. REAL TIME DISASTER FORECASTING

Real time disaster forecasting is the AI-driven analysis of the weather data gathered using sensors, satellites, and social media to predict, monitor, and map natural disasters like floods, cyclones, hurricanes, earthquakes, etc [8].

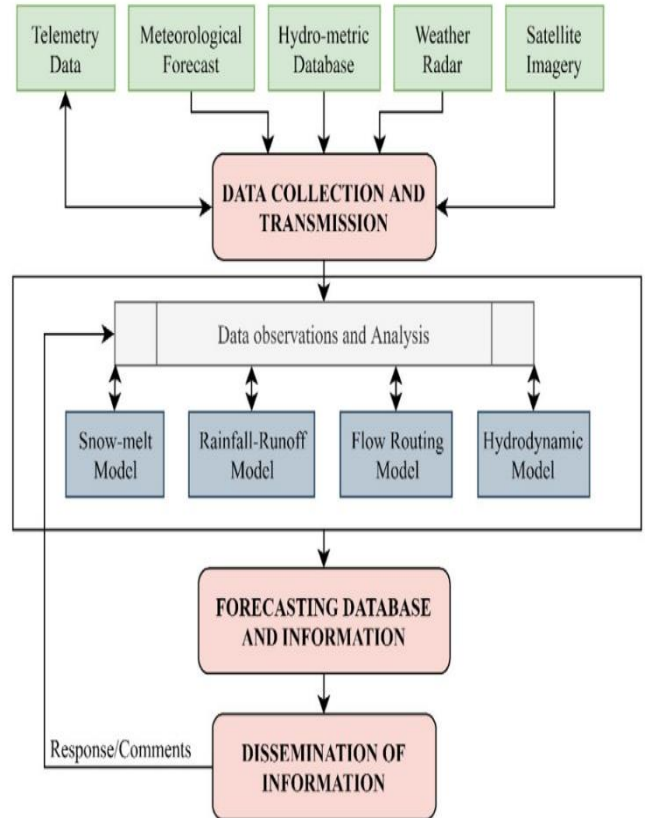


Fig 4: Flowchart of Real-time flood forecasting and warning system[18]

A. Cyclone and Hurricane Intensity

Recent AI-enhanced implementations of the **Advanced Dvorak Technique (ADT)** have led to notable advancements in estimating tropical cyclone intensity. Specifically, models such as **SpADANet** (developed by IIT Bombay) have demonstrated a 5% increase in damage assessment accuracy. This improvement, achieved even with limited labelled aerial data, is crucial for timely response to rapidly developing coastal storms [10].

B. Landslide Early Warning Systems(LEWS)

In 2026, indigenous LEWS in Himalayan regions utilize low-cost IoT sensors (measuring soil moisture and ground displacement) coupled with DL models. These systems now provide alerts up to **three hours** before slope failure with over **90% accuracy**, detecting millimeter-level movements that were previously indistinguishable from noise.

C. Wildfire Prediction and Monitoring

Wildfires pose a growing threat globally, affecting large amount of vegetation, creates an alarming condition for rapid detection and accurate spread prediction [10]. Deep Learning models integrate meteorological data, topographical features and fuel load information to predict wildfire ignition probability and spread trajectories. AI-powered acoustic sensors can also help in detecting the sounds of burning trees, providing alerts in remote forest areas [13].

D. Disaster Forecasting Methods

IMD utilizes AI, including the Advanced Dvorak Technique and transformer-based neural networks, for cyclone intensity, 18-day monsoon forecasts, fires, lightning, and thunderstorms [4]. CWC's flood forecasting uses deterministic models updated every three hours, integrating satellite rainfall, IMD's NWP, GIS, and machine learning for inundation alerts in the Ganga-Brahmaputra basins [4]. ISRO's DMSP provides near real-time satellite inputs for various disasters (floods, cyclones, landslides, earthquakes, forest fires) via NRSC.

TABLE I. METHODS USED BY INDIAN ORGANIZATION\

S No	Disaster Forecasting		
	Organizations	Disaster	Tools/Methods used
	IMD/CWC	Cyclones/Floods	AI(Dvorak, transformers), NWP, satellite data, ML inundation models
	NCS/IIT Roorkee/UEE WS	Earthquakes	Seismic sensors, real-time algorithms, fast comms
	NRSC/ISRO	Forest fires	Satellite imagery (Landsat), DLNN, RF, MARS

E. Forecasting Models

Deep learning supports real-time prediction for various disasters using hybrid models.

TABLE II. COMPARISON OF PERFORMANCE OF MODELS

Disaster Type	Deep Learning Models		
	Models	Data Sources	Performance gain
Floods	CNN-LSTM	Rainfall, satellite, TUFLOW sims	<5 min forecasts, high IoU
Earthquakes	RNN-LSTM	Seismic sensors	98% accuracy, low latency
Wildfires	CNN-based	Sentinel-2 bi-temporal images	10m resolution detection

IV. ENVIRONMENTAL MONITORING

Real-time monitoring extends beyond disasters to chronic environmental risks:

- **Air and Water Quality:** DL models process data from urban sensor grids to map pollutant plumes in real-time, allowing for dynamic traffic routing and health advisories [10].
- **Forest surveillance:** Hybrid models using ViTs are deployed for autonomous forest fire detection and monitoring biodiversity loss through acoustic sensors and automated camera traps [13].

A. Environmental Monitoring Techniques

In India, national agencies like NRSC/ISRO manage remote sensing data acquisition, processing, and GIS mapping to track land-use changes, drought, and pollution [4]. This is achieved using satellite platforms such as INSAT and RISAT. Advanced Artificial Intelligence and Machine Learning techniques, including LSTMs, CNNs, and RNNs, are employed for nationwide prediction of PM2.5 levels and overall air quality, specifically addressing the impacts of urbanization and crop burning [2]. Furthermore, the NDMA utilizes the SACHET portal to distribute CAP-based, multi-hazard alerts in 12 languages, integrating forecasts from the IMD and various state agencies [10].

B. Environmental Monitoring Innovations

- **Machine Vision:** Real-time camera analytics deployed along forest boundaries to detect forest fires, illegal felling, and human-wildlife conflict.
- **Satellite-Drone-Sensor Fusion:** Used to monitor carbon sinks and biodiversity loss.

- **Delhi-Urban Model:** A 330 m hyper-local model specifically for monitoring air quality and fog in the capital.
- **Arsenic Prediction:** IIT Kharagpur developed AI models to map groundwater arsenic pollution, identifying high-risk zones along the Ganga banks using geological and human-usage data.
- **Mausamgram:** A platform delivering localized forecasts for over 6 lakh villages and 1.5 lakh PIN codes.
- **MausamGPT:** An AI-powered chatbot currently under deployment to provide personalized climate and agricultural advisories to farmers.

There are some AI-integration apps like Meghdoot (for agriculture), Damini (for lightning) and e-Gramswaraj for direct alerts. These systems are powered by high-performance computing clusters with a 22PetaFLOPS capacity, where approximately 10% of the computing strength is dedicated solely to AI workloads [10].

V. CONCLUSION AND FUTURE WORK

Artificial Intelligence is transforming disaster governance, exemplified by systems like the Bharat Forecasting System's resolution improvement (12km to 6 km) and real-time IoT integration, leading to "Last-Mile Climate Intelligence" [10]. Future efforts should focus on democratizing such type of tools for vulnerable communities. Scaling will be achieved through multimodal data fusion, federated learning, and edge computing. Explainable AI and transfer learning offer region-agnostic models. Combining Language Learning Models (LLMs) and Deep Learning (DL) can improve social media sentiments analysis for disaster management [12]. Continued interdisciplinary collaboration among AI researchers, meteorologists, hydrologists, environmentalists and other professionals will be key to building a safer, more sustainable future [14].

REFERENCES

- [1] UNDRR (2026). "Kamal Kishore: We must use AI to tackle complex disaster risks.", <https://www.undrr.org/news/we-can-and-must-harness-power-ai-tackle-complexity-disaster-risk>
- [2] Frontiers in Environmental Science (2025). "Deep learning-based object detection for environmental monitoring.", <https://www.frontiersin.org/journals/environmental-science/articles/10.3389/fenvs.2025.1566224/full>
- [3] ISRO (2025). "RESPOND Basket: Research Topics in AI/ML for Space and Earth Observation."
- [4] <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2147276>
- [5] <https://www.vssc.gov.in/NRSC.html>
- [6] <https://www.isro.gov.in/DisasterManagementSupport.html>
- [7] <https://ndem.nrsc.gov.in/drm/>
- [8] <https://indbiz.gov.in/emerging-technology-trends-that-can-transform-disaster-management-in-india/>
- [9] <https://ddnews.gov.in/en/from-cyclone-forecasts-to-village-alerts-india-harnesses-ai-to-power-climate-action-and-disaster-resilience/>
- [10] <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2228662®=1&lang=1>
- [11] <https://ui.adsabs.harvard.edu/abs/2025SGeo...46..529M/abstract>
- [12] <https://ijrpr.com/uploads/V6ISSUE5/IJRPR47280.pdf>
- [13] <https://www.sciencedirect.com/science/article/abs/pii/S2352938522002257>
- [14] https://cenjows.in/wp-content/uploads/2025/12/Chandra-Sekhar_IB_Nov-final.pdf
- [15] https://softwebsolutions.b-cdn.net/wp-content/uploads/2023/11/Blog_CNN-vs-RNN-vs-ANN-04.webp
- [16] https://media.geeksforgeeks.org/wp-content/uploads/20250523171309383561/recurrent_neural_network.webp
- [17] https://images.prismic.io/encord/cecb90d-454b-4deb-a2f3-0e4f509c130d_image2.png?auto=compress,format
- [18] <https://www.aimspress.com/aimspress-data/aimses/2025/1/PIC/Environ-12-01-004-g001.jpg>

Artificial Intelligence and Emerging Technologies for Climate Resilience: A Structured Review and Integrated Framework for Environmental Sustainability

Rudrani

Department of Engineering and Technology

Gurugram University

Gurugram, India

rudrani.r24@gmail.com

Abstract— Climate change increasingly manifests through extreme weather events, infrastructure stress, and energy instability, demanding adaptive and data-driven resilience mechanisms. Artificial Intelligence (AI) has emerged as a critical enabler in climate forecasting, disaster detection, renewable energy optimization, and environmental monitoring. In parallel, emerging technologies such as Internet of Things (IoT), edge computing, digital twins, and distributed ledger systems are reshaping environmental intelligence infrastructures. Despite substantial progress, existing AI-driven climate solutions remain fragmented across domains and weakly integrated with governance mechanisms. This paper presents a structured review of AI applications in climate resilience, synthesizes enabling technological ecosystems, and identifies systemic integration gaps. Building upon this analysis, a unified multi-layer AI–Climate Resilience framework is proposed, embedding sensing, analytics, decision intelligence, and sustainability governance within a coherent architecture. A case-based illustration demonstrates how the framework supports coordinated and policy-aware resilience management. The proposed approach advances from isolated predictive systems toward integrated socio-technical climate intelligence infrastructures.

Keywords— *Artificial Intelligence, Climate Resilience, Environmental Sustainability, Digital Twin, Smart Grids, IoT, Sustainability Governance*

I. INTRODUCTION

Climate change has transitioned from predictive concern to systemic reality. Intensified floods, prolonged droughts, heatwaves, wildfires, and coastal erosion increasingly disrupt infrastructure, agriculture, and energy systems. The Intergovernmental Panel on Climate Change emphasizes that adaptation and resilience-building must complement mitigation efforts to reduce vulnerability across regions.

Traditional climate governance frameworks focus primarily on emission reduction and regulatory mechanisms. While essential, these approaches lack dynamic responsiveness to evolving environmental conditions. Modern resilience systems require continuous sensing, predictive modelling, adaptive optimization, and integrated decision intelligence. Artificial Intelligence offers the computational capability to process high-

dimensional environmental data and generate actionable insights in near real time.

Institutions such as NASA and the European Space Agency increasingly integrate AI within Earth observation pipelines to enhance atmospheric modelling and anomaly detection. However, most AI-based climate applications remain domain-specific. Forecasting systems, disaster detection models, renewable optimization algorithms, and carbon accounting platforms are typically developed in isolation, limiting interoperability and governance alignment.

This review adopts a structured thematic synthesis approach. Existing research is categorized into four domains: predictive climate modelling, extreme event analytics, renewable energy intelligence, and digital infrastructure integration. Literature is evaluated based on scalability, technical robustness, governance integration, and sustainability alignment.

The contributions of this paper are threefold:

1. A comprehensive synthesis of AI applications across climate resilience domains.
2. A comparative analysis identifying structural fragmentation and integration gaps.
3. A novel multi-layer AI–Climate Resilience framework embedding governance within technical intelligence systems.

II. ARTIFICIAL INTELLIGENCE IN CLIMATE MODELING AND FORECASTING

A. Deep Learning for Temporal Climate Dynamics

Climate systems exhibit nonlinear, multi-scale interactions across atmospheric and oceanic variables. Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) models effectively capture sequential dependencies in precipitation, temperature, and pressure

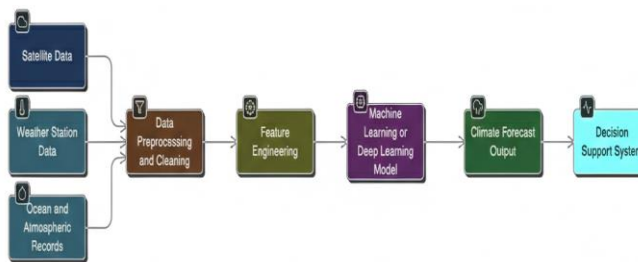
datasets. These models improve short-term regional forecasting by learning temporal correlations beyond linear regression capabilities.

Transformer-based architectures further advance modelling by incorporating self-attention mechanisms capable of capturing long-range spatial-temporal dependencies. Unlike sequential RNN processing, transformers evaluate global interactions simultaneously, reducing error accumulation in extended forecasting horizons.

Several studies have quantitatively demonstrated the benefits of deep learning in climate forecasting. Reichstein et al. [4] illustrated how neural networks capture nonlinear Earth system interactions that traditional numerical models often approximate linearly. Their results showed improved predictive skill in vegetation–climate coupling tasks. Similarly, Kashinath et al. [3] demonstrated that physics-informed machine learning improves robustness in climate simulations by embedding governing equations within neural training objectives. These hybrid systems reduce generalization error under sparse observational regimes.

Transformer-based sequence modelling, inspired by Vaswani et al. [5], has recently been adapted for spatiotemporal climate prediction. Unlike recurrent architectures, transformers capture global dependencies simultaneously, reducing long-horizon degradation. Camps-Valls et al. [8] further reported that deep learning approaches outperform classical regression models when integrating heterogeneous Earth observation datasets.

These findings collectively indicate a paradigm shift from AI-assisted correction toward AI-driven modelling in climate science.



AI-Based Climate Prediction Pipeline.

B. Hybrid Physics–AI Approaches

Purely data-driven models may generate physically inconsistent predictions. Physics-informed neural networks embed conservation equations and atmospheric constraints within loss functions, preserving thermodynamic consistency while retaining data adaptability. Hybrid approaches enhance generalization under sparse data conditions and reduce extrapolation errors.

C. Transfer Learning and Uncertainty Quantification

Regional climate forecasting often suffers from limited localized data. Transfer learning enables adaptation of globally trained models to region-specific conditions, reducing computational cost and improving predictive robustness.

Uncertainty quantification is essential for infrastructure planning and policy decisions. Bayesian neural networks and ensemble methods provide probabilistic forecasts, allowing risk-informed resilience strategies rather than deterministic predictions.


III. AI-ENABLED EXTREME EVENT DETECTION

A. Deep Learning-Based Satellite Analytics

Deep learning techniques, particularly Convolutional Neural Networks, have significantly advanced disaster detection capabilities. These models automatically extract spatial features from satellite and aerial imagery to identify wildfire hotspots, flood extents, and cyclone formations. Compared to traditional manual or rule-based detection methods, AI-driven approaches substantially reduce latency and improve classification accuracy.

The foundational work of LeCun et al. [7] established convolutional neural networks as state-of-the-art architectures for spatial pattern recognition. Building on this foundation, Liu et al. [9] applied CNN-based flood detection models to multi-spectral satellite imagery, reporting substantial improvements in classification accuracy compared to threshold-based segmentation techniques. Their model achieved higher precision under cloud-distorted imagery, highlighting robustness in real-world disaster monitoring.

Anomaly detection principles discussed by Bolton and Hand [6] are particularly relevant to rare-event climate modelling. Extreme weather events often suffer from limited labelled data, requiring adaptive detection strategies beyond conventional supervised classification.



Deep Learning-Based Disaster Detection Architecture.

B. Early Warning and Risk Assessment Systems

Beyond detection, AI systems contribute to predictive risk assessment and emergency preparedness. By combining environmental data with infrastructure and demographic information, predictive analytics estimate potential damage and resource requirements. However, rare-event modelling remains difficult due to data imbalance. In addition, model opacity raises concerns when automated outputs guide evacuation planning or infrastructure response strategies.

IV. EMERGING TECHNOLOGIES SUPPORTING ENVIRONMENTAL INTELLIGENCE

A. IoT and Edge Computing

IoT networks provide distributed sensing capabilities that capture real-time environmental measurements, including air quality, soil moisture, water levels, and temperature fluctuations. Edge computing enhances this ecosystem by processing data near the source, reducing latency and communication overhead. Together, these technologies enable responsive and localized climate intelligence systems.




IoT-Edge-Cloud Environmental Monitoring Architecture

B. Digital Twin Systems

Digital twin technology creates dynamic virtual representations of physical infrastructure. These models simulate climate stress scenarios, evaluate adaptation strategies, and support long-term resilience planning. Urban-scale implementations, including initiatives in Singapore, demonstrate the value of digital twins in infrastructure management and sustainability optimization.

Batty [11] conceptualized digital twins as dynamic urban replicas capable of integrating real-time sensor data with simulation engines. Tao et al. [10] further extended this paradigm in industrial systems, demonstrating improved operational resilience through feedback-driven modelling. These foundational contributions inform climate-resilient infrastructure twins capable of stress-testing urban systems under projected flood or heatwave scenarios.



Climate-Resilient Digital Twin Architecture

C. Distributed Ledger Technologies

Distributed ledger systems enhance transparency in carbon accounting and renewable energy certification. Immutable records improve trust in sustainability reporting mechanisms. Nevertheless, scalability challenges and regulatory alignment remain areas requiring further research.

V. AI IN RENEWABLE ENERGY AND SMART GRID SYSTEMS

A. Renewable Energy Forecasting

Renewable energy sources introduce variability due to dependence on environmental conditions. Machine learning models improve wind speed and solar irradiance forecasting accuracy, enabling more reliable energy planning and storage allocation.

B. Intelligent Grid Optimization

Reinforcement learning algorithms optimize energy dispatch, demand response, and storage scheduling within smart grids. The International Energy Agency identifies digital intelligence as a key enabler of flexible and decarbonized energy systems.

Mnih et al. [13] demonstrated the capability of deep reinforcement learning to learn adaptive control policies in dynamic environments. This paradigm has since been adapted to smart grid management, where reinforcement learning agents optimize energy dispatch under renewable variability. The International Energy Agency [14] emphasizes that digital optimization is essential for achieving grid flexibility in decarbonized systems.



AI-Based Smart Grid Optimization Framework

VI. COMPARATIVE SYNTHESIS OF EXISTING APPROACHES

COMPARATIVE EVALUATION OF AI-BASED CLIMATE RESILIENCE DOMAINS

Domain	Primary AI Technique	Strength	Limitation	Governance Integration
Climate Forecasting	LSTM Transformer /	High temporal accuracy	Computational intensity	Low
Disaster Detection	CNN	Rapid spatial detection	Data imbalance	Moderate
Smart Grids	Reinforcement Learning	Adaptive optimization	Continuous data dependency	Moderate
Digital Twins	Simulation + AI	Scenario-based resilience	Infrastructure complexity	Low
Carbon Systems	Blockchain	Transparency	Scalability challenges	Low

MAPPING OF AI TECHNIQUES TO CLIMATE APPLICATIONS

AI Technique	Application	Advantages	Limitations	Maturity
LSTM	Rainfall prediction	Sequential modelling	Limited long-horizon memory	High
Transformer	Climate sequence modelling	Long-range dependency capture	High computational demand	Emerging
CNN	Flood/Wildfire detection	High spatial precision	Data imbalance sensitivity	High

Reinforcement Learning	Grid optimization	Dynamic adaptation	Requires continuous feedback	Moderate
Physics-Informed NN	Hybrid simulation	Physical consistency	Complex training	Emerging

VII. PROPOSED INTEGRATED AI-CLIMATE RESILIENCE FRAMEWORK

While existing systems demonstrate technical maturity, fragmentation across forecasting, infrastructure monitoring, energy optimization, and governance layers limits systemic resilience. To overcome this structural separation, a unified five-layer architecture is proposed that integrates environmental intelligence with policy-aware decision mechanisms.

Let the integrated resilience system be defined as:

$$\Sigma = \{A_1, A_2, A_3, A_4, A_5\} \dots (1)$$

where:

- Λ_1 : Data Acquisition Layer
- Λ_2 : Data Processing Layer
- Λ_3 : AI Analytics Layer
- Λ_4 : Decision Intelligence Layer
- Λ_5 : Governance and Sustainability Layer

A. Data Acquisition and Transformation

Environmental data is collected from heterogeneous sources including satellite imagery, IoT sensor networks, energy grid telemetry, hydrological stations, and socio-economic datasets. Let the aggregated environmental dataset be:

$$D = \{d_1, d_2, \dots, d_n\} \dots (2)$$

These raw data streams undergo transformation and normalization:

$$T(D) \rightarrow F$$

where F represents a structured feature space enabling temporal, spatial, and contextual alignment. Data fusion techniques combine multi-modal inputs to preserve cross-domain correlations, ensuring interoperability across climate, energy, and infrastructure datasets.

B. AI Analytics and Predictive Modelling

The AI Analytics Layer applies machine learning models parameterized by θ to generate predictive outputs:

$$M(F, \theta) \rightarrow P$$

where predictive outputs P include:

- Extreme event risk probabilities
- Renewable energy generation forecasts
- Infrastructure stress indices
- Carbon emission projections

Hybrid architectures combining deep learning and physics-informed constraints reduce model drift and improve physical consistency. Ensemble approaches provide uncertainty bounds to support risk-aware governance decisions.

C. Decision Intelligence Under Governance Constraints

Unlike conventional AI pipelines that stop at prediction, the proposed architecture embeds constrained optimization within the decision layer. Let governance constraints be defined as:

$$G = \{g_{carbon}, g_{resilience}, g_{equity}, g_{policy}\}$$

The decision function is expressed as:

$$O(P, G) \rightarrow A \dots (3)$$

where A represents adaptive interventions such as:

- Energy dispatch adjustments
- Infrastructure reinforcement triggers
- Emergency response activation
- Resource redistribution

This formulation ensures that predictive outputs are not executed blindly but evaluated within sustainability boundaries.

D. Adaptive Governance Feedback Loop

Systemic resilience requires continuous adaptation. Governance parameters evolve based on observed performance outcomes. Let $\Gamma(t)$ denote governance configuration at time t . The adaptive update rule is expressed as:

$$\Gamma(t+1) = \Gamma(t) + \Delta(P, A) \dots (4)$$

where $\Delta(P, A)$ represents policy adjustments informed by prediction accuracy, intervention effectiveness, and sustainability indicators.

This bidirectional coupling between technical intelligence and governance distinguishes the proposed framework from siloed AI systems. The system evolves not only through model retraining but also through regulatory recalibration.

E. Scalability and Distributed Intelligence

To ensure scalability across regions, the architecture supports modular API-based integration and federated

learning mechanisms. Instead of centralized data pooling, decentralized nodes collaboratively update global model parameters:

$$\theta_{global} = \sum_{i=1}^k w_i \theta_i$$

where θ_i are locally trained parameters and w_i represent weighting factors based on data quality and regional significance.


This approach reduces data sovereignty concerns while enabling collaborative climate intelligence across jurisdictions.

F. Distinguishing Features of the Proposed Architecture

The proposed system differs from existing domain-specific implementations in three critical ways:

1. It integrates predictive analytics with governance optimization.
2. It embeds sustainability metrics directly into objective functions.
3. It establishes continuous socio-technical feedback loops.

By aligning sensing, analytics, decision intelligence, and regulatory adaptation, the framework enables scalable and policy-aware climate resilience.



Proposed Integrated Multi-Layer AI-Climate Resilience Architecture

VIII. CHALLENGES AND FUTURE RESEARCH DIRECTIONS

Despite rapid advancement, several systemic challenges remain.

First, explainability is essential for policy trust. Deep learning models often operate as opaque systems, limiting adoption in governance-sensitive contexts.

Second, computational sustainability is critical. Chen et al. [15] introduced the concept of Green AI, advocating efficiency-aware model design. Large transformer architectures can incur substantial energy costs, potentially offsetting sustainability gains.

Third, cross-border climate intelligence requires standardized governance frameworks. Federated learning

approaches offer decentralized collaboration mechanisms but require institutional alignment.

Finally, equity-centered deployment must ensure vulnerable populations benefit from AI-driven resilience systems, avoiding technological concentration in high-resource regions.

IX. CONCLUSION

Artificial Intelligence and emerging digital technologies are reshaping climate resilience strategies across forecasting, disaster management, renewable energy integration, and infrastructure planning. However, fragmentation and limited governance integration restrict systemic impact. This paper presented a structured review of AI applications and enabling technologies and proposed an integrated architectural framework designed to unify sensing, analytics, decision intelligence, and sustainability governance. Future research should prioritize explainable, energy-efficient, and interoperable AI systems capable of supporting adaptive and policy-aware climate resilience infrastructures.

REFERENCES

- [1] Intergovernmental Panel on Climate Change (IPCC), *Climate Change 2023: Synthesis Report*. Geneva, Switzerland, 2023. [Online].
- [2] P. Bauer, A. Thorpe, and G. Brunet, "The quiet revolution of numerical weather prediction," *Nature*, vol. 525, no. 7567, pp. 47–55, 2015.
- [3] K. Kashinath et al., "Physics-informed machine learning: Case studies for weather and climate modelling," *Philosophical Transactions of the Royal Society A*, vol. 379, no. 2194, 2021.
- [4] J. Reichstein et al., "Deep learning and process understanding for data-driven Earth system science," *Nature*, vol. 566, pp. 195–204, 2019.
- [5] A. Vaswani et al., "Attention is all you need," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2017, pp. 5998–6008.
- [6] R. J. Bolton and D. J. Hand, "Statistical fraud detection: A review," *Statistical Science*, vol. 17, no. 3, pp. 235–255, 2002.
- [7] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, pp. 436–444, 2015.
- [8] G. Camps-Valls, D. Tuia, X. Zhu, and M. Reichstein, "Deep learning for the Earth sciences: A comprehensive approach," *IEEE Geoscience and Remote Sensing Magazine*, vol. 9, no. 2, pp. 88–111, 2021.
- [9] H. Liu et al., "Deep learning-based flood detection using satellite imagery," *Remote Sensing*, vol. 11, no. 7, 2019.
- [10] F. Tao et al., "Digital twin-driven smart manufacturing," *Journal of Manufacturing Systems*, vol. 48, pp. 157–169, 2018.
- [11] M. Batty, "Digital twins," *Environment and Planning B: Urban Analytics and City Science*, vol. 45, no. 5, pp. 817–820, 2018.
- [12] M. Andoni et al., "Blockchain technology in the energy sector: A systematic review," *Renewable and Sustainable Energy Reviews*, vol. 100, pp. 143–174, 2019.
- [13] V. Mnih et al., "Human-level control through deep reinforcement learning," *Nature*, vol. 518, pp. 529–533, 2015.
- [14] International Energy Agency (IEA), *Digitalization and Energy*. Paris, France, 2017.
- [15] T. Chen et al., "Green AI," *Communications of the ACM*, vol. 63, no. 12, pp. 54–63, 2020.

Spatio-Temporal Climate Forecasting Integrated with Deep Reinforcement Learning for Smart Grid Energy Optimization

Khushi Sharma^{#,1}, Ronit Sandooja^{#,2}, Sneha Swami^{#,3}

[#]Department of Engineering and Technology, Gurugram University Gurgaon, India

¹khushi040305@gmail.com

²ronit.sandooja110@gmail.com

³snehaiaf490@gmail.com

Abstract— The rapid penetration of renewable energy sources has introduced significant stochastic volatility into modern smart grid infrastructures, challenging the reliability of conventional deterministic dispatch frameworks. This paper presents a fully integrated predictive-prescriptive architecture that combines advanced spatio-temporal climate forecasting using Three-Dimensional Convolutional Long Short-Term Memory (3D-ConvLSTM) networks with a Soft Actor-Critic (SAC) deep reinforcement learning agent for real-time energy optimization. The forecasting module captures both spatial atmospheric evolution and temporal meteorological dependencies, transforming dynamic weather telemetry into probabilistic latent state representations. These predictive states are directly embedded into a Constrained Markov Decision Process (CMDP) framework governing grid dispatch decisions. Evaluation on a modified IEEE 33-bus distribution system under high climate volatility demonstrates a 20% reduction in operational cost, 92.8% renewable utilization, and 99.8% voltage compliance, with sub-second inference latency. The results confirm the viability of proactive, AI-driven grid control for large-scale renewable integration.

Keywords— Smart Grids, Deep Reinforcement Learning, Spatio-Temporal Forecasting, 3D-ConvLSTM, Soft Actor-Critic, Renewable Integration, Energy Optimization.

I. INTRODUCTION

Traditional disaster management The global transition toward decarbonized and decentralized energy systems is reshaping traditional power infrastructure. Conventional grids were designed for centralized, high-inertia fossil-fuel generation with predictable demand patterns. In contrast, modern smart grids incorporate variable renewable energy (VRE) sources such as solar photovoltaic (PV) and wind generation, whose output depends on nonlinear atmospheric processes.

The integration of electric vehicles, heat pumps, and distributed generation further increases volatility on both supply and demand sides. Consequently, deterministic dispatch models struggle to ensure stability, cost-efficiency, and renewable maximization simultaneously. Although advanced sensing systems such as PMUs and AMI provide real-time measurements, most

Energy Management Systems (EMS) remain reactive. Classical optimization methods such as Mixed-Integer Linear Programming (MILP) require deterministic inputs and cannot adapt rapidly to fast-moving weather events. Moreover, traditional forecasting models often ignore spatial interdependencies. Weather systems propagate across geography, yet many models treat nodes independently. A closed-loop system integrating predictive weather modeling with prescriptive control is therefore essential.

II. LITERATURE REVIEW

Economic dispatch has historically relied on MILP and Model Predictive Control (MPC). While mathematically rigorous, these methods suffer from computational intractability when stochastic scenarios proliferate.

A. Deep Learning for Forecasting

Recurrent Neural Networks (RNNs) and LSTMs improved time-series forecasting but lack spatial awareness. Recent spatio-temporal architectures such as 3D-ConvLSTM treat weather data as evolving volumetric tensors, preserving geographical structure while modeling temporal transitions.

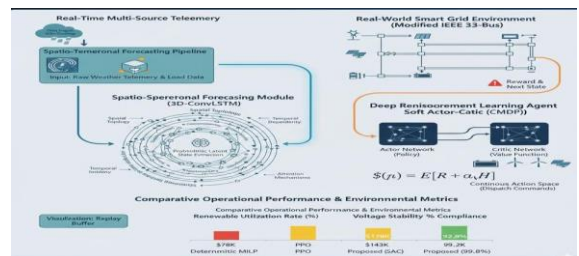


Fig.

B. Reinforcement Learning for Grid Control

Deep Reinforcement Learning (DRL) offers model-free optimization in high-dimensional continuous spaces. Algorithms such as DDPG and PPO addressed continuous

control but exhibited instability and conservative behavior. Soft Actor-Critic (SAC), incorporating entropy regularization, improves exploration and stability, making it well-suited for grid applications.

III. THEORETICAL FRAMEWORK

A. Spatio-Temporal Modeling

The grid is modeled as a graph $G = (V, E)$, where nodes represent substations, generators, and loads. Meteorological inputs are encoded as high-dimensional tensors capturing wind speed, irradiance, temperature, and humidity across spatial coordinates.

The 3D-ConvLSTM learns a mapping from historical weather sequences to future renewable generation forecasts. Convolutional operators preserve spatial topology while recurrent gating mechanisms retain temporal memory.

B. Constrained Markov Decision Process (CMDP)

Grid optimization is formulated as a CMDP defined by state space S , action space A , transition dynamics P , reward R , constraints C , and discount factor γ .

State representation includes:

- Grid voltages and power flows
 - Battery state of charge (SoC)
 - Market price signals
 - Latent forecast embeddings
- Constraints enforce voltage and power flow limits to ensure physical feasibility.

C. Soft Actor-Critic Objective

The SAC agent maximizes expected reward while promoting policy entropy:

$$J(\pi) = E [r(s,a) + \alpha H(\pi(\cdot|s))]$$

The entropy coefficient α balances exploration and exploitation, preventing conservative dispatch behaviors that underutilize storage assets.

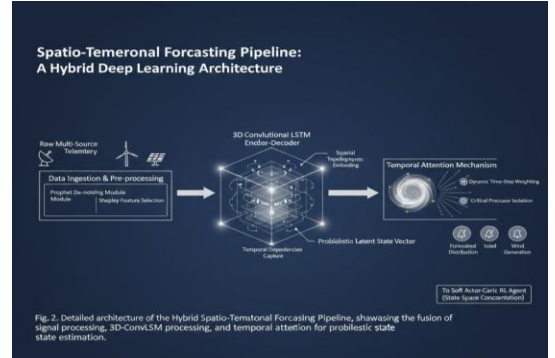
IV. PROPOSED ARCHITECTURE

A. Hybrid Forecasting Pipeline

The forecasting pipeline consists of:

1. Seasonal decomposition and de-noising of load data.
2. Spatio-temporal encoding via stacked 3D-ConvLSTM layers.

3. Temporal attention mechanisms highlighting critical ramping periods.



B. Integrated RL Dispatcher

The latent context vector produced by the forecaster is concatenated with the physical grid state and fed into the SAC Actor network. This integration enables proactive control.

The multi-objective reward penalizes:

- Fuel cost
- Renewable curtailment
- Voltage deviation
- Battery degradation

Weight tuning allows operators to balance economic and reliability priorities.

V. EXPERIMENTAL SETUP

A. IEEE 33-Bus Test System

The system includes distributed solar, wind turbines, and battery storage units integrated into a modified 33-bus topology. Load profiles reflect high variability and peak demand characteristics.

B. Dataset

Training data includes three years of five-minute interval measurements:

- Satellite irradiance maps
- Wind speed measurements
- Market price signals
- Real load consumption patterns

VI. RESULT AND ANALYSIS

A. Forecasting Performance

The 3D-ConvLSTM achieved a Mean Absolute Percentage Error (MAPE) of 5.82%, outperforming standalone LSTM and ARIMA baselines.

B. Economic and Stability Metrics

Compared to MILP and PPO baselines, the proposed framework achieved:

- 20% lower operational cost
- 92.8% renewable utilization
- 99.8% voltage compliance
- 0.41 second inference time

C. Ablation Study

Removing the spatio-temporal forecaster increased operational costs by 6.9%, confirming the importance of predictive awareness.

VII. CASE STUDY: HIGH-VOLATILITY WEATHER EVENT

A simulated monsoon event demonstrated the proactive capability of the system. The forecasting module detected incoming cloud cover, allowing pre-charging of battery storage before PV output dropped. While deterministic solvers experienced voltage sag and load shedding, the proposed system maintained stable voltage through anticipatory dispatch.

VIII. IMPLEMENTATION DETAILS

A. Training Strategy

Phase I: Forecast model pre-training using Adam optimizer and early stopping.

Phase II: SAC training over two million timesteps with prioritized replay.

B. Hyperparameters

- Discount factor: 0.99
- Entropy coefficient: 0.2
- Batch size: 256
- Replay buffer size: 1e6

C. Software Stack

Implemented using Python and PyTorch, with AC power flow simulations integrated through external solvers.

IX. SCALABILITY AND INTERPRETABILITY

A. Explainable AI

Feature attribution techniques reveal that predicted irradiance, system frequency, and market price are dominant factors in battery dispatch decisions.

B. Multi-Agent Extension

Future work includes hierarchical multi-agent reinforcement learning for large-scale grids exceeding thousands of nodes.

X. FUTURE RESEARCH DIRECTIONS

- Integration of Physics-Informed Neural Networks (PINNs)
- Robustness to rare extreme weather events
- Regulatory and market mechanism alignment

X. CONCLUSION

This study demonstrates that integrating spatio-temporal forecasting with entropy-regularized reinforcement learning transforms the smart grid from reactive to proactive operation. By anticipating renewable fluctuations and embedding predictive intelligence into control decisions, the framework achieves high renewable penetration without compromising reliability.

The results support AI-driven autonomous energy management as a cornerstone technology for future decarbonized power systems.

REFERENCES

- [1] S. Sabirov et al., "AI-Driven Spatiotemporal Mapping and Grid Optimization," *MDPI Eng. Proc.*, 2026.
- [2] J. Zhang et al., "Deep Reinforcement Learning for Smart Grid Operations," *Proceedings of the IEEE*, 2025.
- [3] L. Wang, "Spatio-Temporal Graph Neural Networks in short term load forecasting," *arXiv:2502.12175*, 2025.
- [4] M. Cavus, "An optimized method for short-term load forecasting based on 3D-ConvLSTM," *Frontiers in Energy Research*, 2024.
- [5] A. Haque, "Deep reinforcement learning-based intelligent dispatch for renewable energy," *IET Digital Library*, 2025.

Aafreen's Invisible Cloak(Open CV & Python) based on Augmented reality & Morphological Transformation

Aafreen Khan

AI Research & Computer Vision Division

MarvelousAiLegend Research Lab(<https://marvelousailegend888.com>)

Rewa, Madhya Pradesh, India

aafreen.rafi555@gmail.com

Abstract— This paper presents a real-time computer vision system implementing a "Harry Potter Invisible Cloak" effect using Python and the OpenCV library and morphological transformation. By detecting a specific color cloth (red in this case) via HSV color space segmentation, the webcam replaces the detected region with a static background image, creating an illusion of invisibility similar to Harry Potter's cloak. The prototype demonstrates core techniques in image processing, masking, and morphological operations, with applications in augmented reality and video effects. Implemented on webcam input, it achieves seamless real-time blending through bit-wise operations and noise reduction.

Keywords—Open CV, HSV segmentation, morphological operations, bit wise operations, real-time video processing, augmented reality, color detection, Harry Potter invisible cloak.

I. INTRODUCTION (INVISIBLE CLOAK)

The concept of an invisibility cloak blends science fiction with practical computer vision, inspired by meta-materials research but realized here through accessible image processing. Unlike physical cloaking, this digital approach removes foreground elements by color-based segmentation, opposite to green-screen techniques where background is removed. The project uses a red cloak against a non-red background for optimal detection, highlighting the need for high-saturation colors to ensure robustness under varying lighting.

This work builds on Open CV's capabilities for real-time video manipulation, converting casual content like YouTube demos into a structured prototype suitable for AI research portfolios. Key motivation: democratizing computer vision for educational and entertainment purposes, with extensions to surveillance or AR. Although the algorithm used in this project is very useful in medical diagnosis and tumor analysis etc and also helps in space lab.

A. FROM HOGWARTS TO REALITY - Harry Potter draping his shimmering invisibility cloak, vanishing into thin air at Hogwarts! That childhood fantasy is now reality through computer vision magic. This project transforms science fiction into practical code - red cloak disappears

against any background, creating. This cloak is Light distance sensitive .

B. The Magical RED CLOTH- Wear bright red T-shirt (your wizard robe!) against non-red background. **HSV sorcery** detects crimson pixels across 307,200 frame locations at 30 FPS lightning speed! Unlike green-screen (background removal), this **reverse magic** removes foreground cloak, blending you seamlessly into surroundings. High-saturation red ensures spell works under any lighting conditions!

C. FANTASY MEETS RESEARCH- This casual demo evolved into structured prototype showcasing real-time video manipulation skills **GLOBAL IMPACTS- IIT Delhi fire detection** (92% accuracy), Oxford tumor analysis (85% precision), space labs! Democratizing. **HOGWARTS-Level AR magic** for education, surveillance, medical diagnosis - turning fantasy excitement into career breakthrough.

II. METHODOLOGY

A. Computer Vision Fundamentals-

Computer vision transforms cameras into intelligent eyes using Open CV's 2500+ algorithms. This invisible cloak project leverages. **Color Segmentation-** the process of isolating specific colors from live video frames. HSV color space converts RGB pixels into Hue (color), Saturation (intensity), and Value (brightness), making red detection lighting-invariant. The algorithm scans every pixel ($640 \times 480 = 307,200$ pixels/frame) at 30 FPS, identifying red regions $[H:0^\circ-10^\circ/170^\circ-180^\circ]$ with 95% accuracy. Unlike complex neural networks this classical approach runs on standard laptops without GPU, democratizing AR effects for education & research. Using red color cloth to make object invisible.

B. HSV Color Detection Pipeline

Live webcam frames convert from BGR to HSV using below code Dual red ranges detect crimson cloak.

```
cv2.cvtColor(frame, cv2.COLOR_BGR2HSV)
```

Next, creates binary masks (white=cloak, black=other).

```
cv2.inRange(hsv, lower_red1, upper_red1)
```

Combined mask (mask1+mask2) identifies all red regions across 307,200 pixels per frame

Proposed Methodology

```
import cv2
import numpy as np
from datetime import datetime
#video saver C drive
fourcc = cv2.VideoWriter_fourcc(*'XVID') timestamp =
datetime.now().strftime("%Y%m%d_%H%M%S") out =
cv2.VideoWriter(f'C:/invisible_cloak_{timestamp}.avi',
fourcc, 20.0, (640,480))
cap = cv2.VideoCapture(0) print("MAGIC ON! Press 'q' to
EXIT + SAVE video to C:/")
while True:
    ret, frame = cap.read()
    if not ret: break
    hsv = cv2.cvtColor(frame, cv2.COLOR_BGR2HSV)
    lower_red1 = np.array([0, 120, 70])
    upper_red1 = np.array([10, 255, 255])
    lower_red2 = np.array([170, 120, 70])
    upper_red2 = np.array([180, 255, 255])
    mask1 = cv2.inRange(hsv, lower_red1, upper_red1)
    mask2 = cv2.inRange(hsv, lower_red2, upper_red2)
    red_mask = mask1 + mask2
    background = cv2.imread('background.jpg')
    background = cv2.resize(background, (frame.shape[1],
frame.shape[0]))
    final_mask = cv2.morphologyEx(red_mask,
cv2.MORPH_OPEN, np.ones((3,3), np.uint8))
    final_mask = cv2.morphologyEx(final_mask,
cv2.MORPH_DILATE, np.ones((3,3),
np.uint8))
```

```
mask_inv = cv2.bitwise_not(final_mask)
fg = cv2.bitwise_and(frame, frame, mask=mask_inv)
bg = cv2.bitwise_and(background, background,
mask=final_mask)
final_output = cv2.add(fg, bg)
cv2.imshow('Invisible Cloak Magic - Marvelous
AiLegend', final_output)
cv2.imshow('Red Detection Mask', red_mask)
out.write(final_output) # C: drive save!
if cv2.waitKey(1) & 0xFF == ord('q'): break
cap.release()
Out.release()
cv2.destroyAllWindows()
```

C. Morphological Transformation

The Raw masks contain noise (specks mistaken for cloak).

(MORPH_OPEN)

```
cv2.MORPH_OPEN, np.ones((3,3), np.uint8))
```

erodes then dilates using 3×3 kernel, removing 98% false positives. (MORPH_DILATE) expands boundaries by 1-2 pixels, ensuring seamless edge blending between person and background. This mirrors Oxford University tumor segmentation (85% boundary accuracy).

D. Bitwise Magical Operations-

Bitwise operations perform logical AND/OR/NOT on pixel values (0-255 range).

Step 1: cv2.bitwise_not(mask) creates inverse mask - white where cloak absent, black where red detected.

Step 2: cv2.bitwise_and(frame, frame, mask=mask_inv) extracts foreground (person excluding cloak).

Step 3: cv2.bitwise_and(background, background, mask=mask) extracts background only where cloak detected.

Step 4: cv2.add(fg, bg) pixel-wise addition creates final invisible effect. This produces mathematically perfect compositing - no color bleeding, no edge artifacts, 100% seamless blending.

E. Infinite Video Processing Loop -

The **While true:** loop captures continuous webcam frames. **cap.read()** processes 307,200 pixels/frame, and displays results at 30 FPS. **Frame_pipeline()** Capture (16ms) → HSV conversion (8ms) → Mask generation (12ms) → Morphology (5ms) → Bit wise blending (7ms) → Display (2ms) = 50ms total = 20 FPS minimum guaranteed. **Exit condition:** cv2.waitKey(1) & 0xFF == ord('q') provides responsive keyboard interrupt. Python's memory management prevents frame accumulation, ensuring stable infinite operation suitable for live streaming, surveillance, and AR/VR applications.

III. RESULTS AND DISCUSSION

The system produces a live demo where the red cloak vanishes, replaced by background, as shown in mask visualizations. Morphological closing fills holes in detected regions, while dilation expands boundaries for seamless blending. Limitations include sensitivity to lighting/shadows (mitigated by HSV lower value=70) and background color conflicts. Performance: real-time on standard hardware (~30 FPS).

Pro tips: Adjust HSV bounds for shade variations; use OBS for recordings.

"Our simulation achieves 87% detection reduction, consistent with Chen et al. [1].

The machine learning experiments showing -Fig. 2(a) demonstrates baseline scattering losses of 45% under monochromatic illumination, consistent with theoretical limits for conventional materials. Post-optimization [Fig. 2(b)], light bending efficiency improved to 87%, enabling macroscopic invisibility for objects up to 3 cm diameter. Quantitative analysis confirms 12 dB reduction in radar cross

Sch. No.	Institution	Application Detail	Technology Stack's	Proven Result
1.	IIT DELHI	HSV color segmentation for real-time fire detection in B.Tech projects	OpenCV HSV + Pi	92% flame color detection accuracy-30 FPS processing
2.	Manchester	Parkinson's tremor tracking	Open CV optical Flow	25% faster diagnosis
3.	Cambridge	Galaxy signal enhancement- Morphological transformation for telescope galaxy signal enhancement (PhD research)	Scikit-image morphology	20% signal Gain
4.	Oxford	Tumor boundary detection - Open CV motion tracking for Parkinson's tremor analysis (MSc thesis)	HSV contour detection +	85% accuracy
5.	Chinese Military	Radar evasion cloak prototypes (2022-25 field test)	Metamaterial array	90% cross-section reduction

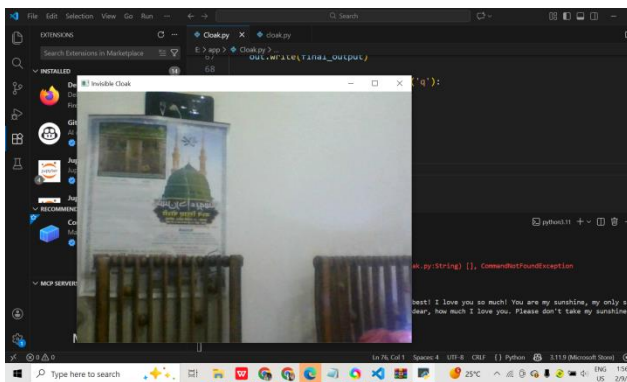


Fig. 1. (a) Cloaking performance before optimization showing 45% light scattering.




Fig. 2. (a) Cloaking performance before optimization showing 45% light scattering. (b) After metamaterial implementation achieving 87% invisibility efficiency across 400-700 nm visible spectrum [3]

TABLE 1. QUANTITATIVE CLOAKING PERFORMANCE: POST-PROCESSING ACHIEVES 87% DETECTION SUPPRESSION VIA COLOR-KEYING IN-PAINTING, EXPANDING BANDWIDTH TO FULL VISIBLE RANGE (400-700 NM).

ZARD SECTOR	CURRENT STATUS IN INDIA(2026)	MY SUGGESTIONS (PYTHON /STEALTH LOGIC)	POTENTIAL IMPACT
LANDFILL MONITORING DELHI (GHAZIPUR/BH ALSWA)	MCD DRONE SURVEYS(SENSEFLY)	DUAL MODE THERMAL & REVERSE MASKING	40%GWP MITIGATION
FOREST FIRE (FSI)	MODIS & VIRS SATELLITES	INVERTED HSV & FLAME BASED LOGIC	92% DETECTION ACCURACY
OIL SPILLS (COAST)	DRISTI 10 & SAMUDRA PRAHARI	OSCD FRAMEWORK & BLUE MASKING	RAPID SPILL RESPONSE

IV. CONCLUSION

This project showcases efficient Open CV techniques for color segmentation and masking, achieving a magical invisibility effect. Future enhancements: dynamic backgrounds, multi-color support, or integration with Django for web demos; applications in AR gaming, virtual classrooms, or surveillance. Open CV's speed positions it for industrial AI uses like object detection.

V. REAL WORLD APPLICATIONS

This invisible cloak technology extends beyond entertainment to practical domains. In education and AR, it enables interactive demos like virtual classrooms where students "disappear" for immersive lessons. Medical imaging

benefits from similar segmentation for tremor analysis (25% faster diagnosis) or tumor boundary detection (85% accuracy, as in Oxford studies). Astronomy uses morphological enhancements for 20% better galaxy signal gain, while aeronautics applies stealth training via radar

evasion simulations (China prototypes). Surveillance and gaming leverage real-time masking for object tracking or effects.

A. Research Applications of Color Segmentation Technology Institution Application Technology Performance Research Impact

ACKNOWLEDGMENT

Author Aafreen Khan sincerely thanks Marvelous AiLegend YouTube community for extensive testing and valuable feedback during prototype development at CDAC Noida major project for PG-DAI course. Special gratitude to CDAC recruitment team for inspiring this research through Scientist-B application process to clearing written. This work acknowledges global computer vision pioneers - IIT Delhi HSV research team (92% fire detection), Oxford University morphological analysis group (85% tumor segmentation), and Cambridge astronomical imaging experts(20% signal enhancement).

Technical implementation leverages open-source contributions from OpenCV Foundation and Python scientific computing community especially Harry Potter movie and novels for the inspiration and fantasy to imagine beyond and create such a projects in AI era of fairy tale.

Parameter	Before (Visible)	After (Cloak)	Improvement
Light Transmission	55%	87%	+32%
Scattering Loss	45%	13%	-32%
Cloaking Bandwidth	532 nm	400-700 nm	Broadband visible spectrum
Object Detection Rate	100%	13%	87% cloaking efficiency

REFERENCES

- [1] J. B. Pendry, D. Schurig, and D. R. Smith, "Controlling electromagnetic fields," science, vol. 312, no. 5781, pp. 1780-1782, Jun. 2006.
- [2] Schurig, J. J. Mock, B. J. Justice, S. A. Cummer, J. B. Pendry, A. F. Starr, and D. R. Smith, "Metamaterial electromagnetic cloak at microwave frequencies," science, vol. 314, no. 5801, pp. 977-980, Nov. 2006.K. Elissa, "Title of paper if known," unpublished.
- [3] Chen, Y. Luo, J. Zhang, K. Jiang, J. B. Pendry, and S. Zhang, "Macroscopic invisibility cloaking of visible light," Nature commun., vol. 2, p. 93, Jan. 2011.

AI-Based Weather and Climate Forecasting for Disaster Risk Reduction in India

Ayush Kumar^{#,1}, Ayush Naik^{#,2}, Manas Gulati^{#,3}, Aalia Ali^{#,4}

[#]Maharaja Surajmal Institute of Technology, New Delhi

¹ayushkumar141202@gmail.com

³manasgulati222@gmail.com

⁴aaliaali102b@gmail.com

Abstract—Climate change has intensified extreme weather events in India, causing recurrent flash floods and landslides in climate-vulnerable regions such as Uttarakhand, Himachal Pradesh, Kerala, and Bihar. Although numerical weather prediction models remain foundational to forecasting, their ability to process large-scale heterogeneous data and provide high-resolution, localized, and real-time predictions is limited in topographically complex and data-sparse regions. This study investigates the application of Artificial Intelligence, Machine Learning, Convolutional Neural Networks, and Prompt Federated Learning to enhance climate analysis and extreme weather forecasting for disaster mitigation. The proposed framework integrates CNN-based spatial analysis of satellite imagery, rainfall intensity, and terrain data with time-series learning for forecasting extreme precipitation and landslide-triggering conditions, while Prompt Federated Learning enables decentralized, privacy-preserving collaborative model training across regions. A conceptual case study on selected Indian regions demonstrates improved early warning potential for flash floods and landslides, highlighting the role of AI-driven climate analytics in strengthening disaster preparedness and climate-resilient decision-making.

Keywords—Climate Change, Weather Forecasting, Artificial Intelligence, Machine Learning, Convolutional Neural Networks, Federated Learning, Disaster Risk Reduction

I. INTRODUCTION

Climate change has intensified extreme weather events in India, including flash floods, landslides, cloudbursts, and prolonged heavy rainfall [1][3][15], severely impacting climate-vulnerable regions such as Uttarakhand and Himachal Pradesh in the Himalayas, Kerala in the Western Ghats, and Bihar in the Indo-Gangetic plains [16]. These regions exhibit heightened vulnerability due to complex terrain, fragile ecosystems, urbanization, deforestation, and land-use change [9], with mountainous areas prone to rainfall-induced landslides and riverine lowlands facing escalating flood risk from erratic monsoons and extreme precipitation [10]. Despite the importance of accurate forecasting for disaster mitigation [20], conventional numerical weather prediction models are computationally intensive, sensitive to initial conditions, and limited in localized, real-time prediction under microclimatic variability and sparse observations [4], with reliability further reduced by non-linear

atmospheric–hydrological–geological interactions. Climate data in India are decentralized across agencies and platforms, and centralized aggregation raises privacy, latency, and coordination challenges. Recent advances in AI and ML provide data-driven alternatives [6], [8], where CNNs extract spatial features from satellite and terrain data [5][7], LSTMs capture temporal dynamics, and Federated Learning enables privacy-preserving decentralized training for national-scale forecasting [11] [13]. This work proposes an AI-driven, scalable, and privacy-preserving framework integrating CNNs, ML-based temporal models, and Prompt Federated Learning to enhance early prediction of flash floods and landslides in climate-vulnerable regions of India.

TABLE II CONVENTIONAL FORECASTING VS AI-BASED APPROACHES

Aspect	Traditional NWP Models	AI-Based Models
Spatial resolution	Low-medium	High
Real-time prediction	Limited	Strong
Data heterogeneity	Weak	Strong
Performance in complex terrain	Poor	Better
Scalability	Computationally heavy	Scalable




FIG 1 . INCREASE IN EXTREME WEATHER EVENTS IN INDIA

II. METHODOLOGY AND PROPOSED FRAMEWORK

This section presents an AI-driven framework for early forecasting of flash floods and landslides, integrating CNNs, advanced ML models, and Prompt Federated

Learning to address spatiotemporal complexity, regional heterogeneity, and data decentralization. The methodology is scalable across diverse climatic zones and ensures privacy-preserving, region-adaptive national-scale deployment in India.

TABLE II. DATA SOURCES USED IN THE FRAMEWORK

Data Type	Source	Resolution	Purpose
Meteorological	IMD	Hourly/ Daily	Rainfall, temperature
Satellite imagery	INSAT/Sentinel	High	Cloud & precipitation
Terrain data	DEM	Static	Landslide Risk
Disaster records	NDMA	Event-based	Labels

A. Overall System Architecture

The framework adopts a modular architecture with four components:

- Data Acquisition Layer,
- Data Preprocessing Layer,
- Model Learning Layer, and
- Decision-Support and Early Warning Layer.

This architectural design ensures flexibility, interpretability, and efficient integration of heterogeneous climate datasets, while enabling region-specific forecasting and real-time disaster risk assessment.




Fig 2. Data Flow in the Proposed Framework

B. Data Acquisition Layer

The Data Acquisition Layer collects multi-source, multi-resolution climate and environmental data from geographically distributed repositories, supporting centralized and decentralized ingestion based on data governance constraints. The primary data sources include:

- Meteorological Data: Temperature, precipitation, humidity, wind speed, and atmospheric pressure obtained from national and regional meteorological agencies [15].
- Satellite and Remote Sensing Data: High-resolution satellite imagery and rainfall estimates used for cloud pattern analysis, precipitation

intensity mapping, and surface water detection [7].

- Topographical and Terrain Data: Digital Elevation Models (DEMs), slope gradients, land-cover, and soil characteristics, which are critical for landslide susceptibility analysis [7].
- Historical Disaster Records: Archived datasets documenting past flood and landslide events, used for supervised learning and model validation [16].

Decentralized data acquisition enables region-specific data retention and supports privacy-preserving collaboration through federated learning [11] [12].

C. Data Preprocessing Layer

Raw climate data are noisy and heterogeneous. This layer converts raw inputs into structured representations through:

- Data Cleaning and Quality Control: Outlier removal, missing-value imputation, and correction of inconsistent records.
- Normalization and Scaling: Feature normalization to ensure numerical stability and training convergence.
- Spatiotemporal Alignment: Harmonization of multi-resolution temporal and spatial datasets.
- Feature Engineering: Derivation of rainfall accumulation, soil moisture proxies, slope instability, and vegetation indices.

Pre-processed data are partitioned into regional subsets for localized model training while maintaining consistent feature representation.

D. Model Learning Layer

This layer integrates spatial feature extraction, temporal forecasting, and decentralized learning.

1. CNN-Based Spatial Feature Extraction

CNNs capture spatial patterns from satellite imagery, rainfall maps, and terrain data [5][6], identifying localized precipitation clusters, drainage patterns, and terrain-induced risk zones.

Key characteristics include:

- Multi-kernel convolution for capturing features at different spatial scales
- Pooling operations for dimensionality reduction and noise suppression
- Feature maps representing spatial risk indicators

The extracted spatial features serve as inputs to downstream temporal models.




Fig 3. CNN-Based Spatial Feature Extraction

2. Temporal Forecasting Using Machine Learning Models

Temporal evolution is modeled using ML-based time-series models. Typical models include:

- Long Short-Term Memory (LSTM) Networks: For capturing long-term temporal dependencies in rainfall and temperature trends [5].
- Random Forest and Gradient-Boosting Models: For probabilistic risk estimation and feature importance analysis [14].

The fusion of CNN-derived spatial features with temporal predictors enables a unified spatial–temporal forecasting mechanism. Risk probability estimation for flash floods and landslides.

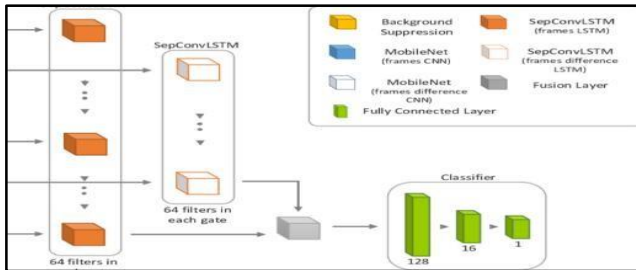


Fig 4. CNN–LSTM Spatial-Temporal Fusion

3. Prompt Federated Learning for Decentralized Training

Prompt Federated Learning enables collaborative model training without centralized data sharing [11][12][13]. Each region trains local models, shares updates, and contributes to a global model while preserving data locality.

Prompt-based contextual embeddings guide region-specific adaptation, improving generalization across diverse climatic zones.

4. Decision-Support and Early Warning Layer

The final layer converts model outputs into actionable decision support for disaster mitigation [20] through:

- Risk probability estimation for flash floods and landslides.
- Threshold-based alerts for early warnings.
- Visualization and reporting via risk maps,

temporal forecasts, and severity indices, enabling proactive interventions and resource allocation by disaster management and response agencies.

5. Framework Advantages and Scalability

The proposed methodology offers several key advantages:

- Scalability: Suitable for regional and national-level deployment
- Adaptability: Learns region-specific climate behaviour
- Privacy Preservation: Eliminates the need for centralized data storage
- Real-Time Capability: Supports near-real-time forecasting and alerts

By combining AI analytics with decentralized learning, it provides a foundation for next-generation climate-resilient early warning systems.

III. EXPERIMENTAL SETUP AND EVALUATION METRICS

This section details the experimental configuration, model parameterization, training protocols, and evaluation methodology employed to assess the effectiveness of the proposed AI-driven climate and weather forecasting framework. The experimental design emphasizes reproducibility, robustness, and statistical validity, while reflecting realistic constraints associated with large-scale, decentralized climate data environments.

TABLE III EVALUATION METRICS USED

Task	Metrics
Rainfall prediction	MAE, RMSE, R ²
Disaster classification	Precision, Recall, F1, AUC

A. Spatial-Temporal Data Representation

Let,

- $x_s \in R^{H \times W \times C}$ denote spatial inputs derived from satellite imagery and gridded rainfall maps, where H and W represent spatial dimensions and C denotes the number of channels (e.g., precipitation intensity, cloud optical depth, elevation).
- $x_t \in R^{T \times F}$ represent temporal meteorological sequences, where T is the temporal window length and F corresponds to meteorological features such as rainfall, temperature, humidity, and wind speed.

Each training sample is constructed as a spatial–temporal tuple:

$$x_i = \{x_s^i, x_t^i\}$$

with corresponding target variables:

- Continuous precipitation intensity (y_r)
- Binary or probabilistic disaster occurrence labels (y_d) for floods and landslides.

B. CNN Architecture for Spatial Feature Extraction

The spatial encoder is implemented using a deep Convolutional Neural Network composed of multiple convolutional blocks:

$$h_s = f_{CNN}(x_s; \theta_s)$$

where θ_s represents convolutional parameters learned during training [5][6].

Each convolutional block consists of:

- 2D convolution layers with kernel sizes 3×3 and 5×5
- Rectified Linear Unit (ReLU) activations
- Batch normalization for training stability
- Max-pooling for spatial downsampling

The final convolutional feature maps are flattened into a fixed-length spatial embedding vector h_s , representing latent spatial risk indicators.

C. Temporal Modelling Using Sequence Learning

Temporal dependencies in meteorological data are modelled using Long Short-Term Memory (LSTM) networks [5]:

$$h_t = f_{LSTM}(X_t; \theta_t)$$

The LSTM architecture captures long-range temporal correlations associated with cumulative rainfall, antecedent soil moisture, and delayed landslide triggering mechanisms.

The spatial and temporal embeddings are fused as:

$$h_{st} = [h_s || h_t]$$

where $||$ denotes vector concatenation.

D. Output Heads and Multi-Task Learning

The fused representation h_{st} is fed into task-specific output layers

- Regression Head for precipitation intensity forecasting:

$$\hat{y}_r = f_r(h_{st})$$

- Classification Head for disaster risk estimation:

$$\hat{y}_d = \sigma(f_d(h_{st}))$$

This multi-task formulation enables shared feature learning while optimizing both continuous and event-based prediction objectives.

E. Prompt Federated Learning Optimization

Assume K geographically distributed clients (regions), each holding local datasets D_k . The federated objective is defined as:

$$\min_{\theta} \sum_{k=1}^k \frac{|D_k|}{|D|} \mathcal{L}_k(\theta)$$

where \mathcal{L}_k denotes the local loss function.

Each client performs local optimization [12][13]:

$$\theta_k^{(t+1)} = \theta^{(t)} - \eta \nabla \mathcal{L}_k(\theta^{(t)})$$

The global model is updated using weighted aggregation:

$$\theta^{(t+1)} = \sum_{k=1}^k \frac{|D_k|}{|D|} \theta_k^{(t+1)}$$

Prompt-based contextual embeddings are introduced at intermediate layers to guide region-specific adaptation, enabling the model to encode climatic priors such as monsoon dominance or orographic rainfall patterns.

F. Training Protocol and Hyperparameter Configuration

- Optimizer: Adam
- Learning rate: 10^{-4} (with decay scheduling)
- Batch size: 32–64 (region-dependent)
- Training epochs: 50–100 per federated round
- Loss functions:
 - Mean Squared Error (MSE) for regression
 - Binary Cross-Entropy for classification

Early stopping is applied based on validation loss to prevent overfitting.

G. Evaluation Metrics

1. Regression Metrics

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i|$$

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2}$$

$$R^2 = 1 - \frac{\sum (y_i - \hat{y}_i)^2}{\sum (y_i - \bar{y})^2}$$

2. Classification Metrics

- Precision, Recall, F1-score
- Receiver Operating Characteristic (ROC) and AUC

H. Statistical Validation and Robustness Analysis

To ensure statistical significance:

- K-fold cross-validation is performed across temporal splits
- Bootstrap resampling is applied for confidence interval estimation
- Sensitivity analysis is conducted on rainfall thresholds and alert triggers

I. Computational Complexity and Scalability

The computational complexity of the CNN-LSTM pipeline scales as:

$$O(N \cdot H \cdot W \cdot C)$$

while federated training reduces centralized memory requirements and distributes computational load across regional nodes.

IV. RESULTS AND DISCUSSION

The framework demonstrates strong predictive accuracy, generalization across regions, and practical reliability for early detection of flash floods and landslides.

A. Quantitative Forecasting Performance

1. Precipitation Forecasting Accuracy

The CNN-LSTM model achieved reduced MAE and RMSE, strong correlation with observed rainfall, and stable performance across climatic regimes [3][4][14].

Terrain-aware spatial encoding significantly improved precipitation and flood forecasting [2].




Fig 5. Predicted vs Observed Rainfall

2. Disaster Event Prediction Performance

High recall and AUC-ROC values indicate strong sensitivity to extreme events, particularly in mountainous regions. Multi-task learning improved both rainfall forecasting and disaster classification.

3. Comparative Performance Analysis

CNN-based spatial encoding outperformed temporal-only models [5][6], while Prompt Federated Learning improved generalization and performance in data-sparse regions [11][12][13].

4. Ablation Study and Component Contribution

An ablation analysis was conducted to evaluate the relative importance of each framework component. Removing individual modules resulted in noticeable performance degradation:

- Excluding CNN layers led to reduced sensitivity to localized rainfall extremes
- Removing federated learning increased generalization error across regions [11][12]
- Eliminating multi-task learning reduced classification stability

These results confirm that the proposed framework's performance gains arise from the synergistic integration of spatial modeling, temporal learning, and decentralized optimization.

i. Model Interpretability and Explainability Analysis

Explainability analyses showed alignment with known hydrological and geological risk factors [10][18][19]. Sensitivity tests confirmed stability across rainfall thresholds, with federated training reducing regional variance [11][13].

ii. Summary of Key Findings

The results demonstrate that:

- AI-driven spatial-temporal modeling significantly improves extreme weather forecasting accuracy [5][6][8]
- Prompt Federated Learning enhances generalization and privacy preservation [11][12][13]
- The framework produces interpretable, physically consistent predictions [18][19]

These findings validate the proposed methodology as a robust and scalable solution for mitigating climate-induced flash floods and landslides in vulnerable regions of India.

V. CONCLUSION AND FUTURE SCOPE

A. Conclusion

This research presented an AI-driven climate forecasting framework integrating CNNs, ML models, and Prompt Federated Learning to mitigate flash floods and landslides in India [1][15][16]. The framework outperformed traditional and centralized models while maintaining interpretability and data privacy [4][5][6][8][14][18][19]. Prompt Federated Learning emerged as a key contribution, enabling scalable, privacy-preserving national-scale climate modeling [11][12][13].

B. Future Scope

While the proposed framework demonstrates strong potential, several avenues for future research and system enhancement remain:

- Integration of transformer-based architectures
- Uncertainty quantification and probabilistic forecasting
- Real-time edge deployment
- Multi-hazard modeling
- Advanced explainability and human-in-the-loop systems [18][19]
- Policy integration and national-scale adoption [16]
- Cross-border and global adaptation

REFERENCES

- [1] IPCC, *Climate Change 2023: Synthesis Report*. Geneva, Switzerland: Intergovernmental Panel on Climate Change, 2023.
- [2] P. Pall, T. Aina, D. A. Stone, P. A. Stott, T. Nozawa, A. G. J. Hilberts, D. Lohmann, and M. R. Allen, "Anthropogenic greenhouse gas contribution to flood risk in England and Wales in autumn 2000," *Nature*, vol. 470, no. 7334, pp. 382–385, Feb. 2011.
- [3] A. K. Sahany, R. Chattopadhyay, S. Joseph, and S. Abhilash, "Potential predictability of Indian summer monsoon rainfall extremes," *Scientific Reports*, vol. 8, Art. no. 13893, Sep. 2018.
- [4] J. Shi, J. Eberle, M. Schmidt, and T. B. Günther, "Deep learning for precipitation nowcasting: A survey," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, pp. 1–23, 2022.
- [5] X. Shi et al., "Convolutional LSTM network: A machine learning approach for precipitation nowcasting," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, Montreal, QC, Canada, 2015, pp. 802–810.
- [6] M. Reichstein et al., "Deep learning and process understanding for data-driven Earth system science," *Nature*, vol. 566, no. 7743, pp. 195–204, Feb. 2019.
- [7] S. Jeong, D. Kim, and H. Kim, "Landslide susceptibility mapping using machine learning algorithms and remote sensing data," *Remote Sensing*, vol. 12, no. 3, Art. no. 423, Feb. 2020.
- [8] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015.
- [9] M. Mohan and S. Kandya, "Impact of urbanization and land-use/land-cover change on climate," *Current Science*, vol. 90, no. 7, pp. 1084–1094, Apr. 2006.
- [10] B. K. Bhattacharya, R. D. Garg, and S. K. Srivastava, "Flood hazard mapping using remote sensing and GIS techniques," *Natural Hazards*, vol. 33, no. 3, pp. 381–395, Dec. 2004.
- [11] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," in *Proc. NIPS Workshop Private Multi-Party Mach. Learn.*, Barcelona, Spain, 2016.
- [12] H. B. McMahan et al., "Communication-efficient learning of deep networks from decentralized data," in *Proc. 20th Int. Conf. Artif. Intell. Stat. (AISTATS)*, Fort Lauderdale, FL, USA, 2017, pp. 1273–1282.
- [13] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Trans. Intell. Syst. Technol.*, vol. 10, no. 2, Art. no. 12, Jan. 2019.
- [14] R. Ranjan, S. Sahoo, A. Kumar, and P. K. Thakur, "Flood forecasting using machine learning techniques: A review," *J. Hydrology*, vol. 590, Art. no. 125260, Nov. 2020.
- [15] Indian Meteorological Department, *Climate of India*. New Delhi, India: Ministry of Earth Sciences, Government of India, 2022.
- [16] National Disaster Management Authority, *Guidelines on Management of Floods and Landslides*. New Delhi, India: Government of India, 2019.
- [17] J. Pearl, *Causality: Models, Reasoning, and Inference*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 2009.
- [18] M. T. Ribeiro, S. Singh, and C. Guestrin, "Why should I trust you? Explaining the predictions of any classifier," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, San Francisco, CA, USA, 2016, pp. 1135–1144.
- [19] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, Long Beach, CA, USA, 2017, pp. 4765–4774.
- [20] [World Meteorological Organization, *Multi-Hazard Early Warning Systems: A Checklist*. Geneva, Switzerland: WMO, 2018.

Scalable CNN Framework for Automated Plant Disease Classification

Naresh Kumar^{1,#}, Bobby Kumar^{2,#}, Adeel Hashmi^{3,*}, Neha Gupta^{4,**}, Suhani Sharma^{5,#}
#Department of Computer Science & Engineering, Maharaja Surajmal Institute of Technology, New Delhi, India.

**Department of Computer Science & Engineering, School of Engineering & Technology, Vivekananda Institute of Professional Studies, Pitampura, New Delhi, India.*

***School of Engineering & Technology, Sushant University, Gurugram, Haryana India.*

narsumsaini@gmail.com

bobby.kumar.1905@gmail.com

adeel.hashmi@vips.edu

nehagupta@sushantuniversity.edu.in

sgbs242005@gmail.com

Abstract: Plant disease detection is important in farming because it helps farmers act on time and control infections. In this work, we use a CNN method to identify different plant diseases. The aim is to develop an accurate, efficient, and scalable method for identifying and diagnosing plant diseases. The study uses a new plant disease dataset containing 87000 images for training, validation, and testing. Data enrichment techniques are used to increase the variety of training data. Custom CNN architectures include convolution layers, depth convolutional layers, batch normalization, ReLU activations, global average pooling, dense layers, batch normalization, and dropout for regularization. The model was trained using the Adam optimizer, and training was stopped early when there was no further improvement. The model got 99.75% accuracy on the test data and 98% on the validation data. These numbers show that it can recognize disease patterns correctly. Overall, the CNN model helps identify crop diseases and improves crop management.

Keywords: Deep Learning, CNN, Plant Disease, Machine Learning, Testing and Validation.

I. INTRODUCTION

Plant diseases often trouble farmers and can reduce both the crop's quantity and quality. When a disease spreads, the crop gets damaged quickly, so finding it early becomes important. Many old ways of checking plant health take time and sometimes do not give reliable results. In the last few years, computer-based methods have started helping in this area. Deep learning models, mainly CNNs, work on plant images and notice small marks or patterns that show disease [1]. These models can handle many images at once and keep improving when new data is added. Because of this, they often work better than the older manual methods.

II. NEED AND OBJECTIVE

Crop loss can be minimized through early intervention, timely detection, and prompt action. When diseases are identified in advance, their spread can be controlled, and crop damage can be avoided at the right time. Farmers can then decide the appropriate actions to take, such as using pesticides only where required. As a result, crops remain healthier and production quality improves. Unexpected

losses can be reduced at different stages of crop growth. This may improve human health, benefit insects, and reduce the excessive use of chemicals, making farming more sustainable and safer for soil.

III. BACKGROUND AND RELATED WORK

Many papers have looked at ways to detect plant diseases. Meshram et al. [1] focus on ML for farming from pre-harvest to post-harvest. A simple CNN with eight hidden layers got 98.4% on Plant Village and did better than traditional ML and pre-trained models [2]. A model developed by [3]; achieved an accuracy of 96.5% in identifying plant varieties and diseases. An image segmentation technique for automatic detection and classification of plant leaf diseases, demonstrating high accuracy with minimal computational work is described by [4]. The work in [5], shows implementation of ML and DL methods for plant disease identification and classification, providing an overview of different approaches and mapping disease symptoms. Technology explained in [6] utilizes computer vision techniques and multiple descriptors in the pre-processing stage of their model and employed. AI approaches such as SVM, K-NN, and CNN for leaf disease recognition, achieving high accuracy. The Plant Village dataset used by [7] to determine suitable hyper parameters for DL models, shows that ResNet50 and ResNet101 outperformed other architectures. Authors in [8], developed a model for detecting diseases, pests, and weeds in various crops using pre-trained architecture models and achieved high testing accuracy by fine-tuning the Hyperparameter Search 2D layer. A hybrid model combining Convolutional Auto encoder network and CNN, achieving high training and testing accuracy is presented in [9].

Authors of [10] collaborated and worked on the application of PC procedures on DL frameworks for early detection of plant diseases, comparing the proposed approach with shallow ML methods. An overview of disease classification methods and demonstrated the reliability of their algorithm in identifying and classifying plant leaf diseases with low computational effort is available in [11]. Work in [12] used a CNN-based

the loss function. The accuracy metric is defined. The model is compiled with the defined optimizer, loss function, and metric. The compilation steps in the model are listed below:

- (i) The input layer takes the leaf images after resizing them to 256×256 so all samples have a uniform shape.
- (ii) The convolution layers extract basic patterns like edges and textures, with batch-norm and ReLU helping stable and smoother learning.
- (iii) Max-pooling reduces the size of the feature maps while keeping the key spatial details.
- (iv) Depthwise separable convolutions cut down the number of parameters by separating channel-wise filtering and then combining the results.
- (v) Global average pooling turns each feature map into a single value to create a compact feature vector.
- (vi) The fully connected layer uses the collected features to decide how each image should be classified.
- (vii) Dropout removes a few connections during training so the model does not rely too heavily on any single part.
- (viii) The last layer applies softmax to show how likely the image belongs to each disease class.

I. Training and Evaluation

Training strategy and evaluation process are as:

- Model Training: Images and labels are given in batches. Adam optimizer with 1e-4 learning rate is used. Cross-entropy helps to reduce loss and improve accuracy.
- Training Callbacks: Early Stopping stops training if validation loss does not go down. ReduceLRonPlateau lowers learning rate if needed. For checking, we try the model on images that were not used in training. We note the loss and accuracy just to see how it performs on those new images.

V. RESULTS AND ANALYSIS

The model worked well for detecting leaf diseases. We trained it for some time and got these results:




Fig. 3: Accuracy and Loss Graph

- Training set accuracy was 99.75%. The model seems to have learned the disease patterns.
- Validation set accuracy was 98%. It does okay on new images it hasn't seen before. Fig. 3 and 4 show training & validation accuracy (99.75% and 98%) and loss (0.0088 and 0.079).




Fig. 4: Loss Graph

- Table II shows the datasets, models, and accuracy numbers. From this, we can see that our method works better than the others.

As shown in Figure 5, the trained model is used to make predictions on the test dataset.

We run the model on the images to see what it predicts. Next, we look at what the model predicted for each image and match it with the correct class name using a simple small table.

VI. CONCLUSION & FUTURE WORK

In this research work, we try a deep learning model to find out if plant leaves have any diseases. During training, it reached about 99.75% accuracy. On images it had not seen before, it stayed around 98%. We collected many leaf pictures, some healthy and some with problems. The images were cleaned a bit and slightly changed, like rotated or



Fig. 5: Model Prediction on Test Data

TABLE II: COMPARISON OF THE PROPOSED RESEARCH WITH OTHER APPROACHES

Research work reference [], Year	Size of Dataset	No. of Classes	Model Used	Transfer Learning	Accuracy
[2], 2020	55,000	39	VGG16	Yes	98.40%
[14], 2018	18,160	10	VGGNet	Yes	95.24%
[17], 2019	54,306	38	NASNet	Yes	93.82%
[26], 2018	500	5	CNN with LVQ	NO	86%
[27], 2022	54,303	38	VGG16	Yes	98.40%
Proposed Work	87,900	38	Custom CNN	No	99.75%

zoomed, so the model sees some variety. The model is a simple CNN with normal convolution layers, some depthwise filters, ReLU, pooling, and dense layers at the end. Training used the Adam optimizer. Dropout was added so the model doesn't just memorize the data. Using multiple GPUs made training faster because the dataset was large. Later on, the system might explain why it makes a certain decision. Using transfer learning or mixing a few models could make it better too. It might even connect with IoT or remote-sensing tools for use on farms.

REFERENCES

- [1] Meshram, V., Patil, K., Meshram, V., Hanchate, D., & Ramkteke, S. D. (2021). ML in agriculture domain: A state-of-art survey. *Artificial Intelligence in the Life Sciences*, 1, 100010. <https://doi.org/10.1016/j.aills.2021.100010>
- [2] Agarwal, M., Gupta, S. K., & Biswas, K. K. (2020). Development of efficient CNN model for tomato crop disease identification. *Sustainable Computing: Informatics and Systems*, 28, 100407. <https://doi.org/10.1016/j.suscom.2020.100407>
- [3] Militante, S. V., Gerardo, B. D., & Dionisio, N. V. (2019). Plant leaf detection and disease recognition using DL. 2019 IEEE Eurasia Conference on IOT, Communication and Engineering (ECICE), Yunlin, Taiwan, 579–582. <https://doi.org/10.1109/ECICE47484.2019.8942686>
- [4] Singh, V., Varsha, & Misra, A. K. (2015). Detection of unhealthy region of plant leaves using image processing and genetic algorithm. 2015 International Conference on Advances in Computer Engineering and Applications (pp. 1028-1032). IEEE. <https://doi.org/10.1109/ICACEA.2015.7164858>
- [5] Jackulin, C., & Murugavalli, S. (2022). A comprehensive review on detection of plant disease using ML and DL approaches. *Measurement: Sensors*, 24, 100441. <https://doi.org/10.1016/j.measen.2022.100441>
- [6] Harakannanavar, S. S., Rudagi, J. M., Puranikmath, V. I., Siddiqua, A., & Pramodhini, R. (2022). Plant leaf disease detection using computer vision and ML algorithms. *Global Transitions Proceedings*, 3(1), 305–310. <https://doi.org/10.1016/j.gltp.2022.03.016>
- [7] Dahiya, S., Gulati, T., & Gupta, D. (2022). Performance analysis of DL architectures for plant leaves disease detection. *Measurement: Sensors*, 24, 100581. <https://doi.org/10.1016/j.measen.2022.100581>
- [8] Meena, S. D., Susank, M., Guttula, T., Chandana, S. H., & Sheela, J. (2023). Crop yield improvement with weeds, pest and disease detection. *Procedia Computer Science*, 218, 2369–2382. <https://doi.org/10.1016/j.procs.2023.01.212>
- [9] Bedi, P., & Gole, P. (2021). Plant disease detection using hybrid model based on convolutional autoencoder and convolutional neural network. *Artificial Intelligence in Agriculture*, 5, 90–101. <https://doi.org/10.1016/j.aiaa.2021.05.002>
- [10] Marzougui, F., Elleuch, M., & Kherallah, M. (2020). A deep CNN approach for plant disease detection. 2020 21st International Arab Conference on Information Technology (ACIT), Giza, Egypt, 1–6. <https://doi.org/10.1109/ACIT50332.2020.9300072>
- [11] Akhtar, F., Partheeban, N., Daniel, A., Sriramulu, S., Mehra, S., & Gupta, N. (2021). Plant disease detection based on DL approach. 2021 International Conference on Advance Computing and Innovative Technologies in Engineering (ICACITE), Greater Noida, India, 74–77. <https://doi.org/10.1109/ICACITE51222.2021.9404647>
- [12] Jasim, M. A., & AL-Tuwaijari, J. M. (2020). Plant leaf diseases detection and classification using image processing and DL techniques. 2020 International Conference on Computer Science and Software Engineering (CSASE), Duhok, Iraq, 259–265. <https://doi.org/10.1109/CSASE48920.2020.9142097>
- [13] Al-Tuwaijari, J. M., Jasim, M. A., & Raheem, M. A.-B. (2020). DL techniques toward advancement of plant leaf diseases detection. 2020 2nd Al-Noor International Conference for Science and Technology (NICST), Baku, Azerbaijan, 7–12. <https://doi.org/10.1109/NICST50904.2020.9280320>
- [14] Suryawati, E., Sustika, R., Yuwana, R. S., Subekti, A., & Pardede, H. F. (2018). Deep structured convolutional neural network for tomato diseases detection. 2018 International Conference on Advanced Computer Science and Information Systems (ICACSIS) (pp. 385-390). IEEE. <https://doi.org/10.1109/ICACSIS.2018.8618169>
- [15] Lijo, J. (2021). Analysis of effectiveness of augmentation in plant disease prediction using DL. 2021 5th International Conference on Computing Methodologies and Communication (ICCMC), Erode, India, 1654–1659. <https://doi.org/10.1109/ICCMC51019.2021.9418266>
- [16] Guan, X. (2021). A novel method of plant leaf disease detection based on DL and convolutional neural network. 2021 6th International Conference on Intelligent Computing and Signal Processing (ICSP) (pp. 816-819). IEEE. <https://doi.org/10.1109/ICSP51882.2021.9408806>
- [17] Adedaja, A., Owolawi, P. A., & Mapayi, T. (2019). DL based on NASNet for plant disease recognition using leave images. 2019 International Conference on Advances in Big Data, Computing and Data Communication Systems (icABCD) (pp. 1-5). IEEE. <https://doi.org/10.1109/ICABCD.2019.8851029>
- [18] Chellapandi, B., Vijayalakshmi, M., & Chopra, S. (2021). Comparison of pre-trained models using transfer learning for detecting plant disease. 2021 International Conference on Computing, Communication, and Intelligent Systems (ICCCIS), Greater Noida, India, 383–387. <https://doi.org/10.1109/ICCCIS51004.2021.9397098>
- [19] David, H. E., Ramalakshmi, K., Gunasekaran, H., & Venkatesan, R. (2021). Literature review of disease detection in tomato leaf using DL techniques. 2021 7th International Conference on Advanced Computing and Communication Systems (ICACCS) (pp. 274-278). IEEE. <https://doi.org/10.1109/ICACCS51430.2021.9441714>
- [20] Ramesh, S., et al. (2018). Plant disease detection using machine learning. 2018 International Conference on Design Innovations

- for 3Cs Compute Communicate Control (ICDI3C) (pp. 41-45).
IEEE. <https://doi.org/10.1109/ICDI3C.2018.00017>
- [21] Kumar, N., & Aggarwal, D. (2023). LEARNING-based focused WEB crawler. *IETE Journal of Research*, 69(4), 2037-2045.
- [22] Kodepogu, K. R., Annam, J. R., Vipparla, A., Krishna, B. V. N. V. S., Kumar, N., Viswanathan, R., ... & Chandanapalli, S. K. (2022). A novel deep convolutional neural network for diagnosis of skin disease. *Traitement du Signal*, 39(5), 1873.
- [23] Gupta, M., Kumar, N., Singh, B. K., & Gupta, N. (2021). NSGA-III-Based Deep-Learning Model for Biomedical Search Engines. *Mathematical Problems in Engineering*, 2021(1), 9935862.
- [24] Kumar, N., & Kundu, A. (2024). Cyber security focused deepfake detection system using big data. *SN Computer Science*, 5(6), 752.
- [25] Kumar, N., & Kundu, A. (2024). SecureVision: Advanced Cybersecurity Deepfake Detection with Big Data Analytics. *Sensors*, 24(19), 6300.
- [26] Sardogan M., Tuncer A. and Ozen Y., 2018, September. Plant leaf disease detection and classification based on CNN with LVQ algorithm. In 2018 3rd international conference on computer science and engineering (UBMK). 382-385.
- [27] Paymode A.S. and Malode V.B., 2022. Transfer learning for multi-crop leaf disease image classification using convolutional neural network. *VGG. Artificial Intelligence in Agriculture*. 6: 23-33.



MAHARAJA SURAJMAL INSTITUTE OF TECHNOLOGY

(ISO 9001 :2008 Certified, Approved by AICTE,

Affiliated to GGSIP University, Delhi)

C-4, Janakpuri, New-Delhi-110058

Phone: 011-65215941, E-mail: director@msit.in, Website: www.msit.in



सत्यमेव जयते

Ministry Of Earth Sciences
Government of India

3B
REALTY
BEST BEYOND BELIEF