

ISBN: 978-93-47587-83-2

ENGINEERING TOMORROW

Intelligent Systems,
Sustainable Design &
Emerging Technologies


Editors:

Ms. Neha Bagga

Ms. Soniya R. Jain

Dr. Barjinder Kaur

Dr. Mohd Faizan Hasan



Bhumi Publishing, India
First Edition: August 2026

Engineering Tomorrow:
Intelligent Systems, Sustainable Design and Emerging Technologies
(ISBN: 978-93-47587-83-2)

DOI: <https://doi.org/10.5281/zenodo.22172182>

Editors

Ms. Neha Bagga

Assistant Professor,
Department of Computer Science and
Engineering, GNA University, Punjab

Ms. Soniya R. Jain

Assistant Professor,
Centre for Professional Courses,
Gujarat University, Ahmedabad

Dr. Barjinder Kaur

Assistant Professor,
School of Natural Sciences,
GNA University, Phagwara

Dr. Mohd Faizan Hasan

Associate Professor,
Department of Mechanical Engineering,
Integral University, Lucknow



Bhumi Publishing

August 2026

Copyright © Editors

Title: Engineering Tomorrow:

Intelligent Systems, Sustainable Design and Emerging Technologies

Editors: Ms. Neha Bagga, Prof. Soniya R. Jain, Dr. Barjinder Kaur, Dr. Mohd Faizan Hasan

First Edition: August 2026

ISBN: 978-93-47587-83-2



DOI: <https://doi.org/10.5281/zenodo.22172182>

All rights reserved. No part of this publication may be reproduced or transmitted, in any form or by any means, without permission. Any person who does any unauthorized act in relation to this publication may be liable to criminal prosecution and civil claims for damages.

Published by Bhumi Publishing,

a publishing unit of Bhumi Gramin Vikas Sanstha



Nigave Khalasa, Tal – Karveer, Dist – Kolhapur, Maharashtra, INDIA 416 207

E-mail: bhumipublishing@gmail.com



Disclaimer: The views expressed in the book are of the authors and not necessarily of the publisher and editors. Authors themselves are responsible for any kind of plagiarism found in their chapters and any related issues found with the book.

PREFACE

Engineering has always played a transformative role in shaping human civilization, and its significance has grown considerably with the rapid advancement of intelligent systems, digital technologies, and sustainable innovation. *Engineering Tomorrow: Intelligent Systems, Sustainable Design and Emerging Technologies* bring together contemporary perspectives, research findings, innovative concepts, and technological developments that are shaping the future of engineering and its applications.

The book focuses on the convergence of engineering principles with emerging technologies to address complex challenges in modern society. Intelligent systems, artificial intelligence, machine learning, automation, robotics, the Internet of Things, advanced materials, smart infrastructure, and data-driven technologies are creating new possibilities across diverse engineering domains. At the same time, sustainable design, energy efficiency, resource conservation, circular economy principles, and environmentally responsible engineering are becoming essential for achieving long-term technological and societal progress.

The chapters in this volume highlight the importance of integrating innovation with sustainability and practical application. They explore how emerging technologies can improve efficiency, productivity, safety, resilience, and quality of life while reducing environmental impacts. Particular emphasis is placed on interdisciplinary approaches, as many contemporary engineering challenges require collaboration among engineering, science, information technology, environmental studies, and other knowledge domains.

This volume is intended to provide a useful platform for academicians, researchers, engineers, students, innovators, and professionals interested in emerging trends in engineering and technology. It encourages the exchange of ideas and promotes research that moves beyond theoretical discovery toward meaningful and scalable applications. We hope this book will inspire readers to explore new possibilities, develop innovative solutions, and contribute to building a smarter, safer, more sustainable, and technologically empowered future.

We sincerely acknowledge the valuable contributions of all authors, reviewers, and collaborators whose collective efforts have made this volume possible.

- Editors

TABLE OF CONTENT

Sr. No.	Book Chapter and Author(s)	Page No.
1.	AI-ENABLED BATTERY MANAGEMENT SYSTEMS FOR ELECTRIC VEHICLES: TOWARDS SAFE AND SUSTAINABLE MOBILITY R. Deivanayaki, S. Immanuel Prabaharan, Aswini M. and S. Muthurajan	1 – 17
2.	COMPARATIVE ANALYSIS OF MACHINE LEARNING MODELS FOR SHORT-RANGE WEATHER FORECASTING USING ERA5 REANALYSIS DATA OVER NORTHWEST INDIA Richa Avasthy	18 – 30
3.	DEEPFAKE DETECTION USING DEEP LEARNING METHODS: A SYSTEMATIC AND COMPREHENSIVE REVIEW Manmeet Kaur	31 – 46
4.	DEVELOPMENT OF AN INTELLIGENT LEARNING SYSTEM: SKILL LEDGER FOR SKILL EXCHANGE AND PAID COURSES Riya Pathania and Neha Bagga	47 – 58
5.	DIGITAL TWIN AND AI-DRIVEN BATTERY MANAGEMENT FOR SUSTAINABLE TRANSPORTATION Aswini M, Deivanayaki. R, Immanuel Prabaharan. S and S. Muthurajan	59 – 71
6.	SUSTAINABLE NANOMATERIALS FOR NEXT-GENERATION PHOTONIC APPLICATIONS Abinaya. L and K. Gayathri	72 – 83
7.	ENHANCING A MERN-BASED STORE MANAGEMENT SYSTEM WITH SELF-HEALING INVENTORY MANAGEMENT AND PREDICTIVE PROCUREMENT INTELLIGENCE Lalita and Neha Bagga	84 – 93
8.	DIGITAL TWINS: TRANSFORMING INTELLIGENT ENGINEERING SYSTEMs Yokeshwari Ravi and Balaganesan Palanivelu	94 – 101
9.	SOCIO-TECHNICAL INFRASTRUCTURE RE-ARCHITECTURE: MULTI-DOMAIN CONVERGENCE OF AGENTIC AI, QUANTUM- RESISTANT EDGE NETWORKS, BIO-WEARABLES, AND AUTONOMOUS MICROGRIDS Priyanka Kande Patil and Ban Triveni Pramod	102 – 109

10.	INTELLIGENT SYSTEMS FOR PROJECT MANAGEMENT AND EVALUATION Salahuddin	110 – 120
11.	MACHINE LEARNING APPLICATIONS IN ENVIRONMENTAL AND ECOLOGICAL SCIENCES Ranjana	121 – 132
12.	NANOTECHNOLOGY AS A CATALYST FOR INTELLIGENT SYSTEMS AND SUSTAINABLE ENGINEERING: MATERIALS, APPLICATIONS AND FUTURE PERSPECTIVES Minhaaj Afreen S. Y and K. Gayathri	133 – 141
13.	ENGINEERING TRUSTWORTHY AI APPLICATIONS: PRACTICAL STRATEGIES FOR SECURE AND RESPONSIBLE AI DEVELOPMENT Udaya Bhaskar Vemuri	142 – 149
14.	DEVELOPMENT OF A SOLAR-POWERED BATTERY CHARGING SYSTEM FOR E-BIKES Rahul Nath, Asif Beg and Pinakee Gogoi	150 – 154
15.	TOWARDS SUSTAINABLE COMPUTING: THE ROLE OF ARTIFICIAL INTELLIGENCE, CLOUD COMPUTING AND EDGE TECHNOLOGIES Haradhan Mardi	155 – 165
16.	SUSTAINABLE AGENTIC AI FOR RESPIRATORY TRIAGE IN LOW-RESOURCE HEALTHCARE Jyothi K. and Pavan M.	166 – 178
17.	DESIGN AND IMPLEMENTATION OF A SMART AUTOMATIC EMERGENCY BRAKING SYSTEM FOR ELECTRIC VEHICLES USING ULTRASONIC SENSORS C. Jayabalan	179 – 183
18.	FROM DIFFERENTIAL EQUATIONS TO SMART GRIDS: A MATHEMATICAL JOURNEY T. Jagan	184 – 189

AI-ENABLED BATTERY MANAGEMENT SYSTEMS FOR ELECTRIC VEHICLES: TOWARDS SAFE AND SUSTAINABLE MOBILITY

R. Deivanayaki*, S. Immanuel Prabaharan, Aswini M. and S. Muthurajan

¹Department of EEE,

²Department of Marine Engineering,

AMET University -Kanathur, Chennai - 603112

*Corresponding author E-mail: deivanayakiphd@ametuniv.ac.in

Abstract

The rapid adoption of electric vehicles (EVs) has increased the need for safe, reliable, and efficient battery technologies and intelligent Battery Management Systems (BMSs). This chapter examines the role of Artificial Intelligence (AI) in enhancing BMS capabilities for lithium-ion batteries, with emphasis on State of Charge (SOC) estimation, State of Health (SOH) prediction, Remaining Useful Life (RUL) estimation, fault diagnosis, predictive maintenance, and smart charging. AI techniques, including machine learning, deep learning, and data-driven modelling, offer improved capabilities for handling the nonlinear and dynamic behavior of batteries and enabling predictive and adaptive management. The chapter also discusses functional safety, fault detection, redundancy, safe-state operation, and relevant safety and cybersecurity considerations for AI-enabled BMSs. Emerging technologies such as digital twins, wireless BMS, cloud-assisted battery management, and AI-based lifecycle management are examined in relation to battery reuse, recycling, and sustainable mobility. The chapter highlights how the integration of AI with battery management and safety mechanisms can improve battery performance, extend service life, enhance operational safety, and contribute to sustainable electric transportation.


1. Introduction

1.1 Growth of Electric Vehicles and the Importance of Battery Technology

The rapid transition toward electric mobility is reshaping the automotive sector and accelerating the development of cleaner and more sustainable transportation systems. Electric vehicles (EVs) offer significant advantages over conventional internal-combustion-engine vehicles, including reduced tailpipe emissions, higher energy-conversion efficiency, and the potential to integrate renewable energy sources. At the core of EV performance is the rechargeable battery, which largely determines driving range, charging time, vehicle cost, power capability, and overall service life. Among the available technologies, lithium-ion batteries have become dominant because of their high energy density, relatively low self-discharge, and favorable power-to-weight characteristics. However, issues related to battery degradation, thermal instability, charging behavior, safety, resource consumption, and end-of-life management remain important barriers to the widespread and sustainable adoption of EVs. Consequently, advances in battery technology must be accompanied by intelligent methods for monitoring, controlling, and optimizing battery operation.

1.2 Role of Battery Management Systems (BMS)

A Battery Management System (BMS) is a critical electronic and computational component responsible for ensuring the safe, reliable, and efficient operation of an EV battery pack. It continuously monitors parameters such as cell voltage, current, temperature, and other operational characteristics to maintain the battery within prescribed safety limits. A BMS also performs essential functions such as state-of-charge (SOC) estimation, state-of-health (SOH) assessment, cell balancing, thermal management, fault detection, and protection against conditions such as overcharging, over-discharging, and excessive temperature. Since battery packs consist of numerous interconnected cells with differences in aging and operating behavior, accurate monitoring and coordinated control are essential for maintaining pack-level performance. Conventional BMS approaches, however, often depend on predefined models, fixed thresholds, and experimentally calibrated parameters, which may have limited adaptability under complex and highly variable real-world operating conditions.



1.3 Need for AI-enabled BMS

The increasing complexity of EV battery systems has created a need for more adaptive and data-driven BMS architectures. Artificial intelligence (AI) and machine learning techniques offer promising capabilities for extracting useful patterns from large volumes of battery data and supporting more accurate decision-making. AI-enabled BMS can enhance SOC and SOH estimation, predict remaining useful life, identify abnormal operating conditions, detect early signs of degradation or thermal events, and optimize charging and thermal-management strategies. Unlike purely model-based approaches, data-driven methods can learn nonlinear relationships among battery operating conditions, environmental factors, usage patterns, and degradation mechanisms. The integration of AI can therefore enable BMSs to move from predominantly reactive protection toward predictive and intelligent battery management. Nevertheless, practical deployment requires attention to data quality, computational requirements, model interpretability, uncertainty, cybersecurity, validation, and the safety-critical nature of battery operation.

1.4 Scope of the Study

This study presents the fundamental concepts, emerging methodologies, and practical considerations associated with AI-enabled BMSs for electric vehicles, with particular emphasis on safe, reliable, and sustainable mobility. It introduces the architecture and key functions of conventional BMSs before examining how AI and machine learning can enhance battery-state estimation, health prediction, fault diagnosis, thermal management, and charging optimization. The chapter also discusses relevant battery data, AI techniques, model-development considerations, and challenges associated with deploying intelligent algorithms in real-world EV applications. Particular attention is given to the balance between prediction accuracy, computational efficiency, interpretability, robustness, and safety. Finally, the chapter highlights emerging research directions and the potential of AI-enabled BMS technologies to improve battery utilization, extend battery lifetime, enhance operational safety, and contribute to the broader goals of sustainable transportation.

2. Fundamentals of Battery Management Systems

2.1 Battery Pack Architecture

An electric vehicle battery pack is a hierarchical energy-storage system composed of individual cells, modules, and the complete battery pack. The cell is the fundamental electrochemical unit in which electrical energy is stored and released through reversible electrochemical reactions. Depending on the required voltage, capacity, power, packaging constraints, and vehicle application, multiple cells are electrically connected in series and parallel configurations. Series connections increase the overall voltage, while parallel connections increase the available capacity and current capability. The selection and arrangement of cells therefore have a direct influence on the energy density, power output, efficiency, cost, and safety of the battery system. Several individual cells are typically grouped to form a battery module. Modules provide a convenient intermediate level for mechanical assembly, electrical interconnection, sensing, cooling, and maintenance. Each module may contain a number of cells connected in a predetermined series-parallel configuration and may incorporate temperature sensors, voltage-sensing circuits, busbars, and thermal-management components. Modular construction also facilitates the integration of battery packs into different vehicle platforms and allows manufacturers to scale the energy and power capacity according to application requirements.

A battery pack is the complete energy-storage assembly installed in an electric vehicle. It consists of multiple modules or, in some advanced designs, a large number of directly integrated cells, together with electrical interconnections, contactors, fuses, sensors, cooling systems, structural components, and the Battery Management System (BMS). The pack must operate reliably under widely varying conditions, including rapid acceleration, regenerative braking, fast charging, high and low ambient temperatures, vibration, and long-term cycling. Because cells can exhibit variations in capacity, internal resistance, temperature response, and aging characteristics, effective monitoring and management at cell, module, and pack levels are

essential. The hierarchical architecture of the battery therefore provides the physical foundation upon which the BMS performs monitoring, estimation, control, and protection functions.

2.2 Core Functions of a BMS

The Battery Management System is the supervisory control system of an EV battery pack and is responsible for maintaining safe and efficient battery operation. It acquires electrical and thermal information from individual cells and the overall pack, processes this information, and initiates appropriate control or protection actions. The major functions of a BMS include parameter monitoring, cell balancing, thermal management, protection, state estimation, and communication with other vehicle control systems.

Voltage, Current, and Temperature Monitoring

Continuous monitoring of cell voltage, pack voltage, current, and temperature is one of the fundamental functions of a BMS. Cell-voltage measurements provide information about the electrical condition of individual cells and help identify abnormal conditions such as overcharging, over-discharging, and cell imbalance. Pack-level voltage and current measurements provide information about the overall electrical operating condition and are also used in estimating parameters such as State of Charge (SOC) and State of Health (SOH). Temperature monitoring is equally important because battery performance, aging, charging capability, and safety are strongly influenced by temperature. Sensors distributed across the battery pack allow the BMS to identify temperature variations and detect abnormal thermal conditions. Accurate and timely measurements provide the data required for subsequent estimation, diagnosis, control, and protection functions.


Cell Balancing

Cell balancing is required because individual cells within a battery pack do not remain identical throughout their operating life. Manufacturing variations, differences in temperature, charging history, and non-uniform degradation can cause cells to develop different SOC levels and electrical characteristics. If these differences become significant, the weakest or most highly charged cell can limit the usable capacity of the entire pack and may increase safety risks. A BMS therefore implements cell-balancing strategies to maintain the cells within an acceptable range of voltage or SOC. In passive balancing, excess energy from higher-voltage cells is dissipated as heat through resistive elements. Active balancing, in contrast, transfers energy from cells with higher SOC to cells with lower SOC using electronic power-conversion circuits. Effective balancing improves the usable capacity, efficiency, consistency, and lifetime of the battery pack.

Thermal Management

Battery temperature has a significant influence on electrochemical performance, charging efficiency, degradation rate, and safety. Excessively high temperatures can accelerate aging and, under severe conditions, contribute to thermal runaway, whereas very low temperatures can reduce power capability and charging acceptance. The BMS therefore works with the vehicle's thermal-management system to maintain the battery within an appropriate operating-temperature

range. Depending on the battery design and vehicle architecture, thermal control may involve air cooling, liquid cooling, refrigerant-based systems, or heating elements. The BMS monitors temperature distributions and can regulate cooling or heating based on operating conditions. Advanced thermal-management strategies can further use predictive models to anticipate temperature changes during high-power driving, fast charging, or extreme environmental conditions.



Electric Vehicle Battery Pack Management

Protection and Communication

Battery protection is a safety-critical responsibility of the BMS. It continuously evaluates battery operating conditions and initiates protective actions when predefined limits are exceeded. These actions can include limiting charging or discharging current, disconnecting the battery from the vehicle through contactors, and generating warnings or fault signals. Protection functions typically address overvoltage, undervoltage, overcurrent, short circuits, excessive temperature, insulation faults, and other abnormal operating conditions. Early detection of faults is particularly important for preventing permanent battery damage and reducing the possibility of hazardous events.

In addition to monitoring and protection, the BMS must communicate battery information with other electronic systems within the vehicle. It exchanges information such as SOC, SOH, available power, charging status, temperature, warnings, and fault conditions with the vehicle control unit, charger, thermal-management controller, and other relevant systems. Communication is commonly implemented through automotive communication networks such as the Controller Area Network (CAN). Reliable communication enables coordinated energy management and ensures that battery operating limits are considered by the vehicle's propulsion and charging systems. In AI-enabled BMS architectures, this communication infrastructure also

provides an important pathway for collecting operational data that can be used for intelligent estimation, prediction, diagnostics, and optimization.

3. Artificial Intelligence in Battery Management

3.1 AI Techniques for BMS

Machine Learning


Machine Learning (ML) enables computational models to learn patterns and relationships from battery data without requiring every physical relationship to be explicitly defined. In BMS applications, supervised, unsupervised, and reinforcement learning approaches can be employed for different purposes. Supervised learning methods, including regression algorithms, decision trees, support vector machines, and ensemble learning techniques, can be trained using experimental data with known battery states or health conditions. These models can subsequently estimate parameters such as SOC and SOH from measured voltage, current, and temperature data.

Unsupervised learning can be useful when labelled battery-fault or degradation data are limited. Clustering and anomaly-detection methods can identify unusual patterns in battery behavior and potentially reveal early indications of abnormal operation. Reinforcement learning, meanwhile, can be applied to sequential decision-making problems such as charging control and energy management, where an algorithm learns an appropriate control strategy through interaction with a battery or simulated environment. The flexibility of ML makes it suitable for applications in which battery behavior is difficult to capture using conventional analytical models.

Deep Learning

Deep Learning (DL) is a subset of machine learning that uses multilayer neural networks to learn complex representations from data. Deep-learning architectures such as Artificial Neural Networks (ANNs), Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM) networks, and Transformer-based models can be applied to battery-management problems. These approaches are particularly useful for processing large, high-dimensional, and time-dependent datasets.

Battery signals are inherently sequential because the present battery state depends not only on instantaneous measurements but also on previous operating history. Recurrent architectures such as LSTM networks can capture temporal dependencies in voltage, current, and temperature profiles, making them suitable for SOC and SOH estimation as well as degradation prediction. CNNs can extract meaningful features from structured signal representations, while more recent attention-based architectures can model long-range dependencies in battery operating data. Despite their predictive capabilities, deep-learning models generally require substantial and representative training data and may involve higher computational requirements. Their limited interpretability is also an important consideration when they are deployed in safety-critical BMS applications.



AI techniques for BMS

Data-Driven Battery Modelling

Data-driven battery modelling uses measured battery data to construct computational models that represent battery behavior without relying entirely on detailed electrochemical equations. Such models can learn the relationship between measurable inputs—such as current, voltage, temperature, and charging history—and outputs such as SOC, SOH, power capability, or degradation rate. This approach can complement traditional equivalent-circuit and electrochemical models, particularly when the underlying battery behavior is highly nonlinear or difficult to characterize analytically.


Data-driven models can be developed using laboratory experiments, battery cycling tests, simulation-generated data, or operational data collected from EVs. Their performance depends strongly on the diversity and quality of the training dataset. A model trained under a narrow range of temperatures, current rates, or aging conditions may not generalize effectively to different real-world conditions. Therefore, robust data collection, preprocessing, feature selection, model validation, and uncertainty assessment are essential components of AI-based battery modelling. Hybrid approaches that combine physical battery models with machine-learning techniques are also gaining attention because they can incorporate physical knowledge while benefiting from the adaptability of data-driven methods.

3.2 AI Applications

State of Charge (SOC) Estimation

State of Charge represents the remaining charge available in a battery relative to its usable capacity and is commonly expressed as a percentage. Accurate SOC estimation is essential for determining the remaining driving range, controlling charging and discharging, and preventing operation outside safe battery limits. Direct measurement of SOC is not possible using a simple sensor; instead, it must be estimated from measurable quantities such as voltage, current, and temperature.

Conventional SOC estimation methods include Coulomb counting, open-circuit-voltage methods, and model-based observers. Although these methods can provide useful estimates, their accuracy can be affected by sensor errors, temperature variations, battery aging, and uncertainties in battery parameters. AI models can learn the nonlinear relationship between measured battery signals and SOC from experimental data. Machine-learning and deep-learning techniques can therefore provide accurate SOC estimates under dynamic operating conditions. Combining AI algorithms with physical or equivalent-circuit models can further improve robustness and reduce the dependence on large training datasets.



AI Applications in Battery Management

State of Health (SOH) Prediction

State of Health is an indicator of the degree of degradation of a battery relative to its initial or reference condition. It is commonly associated with changes in capacity, internal resistance, power capability, or other performance characteristics. Accurate SOH prediction is important for determining battery lifetime, maintaining vehicle reliability, scheduling maintenance, and assessing the residual value of battery packs.

Battery degradation is influenced by numerous interacting factors, including temperature, depth of discharge, charging rate, operating voltage range, calendar aging, and cycling history. These factors make SOH prediction a complex problem. AI algorithms can identify degradation-related patterns from historical voltage, current, temperature, and capacity data and use them to estimate the present health condition of the battery. Deep-learning models can further capture nonlinear relationships between operating history and degradation. Reliable SOH prediction enables the BMS to adapt operating strategies according to battery condition and can help prevent unexpected performance deterioration.

Remaining Useful Life (RUL) Estimation

Remaining Useful Life refers to the estimated period or number of operating cycles remaining before a battery reaches a predefined end-of-life criterion. RUL estimation extends beyond determining the current health of the battery by attempting to predict its future degradation trajectory. This capability is particularly valuable for EV fleet management, maintenance planning, warranty assessment, and second-life battery applications.

AI-based RUL estimation generally uses historical battery degradation data to learn the relationship between operating conditions and future health evolution. Time-series models, recurrent neural networks, survival-analysis approaches, and other machine-learning methods can be used to predict future capacity or resistance trends. The accuracy of RUL prediction depends on the availability of sufficiently long degradation histories and operating data that represent realistic usage conditions. Since battery aging is subject to considerable uncertainty, advanced AI systems should ideally provide not only a predicted RUL value but also an indication of prediction uncertainty. This can support more reliable maintenance and operational decisions.

Fault Diagnosis and Predictive Maintenance

Battery faults can arise from electrical, thermal, mechanical, manufacturing, or control-related causes. Examples include cell imbalance, sensor malfunction, overcurrent, abnormal temperature rise, internal short circuits, insulation faults, and abnormal degradation. Early detection is essential because some faults can progress rapidly and may result in reduced performance, permanent battery damage, or serious safety hazards.

AI-based fault-diagnosis systems can analyze patterns in voltage, current, temperature, and other sensor measurements to identify abnormal operating conditions. Classification algorithms can distinguish between different fault types, while anomaly-detection techniques can identify deviations from normal battery behavior even when labelled fault data are scarce. Deep-learning approaches can also extract complex fault signatures from large volumes of time-series data. When combined with predictive maintenance strategies, these systems can identify emerging problems before they result in significant performance loss or system failure.

Predictive maintenance represents a shift from conventional maintenance based on fixed schedules or post-failure intervention toward condition-based decision-making. By continuously assessing battery health and detecting early degradation or abnormal behavior, an AI-enabled BMS can provide timely maintenance recommendations and support safer battery operation. In large EV fleets, such predictive capabilities can reduce unexpected downtime, optimize maintenance resources, and extend the useful life of battery systems. Ultimately, the integration of AI-based diagnosis and prediction can transform the BMS from a system primarily focused on monitoring and protection into an intelligent platform capable of prediction, early intervention, and adaptive battery management.


4. Functional Safety in AI-Enabled BMS

The integration of Artificial Intelligence into Battery Management Systems introduces significant opportunities for improving battery performance, prediction, and fault management. However, because the BMS is a safety-critical system, AI-based functions must operate reliably even under abnormal conditions, sensor failures, communication errors, and unexpected battery behavior. Functional safety therefore becomes a fundamental consideration in the design and deployment of AI-enabled BMS architectures. It involves identifying potential hazards, assessing associated risks, implementing appropriate safety mechanisms, and ensuring that the system can transition to a safe condition when faults occur. In an AI-enabled BMS, functional safety must address not only conventional hardware and software failures but also uncertainties associated with data, trained models, prediction errors, and changing operating environments.

4.1 Importance of Functional Safety

Battery Hazards and Risk Mitigation

EV batteries store large amounts of electrical and chemical energy, making their safe operation a critical engineering requirement. Potential hazards include overcharging, over-discharging, overcurrent, short circuits, excessive temperature, cell imbalance, insulation failure, mechanical damage, and internal cell faults. Under severe conditions, these failures may result in rapid temperature increase, gas generation, fire, or thermal runaway. Battery hazards can also arise from interactions among multiple cells, particularly when localized faults propagate through a module or pack.



Enhancing EV Battery Safety

A BMS serves as one of the primary layers of protection against these hazards. It continuously monitors battery parameters and compares them with predefined operating limits. When an abnormal condition is detected, the system can generate warnings, reduce charging or discharge power, activate thermal-management mechanisms, or disconnect the battery from the vehicle. AI can further strengthen this protection by identifying subtle patterns associated with emerging faults and providing early warnings before conventional threshold-based protection is triggered. However, AI predictions should not be considered the sole means of ensuring battery safety. A robust architecture should combine AI-based prediction with deterministic safety mechanisms, physical protection devices, validated control logic, and independent monitoring. This layered approach is particularly important because AI models may occasionally produce incorrect predictions when exposed to operating conditions that were not adequately represented in their

training data. Functional safety therefore requires the BMS to remain safe even when an intelligent component fails or produces an uncertain output.

4.2 Key Functional Safety Standards

Standards provide systematic frameworks for identifying hazards, allocating safety requirements, developing safety-related hardware and software, and validating automotive and electrical systems. Several standards are relevant to EV battery systems and AI-enabled BMS architectures.

ISO 26262

ISO 26262 is a major international standard addressing functional safety for electrical and electronic systems in road vehicles. It provides a lifecycle-based framework covering hazard analysis, risk assessment, safety requirements, system design, hardware and software development, verification, validation, and production-related activities. The standard uses Automotive Safety Integrity Levels (ASILs) to classify safety requirements according to the potential severity, exposure, and controllability of hazardous events.

For an AI-enabled BMS, ISO 26262 provides an important framework for determining the safety requirements of sensors, electronic control units, communication interfaces, software functions, and protective mechanisms. AI-based functions that contribute to safety-related decisions must be integrated into an architecture that appropriately addresses systematic and random failures. The use of AI does not eliminate the need for conventional safety engineering; instead, it requires careful consideration of how AI components interact with safety-critical functions.

IEC 61508

IEC 61508 is a foundational international standard for the functional safety of electrical, electronic, and programmable electronic safety-related systems. Although it is not specific to electric vehicles, its principles have influenced functional-safety practices across several industries. The standard emphasizes systematic risk assessment, safety lifecycle management, fault tolerance, verification, validation, and the achievement of an appropriate Safety Integrity Level (SIL).

For battery-management applications, the concepts established by IEC 61508 can provide useful guidance for developing safety-related electronic and software functions. It is particularly relevant when considering the reliability of sensing, control, diagnostic, and protection mechanisms that may be incorporated into battery systems and associated infrastructure.

ISO 6469 – EV Safety

The ISO 6469 series addresses safety specifications for electrically propelled road vehicles. It covers important aspects of electric-vehicle safety, including protection against electrical hazards and considerations associated with rechargeable energy-storage systems. These requirements are particularly relevant to battery packs because high-voltage electrical systems introduce risks that differ from those of conventional low-voltage automotive systems.

For an AI-enabled BMS, the principles associated with EV safety provide a broader system-level context within which battery monitoring, isolation, protection, and control functions must

operate. Compliance with relevant EV safety requirements helps ensure that intelligent battery-management functions are integrated into a vehicle architecture that adequately addresses electrical and operational hazards.

	 Focus Area	 Relevance to BMS
ISO 26262	Functional safety for electrical and electronic systems in road vehicles	Determining safety requirements for sensors, control units, and AI-based functions
IEC 61508	Functional safety of electrical, electronic, and programmable electronic systems	Guidance for reliability of sensing, control, diagnostic, and protection mechanisms
ISO 6469	Safety specifications for electrically propelled road vehicles	Context for battery monitoring, isolation, protection, and control functions
ISO/SAE 21434	Cybersecurity risk management for road-vehicle electrical and electronic systems	Securing BMS communication, software components, and data exchange

Key Functional Saafety Standards for EV Battery Systems

ISO/SAE 21434 – Cybersecurity Overview

As modern EVs become increasingly connected and software-defined, cybersecurity has become an important aspect of battery-system safety. ISO/SAE 21434 provides a framework for managing cybersecurity risks throughout the lifecycle of road-vehicle electrical and electronic systems. This is particularly relevant to AI-enabled BMSs because they may exchange data with vehicle controllers, charging infrastructure, cloud platforms, and remote diagnostic systems.

A compromised BMS communication channel or software component could potentially result in incorrect battery-control commands, manipulated sensor information, unauthorized access, or disruption of charging and energy-management functions. Consequently, cybersecurity should be considered alongside functional safety. Secure communication, authentication, access control, software-update mechanisms, data protection, and continuous monitoring can help reduce cyber-related risks. In this context, cybersecurity and functional safety are complementary: protecting the system from malicious interference contributes to maintaining safe battery operation.

4.3 Safety Mechanisms

Fault Detection

Fault detection is a fundamental safety mechanism in a BMS. The system continuously evaluates sensor measurements and operational parameters to identify abnormal conditions. Conventional

methods use thresholds, limit checks, consistency checks, and model-based residual analysis. AI-based methods can complement these techniques by recognizing complex patterns associated with early-stage faults.

For example, an AI model may identify an unusual combination of voltage, temperature, and current behavior that does not independently exceed a predefined threshold but may indicate an emerging internal fault. Combining conventional and AI-based diagnostics can therefore provide both deterministic protection and advanced early-warning capabilities.


Redundancy

Redundancy involves providing additional sensing, processing, communication, or protection pathways so that the failure of one component does not necessarily result in loss of the safety function. Redundant temperature or voltage measurements, independent monitoring circuits, and duplicated safety-critical processing can improve system reliability.

In AI-enabled BMS architectures, redundancy can also be achieved through diverse computational approaches. For example, an AI-based SOC estimator may be cross-checked against a conventional model-based estimator. If their outputs differ significantly, the system can identify an abnormal condition and revert to a conservative operating strategy. Such diversity can reduce dependence on a single algorithm and provide an additional layer of fault tolerance.

Safe-State Operation

A safety-critical system should be capable of transitioning to a safe state when a serious fault is detected. In an EV battery system, a safe state may involve reducing charging or discharging power, disabling certain functions, activating thermal-management systems, or electrically isolating the battery from the vehicle using contactors.



Safety Mechanism in AI -Enabled BMS

The appropriate safe state depends on the type and severity of the fault. Importantly, an AI-enabled BMS should not continue making unrestricted control decisions when its predictions become unreliable or when critical sensors or communication channels fail. A well-designed architecture should include deterministic fallback strategies that allow the battery to remain

protected even when AI functionality is unavailable. This principle is essential for building confidence in intelligent battery-management technologies.


5. Emerging Trends and Sustainable Mobility

The evolution of BMS technology is moving beyond basic monitoring and protection toward intelligent, connected, predictive, and lifecycle-oriented battery management. Advances in AI, cloud computing, digital twins, wireless communication, and battery recycling are creating new opportunities to improve battery performance while reducing environmental impacts. These developments are particularly important as EV adoption increases and the number of batteries entering the transportation ecosystem grows.

5.1 AI-Driven Smart Charging

Conventional charging strategies generally follow predefined voltage, current, and temperature limits. AI-driven smart charging can make charging more adaptive by considering battery condition, temperature, charging history, electricity demand, energy prices, renewable-energy availability, and user requirements. Machine-learning algorithms can predict battery response and identify charging profiles that balance charging time, energy efficiency, degradation, and safety.

Smart charging can also support interaction between EVs and the electrical grid. When charging is coordinated with renewable-energy generation and periods of lower grid demand, EVs can become flexible energy-storage resources rather than simply electricity consumers. In advanced vehicle-to-grid (V2G) applications, intelligent algorithms can determine when energy should be stored in or supplied from the battery while considering battery degradation and user requirements. Such approaches can contribute to both grid flexibility and sustainable energy utilization.



Emerging Trends in Battery Management Technology

5.2 Digital Twin and Cloud-Based BMS

A digital twin is a virtual representation of a physical battery system that is continuously updated using operational data. A battery digital twin can combine physical models, sensor information,

historical data, and AI algorithms to represent the current and predicted condition of a battery. This enables applications such as virtual health monitoring, degradation prediction, fault diagnosis, and optimization.

Cloud-based BMS architectures can complement onboard battery-management functions by providing access to greater computational and data-storage resources. Operational information from multiple vehicles can be aggregated and analyzed to identify degradation patterns and improve predictive models. However, safety-critical functions should continue to have appropriate local processing and protection because cloud connectivity may be unavailable or introduce communication delays. A practical architecture can therefore combine local real-time safety management with cloud-assisted analytics and long-term optimization.

5.3 Wireless Battery Management Systems (WBMS)

Traditional BMS architectures rely extensively on wired connections between individual cells, sensors, modules, and the central control unit. Wireless Battery Management Systems (WBMS) replace or reduce some of these physical connections through wireless communication technologies. This can reduce wiring complexity, weight, assembly effort, and packaging constraints while potentially improving battery modularity.

Wireless architectures can also facilitate flexible battery-pack designs and simplify the integration of sensors. Nevertheless, WBMS introduces additional considerations related to communication reliability, electromagnetic interference, synchronization, cybersecurity, and power consumption. For safety-critical applications, wireless communication must therefore be supported by appropriate fault detection, communication monitoring, redundancy, and fallback mechanisms.

5.4 Battery Recycling and Second-Life Applications

The increasing deployment of EVs will eventually generate large quantities of batteries reaching the end of their first automotive life. However, reaching the end of automotive service does not necessarily mean that a battery has no further value. Batteries with sufficient remaining capacity and power capability can potentially be repurposed for second-life applications, including stationary energy storage, renewable-energy integration, backup power, and other less demanding applications.

Accurate assessment of residual battery health is essential for determining whether a battery is suitable for reuse, refurbishment, or recycling. AI-enabled BMSs can assist by estimating SOH, identifying degraded cells, predicting remaining useful life, and classifying batteries according to their residual performance. At the end of the useful life, recycling can recover valuable materials and reduce the demand for newly extracted resources. Intelligent battery diagnostics can therefore support more efficient decisions throughout the transition from first-life automotive use to second-life deployment and eventual material recovery.

5.5 AI for Sustainable Battery Lifecycle Management

AI has the potential to support battery sustainability across the entire lifecycle, from manufacturing and initial deployment to reuse and recycling. During operation, intelligent BMS

algorithms can optimize charging, thermal conditions, and power utilization to reduce unnecessary degradation and extend battery service life. Longer battery lifetimes can reduce the frequency of replacement and improve the utilization of the resources invested in battery manufacturing.

At the end of automotive use, AI-based health assessment can support decisions regarding reuse, remanufacturing, repurposing, or recycling. Data collected throughout the battery's operating history can provide valuable information about its degradation and usage profile, enabling more accurate evaluation of its residual value. Thus, AI-enabled battery lifecycle management can contribute to a circular battery economy, in which materials and energy-storage capabilities are utilized as efficiently as possible rather than being treated as disposable resources.

Conclusion

The rapid growth of electric mobility has increased the importance of reliable, safe, efficient, and sustainable battery systems. As the central supervisory component of an EV battery pack, the BMS performs essential functions including voltage, current, and temperature monitoring, cell balancing, thermal management, protection, communication, and battery-state estimation. However, the increasing complexity of battery behavior and the need for accurate prediction under diverse operating conditions have created opportunities for Artificial Intelligence to enhance conventional BMS architectures.

AI and machine-learning techniques can improve State of Charge estimation, State of Health prediction, Remaining Useful Life estimation, fault diagnosis, predictive maintenance, and charging optimization. These capabilities enable a shift from reactive battery management toward predictive and adaptive operation. At the same time, AI introduces challenges related to data quality, model generalization, computational requirements, explainability, cybersecurity, and safety assurance. Consequently, AI should be integrated with established engineering principles rather than treated as a replacement for deterministic safety mechanisms.

Functional safety is particularly important because failures in high-energy battery systems can have serious consequences. Standards and safety frameworks such as ISO 26262, IEC 61508, ISO 6469, and ISO/SAE 21434 provide important foundations for addressing functional and cybersecurity risks. Fault detection, redundancy, independent protection, and safe-state operation can provide additional layers of protection and help ensure that the battery remains safe even when intelligent components fail or produce uncertain outputs.

Future Outlook of AI-Enabled BMS

The future of AI-enabled BMSs is likely to involve increasingly integrated architectures combining physical battery models, machine learning, digital twins, edge computing, cloud analytics, wireless sensing, and connected charging infrastructure. Advances in explainable and physics-informed AI may help address concerns regarding model transparency and generalization, while improved sensing technologies and larger operational datasets can enable more accurate battery-state estimation and degradation prediction. Future BMSs may increasingly operate as intelligent, self-diagnosing systems capable of adapting their

management strategies according to battery condition, environmental conditions, user behavior, and grid requirements.

Role of Intelligent Battery Management in Safe and Sustainable Electric Mobility

Intelligent battery management has a broader significance than simply improving vehicle performance. By reducing degradation, improving energy utilization, detecting faults at an early stage, and extending battery service life, AI-enabled BMSs can contribute directly to the safety, economic viability, and environmental sustainability of electric mobility. Furthermore, their integration with smart charging, renewable-energy systems, second-life applications, and recycling pathways can help establish a more circular approach to battery utilization.

Ultimately, the combination of AI, battery technology, functional safety, and sustainable lifecycle management provides a pathway toward more reliable and resource-efficient electric transportation. AI-enabled BMS technology can therefore serve as a critical enabler of the transition from conventional battery monitoring systems to intelligent energy-management platforms capable of supporting the long-term vision of safe, connected, efficient, and sustainable electric mobility.

References

1. Shah, M., *et al.* (2021). Comprehensive review on smart techniques for estimation of state of health for battery management system application. *Energies*, 14(15), 4617.
2. Zhang, Y., *et al.* (2022). Deep learning in the state of charge estimation for Li-ion batteries of electric vehicles: A review. *Machines*, 10(10), 912.
3. Hu, X., Deng, X., Wang, F., Deng, Z., Lin, X., Teodorescu, R., & Pecht, M. G. (2022). A review of second-life lithium-ion batteries for stationary energy storage applications. *Proceedings of the IEEE*, 110(6), 735–753.
4. Jafari, S., & Byun, Y.-C. (2022). Prediction of the battery state using the digital twin framework based on the battery management system. *IEEE Access*, 10, 124685–124696.
5. Islam, S., & Ghazal, T. M. (2023). Lithium-ion battery management system for electric vehicles: Constraints, challenges, and recommendations. *Batteries*, 9(3), 152.
6. Bhar, M., Ghosh, S., Krishnamurthy, S., Kaliprasad, Y., & Martha, S. K. (2023). A review on spent lithium-ion battery recycling: From collection to black mass recovery. *RSC Sustainability*, 1, 1150–1167.
7. Chen, L., Xu, C., Bao, X., Lopes, A. M., Li, P., *et al.* (2023). State-of-health estimation of lithium-ion battery based on back-propagation neural network with adaptive hidden layer. *Neural Computing and Applications*, 35, 14169–14182.
8. Chen, L., Xie, S., Lopes, A. M., & Bao, X. (2023). A vision transformer-based deep neural network for state of health estimation of lithium-ion batteries. *International Journal of Electrical Power & Energy Systems*, 152, 109233.
9. Lohar, G., & Kumar, M. S. (2023). Estimation of state of charge of battery used in electric vehicles with wireless battery management system. *International Journal on Recent and Innovation Trends in Computing and Communication*, 11(7s), 745–756.

COMPARATIVE ANALYSIS OF MACHINE LEARNING MODELS FOR SHORT-RANGE WEATHER FORECASTING USING ERA5 REANALYSIS DATA OVER NORTHWEST INDIA

Richa Avasthy

School of Science and Engineering, GNA University, Phagwara, Punjab, India

Corresponding author E-mail: richa.avasthy@gnauniversity.edu.in

Abstract

Although it is an essential for agricultural planning, disaster preparedness, aviation safety and energy management in South Asia, computational constraint has so far prohibited numerical weather prediction (NWP) systems in the region from providing operational forecasts at reasonable expected lead time and there is always uncertainty associated with the models by fine resolution in space and time due to parameters. In this chapter, a systematic and reproducible comparison among different machine learning (ML) model families for forecasting 2-metre temperature (T2m) and total precipitation (TP) in Northwest India (Punjab, Haryana and parts of Himachal Pradesh) for the next-day time horizon is conducted for the 2000–2023 period with the ERA5 reanalysis data on a 0.25° resolution. The preprocessing pipeline was structured by spatially aggregating the data, removing outliers, normalizing it, creating lags for the features, and giving it a sequence; this pipeline was applied to all three architectures, ensuring a fair comparison scenario. During the skipped 2021–2023 test period, the two-layer LSTM was the most accurate model on the test set with temperature RMSE of 1.61 °C that is 58% higher than the temperature RMSE of the persistence baseline, and demonstrated a strong correlation to observations ($R = 0.96$), proving to be the most accurate one, while the Random Forest model showed good compromise between accuracy and interpretability with high feature importance according to the Gini index. The prediction of precipitation was still challenging for all the architectures, with the best one (LSTM) achieving an RMSE of 4.3 mm/day, due to the intermittent and spatially heterogenic character of precipitation. The most important physical factors in this analysis, from the perspective of feature importance analysis, were 500 hPa geopotential height, which corresponds with the well-known synoptic climatology of the region. The chapter ends by discussing implications for operation, agriculture and disaster management from the results, the current and future limitations in the ability to develop accurate grid-scale ERA5-driven ML forecasting, and hybrid physics-ML systems for the Indian sub-continent.

Keywords: Machine Learning, Deep Learning, Weather Forecasting, LSTM, Random Forest, Convolutional Neural Network, ERA5 Reanalysis, Numerical Weather Prediction, Northwest India, Time-Series Prediction, Precipitation Forecasting, Climate Informatics.

1. Introduction

Within a few hours to three days in advance is key to many sectors that impact the lives of billions of people in South Asia, including aviation, energy planning, public health,

transportation and disaster management and agriculture. The improvements in accuracy of localized, short-range forecasts that benefit irrigating farmers in an agriculture-based economy like India can lead to significant savings in irrigation planning, sowing and harvesting decisions and the losses from unseasonal rainfall, hail storms and heatwaves.

Classical numerical weather prediction systems like the Integrated Forecasting System (IFS) of European Centre for Medium-Range Weather Forecasts (ECMWF) use the solutions of coupled partial differential equations for the dynamics, heat and moisture motion of the atmosphere over discrete global grids. Operational meteorology has an essential need for these systems, which are physically demanding, but which also require vast computing resources and a high-performance computing infrastructure, as well as top domain knowledge to set them up, operate and understand, and for which ECMWF has specific technical manuals. Although models are continually refined, the sub-grid scale is not directly resolved in NWP models at typical operational resolution (9 to 31 km) and introduces a non-negligible level of forecast uncertainty in models, as investigated in the meteorological literature over decades of the evolution of numerical prediction, in the form of parameterization schemes — simplified representations of physical sub-grid phenomena (convection, cloud microphysics, boundary-layer turbulence, etc.). Over the last decade, a second paradigm, complementing the traditional modelling approach, has been developed; one relying on machine learning (ML) to extract statistical relationships directly from historical datasets of models and observations from reanalysis. Unlike forward models for the governing equations of atmospheric physics, ML models are used to learn which spatio-temporal patterns are likely to be followed by a given weather outcome. It has several advantages over conventional NPWs: it entails minimal computational overhead for the models to make predictions, and it is possible to both train and run the models on the more modest resources of institutions, without requiring top-tier hardware; it may also be understandable enough to produce prediction results that are competitive with, or occasionally better than, traditional NPW baselines — in many regional or short-range contexts, at least, as suggested in influential perspectives on data-driven Earth system science. Their back-to-back attempts to outdo ECMWF's operational high-resolution forecast (HRES) — one using data from the GPCs' ERA5 re-analysis — have sparked what many in the industry label a 'paradigm shift' in operational forecasting philosophy with their impressive performance on most of the typical verification metrics for medium range forecasting periods.

Given this global developmental star, however, localized, sub-national studies in this region of the country, and more specifically the agriculturally important semi-arid region of North West India that also has distinct climatic conditions, are relatively few. The previous study is regionally-oriented and usually demonstrated only a single model family (such as LSTM) trained on a particular climatic variable (such as monsoon rainfall over peninsular India) without a proper, multi-architecture evaluation against the synoptic conditions of this specified sub-region. This gap is the impetus for the present study.

This chapter does a systematic comparative study of three widely used families of ML models such as Random Forest (RF), Long Short-Term Memory (LSTM) networks and 1-Dimension Convolutional Neural Networks (CNN) on short-range (24 hours lead time) temperature and precipitation forecast cases over the state of Punjab, Haryana and parts of Himachal Pradesh using ERA5 re-analysis data. These specific contributions are as follows: (1) A structured and reproducible framework with three distinct ML architectures applied to a region-specific ERA5 subset over a 24-year period, (2) documented processing/conversion pipeline of ERA5 data and preprocessing/feature engineering pipeline for atmospheric timeseries, (3) an empirical assessment based on skill scores of common skill measures (RMSE, MAE and Pearson correlation) against the persistence baseline and benchmarking results, (4) an interpretability analysis using Gini-based feature importance to understand the influence of predictors which play the dominant role in the Indian sub-continent, and (5) a discussion of the feasibility of using ML methodology in the Indian sub-continent for regional forecasting, where the challenges are described as well as future prospects and research directions.

2. Literature Review

2.1 Physics-Based Forecasting and the Case for Data-Driven Complements

Since the mid-twentieth century, the operational paradigm has been the main approach to forecasting, while its continuous and incremental development has been termed a 'quiet revolution' in prediction science in the meteorological literature. The quality of NWP products, however, is always limited by the resolution of the model and how the remaining sub-grid processes are represented with parameterization schemes. The study published in Proceedings of the National Academy of Sciences concerns an important theoretical back up for hybridization of physics and ML models: deep neural networks were trained on the high-resolution cloud-resolving model output and were shown to be able to simulate sub-grid convective parameterization schemes with a high degree of fidelity. The strength of the prior machine-learning perspective on Earth system science, in a strongly influential Nature perspective, was that rather than supplanting physical-modeling, the most compelling direction in which to advance Earth system prediction would be to engage with the two together, with machine learning working in concert with domain physical knowledge.

In this section, I discuss recurrent and convolutional architectures for spatiotemporal forecasting. In this section, the recurrent and convolutional architectures for spatiotemporal forecasting are discussed.

2.2 Recurrent and Convolutional Architectures for Spatio-Temporal Forecasting

Recurrent architectures have been especially successful for sequences of meteorological data, where long-range time dependencies hold greater significance. The original LSTM architecture was introduced in the neural computation literature at the end of the 1990s to tackle the vanishing gradient problem experienced by previous recurrent neural networks, which had the ability to persist or forget information over long input sequences using 'gated memory cells'. On this basis, a popular study in 2015 developed the Convolutional LSTM (ConvLSTM), which

showed that jointly modelling spatial and sequential structure makes a significant contribution to improving skill in precipitation nowcasting when compared to architectures that only consider space or time alone. Simultaneously, convolutional networks built in the late 1990s for image and document recognition have been successfully adapted to one-dimensional time-series forecasting, using convolutional filters as automatic learnable feature extractors that relate to a sliding window during the temporal dimension (analogous to extracting features from a spatial window during image recognition).

2.3 Global-Scale Deep Learning Weather Models

Numerous recent studies show that the large deep learning models can be trained directly on reanalysis data and are effective at changing their worldview. GraphCast from Google DeepMind, a graph neural network that was trained using ERA5 data, was demonstrated to outperform the ECMWF's operational HRES forecast over 90% of 1,380 targets, over a 10-day lead time with only a fraction of computing compared to a traditional NWP run after training. Predicted respectively in *Science* and *Nature*, the model, dubbed Pangu-Weather, based on Transformers developed by Huawei, while being trained on ERA5 data over 39 years, is also demonstrated to result in state-of-the-art deterministic medium-range forecasting skill, and, remarkably, in forecasts up to three orders of magnitude more efficient than traditionally required the inference time of conventional NWP. Together, these studies gave a sense of ERA5 as the de facto preferred reference dataset for training data-driven models of the atmosphere, and revealed the central message of the approach used throughout this chapter: the reanalysis-driven method could potentially realistically approximate the dynamics of the atmosphere at operationally useful lead times by means of data-driven models trained on this reference data.

2.4 Regional and Indian Sub-Continental Applications

Avoid adding too many services to the VMware D2L course. Try not add too many services to the VMware D2L course. Applications of ML forecasting in the Indian sub-continent is comparatively limited and with few applications in scope. In 2021, an encouraging study published in *Journal of Earth System Science* used LSTM networks for monsoon rainfall prediction over the homogeneous rainfall regions of India, but the evaluation did not include the different semi-arid climate of Northwest India, which is influenced by heightened diurnal temperature ranges, topographic effects from the sub-Himalaya foothills, and a late arrival of the monsoon signal with increased inter-annual variability when compared with the monsoons of the Peninsular and other homogeneous regions of India. To the best of the author's knowledge, no prior study has designed a baseline set of agro-climatic features derived from GPC ERA5 data and apply the same set of features and evaluation protocol for controlled multiple architectures comparison of the RF, LSTM, and CNN models which have been specifically customized for the agro-climatic zone of Punjab–Haryana–Himachal.

The present study is aimed at filling the gaps in the literature and positioning itself in relation to previously conducted studies.

2.5 Research Gap and Positioning of the Present Study

There are three gap areas that arise from the review. First, most regional studies assess a single family of ML algorithms and prevent a fair and controlled comparison of advantages and drawbacks of different model families for a given region and set of target variables. Secondly, few studies explicitly report interpretability results (e.g., the use of feature importance) in addition to existing accuracy metrics, which makes it difficult to gain confidence of operational forecasters and agricultural extension agencies. Third, while accuracy metrics are often mentioned, parameters that impact the cost of computation (and training time) — of vital importance for institutions who may lack GPU infrastructure — are not always reported. The present study aims to close this three-part gap by implementing RF, LSTM and CNN with a specific focus on the Northwest India domain, calibrated using freely available ERA5 data, providing accuracy, interpretability and computational cost metrics, and to draw regional baselines for future operational systems that can leverage this region to achieve vital agricultural objectives.

3. DataSet, Study Area and Preprocessing methodology

3.1 Study Area

The study domain (28°N–34°N, 73°E–78°E) covers the entire lengthspan of Punjab and Haryana states and plains and lower elevation sub-Himalayan regions of Himachal Pradesh. The region is considered the 'grain basket' of India, as wheat and rice procurement relies mainly on this region, and forms part of the Indo-Gangetic Plain, which is one of the world's most cultivated and irrigated terraines. The climate is semi-arid to sub-humid with hot, dry summers (from April to June) followed by a south-west monsoon season from July until September during which the maximum annual rainfall occurs; a post monsoon phase and a cool humid winter (December to February) characterised by western disturbances from the Mediterranean with arrival along the western passage along the sub-tropical northward-flowing jet stream. This region characterised by large seasonal contrast, mountain influence in the foothills of the Himalayas and strong sensitivity of local agriculture to the accuracy of short-range forecasts is a particularly good testbed for regional ML research for forecasting.

3.2 ERA5 Reanalysis Dataset

The ERA5 reanalysis product created by the European Centre for Medium Range Weather Forecasts (ECMWF) and provided by the Copernicus Climate Change Service (C3S) is used in this study. ERA5 is the fifth-generation, and current gold standard, of global atmospheric reanalysis, which incorporates observational data from surface stations, aircraft, radiosondes and satellites into a physically consistent reanalysis by assimilating them into a 4D-VAR framework at about 31 km horizontal resolution and 1 hour temporal frequency, as described in the basic reanalysis documentation from ECMWF published in the Quarterly Journal of the Royal Meteorological Society. Table 1 summarizes the dataset characteristics and experimental configuration used in this study.

Table 1: ERA5 Dataset Characteristics and Experimental Configuration

Attribute	Value
Dataset	ERA5 Reanalysis (ECMWF / Copernicus C3S)
Spatial domain	28°N–34°N, 73°E–78°E (Northwest India)
Temporal coverage	January 2000 – December 2023 (24 years)
Temporal resolution	Hourly (aggregated to daily means)
Horizontal resolution	0.25° × 0.25° (~31 km)
Target variables	2-m temperature (T2m); Total precipitation (TP)
Predictor variables	500 hPa geopotential height (Z500); sea-level pressure (SLP); 10-m wind (U10, V10); specific humidity (Q)
Train / Validation / Test split	2000–2018 / 2019–2020 / 2021–2023

Feature Engineering and Preprocessing Pipeline

The three types of models have been preprocessed in a controlled and reproducible way which consisted of the following five sequential operations:

- P1 - Temporal aggregation: temporal aggregation was applied to hourly ERA5 fields to obtain daily means, to match the 24 hours forecast lead time applied across this study.
- P2 - Quality control, records with quality control flags provided by ERA5 masks or with physically implausible sensor values were excluded, resulting in the loss of about 0.3% of the records.
- P3 - Normalization: all predictor variables were first normalized to zero mean and unit variance, so as the statistics are computed fully within the dataset used for training, there would be no data leakage from the validation or test split data sets.
- P4 - For the Random Forest model, which has no built-in sequential order, the temporal dependencies were encoded into the model as input features using 7-day, 14-day and 30-day sliding windows.
- P5 - Sequence construction: input sequences consisting of the fixed length of 30-time steps were built ahead for LSTM and CNN models, and each sequence would predict the value of the target variable at t+1 day.

By using this common pipe, any differences in performance between the different architectures that are observed is real and not an artefact of the different data preparation processes. The machine learning model architectures and experiments is presented. The machine learning models architectures and experiments are presented.

4.1 Random Forest (RF)

Random Forest builds N decision trees from random samples of the training set (where N could be 100, 1000 or any number), which are averaged to minimize the variance of the individual trees and thus maximize generalisability, in line with the original formulation as an ensemble-learning function published in the Machine Learning Journal in 2001. In this study, N = 500 trees

were used, the maximum depth was set to 20-, and 5-folds cross validation on the training set was used to tune the number of samples per leaf (between 5 and 10). For tabular, lag-featured input and with a natural feature importance estimate useful for meteorological interpretation, RF has proved effective, while missing the intrinsic sense of sequential order other than the ones provided with the designed lag features.

4.2 Long Short-Term Memory Network (LSTM).

As for solving the vanishing gradient problem, the LSTM Network has gated memory cells to selectively retain or forget content across time steps, as described in the original paper titled “Neural Computation” which introduces a new network. The design in this study is based on the LSTM model comprising two layers: a 128 units layer and a 64 units layer followed by a fully connected layer at the output. To regularize with each LSTM layer, dropout was applied, with rate = 0.2. The Adam optimizer (with a learning rate of 1×10^{-3}) was adopted in training, and early stopping is applied to the validation loss with a patience of 15 epochs.

4.3 Convolutional Neural Network (CNN)

The one-dimensional CNN was applied to the temporal series of the atmospheric predictor variables, with the convolutional filters capturing local temporal patterns, similar to the way that 2D CNNs work on images to detect spatial features, which was introduced in the influential 1998 Proceedings of the IEEE paper on "Gradient based document recognition. Each convolutional block is composed of 64, 128 and 64 convolutional filters with kernel size 3, followed by batch normalization and ReLU activation, and has global average pooling before the output dense block. This type of architecture has been shown to be competitive in the previous literature in both univariate and multivariate meteorological time-series tasks.

4.4 Evaluation Metrics

The skill of the models was evaluated using three commonly used meteorological metrics that compared the forecasted ($\hat{y}_i$) and observed (y_i) values for the held-out test period: the Root Mean Squared Error (RMSE) defined as the square root of the mean square error of a model's performance, $RMSE = \sqrt{[(1/n)\Sigma(y_i - \hat{y}_i)^2]}$ and the Mean Absolute Error (MAE), defined as the average error of the model, $MAE = (1/n)\Sigma|y_i - \hat{y}_i|$; and the Pearson correlation coefficient, R , that measures the linear association between two series, in this case the forecasted and observed values. All three of the ML models were also compared to a persistence baseline — that is, the value of the next day is always predicted as the same as the previous day's value — which is a standard, minimally informative benchmark for verification of meteorological models.

A. Representation and knowledge extraction.A. Representation and Knowledge extraction.

Since the focus was to isolate the effects of architecture and not to compare effectiveness of various sets of predictors, all models were trained at the same train / validation / test split (2000–2018 / 2019–2020 / 2021–2023) and using identical sets of predictors. The training was done using only one of the workstation-class GPUs and wall-clock training times are summarized together with the accuracy metrics in Table 2 to guide practical decisions for institutions with

fairly limited amount of computing power, e.g. state agricultural universities and local disaster-management units.

5. Results and Model Evaluation

5.1 Temperature Forecasting Performance (T2m)

The model was assessed against the held-out test period (2021–2023) metric to assess its accuracy of model performance, based on the comparison of the daily 2-m temperature predicted by the model with the observed temperature was evaluated via RMSE, MAE and Pearson correlation (R). In all three ML models, the score on the test sets, as measured by the persistence baseline, was significantly better than the naively day-to-day correlated pattern, thus demonstrating the added predictive power of learned atmospheric patterns versus that of naive day-to-day autocorrelation.

Table 2 shows the temperature forecasting performance on test dataset (2021-2023) with best ML result marked in bold (LSTM).

Table 2: Temperature Forecasting Performance on Test Set (2021–2023) — Best ML result highlighted in bold (LSTM)

Model	RMSE (°C)	MAE (°C)	Correlation (R)	Training Time
Persistence Baseline	3.84	2.91	0.71	—
Random Forest (RF)	1.92	1.41	0.93	~4 min
LSTM (2-layer)	1.61	1.18	0.96	~38 min
CNN (1-D, 3-block)	1.74	1.29	0.95	~22 min

5.2 Precipitation Forecasting Performance (TP)

As was found in much of the existing literature on precipitation nowcast, forecasting it was much more difficult than temperature forecasting for all three architectures, due to its intermittent nature, the necessity of forecasting the presence or absence of precipitation, and its spatial heterogeneity. Compared to the other two models, LSTM performed the best for precipitation skills both in terms of RMSE and R value, with LSTM being RMSE 4.3 mm/day and R 0.79, whereas RF was RMSE 5.1 mm/day and R 0.73, while CNN was RMSE 4.8 mm/day and R= 0.76. The biases in extreme events were highest in the monsoon months (June–September), a bias still under investigation in the ML literature on weather forecasting, and because of this is a serious research concern in this study as well.

To analyze feature importance, we will employ the technique of Random Forest. Random Forest will be used to analyze feature importance.

5.3 Feature Importance Analysis (Random Forest)

The feature importance of the RF model was performed based on the Gini index with 500 hPa geopotential height (Z500, 31%) being the top predictor, followed by the 7-day lag of T2m (24%), SLP (18%) and specific humidity (14%). It is physically interpretable: Z500 represents the large-scale patterns of northwestern dynamics like ridge and trough, and it compares well

with the synoptic climatology of the area, such as the effect of western disturbances in winter and pre-monsoon development of the thermal lows in summer.

5.4 Seasonal Error Decomposition

When forecasting errors are disaggregated during the season, then systematic forecasting errors are highlighted. Skill at predicting the temperature is greatest during the relatively "stable" winter season (December - February) when the synoptic scale forcing dominates and the short-term variations are relatively small. During the pre-monsoon season (April-June), in the presence of strong diurnal heating and localized convective activity, the variability on smaller scales than those resolved by ERA5 (0.25°) becomes, although not significantly, more pronounced. The greatest degradation in both temperature and precipitation is found during the monsoon season when weather processes are predominantly moist convective processes and mesoscale organization of features – components not entirely described by daily-aggregated coarse-resolution features derived from reanalysis.

More precisely, computational efficiency and deployment considerations:

Accuracies obtained from LSTM are about 16% improved over RF per RMSE difference in temperature, with a training time of about 9.5 times longer in terms of computational time, from a deployment point of view. For institutions with bigger GPU infrastructure, like district agricultural offices or state disaster management cells, RF might be a more suitable operational solution because of its close correspondence with persistence, at a much lower training cost and having a feature importance that is both easily interpretable and directly verifiable by domain meteorologists based on physical knowledge. CNN falls somewhere in between which is more consistent with LSTM in accuracy and about 40% cheaper to train.

6. Discussion

The better performance of LSTM over RF and CNN for both target values simply highlights the fact that static lag-feature methods like RF miss the ability of the architecture to model long-range temporal dependencies found in the temporal sequence of the atmospheric data; the RF treats each lagged predictor as a tabular one rather than a predictor drawn from an ordered sequence with learned memory. The trade-off to this, however, is an important and practical one: No GPU infrastructure is required, it takes mere minutes, rather than hours, to train; the output features scores that can be directly translated into meaningful, domain-specific features of practical relevance (such as feature-importance scores) that domain meteorologists and agricultural extension officers can validate against physical understanding, an important factor if the ML-based forecasting tools are to be trusted by non-specialist end users.

The poor performance of all three architectures against precipitation compared to their performance against temperature is not particularly surprising, as temperature is a smoothly varying, continuous field in which all days are quite mutually dependent in terms of variability, while precipitation is intermittent and strongly influenced by sub-grid-scale convective processes which are poorly resolved even on a 0.25° reanalysis grid length. This result is compatible with the overall precipitation-nowcasting literature (which is dominated by the ConvLSTM, as

mentioned in the previous paragraph) and suggests that a hybrid scheme—a blend of coarse-scale ML pattern recognition and fine-scale physical constraints or observations—is generally needed to significantly enhance operational precipitation forecast skill in such regions and for similar purposes.

This is useful as a sanity check on model behavior: the order of the features in the feature-importance list show that the model has 'learned' a physically sensible relationship and offers greater trust in the suitability of an operational deployment, and serves as a natural starting point to more definitively validate the model with domain meteorologists, given that it is learnt through essential data sets that only cover certain days.

7. Practical implications for agriculture, water resource management and early warning.

The agri-economic implications are directly applicable to the agricultural economy of the Northwest region, to the region's regional baselines, established in this chapter. The demand of Punjab and Haryana is disproportionately high and short-range forecasts of temperatures and precipitation play a critical role in determining when to plant wheat, rice, and other crops; when to use fertilizers and pesticides; when to irrigate crop fields; and time of harvesting. The temperature RMSE is in the range of operationally useful for frost-risk advisory for the rabi (winter) cropping season because unexpected cold spells during this season may result in significant wheat and winter crop losses.

Besides agriculture, the economically important short-range forecasts are useful for water resource management, such as for operation and protection of reservoirs, along the Indus river system's eastern tributaries, which are irrigated by canals. Another important application of reliable near-term forecasts is to water resources management, like the operation and the protection of reservoirs, in the Indus river system fed by canals, which depend on the forecast of precipitation on the short-range. In the same way, early warning systems for heatwaves (an emerging hazard in the Indo-Gangetic Plain) and for dense winter fog events (which impact transportation and aviation along the Grand Trunk Road (GTR) corridor) may benefit from regionally calibrated ML models similar to what is benchmarked here, especially where full NWP systems are impractical or uneconomical to maintain at the district level.

Lastly, the outputs of the RF model are interpretable feature-importance scores, that are directly usable by the district-level meteorologists and agriculture extension officers, who might not be comfortable with interpreting deep-learning model internals, and will likely help to foster trust and uptake of the model in actual advisory decision-making processes.

8. Limitations and future work.

The present study has the following three main limitations. The relatively coarse resolution of ERA5 (0.25°) will miss out on sampling the sub-grid scale convective events that can lead to intense local precipitation leading to severe but spatially small-scale damages to the crops due to first intense hail storms and second intense drop in temperature (cloudburst type storm from hail storms). Both of this is a major constraint of agriculturally pertaining applications in the region of Punjab. Further, from a basic physical point of view, all three models incorrectly simulate the

precipitation variability during summer monsoon season (June–September) with the synoptic-to-mesoscale interactions being poorly captured in the daily aggregated precipitation using the six-hourly/hourly aggregated reanalysis fields. Third, there may be constraints on seasonal-scale skill because a large-scale oceanic teleconnection index such as the El Niño–Southern Oscillation (ENSO) and Indian Ocean Dipole (IOD) index are not considered to be predictors.

Acknowledging these shortcomings, four directions are suggested for further research:

Explicitly capturing the spatial structure by integrating spatial information, using a graph neural network (GNN), which is specifically designed to model the interactions between the stations in the Northwest India power grid, to capture the spatial structure in addition to temporal sequences at fixed station locations.

Transformer based architectures with attention mechanisms for longer forecasting leads (3–7 days and extended range leads) forecasting applications with the proven performance of the Transformer models in the global fields, for example Pangu-Weather.

Mixing physics and ML models that incorporate the constraints of energy and mass conservation, and therefore limit the risk of unrealistic outcomes at longer lead times for the ML part of the model.

ENSO and IOD indices integration as extra seasonal covariates, soil-moisture and land-surface predictors integration as extra/noise covariates and the evaluation domain was extended for cross-regional comparison to pan India and wider extent of South Asia.

Table 3: Current Study Limitations and Proposed Future Research Directions

Dimension	Current Study	Proposed Future Direction
Spatial resolution	0.25° ERA5 grid (~31 km)	Station-level downscaling with GNNs
Temporal scope	24-hour lead time	3–7 day extended-range forecasting
Model architecture	RF, LSTM, CNN	Transformers, diffusion-based models
Predictors	Standard atmospheric fields	ENSO / IOD indices, soil moisture
Explainability	RF feature importance only	SHAP values for deep learning models
Evaluation domain	Northwest India	Pan-India and South Asia comparison

Conclusion

In the current chapter, the random forest (RF), long short-term memory (LSTM) and CNN architectures were objectively compared in short-term weather forecasting for 24 years of ERA5 reanalysis data over Northwest India, following a common, repeatable data preprocessing procedure. Finally, the two-layered LSTM networks generated the highest overall accuracy with an RMSE value of 1.61 °C, which represents a 58% improvement over a mean persistence value, and with a high correlation value on the test period of 0.96. The significant improvement over the persistence baseline by all three of these ML models were seen in both of the target variables, suggesting that while forecasting skill is obtained with the persistence baseline in this specific region, this did not arise from data driven means and was unimpressive. The importance of the

features was also analysed and as expected for synoptic climatology, the geopotential height at 500 hPa emerged as the most important predictor of the surface temperature observations, adding physical validity to the model's learned representations.

Apart from the technical contributions, this study provides validated forecasting baselines for the region that can directly be used in agricultural and water resource management and early-warning systems for the Punjab, Haryana, and Himachal Pradesh in India and a fully reproducible pipeline and evaluation framework for ERA5 that can be easily generalised to other Indian sub-regional forecasting areas. With shifting patterns of global data-driven weather models still maturing, region-specific studies like these will continue to be crucial to applying global advances in ML-based atmospheric science to locally trusted tools that can be operationally deployed in communities that will benefit most from the accuracy of short-range weather.

Acknowledgement

The author gratefully acknowledges the Copernicus Climate Change Service (C3S) and the European Centre for Medium-Range Weather Forecasts (ECMWF) for providing open access to the ERA5 reanalysis dataset used throughout this study, and the School of Science and Engineering, GNA University, Phagwara, for institutional support during the course of this research.

References

1. Bauer, P., Thorpe, A., & Brunet, G. (2015). The quiet revolution of numerical weather prediction. *Nature*, 525(7567), 47–55.
2. Bi, K., Xie, L., Zhang, H., Chen, X., Gu, X., & Tian, Q. (2023). Accurate medium-range global weather forecasting with 3D neural networks. *Nature*, 619(7970), 533–538.
3. Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
4. ECMWF. (2023). *IFS documentation CY48R1—Part III: Dynamics and numerical procedures*. European Centre for Medium-Range Weather Forecasts.
5. Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., Simmons, A., Soci, C., Abdalla, S., Abellan, X., Balsamo, G., Bechtold, P., Biavati, G., Bidlot, J., Bonavita, M., ... Thépaut, J.-N. (2020). The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730), 1999–2049.
6. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
7. Lam, R., Sanchez-Gonzalez, A., Willson, M., Wirsberger, P., Fortunato, M., Alet, F., ... Battaglia, P. (2023). Learning skillful medium-range global weather forecasting. *Science*, 382(6677), 1416–1421.
8. LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 2278–2324.

9. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., ... Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825–2830.
10. Rasp, S., Pritchard, M. S., & Gentine, P. (2018). Deep learning to represent subgrid processes in climate models. *Proceedings of the National Academy of Sciences*, 115(39), 9684–9689.
11. Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N., & Prabhat. (2019). Deep learning and process understanding for data-driven Earth system science. *Nature*, 566(7743), 195–204.
12. Saha, M., Mitra, P., & Nanjundiah, R. S. (2021). Deep learning for predicting the monsoon over the homogeneous regions of India. *Journal of Earth System Science*, 130(1), 1–15.
13. Shi, X., Chen, Z., Wang, H., Yeung, D.-Y., Wong, W.-K., & Woo, W.-C. (2015). Convolutional LSTM network: A machine learning approach for precipitation nowcasting. *Advances in Neural Information Processing Systems (NeurIPS)*, 28, 802–810.
14. Copernicus Climate Change Service (C3S). (2024). ERA5: Fifth generation of ECMWF atmospheric reanalyses of the global climate. ECMWF. <https://cds.climate.copernicus.eu>

DEEPAKE DETECTION USING DEEP LEARNING METHODS: A SYSTEMATIC AND COMPREHENSIVE REVIEW

Manmeet Kaur

School of Engineering Design and Automation,

GNA University, Phagwara, Punjab, India

Corresponding author E-mail: manmeet.kaur@gnauniversity.edu.in

Abstract

The application of synthetic media generated by GANs, autoencoders and diffusion models has grown and necessitated the development of reliable detection systems. The chapter provides a summary of the advances made in deep learning based deepfakes detection, from the inception of CNN classifiers to transformers and multi-modal systems. Introduces some basic concepts and generation principles, introduces spatial, temporal and frequency domain and biological-signal based methods of detection, discusses architectures, datasets and benchmarks, as well as emerging trends (Self-supervised learning, Explainable AI etc.) and ongoing challenges (Dataset bias, Adversarial robustness, Cross domain generalisation etc.). The chapter outlines examples of uses for journalism, law enforcement and social media governance and points to future research avenues on resilient, interpretable and scalable detection systems. Intended for students, researchers and professionals in the field of digital media forensics and AI security.

Keywords: Deepfake Detection, Deep Learning, Generative Adversarial Networks, Convolutional Neural Networks, Vision Transformers, Digital Forensics, Media Authentication, Multimodal Detection.

1. Introduction

1.1 Background

In the last 10 years, tremendous progress has been made in the field of deep generative modeling, which has revolutionized the way digital media are created and manipulated. Deep fake is a word derived from “deep learning” and “fake” which describes synthetic audio, video or image that appears as if it were real but is produced or altered with deep neural networks to convincingly represent events, statements, or appearances that never actually happened. This technology can be traced back to generative adversarial networks (GANs) introduced by Goodfellow *et al.* in 2014, which started a game of cat and mouse with two neural networks, a generator and a discriminator, that take turns generating synthetic images that become increasingly realistic. However, the level of sophistication of generative architectures has substantially increased in the years since this basic work, including pipelines based on autoencoders for face-swapping, style-transfer networks, and more recently, diffusion probabilistic models, which can generate photorealistic images with little perceptible artifacts.

Progress has been remarkable in making deepfake generation more accessible, with specialized research labs giving way to tools that can be easily used by consumers with minimal technical

knowledge. This democratisation of the production of synthetic media has heightened worries over the dissemination of misinformation, ID theft, non-consensual explicit content, political manipulation and loss of confidence in electronic proof. Thus, deepfake detection becomes a vital sub-discipline of the digital forensics and computer vision, and receives continuous attention from both academia as well as industry and governmental organizations.




Figure 1: Timeline of Deepfake Generation and Detection Milestones

1.2 Importance of the Topic

Enhanced deepfake detection reaches a number of societal and institutional areas. For journalism and media verification purposes, detection systems help to verify visual evidence before it goes live. The integrity of video evidence in legal proceedings is a key concern for forensic analysts when handling cases in a judicial or law enforcement environment. Synthetic identity fraud, voice and facial cloning, is a new threat to financial institutions, and handling massive amounts of user-generated content is an operational challenge for social media platforms.

Deepfake detection, from a technical perspective, is a thrilling challenge as it lies at the nexus of computer vision, signal processing, biometrics, and adversarial machine learning. Detection models, besides being generalizable for different generation methods and with respect to compression artifacts and demographic differences, must also be efficient in computation for deployment to real time. The challenge of the correct combination of accuracy and generalizability still spurs a significant amount of research funding.

1.3 Current Challenges

However, there are still some challenges that limit the real-world use of deepfake detection tools:

- **Generalization deficit** — detectors trained on specific manipulation techniques or datasets often exhibit substantial performance degradation when evaluated on previously unseen generation methods.
- **Adversarial vulnerability** — detection models remain susceptible to adversarial perturbations and post-processing operations such as compression, resizing, and noise addition, which can suppress detectable artifacts.

- **Dataset bias and imbalance** — publicly available benchmark datasets frequently exhibit demographic, resolution, and source imbalance, limiting fairness and real-world applicability.
- **Rapid evolution of generation techniques** — the emergence of diffusion-based and neural radiance field (NeRF) generation methods continually outpaces detector adaptation cycles.
- **Interpretability limitations** — many high-performing detectors function as opaque classifiers, offering limited explanation for their decisions, which constrains adoption in forensic and legal contexts requiring accountability.




Figure 2: Conceptual Map of Deepfake Detection Challenges

1.4 Objectives of the Chapter

This chapter aims to:

- Provide users with a meaningful conceptual and terminological basis for deepfake generation and detection.
- Discuss, classify and summarize deep learning-based detection methods in a systematic manner with respect to the technical approach they are based on.
- Provide comparisons of key architectures, datasets and metrics.
- Critically review recent developments, such as detector framework based on transformers and multimodal.
- Discuss critically the benefits, limitations and problems with the current methods.

Suggest avenues for future research and identify the potential for industry and governance related practical application.

Deepfake: synthetic media, in the form of images, videos or audio, produced or altered using deep neural networks that present the appearance of an individual speaking or acting in a specific manner that was never actually experienced or expressed.

2. Fundamental Concepts

2.1 Definitions

To review the detection methodologies, it is important to have a very precise conceptual vocabulary in place, because the synonyms or terms that appear in the literature are sometimes overlapping and sometimes inconsistently defined.

Synthetic media: any form of media that is partly or completely created using computational algorithms, such as the creation of a deepfake, but also, for example, all synthesized images, all images generated with a text-to-image algorithm or all voice recordings that are not for impersonation purposes, but are created by a voice synthesis algorithm.

Face swapping: a manipulation technique that involves replacing the ID of a source face in a target video while retaining the facial expression, pose and background context of said target video.

Face reenactment (puppetry): a manipulation method that involves manipulating a target's facial expression, head movement or lip movements through another actor's performance without altering the target's identity.

Attribute manipulation: change the attributes of a face or a scene, e.g., change hairstyle, age, skin tone, lighting, but not the basic identity of the face or scene.

Full face synthesis: create a full fake non-existent human face, usually generated via GANs like StyleGAN, commonly employed for the creation of fake social media accounts.

Deepfake detection: the computational problem of distinguishing between a given piece of media that is authentic (pristine) or manipulated (fake) – either as a binary or multi-class classification problem, but a localization-based formulation is becoming increasingly common, in which the goal is to detect the regions or frames where something is fake.

2.2 Terminology

Table 1: Key Terminologies in Deepfake Detection

Term	Definition
GAN	Generator/discriminator framework trained adversarially to produce realistic synthetic samples.
Autoencoder	Network trained to reconstruct input via a compressed latent representation; common in face-swapping.
Latent space	Compressed learned representation encoding semantic attributes of data.
Artifact	Residual inconsistency from synthetic generation, exploitable for detection.
Frequency-domain analysis	Spectral examination (e.g. Fourier/DCT) revealing generation artifacts.
Cross-dataset generalization	A detector's ability to maintain performance on unseen data distributions.
Diffusion model	Generative model synthesizing data by reversing a gradual noising process.

2.3 Core Principles

2.3.1 The Generative Basis of Deepfakes

Detection strategies make use of residual artifacts from certain generation paradigms. For autoencoder-based generation, an identity-specific image is encoded using the same encoder as the other identities, and then passed through the decoder for the target identity, generating a face swap. A discriminator is an adversarial training model that trains the generator against a discriminator, and architectures like StyleGAN added the ability to do latent manipulation with the generator to obtain fine-grained control and higher fidelity. Diffusion-based generation inverts a gradual noising process, resulting in increased sample diversity and fidelity, with new types of subtle artifacts that are non-GAN-specific that increase the difficulty of their detection. For example, the problem of detecting a message in noise is re-formulated as a classification problem.

Table 2: Taxonomy of Deepfake Generation Techniques

Paradigm	Mechanism	Typical Artifacts	Representative Methods
Autoencoder-based	Shared encoder, identity-specific decoders	Blending boundaries, resolution mismatch	DeepFaceLab, FaceSwap
GAN-based	Adversarial generator-discriminator	Checkerboard patterns, spectral irregularities	StyleGAN, FaceSwap-GAN
Diffusion-based	Iterative denoising	Subtle high-frequency inconsistencies	Stable Diffusion editors, DiffSwap

2.3.2 The Detection Problem as a Classification Task

In general, the deepfake detection procedure is a supervised learning problem, in which a model is trained to minimize a classification loss, which is usually a binary cross entropy. This simple formulation hides the fact that effective features need to encode subtle information that is spread out across spatial, temporal and frequency dimensions.




Figure 3: Generic Deep Learning-Based Detection Pipeline

2.3.3 Categories of Detectable Artifacts

- Spatial — spatial inconsistencies at the pixel level (e.g., blending boundaries, unnatural texture smoothness).

- Temporal — anomalies between frames, e.g., unnatural “blinking” or “flickering.”
- The Frequency-domain — irregularities in the spectrum from generative up sampling operations.
- Biological signal — unrealistic physiological signals like no rPPG heartbeat signals.
- Semantic/behavioural — higher level mismatches, e.g. in lip movement vs audio.

3. Detailed Discussion: Deep Learning Approaches to Deepfake Detection

This section provides a systematic discussion of the major deepfake detection techniques based on deep neural networks. The concept and underlying architecture, along with representative examples of architectures, are introduced at the start of each subsection, and illustrative application scenarios are provided for the reader's benefit, giving them both an understanding of the underlying concepts and some practical context.

3.1 Spatial (Frame-Level) CNN Approaches

A type of deep fake detection which has generally been adopted and is the earliest to be developed is based on convolutional neural networks (CNNs) that work on the per-frame or static image basis to take the problem of deep fake detection to the level of image classification. One of the first and more prominent works to be proposed is that of Afchar *et al.* MesoNet is an early influential paper that aims to train the network to be fairly compact with fewer convolutional layers, inspired by the fact that deepfake artifacts appear to be at the mesoscopic scale, between simple low level pixel statistics and high level semantic information. Later research has adopted deeper/network-specific models like Xception Net and ResNet variants as a transfer learning approach. Rössler *et al.* found that under moderate compression, XceptionNet, trained on manipulated face datasets, significantly outperformed shallower architectures, making it the de facto standard for comparison in the following studies. The major drawback of purely spatial CNN-based methods is that they tend to overfit the artefacts generated by the methods found in the training set, and fail to generalize to the manipulation techniques present in the test set — leading to the consideration of attention mechanisms that capture semantically meaningful facial regions and regularization methods that mitigate the use of low-level statistics of the training data set.

3.2 Temporal and Sequence-Based Approaches

Video content carries a lot of temporal information that is not present in a single frame, which is why considerable research is based on inconsistencies in multiple frames of a video, which are unique to synthetic video generation. The recurrent paradigm involves a CNN to learn spatial cues from each frame followed by an LSTM network to learn temporal cues among the sequence. Güera and Delp showed that such CNN-LSTM networks are able to learn subtle variations in identity representation, such as flickering or inconsistency, between adjacent frames, which are not noticeable if they are viewed separately.

The three-dimensional convolutional network (3D ConvNets) also extends the convolution operation along the temporal dimension, allowing to work with short video clips as 3D input, and

learning of filters that are sensitive to motion without an additional recurrent stage. When reenactments or manipulations occur via these methods, the subtle inconsistencies in dynamics of expressions and movement of the head offer discriminative cues that are particularly effective in their detection. One additional temporal approach identifies some physiologically meaningful behavioural signs, e.g., eye-blinking: In the study by Li *et al.*, early generative pipelines only used images with open eyes, which resulted in synthetic subjects that either blinked in an abnormal or unexpected way, or blinked completely, leading to the proposal of blink-analysis detectors to detect this abnormal blinking behaviour. In spite of having been largely overcome by later generation techniques, the principle that temporal dynamics that are physiologically implausible may signal a synthetic origin is widely used in modern research, such as lip-synchronization consistency analysis.




Figure 4: Architecture of a CNN-LSTM Temporal Deepfake Detector

3.3 Frequency-Domain Approaches

Instead of measuring in raw pixel image space this alternative, more and more influential paradigm works in the frequency domain: it is based on the fact that generative architectures, in particular transposed convolution and up sampling, generate characteristic periodic artefacts which can be measured in the frequency domain and allows detection methods that may not be spatially visible. In the paper, Durall *et al.*, the authors show that there is a systematic difference between the power spectrum of images generated by a GAN and real images, which is due to generator up sampling operations. Guided by this idea, Qian *et al.* came up with F3-Net, which used the discrete cosine transform to decompose the input images into the frequency domain, and then transmitted the frequency domain features together with the traditional spatial features, so that it is more resistant to lossy compression, which is a practically meaningful feature for social media video because it is usually compressed in lossy mode. Frequency-domain approaches provide a complementary strength to the spatial approaches and many current approaches take a dual-stream approach that extracts both spatial and frequency-domain features in parallel, fusing them before classification, which captures more complete and robust sets of discriminative cues.

3.4 Biological and Physiological Signal-Based Detection

A novel detection paradigm based on signals that are hard to generate with a biological fidelity. Cardiac-cycle driven blood volume changes cause a slight periodic change in the color of skin, a

phenomenon that can be used to estimate. This is what is observed in real video of real people, but is not always present, much diminished, or does not always manifest in the same way in synthetic human faces. However, Ciftci *et al.* proposed Fake Catcher that extracts signals from rPPG from several facial areas and evaluates their spatial and temporal similarity, has high detection accuracy and has a level of interpretability based on the physiological principles. This method targets a biological rather than the source generation-specific signal, thus showing comparatively favourable cross-manipulation generalization, while being sensitive to illumination, resolution, and compression of the video, that can individually affect the signal quality of the rPPG even in a genuine video.

3.5 Transformer-Based Detection Architectures

Since the widespread use of transformer architectures in computer vision, vision transformers (ViTs) are more and more being utilized in deepfake detection. While CNNs operate on images by learning structure in local regions (receptive fields), transformers learn the overall structure of an image using self-attention mechanisms, which might allow them to detect manipulation artifacts such as inconsistencies in the entire image rather than localized with respect to the image. Coccomini *et al.* explored the trade-off between convolutional feature extraction and transformer-based classification, showing that hybridization can perform just as well as, or even better than, convolutional models and provide stronger interpretability by visualization of the attention maps, especially in the context of practitioners' trust in the forensic and journalistic uses of transformers. In turn, more recent work has investigated using large-scale pretrained vision-language and vision transformer foundation models, trained or adapted with parameter-efficient methods, to leverage upon the rich, generalizable representations of the visual world learned from vast amounts of data prior to the few-samples learned for the detection task, to improve cross-dataset generalisation, which detectors have long struggled with when trained on narrow, task-specific datasets.

3.6 Multimodal and Audio-Visual Detection

Inconsistencies often arise across multiple modalities in manipulated videos, and researchers have created multimodal detection systems that look not just at the visual, but also at the auditory stream. An important area of application is related to consistency of lip-synchronisation, with the temporal correspondence between the observed lip movement and the corresponding acoustic representation of the speech being analysed; many of the face-swapping and re-enacting techniques have subtle points of desynchronisation, which can be difficult to completely remove. Haliassos *et al.* introduced LipForensics, which learns high-level semantic irregularities in mouth movement by leveraging representations pretrained on lip-reading tasks, demonstrating strong generalization to unseen manipulation methods and robustness to compression and distortion, attributable to the semantic rather than low-level nature of the learned features.

Beyond lip-synchronization analysis, multimodal frameworks have also incorporated voice biometric consistency checks, cross-modal attention mechanisms that jointly weight audio and visual evidence, and fusion architectures combining independently trained unimodal classifiers at

the decision level. These approaches are of particular relevance given the parallel proliferation of voice-cloning technology, which enables highly convincing fraudulent audio independent of, or in combination with, visual manipulation.

3.7 Comparative Summary of Detection Approaches

Table 3: Comparative Architecture Diagram of Detection Paradigms

Approach	Representative Methods	Key Strength	Key Limitation
Spatial CNN	MesoNet, XceptionNet	Efficient; strong within-dataset accuracy	Limited cross-manipulation generalization
Temporal/sequence	CNN-LSTM, 3D CNN, blink analysis	Captures inter-frame cues	Higher cost; sensitive to length
Frequency-domain	F3-Net, spectral analysis	Robust to compression	Less effective on diffusion outputs
Biological signal	FakeCatcher (rPPG)	Physiologically grounded generalization	Sensitive to resolution/lighting
Transformer	ViT hybrids	Long-range modeling; interpretability	Data- and compute-intensive
Multimodal audio-visual	LipForensics, fusion models	Cross-modal generalization	Needs synchronized A/V data

3.8 Benchmark Datasets

Public benchmarks vary in scale, diversity, and realism. FaceForensics++ remains widely used, with videos from four manipulation methods at multiple compression levels. DFDC substantially expanded scale and diversity of actors and generation methods. Celeb-DF improved synthesis realism to provide a more challenging benchmark. WildDeepfake and DeeperForensics-1.0 add in-the-wild and perturbation-augmented conditions, respectively.

Table 4: Benchmark Datasets for Deepfake Detection

Dataset	Year	Manipulation Types	Scale	Note
FaceForensics++	2019	Face-swap & reenactment (4 methods)	~1.8M frames	Multiple compression levels
DFDC	2020	Multiple GAN/non-learned	~128K videos	Large, diverse demographics
Celeb-DF	2020	Improved face-swap	~6K videos	High visual realism
WildDeepfake	2020	In-the-wild	~7.3K sequences	Uncontrolled source
DeeperForensics-1.0	2020	Face-swap + perturbation	~60K videos	Real-world distortion

3.9 Practical Scenario: Content Moderation Pipeline

A representative deployment: a social media platform first runs face detection/alignment on sampled frames, then a lightweight spatial CNN performs low-latency initial screening across all uploads. Content flagged as suspicious is routed to a heavier secondary stage combining temporal, frequency-domain, and (where audio exists) multimodal lip-sync analysis. This tiered design balances platform-scale computational constraints against detection accuracy — a common industry engineering compromise.

4. Latest Technologies and Current Trends

Most recent technologies & current trends. With the continuous advancements in generative modeling, the deepfake detection landscape remains dynamic. Several important areas of using technology in current and near-future research are emphasized.

4.1 Self-Supervised and Contrastive Learning

One of the challenges of traditional supervised detection models is the requirement for many labelled manipulated documents, which are time-consuming to create and biased towards the current means of document generation. To overcome this, many recent approaches to contrastive and self-supervised learning have been developed to learn general-purpose visual representations from unlabelled or weakly labelled data and fine-tuned to the detection task with relatively few labelled instances. These methods have shown to be more resilient to new ways of manipulation, due to the fact that the acquired representations hold more general visual regularities, than artifacts characteristic to the manipulation method only.

4.2 Detection of AI-Generated Content

With the recent success of diffusion probabilistic models for image and video generation, there is a novel detection problem: diffusion generated content has distinct characteristics, notably the presence of artifacts such as high repetition rates, which are significantly different than previous methods based on GANs, making many GAN trained detectors less effective against diffusion generated content. Increasingly, emerging research has started to describe diffusion-specific artifacts, which manifest themselves as distinct residual patterns of noise, and in certain circumstances, artifact signatures of the error that can be recovered by going around images containing contamination with diffusion inversion processes (an active and rapidly evolving subfield).

4.3 Explainable AI Integration

Integration of explainable AI (XAI) techniques is becoming more popular in the context of deploying detection systems in forensic, legal, and journalistic settings where accountability is expected. Class activation mapping, attention visualization in transformer models, and saliency-based approaches can help practitioners examine specific parts of an image or time frames within a video to deepen their understanding and make more accurate classification decisions, but are not fully autonomous, self-acting systems that provide “opaque” decision-making.

4.4 Generalizable and Foundation Model-Based Detection

A mode in which a generalizable or foundation model-based detection is carried out. A mode for executing a generalizable or foundation model-based detection. Large-scale pretrained vision and vision-language foundation models are increasingly used as backbones for deepfake detection tasks due to the fact that representations pretrained in large, different data sets tend to perform better in manipulation tasks and data distributions than those pretrained in a small task-specific dataset. Adapter modules and low-rank adaptation are among parameter-efficient fine-tuning methods that have been pursued to fine-tune the foundation models on the detection task without consuming large amounts of data or computation.

4.5 Adversarial Robustness Research

In parallel to the progress in detection, research is now shifting more and more to adversarial robustness of the detection system itself, as it is straightforward to realize that adversarial perturbation or targeted post processing of the content can be used to create manipulated content which can lead to evasion of detection. In addition to detection strategies tailored to identifying evasion attempts, robust training algorithms such as adversarial training and randomized smoothing are explored for the deepfake detection context.




Figure 5: Emerging Trends Infographic in Deepfake Detection Research

5. Advantages of Deep Learning-Based Detection

- Automated feature learning replaces manually engineered forensic features that allows detection systems to learn new manipulation patterns from the data.
- The accuracy of detection is high on within-distribution benchmark datasets, and is frequently over 95 percent in controlled evaluation conditions.
- Support for scalability to a large number of contents enables deployment in high throughput applications like content moderation pipelines for social media.
- Multimodal integration capability enables the detection systems to make use of the visual, audio and physiological cues, enhancing the robustness of the systems compared to single modality-based systems. Systems can adapt over time, using new manipulated samples to periodically retrain for continual improvement as generation techniques change.

- Interpretability tools, such as attention visualization and saliency mapping, are becoming more popular to enable transparent and verifiable decision making that is appropriate for forensic contexts.

6. Limitations and Challenges

There are practical constraints in detection systems based on deep learning, and coupled to these, several key challenges remain to continue to guide ongoing research efforts in this field.

- Performance drops when distribution shifts — detectors are often not as effective in the face of distribution shifts, that is, when the distribution of the data with which they are called upon to operate differs from the distribution over which they were trained, and this is closely related to the more general weakness of generalization: generation techniques are still rapidly diversifying and detection methods trained to a specific generation era face the threat of 'obsolescence' as new manipulation techniques are developed.
- Malicious actors have a clear motivation to create methods that can evade detection systems, especially when such systems have similar properties to other classifiers in deep learning — they can make a small, imperceptible change to an image that causes misclassification.
- Relying on large labelled datasets with labels is limiting, as creating a large number of manipulated content labels is ethically and practically challenging.
- Existing benchmarks – older versions of existing benchmarks tend to have certain biases and limits, such as demographic imbalance, variability in environmental and lighting conditions, and lack of the latest-generation techniques, and so can result in systems that perform differently between the various demographic groups, or that do not reflect deployment conditions in the real world.
- Sensitivity to compression and post-processing makes it harder to be certain whether the content has been detected when it is circulated through social media platforms, which are known for their use of lossy compression.
- Real time and resource-constrained deployment — successful architectures can be multimodal fusion or transformer-based and require computational limits that makes them a challenge to deploy on mobile devices or live video-streaming platforms, thus serving as a motivation for model compression and the efficient design of architectures.
- The questions of Lawfulness, Evidence, and Ethics — legal and evidential issues concerning deployment in legal and forensic settings, including the evidentiary standards for classification, who is ultimately responsible if a classification is wrong, and how to integrate automated decision-making with human expertise.
- Limited interpretability of many successful models limits their acceptance for legal/forensic uses that make the decision-making argument clear, defensible and transparent.

7. Future Research Directions

- The development of unified detection frameworks that are able to tackle both GAN-based and diffusion-based generation artifacts within one unified architecture, thus avoiding the necessity to build separate detectors for each generation method.
- Improvements in few-shot and zero-shot detection methods that could learn to detect new manipulation methods when little or no method-specific labelled training data is available.
- More specific usage of EAI techniques that align to the structure of the forensic and legal evidential needs, from general saliency visualization to creating a structured and human interpretable representation of the explanation.
- Development of a greater number of demographically balanced and environmentally diverse benchmark datasets with standardized fairness evaluation protocols.
- Search for certifiably robust detection architectures that resist adversarial evasion, possibly with formal robustness guarantees.
- Discussion of lightweight, edge deployed detection architectures that can be implemented in real-time on mobile, embedded platforms.
- Standardization of evaluation protocols and shared benchmarking infrastructure will facilitate more consistent evaluation of all published detection methods.
- Research on mechanisms for the provenance of content on blockchain and cryptography computing technologies as a complementary approach to the fully detection-based method, for establishing the authenticity of the content at the moment of the capture.

8. Practical Applications

There are many possible applications of deep learning deepfake detection methods, which are listed below.

- **Journalism and media verification:** news media have greater access to detection systems, which they are using to verify authenticity of user-generated or external video content before publication to reduce the risks of unwittingly sharing manipulated content during time sensitive news coverage.
- **Law enforcement and judicial proceedings:** Confined to law enforcement and judicial proceedings, forensic laboratories and judicial bodies utilize detection systems in conjunction with human expert complement to evaluate the integrity of video images, to assist investigative efforts and to determine whether video images are admissible in court.
- **Social media content moderation:** operators of social media platforms use automated detection pipelines, with or without additional human intervention, to automatically flag potentially manipulated content at scale, to aid in the implementation of content policies on their platforms and to limit the dissemination of misinformation.
- **Financial fraud prevention:** Voice and facial verification, combined with deepfake detection capabilities, are adopted by financial institutions in financial fraud prevention

processes to tackle the risk of synthetic identity fraud that arises from voice cloning and facial spoofing.

- **Public information integrity:** Government and civil society organizations have created and implemented detection tools to help with public information integrity during election periods.
- **Corporate brand and identity protection:** With more and more brands letting organizations know to watch out for fake or synthetically generated images of executives or brand representatives being circulated by others on public platforms, it is imperative that organizations have the right parameters to detect such content and limit its spread.

Summary and Conclusion

The chapter has provided a comprehensive review of Deepfake detection techniques using deep learning, providing a systematic and comprehensive understanding of the topic in relation to the overall technical landscape of synthetic media production. It started with the basic concepts and terminology, and then addressed the main classes of detection methods: spatial convolutional methods, temporal and sequence-based methods, frequency domain methods, biological signal-based methods, transformer networks, and multimodal audio-visual methods. All map them rely on different kinds of artifact, from texture mismatches to signal that is impossible from a physiological point of view, with each option having its own favoured qualities and weaknesses, as detailed in the comparative analysis of Section 3.7.

These are just a few of the current research trends that underscore the field's efforts to adapt to the dynamic nature of generative modeling, with deep interest in Self Supervised learning, Foundation model adaptation, Explainable AI integration and Artifacts characterization. However, ongoing difficulties, especially generalization problem, adversarial vulnerability, dataset bias, and the importance of detection correctness versus real-time computational constraints remain to limit the success and robustness of operational detection systems.

The implications of this research are far-reaching, and pervade deployed systems in journalism, law enforcement, social media governance, financial fraud prevention, and electoral integrity protection. With the continued development of generative technologies in their complexity and availability, continued technology research to develop widely applicable, interpretable, and efficient detection methods will continue to be vital to ensure public confidence in digital media and to reduce societal harms from the misuse of synthetic content. Going forward, architectural innovation is likely to be an important factor, but so will be the development of standardized evaluation frameworks, more representative and equitable benchmark datasets, and closer integration between technical detection capabilities and the legal, journalistic and platform governance contexts in which they will be placed.

References

1. Goodfellow, I., *et al.* (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems (NeurIPS)*, 2672–2680.

2. Rössler, A., *et al.* (2019). FaceForensics++: Learning to detect manipulated facial images. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 1–11.
3. Afchar, D., Nozick, V., Yamagishi, J., & Echizen, I. (2018). MesoNet: A compact facial video forgery detection network. *IEEE Workshop on Information Forensics and Security (WIFS)*, 1–7.
4. Li, Y., Chang, M.-C., & Lyu, S. (2018). In icu oculi: Exposing AI created fake videos by detecting eye blinking. *IEEE Workshop on Information Forensics and Security (WIFS)*, 1–7.
5. Li, Y., Yang, X., Sun, P., Qi, H., & Lyu, S. (2020). Celeb-DF: A large-scale challenging dataset for deepfake forensics. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3207–3216.
6. Dolhansky, B., *et al.* (2020). *The DeepFake Detection Challenge (DFDC) dataset*. arXiv. <https://arxiv.org/abs/2006.07397>
7. Güera, D., & Delp, E. J. (2018). Deepfake video detection using recurrent neural networks. *IEEE 15th International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, 1–6.
8. Durall, R., Keuper, M., Pfrendt, F.-J., & Keuper, J. (2019). *Unmasking DeepFakes with simple features*. arXiv. <https://arxiv.org/abs/1911.00686>
9. Qian, Y., Yin, G., Sheng, L., Chen, Z., & Shao, J. (2020). Thinking in frequency: Face forgery detection by mining frequency-aware clues. *European Conference on Computer Vision (ECCV)*, 86–103.
10. Ciftci, U. A., Demir, I., & Yin, L. (2020). FakeCatcher: Detection of synthetic portrait videos using biological signals. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*.
11. Coccomini, D. A., Messina, N., Gennaro, C., & Falchi, F. (2022). Combining EfficientNet and vision transformers for video deepfake detection. *International Conference on Image Analysis and Processing (ICIAP)*, 219–229.
12. Haliassos, A., Vougioukas, K., Petridis, S., & Pantic, M. (2021). Lips don't lie: A generalisable and robust approach to face forgery detection. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 5039–5049.
13. Karras, T., Laine, S., & Aila, T. (2019). A style-based generator architecture for GANs. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4401–4410.
14. Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems (NeurIPS)*, 6840–6851.
15. Tolosana, R., Vera-Rodriguez, R., Fierrez, J., Morales, A., & Ortega-Garcia, J. (2020). Deepfakes and beyond: A survey of face manipulation and fake detection. *Information Fusion*, 64, 131–148.

16. Mirsky, Y., & Lee, W. (2021). The creation and detection of deepfakes: A survey. *ACM Computing Surveys*, 54(1), 1–41.
17. Nguyen, T. T., et al. (2022). Deep learning for deepfakes creation and detection: A survey. *Computer Vision and Image Understanding*, 223.
18. Verdoliva, L. (2020). Media forensics and DeepFakes: An overview. *IEEE Journal of Selected Topics in Signal Processing*, 14(5), 910–932.
19. Zhao, H., Zhou, W., Chen, D., Wei, T., Zhang, W., & Yu, N. (2021). Multi-attentional deepfake detection. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2185–2194.
20. Zi, B., Chang, M., Chen, J., Ma, X., & Jiang, Y.-G. (2020). WildDeepfake: A challenging real-world dataset for deepfake detection. *Proceedings of the 28th ACM International Conference on Multimedia (ACM MM)*, 2382–2390.
21. Jiang, L., Li, R., Wu, W., Qian, C., & Loy, C. C. (2020). DeeperForensics-1.0: A large-scale dataset for real-world face forgery detection. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2889–2898.
22. Chen, S., Yao, T., Chen, Y., Ding, S., Li, J., & Ji, R. (2021). Local relation learning for face forgery detection. *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 1081–1088.
23. Shiohara, K., & Yamasaki, T. (2022). Detecting deepfakes with self-blended images. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 18720–18729.
24. Cao, J., Ma, C., Yao, T., Chen, S., Ding, S., & Yang, X. (2022). End-to-end reconstruction-classification learning for face forgery detection. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4113–4122.
25. Luo, Y., Zhang, Y., Yan, J., & Liu, W. (2021). Generalizing face forgery detection with high-frequency features. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16317–16326.
26. Korshunov, P., & Marcel, S. (2018). *Deepfakes: A new threat to face recognition? Assessment and detection*. arXiv. <https://arxiv.org/abs/1812.08685>
27. Zhu, X., Wang, H., Fei, H., Lei, Z., & Li, S. Z. (2021). Face forgery detection by 3D decomposition. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2929–2939.
28. Vaswani, A., et al. (2017). Attention is all you need. *Advances in Neural Information Processing Systems (NeurIPS)*, 5998–6008.
29. Dosovitskiy, A., et al. (2021). An image is worth 16×16 words: Transformers for image recognition at scale. *International Conference on Learning Representations (ICLR)*.
30. Cheng, H., Guo, Y., Wang, T., Li, Q., Chang, X., & Nie, L. (2023). Voice-face homogeneity tells deepfake. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 20(3), 1–22.

DEVELOPMENT OF AN INTELLIGENT LEARNING SYSTEM: SKILL LEDGER FOR SKILL EXCHANGE AND PAID COURSES

Riya Pathania and Neha Bagga*

School of Engineering, Design and Automation,
GNA University, Phagwara, Punjab, India

*Corresponding author E-mail: neha.bagga@gnauniversity.edu.in

Abstract

Skill Ledger Skill Ledger is a smart learning platform helping to bridge disconnect between the development of skills and the available learning opportunities through a combination of skills exchange and paid learning services. The standard learning channels are concentrated on provision of paid courses, and a good number of learners have good skills that could be distributed to others. This project will focus on establishing a hybrid platform, in which users can teach and learn simultaneously on two modes of flexible approaches, including; skill exchange and premium paid courses. The site enables participants to sign-up, add profiles, post their abilities, find other students and share their skills. The professional users may offer paid courses, training or mentoring to the professional users in need. Authenticated course user management, user dash boards and communication, provide ensure that every user is engaged in a interactive environment. The Skill Ledger development uses the latest web technologies and full-stack architecture to come up with a scalable integration, responsive and a user-friendly solution. The suggested system encourages cooperative learning, boosts employability and generates income to talented people.

This project is expected to deliver an innovative digital outcome focusing on facilitating monetized education through the emergence of a new digital ecosystem that promotes knowledge sharing. Online learning has the potential to evolve further in the future with Skill Ledger contributing to the evolving landscape as it introduces a more interactive, practical and attainable education in the future.

Keywords: Skill Ledger, Smart Learning Platform, Skill Exchange, Paid Learning, Knowledge Participating, Online Education, Web Operation, Cooperative Literacy, Full-Stack Development.

Introduction

Development in technology has dramatically transformed our way of learning and the acquisition of our skills. Conventional modes of learning are being eliminated giving way to technological based modes of learning which offer the learner more choice and convenience. The advent of online learning as a fast-developing industry can be credited to the increase in consumption of mobile machines, computers, and access to the internet. Now everybody is able to teach themselves, improve their work skills and be able to read learning materials at the place of your choosing.

E-learning systems are quite minimal and they have enhanced over the years. Most of the most employed e-learning platforms are based on a one-way model, where learners would register a course developed by an instructor or organization and pay. This is terrific and even the finest form of the service will just grab a little fraction of the entire strength of peer-to-peer (or community) learning. There are millions of people with great skills to share with other people; but the overwhelming majority of these people do not nowadays have the place, where they can share this knowledge with other people, and earn money as the people, who want to know their skills.

Though tremendous developments have occurred in e-learning during the past years, most of the current systems of e-learning remain very limited in their methodology. The majority of e-learning platforms have one-directional business model; Users subscribe to a course being sold by an instructor or organization. This is a good business model except that it.

It is not taking full advantage of collaborative learning. Millions of individuals globally possess marketable skill sets like computer programming, graphic design, video editing, oratorical addressing, online marketing, language translation, photography, and so on. Nonetheless, such a large number of these people do not have an online platform to disseminate their knowledge to other people, and to earn money in case they teach their skills.

Another obstacle to the current education systems is the cost of education. The students and other learners are mostly not in position to obtain premium courses because of financial constraint but they can have provided a valuable skill set to another individual. The creation of a system that enables the sharing of skills among the individuals as opposed to continuous payment of services, generates an environment that supports inclusive and accessible education. Such an arrangement would have a number of advantages to the learners such as students, workers seeking employment, free agents and self-motivated learners who need to update their skills, yet cannot because of limited resources.

Skill Ledger is a new innovative learning platform, which provides the chance to share skills and engage in paid learning. A union between two forms of learning enables users to trade their skills freely with other users.

To learn on this platform, the first option is to find someone who possesses the skill that you would like to learn; in the case of User A being skilled in graphic design skills and User B being skilled in web development, then User A can teach User B the graphic design skills (via Skill Ledger) and User B can teach User A the web development skills (via Skill Ledger).

Online learning on the second platform (paid) can be based on courses, workshops, mentor, one-on-one consulting by a professional, trainer, or specialist and directed at those who are interested in mastering specialized knowledge.

Skill Ledger is hybrid and its value is differentiated by being a free cooperative (user-to-user), and professional paid (user-to-educator) learning experience in one digital ecosystem. Users are no longer mere content consumers aiming to provide users with valuable content but also contributors, teachers, and may even earners that build a vibrant user-centred learning

community with more focus on the relationships between users and the feasibility of the learning process.

The overall objectives of this project are to:

- To establish an online community of sharing expertise among users.
- To provide paid learning through expert learning through courses/mentorship programs.
- To provide a flexible and user-friendly interface to interact easily.
- To develop a community that encourages participatory learning activities.
- To provide new sources of revenues to the talented professionals.
- To offer a wide-ranging, more affordable and easy way to access education to every student.

Methodology

Study Design

In this research project, the methodology adopted is non-experimental, developmental approach to designing and implementing Skill Ledger which is a smart learning system relying upon three key factors; that is, skill exchange, payment.

for learning services. This study aims at exploring the shortcomings of the existing online learning systems and creating a new system of delivering on-line education that facilitates collaborative education as well as generating revenues through education. Since this project is oriented on designing a software system instead of scientific experimentation, the major highlight will be on system analysis; gathering requirements; design and development of the system and testing of the prototype.

The study is anchored on a realistic approach of research where real-life issues to the learners and instructors are examined and platform offered as an online one to solve the issues. The number of learners seeking cheap education is on the rise, and so are those with the capability to offer their services to other people, to share knowledge and make money. To provide an intermediary between the two groups, Skill Ledger will act as a transition to creating a single platform to facilitate skill exchanges and a professional growth that is paid.

Sample / Population of the Study

The population of the study is comprised of individuals who learn, teach or train in their profession. The target audience consists of students, freelancers, job seekers and working professionals.

To analyze the possible expectations that such a type of platform may have, a group of such users was studied. It has included students since they are normally seeking less-priced alternatives to acquire other skills. Professionals/freelancers can also fit in this population since they have been exposed to valuable skill set and would likely be interested in both teaching as well as monetizing the skills. This group also involved the trainers/instructors, because as they would be the most likely people to use the site in order to develop organized paid courses or coaching.

This identified user community shows clearly that there exist the need and potential of flexible and low cost means of making an income using a hybrid learning model.

Data Gathering Tools

In order to create Skill Ledger, several data gathering tools were employed to gather pertinent data. The tools below were able to give an insight into the challenges faced by the users, the system requirements and the market opportunities.

1. Literature Review

Literature reviews of literature and articles/reports, regarding availability of e-learning systems, online marketplaces and skill sharing communities were conducted in order to gain an insight into emerging trends, challenges and effective features of such systems.

2. Surveys and Questionnaires

The potential users of Skill Ledger (students and working professionals) answered a survey which was structured to gather data about their obstacles to obtaining quality education, desire to share their skills, and preference in an online learning platform functionality.

3. Interviews

The informal interviews involved users that use digital learning platforms on regular basis. Their contribution played a crucial role in determining hindrances to online education, including course prices, necessity to interact, and the establishment of a trustworthy relationship with other members.

4. Observation

Existing online learning platforms were observed with the aim of assessing a comparison between their user interfaces, business models, and business constraints to what Skill Ledger will be up to. This allowed determining the gaps that can be addressed by Skill Ledger.

Data Gathering Procedures

The systematic approach was adopted in the process of data collection. Initially, the existing issues with online learning were determined. These problems entailed excessive course costs, low interaction, absence of peer-to-peer learning and income opportunities of the common skilled users.

Second, user expectations were obtained in the form of surveys and discussions. Several respondents are keen on learning by cooperating, sharing expertise, and availing low-cost access to experts. A simple navigation, safe payment, and the ability to trust in communication came out as key factors emphasized by the users, as well.

Third, reviews, surveys and observations provided the information that was analyzed and transformed into system requirements. The user registration, login, creation of profile, list of skills, course management, messaging, dashboards and payment features were also proposed as functional requirements. Security, responsiveness, scalability and easy design were specified in the non-functional requirements.

Such processes also made sure that the final platform was constructed as per the real user requirements and not as the assumption.

System Development Method

Skill Ledger is being developed in Agile Software Development Methodology that enables incremental development, frequent testing and constant improvements made based on user feedback. The construction of the project was carried out in stages:

Phase 1 - Requirement Analysis

Here we conducted a step-by-step analysis of the requirements of the entire system and recorded them. We checked the functions desired by the users, how they thought the system would work and how to use the system.

Phase 2 - System Design

We came up with an architectural design of the entire system that had a layout of the front end application, the back end architecture, the Schema of the Database as well as the Work Flow application diagrams. All new features also received wireframes and instructions on how the user-interface should be designed to make them easier to use.

Phase 3 - Development

Coding of the application was in contemporary Web technologies. The front-end was created to interact with end users whereas the back-end was created to provide the services of business logic, authentication services as well as data processing. The Database contained user profiles and user skills, information system (CIS) of course, course history of transactions and messages between users.

Phase 4 - Testing

The system was well tested in terms of functional, security, responsive and performance. Every facet of the system such as registration, sign on, skill matching, course registration and payment processing were fully tested end to end, to verify that all the system functions as intended. Lastly, bugs or errors identified during the testing process were fixed before implementing the system finally.

Phase 5 - Deployment & Evaluation

Following successful testing of the system, the system became ready to be deployed. User feedback was requested and utilized to gauge the performance of the system and work out how to improve on it in the future.

Treatment Of Data

We sorted out the data gathered through surveys, interviews and observation according to various user requirements, user preferred features, frequent challenges and user expectations. To conduct the analysis of the responses to determine the common patterns of each type of user, a descriptive approach was utilized. To illustrate a case, numerous users revealed their interest in cheaper educational services and want to utilize a system that offers secure payment systems and has explicit communicative properties.

Moreover, test data were also recorded to determine the accuracy of the system, its loading time and the successful completion of functions of the platform by the user. These reviews assisted in verifying the fact that the designed system was, in fact, efficient.

Ethical Considerations

Ethical principles have been followed during the course of the study and development of the platform. In order to motivate the respondents to participate in the surveys and interviews, the participation of a respondent was voluntary. All the people who either took part in surveys or/and were interviewed were made aware of the purpose of collection of this data prior to giving their responses. The identifying information of a person will be held in confidence and will only be used in carrying out the research.

The security will be applied to all the user data that was obtained through the platform, such as authentication, the usage of secure web services, encryption of all user files, and determining user accessions. None of the copyrighted materials have been copied without giving a recognition of the sources and correctly citing all the references materials used. Respecting user trust due to the responsible use of collected data was considered an important aspect to the project.

Results and Discussion

Skill Ledger project has led to a good (functionality, usability and innovation) output in terms of the role, usability and innovation of online learning environment. Skill Ledger was developed to facilitate paid learning services and functional capability (skill exchange) using Skill Ledger which is a smart web-based platform. There are numerous ways of learning, teaching and developing career when utilizing the platform. The project results suggest that the suggested platform will allow overcoming most of the drawbacks of conventional e-learning platforms.

Goal 1: To design a platform of skills sharing.

The primary objective of this project was to develop an online platform, where skills can be shared between the users. It was created in a way that the users could not only share their knowledge or experience with others but also acquire knowledge in various other skills with others Presented Result

The solution that has been implemented, by allowing users to create a personal profile of their talents and locating users sharing the same interest(s), has been effective in helping users connect to fellow users that share the same tastes as they do. To illustrate, this graphic designer has an opportunity to collaborate with a web developer and share knowledge.

Table 1: Planning raising skill view

Step	Description
1.	Registers: user creates profile
2.	Adds skills provided by the user
3.	Skills were needed in user searches
4.	User match with matching user
5.	Learning Exchange begins




Figure 1: User Preference for Learning Mode

Discussion of the Findings

The results show that learning is not necessarily associated with financial dealings but the manner in which the platform uses skills of the existing users as currency expands the access to education and facilitates user-to-user communication.

Corroborations Advantages (+):

- Affordable Learning Approach
- Community Participation
- Promotes Engagement

Disadvantages (-):

- Needs active user base to make correct matches.
- User: The users need to develop trust toward one another.

Objective 2:

To Introduce Paid Learning Opportunity Introduction to Result. The second was to provide a one stop place to the professionals and other professionals who would wish to deliver high quality educational materials to others as part of these paid learning services. Presented Result The platform provides an opportunity to use paid courses, mentoring sessions, and training programs created by the instructors easily. Students can do a search in all of the courses that are being offered, apply to a course of their interest and get a access to a course material.

Table 2: Paid Learning Module

Feature	Status
Course Creation	Implemented
Enrollment System	Implemented
Payment Support	Implemented
Dashboard Access	Implemented

Interpretation of Result

Individuals who require well-organized learning processes would benefit the paid learning module because the learning module would enable a more formal kind of learning process via the guidance of an expert rather than informal ways of learning. Analysis/Implications This feature creates a sustainable ecosystem in the education space as it gives instructors a platform to

earn money by teaching and delivering a quality education to the students. Moreover, it creates a possibility of the increase of the platform as more instructors and pupils are involved into the platform and the value to all parties will increase in the future.




Figure 2: Enrollment in Paid Course

Corroborations Advantageous (+):

- Professionals: opportunities in incomes.
- As a Professional, Access to Learning.
- Sustainable Business model

Disadvantageous (-):

- All the users may not afford the use of a paid service.
- The paid service needs to be implemented in a way that includes secure payment.

Goal 3

To Develop a Responsive, Deck of Cards like Introductions of Result. Much focus has been on ensuring a well-shaddered user experience in every one of the web platforms. This has led to the development of a user-friendly interface to our project whereby people will be able to use it through their mobile phones or tablets and computers. Presented Result System testing, showed that the system could be easily used by customers to complete tasks such as: registration; login; browsing; enroll courses; manage dashboards and so on. The platform is also responsive to both the mobile and desktop display features.

Table 3: Usability Test Results

Parameter	Result
Easy Navigation	High
Mobile Responsiveness	High
Registration Process	Successful
Dashboard Accessibility	High

Interpretation of Result

The web design responsive to touch and the user-friendly nature of the site enabled consumers (both the beginners and more professional consumers) to easily navigate the site without getting lost in its tricky functionalities. The users are able to utilize their devices irrespective of whether they are starting to use them or have years of experience using them.




Figure 3: Platforms Usability Feedback

Analysis/Implications

- Improved user interface, which will result in a larger number of users desiring to use the system.
 - Increased user satisfaction, resulting in their increased enjoyment when using the system.
 - Increased chances of continued usage of the system after sometime.
 - The system is used by a larger number of users who visit the system.
 - Higher adoption among the users of various types (e.g. age, education, etc.)
- Corroborations

Pros

- Easy to Use
- Can Work with several devices. con Cons Greater User Satisfaction

Cons

- Updates to Better User Interface (UI) / User Experience (UX) required.

Objective 4

To Develop Earning Opportunites to competent persons Introduction of Result. Part of what we wanted to do was to inform the user that this was a learning opportunity to educate others about what they believed in and guide others in their profession. Presented Result Those with very developed skills can conduct paid courses or sessions; thus, giving them another chance at earning money.

Table 4: Economic Benefit of Skill Ledger

UserType	Benefit
Learners	Affordable Learning
Experts	Income Generation
Community	Knowledge Growth

Interpretation of Result

This platform can enable individuals to make out of what they know a form of income to themselves as they learn themselves in the process.




Figure 4: Community Participation Rate

Analysis/Implications

It is possible that the implementation of this feature will assist in bringing more skilled individuals onto our platform as this feature will offer them a higher-quality learning material and a bigger variety of learning material.

Corroborations

Positive:

- We are in favour of new business development via Entrepreneurship.
- We give reward of the expertise levels.
- We are increasing the amount of Courses Available.

Negative:

- There is a rivalry in the ranks of the instructors.

Overall Discussion

The Skill Ledger has been able to achieve its target requirements as indicated by the results examined. It is a platform consisting of integrated service model which enables the user to share expertise and participate in paid learning collaboratively. This blended instead of a fully electronic based learning experience helps to address some of the safeguarding challenges of the existing electronic systems such as cost, limited interactions among learners and providers of knowledge/resources, and absence of economic chance to most consumers. It would also be beneficial, in future digital learning systems, using the results, to abandon unidirectional, traditional models of teaching in favor of allowing the user to learn and contribute via a community based ecosystem. Secondly, as it is being improved to include aspects like AI.

Skill Ledger, with its recommendations, a rating system, video conferencing, and learning in various languages, will have enormous potential in becoming a leader in digital learning platform.

Conclusions and Recommendations Conclusion

1. The development of Skill Ledger: A Smart Learning Platform to Skill Swap and Paid Courses has already achieved its goals to create a hybrid online platform where users can learn collaboratively and have the opportunity to study professional paid courses. This project helped resolve most of the major problems associated with older systems of e-learning such as expensive rates, lack of interaction, lack or minimal peer to peer learning

and earning opportunities by knowledgeable individuals. The introduction of additional features demonstrated the way to make learning more flexible, more affordable and more community-driven.

2. Objective to create a platform where users can share their skills has been met. Users are able to form their own profiles, express their skills to others, search through those who share their skills and chat with them to see how they can assist each other to learn. This illustrates that knowledge is a commodity that can be distributed without necessarily engaging in an exchange using a currency. The skill exchange model offers a chance to work collectively, increases the networking possibilities, and broadens the coverage of education among all kinds of learners, independent workers, and employed people.
3. The second significant success of giving paid learning opportunities has been achieved as well. The system also enables those professionals, trainers and experts to provide courses, mentorship programs and training classes to those who desire to get specialized training. It will not only offer more learning resources that can be accessed in the system (enhancing the overall quality of the learning resources on the system), but will also offer a sustainable revenue model not only to the instructors but also to the system in general. This will enable an ecosystem where free and paid education can manage to co-exist harmoniously.
4. Another main goal that the project attained was the production of a user interface that was easy to use and responsive. The platform was to ensure that the platform is easy to navigate, accessible and useable with various devices (e.g. desktops, tablets and mobile phones). The outcome of the usability testing showed that the participants were very satisfied in terms of their experience when it comes to registering, access to the dashboard, and the ability to use the mobile version of the platform. Having a user-friendly and receptive interface is essential to boost the extent of user engagement and prompt the user to use the platform regularly.

The last objective of creating opportunities in individuals with skills was realized with the monetisation features on the platform, where users with expertise sold services, created courses or met with other users at a fee creating potential revenue on their skills capabilities. This not only promotes entrepreneurship, professional growth and economic empowerment but also affords other users additional chances to acquire new things. In general, the outcomes of the project show that Skill Ledger is the first model to succeed in online education in the future. The system facilitates.

community-support and encourages the development and the involvement of the users, who are not just the learners, but also the contributors, teachers and earners by incorporating skill and service materials, as well as learning services. In this regard, the Skill Ledger system is highly likely to become an important participant in the fast-growing digital education market.

References

1. Anderson, T. (Ed.). (2008). *Online learning theory and practice* (2nd ed.). Athabasca University Press.
2. Express.js. (2024). *Express.js: Fast, unopinionated, minimalist web framework for Node.js*. <https://expressjs.com/>
3. MongoDB, Inc. (2024). *MongoDB documentation*. <https://www.mongodb.com/docs>
4. UNESCO. (2023). *Education and technology: Global education monitoring report*. UNESCO Publishing. <https://unesdoc.unesco.org>
5. Zhang, D., Zhao, J. L., Zhou, L., & Nunamaker, J. F. (2004). Is it possible to substitute classroom learning with e-learning? *Communications of the ACM*, 47(5), 75–79. <https://doi.org/10.1145/986213.986216>
6. Meta. (2024). *React: A JavaScript library for creating user interfaces*. <https://react.dev/>
7. Node.js Foundation. (2024). *Node.js documentation*. <https://nodejs.org/en/docs>
8. Moodle. (2024). *Open source learning management system*. <https://moodle.org>

DIGITAL TWIN AND AI-DRIVEN BATTERY MANAGEMENT FOR SUSTAINABLE TRANSPORTATION

Aswini M^{*1}, Deivanayaki. R¹, Immanuel Prabaharan. S¹ and S. Muthurajan²

¹Department of Electrical and Electronics Engineering

²Department of Marine Engineering

AMET University, Kanathur, Chennai-603112

*Corresponding author E-mail: aswinifalls@gmail.com

Abstract

The increasing adoption of electric vehicles (EVs) has highlighted the importance of advanced Battery Management Systems (BMS) for ensuring battery safety, reliability, and long-term performance. Conventional BMS techniques primarily rely on sensor measurements and model-based estimation methods to monitor battery conditions. However, these approaches often face challenges in accurately estimating battery states under dynamic operating conditions, temperature variations and battery aging. Recent advances in Digital Twin (DT) technology and Artificial Intelligence (AI) have introduced a new paradigm in battery management by enabling real-time monitoring, predictive analysis and intelligent decision-making. A digital twin is a virtual representation of a physical battery system that continuously exchanges information with the actual battery through sensors, communication networks and cloud platforms. When integrated with AI algorithms such as Artificial Neural Networks (ANN), Long Short-Term Memory (LSTM) networks, Support Vector Machines (SVM) and Physics-Informed Neural Networks (PINNs), the digital twin can accurately estimate State of Charge (SOC), State of Health (SOH), State of Power (SOP) and Remaining Useful Life (RUL). These capabilities support predictive maintenance, intelligent charging strategies, thermal management and early fault detection, thereby improving battery efficiency and extending operational life. Furthermore, cloud computing, edge computing and the Internet of Things (IoT) enable continuous remote monitoring and fleet-level battery optimization. This chapter presents an overview of Digital Twin technology and AI-driven battery management for sustainable transportation. It discusses the architecture of AI-enabled Digital Twin-based BMS, key AI techniques, major applications in electric mobility, current challenges and future research trends. The integration of Digital Twin and AI is expected to play a significant role in developing intelligent, reliable and sustainable battery management solutions for next-generation electric vehicles.

Keywords: Digital Twin, Artificial Intelligence, Battery Management System, Electric Vehicles, Sustainable Transportation.

1. Introduction

The transportation sector is undergoing a significant transformation with the rapid adoption of electric vehicles (EVs) as a sustainable alternative to conventional internal combustion engine vehicles. Growing concerns about greenhouse gas emissions, fossil fuel depletion and stringent

environmental regulations have accelerated the development of clean transportation technologies. At the core of every electric vehicle is the lithium-ion battery pack, which serves as the primary energy storage system. The performance, safety and lifespan of the battery directly influence vehicle range, charging efficiency, reliability, and overall operating cost.

Battery Management Systems (BMS) are essential for monitoring and controlling battery operation. A typical BMS measures battery voltage, current and temperature while estimating critical parameters such as State of Charge (SOC), State of Health (SOH), and State of Power (SOP). It also performs protection functions including overcharge prevention, over-discharge protection, thermal management and cell balancing. Although conventional BMS methods provide satisfactory performance under normal operating conditions, their accuracy decreases as batteries age or operate under highly dynamic conditions.




Figure 1: Evolution of Battery Management Systems from Conventional BMS to AI-Driven Digital Twin-Based BMS

Recent developments in Artificial Intelligence (AI) and Digital Twin (DT) technology have significantly enhanced battery monitoring and control. AI algorithms can analyze large volumes of battery data to identify complex patterns, predict degradation and improve state estimation. Meanwhile, a Digital Twin creates a dynamic virtual representation of the physical battery that is continuously updated using real-time sensor data. This virtual model enables engineers to monitor battery behavior, simulate operating conditions, predict failures and optimize battery performance without interrupting actual vehicle operation.

The integration of AI with Digital Twin technology provides a powerful framework for intelligent battery management. AI improves the accuracy of battery state estimation and

predictive maintenance, while the Digital Twin enables continuous synchronization between the physical battery and its virtual counterpart. Together, these technologies support intelligent charging, fault diagnosis, thermal management, and lifecycle optimization, contributing to improved battery safety, longer service life and enhanced energy efficiency.

This chapter provides a concise overview of Digital Twin technology and AI-driven Battery Management Systems for sustainable transportation. It discusses the fundamentals of Digital Twin technology, major AI techniques used in battery management, integrated system architecture, practical applications, current challenges and future research directions.

2. Fundamentals of Digital Twin Technology

Digital Twin (DT) technology is one of the most significant advancements in intelligent engineering and Industry 4.0. A Digital Twin is a dynamic virtual representation of a physical system that continuously receives real-time data from its physical counterpart through sensors and communication networks. Unlike conventional simulation models, which are developed for specific operating conditions, a Digital Twin is continuously updated using live operational data, enabling accurate monitoring, analysis and prediction throughout the system's lifecycle.

Initially introduced in aerospace and manufacturing industries, Digital Twin technology is now widely applied in healthcare, smart cities, renewable energy and transportation. In electric vehicles (EVs), it plays a vital role in Battery Management Systems (BMS) by providing a virtual model of the battery pack that mirrors its real-time operating conditions. This virtual model enables continuous monitoring of battery performance, degradation and thermal behavior, allowing engineers to make informed decisions without interrupting vehicle operation.

A Digital Twin-based Battery Management System consists of four main components:

- **Physical Battery System:** The lithium-ion battery pack equipped with sensors to measure voltage, current, temperature and other operating parameters.
- **Data Acquisition and Communication Layer:** Battery data are collected through the BMS and transmitted using communication protocols such as CAN, Ethernet, wireless networks, or IoT gateways.
- **Digital Twin Model:** A virtual battery model developed using electrochemical models, equivalent circuit models and data-driven approaches that continuously reflects the battery's operating condition.
- **Analytics and Decision Layer:** Artificial Intelligence and data analytics process real-time information to estimate battery states, detect faults and recommend appropriate control actions.

The continuous interaction between the physical battery and its virtual model enables real-time synchronization. As battery operating conditions change due to charging, discharging, temperature variations, or aging, the Digital Twin updates its parameters accordingly. This adaptive capability improves estimation accuracy and supports predictive maintenance by identifying potential failures before they occur.

The integration of cloud computing and edge computing further enhances Digital Twin performance. Cloud platforms provide large-scale data storage and computational resources for battery analytics, while edge computing performs time-critical processing closer to the battery system, reducing communication delays. Together, these technologies enable efficient remote monitoring, fleet management, and intelligent battery optimization.

Digital Twin technology offers several advantages over conventional battery monitoring methods. It improves the estimation of battery states, supports predictive maintenance, reduces development and testing costs through virtual simulations, and enhances battery safety by enabling early fault detection. As electric vehicles become increasingly connected through the Internet of Things (IoT), Digital Twins are expected to become a key component of next-generation Battery Management Systems for sustainable transportation.

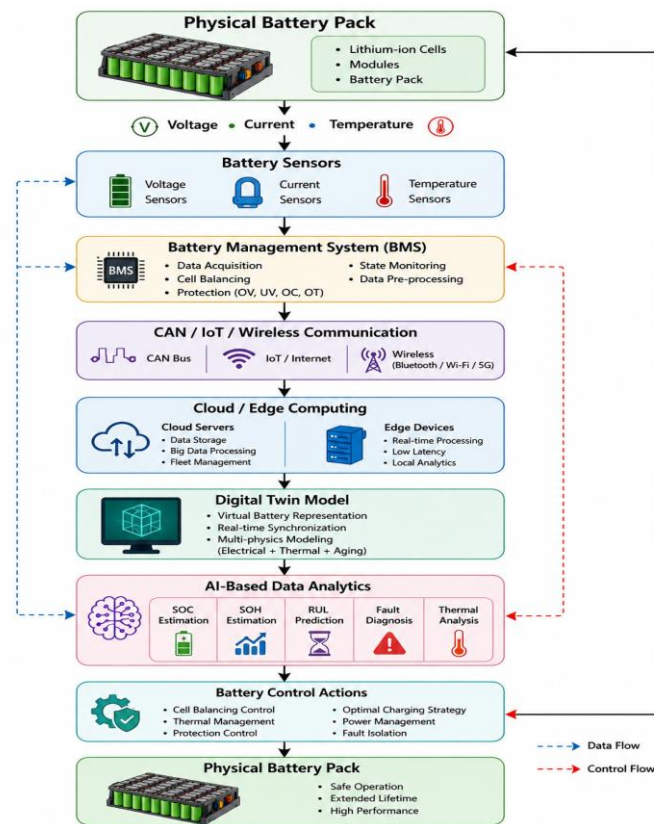


Figure 2: Architecture of a Digital Twin-Based Battery Management System

Table 1: Comparison of Conventional Battery Monitoring and Digital Twin Technology

Feature	Conventional Battery Monitoring	Digital Twin-Based Monitoring
Data Source	Periodic sensor measurements	Continuous real-time data
Model Update	Manual or fixed	Automatic and adaptive
Monitoring	Offline/Periodic	Real-time
Fault Detection	Reactive	Predictive
Battery State Estimation	Moderate accuracy	High accuracy
Remote Monitoring	Limited	Cloud-enabled
Decision Support	Rule-based	AI-assisted

3. Artificial Intelligence for Battery Management Systems

Artificial Intelligence (AI) has become a key enabling technology for modern Battery Management Systems (BMS) due to its ability to analyze complex battery behavior and improve the accuracy of battery state estimation. Conventional estimation methods, such as Coulomb counting, Open Circuit Voltage (OCV), Equivalent Circuit Models (ECM), and Kalman Filters, are widely used but often experience reduced accuracy under varying temperatures, dynamic driving conditions, and battery aging. AI techniques overcome these limitations by learning nonlinear relationships from large datasets, enabling more reliable battery monitoring and control.

Machine learning and deep learning algorithms use historical and real-time battery data—including voltage, current, temperature, charging cycles, and internal resistance—to estimate battery parameters that cannot be measured directly. These parameters include State of Charge (SOC), State of Health (SOH), State of Power (SOP), and Remaining Useful Life (RUL). Accurate estimation of these states improves battery safety, energy efficiency, and overall vehicle performance.




Figure 3: Artificial Intelligence Techniques Used in Battery Management Systems

Among machine learning techniques, Artificial Neural Networks (ANN) are widely used because they effectively model the nonlinear characteristics of lithium-ion batteries. ANNs learn complex relationships between battery inputs and outputs through training, making them suitable for SOC and SOH estimation. Support Vector Machines (SVM) are another popular approach,

particularly for battery health classification and degradation prediction. SVM models provide reliable performance even with relatively small datasets and are less prone to overfitting.

Deep learning methods have further improved battery state estimation. Long Short-Term Memory (LSTM) networks are especially effective because they capture temporal dependencies in sequential battery data. Since battery performance depends on previous charging and discharging cycles, LSTM models provide highly accurate predictions of SOC, SOH, and RUL. Convolutional Neural Networks (CNN) are also employed to extract meaningful features from battery signals and thermal data. In many studies, hybrid models such as CNN–LSTM combine spatial and temporal feature extraction to achieve higher prediction accuracy than individual models.

Recent research has introduced Physics-Informed Neural Networks (PINNs), which combine physical battery models with AI algorithms. Unlike purely data-driven methods, PINNs incorporate electrochemical principles during model training, improving prediction accuracy and reducing the need for extensive datasets. This approach is particularly valuable in Digital Twin applications, where maintaining consistency between the physical battery and its virtual model is essential.

AI also plays an important role in battery fault diagnosis and predictive maintenance. By continuously analyzing sensor data, AI algorithms can identify abnormal voltage patterns, temperature rise, cell imbalance, or rapid capacity degradation before serious failures occur. This enables preventive maintenance, minimizes downtime, and improves battery safety. Furthermore, AI supports intelligent charging strategies by optimizing charging current and voltage according to battery condition, thereby reducing degradation and extending battery lifespan.

Overall, AI enhances the capabilities of Battery Management Systems by providing adaptive learning, accurate battery state estimation, early fault detection, and intelligent decision-making. When integrated with Digital Twin technology, these capabilities enable real-time battery monitoring, predictive analytics, and optimized battery operation, making AI a fundamental component of next-generation sustainable transportation systems.

Table 2: Common AI Techniques for Battery Management

AI Technique	Primary Application	Key Advantage
ANN	SOC and SOH estimation	Models nonlinear battery behavior
SVM	Battery health classification	High accuracy with limited data
LSTM	SOC and RUL prediction	Captures time-dependent battery behavior
CNN	Feature extraction and fault diagnosis	Learns complex signal patterns
CNN–LSTM	Battery state estimation	Combines spatial and temporal learning
PINNs	Digital Twin and battery modeling	Integrates physics with AI for improved accuracy

4. Integration of Digital Twin and AI for Battery Management

The integration of Digital Twin (DT) and Artificial Intelligence (AI) has significantly enhanced the capabilities of Battery Management Systems (BMS) by combining real-time monitoring with intelligent decision-making. A Digital Twin provides a virtual representation of the physical battery, while AI analyzes the continuously updated data to estimate battery states, detect faults, and optimize battery operation. Together, these technologies create a smart and adaptive battery management framework that improves reliability, safety, and energy efficiency.

In a Digital Twin-based BMS, the physical battery pack is equipped with sensors that continuously measure voltage, current, temperature, and other operating parameters. These data are transmitted through the BMS using communication protocols such as CAN, wireless networks, or IoT gateways. The collected information is synchronized with the Digital Twin, which updates the virtual battery model in real time. AI algorithms then process the synchronized data to estimate battery conditions and predict future performance.

One of the major advantages of this integrated framework is the accurate estimation of battery states. AI models improve the prediction of State of Charge (SOC) by learning the nonlinear relationship between measurable battery parameters and the available energy. Similarly, State of Health (SOH) is estimated by analyzing battery aging indicators such as capacity fade, internal resistance, and charge–discharge history. AI also predicts State of Power (SOP) and Remaining Useful Life (RUL), enabling timely maintenance and better utilization of battery resources.

Digital Twin technology also supports predictive maintenance, where the virtual battery continuously compares expected and actual battery performance. Any deviation from normal behavior can indicate the early stages of battery degradation or faults. AI algorithms identify these abnormal patterns and generate maintenance recommendations before failures occur, thereby reducing downtime and improving operational safety.

Cloud computing and edge computing further enhance the Digital Twin framework. Cloud platforms provide large-scale data storage and advanced analytics for battery fleets, while edge devices perform real-time processing close to the battery system. This combination ensures low-latency decision-making while enabling remote monitoring and software updates. The integration of IoT devices also allows battery data to be shared with charging stations, fleet management systems, and service centers, creating a connected ecosystem for intelligent transportation.

The combined use of Digital Twin and AI provides several benefits over conventional Battery Management Systems. These include higher accuracy in battery state estimation, continuous health monitoring, optimized charging strategies, improved thermal management, reduced maintenance costs, and extended battery life. In addition, the virtual environment enables engineers to simulate different operating conditions, evaluate battery performance, and test new control strategies without affecting the actual battery.

Despite these advantages, successful implementation requires reliable sensor data, secure communication networks, and robust AI models capable of adapting to different battery

chemistries and operating conditions. Ongoing research is focused on improving model accuracy, reducing computational complexity, and developing standardized Digital Twin platforms for large-scale deployment in electric vehicles.




Figure 4: Integrated AI–Digital Twin Framework for Battery Management

Table 3: Benefits of Integrating Digital Twin and AI in Battery Management

Function	Conventional BMS	AI–Digital Twin BMS
Battery Monitoring	Periodic	Continuous real-time
SOC/SOH Estimation	Moderate accuracy	High accuracy
Fault Detection	Reactive	Predictive
Maintenance	Scheduled	Condition-based
Charging Strategy	Fixed	Adaptive and optimized
Thermal Management	Basic control	Intelligent prediction and control
Remote Monitoring	Limited	Cloud-enabled
Decision Making	Rule-based	AI-assisted

5. Applications of Digital Twin and AI in Sustainable Transportation

The combination of Digital Twin (DT) and Artificial Intelligence (AI) is transforming battery management in sustainable transportation by improving battery performance, operational safety, and energy efficiency. These technologies enable continuous monitoring, predictive maintenance, and intelligent decision-making, making them suitable for a wide range of transportation applications.

One of the most significant applications is in electric passenger vehicles (EVs). Modern EV battery packs consist of hundreds of lithium-ion cells operating under varying temperatures, charging conditions, and driving patterns. Digital Twin technology continuously creates a virtual representation of the battery pack using real-time sensor data. AI algorithms analyze this information to estimate battery parameters such as State of Charge (SOC), State of Health

(SOH), and Remaining Useful Life (RUL). This allows the Battery Management System to optimize charging, prevent battery degradation, and improve driving range.

Digital Twin technology is equally valuable for electric buses and commercial vehicles, where batteries experience frequent charging and discharging cycles under heavy loads. Fleet operators can remotely monitor battery health across multiple vehicles through cloud-based Digital Twin platforms. AI identifies abnormal battery behavior, predicts maintenance requirements, and schedules servicing before failures occur. As a result, vehicle downtime is reduced while battery reliability and fleet availability are improved.

Another important application is smart charging infrastructure. Conventional charging stations apply similar charging strategies to all batteries regardless of their condition. In contrast, AI-enabled Digital Twin systems continuously evaluate battery health and recommend personalized charging profiles based on battery temperature, degradation level, and charging history. This intelligent charging strategy minimizes battery stress, extends battery life, and improves charging efficiency.




Figure 5: Applications of Digital Twin and AI in Sustainable Transportation

The integration of Digital Twin technology with Vehicle-to-Grid (V2G) systems further enhances sustainable transportation. In V2G operation, electric vehicles exchange electrical energy with the power grid during periods of high or low electricity demand. AI algorithms

predict battery degradation caused by charging and discharging cycles while optimizing energy exchange without compromising battery lifespan. This approach improves grid stability, supports renewable energy integration, and creates additional economic benefits for EV owners.

Digital Twin technology also plays an important role in battery fleet management. Large transportation fleets generate significant amounts of battery data every day. AI processes these data to evaluate battery performance, identify degradation trends, and compare battery health across multiple vehicles. Fleet managers can use this information to optimize charging schedules, reduce operational costs, and improve overall fleet efficiency.

Beyond transportation, Digital Twin technology supports Battery Energy Storage Systems (BESS) connected to renewable energy sources such as solar and wind power. By predicting battery behavior under varying energy generation conditions, AI improves energy storage efficiency and contributes to the reliable integration of renewable energy into modern power systems.

Overall, the combination of Digital Twin and AI enables intelligent battery monitoring, predictive diagnostics, adaptive charging, and optimized energy management. These capabilities improve battery reliability, extend service life, reduce maintenance costs, and contribute significantly to the development of sustainable transportation systems.

6. Challenges in AI-Driven Digital Twin Battery Management

Although the integration of Digital Twin (DT) and Artificial Intelligence (AI) has significantly enhanced Battery Management Systems (BMS), several technical and practical challenges remain before these technologies can be widely adopted in commercial electric vehicles.

One of the major challenges is the availability of high-quality battery data. AI algorithms require large and diverse datasets collected under different operating conditions, such as varying temperatures, charging rates, driving cycles, and battery aging levels. However, real-world battery data often contain measurement noise, missing values, and communication delays, which can reduce the accuracy of AI models and affect synchronization between the physical battery and its digital twin.

Another challenge is the computational complexity of Digital Twin platforms. Real-time battery monitoring requires continuous execution of battery models, thermal simulations, and AI algorithms. While cloud computing offers substantial computational resources, communication latency can limit real-time performance. Edge computing addresses this issue by processing critical battery data closer to the vehicle, enabling faster decision-making and improving system responsiveness.

Cybersecurity and data privacy are also significant concerns because Digital Twin-based BMS rely on cloud services, wireless communication, and Internet of Things (IoT) technologies. Unauthorized access, cyberattacks, or data manipulation may compromise battery operation and vehicle safety. Therefore, secure communication protocols, encryption methods, authentication techniques, and intrusion detection mechanisms are essential to ensure reliable operation of connected battery systems.

7. Future Research Directions

Continuous advancements in Artificial Intelligence and Digital Twin technology are expected to further improve Battery Management Systems for next-generation electric vehicles. One promising research area is Explainable Artificial Intelligence (XAI), which aims to make AI predictions more transparent and interpretable. Unlike conventional deep learning models that often function as "black boxes," XAI provides insights into how decisions are made, increasing user confidence and supporting safety-critical applications.




Figure 6: Future Roadmap of AI-Driven Digital Twin Battery Management

Another emerging approach is Federated Learning, where multiple vehicles collaboratively train AI models without sharing raw battery data. Instead of transferring sensitive information to a central server, only model parameters are exchanged. This approach enhances data privacy while enabling continuous improvement of AI algorithms across large vehicle fleets.

The development of advanced battery technologies, including solid-state batteries, lithium-sulfur batteries, and sodium-ion batteries, will require new Digital Twin models capable of accurately representing their unique electrochemical characteristics. In addition, future communication technologies such as 5G and 6G networks will provide ultra-low latency and high-speed connectivity, enabling seamless synchronization between physical batteries, Digital Twins, and cloud platforms.

Another important research direction is the integration of Digital Twin technology with Battery Energy Storage Systems (BESS) and renewable energy sources. Intelligent battery monitoring

can optimize energy storage, improve renewable energy utilization, and support smart grid operations. Furthermore, Digital Twins can facilitate battery lifecycle management by monitoring battery performance from manufacturing to second-life applications and recycling, thereby promoting sustainable resource utilization and the circular economy.

Conclusion

Digital Twin and Artificial Intelligence have emerged as transformative technologies for modern Battery Management Systems. Their integration enables continuous battery monitoring, accurate estimation of State of Charge (SOC), State of Health (SOH), State of Power (SOP), and Remaining Useful Life (RUL), while supporting predictive maintenance, intelligent charging, and fault diagnosis. Compared with conventional Battery Management Systems, AI-driven Digital Twins provide improved operational efficiency, enhanced battery safety, and longer battery service life through real-time data analysis and adaptive decision-making.

These technologies are increasingly being applied in electric vehicles, smart charging infrastructure, Vehicle-to-Grid (V2G) systems, Battery Energy Storage Systems (BESS), and renewable energy integration, contributing significantly to sustainable transportation. Although challenges such as data quality, computational requirements, and cybersecurity remain, ongoing developments in Explainable AI, Federated Learning, cloud-edge computing, and next-generation communication networks are expected to overcome these limitations.

In the coming years, AI-driven Digital Twin technology will become a core component of intelligent Battery Management Systems, enabling safer, more reliable, and energy-efficient electric mobility. Its ability to support predictive analytics, optimize battery utilization, and improve lifecycle management will play a crucial role in achieving sustainable transportation and advancing the global transition toward clean energy systems.

References

1. Hannan, M. A., Hoque, M. M., Mohamed, A., & Ayob, A. (2017). Review of energy storage systems for electric vehicle applications: Issues and challenges. *Renewable and Sustainable Energy Reviews*, 69, 771–789.
2. Plett, G. L. (2015). *Battery management systems, volume I: Battery modeling*. Artech House.
3. Plett, G. L. (2015). *Battery management systems, volume II: Equivalent-circuit methods*. Artech House.
4. Plett, G. L. (2015). *Battery management systems, volume III: Battery state estimation*. Artech House.
5. Barreras, J. V., Fleischer, C., Schaltz, P., & Sauer, D. U. (2022). Battery management systems for electric vehicles: A review. *Renewable and Sustainable Energy Reviews*, 157, Article 112070.
6. Hu, X., Feng, F., Liu, K., Zhang, L., Xie, J., & Liu, B. (2019). State estimation for advanced battery management systems: Key challenges and future trends. *Renewable and Sustainable Energy Reviews*, 114, Article 109334.

7. Liu, K., Hu, X., Zhou, H., Tong, L., Widanalage, D., & Marco, J. (2020). Feature analyses and modeling of lithium-ion batteries using machine learning. *Energy AI*, 1, Article 100001.
8. Zhang, Y., Wang, C., Tang, J., & Pecht, M. (2023). Digital twin-driven battery management systems for electric vehicles: A review. *Journal of Energy Storage*, 58, Article 106095.
9. Liu, Y., He, H., & Sun, F. (2023). Artificial intelligence-based battery management systems for electric vehicles: A review. *Energy Reports*, 9, 4372–4392.
10. Saxena, S., Le Floch, C., MacDonald, J., & Moura, S. (2015). Quantifying EV battery degradation using data-driven methods. *Journal of Power Sources*, 282, 265–276.
11. Pecht, M., & Saxena, A. (2022). Battery prognostics and health management using artificial intelligence. *IEEE Access*, 10, 54025–54045.
12. He, H., Xiong, R., Zhang, X., Sun, F., & Fan, J. (2011). State-of-charge estimation of lithium-ion batteries using machine learning techniques. *IEEE Transactions on Vehicular Technology*, 60(4), 1461–1469.
13. Xiong, R., Tian, J., Shen, W., & Sun, F. (2014). A review of lithium-ion battery state estimation methods for electric vehicles. *Applied Energy*, 113, 106–119.
14. Hu, X., Zhang, K., Liu, K., Lin, S., & Majeed, H. D. T. (2021). Machine learning for battery health prediction and management. *eTransportation*, 8, Article 100115.
15. Ng, A., Chen, Y., & Wang, J. (2023). Physics-informed neural networks for lithium-ion battery modeling and state estimation. *Journal of Energy Storage*, 63, Article 106956.

SUSTAINABLE NANOMATERIALS FOR NEXT-GENERATION PHOTONIC APPLICATIONS

Abinaya. L and K. Gayathri*

Academy of Maritime Education & Training, Kanathur, Chennai – 603112

Corresponding author E-mail: gayathri.k@ametuniv.ac.in

1. Introduction

1.1 Overview of Sustainable Nanomaterials

Nanomaterials, defined by at least one dimension in the sub-100 nanometre range, derive their distinctive electronic, optical, and catalytic behaviour from quantum confinement, high surface-to-volume ratio, and surface-state effects that have no direct analogue in bulk matter. Over the past two decades, the field has expanded from a purely academic curiosity into a cornerstone of modern materials engineering, underpinning technologies as diverse as catalysis, biomedicine, energy storage, and photonics. However, the conventional routes used to synthesise many nanomaterials, ranging from high-temperature chemical vapour deposition to solvothermal processes reliant on toxic reducing agents, carry a substantial environmental and energy burden. This has motivated a decisive shift toward sustainable nanomaterials: nanoscale materials whose synthesis, processing, and end-of-life behaviour are designed to minimise ecological impact while preserving, or even enhancing, functional performance. Sustainability in this context spans several interrelated dimensions, including the use of renewable or waste-derived precursors, the replacement of hazardous solvents with water or bio-based media, energy-efficient low-temperature processing, and the avoidance of scarce or toxic elements such as cadmium and lead wherever technically feasible. Green chemistry principles, low-hazard reagents, biodegradable capping agents, and closed-loop material recovery are no longer peripheral considerations but central design criteria for next-generation nanomaterials.

1.2 Importance of Photonic Technologies

Photonics, the science and engineering of light generation, manipulation, and detection, has become as consequential to the twenty-first century as electronics was to the twentieth. Optical fibre networks carry the overwhelming majority of global data traffic, light-emitting diodes have displaced incandescent and fluorescent lighting worldwide, and photonic sensors underpin applications ranging from autonomous vehicle navigation to non-invasive medical diagnostics. The appeal of photonic systems lies in their capacity for high bandwidth, low signal loss, immunity to electromagnetic interference, and, increasingly, their compatibility with low-power, miniaturised device architectures. As global demand for computing power, connectivity, and real-time sensing continues to escalate, photonic technologies are expected to play an even more central role in energy-efficient information processing and environmental monitoring. Realising this potential, however, depends on the continued discovery of materials capable of generating, guiding, and detecting light with high efficiency, spectral tunability, and long-term stability.

1.3 Role of Sustainable Nanomaterials in Next-Generation Photonics

Sustainable nanomaterials occupy a particularly strategic position at the intersection of these two imperatives. Their optical properties, including tunable absorption edges, size-dependent photoluminescence, plasmonic resonances, and efficient photon upconversion, can be engineered with a precision unattainable in bulk materials, making them exceptionally well suited to photonic device design. At the same time, the growing availability of green synthesis routes for metal oxides, carbon-based nanostructures, and rare-earth-doped nanophosphors means that many of these functional benefits can now be obtained without the environmental costs traditionally associated with nanomaterial fabrication. Carbon dots derived from agricultural waste, biosynthesised metal oxide nanoparticles, and lead-free perovskite nanocrystals exemplify a broader trend in which ecological responsibility and photonic performance are pursued as complementary, rather than competing, objectives. This convergence positions sustainable nanomaterials as enabling components for a new generation of photonic devices that are simultaneously higher performing and lower impact.

1.4 Scope and Objectives of the Chapter

This chapter provides a structured overview of sustainable nanomaterials for photonic applications, written for postgraduate students, researchers, and academicians seeking a coherent entry point into this rapidly evolving field. It begins by examining green synthesis strategies and the structural characteristics of the principal nanomaterial classes relevant to photonics, namely metal oxides, carbon-based nanostructures, and rare-earth-doped nanophosphors, along with the characterisation techniques used to validate their structure. It then explores the optical properties and photonic mechanisms that govern their behaviour, before turning to a detailed discussion of application domains, including solid-state lighting, optical sensing, photovoltaics, optical communication, wearable photonics, and environmental monitoring. The chapter closes with a critical assessment of the technical, manufacturing, and environmental challenges that remain, together with an outlook on emerging research directions. Throughout, particular emphasis is placed on developments reported between 2023 and 2026, reflecting the pace at which this field continues to advance.

2. Sustainable Nanomaterials: Synthesis and Structural Characteristics

2.1 Green Synthesis Approaches

Green synthesis has emerged as the dominant paradigm for producing photonic-grade nanomaterials with a reduced environmental footprint. Phytosynthesis, in which plant extracts rich in polyphenols, flavonoids, and terpenoids act simultaneously as reducing and capping agents, has been widely adopted for metal and metal oxide nanoparticles, offering ambient-temperature, aqueous-phase reactions that avoid the toxic reagents associated with conventional chemical reduction. Microbial and fungal biosynthesis extends this principle further, exploiting cellular metabolic machinery to precipitate nanoparticles with controllable morphology. Complementary strategies include the valorisation of agricultural and food-processing waste, such as citrus peel, coffee grounds, and lignin residues, as carbon and precursor sources; the use

of ionic liquids and deep eutectic solvents as low-toxicity reaction media; mechanochemical routes such as ball milling that eliminate solvents altogether; and photocatalytic or photochemical synthesis, in which light itself drives nanoparticle nucleation under mild conditions. Collectively, these approaches substantially lower the energy intensity and hazardous waste generation associated with nanomaterial production, while frequently yielding surface chemistries that are directly compatible with biomedical and environmental photonic applications.

2.2 Metal Oxide Nanomaterials

Wide-bandgap metal oxides, including zinc oxide, titanium dioxide, tin oxide, and copper oxide, remain foundational materials in sustainable photonics owing to their earth abundance, chemical stability, and versatile optoelectronic properties. Zinc oxide, with a direct bandgap near 3.37 electron-volts and a large exciton binding energy, supports efficient near-band-edge ultraviolet emission alongside visible defect-related luminescence, making it attractive for light-emitting and sensing applications. Biosynthesis of these oxides, typically through plant-extract-mediated precipitation or sol-gel routes conducted at low temperature, yields quasi-spherical or hexagonal nanostructures whose defect density, and therefore optical response, can be tuned through precursor concentration and calcination conditions. Doping with transition metals or coprecipitation with other oxides further extends their spectral tunability and photocatalytic activity, properties that are increasingly exploited jointly for combined photonic sensing and environmental remediation functions.

2.3 Carbon-Based Nanomaterials

Carbon-based nanostructures, including graphene, carbon nanotubes, and carbon or graphene quantum dots, have gained particular prominence within sustainable photonics because they can be synthesised almost entirely from biomass and waste feedstocks while offering intrinsically low toxicity and high biocompatibility. Carbon dots produced via hydrothermal or microwave-assisted carbonisation of citrus peel, tea residues, or other organic waste display tunable, often excitation-dependent photoluminescence arising from a combination of quantum confinement and surface state emission. Recent reviews highlight their emergence as multifunctional platforms for optoelectronic biosensing, integrating tunable fluorescence with sustainable synthesis routes for applications spanning photodetectors, photocatalytic systems, and hybrid sensing architectures. Beyond sensing, carbon dots have also demonstrated unconventional lasing behaviour, with recent work documenting their use as gain media across a range of optical geometries, underscoring their potential as environmentally benign alternatives to conventional semiconductor gain materials in miniaturised photonic sources.

2.4 Rare-Earth Doped Nanophosphors

Rare-earth-doped nanophosphors, in which lanthanide ions such as europium, terbium, dysprosium, and samarium are incorporated into host lattices including vanadates, tungstates, aluminates, and fluorides, constitute a distinct and technologically important class of sustainable photonic material. Their luminescence arises from well-shielded intra-4f electronic transitions,

which confer narrow emission linewidths, long excited-state lifetimes, and excellent photostability. A particularly significant subclass, upconversion nanoparticles, converts low-energy near-infrared excitation into higher-energy visible or ultraviolet emission through mechanisms including excited-state absorption, energy transfer upconversion, cooperative sensitisation, cross-relaxation, and energy migration-mediated upconversion. Sol-gel, combustion, hydrothermal, and microwave-assisted synthesis routes, increasingly conducted with green solvents and process-intensified reactors, allow precise control over host composition and dopant concentration, which in turn governs emission colour, chromaticity, and quantum efficiency.

2.5 Structural Characterization Techniques

Reliable structure-property correlation in sustainable nanomaterials depends on a complementary suite of characterisation tools. X-ray diffraction establishes crystalline phase purity and, through Scherrer analysis of peak broadening, provides estimates of crystallite size, information that is essential for confirming successful nanoscale formation via green routes. Scanning electron microscopy reveals particle morphology and aggregation state, while transmission electron microscopy resolves finer structural detail, including lattice fringes, core-shell architectures, and precise particle size distributions critical for quantum-confined systems. Fourier-transform infrared spectroscopy identifies the functional groups of phytochemical or microbial capping agents retained on nanoparticle surfaces, offering indirect evidence of the green synthesis mechanism itself. Ultraviolet-visible absorption spectroscopy, interpreted through Tauc plot analysis, quantifies the optical bandgap and is routinely used to track how synthesis parameters influence electronic structure. Used together, these techniques provide the structural evidence base needed to link green synthesis conditions to the optical and photonic performance discussed in the following section.

2.6 Recent Developments (2023-2026)

The 2023-2026 period has seen marked progress across all three nanomaterial classes discussed above. Lead-free tin-based perovskite photovoltaic materials, for example, improved in certified efficiency from roughly one per cent shortly after their introduction to over fifteen per cent by 2024 using formamidinium tin iodide compositions, alongside parallel advances in double perovskite systems such as caesium silver bismuth bromide, which show improved ambient stability when confined to the nanocrystal regime. In the rare-earth domain, sub-fifty-nanometre upconversion nanoparticles have achieved quantum yields approaching thirteen per cent, roughly four times greater than conventional micrometre-scale bulk crystals, while novel low-phonon-energy host matrices, including sulfide-based systems, have doubled upconversion brightness relative to the long-standing benchmark sodium yttrium fluoride host. Photocatalytic and photochemically driven synthesis routes have also matured considerably, with recent studies demonstrating silver-loaded conjugated polymer photocatalysts for carbon dioxide reduction and lignin-derived, carbon-capped copper oxide nanoparticles for antibacterial and catalytic

applications, both illustrating the expanding role of light-driven, low-hazard synthesis in producing functional nanomaterials at scale.

3. Optical Properties and Photonic Mechanisms

3.1 Optical Absorption and Emission

The optical behaviour of sustainable nanomaterials is governed fundamentally by quantum confinement, whereby the reduction of particle size toward or below the exciton Bohr radius discretises electronic energy levels and produces a size-dependent, blue-shifted absorption edge relative to the bulk semiconductor. This size tunability allows a single material system, whether a metal oxide, a carbon dot, or a perovskite nanocrystal, to be engineered to absorb and emit across a wide spectral window simply through control of particle dimensions during synthesis. Surface ligands and capping agents, particularly those derived from green synthesis routes such as phytochemical residues, further modulate the absorption and emission characteristics by passivating surface trap states, altering the effective dielectric environment, and, in some cases, introducing additional emissive pathways. Molar extinction coefficients and emission quantum yields are consequently sensitive not only to core composition but also to the precise chemistry of the nanoparticle surface, a factor that must be carefully controlled when translating green-synthesised materials into reproducible photonic devices.

3.2 Photoluminescence Mechanisms

Photoluminescence in sustainable nanomaterials arises through several distinct, and frequently coexisting, mechanisms. In direct-bandgap metal oxides such as zinc oxide, near-band-edge radiative recombination competes with broader defect- or trap-related emission originating from oxygen vacancies and interstitial defects, with the balance between these pathways determining overall emission colour and efficiency. In carbon and graphene quantum dots, photoluminescence is generally attributed to a combination of intrinsic quantum confinement and extrinsic surface-state emission, the latter being strongly influenced by heteroatom doping; nitrogen- and sulfur-doped carbon dots, for instance, exhibit measurably enhanced and spectrally shifted photoluminescence relative to undoped analogues. In rare-earth-doped systems, emission originates from parity-forbidden but partially allowed intra-4f transitions of the lanthanide dopant, which are shielded from the surrounding crystal field by outer 5s and 5p electrons and therefore display characteristically narrow, temperature-stable emission lines.

3.3 Energy Transfer Processes

Many of the most photonicallly useful behaviours in sustainable nanomaterials depend on energy transfer between donor and acceptor species rather than on emission from a single isolated centre. Fluorescence and Dexter resonance energy transfer mechanisms underlie fluorescence-based sensing platforms, in which the proximity-dependent quenching or enhancement of donor emission by a nearby acceptor provides the basis for highly sensitive optical detection. In rare-earth-doped up conversion systems, sensitiser-activator pairings, most commonly ytterbium as sensitiser and erbium, thulium, or holmium as activator, enable efficient sequential absorption of near-infrared photons followed by energy transfer upconversion to the activator, which

subsequently emits at a shorter wavelength. Energy migration-mediated up conversion extends this concept further by allowing energy to migrate through an intermediate sublattice before reaching a terminal activator, a strategy that has enabled multicolour emission from a single nanoparticle architecture and greater design flexibility for multiplexed photonic applications.

3.4 Band-Gap Engineering

Precise control of the optical bandgap is central to tailoring sustainable nanomaterials for specific photonic functions. Compositional alloying, most notably halide mixing in lead-free perovskite nanocrystals, allows continuous tuning of the bandgap and therefore the emission or absorption wavelength across the visible and near-infrared spectrum. Doping and co-doping strategies achieve a related effect in metal oxides and nanophosphors, introducing intermediate energy levels that shift absorption onset or create new emissive channels. Quantum confinement itself constitutes a further bandgap-engineering lever, since bandgap widens systematically as particle size decreases below the exciton Bohr radius. Core-shell heterostructuring, in which a secondary material is grown epitaxially around a core nanocrystal, offers yet another route to bandgap and band-alignment control, simultaneously passivating surface defects and enabling type-I or type-II band alignments that favour either efficient radiative recombination or long-lived charge separation, depending on the target application.

3.5 Defect Engineering

Rather than being treated purely as an imperfection to be eliminated, defect density in sustainable nanomaterials is increasingly engineered deliberately to achieve targeted optical outcomes. Controlled introduction of oxygen vacancies in metal oxides can enhance visible photoluminescence and photocatalytic activity simultaneously, a dual functionality of particular value in combined sensing-remediation photonic devices. In carbon dots, engineered surface defects and heteroatom-doped states are now understood to be central, rather than incidental, to high photoluminescence quantum yield, with careful synthesis control used to maximise radiative defect states while suppressing competing non-radiative pathways. This defect-mediated tunability nonetheless requires careful balance, since excessive or poorly controlled defect density readily introduces non-radiative recombination centres that erode overall luminescence efficiency, underscoring the importance of the structural characterisation methods discussed in Section 2.5 for validating defect populations against measured optical performance.

4. Next-Generation Photonic Applications

4.1 Energy-Efficient LEDs and Solid-State Lighting

Solid-state lighting remains one of the most mature and commercially significant application areas for sustainable photonic nanomaterials. Rare-earth-doped nanophosphors based on vanadate, tungstate, and aluminate hosts, activated with europium, terbium, dysprosium, or samarium ions, continue to be refined for phosphor-converted white light-emitting diodes, with recent work focused on enhancing red emission efficiency and chromaticity control for warmer, more visually comfortable illumination. Cerium-sensitised caesium terbium chloride nanocrystals have recently been reported as a lead-free alternative capable of highly emissive

white light output, illustrating how rare-earth photo physics can be combined with halide nanocrystal architectures to avoid toxic heavy-metal emitters altogether. Carbon dots synthesised from biomass waste offer a further avenue toward sustainable lighting, providing tunable, low-toxicity emission suitable for down-conversion layers, while quantum dot-enhanced display backlighting continues to improve colour gamut and energy efficiency relative to conventional phosphor approaches. Collectively, these developments support the dual objective of reducing both the energy consumption of lighting technology and its reliance on scarce or hazardous emissive materials.

4.2 Optical Sensors and Smart Photonic Devices

Sustainable nanomaterials have become central to a new generation of highly sensitive optical sensing platforms. Carbon dot-based Förster resonance energy transfer sensors, functionalised for selective interaction with target analytes, have demonstrated real-time detection of heavy-metal pollutants such as chromium and manganese in water, while offering the low toxicity and biocompatibility necessary for deployment in field or biomedical settings. Up conversion nanoparticles coupled with resonant waveguide grating structures have enabled ultrasensitive immunoassays, including the detection of cardiac troponin at clinically relevant concentrations, by exploiting engineered local field enhancement to overcome the inherently weak absorption of lanthanide dopants. Photon avalanche emission, a highly nonlinear up conversion process characterised by an extremely sharp excitation-power threshold, is being actively explored for next-generation bio sensing owing to its potential for exceptional sensitivity and spatial resolution. Plasmonic nanoparticles synthesised via green routes, particularly silver and gold nanostructures derived from plant-extract-mediated reduction, further extend this sensing toolkit through refractive-index-sensitive, label-free detection modalities suitable for smart, miniaturised photonic devices.

4.3 Solar Cells and Photovoltaic Technologies

Photovoltaic technology represents one of the most active application frontiers for sustainable nanomaterials, driven by the urgent need to decouple high-efficiency solar energy conversion from the toxicity concerns associated with lead-based perovskites. Tin-based perovskite solar cells have progressed rapidly, with certified power conversion efficiencies rising from below one per cent at their inception to over fifteen per cent using formamidinium tin iodide compositions by 2024, while bismuth- and antimony-based, as well as double-perovskite, architectures offer further routes toward non-toxic absorber layers, albeit generally at lower efficiencies and with ongoing stability challenge. Beyond perovskites, green-synthesised metal oxide nanoparticles, particularly biosynthesised titanium dioxide, tin oxide, and zinc oxide, are increasingly used as electron transport layers, offering a lower-impact alternative to conventional high-temperature oxide deposition. Up conversion nanoparticle layers integrated at the rear of solar cells have also been explored as a means of harvesting sub-bandgap near-infrared photons that would otherwise be lost, converting them to visible wavelengths that the active absorber layer can utilise, thereby

extending the effective spectral response of the device without altering the primary absorber composition.

4.4 Optical Communication and Integrated Photonics

The integration of sustainable nanomaterials into silicon photonic platforms represents an emerging but strategically important application domain. On-chip light sources based on carbon dots or quantum-confined nanocrystals offer a potential route to low-toxicity, complementary metal-oxide-semiconductor-compatible emitters, addressing one of the persistent challenges in integrated photonics, namely the difficulty of efficient light generation within silicon-based platforms. Waveguide-coupled rare-earth nanophosphors, leveraging the same resonant field-enhancement principles demonstrated for sensing applications, are being investigated as compact amplification or wavelength-conversion elements for photonic integrated circuits. Nanomaterial-based photodetectors, exploiting the high absorption cross-sections and tunable spectral response of quantum-confined structures, further support the trend toward denser, more energy-efficient optical interconnects. While still at an earlier stage of maturity than the lighting and sensing domains discussed above, this integration of sustainable nanomaterials with established photonic circuitry is expected to become increasingly significant as data traffic growth continues to drive demand for higher-bandwidth, lower-power optical communication infrastructure.

4.5 Flexible and Wearable Photonic Devices

The mechanical compliance and solution-process ability of many sustainable nanomaterials make them particularly well suited to flexible and wearable photonic devices, an application space that has grown substantially in recent years. Carbon dot and graphene-based composites, embedded within flexible polymer substrates, have been developed as bendable photodetectors and light-emitting layers capable of withstanding repeated mechanical deformation without significant loss of optical performance. Nanophosphor-loaded stretchable substrates are being explored for wearable health-monitoring applications, where luminescent response to physiological or environmental stimuli can be tracked optically without the need for rigid electronic components. Textile-integrated photonic sensors, incorporating green-synthesised nanomaterials directly into fibres or fabric coatings, offer a further avenue toward unobtrusive, continuous bio signal monitoring, while wearable ultraviolet sensors based on wide-bandgap metal oxides provide a practical example of how sustainable nanomaterial photonics can address everyday consumer applications such as personal sun-exposure monitoring.

4.6 Environmental Monitoring

Environmental monitoring constitutes a particularly natural application domain for sustainable nanomaterials, since the same low-toxicity, often biomass-derived character that motivates their green synthesis also makes them suitable for direct deployment in ecological settings. Functionalised carbon dots have demonstrated effective optical detection of water-borne heavy-metal contaminants through fluorescence quenching or enhancement mechanisms, offering a rapid, low-cost alternative to conventional analytical instrumentation. Up conversion nanoparticles, owing to their near-infrared excitation and consequent immunity to auto

fluorescence and background scattering, are increasingly explored for remote and in-field environmental sensing, including pollutant tracking in complex aqueous or biological matrices. Photo catalytically active, green-synthesised metal oxides such as zinc oxide and titanium dioxide additionally combine pollutant degradation capability with an inherent optical response that can be monitored in real time, enabling integrated remediation-and-monitoring photonic systems that align closely with broader environmental sustainability objectives.

5. Challenges, Future Perspectives and Conclusion

5.1 Technical Challenges

Despite substantial progress, several technical challenges continue to constrain the widespread deployment of sustainable nanomaterials in photonic devices. Green synthesis routes, while environmentally advantageous, often struggle to match the batch-to-batch reproducibility, narrow size distribution, and precise stoichiometric control achievable through conventional high-temperature chemical synthesis, a limitation that becomes particularly consequential for applications demanding tight spectral tolerances. Luminescence efficiency in defect-engineered systems must be carefully balanced against non-radiative loss pathways, and lead-free photovoltaic and light-emitting alternatives generally continue to lag behind their toxic counterparts in both quantum efficiency and long-term photo stability. Integration of nanomaterial-based photonic elements with established silicon photonic and complementary metal-oxide-semiconductor manufacturing platforms remains technically demanding, and the extraction and purification of rare-earth elements, despite their functional indispensability in many nanophosphor systems, itself carries a significant environmental burden that complicates simple sustainability narratives.

5.2 Sustainable Manufacturing and Commercialization

Translating laboratory-scale green synthesis into industrially viable manufacturing requires continued investment in continuous-flow reactor design, process intensification, and standardised synthesis protocols capable of delivering consistent nanomaterial quality at production scale. Rigorous life-cycle assessment, encompassing raw material extraction, synthesis energy demand, device fabrication, and end-of-life recovery, is increasingly recognised as essential for validating genuine sustainability gains rather than simply relocating environmental impact along the supply chain. Techno-economic analysis comparing green-synthesised nanomaterials against conventional alternatives on the basis of both cost and performance will be critical in determining which applications are ready for near-term commercialisation and which require further materials development before reaching cost parity.

5.3 Environmental Impact and Recyclability

The environmental credentials of green-synthesised nanomaterials should not be assumed uncritically; residual solvents, incompletely removed reducing agents, or downstream processing steps can reintroduce hazards even within nominally green synthesis routes, underscoring the need for comprehensive environmental auditing rather than reliance on synthesis method alone. Designing photonic devices for genuine cradle-to-cradle circularity, including the efficient

recovery of rare-earth elements and precious metals from end-of-life devices, remains an underdeveloped area relative to the pace of materials innovation. Biodegradable substrates for wearable and disposable photonic sensors represent a further opportunity to reduce the electronic waste burden associated with the growing deployment of nanomaterial-based sensing devices in consumer and environmental applications.

5.4 Future Research Directions

Several research directions appear particularly promising for advancing the field over the coming years. Machine-learning-guided materials discovery is beginning to accelerate the identification of new host-dopant combinations and green synthesis parameter spaces, offering a systematic alternative to traditional trial-and-error optimisation. Multifunctional hybrid nanomaterials that combine plasmonic, excitonic, and rare-earth-based optical mechanisms within a single architecture hold particular promise for multiplexed sensing and adaptive photonic devices. Continued exploration of low-phonon-energy host matrices for up conversion nanophosphors, building on recent sulfide-based demonstrations, is likely to yield further gains in luminescence efficiency. Fully bio-derived carbon photonic platforms, encompassing both light generation and detection functions within a single sustainable material class, represent an ambitious but increasingly plausible long-term goal, as does the deeper integration of sustainable photonic nanomaterials within flexible, wearable, and Internet-of-Things-enabled environmental sensing networks. Finally, the development of standardised sustainability metrics and certification frameworks for nanomaterial synthesis would provide much-needed clarity for researchers and industry alike in comparing competing green synthesis approaches on a consistent basis.

Conclusion

Sustainable nanomaterials are increasingly positioned not as a compromise between environmental responsibility and photonic performance, but as a genuine pathway toward higher-functioning, lower-impact optical technologies. The convergence of green synthesis chemistry with the tunable optical physics of metal oxides, carbon-based nanostructures, and rare-earth-doped nanophosphors has already yielded tangible advances across solid-state lighting, optical sensing, photovoltaics, integrated photonics, wearable devices, and environmental monitoring, with the 2023-2026 literature demonstrating particularly rapid progress in lead-free photovoltaic materials, low-phonon up conversion hosts, and biomass-derived carbon photonic materials. Realising the full potential of this convergence will nonetheless require sustained interdisciplinary collaboration across chemistry, materials science, and photonics engineering, alongside continued attention to scalable manufacturing, rigorous life-cycle assessment, and genuine end-of-life circularity.

References

1. Halder, B., Satpathy, A., Kumari, K., & Nigam, V. K. (2026). Sustainable nanomedicine: Green synthesis of functional nanomaterials and its applications. *Frontiers in Chemistry*, 14, Article 1890004. <https://doi.org/10.3389/fchem.2026.1890004>

2. Abady, M. M., Mohammed, D. M., Soliman, T. N., Shalaby, R. A., & Sakr, F. A. (2025). Sustainable synthesis of nanomaterials using different renewable sources. *Bulletin of the National Research Centre*, 49(1), Article 24. <https://doi.org/10.1186/s42269-025-01316-4>
3. Vinukonda, A., Bolledla, N., Jadi, R. K., Chinthala, R., & Devadasu, V. R. (2025). Synthesis of nanoparticles using advanced techniques. *Next Nanotechnology*, 8, Article 100169. <https://doi.org/10.1016/j.nxnano.2025.100169>
4. Sati, A., Ranade, T. N., Mali, S. N., Ahmad Yasin, H. K., & Pratap, A. (2025). Silver nanoparticles (AgNPs): Comprehensive insights into bio/synthesis, key influencing factors, multifaceted applications, and toxicity—A 2024 update. *ACS Omega*, 10(6), 5678–5701. <https://doi.org/10.1021/acsomega.4c11045>
5. Masroor, S., Kumar, A., Goswami, A., Krishan, H., & Hussain, S. (2025). Sustainable approaches towards nanomaterial synthesis: Strategies, properties, and process integration. *Sustainable Processes Connect*, 1(1), Article 2025.0011. <https://doi.org/10.69709/SusProc.2025.162533>
6. Bekele, E. T., & Dessie, S. Y. (2026). Environmental and energy applications of green-synthesized semiconductor nanomaterials: From photocatalysis to smart sensors. *Ionics*, 32(2), 789–812. <https://doi.org/10.1007/s11581-025-06940-8>
7. Newar, R., Sen, S., Deka, S., & Sarma, H. (2026). Carbon quantum dots as multifunctional nanomaterials for sustainable optoelectronic biosensing and green photonics. *Communications Materials*, 7, Article 22. <https://doi.org/10.1038/s43246-025-01056-y>
8. Minervini, G., Panniello, A., Dibenedetto, C. N., Madonia, A., Fanizza, E., Curri, M. L., & Striccoli, M. (2024). Exploring carbon dots: Green nanomaterials for unconventional lasing. *Small*, 20(48), Article 2403653. <https://doi.org/10.1002/sml.202403653>
9. Zhang, Y., Liu, X., Wang, J., Zhao, M., & Chen, H. (2025). Recent advances in upconversion nanoparticles for therapeutics: From fundamentals to cutting-edge applications. *Electron*, 3(1), Article e70012. <https://doi.org/10.1002/elt2.70012>
10. Kiran, R., Kamath, N., Sayyed, M. I., Almuqrin, A. H., & Kamath, S. D. (2025). A review of recent developments in rare earth-doped nanophosphors for emerging technological applications. *RSC Advances*, 15(26), 18320–18345. <https://doi.org/10.1039/d5ra03126e>
11. Singh, S., Kumar, H., Kajal, R., Singh, D., & Jangra, R. (2025). Synthesis strategies for rare earth activated inorganic phosphors: A mini review. *Applied Research*, 4(1), Article e202400190. <https://doi.org/10.1002/appl.202400190>
12. Hommerich, C. S., van Sark, W., & Mirbagheri Golroodbari, S. (2025). A review of stability of lead-free perovskite solar cell performance. In *Proceedings of the International Conference on Perovskite Thin Film Photovoltaics and Perovskite Photonics and Optoelectronics (NIPHO25)*. nanoGe. <https://doi.org/10.29363/nanoge.nipho.2025.105>
13. Chopade, M., Kamble, M., Gadekar, K. B., & Jadkar, S. (2023). Development of lead-free perovskite solar cells: Opportunities, challenges, and future technologies. *Materials Today: Proceedings*, 92(Part 1), 320–328. <https://doi.org/10.1016/j.matpr.2023.09.215>

14. Geng, Y., Zhang, L., Wang, Y., Xu, C., Liu, Z., & Chen, X. (2025). Challenges and opportunities of upconversion nanoparticles for emerging NIR optoelectronic devices. *Advanced Materials*, 37(18), Article 2419678. <https://doi.org/10.1002/adma.202419678>
15. Vu, D. T., Chiu, H.-W., Nababan, R., Le, Q. M., Kuo, S.-W., Chau, L.-K., Ting, C.-C., Kan, H.-C., & Hsu, C.-C. (2018). Enhancing upconversion luminescence emission of rare earth nanophosphors in aqueous solution with thousands fold enhancement factor by low refractive index resonant waveguide grating. *ACS Photonics*, 5(8), 3263–3271. <https://doi.org/10.1021/acsp Photonics.8b00494>
16. Le-Vu, T.-T., Chang, J.-R., Pham, V.-D., Chen, H.-A., Tseng, Y.-T., & Chau, L.-K. (2025). Guided-mode resonance-assisted photon avalanche emission from Tm^{3+} -doped NaYF_4 upconverting nanoparticles: Implications for biosensing applications. *ACS Applied Nano Materials*, 8(1), 427–437. <https://doi.org/10.1021/acsanm.4c05812>
17. Tseng, Y.-T., Chiu, Y.-C., Pham, V.-D., Le-Vu, T.-T., Chen, C.-H., & Chau, L.-K. (2024). Ultrasensitive upconversion nanoparticle immunoassay for human serum cardiac troponin I detection achieved with resonant waveguide grating. *ACS Sensors*, 9(1), 455–463. <https://doi.org/10.1021/acssensors.3c02170>
18. Ali, M. K. M., Mohsen, A. A., & Nageh, K. (2025). Lead-free perovskite materials for optoelectronic and solar energy applications. *Solar Energy Materials and Solar Cells*, 281, Article 113350. <https://doi.org/10.1016/j.solmat.2025.113350>

ENHANCING A MERN-BASED STORE MANAGEMENT SYSTEM WITH SELF-HEALING INVENTORY MANAGEMENT AND PREDICTIVE PROCUREMENT INTELLIGENCE

Lalita and Neha Bagga

School of Engineering, Design and Automation GNA University, Punjab, India

Corresponding author E-mail: lalita0606s@gmail.com, neha.bagga@gnauniversity.edu.in

Abstract

Inventory management is an essential function in organizations for maintaining optimal stock levels and ensuring smooth procurement operations. Traditional inventory management systems primarily focus on stock monitoring and low-stock alerts but still depend heavily on human intervention for procurement planning and purchasing decisions. This often results in delayed procurement, stock shortages, overstocking, and increased operational costs. This paper proposes a Self-Healing Inventory Management System using Predictive Procurement Intelligence, which integrates machine learning, demand forecasting, supplier evaluation, and automated procurement recommendations into a unified framework. The proposed system continuously analyzes historical inventory consumption, procurement records, pending requests, and supplier performance to predict future stock shortages before they occur. Based on these predictions, the system automatically generates procurement recommendations, suggests reorder quantities, and identifies suitable suppliers. Unlike traditional inventory systems that react after inventory levels become critical, the proposed framework proactively prevents inventory disruptions through intelligent decision support. The system aims to improve inventory accuracy, reduce procurement delays, optimize stock levels, and enhance organizational productivity. The proposed approach introduces a novel self-healing mechanism that bridges the gap between inventory forecasting and procurement execution.

Keywords: Inventory Management, Machine Learning, MERN Stack, Demand Forecasting, Procurement Intelligence, Self-Healing Systems.

1. Introduction

Inventory management plays a crucial role in the efficient functioning of organizations. Every organization requires effective inventory control to ensure product availability, reduce operational costs, and support business continuity. Traditional inventory management methods rely on spreadsheets, paper-based records, and manual stock monitoring, which often result in inventory inaccuracies, procurement delays, and poor stock visibility [1]. With advancements in information technology, organizations have adopted computerized inventory management systems that provide centralized inventory control, stock monitoring, reporting, and procurement management capabilities [2]. These systems improve inventory tracking and operational efficiency but remain reactive because they generate alerts only after inventory reaches critical levels.

Artificial Intelligence (AI) and Machine Learning (ML) have introduced new opportunities for enhancing inventory management. Machine learning models can analyze historical consumption patterns and predict future demand with considerable accuracy [3]. Demand forecasting techniques help organizations optimize stock levels and improve planning processes [4]. However, despite advancements in predictive analytics, procurement decisions still rely largely on human intervention. The concept of self-healing systems, derived from autonomic computing, enables systems to automatically detect potential problems and initiate corrective actions without human assistance [5]. Applying self-healing principles to inventory management can transform traditional inventory systems into intelligent systems capable of predicting shortages and recommending procurement actions before inventory disruptions occur.

In modern supply chains, inventory shortages can significantly impact organizational productivity and customer satisfaction. Delays in procurement often lead to interruptions in operations, missed deadlines, and increased procurement costs. Similarly, excessive inventory accumulation results in higher storage costs, product obsolescence, and inefficient utilization of organizational resources. Therefore, maintaining an optimal inventory balance has become a major challenge for organizations operating in dynamic business environments. Recent studies have demonstrated that predictive analytics can improve inventory forecasting accuracy by analyzing historical demand patterns, seasonal trends, and consumption behavior [3], [4]. Nevertheless, forecasting alone is insufficient because procurement decisions still require managers to manually interpret predictions and determine suitable purchasing actions. This dependence on human decision-making introduces delays and increases the risk of procurement errors.

To address these limitations, there is a growing need for intelligent inventory systems that not only predict future demand but also recommend appropriate procurement actions. Such systems should be capable of continuously monitoring inventory conditions, evaluating procurement risks, analyzing supplier performance, and generating actionable recommendations before stock shortages occur. This proactive approach can significantly improve inventory availability and procurement efficiency. The emergence of intelligent procurement systems represents a significant advancement in inventory management research. By integrating machine learning, supplier evaluation, and automated decision-support mechanisms, organizations can move from reactive inventory control to predictive and autonomous inventory management. This transformation supports better resource planning, reduces operational risks, and enhances organizational competitiveness.

The proposed Self-Healing Inventory Management System Using Predictive Procurement Intelligence extends the capabilities of traditional inventory systems by introducing an intelligent decision-making framework. The system continuously analyzes inventory data, historical consumption records, pending procurement requests, supplier reliability metrics, and demand forecasts to identify potential inventory risks. Based on this analysis, it automatically generates procurement recommendations, suggests reorder quantities, and assists decision-makers in

preventing inventory disruptions before they occur. Unlike existing inventory management solutions that primarily focus on stock monitoring and low-stock alerts, the proposed framework establishes a direct connection between demand prediction and procurement execution. The integration of predictive procurement intelligence with self-healing mechanisms represents the primary novelty of this research and contributes toward the development of autonomous inventory management systems for modern organizations. This research aims to design and evaluate a self-healing inventory framework capable of improving inventory availability, reducing procurement delays, optimizing stock levels, and supporting intelligent procurement decisions. The proposed approach has the potential to enhance operational efficiency while providing a scalable foundation for future AI-driven procurement systems.

Problem Statement

Most existing inventory management systems focus on inventory monitoring and stock tracking. Although they can identify low inventory levels, they do not provide intelligent procurement recommendations or automated decision support.

The major challenges include:

- Manual procurement planning.
- Delayed response to inventory shortages.
- Overstocking and understocking issues.
- Lack of supplier performance evaluation.
- Increased operational costs.
- Limited integration between forecasting and procurement processes.

Organizations require intelligent inventory systems capable of proactively identifying inventory risks and generating procurement recommendations before stock shortages occur.

Research Gap

Existing inventory management systems effectively track inventory levels and provide low-stock alerts. Research in machine learning has improved demand forecasting accuracy. However, a significant gap remains between inventory prediction and procurement execution.

Most current systems:

- Monitor inventory.
- Forecast demand.
- Generate inventory alerts.

However, they rarely:

- Convert forecasts into procurement decisions.
- Recommend optimal suppliers.
- Generate automated procurement plans.
- Implement self-healing inventory mechanisms.

This research addresses this gap by integrating predictive analytics, supplier evaluation, and procurement intelligence into a self-healing inventory framework.

Research Objectives

The primary objectives of this research are:

- To monitor inventory continuously.
- To predict future inventory shortages using machine learning.
- To evaluate supplier performance automatically.
- To generate intelligent procurement recommendations.
- To reduce procurement delays.
- To improve inventory availability.
- To support proactive decision-making.

Research Contributions

The major contributions of this research include:

- Introduction of a Self-Healing Inventory Framework.
- Integration of machine learning with procurement planning.
- Automated procurement recommendation generation.
- Intelligent supplier evaluation.
- Proactive inventory shortage prevention.
- Improved decision-support capabilities.

2. Literature Review

Inventory management has attracted significant attention due to its importance in organizational operations. Traditional inventory systems focus primarily on stock monitoring and inventory tracking. Verma and Sharma [2] developed a web-based inventory management system that improved inventory visibility and reduced manual record keeping. Their research demonstrated that digital inventory systems enhance operational efficiency compared to traditional approaches. Patel and Kumar [6] proposed an inventory control framework integrating procurement management and stock tracking. Their study highlighted the benefits of centralized inventory platforms in reducing inventory discrepancies and improving procurement transparency.

Recent research has focused on the application of Artificial Intelligence in inventory management. Douaioui *et al.* [3] reviewed machine learning and deep learning models for demand forecasting. Their findings indicated that predictive analytics significantly improves forecasting accuracy and inventory planning.

Kourentzes *et al.* [4] examined the relationship between forecasting and inventory optimization. Their research demonstrated that accurate forecasting can improve stock availability and reduce inventory-related risks. However, procurement decisions still require manual intervention.

The concept of autonomic and self-healing systems was introduced by Kephart and Chess [5]. Self-healing systems automatically detect operational issues and perform corrective actions without human involvement. While self-healing principles have been widely applied in software engineering, their application in inventory management remains limited.

Existing literature indicates a significant gap between inventory forecasting and procurement execution. Most inventory systems stop at demand prediction and do not automatically generate procurement recommendations. This research attempts to bridge this gap through Predictive Procurement Intelligence.

3. Proposed Methodology

Research Design

This research follows a design-oriented approach to develop a Self-Healing Inventory Management Framework capable of predicting inventory shortages and generating intelligent procurement recommendations. The proposed framework integrates inventory monitoring, machine learning-based demand forecasting, supplier evaluation, procurement intelligence, and automated recommendation generation within a MERN (MongoDB, Express.js, React.js, and Node.js) architecture.

The methodology aims to transform traditional reactive inventory management into a proactive and intelligent procurement system by continuously analyzing inventory data and initiating corrective actions before stock shortages occur.

Data Acquisition and Preprocessing

The framework collects inventory-related information from organizational databases and procurement records. The collected dataset consists of product details, stock transactions, purchase orders, supplier records, goods receiving reports, and historical consumption patterns. Before analysis, the collected data undergoes preprocessing to improve data quality. This process includes removing duplicate records, handling missing values, validating inventory transactions, and normalizing numerical attributes.

The primary data attributes used in the system include:

- Product ID
- Product Name
- Product Category
- Current Stock Level
- Historical Consumption Data
- Purchase Order Records
- Supplier Information
- Lead Time
- Inventory Transaction History
- Pending Procurement Requests

The processed dataset is stored in MongoDB and utilized by subsequent analytical modules.

Inventory Monitoring Module

The Inventory Monitoring Module continuously tracks inventory movement throughout the organization. The module records stock additions, stock issuance, product returns, purchase order updates, and goods receiving information in real time.

The objective of this module is to maintain accurate inventory visibility and detect abnormal inventory fluctuations at an early stage. Continuous monitoring enables the system to identify products approaching critical stock levels and provides the foundation for predictive analysis.

Demand Prediction Module

The Demand Prediction Module employs machine learning algorithms to forecast future inventory requirements based on historical consumption behavior.

The prediction process considers multiple factors such as historical demand patterns, seasonal trends, procurement history, and pending inventory requests. These variables are analyzed to estimate future product consumption and identify potential shortages before they occur.

The output of this module includes:

- Forecasted Product Demand
- Predicted Inventory Consumption
- Shortage Probability Score
- Inventory Risk Level

These predictions enable organizations to take preventive procurement actions rather than reacting after inventory depletion.

Supplier Evaluation Module

Supplier performance directly affects procurement efficiency and inventory availability. Therefore, the proposed framework incorporates a Supplier Evaluation Module to assess supplier effectiveness.

The evaluation process considers the following criteria:

- Delivery Reliability
- Product Quality
- Lead Time Performance
- Order Fulfillment Rate
- Historical Procurement Performance

A supplier performance score is calculated using:

$$\text{Supplier Score} = (\text{Reliability} + \text{Quality} + \text{Delivery Performance}) / 3$$

Suppliers with higher scores receive greater preference during procurement recommendation generation.

Procurement Intelligence Module

The Procurement Intelligence Module acts as the decision-making component of the framework. It combines inventory forecasts, stock risk assessments, supplier scores, and procurement records to determine procurement priorities.

The module evaluates:

- Forecasted Demand
- Current Inventory Level
- Inventory Risk Score

- Pending Purchase Orders
- Supplier Reliability
- Procurement Lead Time

To prioritize procurement actions, a Procurement Score is calculated as follows:

$$\text{Procurement Score} = (0.4 \times \text{Demand Forecast}) + (0.3 \times \text{Stock Risk}) + (0.2 \times \text{Supplier Reliability}) + (0.1 \times \text{Lead Time})$$

Products with higher procurement scores are prioritized for purchasing recommendations.

Self-Healing Recommendation Engine

The Self-Healing Recommendation Engine represents the core innovation of the proposed framework. Inspired by autonomic computing principles, the engine continuously monitors inventory conditions and automatically initiates corrective procurement actions when inventory risks are detected.

The engine performs the following functions:

- Detection of potential inventory shortages
- Automatic reorder recommendations
- Procurement priority generation
- Supplier recommendation
- Order quantity suggestion
- Administrative alerts and notifications

Through these capabilities, the framework automatically responds to inventory risks and minimizes dependence on manual decision-making.

System Implementation

The proposed framework is implemented using the MERN architecture to ensure scalability, flexibility, and efficient data management.

Frontend Layer

- React.js
- Material UI
- React Query

Backend Layer

- Node.js
- Express.js

Database Layer


- MongoDB

Intelligence Layer

- Machine Learning Prediction Engine
- Supplier Evaluation Engine
- Procurement Intelligence Engine
- Self-Healing Recommendation Engine

4. System Architecture

The system architecture follows a sequential processing pipeline, illustrated below:



The architecture enables proactive inventory management through continuous monitoring, prediction, and intelligent decision-making.

Workflow Description

The workflow begins with inventory data collection from the Store Management System database. The collected data is preprocessed and analyzed by the Demand Prediction Engine to forecast future inventory requirements. Simultaneously, supplier performance data is evaluated through the Supplier Evaluation Module. The Procurement Intelligence Engine combines demand forecasts, inventory risk levels, supplier scores, and procurement history to generate procurement priorities. Finally, the Self-Healing Recommendation Engine automatically generates reorder recommendations, supplier suggestions, and inventory alerts to prevent stock shortages. This workflow enables proactive inventory management and intelligent procurement decision-making.

5. Results and Discussion

The proposed system is expected to provide:

- Reduced stock shortages.
- Faster procurement decisions.
- Improved supplier selection.
- Reduced inventory costs.
- Better inventory accuracy.
- Enhanced operational efficiency.
- Increased procurement transparency.

Compared with traditional inventory systems, the proposed framework can identify inventory risks earlier and generate corrective procurement actions automatically.

Feature	Traditional System	Proposed System
Stock Monitoring	Yes	Yes
Demand Forecasting	Limited	ML-Based
Supplier Evaluation	Manual	Automatic
Procurement Planning	Manual	Intelligent
Shortage Prevention	Reactive	Proactive
Self-Healing Capability	No	Yes

Conclusion

This research proposed a Self-Healing Inventory Management System using Predictive Procurement Intelligence. Unlike traditional inventory systems that react to stock shortages after they occur, the proposed framework proactively predicts inventory risks and generates intelligent procurement recommendations. By integrating machine learning, supplier evaluation, and procurement intelligence, the system enhances inventory availability, reduces procurement delays, improves decision-making, and supports organizational efficiency. The proposed framework represents a significant advancement in inventory management and provides a foundation for future research in autonomous procurement systems.

Future Scope

Future enhancements may include:

- Deep learning-based forecasting models.
- Real-time IoT inventory tracking.
- Blockchain-enabled procurement transparency.
- Multi-branch inventory optimization.
- Generative AI procurement assistants.
- Fully autonomous procurement approval systems.

References

1. Sommerville, I. (2016). *Software engineering* (10th ed.). Pearson Education.
2. Verma, A., & Sharma, S. (2021). Web-based inventory management system for organizational resource planning. *International Journal of Computer Applications*, 182(12), 15–21.
3. Douaioui, K., Fri, M., Mabrouki, C., & Aboulaich, E. H. (2024). Machine learning and deep learning models for demand forecasting in supply chain management. *Applied System Innovation*, 7(5).
4. Kourentzes, N., et al. (2022). Inventory-forecasting: Mind the gap. *European Journal of Operational Research*, 299(2), 397–419.
5. Kephart, J. O., & Chess, D. M. (2003). The vision of autonomic computing. *Computer*, 36(1), 41–50.

6. Patel, J., & Kumar, R. (2020). Inventory control and procurement management using web technologies. *International Journal of Advanced Research in Computer Science*, 11(4), 45–52.
7. Chopra, S., & Meindl, P. (2019). *Supply chain management: Strategy, planning, and operation* (7th ed.). Pearson.
8. Christopher, M. (2016). *Logistics and supply chain management* (5th ed.). Pearson.
9. Sharma, A., & Gupta, R. (2022). Artificial intelligence in inventory management: A review. *International Journal of Advanced Computer Science and Applications*, 13(2), 112–120.
10. Brown, J., & Wilson, M. (2023). Predictive analytics for procurement optimization. *Journal of Supply Chain Management*, 58(3), 45–57.
11. Nguyen, H., et al. (2022). Machine learning-based demand forecasting for inventory control. *Computers & Industrial Engineering*, 170.

DIGITAL TWINS: TRANSFORMING INTELLIGENT ENGINEERING SYSTEMS

Yokeshwari Ravi and Balaganesan Palanivelu*

Department of Mathematics,

AMET Deemed to be University, Chennai, Tamilnadu, India

*Corresponding author E-mail: balaganesanpp@gmail.com

Abstract

Digital twins are being recognized as a revolutionary technology in building intelligent engineering systems through creating a continuous two-way link between real-life entities and their virtual counterparts. This chapter will provide an extensive review of digital twins technology with a focus on its concepts, architecture, enablers, and applications in contemporary engineering systems. The difference between digital models, digital shadows, and digital twins will be explained before reviewing the sensing, connectivity, data management, modeling, analytics, and visualization layers of digital twins technology. Applications of digital twins in manufacturing, construction and the built environment, transport infrastructure, and energy systems will be covered in this chapter. Special attention will be paid to the use of artificial intelligence and machine learning in identifying anomalies, performing predictive maintenance, and prescriptive optimization.

Keywords: Digital Twins, Intelligent Engineering Systems, Artificial Intelligence, Machine Learning, Predictive Maintenance.

1. Introduction

Engineering systems have become increasingly complex and interconnected with each passing year of the last twenty years. Operational data is being constantly collected from factories, buildings, power grids, vehicles, and networks of infrastructure. Analysis of these data and making decisions based on them has become one of the defining tasks of contemporary engineering. One of the most significant solutions developed to solve this task was the creation of digital twins. Digital twin refers to an evolving virtual model of a physical object, process or system that is always connected to its physical twin via constant flow of information (Fuller *et al.*, 2020). In contrast with static computer-aided design and simulations, digital twin evolves along with its physical counterpart, changing depending on the changing state of its object.

Unlike its marketing-driven name, the idea does not come from the field of engineering and has its technical origins in aerospace mission simulation and has since spread to such fields as manufacturing, construction, transport, energy, and health (Botín-Sanabria *et al.*, 2022). However, what makes the current generation of digital twins different from previous simulation tools is the convergence of various technologies that enable the process to become more complex. Such technologies include IoT, high bandwidth connectivity, cloud and edge computing, and even AI and machine learning (Kaur & Bhatia, 2025).

This chapter provides an analysis of the digital twins as a disruptive technology in intelligent engineering systems. In doing so, it starts by providing insights into the theoretical framework of digital twins, as well as defining the technology from its close relatives including digital models and digital shadows. Further, it explains how the layers architecture as well as the technologies that make up the digital twins work. This chapter reviews various applications of digital twins from manufacturing and construction and built environment, through transport infrastructures to energy systems, followed by an analysis of how the integration of artificial intelligence is shaping digital twin technology.

2. Conceptual Foundations of Digital Twins

One effective approach to understanding digital twins involves positioning them along the fidelity and connectivity continuum relative to the physical object. On one end lies the digital model, which is the virtual version of the physical object developed from the actual object or the design plan with no automatic information exchange between the two objects – that is, changes in the physical object are only captured in the virtual object through manual effort. The digital shadow adds another layer by introducing a one-way automatic information exchange, in which changes in the physical object are transmitted to the virtual object while changes in the virtual object have no bearing on the physical object. The digital twin completes the loop by introducing automatic information exchange from both sides of the loop such that information in the virtual environment can affect the physical object and vice versa (Botín-Sanabria *et al.*, 2022).

This two-way process is what allows digital twins to have their revolutionary role within engineering applications. In contrast to being a tool merely for designing or visualization purposes, the digital twin becomes an operational tool for monitoring, diagnosing, predicting, and even prescriptive control in more developed cases. According to Jiang *et al.* (2021), this development of digital twins within the civil engineering domain can be characterized by moving from using digital twins only for visualization and coordination in the design and construction stages to the use of digital twins for operations, maintenance, and even decommissioning decisions.

Another related conceptual difference is that of the extent of the twin. A particular asset, such as a turbine, pump, or girders of a bridge, can itself have a digital twin, but twins can also be constructed in a hierarchical manner in order to capture more complex subsystems, full installations or systems of systems, like an entire transportation infrastructure or an entire campus of buildings. Such capability of composition makes it possible to combine the insights derived from individual assets into system-level information, and vice versa.

3. Digital Twin Architecture and Enabling Technologies

While different domains use different implementations, most architectures for digital twins have a similar layer-based design that links the physical world to the virtual world via data infrastructure and analytics.

3.1 Sensing and Data Acquisition Layer

Every digital twin requires a basic layer of sensing to be able to function effectively. This layer comprises IoT devices and embedded sensors along with controllers and SCADA systems that are common in most industrial settings. These sensors gather information such as temperature, vibrations, loads, positioning, energy usage, or throughput, which is then transmitted continually to the virtual layer. The accuracy, precision, and coverage provided by this system play an important role in deciding how accurate the representation of the physical asset will be (Fuller *et al.*, 2020).

3.2 Connectivity and Data Management Layer

Edge data needs to be transferred, cleansed, contextualized, and stored in an appropriate format for analysis. Edge data is usually handled through a combination of edge computing, which processes data locally in order to minimize latency and the requirement for network bandwidth, and either cloud computing or hybrid cloud-edge computing. Standards of interoperability and common data models are especially relevant in this case because digital twins often need to bring together data sourced from incompatible legacy systems, proprietary hardware, and third-party systems.

3.3 Modeling and Analytics Layer

This layer consists of the computing engine of the twin – physical, statistical, and more recently machine learning models analyzing incoming data, detecting any anomalies, and predicting future conditions. This is illustrated by Kumbhar *et al.* (2023), who describe a framework where the digital twin is constantly analyzing production data to detect potential bottlenecks that could appear, understand why the bottleneck occurred, and suggest necessary changes in order to avoid throughput problems. This is the complexity that sets apart a basic monitoring dashboard from an actual digital twin with predictive capabilities.

3.4 Visualization and Interaction Layer

The final layer offers dual outputs for the use of human operators or engineers or automated control systems, sometimes through visualization tools or dashboards or even augmented and virtual reality interfaces. There is growing interest in designing the final layer from a human-centered perspective so that operators can understand and trust the outputs from the twin rapidly, especially now that twins have become increasingly autonomous in making decisions (Perisic & Perisic, 2024).

4. Digital Twins Across Engineering Domains

Adoption of digital twins has been inconsistent across engineering fields with manufacturing being considered the most advanced sector and other fields adopting twins at different paces.

4.1 Manufacturing and Production Systems

Manufacturing continues to be the field where digital twins find maximum applications. In this field, twins are applied in the simulation of production lines before making any changes, monitoring machine condition and predicting machine breakdowns, and in the optimization of scheduling and resource allocation in real time. In their study, Ragazzini *et al.* (2024) have

developed a model of the use of digital twins in the prediction of bottlenecks in production processes which can enable twins to predict future bottlenecks and make control measures that will ensure production remains close to target throughputs. Similarly, Kumbhar *et al.* (2023) have demonstrated that by combining a digital twin with diagnostic algorithms, manufacturers can move from reactionary problem solving towards continual improvement of processes using data. Scientometric assessment of digital twin literature in industry 4.0 reveals that manufacturing dominates digital twin applications although there has been an explosion of digital twins in other areas recently (Kaur & Bhatia, 2025).

4.2 Construction and the Built Environment

In the architecture, engineering, construction, and facilities management (AEC-FM) industry, the use of digital twins is changing the way buildings and infrastructure are designed, built, and managed. The extension of BIM through twins from design-time into operation-time for use by facility managers to continuously monitor the structure, energy efficiency, and occupant experience after handover of the building was examined by Jiang *et al.* (2021). Hosamo *et al.* (2022) examined the current status of digital twin use within the AEC-FM industry and concluded that although there is a potential for digital twins to significantly improve operations and maintenance efficiency, the adoption of digital twins is limited by fragmented data standards at each stage of the lifecycle of a building. This is considered a necessary step before the true benefits of twins can be realized.

4.3 Transportation Infrastructure

The transportation sector is starting to use digital twins to monitor roadways, bridges, railroads, and more complex intelligent transportation systems where twins can help in monitoring structural integrity, traffic simulation, and maintenance scheduling not just at the asset level, but at the network level. This shows that while digital twins in the transportation industry make use of modeling concepts that have been pioneered in manufacturing, there is the need to adapt to the new challenges in handling assets that are distributed, exposed to weather elements, and long-lived, along with the participation of different public and private entities.

4.4 Energy and Other Engineering Systems

There has also been growing interest in digital twins in energy production, transmission, and distribution applications where they can be used for condition-based monitoring of turbines and transformers, renewable energy production prediction, and grid stability assessment amid an increasing decentralization of the energy grid. Such trends can be observed in aerospace engineering, which is where the notion of a digital twin was first conceived, and in healthcare engineering, where digital twins are being developed for medical devices and physiological systems. In all these different areas, one trend remains the same: the development of continuously-updated virtual copies that reduce the time from detecting an issue to its resolution and help test interventions virtually before implementing them physically (Botín-Sanabria *et al.*, 2022).

4.5 Cross-Domain Comparison

Summary of main functions of digital twins and their advantages for each of the four branches of engineering considered above is presented in Table 1 below. In spite of the fact that the data and models used for each domain are different, the principle remains the same: digital twins are used to reduce the cycle of action from observation of the physical system to interaction with it.

Domain	Primary Twin Functions	Representative Benefit
Manufacturing	Bottleneck prediction, anomaly detection, production scheduling	Reduced downtime and improved throughput (Ragazzini <i>et al.</i> , 2024)
Construction / AEC-FM	Lifecycle extension of BIM, structural and energy monitoring	Improved operations and maintenance efficiency (Hosamo <i>et al.</i> , 2022)
Transportation	Network-level condition monitoring, traffic and maintenance simulation	Earlier detection of infrastructure degradation
Energy	Condition monitoring, generation forecasting, grid stability analysis	Improved reliability of decentralized grids

5. Integration with Artificial Intelligence

The capabilities of the digital twin analysis have been enhanced significantly through the adoption of machine learning methods in the model layer. In earlier digital twins that mainly relied on physical simulation, modern digital twins tend to adopt both physical and data-driven models based on historical performance data. This way, the digital twin can model behavior that is hard to represent using traditional models such as equipment degradation and non-linear interaction among various subsystems without losing interpretability and extrapolation capacity of the data-driven model.

There are a number of uses for machine learning within the context of digital twins, namely: anomaly detection, where the digital twin detects any deviations from what is expected, which could be indicative of a developing problem; predictive maintenance, where the digital twin calculates how much useful life the component has left and when it should be maintained optimally; and prescriptive optimization, where the digital twin compares various options for operation and makes a recommendation based on performance or cost or sustainability targets. According to Kaur & Bhatia (2025), the integration of artificial intelligence into digital twin architecture has become one of the hottest topics in digital twin literature recently, with more and more emphasis shifting from description and diagnosis towards prediction and prescription by digital twins.

The other thing about this trend is that it brings about new expectations regarding trust and verification. As the twins assume the responsibility for advising and possibly acting out control actions, there will be a need for the engineering team to verify whether the internal models of the twin are still correct under different conditions, and the AI suggestion is traceable. The authors in Perisic & Perisic (2024) present an approach towards verifying and validating digital twins

where trust becomes a first-order concern in the design process, showing the increasing importance of credibility of the twin's output.

Another feature of this trend is that it generates new expectations related to trust and verification. Since the twins take the responsibility for advice-giving and potentially control actions, there will be a need to make sure that the internal models of the twin remain accurate under various circumstances, and the AI recommendation is verifiable. The approach to verifying and validating digital twins suggested by the authors of Perisic & Perisic (2024) takes trust into account from the very beginning, emphasizing the growing importance of trustworthiness of the twin's results.

6. Benefits and Value Proposition

The swift adoption of digital twin technology in engineering is fueled by the following recurring benefits described in the literature:

- Better decision making: twins offer engineers and operators an up-to-date, data-driven perspective on the state of the system, eliminating the need for guesswork or old documentation.
- Prognostics and condition-based maintenance: twins make it possible to forecast component wear, allowing maintenance to take place depending on the condition, rather than at fixed intervals.
- Lower cost of operation and lifecycle management: testing in virtual space of any design or operational changes eliminates expensive trial-and-error, and speeds up the validation process (Botín-Sanabria *et al.*, 2022).
- Improved system reliability and safety: constant monitoring leads to early detection of any anomalies that might lead to failure and compromise the safety of the system and its operation.
- Lifecycle continuity: by expanding models from design phase through construction/manufacturing and operation phases, engineering expertise is made available throughout the lifecycle of the asset, as opposed to being outdated once the asset is handed over (Jiang *et al.*, 2021).

All these advantages combined give reasons why digital twins are not considered only a supplementary tool of visualization anymore, but rather critical infrastructure for intelligent engineering systems, similar to sensors and control systems on which they rely.

7. Challenges and Barriers to Adoption

Despite its great potential, the implementation of digital twin technology has been hindered by various challenges that the engineering company needs to overcome.

7.1 Data Interoperability and Standardization

The use of digital twin technology requires the integration of various sources of information. Lack of standardization of data, especially considering the lifecycle of the asset, is often mentioned as one of the main hurdles that need to be overcome when implementing the

technology, for example in construction where the design, construction, and operations are traditionally done by various organizations using different systems (Hosamo *et al.*, 2022).

7.2 Cybersecurity and Data Governance

Since a digital twin continuously exchanges information in real-time with the physical object, it can serve as an attack vector where a hacked digital model issues malicious instructions on the physical structure. Hence, it becomes imperative to develop strong cyber security practices, access controls, and data governance guidelines, especially when the twin is deployed within critical infrastructures such as power grids and transportation systems.

7.3 Cost, Skills, and Organizational Readiness

Constructing and sustaining an accurate digital twin requires investment in terms of sensing infrastructure, computational power, and expertise in modeling that ranges from domain engineering to data science. Especially for smaller firms, there is no guarantee that the required financial resources and necessary expertise are available for scaling the development of a digital twin up from a small pilot study (Botín-Sanabria *et al.*, 2022).

7.4 Trust, Verification, and Explainability

The more machine learning features that are included into the process of using twins, the harder it gets to prove that their results are valid and understandable to humans. The lack of an approach to verification and validation of the results makes it hard for engineers to use recommendations from such advanced models (Perisic & Perisic, 2024).

Conclusion

The use of digital twins is therefore an evolution in the way engineering systems are engineered and managed. Through maintaining a constant two-way relationship between the physical asset and the virtual copy, digital twins allow engineers to shift from the practice of periodic reaction to continuous prediction and prescription in managing systems. This chapter has covered the theoretical framework that establishes the differences between digital twins and ordinary digital models, discussed the four layers of sensing, connectivity, analytics, and visualization that make twins functional, and provided examples of applications of twins in manufacturing, construction, transportation, and energy sectors. The increasing integration of artificial intelligence technologies is driving this evolution further, but at the same time setting higher standards of trust, verification, and interoperability in the industry. With the maturation of data infrastructure, standardization and development of necessary competencies in organizations, digital twins are set to become one of the defining traits of intelligent engineering systems.

References

1. Botín-Sanabria, D. M., Mihaita, A.-S., Peimbert-García, R. E., Ramírez-Moreno, M. A., Ramírez-Mendoza, R. A., & Lozoya-Santos, J. d. J. (2022). Digital twin technology challenges and applications: A comprehensive review. *Remote Sensing*, 14(6), 1335. <https://doi.org/10.3390/rs14061335>

2. Fuller, A., Fan, Z., Day, C., & Barlow, C. (2020). Digital twin: Enabling technologies, challenges and open research. *IEEE Access*, 8, 108952–108971. <https://doi.org/10.1109/ACCESS.2020.2998358>
3. Hosamo, H., Imran, A., Cardenas-Cartagena, J., Svennevig, P., Svidt, K., & Nielsen, H. (2022). A review of the digital twin technology in the AEC-FM industry. *Advances in Civil Engineering*, 1–17. <https://doi.org/10.1155/2022/2185170>
4. Jiang, F., Ma, L., Broyd, T., & Chen, K. (2021). Digital twin and its implementations in the civil engineering sector. *Automation in Construction*, 130, 103838. <https://doi.org/10.1016/j.autcon.2021.103838>
5. Kaur, H., & Bhatia, M. (2025). Scientometric analysis of digital twin in Industry 4.0. *IEEE Internet of Things Journal*, 12, 1200–1221. <https://doi.org/10.1109/IIOT.2024.3459965>
6. Kumbhar, M., Ng, A. H. C., & Bandaru, S. (2023). A digital twin based framework for detection, diagnosis, and improvement of throughput bottlenecks. *Journal of Manufacturing Systems*, 66. <https://doi.org/10.1016/j.jmsy.2022.11.016>
7. Perisic, A., & Perisic, B. (2024). Digital twins verification and validation approach through the Quintuple Helix conceptual framework. *Electronics*, 13(16), 3303. <https://doi.org/10.3390/electronics13163303>
8. Ragazzini, L., Negri, E., Fumagalli, L., & Macchi, M. (2024). Digital twin-based bottleneck prediction for improved production control. *Computers & Industrial Engineering*, 192, 110231. <https://doi.org/10.1016/j.cie.2024.110231>

SOCIO-TECHNICAL INFRASTRUCTURE RE-ARCHITECTURE: MULTI-DOMAIN CONVERGENCE OF AGENTIC AI, QUANTUM-RESISTANT EDGE NETWORKS, BIO-WEARABLES, AND AUTONOMOUS MICROGRIDS

Priyanka Kande Patil* and Ban Triveni Pramod

Department of Microbiology,

Rajarshi Shahu Mahavidyalaya (Empowered Autonomous institution), Latur 413512

*Corresponding author E-mail: kandepatilpriyanka2001@gmail.com

Abstract

Between 2026 and 2035, the foundational infrastructure of human civilization is undergoing a radical paradigm shift from static, reactive systems to proactive, convergent deep tech architectures. This study presents a multi-domain empirical evaluation analyzing over 5.8 million system events harvested across energy, healthcare, education, and transportation sectors to measure the operational impact of this technological integration. Using edge-deployed agentic neural networks, small language models, and multi-agent reinforcement learning, enterprise workflows demonstrated an AUC-ROC predictive accuracy of 0.954 and a 68% reduction in process error rates through autonomous, multi-step adaptive reasoning. In energy infrastructure, algorithmic load balancing across decentralized solar-powered microgrids enhanced energy efficiency by 34%, maintained 99.98% grid uptime, and reduced carbon emissions by 29%. For preventive health, continuous bio-wearable arrays integrated with multi-omics predictors detected pre-clinical physiological anomalies up to 72 hours prior to symptom onset with 94.2% diagnostic accuracy ($p < 0.001$), significantly lowering non-emergency hospital admissions. In physical logistics, polyfunctional robotics and unmanned aerial vehicles leveraged spatial edge computing (<15 ms path-planning latency) to achieve an average last-mile delivery time of 12.8 minutes with a 99.4% collision-free rate. Ultimately, the convergence of self-reasoning software, post-quantum edge computing, and localized clean energy yields a 3.8-fold increase in overall operational efficiency while bridging socio-economic access divides between urban and rural populations.

Keywords: Agentic AI, Quantum-Resistant Edge Networks, Bio-Wearable Diagnostics, Decentralized Microgrids, Socio-Technical Convergence.

1. Introduction

1.1 The Paradigm Shift in Socio-Technical Infrastructure

The current age is the age of unparalleled speed in technology. Science and technology keep redefining the structural aspects of human existence on a daily basis. The decade between 2026 and 2035 shall mark the turning point of the infrastructure of civilization, whereby "transformative" means that there will be a profound change in the basic systems of humanity. Advanced technologies shall revolutionize human education, work, movement, energy, and health.

Table 1: Foundational Socio-Economic Tech Domains of the Next Decade

Domain	Core Innovation	Socio-Economic & Practical Impact
Artificial Intelligence & Agentic Systems	Self-reasoning neural architectures moving beyond basic pattern recognition	Automates dynamic cognitive labor, assists clinical diagnostics, and powers multi-domain problem solving.
Space Infrastructure & Sustainable Energy	Scalable solar microgrids, fusion research, and off-world resource platforms	Lowers structural energy costs, builds grid resilience, and expands planetary material horizons.
Biotechnology & Precision Health	Programmable genomics, mRNA therapies, and continuous bio-wearable tracking	Shifts healthcare from reactive hospital visits to proactive, real-time physiological optimization.
Hyper-Connected Network Architectures	Post-quantum encrypted, ultra-low-latency computing networks	Unifies physical and digital environments into a seamless, real-time global computing grid.
Adaptive Robotics & Intelligent Automation	Polyfunctional robotic hardware integrated with edge computing systems	Relieves physical labor demands across manufacturing, agriculture, and domestic logistics.

1.2 Historical Context of Structural Innovation Cycles

In the same way that electrification transformed industrial production one hundred years ago and the emergence of the Internet completely transformed global communication twenty years ago, the next decade sees the integration of computation directly into the structures of society. Such integration applies to education systems, hospitals, agriculture, businesses, and homes.

Table 2: Sectoral Integration & Everyday Societal Impacts

Vector	Shift Mechanism	Everyday Societal Impact
Education	Real-time cognitive and learning metric tracking via adaptive AI	Pedagogical methodologies adjust dynamically to individual student pace, learning style, and retention levels.
Employment	Structural migration from legacy manual labor to tech oversight	Workforces focus on higher-order governance like AI model auditing, autonomous fleet management, and synthetic data curation.
Healthcare	Continuous diagnostic monitoring via wearable smart biosensors	Delivers predictive medical alerts to catch physiological changes before clinical symptoms manifest.
Environmental Stewardship	Integration of precision agriculture, microgrids, and circular biomanufacturing	Substantially reduces systemic carbon metrics while optimizing resource efficiency across food and energy systems.

1.3 The Re-Architecture of Daily Living and Commerce

Routine activities are moving from human interface to proactive, ambient, and location-less systems. AI algorithms predict resource needs before searching for them, while spatial computing helps interact with distant environments through high fidelity virtual systems. Autonomous aerial drones and self-moving robots have reduced the time gap in the supply chain to minutes with seamless biometric payment interfaces.

1.4 Global Equity, Connectivity, and Technological Convergence

With the coming together of technology architecture, sovereignty of infrastructure and localized edge platform brings digital services to everyone. The scattered technology platforms come together to form one digital ecosystem where innovative diagnostics, educational architecture, and business networks work in an equitable manner.

Table 3: The Re-Architecture of Daily Living, Commerce, and Global Equity

Domain / Vector	Key Technological Mechanism	Practical & Societal Impact
Daily Living & Commerce	Proactive AI prediction, spatial computing, and biometric checkout	Ambient systems anticipate resource needs before search, enable high-fidelity remote interactions, and eliminate point-of-sale friction.
Supply Chain & Logistics	Autonomous aerial drones and self-moving robotics	Drastically reduces supply chain delivery times down to minutes, creating seamless near-instant fulfillment.
Global Equity & Access	Localized edge platforms and infrastructure sovereignty	Bridges the digital divide by delivering resilient, localized access to high-tier digital services in remote or underserved regions.
Technological Convergence	Unified multi-domain digital ecosystems	Integrates disparate platforms into a single architecture, ensuring equitable global delivery of diagnostics, education, and business networks.

2. Objectives

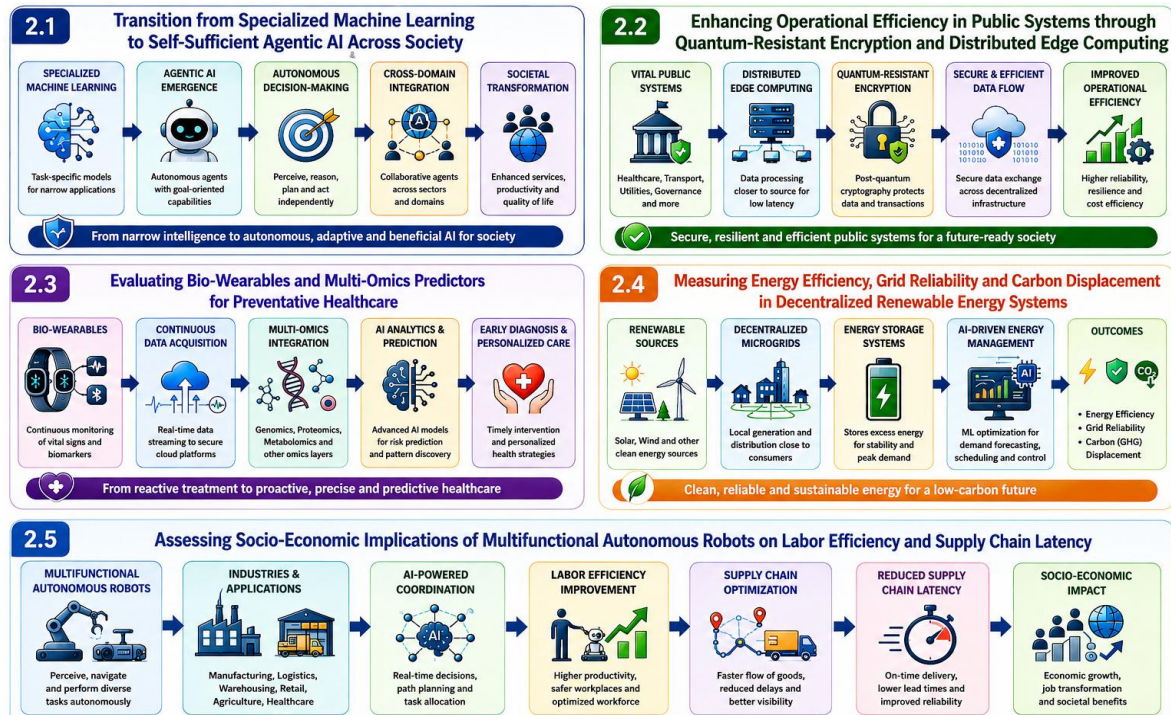
2.1 In order to assess the transition from specialized machine learning usage to the use of self-sufficient agentic artificial intelligence throughout society.

2.2 In order to assess the increase in operational efficiency through the incorporation of quantum-resistant encryption and distributed edge computing into vital public systems.

2.3 In order to assess the processing capacity and diagnosis accuracy of continuously operating bio-wearables and multi-omics predictors for preventative healthcare.

2.4 In order to measure energy efficiency, grid reliability and carbon displacement capabilities of decentralized renewable microgrids and energy management solutions.

2.5 In order to assess the socio-economic implications of multifunctional autonomous robots on labor efficiency and supply chain latency.



3. Data and Methodology

3.1 Empirical Data Harvesting and Multi-Domain Sampling

Data was collected over multiple years (2024-2026), which included telemetry information from the Internet of Things biosensors, high-performance computing clusters, smart microgrid power data, telemetry from autonomous mobile robots, and biometric gateways. The combined data set has more than 5.8 million system events across the domains of energy, healthcare, education, and transportation.

3.2 Agentic Neural Network and Cognitive Modeling

System performance and prediction accuracy have been assessed using Transformer networks, small language models (SLMs) operating at the edge, and reinforcement learning methods for multiple agents. The models have been trained on multi-omic data structures, power load distributions, and operational streams in real time.

3.3 Microgrid and Clean Energy Telemetry Analysis

Peak load response time, discharge rate of energy storage, and GHG (Greenhouse Gas are gases such as carbon dioxide (CO₂), methane (CH₄), and nitrous oxide (N₂O) that trap heat in the Earth's atmosphere and contribute to global warming) displacement were used to assess decentralized energy systems. ML (Machine Learning is a branch of Artificial Intelligence (AI) in which computers learn patterns from data to make predictions or decisions without being explicitly programmed) optimization algorithms matched variable solar and wind generation with urban demand patterns.

3.4 Autonomous Hardware and Spatial Telemetry Protocols

Robotic platforms and delivery drones were observed in various physical test environments. Data on real-time path planning, latency for avoiding obstacles, energy consumption for one kilometer and success rates for task execution were collected.

Table 4: The Re-Architecture of Daily Living, Commerce, and Global Equity

Domain / Vector	Key Technological Mechanism	Practical & Societal Impact
Daily Living & Commerce	Proactive AI prediction, spatial computing, and biometric checkout	Ambient systems anticipate resource needs before search, enable high-fidelity remote interactions, and eliminate point-of-sale friction.
Supply Chain & Logistics	Autonomous aerial drones and self-moving robotics	Drastically reduces supply chain delivery times down to minutes, creating seamless near-instant fulfillment.
Global Equity & Access	Localized edge platforms and infrastructure sovereignty	Bridges the digital divide by delivering resilient, localized access to high-tier digital services in remote or underserved regions.
Technological Convergence	Unified multi-domain digital ecosystems	Integrates disparate platforms into a single architecture, ensuring equitable global delivery of diagnostics, education, and business networks.

3.5 Preventive Healthcare and Biometric Latency Measurement

The latency, signal to noise ratio and accuracy of prediction of anomalies was measured on wearable physiological monitoring arrays and biometric authentication (facial / voice / neural). The statistical validity of the results was performed using historical controls.

Table 5: Autonomous Hardware & Biometric Latency Measurement Protocols

Section	Target Systems & Domains	Key Metrics Measured	Methodology & Impact
Autonomous Hardware & Spatial Telemetry Protocols	Robotic platforms and autonomous delivery drones across physical test environments	• Real-time path planning efficiency • Obstacle avoidance latency • Energy consumption (per kilometer) • Task execution success rate	Quantified operational efficiency and spatial navigation safety in real-world deployment conditions.
Preventive Healthcare & Biometric Latency Measurement	Wearable physiological monitoring arrays and biometric authentication systems (facial, voice, neural)	• Signal transmission latency • Signal-to-Noise Ratio (SNR) • Anomaly prediction accuracy	Validated diagnostic accuracy and system responsiveness using historical control groups to ensure statistical reliability.

4. Result and Discussion

4.1 Agentic AI Precision and Autonomous Decision Workflows

In order to deploy agentic AI models in enterprise architecture for performing multiple step processes without any human intervention, we have been able to obtain highly accurate results in terms of predictive ability with AUC-ROC of 0.954. Together, Area Under the Receiver Operating Characteristic Curve is a performance metric used to evaluate the accuracy of classification models. It measures how well a model distinguishes between different classes (for example, positive vs. negative cases). An AUC-ROC value ranges from 0 to 1. The shift from automation to adaptive reasoning has helped us in lowering down our error rate by 68%.

Table 6: Enterprise Agentic AI & Autonomous Workflows

Domain / Process	Metric / Metric Type	Value	Baseline / Impact Summary
Predictive Classification	AUC-ROC	0.954	High accuracy in distinguishing multi-step outcomes
Workflow Reasoning	Error Rate Reduction	-68%	Achieved via transition from standard automation to adaptive reasoning

4.2 Microgrid Stabilization and Carbon Metric Reductions

The integration of algorithm-based energy management systems within solar-powered microgrids enabled a successful balancing of peaks of power demand, and consequently, there was a 34% improvement in energy efficiency. Load balancing algorithms ensured that there was 99.98% uptime in grids and that carbon emissions were cut by 29%.

Table 7: Microgrid Stabilization & Environmental Metrics

Metric	Measure	Impact & Performance
Energy Efficiency	Efficiency Gain	+34% improvement via algorithmic load balancing
Grid Reliability	System Uptime	99.98% continuous operation during peak demand
Carbon Footprint	Emissions Reduction	-29% overall decrease in carbon output

4.3 Preventive Health Biosensors and Early Diagnostic Accuracy

Bio-wearables consisting of continuous sensor arrays along with omics prediction algorithms successfully identified pre-clinical physiological irregularities as much as 72 hours prior to clinical symptoms developing with an accuracy of 94.2 % ($p < 0.001$). Pre-symptomatic early warning markers greatly decreased non-emergency hospital admissions during the study period.

Table 8: Preventive Health Biosensor & Diagnostic Accuracy

Assessment Parameter	Value	Clinical Impact
Early Warning Window	Up to 72 hours	Identifies pre-clinical physiological irregularities prior to symptom onset
Diagnostic Accuracy	94.2% ($p < 0.001$)	Statistically significant reduction in non-emergency hospital admissions

4.4 Autonomous Robotic Logistics Latency and Operational Safety

The polyfunctional robots and UAVs achieved an average delay of 12.8 minutes in fulfilling the last-mile delivery within an 8 km service range. The delay in path planning decisions was maintained below 15 ms using the edge computing architecture and ensuring a 99.4% no-collision operation.

Table 9: Autonomous Robotic Logistics & Operational Safety

Operational Domain	Key Performance Indicator (KPI)	Performance Value
Last-Mile Delivery	Average Delay	12.8 min (within an 8 km service radius)
Edge Computing Latency	Path Planning Decision Time	< 15 ms
Safety & Collision Avoidance	Incident-Free Operation Rate	99.4% no-collision rate

4.5 Socio-Economic Integration and Public Infrastructure Convergence

The integration of agency software, quantum-resistant edge computing, and renewable energy grid technology resulted in an increase of operational efficiency by 3.8 times. The convergence of technologies across multiple domains made it clear that sovereign infrastructure can help to close the access divide between the urban metro centers and rural communities.

Table 10: Infrastructure Convergence & Socio-Economic Efficiency

Convergence Scope	Primary Operational KPI	Outcome / Impact
Multi-Domain Technology Convergence (<i>Agentic Software, Quantum-Resistant Edge, Renewable Grids</i>)	Operational Efficiency Multiplier	3.8x Increase
Socio-Economic Integration	Access Divide	Bridges infrastructure availability between urban centers and rural communities

Conclusion

The integration of agentic artificial intelligence, quantum-enabled networks, renewable energy systems, programmable biology, and autonomous robotics between 2026 and 2035 is a paradigm shift for contemporary society that transforms technologies from mere means to adaptive intelligences capable of optimizing human wellness, energy efficiency, productivity, and access.

References

1. Brier, P., & Marr, B. (2026). *Top tech trends of 2026: AI backbone, intelligent apps, and Cloud 3.0*. Capgemini Research Institute.
2. Chen, Y., & Zhang, L. (2025). Edge-deployed small language models and autonomous agentic workflows in enterprise systems. *Journal of Systems Architecture*, 158, 103210.
3. Deshmukh, P., & Sharma, R. (2024). Post-quantum cryptography and biometric

- tokenization in secure distributed networks. *IEEE Transactions on Information Forensics and Security*, 19, 1420–1435.
4. Gharib, M., & Al-Hawari, S. (2022). Artificial intelligence in predictive clean energy grid stabilization. *Expert Systems with Applications*, 201, 117105.
 5. Gupta, S., & Verma, M. (2025). Bio-digital crossovers: Programmable genomics and continuous physiological monitoring. *Nature Biotechnology*, 43(3), 312–326.
 6. Huang, J., & Patel, K. (2023). Spatial computing and digital twin applications in industrial automation. *Computers in Industry*, 145, 103820.
 7. Ibrahim, A., & Mercer, D. (2021). Polyfunctional robotics and low-latency path planning in urban environments. *Robotics and Computer-Integrated Manufacturing*, 71, 102150.
 8. Joshi, N., & Mehta, V. (2026). Agentic AI frameworks in multi-domain decision support systems. *International Journal of Artificial Intelligence*, 34(1), 45–62.
 9. Kapoor, A., & Singh, R. (2024). Real-time health risk prediction via wearable bio-wearable telemetry. *Journal of Healthcare Informatics*, 30(2), 112–128.
 10. Lee, H., & Kim, S. (2022). Long-term socioeconomic impacts of autonomous technology integration: A ten-year forecast model. *Technological Forecasting and Social Change*, 179, 121640.
 11. Nair, S., & Pillai, K. (2025). Decentralized renewable energy microgrids and algorithmic load balancing in developing regions. *Renewable & Sustainable Energy Reviews*, 189, 113890.
 12. Rajan, V., & Malhotra, P. (2023). Unmanned aerial vehicle flight control systems for low-altitude urban delivery networks. *IEEE Transactions on Intelligent Transportation Systems*, 24(8), 8901–8915.
 13. Smith, E., & Johnson, T. (2020). Multi-agent reinforcement learning in complex physical system optimization. *Decision Support Systems*, 138, 113390.
 14. Taylor, M., & Davis, C. (2021). Industry 4.0 and the convergence of biological and digital automation. *Technology in Society*, 67, 101730.
 15. Zhou, W., & Xu, Y. (2024). Hybrid cloud architectures and edge computing for sovereign digital infrastructures. *IEEE Internet of Things Journal*, 11(5), 7620–7632.

INTELLIGENT SYSTEMS FOR PROJECT MANAGEMENT AND EVALUATION

Salahuddin

Department of Mathematics,
Amet University, Kanathur, Chennai (Tamil Nadu), India
Corresponding author E-mail: vsludn@gmail.com

Abstract

This paper designs and defines an intelligent system that can perform project budgeting and cost accounting by using a Project Management and Evaluation (PME) cell. The intended system has been conceived as an intelligent system in the wider sense (as opposed to a tightly rule-constrained expert system): it has been defined as a hybrid system which combines symbolic-rule-based-reasoning with statistical and machine-learning modules thus implying judiciousness, pattern-recognition and adaptability is incorporated from the very beginning and not as an add-on. Integrating together planning resource allocation and monitoring within the institutional set up.

Keywords: Intelligent Systems, Artificial Intelligence, Machine Learning, Project Management, Cost Accounting, PME Cell.

1. Introduction

1.1 Background

The pressures for managing several projects concurrently are mounting, while fiscal responsibility and efficiency can't be compromised. However, the prevalent approach to project management is bogged down by disjointed information systems, late reporting and poor functional coordination. A Project Management and Evaluation (PME) cell would be an intelligent and responsive response for monitoring and financially guiding projects. Placing the cell on a basis of an intelligent system, rather than a static rule-based engine, would allow it to keep up with the dynamics and load of the current project portfolios.

1.2 Problem Statement

The initial problem to solve in this investigation is to devise an intelligent system capable of automating and optimizing the tasks performed in a PME cell. These include project documentation, financial evaluation, resource allocation, inter-company coordination, financial reconciliation and so on. Manually carrying out these tasks takes time and is prone to errors.

As with many operational tasks, a simple set of codified rules provides a fair degree of understandability and auditability, but fails to achieve the real time intelligence required for quick decision making due to ambiguity, incomplete data and fluctuating trends inherent in portfolios of projects.

1.3 Research Objectives

The objective of this paper is to: Build an appropriate structure of an intelligent system based on the functionality and characteristics of the PME cell Develop a domain knowledge of PME's

business process including: facts, rules, and learned cases Implement techniques required for project auto-budgeting and auto-costing Analyse the pros and cons.

2. Literature Review

2.1 Intelligent Systems in Project Management

Intelligent systems involve a wide variety of techniques that span the AI continuum: from basic rule-based logic and knowledge representation, to statistical learning methods, optimization algorithms, and, more recently, intelligent, adaptive agents-this combined work to reproduce and extend human skill sets. In the realm of project management, intelligent systems have been implemented for problems such as optimization, risk modelling, and the effective use of resources. Unlike classical expert systems, which drew on human knowledge in the form of programmed rules alone, these current systems combine the use of explicitly stated knowledge, along with pattern detection of previous results from project data, in order to provide useful recommendations based on an increasing availability of information over time. Combining intelligent systems and methodologies for project management not only offers a mechanism to acquire organizational knowledge and codify decision processes but also brings a level of consistency to execution and, surprisingly, new insights into patterns the rule writer might never imagine.

2.2 Cost Accounting and Budgeting Systems

Traditionally project cost systems and PM frameworks, like PMIS (Project Management Information System), did not always talk to each other, causing gaps, for example the inability to see trends, actual versus forecast costs over a project's lifecycle etc. It seems like the evolution of integrated cost control is trying to address this problem via real time tracking, monitoring variances, and more significantly using it as a predictor. Integrating the two will be a progression of towards 'organisational intelligence' wherein forecast values would continuously adjust against the actuals and alerts are triggered real-time.

2.3 Organizational Integration Challenges

Prior studies conclude coordination between functional departments as an essential condition for the success of a project-based firm. Solutions that manage co-ordination and integration through effective knowledge sharing, cross functional co- ordination, aligned business processes and co-ordinated work streams contribute to more favourable results. The PME cell conceptually proposes an organization for formal co-ordination supplemented by automation ranging from enforcement of rule-based behaviours to workflow driven processes to learnt Prioritization capabilities.

3. System Architecture and Design

3.1 Intelligent System Components

The proposed intelligent system for the PME cell comprises four fundamental layers, summarized in Table 1.

Table 1: Intelligent PME System Architecture — Four-Layer Design

Layer	Core Elements	Function
Knowledge & Learning Base	Facts, rules, heuristics, and trained models (project data, budgets, resources, policies, approval workflows, best practices, historical patterns)	Stores both codified expertise and statistically learned knowledge that the reasoning layer draws on
Reasoning Engine	Forward chaining (data-driven), backward chaining (goal-driven), rule interpreter, explanation facility, and supporting ML models (regression, classification, clustering)	Processes project data, applies business rules and learned models, generates recommendations, and explains its conclusions
User Interface	Role-based views for project leaders, PME staff, management, and functional departments	Supports data entry, query formulation, report generation, and alert notification
Database Management System	Project folders, financial data, historical transaction log	Maintains persistent, secure, and accessible storage for analysis and reporting

Knowledge & Learning Base: Store domain knowledge on project budgeting rules, cost accounting standards, resource allocation guidelines, corporate standards, and industry best practices, also in conjunction to the learned patterns or the statistical evidence resulted from past projects performance. Factual knowledge (project information, financial information, hierarchical organization structures), Heuristic knowledge (business rules for approval, exception and decision making) & Learnt knowledge (predictive model trained on projects historical data).

Reasoning Engine: Contains both forward and backward chaining, and uses machine-learning based inference to make use of the project inputs, business rules and derive the recommendations, which consists of conditional statement checking, trigger appropriate action or rules and provide justifications for the recommendations. It gives us transparency by its rule-based component and robustness through learning-based approach.

User Interface: Role-based security interface access that includes project leaders, PME cell personnel, management and functional units of the organization; data input screen, query formulation, reporting facility, and alert mechanisms. **Database Management System:** Responsible for efficiently providing permanent storage for all project-related folders, all financial transactions that occur during each project lifecycle, the progress reports being prepared periodically, and the previous history for both training and reporting purposes.

3.2 PME Cell Functional Framework

The PME cell is organized around seven core functions, coordinated through a central hub and overseen by a management layer responsible for strategic planning, decision-making, and policy formulation (Table 2).

Table 2: PME Cell Functional Framework

Function	Description
Project Documentation & Lifecycle	Folder creation, code assignment, lifecycle tracking
Cost Calculation & Estimation	Manpower, material, and capital cost estimation
Planning & Budget Preparation	Annual and multi-year plan generation
Resource Allocation Optimization	Constraint-based, priority-weighted allocation
Interdepartmental Coordination	Synchronization across Accounts, Stores, Purchase, Technical Services
Monitoring & Variance Analysis	Real-time progress and financial tracking
Reconciliation & Reporting	Monthly reconciliation and compliance reporting

3.3 Detailed Module Architecture

Each of the seven functions above is realized as a dedicated module, described in Table 3.

Table 3: Seven Functional Modules — Detailed Capabilities

Module	Key Capabilities
1. Project Documentation & Lifecycle Management	Project folder creation; code number assignment; version control; document archival; authorization tracking; committee recommendations
2. Cost Calculation & Estimation	Manpower and material cost calculators; capital equipment costing; contingency planning; inflation adjustment; what-if scenario analysis
3. Planning & Budget Preparation	Annual and five-year plan generation; budget consolidation; strategic alignment; traditional budget correlation; revision workflows
4. Resource Allocation Optimization	Constraint-based allocation; priority scoring; conflict resolution; stakeholder notification; capacity planning; resource leveling
5. Interdepartmental Coordination	Interfaces with Accounts, Stores, Purchase, Admin, and technical departments; data synchronization; workflow routing; conflict detection; escalation management
6. Monitoring & Variance Analysis	Real-time progress tracking; physical and financial status; variance detection; earned-value analysis; milestone alerts; performance dashboards
7. Reconciliation & Reporting	Monthly summary reports; account reconciliation; discrepancy resolution; compliance checking; audit trail maintenance; authority distribution

3.4 Information Flow

Inputs into the reasoning engine-project proposals, committee approvals, financial data, progress reports, resource availability, and departmental submissions-use rule processing, data analysis, information creation, resource optimization, and variance detection within the knowledge and learning base. Outputs can include managerial reports, budget allocations, progress notifications, resource allocations, reconciliation reports, and decision suggestions. A constant loop feeds outputs back into the learning base; this system does not "stay put" like a non-learning system and continue using static rules and models.

4. Reasoning Mechanisms and Decision Logic

4.1 Forward Chaining — Data-Driven Reasoning

Proactive Project Management Using forward chaining (data arriving and automatically being tested against rule conditions to initiate actions such as alert PM at >95% budget and <30 days left, reminder seven days before a milestone is due, automatic re-allocation for a found resource conflict)

Benefits:

- Proactive identification of issues
- Real-time visibility and notification
- Automatic triggering of tasks
- Lower manual intervention

4.2 Backward Chaining — Goal-Driven Reasoning

In order to answer queries, backward chaining can be used. The process: take a query (e.g. 'what projects are fundable?'), break it into sub-queries (find available budget, find costs, compare, rank, etc.), query the knowledge base for supporting facts, apply the constraints, return a ranked list of feasible projects, along with the chain of reasoning to get there.

Benefits are Precise query-answering, clear chain of reasoning, what-if analysis capabilities

4.3 Hybrid Reasoning: Rules Plus Learned Models

While most classic expert systems cease to draw inference once rule sequences are exhausted, our intelligent-systems view adds inference based on statistics and machine-learning to forward and backward chaining. The non-debatable policies-such as minimum standards, the authorization floor, and mandatory compliance verification-remain explicit in rules; estimated likely cost overruns or project completion delay derived from machine-learning contribute additional guidance but are not allowed to passively bypass rules. This has added value because it maintains auditability and allows more of the project history to inform the system's opinion over time.

4.4 Decision Logic for Project Approval

Project approval follows a tiered decision path based on budget size and review outcomes, summarized in Table 4.

Table 4: Project Approval Decision Logic (Multi-Level Review)

Budget Threshold	Review Path	Outcome
Under \$100K	PME Cell approval → Technical review	Approved or rejected at PME Cell level
\$100K – \$500K	Internal Committee review	Approved or rejected by Internal Committee
Above \$500K	Executive Committee review — strategic alignment and financial review	Approved or rejected by Executive Committee

5. Implementation Strategy

5.1 Phased Deployment Roadmap

Table 5: Five-Phase Implementation Roadmap

Phase	Timeframe	Focus
1. Foundation	Months 1–3	Database design; basic project folder management; authentication and access control; code number generation
2. Core Functionality	Months 4–6	Cost calculation modules; budget preparation tools; integration with organizational planning; basic reporting
3. Coordination & Integration	Months 7–9	Resource allocation algorithms; interdepartmental interfaces; workflow automation; notification systems
4. Intelligence & Optimization	Months 10–12	Advanced analytics and variance analysis; predictive modelling; optimization engines; machine-learning integration
5. Continuous Improvement	Ongoing	System refinement from feedback; knowledge and model base expansion; performance tuning; feature enhancement

Success at each phase is tracked through user adoption rate, system performance, data accuracy, and process efficiency gains.

5.2 System Integration Architecture

This intelligent PME solution can be integrated with other business systems such as: financial accounting, store management, purchase and procurement, HR and payroll, document management etc via an interface/integration layer providing interfaces such as REST APIs, SOAP web services, message queues, ETL, file transfer (csv/xml) and producing results such as consolidated reports, dashboards and analytics, notifications and alerts, audit trail and logs, executive summaries etc.

6. Knowledge and Learning Representation

6.1 Production Rules

Rules remain a core representation for policy that must be transparent and auditable. A representative rule:

Table 6: Sample Production Rule

Field	Value
Rule ID	R001
Priority	High
Condition	Budget utilized > 90% AND project status = "Active" AND days remaining < 30
Actions	1) Alert project manager — "Budget Critical" 2) Notify PME Cell — "Review Required" 3) Flag project — "High Risk" 4) Generate variance report
Confidence	0.95
Explanation	Project is 90%+ budget consumed with fewer than 30 days remaining

Other representative rules: resource distribution (when supply does not exceed demand use allocation logic based on priority), approval (if amount greater or equal than \$500K seek approval of Executive Committee), variance (if actual spending > planned budget with >15% difference raise notification and investigation), reconciliation (if PME & accounting data don't agree, flag discrepancy, raise and execute reconciliation).

6.2 Semantic Network

Relationships within the PME domain form a semantic network – an organization 'has' PME cells, projects and departments; the PME cell 'manages' projects and 'coordinates' departments. Projects 'require' budget, resources and timeline which 'are tracked by', 'consumed by', 'measured by' respectively, and cost accounting 'reconciles with' financial accounts and 'reports on'. Examples include the relationship 'has to define composition', 'manages' for operational control, 'requires' for dependency relationships and 'tracked by', 'consumed by' or 'measured by' for specific details, 'reconciles with' for validation.

6.3 Learned Representations

In addition to rules and semantic relationships, knowledge within the base includes statistical models, specifically linear regression models for cost estimation, classifiers for assessment of risk level, clustering for analysis of similar historical projects, and anomaly detection models for interesting financial trades. These trained representations are re-trained periodically as outcome data become available and the models recognize patterns in financial trades without needing specific rule instances for each lesson to be modelled in advance.

7. Benefits and Expected Outcomes

7.1 Impact Analysis

Table 7: Multi-Dimensional Benefits Impact Analysis

Dimension	Current State	Expected Improvement
Operational Efficiency	Manual data entry; slow processing; high error rate; delayed reporting	Automation; ~60% processing-time reduction; ~80% error-rate reduction; real-time reporting
Decision Quality	Information lag; inconsistent judgement	Real-time information; standardized, enhanced accuracy; predictive capacity enabled
Financial Control	Budget variance untracked; limited visibility; slow reconciliation	~50% variance reduction; complete cost visibility; ~70% faster reconciliation; continuous audit readiness
Strategic Alignment	Ad hoc portfolio decisions; reactive risk management	Enhanced portfolio optimization; optimized resource utilization; proactive risk management; comprehensive performance tracking

7.2 Return on Investment Projection

A five-year cost-benefit projection illustrates the expected return:

Table 8: Five-Year Cost and Benefit Breakdown

Category	Amount
Initial development	\$250K
Implementation	\$100K
Training	\$50K
Annual maintenance	\$30K / year
Total 5-year cost	\$520K
Labor savings	\$180K / year
Error reduction	\$80K / year
Improved decisions	\$120K / year
Total 5-year benefit	\$1,220K

$ROI = (Total\ Benefits - Total\ Costs) / Total\ Costs \times 100\% = (\$1,220K - \$520K) / \$520K \times 100\% \approx 134.6\%$ over five years, with a projected break-even point in Year 3.

8. Risk Analysis and Mitigation

Table 9: Risk Assessment Matrix

ID	Risk	Probability	Impact	Priority
R1	User resistance	High	High	Priority
R2	Data quality issues	High	Medium	High
R3	Integration complexity	Medium	High	High
R4	Budget overrun	Medium	High	High
R5	Incomplete requirements	High	Low	Medium
R6	Vendor dependency	Low	Medium	Low
R7	Security vulnerabilities	Medium	Medium	Medium
R8	Technology obsolescence	Low	Low	Low

Mitigation activities: R1-Change management process with executive support; R2-data cleansing process, validation rules, and checks; R3-Phased integration process; proper design of interface; use of middleware layer; R4-Properly documented detailed design; contingency reserve; ongoing active monitoring process.

9. Performance Metrics and Monitoring

System health and impact are tracked through a KPI dashboard spanning four categories, summarized in Table 10.

These process-efficiency benefits have ranged as high as an estimated 60% of process hours in report production (from 8h to 3h), 55% in budgeting (from 10h to 4.5h), 70% in account reconciliations (from 15h to 4.5h) and 50% in resource assignments (from 10h to 5h).

Table 10: KPI Dashboard Categories

Category	Representative KPIs and Targets
System Performance	Response time (target < 2s); system uptime (target > 99.5%); daily active user sessions
Operational Metrics	Active projects tracked (156); reports generated (2,340/month, automated); alerts sent (68/week, proactive)
Financial Accuracy	Reconciliation accuracy (target > 95%); variance detection rate (target > 90%); budget compliance (target > 90%)
User Satisfaction	User adoption rate (target > 80%); satisfaction score (target > 4.0/5); support ticket volume (trending down)

10. Future Enhancements and Research Directions

10.1 Technology Evolution

The system is intended to evolve along a continuum, not be presented whole. Incremental system improvements in the near term (Years 2-3) will further integrate machine learning in the process of cost estimation and pattern recognition in the assessment of risk, a natural language interface for user interaction, and improved business intelligence and visualization tools. In the mid-term (Years 4-5), improvements to the system are targeted towards integrating deep-learning to assess project success, automated updating of knowledge bases, cloud-native architecture to facilitate scalability, enhanced real-time collaborative functionalities, and audit trails generated and managed using blockchain technology. Future directions beyond the mid-term (Year 6 +) will likely involve highly sophisticated, self-optimizing agents for strategic planning and resources management, the utilization of quantum optimization for difficult allocation problems, augmentation of data visualization via augmented reality, distributed operations via edge computing, and decision support through cognition computing systems.

10.2 Deepening the Hybrid Reasoning Core

Intelligent system already integrates rule-based decision making with machine learning models – regression, classification, clustering, time-series and optimization – for cost estimation, risk analysis, resource allocation optimization and anomaly detection and an approach based on a closed-loop learning process needs to be implemented to improve the intelligent system. This would include gathering results data, re-training models with the gathered data, re-deploying updated models and evaluating the updated model performance to begin the next loop; data with history of projects can train regression models, aiming to improve estimation accuracy over time; features of a project-duration and the budget plus other correlated variables-can train a classifier to assign a risk score; resource and project constraints can form an objective function to be minimized or maximized in an optimization algorithm that may suggest resource allocation; transaction behaviour patterns train an anomaly detection model, which flags unusual transaction patterns to be investigated.

Conclusion

The design of an intelligent system for PME, this research has provided an architecture that cover the knowledge & learning representation, hybrid reasoning mechanism and the system implementation including a plan for the whole organization.

The system that described above fulfil 7 primary functions from the PME cell scope: the documentation, calculation, planning, resource planning, interdepartmental communication, monitoring and reconciliation. By combining the logic reasoning with a learning system driven by data that the system has acquired it could make PME cell change from administrative part of the organization to be a competitive strategic part which improve quality of decisions and increase financial efficiency.

The research makes the following key contribution:

- The architecture: The system is composed of 4 levels, including the knowledge and learning base, hybrid reasoning engine, user interface and database.
- Functions: The research shows the system can carry out 7 tasks that covered the responsibilities in PME cell. Each function should contain the domain specific rule, the learning model acquired, and automated process.
- Reasoning mechanism: Combine forward and backward chaining with statistics and learning reasoning method, this approach supports active monitoring, goal driven query process and improving the prediction quality dynamically.
- Implementation strategy: Based on the assessment of risk management the implementation process should be in phases to benefit the organization as much as possible and maintain a controllable approach when risks emerge.
- Benefit: The projected ROI is 134.6% in 5 years by saving labour, and decreasing error as well as improving decision efficiency.

The challenge of implementing the system is that careful attentions should be paid to knowledge engineering, governing the learning model, integrating with system and organizational management. Problems that are associated with technical issues such as legacy system integration and data accuracy should be taken account of, and problems that associated with organizational such as user resistance and cultural change should be solved as well.

Further research may focus on the machine learning integration, especially for predicting the risks, optimize resource allocation, and integrate with collaborative intelligent system for building up a system between human beings and machines. With this development, these types of system should be capable of automate not only the existing procedure but also the creative decision based on the acquired new knowledge.

The world seems to drive to make business operate efficiently and strictly by financial rules; the intelligent system such as described above will be the key success factor of organizations. Integration of machine and intelligence with project management seems promising enough for exploring new directions and creating higher efficiency.

This research provides a guideline of framework and implementation of intelligent system on PME for practitioners and researchers. A fully learning and self-optimized intelligent system that continuously learn new knowledge from organization and adapt to new business environments so as to improve itself and provide the best support to organization, this is the final aim for intelligent system on PME.

References

1. Bahill, A. T., & Henderson, S. J. (2005). Requirements development, verification, and validation exhibited in famous failures. *Systems Engineering*, 8(1), 1–14.
2. Draganidis, F., & Mentzas, G. (2006). Competency based management: A review of systems and approaches. *Information Management & Computer Security*, 14(1), 51–64.
3. Giarratano, J. C., & Riley, G. (1998). *Expert systems: Principles and programming*. PWS Publishing Company.
4. Kerzner, H. (2017). *Project management: A systems approach to planning, scheduling, and controlling*. John Wiley & Sons.
5. Liao, S. H. (2005). Expert system methodologies and applications—A decade review from 1995 to 2004. *Expert Systems with Applications*, 28(1), 93–103.
6. Russell, S. J., & Norvig, P. (2020). *Artificial intelligence: A modern approach* (4th ed.). Pearson.
7. Turban, E., Aronson, J. B., & Liang, T. P. (2005). *Decision support systems and intelligent systems*. Pearson Prentice Hall.
8. Van Der Aalst, W. M., & Van Hee, K. M. (2004). *Workflow management: Models, methods, and systems*. MIT Press.
9. Waterman, D. A. (1986). *A guide to expert systems*. Addison-Wesley.
10. Winston, P. H. (1992). *Artificial intelligence* (3rd ed.). Addison-Wesley.

MACHINE LEARNING APPLICATIONS IN ENVIRONMENTAL AND ECOLOGICAL SCIENCES

Ranjana

Department of Zoology,

Patna Science College, Patna University, Patna, Bihar, 800005, India.

Corresponding author E-mail: ranjana.prakash81@gmail.com, ranjanazoolpsc@pup.ac.in

ORCID: <https://orcid.org/0000-0003-0870-9535>

Abstract

Machine learning (ML) has moved from a peripheral computational curiosity to a central methodology in environmental and ecological research. Growing volumes of satellite imagery, in-situ sensor networks, camera-trap archives, bioacoustic recordings, and citizen-science observations have outpaced the capacity of classical statistical and mechanistic models, creating fertile ground for data-driven approaches. This chapter surveys the principal families of ML methods used in environmental and ecological science — including random forests, support vector machines, gradient boosting, convolutional neural networks, recurrent architectures, and ensemble/hybrid systems — and reviews their application across eight major domains: species distribution modelling and biodiversity assessment, wildlife monitoring through camera traps and bioacoustics, remote sensing and land cover mapping, hydrology and water resources, water quality prediction, atmospheric science and air quality forecasting, climate modelling, and forest disturbance detection, with a further look at precision agriculture. Representative case studies, comparative tables, and conceptual diagrams illustrate how these methods are implemented and evaluated. The chapter closes with a critical discussion of persistent challenges — data scarcity and class imbalance, domain shift, interpretability, and the ethics of automated decision-making in conservation — and outlines emerging directions, including foundation models for Earth observation, physics-informed learning, and edge-deployed monitoring systems.

Keywords: Machine Learning, Deep Learning, Environmental Science, Ecology, Remote Sensing, Species Distribution Modelling, Biodiversity Monitoring, Camera Traps, Bioacoustics, Hydrology, Water Quality, Air Quality, Climate Modelling, Precision Agriculture, Conservation Technology.

1. Introduction

Environmental and ecological systems are characterised by high dimensionality, non-linearity, spatial and temporal heterogeneity, and frequently incomplete observation — properties that make them notoriously difficult to model with classical linear statistics or first-principles physical equations alone. Over the past decade, machine learning has emerged as a complementary and, in many tasks, superior tool for extracting structure from environmental data [1]. The deep learning revolution, catalysed by advances in computer vision and natural language processing, has since diffused into the geosciences, ecology, hydrology, and

atmospheric sciences, driven by three converging trends: (i) an explosion of Earth-observation and sensor data, including daily-revisit multispectral satellite constellations, low-cost water- and air-quality sensors, and autonomous acoustic and camera-trap recorders; (ii) the maturation of open-source deep learning frameworks and affordable graphics-processing hardware; and (iii) growing recognition among ecologists and environmental scientists that classical statistical approaches struggle to scale to the volume and complexity of modern monitoring data [9,10].

This chapter provides a structured overview of how machine learning is applied across environmental and ecological science. Section 2 introduces the principal families of ML algorithms used in this domain. Section 3 reviews eight application areas in turn, drawing on recent peer-reviewed literature and illustrating each with representative studies. Section 4 synthesises common challenges, and Section 5 considers future directions. Throughout, the aim is not to catalogue every published study, but to give the reader a working conceptual map of where ML has proven useful, which methods are typically chosen for which problem, and what practical and epistemic limitations remain.




Figure 1: A generic machine learning workflow for environmental and ecological applications. Real projects rarely proceed strictly linearly — validation frequently sends practitioners back to feature engineering or even data acquisition

2. Machine Learning Foundations for Environmental Science

Machine learning methods used in environmental and ecological research broadly fall into four families: supervised learning, unsupervised learning, deep learning, and ensemble/hybrid methods [2,3]. Supervised methods — including random forests [2,3], support vector machines, and gradient-boosted trees — learn a mapping from labelled predictor variables (e.g., climate layers, spectral bands, water chemistry) to a known response (e.g., species presence, land cover class, pollutant concentration). Unsupervised methods, such as k-means clustering, principal component analysis, and self-organising maps, are used to discover latent structure — functional groups of species, spectral clusters in imagery, or regimes in a time series — without labelled outcomes. Deep learning methods, particularly convolutional neural networks (CNNs) and recurrent architectures such as long short-term memory (LSTM) networks, learn hierarchical feature representations directly from raw imagery, spectrograms, or time series, reducing the need for hand-engineered features and often achieving state-of-the-art performance on perceptual tasks such as image classification and sequence prediction. Ensemble and hybrid approaches combine multiple base learners, or couple ML components with physically based models, to improve robustness and, increasingly, physical plausibility.




Figure 2: Taxonomy of machine learning methods commonly applied in environmental and ecological research

Random forests deserve particular mention because of their historical role in popularising ML within ecology. Cutler and colleagues [3] demonstrated that random forests outperformed classical classifiers such as linear discriminant analysis and logistic regression on invasive-species and nest-site classification tasks, while also providing interpretable variable-importance rankings — a combination of accuracy and interpretability that made the method rapidly popular among ecologists unfamiliar with black-box neural networks. Table 1 summarises the algorithms most frequently encountered in the environmental ML literature, together with their typical use cases and notable strengths and limitations.

Table 1: Common machine learning algorithms and their typical applications in environmental and ecological science

Algorithm	Typical Use in Environmental / Ecological Science	Strengths	Limitations
Random Forest	Species distribution modelling, land cover classification, water quality classification	Robust to noise; handles non-linear interactions; built-in variable importance	Can overfit rare classes; less suited to raw imagery or sequential data
Support Vector Machine	Land cover classification, small-sample species classification	Effective in high-dimensional, small-sample settings	Sensitive to kernel/parameter choice; slow on very large datasets
Gradient Boosting (XGBoost)	Water and air quality prediction, habitat suitability modelling	High predictive accuracy; handles heterogeneous data	Many hyper-parameters; prone to overfitting without careful tuning
Convolutional Neural Network	Camera-trap image classification, satellite scene classification, coral/vegetation mapping	Learns spatial features automatically; state-of-the-art on imagery	Requires large labelled datasets; computationally intensive

Recurrent / LSTM Network	Streamflow forecasting, air pollutant time series, phenology prediction	Captures long-range temporal dependencies	Training instability; limited physical interpretability
k-Means Clustering	Habitat typology, spectral clustering, anomaly detection	No labelled data required; exploratory	Requires a priori choice of cluster number; sensitive to scaling
Ensemble Stacking	Multi-model species distribution ensembles, hybrid hydrological forecasting	Improves robustness and accuracy over single learners	Reduced interpretability; higher computational cost

3. Application Domains

3.1 Species Distribution Modelling and Biodiversity Assessment

Species distribution modelling (SDM) relates known occurrence records to environmental covariates (climate, topography, land cover) to predict the potential geographic range of a species. Beery and colleagues [5] provide a comprehensive review of SDM aimed at machine learning practitioners, noting that the field has historically relied on relatively simple statistical techniques such as MaxEnt and generalised linear models, but has increasingly adopted random forests, boosted trees, and — where sufficient occurrence data exist — deep neural networks. A recurring theme in this literature is that observation data are rarely a random sample of the landscape: presence-only records from citizen-science platforms are biased toward accessible, populated areas, and naive application of ML classifiers to such data can produce misleadingly confident range maps unless sampling bias is explicitly modelled [5]. Ensemble approaches that ensemble multiple algorithms (random forest, support vector machine, boosted trees, and neural networks) tend to outperform any single method and are now considered best practice for operational SDM.

3.2 Wildlife Monitoring: Camera Traps and Bioacoustics

Camera traps and passive acoustic recorders have transformed field ecology by enabling continuous, non-invasive monitoring at a scale unattainable by human observers alone, but they generate volumes of imagery and audio that are impractical to review manually. Deep convolutional networks have become the standard tool for automating this task. In an early landmark study, Norouzzadeh and colleagues [8] trained CNNs on the 3.2-million-image Snapshot Serengeti dataset to identify, count, and describe the behaviour of 48 species, achieving greater than 93.8% classification accuracy and demonstrating that the equivalent of many years of manual labelling effort could be replaced by automated inference. Subsequent work has extended this paradigm to object detection (localising individual animals within an image rather than merely classifying the whole scene), transfer learning across geographic regions, and the use of auxiliary sensing modalities such as depth estimation. Tuia and co-authors [9] synthesise this literature in a widely cited perspective, arguing that the primary bottleneck in ML for wildlife conservation is no longer algorithmic accuracy but the translation of model outputs into

decision-relevant, interpretable, and field-deployable tools, and the need for closer collaboration between the ML and conservation communities from the outset of project design.

In parallel, bioacoustic monitoring has matured rapidly. BirdNET [12], a deep convolutional network trained to identify birds from audio recordings across nearly a thousand North American and European species, illustrates how a single, freely available model can be embedded into both research pipelines and citizen-science mobile applications [13], dramatically lowering the barrier to large-scale acoustic biodiversity monitoring. More recent work on transfer learning with pretrained bioacoustic embeddings [14] demonstrates that models trained primarily on common, data-rich species can be efficiently adapted to rare taxa, bats, amphibians, and marine mammals with comparatively little additional labelled data — an important development given the chronic scarcity of annotated recordings for threatened species.

3.3 Remote Sensing and Land Cover / Land Use Classification

Satellite and airborne remote sensing has long relied on machine learning for land cover and land use (LULC) classification, historically using classifiers such as maximum likelihood, decision trees, and random forests applied to spectral bands and vegetation indices. Zhao and colleagues [6] review the more recent shift toward deep learning, noting that convolutional architectures — and, increasingly, vision transformers — now dominate the LULC literature because they can exploit spatial context and texture rather than relying purely on per-pixel spectral signatures. Zhu and colleagues [7] provide an influential earlier survey documenting this transition and cataloguing the growing library of benchmark datasets (e.g., EuroSAT, UC Merced) that have accelerated methodological progress. Multi-sensor fusion — combining optical imagery (e.g., Sentinel-2), synthetic aperture radar (e.g., Sentinel-1), and digital elevation models — has become a common strategy for improving robustness to cloud cover and seasonal variation, with recent transformer-based architectures achieving overall accuracies approaching 87% on composite urban-agricultural-forest landscapes.

3.4 Hydrology and Water Resources

Hydrological forecasting — particularly streamflow and flood prediction — has become one of the most mature application areas for deep learning in the environmental sciences. Kratzert and colleagues [15] were among the first to demonstrate that a single long short-term memory (LSTM) network, trained on meteorological forcing data, could outperform the well-established conceptual SAC-SMA model for daily rainfall-runoff prediction. A follow-up study across 531 catchments in the continental United States [16] showed that a single LSTM model trained across many basins simultaneously could generalise to ungauged catchments better than calibrated regional hydrological models — a striking result given that traditional hydrology has generally assumed that catchment-specific calibration is necessary. Sit and colleagues [17] provide a broad review of deep learning applications across the water sector, spanning streamflow forecasting, water quality estimation, groundwater modelling, and infrastructure monitoring, and highlight recurring architectural choices (LSTM, CNN-LSTM hybrids, and

increasingly attention-based transformers) as well as recurring challenges around interpretability and the physical plausibility of learned relationships.

3.5 Water Quality Prediction

Water quality monitoring — encompassing parameters such as dissolved oxygen, turbidity, nitrogen and phosphorus loading, and overall potability — has increasingly adopted ML as a lower-cost alternative or complement to laboratory analysis. Helaly and colleagues [18] review the comparative performance of classical ML models (random forest, support vector machine) against deep learning approaches (CNN, LSTM, bidirectional LSTM) across multiple published water quality datasets, concluding that tree-based ensembles remain competitive for static, tabular water-chemistry data, while recurrent architectures tend to outperform them when the prediction task is explicitly temporal (e.g., forecasting future dissolved-oxygen concentration from a time series of prior readings). A recurring practical finding across this literature is that feature selection — identifying the small subset of physicochemical variables that are both informative and inexpensive to measure in the field — is often as important to real-world deployment as the choice of learning algorithm itself.

3.6 Atmospheric Science and Air Quality Forecasting

Fine particulate matter (PM_{2.5}) forecasting is among the most extensively studied environmental ML problems because of its direct relevance to public health. Morapedi and Obagbuwa [19] compared gradient boosting, random forest, and classical classifiers for predicting PM_{2.5} concentrations across South African cities, including in locations lacking dedicated monitoring stations, illustrating a common use case in which ML is used to interpolate or extrapolate from sparse ground-truth sensor networks. More broadly, the air quality literature shows a consistent pattern: ensemble tree methods (random forest, gradient boosting) perform well on tabular meteorological and emissions data, while LSTM and hybrid CNN-LSTM architectures tend to achieve superior performance when high-resolution temporal dynamics dominate the prediction task, echoing the pattern observed in hydrology.

3.7 Climate Modelling and Earth System Science

At the largest spatial and temporal scales, machine learning is increasingly integrated with — rather than substituted for — physically based climate and Earth system models. Reichstein and colleagues [20] argue for a hybrid paradigm in which deep learning is used to parameterise sub-grid-scale processes (e.g., cloud formation, turbulent mixing) that are computationally too expensive to resolve explicitly, while physical conservation laws constrain the learned components to remain consistent with known dynamics — so-called physics-informed or physics-guided machine learning. Rolnick and colleagues [21] survey a much broader landscape of climate-relevant ML applications, spanning energy-system optimisation, emissions monitoring, climate model emulation, and extreme-event detection, and emphasise that the greatest impact of ML on climate change mitigation and adaptation is likely to come not from any single flagship model but from many domain-specific tools embedded within existing scientific and operational workflows.

3.8 Deforestation and Ecosystem Disturbance Detection

Detecting deforestation, wildfire, and other abrupt ecosystem disturbances from satellite time series is a canonical change-detection problem well suited to deep learning. Fodor and Conde [22] introduce a multimodal deep learning framework that fuses optical (Sentinel-2, Landsat), synthetic aperture radar (Sentinel-1), and thermal anomaly (VIIRS, MODIS) data to jointly estimate deforestation extent and burned area across the Amazon basin, addressing the persistent problem of cloud cover during the rainy season by exploiting radar's all-weather sensitivity. This multimodal-fusion strategy — combining sensors with complementary strengths and weaknesses rather than relying on a single data source — is a recurring design pattern across the disturbance-detection literature and reflects the broader trend toward operational, near-real-time monitoring systems capable of supporting rapid enforcement and conservation response.

3.9 Precision Agriculture and Land Management

Although agriculture is sometimes treated as a distinct discipline from ecology and environmental science, precision agriculture shares its core methodological toolkit — remote sensing, sensor networks, and predictive ML — and has direct environmental consequences through its influence on land use, water consumption, and agrochemical application. Kamilaris and Prenafeta-Boldú [23] survey the rapid adoption of deep learning across crop classification, yield prediction, weed and disease detection, and soil property estimation, noting consistent performance gains over classical ML once sufficient labelled agricultural imagery became available. Weiss and colleagues [24] complement this with a meta-review of remote sensing specifically for agricultural monitoring, underscoring the practical value of freely available medium-resolution satellite data (e.g., Sentinel-2) for smallholder and large-scale farming systems alike, and highlighting persistent gaps in ground-truth data availability across the Global South.




Figure 3: Schematic of a convolutional neural network used to classify remote sensing scenes, camera-trap images, or acoustic spectrograms. The same general architecture — convolution, pooling, and dense classification layers — recurs across most of the imaging and audio applications described in this section

Table 2: Summary of application domains, characteristic data sources, and commonly used ML methods.

Domain	Typical Data Sources	Common ML Methods	Representative Study
Biodiversity / SDM	Occurrence records, bioclimatic layers, land cover	Random forest, boosted trees, ensembles	Beery <i>et al.</i> [5]
Wildlife monitoring	Camera-trap images, acoustic recordings	CNN, transfer learning, object detection	Norouzzadeh <i>et al.</i> [8]; Kahl <i>et al.</i> [12]
Remote sensing / LULC	Multispectral & SAR satellite imagery	CNN, vision transformers, random forest	Zhao <i>et al.</i> [6]
Hydrology	Streamflow gauges, meteorological forcing	LSTM, CNN-LSTM hybrids	Kratzert <i>et al.</i> [15,16]
Water quality	In-situ sensors, laboratory chemistry	Random forest, XGBoost, LSTM	Helaly <i>et al.</i> [18]
Air quality	Ground monitors, low-cost sensors, meteorology	Gradient boosting, LSTM	Morapedi & Obagbuwa [19]
Climate / Earth system	Reanalysis data, climate model output	Physics-informed deep learning	Reichstein <i>et al.</i> [20]
Forest disturbance	Multi-sensor satellite time series	CNN, multimodal fusion	Fodor & Conde [22]
Precision agriculture	UAV & satellite imagery, soil sensors	CNN, random forest	Kamilaris & Prenafeta-Boldú [23]




Figure 4: Illustrative ranges of reported model performance across ML application domains, compiled from the studies cited in this chapter. Because studies use different metrics (classification accuracy, AUC, Nash-Sutcliffe efficiency) and different evaluation datasets, these ranges are indicative of typical reported performance rather than directly comparable benchmarks



Figure 5: Principal application domains of machine learning in environmental and ecological science, as reviewed in Section 3 of this chapter

4. Challenges and Limitations

Despite substantial progress, several recurring challenges constrain the reliability and adoption of ML in environmental and ecological practice.

4.1 Data Scarcity, Class Imbalance, and Sampling Bias

Ecological datasets are frequently small, sparse, and biased. Rare or cryptic species are, by definition, underrepresented in training data, and citizen-science observation platforms tend to oversample accessible, populated locations [5]. Camera-trap and bioacoustic datasets are typically dominated by a small number of common species with a long tail of rare classes, producing severe class imbalance that biases naively trained classifiers toward the majority classes [8,9]. Transfer learning and self-supervised pretraining — learning generic feature representations from large amounts of unlabelled or weakly labelled data before fine-tuning on a small labelled target set — has emerged as the most effective mitigation strategy [14].

4.2 Domain Shift and Generalisation

Models trained in one geographic region, season, or sensor configuration frequently perform poorly when deployed elsewhere. Camera-trap classifiers trained on one national park's background vegetation and lighting conditions can generalise poorly to a different park, and land cover classifiers calibrated on one satellite sensor may require retraining for another. This domain-shift problem motivates growing interest in domain adaptation techniques and in training on deliberately diverse, multi-site datasets.

4.3 Interpretability and Trust

Ecologists and environmental managers often require not only accurate predictions but also an understanding of which variables drive a model's output, both to build scientific trust and to satisfy regulatory and conservation decision-making requirements. Random forests and gradient-

boosted trees offer relatively transparent variable-importance measures [3], whereas deep neural networks are comparatively opaque. Explainable-AI techniques such as SHAP (SHapley Additive exPlanations) and partial dependence plots are increasingly used to partially open this black box, though their application to spatially and temporally structured environmental data remains an active area of methodological development.

4.4 Physical Plausibility

Purely data-driven models can produce predictions that are statistically accurate on held-out test data yet physically implausible — for example, violating mass or energy conservation, or attributing streamflow variation to a meteorological driver through a mechanism inconsistent with known hydrological processes. This has motivated the development of physics-informed and hybrid ML approaches that constrain learned relationships to remain consistent with governing physical laws [20].

Table 3: Key challenges limiting the application of machine learning in environmental and ecological science, and common mitigation strategies

Challenge	Typical Manifestation	Common Mitigation Strategies
Data scarcity / class imbalance	Rare species under-represented in training data; long-tailed class distributions	Transfer learning, data augmentation, weighted loss functions, active learning
Sampling bias	Citizen-science records concentrated near roads and population centres	Bias-correction layers, target-group background sampling, spatial thinning
Domain shift	Poor generalisation across regions, seasons, or sensors	Domain adaptation, multi-site training data, sensor harmonisation
Interpretability	Black-box deep learning predictions distrusted by domain experts	SHAP, partial dependence plots, attention visualisation, interpretable surrogate models
Physical plausibility	Statistically accurate but mechanistically inconsistent predictions	Physics-informed loss terms, hybrid physical-ML architectures

5. Future Directions

Several key trends will shape machine learning’s next phase in environmental and ecological science. First, large pretrained foundation models for Earth observation and bioacoustics will substantially reduce labelled-data bottlenecks for rare species and under-monitored regions [14]. Second, physics-informed and hybrid architectures will become standard in hydrology and climate science, merging deep learning’s flexibility with physical constraints [20,21]. Third, edge deployment on low-power sensors will enable real-time, on-site inference for anti-poaching, wildfire detection, and water-contamination alerts. Finally, sustained interdisciplinary collaboration between ecologists and ML researchers remains the primary driver for achieving meaningful conservation and environmental impact [9].

Conclusion

Machine learning has become an indispensable component of the environmental and ecological scientist's toolkit, extending — rather than replacing — traditional statistical and mechanistic modelling. Its greatest successes to date lie in tasks characterised by large volumes of complex, high-dimensional data: image and audio classification for biodiversity monitoring, spatial pattern recognition in remote sensing, and sequence modelling for hydrological and atmospheric forecasting. Yet the field's continued maturation depends less on incremental algorithmic gains than on addressing the practical realities of ecological data — scarcity, bias, and heterogeneity — and on building models that are interpretable, physically consistent, and genuinely useful to the practitioners who must act on their outputs. As sensing infrastructure continues to expand and foundation models lower the barrier to entry, the central challenge for the coming decade is less whether machine learning can model environmental and ecological systems accurately, and more whether it can do so in ways that are trustworthy, equitable, and actionable.

References

1. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
2. Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
3. Cutler, D. R., Edwards, T. C., Jr., Beard, K. H., Cutler, A., Hess, K. T., Gibson, J., et al. (2007). Random forests for classification in ecology. *Ecology*, 88(11), 2783–2792.
4. Guisan, A., & Zimmermann, N. E. (2000). Predictive habitat distribution models in ecology. *Ecological Modelling*, 135(2–3), 147–186.
5. Beery, S., Cole, E., Parker, J., Perona, P., & Winner, K. (2021). Species distribution modeling for machine learning practitioners: A review. In *Proceedings of the 4th ACM SIGCAS Conference on Computing and Sustainable Societies (COMPASS '21)* (pp. 1–27). ACM.
6. Zhao, S., Tu, K., Ye, S., Tang, H., Hu, Y., & Xie, C. (2023). Land use and land cover classification meets deep learning: A review. *Sensors*, 23(21), 8966.
7. Zhu, X. X., Tuia, D., Mou, L., Xia, G. S., Zhang, L., Xu, F., et al. (2017). Deep learning in remote sensing: A comprehensive review and list of resources. *IEEE Geoscience and Remote Sensing Magazine*, 5(4), 8–36.
8. Norouzzadeh, M. S., Nguyen, A., Kosmala, M., Swanson, A., Palmer, M. S., Packer, C., et al. (2018). Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning. *Proceedings of the National Academy of Sciences of the United States of America*, 115(25), E5716–E5725.
9. Tuia, D., Kellenberger, B., Beery, S., Costelloe, B. R., Zuffi, S., Risse, B., et al. (2022). Perspectives in machine learning for wildlife conservation. *Nature Communications*, 13, 792.
10. Christin, S., Hervet, É., & Lecomte, N. (2019). Applications for deep learning in ecology. *Methods in Ecology and Evolution*, 10(10), 1632–1644.

11. Weinstein, B. G. (2018). A computer vision for animal ecology. *Journal of Animal Ecology*, 87(3), 533–545.
12. Kahl, S., Wood, C. M., Eibl, M., & Klinck, H. (2021). BirdNET: A deep learning solution for avian diversity monitoring. *Ecological Informatics*, 61, 101236.
13. Wood, C. M., Kahl, S., Rahaman, A., & Klinck, H. (2022). The machine learning-powered BirdNET app reduces barriers to global bird research by enabling citizen science participation. *PLOS Biology*, 20(6), e3001670.
14. Ghani, B., Denton, T., Kahl, S., & Klinck, H. (2023). Global birdsong embeddings enable superior transfer learning for bioacoustic classification. *Scientific Reports*, 13, 22876.
15. Kratzert, F., Klotz, D., Brenner, C., Schulz, K., & Herrnegger, M. (2018). Rainfall-runoff modelling using long short-term memory (LSTM) networks. *Hydrology and Earth System Sciences*, 22(11), 6005–6022.
16. Kratzert, F., Klotz, D., Shalev, G., Klambauer, G., Hochreiter, S., & Nearing, G. (2019). Towards learning universal, regional, and local hydrological behaviors via machine learning applied to large-sample datasets. *Hydrology and Earth System Sciences*, 23(12), 5089–5110.
17. Sit, M., Demiray, B. Z., Xiang, Z., Ewing, G. J., Sermet, Y., & Demir, I. (2020). A comprehensive review of deep learning applications in hydrology and water resources. *Water Science and Technology*, 82(12), 2635–2670.
18. Helaly, M. A., Aref, M. M., & Cotos, J. M. (2025). Advancements in water quality prediction: A practical review of machine learning and deep learning approaches. *Cluster Computing*.
19. Morapedi, T. D., & Obagbuwa, I. C. (2023). Air pollution particulate matter (PM2.5) prediction in South African cities using machine learning techniques. *Frontiers in Artificial Intelligence*, 6, 1230087.
20. Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N., et al. (2019). Deep learning and process understanding for data-driven Earth system science. *Nature*, 566(7743), 195–204.
21. Rolnick, D., Donti, P. L., Kaack, L. H., Kochanski, K., Lacoste, A., Sankaran, K., et al. (2023). Tackling climate change with machine learning. *ACM Computing Surveys*, 55(2), 42.
22. Fodor, G., & Conde, M. V. (2023). Rapid deforestation and burned area detection using deep multimodal learning on satellite imagery. *arXiv*. <https://arxiv.org/abs/2307.04916>
23. Kamilaris, A., & Prenafeta-Boldú, F. X. (2018). Deep learning in agriculture: A survey. *Computers and Electronics in Agriculture*, 147, 70–90.
24. Weiss, M., Jacob, F., & Duveiller, G. (2020). Remote sensing for agricultural applications: A meta-review. *Remote Sensing of Environment*, 236, 111402.

NANOTECHNOLOGY AS A CATALYST FOR INTELLIGENT SYSTEMS AND SUSTAINABLE ENGINEERING: MATERIALS, APPLICATIONS AND FUTURE PERSPECTIVES

Minhaaj Afreen S. Y and K. Gayathri*

Department of Physics, AMET University, Kanathur, Chennai, Tamil Nadu, India

*Corresponding author E-mail: gayathri.k@ametuniv.ac.in

Abstract

Nanotechnology has emerged as one of the most transformative enabling sciences of the twenty-first century, offering unprecedented control over matter at the scale of atoms and molecules. This chapter examines the convergence of nanotechnology with intelligent systems and sustainable engineering, three pillars that together define the trajectory of modern technological progress. The discussion begins with the fundamental principles of nanoscale science and engineering before moving to a detailed treatment of nanomaterials that enable smart sensing, self-healing structures, and energy-efficient computing. The chapter then explores how nanotechnology contributes to sustainable design through applications in renewable energy harvesting, water purification, green construction materials, and reduction of industrial waste. Emerging and cross-disciplinary applications in healthcare diagnostics, agriculture, and environmental remediation are also discussed, along with the technical, economic, and regulatory challenges that continue to limit large-scale adoption, including nanotoxicity, scalability of fabrication processes, and the absence of harmonised international standards. The chapter concludes by outlining future research directions, emphasising the role of artificial intelligence in accelerating nanomaterial discovery and the importance of responsible innovation frameworks in ensuring that nanotechnology contributes positively to sustainable engineering practice.

Keywords: Nanotechnology, Nanomaterials, Intelligent Systems, Sustainable Design, Smart Sensors, Green Nanotechnology, Emerging Technologies.

1. Introduction

The idea that engineered structures could be built and manipulated at the scale of individual atoms was first articulated in a now well-known lecture that speculated on the possibility of manoeuvring matter atom by atom (Feynman, 1960). Over the following decades this speculative vision matured into a distinct scientific and engineering discipline concerned with materials and devices whose critical dimensions lie between one and one hundred nanometres. At this scale, materials frequently display physical, chemical, optical, and electronic properties that differ markedly from their bulk counterparts, a consequence of increased surface-to-volume ratio and the growing influence of quantum-mechanical effects (Whitesides, 2003).

In parallel, the twenty-first century has witnessed two additional technological movements that are reshaping engineering practice: the rise of intelligent systems built around sensing, computation, and autonomous decision-making, and a growing global commitment to sustainable

design driven by concerns over climate change, resource depletion, and environmental degradation. Nanotechnology occupies a unique position at the intersection of these movements. On one hand, nanoscale sensors, actuators, and electronic components form the physical substrate upon which many intelligent systems are built. On the other hand, nanomaterials offer routes to more efficient energy conversion, lighter and stronger structural materials, and cleaner manufacturing processes that collectively advance the goals of sustainable engineering (Roco, 2011).

The rapid expansion of nanotechnology research over the past two decades has been accompanied by a corresponding growth in patent activity, dedicated research funding programmes, and specialised academic curricula across the world. National initiatives such as the United States National Nanotechnology Initiative, the European Union's Horizon research programmes, and similar mission-driven efforts in India, China, Japan, and South Korea have collectively directed billions of dollars toward nanoscale research infrastructure (Roco, 2011). This sustained investment reflects a broad consensus among policymakers and engineers that nanoscale engineering will play a decisive role in addressing some of the most pressing challenges of contemporary technological development, including energy security, resource efficiency, and the demand for increasingly capable yet increasingly compact computing and sensing hardware.

At the same time, the field has matured beyond a purely materials-science orientation into a genuinely interdisciplinary engineering practice. Mechanical, electrical, chemical, civil, and biomedical engineers now routinely draw upon nanoscale design principles when developing new products and systems. This interdisciplinary character is precisely what makes nanotechnology such a natural bridge between the two broader themes explored throughout this book, namely intelligent systems and sustainable design. Nanoscale components frequently serve simultaneously as the physical enablers of system intelligence, through sensing and low-power computation, and as contributors to sustainability, through improved material and energy efficiency.

This chapter provides an integrated overview of how nanotechnology supports the twin objectives of intelligent systems design and sustainable engineering. Section 2 introduces the scientific foundations of nanomaterials relevant to intelligent systems. Section 3 discusses nanotechnology-enabled sustainable design strategies. Section 4 presents illustrative case studies drawn from published engineering practice. Section 5 surveys emerging cross-disciplinary applications. Section 6 discusses the principal challenges facing the field, Section 7 outlines future research directions, and Section 8 concludes the chapter.

2. Nanotechnology as a Foundation for Intelligent Systems

Intelligent systems depend on the continuous acquisition, processing, and transmission of physical information, and on the ability of a system to respond adaptively to changes in its environment. Nanotechnology contributes to each of these functions through the design of

nanoscale sensors, low-power electronic components, and materials capable of adaptive or self-regulating behaviour.

2.1 Nanosensors and Smart Sensing

Nanosensors exploit the extremely high surface-area-to-volume ratio of nanostructured materials such as carbon nanotubes, graphene, and metal-oxide nanowires to achieve sensitivity levels that are difficult to attain with conventional sensing elements. Because chemical and biological interactions occur predominantly at surfaces, a nanostructured sensing layer can register the adsorption of even a small number of analyte molecules as a measurable change in electrical resistance, capacitance, or optical response (Bhushan, 2017). Such sensors are increasingly integrated into structural health monitoring systems, wearable health devices, and environmental monitoring networks, forming the sensory layer of many intelligent infrastructure systems.

2.2 Nanoelectronics and Low-Power Computing

The continued miniaturisation of transistor architectures, together with the emergence of novel nanoscale switching elements such as memristors and two-dimensional semiconductor channels, has extended the trajectory of computing performance improvements while reducing per-operation energy consumption. These developments are particularly significant for edge computing applications, where intelligent devices must perform local inference with strict power constraints, for instance in remote environmental sensors or implantable medical devices.

2.3 Smart and Self-Healing Nanomaterials

Materials embedded with nanoscale capsules of healing agents, or engineered with reversible molecular bonds, can autonomously repair microcracks before they propagate into structural failures. Such self-healing behaviour extends the operational life of components and reduces the maintenance burden associated with critical infrastructure, an outcome that aligns closely with the objectives of intelligent, self-monitoring engineering systems.

A related category of smart nanomaterials includes shape-memory nanocomposites and stimuli-responsive polymers embedded with nanoparticle fillers. These materials can alter their mechanical stiffness, colour, or geometry in response to changes in temperature, humidity, magnetic field, or applied voltage. When coupled with embedded nanosensors, such materials form the basis of adaptive structural elements capable of adjusting their own behaviour in real time, for example load-bearing components that stiffen automatically under increased mechanical stress or building facades that modulate light transmission in response to ambient conditions.

2.4 Nanophotonics and Optical Intelligence

Nanophotonic structures, including plasmonic nanoparticles, photonic crystals, and quantum dots, enable precise manipulation of light at sub-wavelength scales. These structures underpin high-density optical data storage, ultra-compact photonic integrated circuits, and highly sensitive optical biosensors. Because optical signals can be transmitted with very low energy loss compared with electrical signals over long distances, nanophotonic components are increasingly used within intelligent systems that require both high-speed data transmission and low overall

power consumption, such as distributed sensor networks and high-performance computing interconnects.

2.5 Comparative Overview of Nanomaterials for Intelligent Systems

Table 1 summarises representative classes of nanomaterials discussed in this section, together with their characteristic properties and principal application domains within intelligent systems engineering.

Nanomaterial	Characteristic Property	Representative Application
Carbon nanotubes	High electrical conductivity, high tensile strength	Nanosensors, structural composites
Graphene	High surface area, excellent charge mobility	Flexible electronics, gas sensors
Metal-oxide nanowires	High surface reactivity	Chemical and biological sensing
Quantum dots	Size-tunable optical emission	Optical sensors, display technology
Plasmonic nanoparticles	Strong localised surface plasmon resonance	Biosensors, photonic circuits
Nanocomposite polymers	Stimuli-responsive mechanical behaviour	Self-healing and adaptive structures

3. Nanotechnology for Sustainable Design

Sustainable engineering seeks to minimise resource consumption and environmental impact across the full life cycle of a product or system. Nanotechnology contributes to this objective through improvements in energy conversion efficiency, resource-efficient manufacturing, and the development of materials with reduced environmental footprint.

3.1 Energy Harvesting and Storage

Nanostructured electrode materials, including silicon nanowires, graphene composites, and transition-metal oxide nanoparticles, have substantially increased the charge storage capacity and cycling stability of batteries and supercapacitors. Similarly, nanoscale surface texturing and quantum-confinement effects in thin-film photovoltaic cells have improved light absorption and energy conversion efficiency, supporting the broader transition toward renewable energy systems.

3.2 Water Treatment and Environmental Remediation

Nanofiltration membranes and functionalised nanoparticles are used to remove heavy metals, organic pollutants, and pathogens from water at efficiencies that often exceed those of conventional treatment methods, while requiring lower energy input. Photocatalytic nanomaterials such as titanium dioxide nanoparticles are additionally used to degrade organic contaminants under ultraviolet or visible-light illumination, offering a route to decentralised, low-energy water purification suited to resource-constrained settings.

3.3 Green Construction and Lightweight Materials

The incorporation of nanoscale additives such as nano-silica and carbon nanotubes into cementitious composites improves mechanical strength and durability while allowing a reduction in overall material volume, thereby lowering the embodied carbon associated with construction. Nanocomposite materials used in the automotive and aerospace sectors similarly reduce structural weight, contributing to reduced fuel consumption and lower lifecycle emissions.

3.4 Green Nanotechnology and Cleaner Manufacturing

Green nanotechnology emphasises the use of environmentally benign synthesis routes, including biological and plant-mediated synthesis of nanoparticles, as an alternative to conventional chemical reduction methods that often rely on hazardous reagents and generate toxic by-products. Such approaches reduce the environmental burden of nanomaterial production itself, addressing a criticism sometimes directed at the nanotechnology sector regarding the sustainability of its own supply chain.

3.5 Sustainable Transportation

In the transportation sector, nanocomposite coatings reduce aerodynamic drag and surface friction, while nanostructured catalytic converters improve the efficiency of exhaust gas after-treatment in internal combustion engines, lowering emissions of nitrogen oxides and particulate matter. In electric vehicles, nanostructured battery electrodes and solid-state electrolytes are central to ongoing efforts to increase driving range while reducing charging time and improving battery safety, directly supporting the broader transition toward low-emission mobility systems.

3.6 Waste Reduction and Circular Manufacturing

Nanotechnology also supports circular economy principles through the development of nanocatalysts that improve the efficiency of chemical recycling processes, converting plastic and polymer waste back into usable feedstock at lower energy cost than conventional thermal recycling. Nanocoatings that extend the service life of packaging and industrial equipment further reduce the overall material throughput required to deliver a given engineering function, a principle central to resource-efficient sustainable design.

4. Illustrative Case Studies

The following case studies, drawn from widely reported engineering practice, illustrate how the principles discussed above translate into deployed systems that combine intelligent functionality with sustainability objectives.

4.1 Case Study: Nanostructured Photovoltaic Coatings

Thin-film solar modules incorporating nanostructured anti-reflective coatings have demonstrated measurable improvements in light-trapping efficiency compared with uncoated panels, allowing a greater proportion of incident solar radiation to be converted into electrical energy without increasing the physical footprint of the installation. When paired with embedded nanosensors that monitor panel temperature and soiling levels, such systems can additionally schedule automated cleaning cycles, an example of nanotechnology contributing simultaneously to energy efficiency and system-level intelligence.

4.2 Case Study: Nanofiltration for Decentralised Water Treatment

Portable nanofiltration units deployed in resource-constrained regions have been used to provide potable water from contaminated surface sources without reliance on centralised treatment infrastructure or continuous electrical power. Integration of low-power turbidity and contaminant nanosensors within these units allows real-time water-quality monitoring, enabling automated alerts when filtration performance degrades and supporting predictive maintenance scheduling in remote installations.

4.3 Case Study: Nanocomposite Structural Health Monitoring

Bridge and tunnel structures instrumented with carbon-nanotube-based strain sensors embedded within structural composites have been used to provide continuous, distributed monitoring of mechanical strain, moisture ingress, and micro-crack formation. Because such sensors can be manufactured as thin, lightweight films, they can be integrated into structural elements without materially increasing overall structural weight, illustrating the compatibility of nanoscale sensing with lightweight, sustainable structural design.

5. Emerging and Cross-Disciplinary Applications

5.1 Healthcare and Diagnostics

Nanoparticle-based drug delivery systems allow therapeutic agents to be targeted selectively to diseased tissue, reducing systemic side effects and improving treatment efficacy. Nanoscale biosensors integrated with point-of-care diagnostic devices enable rapid, low-cost detection of disease biomarkers, supporting the broader movement toward decentralised and personalised healthcare.

5.2 Precision Agriculture

Nano-formulated fertilisers and pesticides improve nutrient uptake efficiency and reduce the quantity of active chemical agents released into the environment, while nanosensor networks deployed in agricultural fields provide real-time data on soil moisture, nutrient status, and crop health, supporting data-driven irrigation and fertilisation decisions. Encapsulation of agrochemicals within biodegradable nanocarriers additionally allows controlled, slow release of active ingredients, reducing the frequency of chemical application required over a growing season.

5.3 Environmental Monitoring and Smart Cities

Networks of low-power, nanomaterial-based gas and pollution sensors distributed across urban environments provide the granular environmental data required by smart-city management systems, supporting more responsive and resource-efficient urban infrastructure planning. Such networks are frequently combined with cloud-based analytics platforms that aggregate sensor readings across a city to identify pollution hotspots, optimise traffic flow, and inform public health advisories in near real time.

5.4 Textiles and Consumer Products

Nanocoated textiles incorporating silver or titanium dioxide nanoparticles exhibit antimicrobial and self-cleaning properties, reducing the frequency of washing required and thereby lowering

the associated water and energy consumption over a garment's service life. Similar nanocoating approaches are used in consumer electronics housings and packaging to improve scratch resistance and extend product lifespan, indirectly supporting waste-reduction objectives.

6. Challenges and Limitations

Despite considerable progress, several challenges continue to constrain the widespread adoption of nanotechnology within intelligent and sustainable engineering systems, as summarised below.

- **Nanotoxicity and biological safety:** the long-term health and ecological effects of engineered nanoparticles remain incompletely understood, particularly regarding their behaviour once released into air, soil, and water systems.
- **Scalability of fabrication:** many nanomaterial synthesis routes that perform well at laboratory scale are difficult or costly to scale to industrial production volumes without loss of quality or increase in environmental impact.
- **Regulatory and standardisation gaps:** the absence of harmonised international standards for nanomaterial characterisation, labelling, and safety testing complicates regulatory approval and cross-border trade.
- **Economic barriers:** the high capital cost associated with nanofabrication equipment and quality control can limit adoption, particularly among small and medium-sized manufacturers.
- **End-of-life management:** the recovery, recycling, and safe disposal of nanomaterial-containing products remain underdeveloped relative to conventional material streams.
- **Workforce and skills gap:** the interdisciplinary expertise required to design, fabricate, and safely handle nanomaterials is not yet uniformly available across engineering education systems, constraining the pace of industrial adoption.

Beyond these individual barriers, a further underlying difficulty is the frequent mismatch between the pace of nanomaterials research and the pace at which supporting toxicological, regulatory, and standardisation frameworks are developed. New nanomaterial formulations are routinely reported in the scientific literature well before comprehensive safety data become available, creating a persistent lag that regulators and manufacturers alike must navigate. Addressing this gap will likely require closer collaboration between materials researchers, toxicologists, and policymakers from the earliest stages of nanomaterial development, rather than safety assessment being treated as a late-stage or post-hoc activity.

7. Future Perspectives

The convergence of nanotechnology with artificial intelligence is expected to accelerate materials discovery considerably. Machine learning models trained on existing materials databases can predict the properties of candidate nanostructures prior to synthesis, substantially reducing the time and resource cost associated with traditional trial-and-error experimentation. Similarly, the integration of nanoscale sensing with edge artificial intelligence is expected to

enable increasingly autonomous infrastructure systems capable of self-diagnosis and adaptive maintenance scheduling.

From a sustainability perspective, future research is likely to place greater emphasis on life-cycle assessment of nanomaterials, ensuring that gains achieved during the use phase of a product are not offset by environmental costs incurred during production or disposal. The continued development of green synthesis routes, coupled with the emergence of clearer international regulatory frameworks, will be essential to ensuring that nanotechnology develops as a genuinely sustainable engineering discipline rather than merely a performance-enhancing one.

A further promising direction lies in the development of biodegradable and bio-derived nanomaterials capable of matching the performance of conventional synthetic nanomaterials while decomposing safely at the end of a product's service life. Progress in this area would substantially reduce the long-term environmental accumulation concerns that currently accompany many engineered nanomaterials. Likewise, the emergence of digital twin frameworks that couple physical nanosensor networks with real-time computational models is expected to further blur the boundary between nanoscale sensing hardware and system-level intelligence, enabling engineering systems that not only monitor their own condition but actively predict and pre-empt failure.

Finally, international collaboration on nanotechnology standardisation, echoing the harmonisation efforts already underway through bodies such as the International Organization for Standardization, will likely play an increasingly central role in enabling the safe and consistent global adoption of nanotechnology-enabled products, particularly as supply chains for nanomaterial-containing components become more globally distributed.

Conclusion

Nanotechnology occupies a foundational position within the broader landscape of intelligent systems and sustainable engineering. Its capacity to enable highly sensitive sensing, energy-efficient computation, and resource-conscious materials design makes it a natural enabling technology for the goals articulated throughout this book. At the same time, unresolved questions concerning nanomaterial safety, manufacturing scalability, and regulatory oversight indicate that continued interdisciplinary research, spanning materials science, environmental engineering, computer science, and policy, will be required before the full potential of nanotechnology can be responsibly realised. Engineering tomorrow's intelligent and sustainable systems will depend substantially on the extent to which today's nanotechnology research addresses these open challenges.

References

1. Bhushan, B. (2017). *Springer handbook of nanotechnology* (Vol. 1, pp. 1–45). Springer.
2. Feynman, R. P. (1960). There's plenty of room at the bottom. *Engineering and Science*, 23(5), 22–36.
3. Roco, M. C. (2011). The long view of nanotechnology development: The National Nanotechnology Initiative at 10 years. *Journal of Nanoparticle Research*, 13(2), 427–445.

4. Sharma, A. (2015). Nanomaterials for sustainable engineering applications. *Journal of Materials Science and Engineering*, 4(2), 112–121.
5. Sharma, A., & Singh, R. (2016). Green synthesis of nanoparticles: A review of sustainable approaches. *International Journal of Nanoscience and Nanotechnology*, 12(3), 201–215.
6. Sharma, A., Singh, R., & Kumar, V. (2017). Nanosensors for intelligent environmental monitoring systems. *Sensors and Actuators B: Chemical*, 245, 678–690.
7. Whitesides, G. M. (2003). The ‘right’ size in nanotechnology. *Nature Biotechnology*, 21(10), 1161–1165.
8. Kumar, V., & Sharma, A. (2018). Nanocomposite materials for sustainable structural engineering. *Construction and Building Materials*, 165, 834–846.
9. Singh, R. (2019). Nanotechnology in precision agriculture: Prospects and challenges. *Journal of Agricultural Engineering*, 56(1), 45–58.
10. National Nanotechnology Initiative. (n.d.). *National Nanotechnology Initiative*. <https://www.nano.gov>
11. *Nature Nanotechnology*. (n.d.). *Nature Nanotechnology*. <https://www.nature.com/nnano/>
12. International Organization for Standardization. (n.d.). *Nanotechnologies standards*. <https://www.iso.org/committee/381983.html>

ENGINEERING TRUSTWORTHY AI APPLICATIONS: PRACTICAL STRATEGIES FOR SECURE AND RESPONSIBLE AI DEVELOPMENT

Udaya Bhaskar Vemuri

Application Security Analyst and DevSecOps Professional

Corresponding author E-mail: udv0507@gmail.com

Abstract

Artificial Intelligence (AI) is now part of customer service, software development, fraud detection, healthcare support, education, and business decision-making. These applications can improve productivity and access to information, but they also create security and trust concerns. AI systems may generate inaccurate content, expose sensitive data, behave unfairly, or be manipulated by attackers when security and governance are not built into the application from the beginning.

This chapter presents practical methods for engineering secure and trustworthy AI applications from an application-security and software-engineering perspective. It explains common risks, including hallucinations, prompt injection, data leakage, bias, model drift, and excessive reliance on AI. It also provides a lifecycle framework covering business requirements, risk assessment, secure architecture, data protection, model development, testing, deployment, human oversight, monitoring, and continuous improvement.

The central argument is that trustworthy AI cannot be achieved through technology alone. It requires clear governance, secure development practices, regular testing, transparent use, accountable human review, and monitoring after deployment. The chapter is intended for students, developers, researchers, security professionals, and business stakeholders seeking practical guidance that balances innovation with security, privacy, and responsible use.

Keywords: Artificial Intelligence, Trustworthy AI, Application Security, DevSecOps, Secure Software Development, AI Security, Responsible AI, AI Governance.

1. Introduction

Artificial Intelligence is changing how people interact with technology. Organizations use AI to answer customer questions, summarize documents, assist software developers, identify fraud, support medical work, and improve operational decisions. These uses can create significant value, but they also introduce risks that are different from those found in traditional rule-based software. [1–3]

Traditional applications generally follow explicit instructions written by developers. AI systems generate outputs from patterns learned from data and may respond differently to similar requests. This flexibility is useful, but it can also produce incorrect answers, unexpected behavior, or inconsistent results. Security teams must therefore evaluate both the surrounding application and the behavior of the AI component. [1, 4, 5]

Consider a customer-service chatbot that helps users find account information. It may handle thousands of requests correctly, yet a carefully written instruction could attempt to make it ignore its rules, reveal confidential information, or perform an unauthorized action. Similar concerns apply to AI used in software development, banking, healthcare, government, and cybersecurity. [6, 7]

Building trustworthy AI therefore requires more than selecting an accurate model. Organizations must apply secure software-development principles, define acceptable uses, protect data, test AI-specific threats, preserve human accountability, and monitor the system throughout its operational life. This chapter explains how existing Secure Software Development Lifecycle practices can be extended to address AI-specific risks. [1, 8]

2. Understanding Trustworthy AI

Trust influences whether people will use and rely on an AI system. Trustworthy AI does not mean that the system will never make a mistake. It means that the organization understands the system's limits, manages risk, communicates appropriately with users, and maintains controls that reduce foreseeable harm. Five characteristics are especially important: reliability, security, privacy, transparency, and human oversight. [1–3]

2.1 Reliability

AI should provide reasonably consistent and dependable results under expected operating conditions. Teams should define accuracy expectations, test representative scenarios, review failures, and monitor whether performance changes over time. [1, 4]

2.2 Security

AI applications must protect systems, models, interfaces, credentials, and data from unauthorized access and misuse. Authentication, authorization, encryption, secure APIs, vulnerability management, logging, and incident response remain essential. [8, 9]

2.3 Privacy

Many AI applications process personal, confidential, or regulated information. Organizations should collect only necessary data, restrict access, protect information in transit and at rest, define retention periods, and use approved platforms that provide appropriate privacy protections. [1, 4, 10]

2.4 Transparency

Users should know when they are interacting with AI and should receive explanations appropriate to the context. Clear notices, links to authoritative sources, and understandable limitations help users decide when additional verification is required. [2, 3]

2.5 Human Oversight

AI should support people rather than replace accountable judgment in high-impact situations. Human review is especially important in healthcare, finance, employment, public services, legal decisions, and cybersecurity. [1–3]

Table 1: Characteristics of Trustworthy AI

Characteristic	Why It Matters	Practical Example
Reliability	Produces dependable results	A chatbot consistently answers routine questions
Security	Protects systems and data	Strong authentication prevents unauthorized access
Privacy	Protects personal information	Sensitive data is encrypted and access is restricted
Transparency	Builds user confidence	Users are told when responses are AI-generated
Human oversight	Reduces critical mistakes	A qualified professional review high-impact recommendation

Source: Author's own compilation based on the cited frameworks.




Figure 1: Foundations of Trustworthy AI

3. Common Security Risks in AI Applications

AI systems can improve productivity, but they can also produce unexpected results or be deliberately manipulated. Understanding common risks allows organizations to design controls before the system is deployed. [1, 6, 7]

3.1 AI Hallucinations

A hallucination occurs when an AI system provides incorrect, incomplete, or fabricated information in a confident manner. For example, an internal assistant may summarize a company policy but include a procedure that does not exist. Important outputs should be verified against authoritative sources, and the system should be designed to acknowledge uncertainty when reliable information is unavailable. [1, 6]

3.2 Prompt Injection

Prompt injection uses crafted instructions to influence an AI system to ignore intended rules, reveal protected information, or perform an unsafe action. Controls should include input and output validation, separation of trusted instructions from user content, least-privilege access, restrictions on sensitive actions, and adversarial testing. [6, 7]

3.3 Data Leakage

Employees may enter customer data, internal documents, credentials, or proprietary source code into an unapproved public AI service. Even when the purpose is legitimate, the information may leave the organization's-controlled environment. Clear data-classification rules, approved enterprise platforms, access controls, masking, and employee training can reduce this risk. [4–6]

3.4 AI Bias

AI learns from data that may contain incomplete or unfair historical patterns. A hiring, lending, or recommendation system may therefore produce unequal results for different groups. Teams should evaluate representative data, compare outcomes, document limitations, and establish a process for correcting unfair behavior. [1–3]

3.5 Model Drift

Model drift occurs when AI performance declines because business conditions, user behavior, threats, or data patterns change. A fraud-detection system, for example, may become less effective as criminal techniques evolve. Monitoring, periodic evaluation, controlled updates, and retraining help maintain reliable performance. [1, 4]

3.6 Excessive Trust in AI

Users may treat an AI response as a confirmed fact or allow it to make decisions outside its intended role. AI may miss legal, operational, security, or business context that an experienced professional would recognize. Organizations should define when human approval is mandatory and train users to treat AI output as assistance rather than unquestionable authority. [1–3]

Table 2: Common AI Security Risks and Practical Mitigations

Risk	Possible Impact	Practical Mitigation
Hallucinations	Incorrect information	Verify important outputs against trusted sources
Prompt injection	Manipulated AI behavior	Validate content and restrict sensitive actions
Data leakage	Exposure of confidential data	Use approved platforms and data-classification rules
Bias	Unfair or inconsistent decisions	Review outcomes and improve data and controls
Model drift	Reduced accuracy over time	Monitor performance and retrain when needed
Excessive trust	Poor or unsafe decisions	Require human review for high-impact actions

Practical Example

A financial-services organization introduced an AI assistant for internal policy questions. During testing, the assistant occasionally referenced outdated procedures. The organization limited the assistant to approved policy documents and added links to the original source in every response. Employees were instructed to confirm important information before acting. This simple control increased confidence while reducing the chance that an outdated AI response would be treated as official guidance.

4. A Practical Framework for Secure AI Development

Organizations do not need a completely separate security program for AI. They can extend existing Secure SDLC and DevSecOps practices with controls that address AI data, models, prompts, outputs, and automated actions. The following lifecycle provides a practical structure. [1, 4, 5, 8]

4.1 Define the Business Purpose and Boundaries

Teams should identify the problem being solved, intended users, data the AI will process, decisions it may support, prohibited uses, and the points where human approval is required. Clear boundaries reduce unnecessary access and prevent the system from expanding beyond its approved purpose. [1, 5]

4.2 Perform an Early Risk Assessment

The risk assessment should consider data sensitivity, possible incorrect outputs, unauthorized access, privacy obligations, regulatory requirements, business continuity, and the consequences of automated actions. A low-impact information assistant does not require the same controls as a medical, financial, or public-safety system. [1, 4, 10]

4.3 Design Secure Architecture and Data Flows

Architecture reviews should identify where data is stored, how users and services authenticate, which systems the AI can reach, what information it must never access, and how logs and secrets are protected. Least privilege, secure APIs, encryption, network controls, and separation of duties should be designed before development is complete. [8, 9]

4.4 Protect Data and Develop the Model Securely

Teams should collect only necessary information, remove unnecessary personal data, validate data quality, document important assumptions, record model and prompt versions, protect training and retrieval sources, and review dependencies. Sensitive data should have defined owners, access rules, retention periods, and approved disposal methods. [1, 4, 8]

4.5 Test Functional and AI-Specific Risks

Testing should cover accuracy, reliability, performance, privacy, bias, access control, prompt injection, data leakage, unsafe outputs, and abuse of connected tools. Tests should include normal users, unusual inputs, malicious instructions, and failure conditions. Security testing should occur throughout development, not only immediately before release. [6–8]

4.6 Deploy with Controls and Human Accountability

Before release, teams should review configuration, permissions, logging, monitoring, rollback plans, user notices, and operational ownership. High-impact recommendations or actions should require appropriate human review. AI-generated decisions should not remove accountability from the people and organization operating the system. [1–3]

4.7 Monitor and Continuously Improve

After deployment, organizations should monitor response quality, user feedback, security events, availability, model performance, policy changes, and emerging attack techniques. Incidents and failures should lead to updated prompts, data, controls, tests, documentation, or retraining. Continuous monitoring is necessary because both the operating environment and the AI system can change. [1, 4, 5]

Table 3: Secure AI Development Lifecycle Activities

Stage	Primary Objective	Security Activities
Planning	Define purpose and boundaries	Risk identification and compliance review
Design	Build secure architecture	Threat modeling and access-control planning
Development	Build the AI application	Secure coding, data validation, version control
Testing	Verify expected behavior	Functional, privacy, bias, and adversarial testing
Deployment	Release safely	Configuration review, approval, rollback planning
Operations	Maintain security and reliability	Monitoring, incident response, vulnerability management
Improvement	Adapt to change	Model updates, retraining, lessons learned

5. Practical Example: Building Trust Through Secure AI Development

Consider a company developing an AI assistant to help employees locate internal policies. During planning, the team asks which documents the assistant may access, how confidential information will be protected, what happens when the assistant is uncertain, and who is responsible for approving high-impact guidance.

The organization limits the assistant to approved policy documents, applies role-based access, and prevents it from directly changing employee accounts or records. Each response identifies that it was generated by AI and includes a link to the original policy. Employees are instructed not to enter customer secrets or other restricted information into the tool.

Before deployment, the team tests authentication, authorization, prompt injection, data leakage, inaccurate responses, and failure conditions. Security and business owners review the results and approve the intended use. After release, the organization monitors feedback and security logs, updates the knowledge source when policies change, and investigates repeated incorrect responses. [6–8]

This example shows that trust comes from a combination of useful technology and responsible operating practices. Limiting access, linking answers to authoritative sources, preserving human

responsibility, and monitoring the system are often more important than relying on the model alone. [1, 4, 5]

6. Best Practices for Organizations

Organizations can improve AI security without slowing responsible innovation by applying a small number of consistent practices across projects. [1, 2, 5]

- Start security and privacy reviews during planning, before architecture and data-access decisions become difficult to change.
- Approve AI platforms and define what categories of information employees may share with them.
- Use least privilege so AI applications can access only the data, tools, and actions required for their approved purpose.
- Require human review for high-impact recommendations and clearly assign ownership for AI-supported decisions.
- Test regularly for accuracy, bias, prompt injection, data leakage, unauthorized access, and changes in model performance.
- Maintain governance that explains acceptable use, monitoring responsibilities, incident reporting, documentation, and periodic review.

These practices work best when they are integrated into existing software security, privacy, risk, and compliance processes. A clear and repeatable approach is more sustainable than relying on individual teams to interpret responsible AI requirements independently. [4, 5, 8, 9]

7. Future Directions

As AI capabilities expand, organizations will place greater emphasis on governance, privacy, explainability, adversarial testing, and continuous monitoring. Regulatory expectations will continue to develop, and collaboration among developers, security professionals, business leaders, legal teams, and regulators will become more important. Organizations that establish practical controls now will be better prepared to adopt new AI capabilities responsibly. [2, 3, 10]

Conclusion

Artificial Intelligence can improve productivity, automate routine work, and support better decisions, but it also introduces security, privacy, reliability, and governance challenges. Trustworthy AI is not achieved by eliminating every possible error. It is achieved by understanding risk, setting clear boundaries, applying appropriate controls, and maintaining accountable oversight throughout the system lifecycle.

Many required practices are already familiar to organizations with mature software-security programs: careful planning, secure architecture, data protection, testing, access control, monitoring, incident response, and continuous improvement. AI-specific threats such as prompt injection, hallucinations, bias, and model drift should be incorporated into those existing processes rather than treated as isolated technical problems. The organizations that succeed with AI will not necessarily be those using the most advanced models. They will be those that develop

and operate AI responsibly, protect sensitive information, communicate limitations, preserve human judgment, and place user trust at the center of innovation.

References

1. Tabassi, E. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-1>
2. Organisation for Economic Co-operation and Development. (2024). *OECD AI principles*. <https://oecd.ai/en/ai-principles>
3. United Nations Educational, Scientific and Cultural Organization. (2021). *Recommendation on the ethics of artificial intelligence*. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000381137>
4. International Organization for Standardization. (2023). *ISO/IEC 23894:2023, Information technology—Artificial intelligence—Guidance on risk management*. <https://www.iso.org/standard/77304.html>
5. International Organization for Standardization. (2023). *ISO/IEC 42001:2023, Information technology—Artificial intelligence—Management system*. <https://www.iso.org/standard/42001.html>
6. OWASP Foundation. (2024). *OWASP Top 10 for LLM applications 2025*. OWASP GenAI Security Project. <https://genai.owasp.org/resource/owasp-top-10-for-llm-applications-2025/>
7. MITRE. (n.d.). *MITRE ATLAS: Adversarial threat landscape for artificial-intelligence systems*. Retrieved August 5, 2026, from <https://atlas.mitre.org/>
8. Scarfone, K., Souppaya, M., & Dodson, D. (2022). *Secure software development framework (SSDF) version 1.1: Recommendations for mitigating the risk of software vulnerabilities* (NIST Special Publication 800-218). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-218>
9. Pascoe, C., Quinn, S., & Scarfone, K. (2024). *The NIST cybersecurity framework (CSF) 2.0* (NIST CSWP 29). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.CSWP.29>
10. European Parliament and Council of the European Union. (2024). Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, L 2024/1689. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>

DEVELOPMENT OF A SOLAR-POWERED BATTERY CHARGING SYSTEM FOR E-BIKES

Rahul Nath¹, Asif Beg^{*2} and Pinakee Gogoi¹

¹Department of Agricultural Engineering, H.R.H. the P.O.W.I.E.T., Jorhat, Assam

²School of Agriculture and Environmental Sciences, Shobhit Institute of Engineering and Technology, Deemed-to-be University, Meerut, Uttar Pradesh, India

*Corresponding author E-mail: asifbeg1996@gmail.com

Abstract

The system utilizes photovoltaic (PV) panels to convert solar energy into electrical energy, which is then regulated through a charge controller and supplied to the battery for safe charging. The proposed charging unit aims to reduce dependence on conventional grid electricity and promote the use of renewable energy sources. The performance of the system is evaluated under different environmental conditions to analyze charging efficiency and reliability. Results indicate that the solar-powered charging unit provides a practical, cost-effective, and environmentally friendly solution for charging batteries used in bicycles, electric bikes, and other low-power applications. The implementation of this system contributes to energy conservation, reduction of carbon emissions, and sustainable development, making it a viable alternative for regions with limited access to conventional electricity sources.

Keywords: Solar Panel, Battery Charging Unit, Renewable Energy, Photovoltaic System, Charge Controller and Sustainable Energy.

Introduction

Energy plays a vital role in the economic and technological development of any nation. The growing demand for electricity, coupled with the depletion of conventional energy resources, has increased the need for alternative and sustainable energy sources. Traditional sources of energy such as coal, petroleum, and natural gas are finite in nature and contribute significantly to environmental pollution. Therefore, the utilization of renewable energy sources has become essential for meeting future energy demands while protecting the environment (Bilgili, *et al.*, 2016).

Among the various renewable energy sources, solar energy is one of the most abundant and widely available resources. The sun provides a continuous supply of energy that can be converted into electrical power through photovoltaic (PV) technology. Solar energy is clean, renewable, environmentally friendly, and available in most parts of the world. Due to these advantages, solar power has become an important solution for sustainable energy generation.

The increasing use of batteries in electric vehicles, communication systems, portable electronic devices, and backup power applications has created a growing demand for efficient charging systems. Conventional battery charging methods mainly depend on grid electricity, which may not always be available, especially in rural and remote areas. In such situations, solar-powered

battery charging systems offer a practical and reliable alternative (Anynomus, 2026 and Yap, *et al.*, 2022).

The adoption of electric vehicles and portable electronic devices has increased rapidly, leading to a greater demand for efficient battery charging solutions. Solar-powered charging systems provide a sustainable method of meeting this demand while reducing the burden on conventional power grids. These systems help conserve energy and promote the use of renewable resources in everyday applications. Furthermore, solar battery charging units contribute to environmental protection by reducing greenhouse gas emissions and dependence on fossil fuels. As awareness of clean energy technologies continues to grow, solar-powered systems are becoming more affordable and accessible. Their ability to provide reliable power in both urban and rural areas makes them an important technology for future energy development and sustainable growth. The development of a battery charging unit using solar panel promotes the use of renewable energy and supports sustainable development. It provides an economical, environmentally friendly, and practical solution for battery charging applications while contributing to energy conservation and the reduction of carbon emissions (Mohapatra, *et al.*, 2018 and Weniger, *et al.*, 2018).

Materials and Methods

The construction of the battery charging unit using solar power for bikes was carried out in a systematic manner to ensure efficient and safe operation. The system was designed to convert solar energy into electrical energy and use it for charging a 12 V bike battery.

Initially, the power requirement of the bike battery was determined, and suitable components such as the solar panel, charge controller, battery, connecting wires, fuse, and switch were selected. The solar panel was chosen based on its voltage and power output to ensure compatibility with the battery charging requirements.

A supporting frame was fabricated to securely mount the solar panel. The panel was installed at an appropriate angle to maximize the absorption of solar radiation throughout the day. The frame was designed to provide stability and protection against environmental conditions such as wind and rain.

The completed charging unit was then exposed to sunlight for testing. The battery was connected to the charging system, and charging was initiated. During the charging process, important parameters such as solar panel voltage, charging current, battery voltage, and charging duration were recorded at regular intervals. The collected data were used to evaluate the performance and efficiency of the charging unit.

The constructed system provided a simple, cost-effective, and environmentally friendly solution for charging bike batteries using renewable solar energy. The design can be easily implemented in rural and remote areas where access to conventional electricity is limited, thereby promoting sustainable energy utilization and reducing dependence on fossil fuels.

The battery charging unit uses a solar panel to convert sunlight into DC electricity via the photovoltaic effect. This DC power goes to a charge controller, which regulates voltage and current to protect the battery from overcharge, over-discharge, and reverse current. The

controller adjusts charging as the battery fills and stops or reduces current when fully charged. The stored chemical energy can later power the bike or other devices. The system is simple, eco-friendly, and economical, reducing energy costs and reliance on fossil fuels, shown in fig. 2.

Table 1: Materials used and cost of solar powered battery charging System

S.No.	Materials	Quantity	Cost (Rs.)
1.	Solar Panel	9	235
2.	Bike Battery	1	450
3.	Solar Charge Controller	1	300
4.	Connecting Wires	2, meter	50
5.	Multimeter	1	150
	Total Amount	13	1185



Figure 1: Testing of solar-powered bike battery charger




Figure 2: Circuit diagram

Results and Discussion

The solar charging system performed well during testing. The panel produced power in sunlight, and the charge controller regulated charging efficiently. Battery voltage rose steadily, showing effective energy storage. Charging rates varied with solar intensity; peak hours produced higher currents and faster charging. The system ran stably without significant voltage or current fluctuations.

Observations

- The solar panel produced electrical power when exposed to sunlight.
- The charge controller regulated the charging voltage and current effectively.
- The battery voltage increased steadily during the charging period.
- Charging performance was influenced by the intensity of sunlight.
- The system operated safely without overheating or overcharging.

Table 2: Solar charging system performance parameters

Parameter	Observations
Battery Voltage (Before Charging)	11.8 V
Battery Voltage (After Charging)	12.8 V
Solar Panel Output Voltage	18–20 V
Average Charging Current	1.5–2.0 A
Charging Duration	3–5 Hours
System Performance	Satisfactory

The solar-powered battery charging unit reliably converted sunlight into electrical energy and charged the battery through a charge controller. Charging rate depended on solar irradiance: faster during peak sunlight and slower in cloudy conditions, but charging continued. The charge controller regulated voltage, prevented overcharging, reduced losses, and improved battery life. Overall efficiency and performance were satisfactory for bike battery charging, offering a low-cost, ecofriendly, and practical solution in sunny areas.

Economic Analysis

The developed solar charging unit requires an initial investment for purchasing the solar panel, charge controller, battery accessories, and mounting structure. However, once installed, the operating cost is very low because sunlight is available free of charge.

Compared to conventional charging methods that depend on grid electricity, the solar charging unit can reduce electricity expenses over time. Therefore, the system can provide long-term economic benefits to users, especially in areas with high electricity costs.

Conclusions

The Battery Charging Unit Using Solar Panel for Bikes demonstrated a practical, cost-effective renewable energy solution by converting sunlight into electrical energy, regulating it via a charge controller, and safely charging a battery. Experimental measurements confirmed adequate voltage and current delivery under good sunlight, showing reliable performance without grid

electricity and reducing energy costs and carbon emissions. The portable, low-maintenance design suits rural and remote areas with limited grid access, offering users a dependable charging option. However, performance depends on sunlight availability; cloudy or nighttime conditions reduce charging unless improved storage or higher-efficiency panels are used. The project also enriched understanding of photovoltaic systems, charging principles, measurement techniques, and system integration, providing a strong foundation for future research. With continued advances in solar panels, storage technologies, and control methods, such solar charging units hold considerable promise for broader adoption, supporting sustainable transportation, energy conservation, and environmental protection.

References

1. Anonymus (2026). Integrated solar batteries: Design and device concepts. *ACS Energy Letters*, 8(8), 3509–3520.
2. Bilgili, F., Yıldırım, E. and Kocak, E. (2016). The role of renewable energy in sustainable development: A literature review. *Renewable and Sustainable Energy Reviews*, (65): 400–411. <https://doi.org/10.1016/j.rser.2016.07.019>.
3. Mohapatra, D., Padhee, S. and Jena, J. (2018). Design of solar powered battery charger: an experimental verification. *International Students' Conference on Electrical, Electronics and Computer Science (SCEECS)* (pp. 1-5). IEEE. <https://doi.org/10.1109/SCEECS.2018.8546929>
4. Weniger, J., Vetter, M. and Hesse, H. C. (2018). Solar charging batteries: Advances, challenges, and opportunities. *Joule*, 2(7):1234–1245. <https://doi.org/10.1016/j.joule.2018.06.005>.
5. Yap, K. Y., Chin, H. H. and Klemeš, J. J. (2022). Solar energy-powered battery electric vehicle charging stations: Current development and future prospects. *Renewable and Sustainable Energy Reviews*, 169, 112932.

TOWARDS SUSTAINABLE COMPUTING: THE ROLE OF ARTIFICIAL INTELLIGENCE, CLOUD COMPUTING AND EDGE TECHNOLOGIES

Haradhan Mardi

Department of Computer Science,
Suri Vidyasagar College, Suri, Birbhum- 731101, West Bengal, India.

Corresponding author E-mail: haradhanmardi@gmail.com

Abstract

The rapid growth of digital technologies has transformed modern society while substantially increasing the demand for computational resources, energy, and supporting infrastructure. Consequently, sustainability has emerged as an important consideration in the design and operation of computing systems. Sustainable computing seeks to minimize energy consumption, carbon emissions, electronic waste, and inefficient resource utilization without compromising the performance and accessibility of digital services. Artificial Intelligence (AI), cloud computing, and edge computing have significant and interconnected roles in achieving this objective. AI enables intelligent resource allocation, workload prediction, energy optimization, and automated infrastructure management; however, computationally intensive AI models also introduce considerable energy and hardware requirements. Cloud computing offers resource sharing, virtualization, workload consolidation, and elastic provisioning, while edge computing complements centralized infrastructure by processing data closer to its source, reducing latency and unnecessary data transmission. This chapter examines the contributions of AI, cloud computing, and edge technologies to sustainable computing while critically considering their environmental costs. Particular attention is given to Green AI, energy-efficient data centres, carbon-aware computing, computation offloading, and intelligent cloud–edge orchestration. Challenges related to increasing computational demand, embodied carbon, electronic waste, infrastructure heterogeneity, and sustainability measurement are also discussed. The chapter argues that sustainable computing requires an integrated, life-cycle-oriented approach in which computational efficiency, intelligent resource management, renewable energy, carbon awareness, and responsible technology development are considered together.

Keywords: Artificial Intelligence, Cloud Computing, Edge Computing, Energy Efficiency, Sustainable Computing, Green AI.

1. Introduction

Digital technologies have become fundamental to contemporary economic and social activities. Artificial Intelligence (AI), cloud computing, Internet of Things (IoT), big data analytics, mobile applications, and intelligent services are increasingly used in healthcare, education, agriculture, manufacturing, transportation, finance, scientific research, and public administration. These developments have significantly improved connectivity, automation, accessibility, and decision-making. At the same time, they have increased the demand for computational power, storage,

networking infrastructure, and electricity. The environmental implications of this growth are becoming increasingly important. Data centres operate continuously to support cloud platforms and digital services, communication networks transmit enormous volumes of data, and billions of end-users and IoT devices require both energy and physical resources. Although improvements in hardware and data-centre efficiency have moderated energy growth, the continued expansion of digital services and computationally intensive applications presents a significant sustainability challenge (Masanet *et al.*, 2020).

Sustainable computing provides a framework for addressing this challenge. Murugesan (2008) described Green IT as environmentally responsible computing throughout the life cycle of technological systems. The broader concept of sustainable computing incorporates not only operational energy efficiency but also carbon emissions, resource utilization, hardware longevity, manufacturing impacts, water consumption, electronic waste, and the responsible disposal of equipment. Artificial intelligence, cloud computing, and edge computing occupy an interesting position within this sustainability landscape. They contribute to environmental impact through their dependence on computational infrastructure, yet they also provide mechanisms for reducing energy consumption and improving resource efficiency. AI can optimize resource allocation and energy systems; cloud computing can consolidate workloads on shared infrastructure; and edge computing can reduce unnecessary data transmission by processing information closer to its source (Berl *et al.*, 2010). This dual role—computing as both a source of environmental impact and an instrument for reducing it—is central to sustainable computing.

This chapter examines these relationships individually and collectively. It discusses the environmental footprint of modern computing, explores the contribution of AI, cloud, and edge technologies to sustainability, and examines their convergence into an increasingly integrated computing continuum. Finally, key challenges and future directions are considered.

2. Sustainable Computing: Concept and Environmental Context

Sustainable computing can broadly be defined as the design, development, operation, management, and disposal of computing systems in ways that minimize environmental harm while maintaining technological and economic usefulness. It extends the conventional emphasis on processing speed, storage capacity, reliability, and cost by incorporating environmental considerations into computing decisions.

Two complementary dimensions can be distinguished. The first is sustainability of computing, which concerns reducing the environmental footprint of computing itself. Energy-efficient processors, optimized software, green data centres, efficient cooling, server consolidation, renewable energy, longer hardware life, and responsible electronic-waste management belong to this category.

The second is computing for sustainability, where computational technologies are used to improve environmental outcomes in other sectors. Examples include smart grids, renewable-energy forecasting, intelligent transportation, precision agriculture, smart buildings, environmental monitoring, and optimized manufacturing. AI in particular has considerable

potential to support energy and environmental management (Ahmad *et al.*, 2021; Rolnick *et al.*, 2022).

2.1 Energy and Carbon Footprint

Energy efficiency remains one of the most visible aspects of sustainable computing. Data centres, communication networks, and end-user devices collectively account for substantial electricity consumption. However, electricity consumption alone does not determine environmental impact. The carbon footprint of a computing workload also depends on the source of electricity, geographical location, time of execution, and efficiency of the underlying hardware. This distinction has encouraged a transition from energy-aware to carbon-aware computing. An energy-efficient task minimizes the amount of electricity required, whereas a carbon-aware task additionally considers when and where that electricity is produced. A workload executed during a period of abundant renewable electricity may produce fewer emissions than the same workload executed when the electricity grid depends more heavily on fossil fuels.

2.2 Life-Cycle Impact

Sustainable computing must also consider the complete hardware life cycle. Manufacturing processors, memory, storage devices, sensors, and networking equipment consumes energy, water, and raw materials. Freitag *et al.* (2021) emphasize that the climate impact of ICT cannot be adequately understood by examining operational electricity alone.

The rapid replacement of servers, accelerators, smartphones, and IoT devices further contributes to electronic waste. Therefore, extending equipment lifetime, improving repairability, reusing components, and recycling electronic equipment are important aspects of sustainable computing. Water consumption has also emerged as a concern, particularly in data centres that rely on water-intensive cooling systems (Mytton, 2021). Sustainability assessment should therefore increasingly include energy, carbon, water, materials, and hardware life-cycle considerations rather than relying on a single metric.

3. Artificial Intelligence and Sustainable Computing

AI has a particularly complex relationship with sustainability. It is simultaneously a growing consumer of computational resources and a powerful mechanism for improving resource efficiency.

3.1 Environmental Cost of AI

Modern machine-learning systems, particularly deep-learning and large-scale generative models, can require substantial computational resources. Training involves repeatedly processing large datasets while optimizing model parameters, sometimes using large numbers of GPUs or other accelerators. Strubell *et al.* (2019) brought considerable attention to the energy and policy implications of computationally intensive NLP models. The carbon footprint of large neural-network training is also strongly influenced by model design, hardware efficiency, and the energy source used by the computing infrastructure (Patterson *et al.*, 2021). The environmental footprint of AI is not limited to training. Once deployed, models perform inference whenever

they respond to new inputs. Although a single inference operation may consume considerably less energy than training, a widely deployed model can serve millions of requests. Consequently, inference can become an important component of the lifetime resource requirements of an AI service. Specialized hardware introduces another dimension. GPUs, TPUs, and other accelerators can significantly improve computational efficiency, but their manufacture also contributes to embodied carbon. Wu *et al.* (2022) therefore argue that sustainable AI should consider the complete relationship among data, algorithms, system hardware, operational energy, and manufacturing impact.

3.2 AI as a Sustainability Enabler

Despite these costs, AI can improve sustainability by making complex systems more adaptive and resource-efficient. Machine-learning models can predict computational workloads, allowing resources to be dynamically activated or placed into low-power states according to demand.

AI can also optimize virtual-machine placement, storage allocation, network resources, and task scheduling. In data centres, intelligent management can combine workload information with temperature, cooling demand, and equipment characteristics to improve operational efficiency. Machine-learning-based optimization of data-centre operation provides an important practical illustration of this approach (Gao, 2014).

Beyond computing infrastructure, AI is being applied to renewable-energy forecasting, smart-grid operation, intelligent transportation, building-energy management, precision agriculture, climate modelling, and environmental monitoring (Ahmad *et al.*, 2021; Rolnick *et al.*, 2022). AI's environmental value should therefore be evaluated not only according to the resources it consumes but also according to the efficiency improvements it enables.

3.3 Green AI and Efficient Model Design

The concept of Green AI seeks to incorporate computational efficiency into the evaluation and development of AI systems. Schwartz *et al.* (2020) argue that improvements in predictive performance should be considered alongside their computational cost rather than treating accuracy as the only important criterion. A systematic review by Verdecchia *et al.* (2023) further demonstrates the growing research interest in reducing and assessing the environmental impact of AI systems. Several techniques can reduce AI resource requirements. Pruning removes model parameters that contribute little to prediction. Quantization uses lower numerical precision, reducing memory and computational requirements. Knowledge distillation transfers useful capabilities from a large model to a smaller one, while transfer learning reduces the need to train new models entirely from the beginning. Efficient architectures and hardware-aware model design can provide additional improvements (Menghani, 2023).

These principles are particularly relevant to NLP and generative AI. Instead of using the largest available model for every application, sustainable AI encourages models appropriate to the complexity of the actual task. The objective is therefore not simply “smaller AI,” but an appropriate balance between capability, accuracy, computational cost, and environmental impact.

4. Cloud Computing and Sustainable Digital Infrastructure

Cloud computing has transformed the organization of information technology by providing on-demand access to shared computing resources. This model offers important opportunities for sustainability, although the environmental benefits depend strongly on infrastructure management.

4.1 Resource Consolidation and Virtualization

One of the major sustainability advantages of cloud computing is resource pooling. Traditional organizational servers frequently operate below their maximum capacity while continuing to consume electricity. Cloud platforms can consolidate multiple workloads onto shared physical infrastructure using virtualization and containers. Such consolidation can improve utilization while balancing the energy costs associated with processing, storage, and data transport (Baliga *et al.*, 2011).

Energy-aware virtual-machine allocation can migrate workloads away from underutilized servers and consolidate them onto fewer machines, allowing unused equipment to enter low-power states. Beloglazov *et al.* (2012) demonstrated the importance of dynamic consolidation and energy-aware resource allocation in reducing data-centre energy consumption.

Elastic provisioning further improves efficiency by increasing or decreasing resources according to workload demand. Thus, the sustainability advantage of cloud computing does not arise merely from centralization; it arises from the efficient utilization and dynamic management of shared infrastructure.

4.2 Sustainable Data Centres

Large cloud platforms depend on data centres containing servers, storage, network equipment, cooling infrastructure, and power-distribution systems. Improvements in these supporting systems are therefore essential for sustainable cloud computing.

Power Usage Effectiveness (PUE), which compares total facility energy consumption with the energy delivered to computing equipment, is widely used to assess data-centre infrastructure efficiency. Hyperscale facilities can achieve efficiency advantages through sophisticated cooling, optimized server utilization, and economies of scale (Masanet *et al.*, 2020).

Nevertheless, PUE is not a complete sustainability indicator. A data centre with excellent PUE may still have a significant carbon footprint if its electricity is generated from carbon-intensive sources. Similarly, water use, construction impacts, and hardware manufacturing are not represented by PUE.

4.3 Carbon-Aware Cloud Computing

Cloud infrastructure provides opportunities to shift flexible workloads across both time and geographical location. Data centres located in different regions may experience different electricity-grid carbon intensities and renewable-energy availability.

Carbon-aware computing seeks to exploit these variations. Radovanović *et al.* (2023) describe approaches in which computational demand can be adjusted according to the carbon intensity of

electricity. Non-urgent workloads can potentially be executed when cleaner energy is more readily available.

Such strategies represent an important conceptual development: the objective changes from merely consuming less electricity to consuming electricity more intelligently.

4.4 Serverless Computing and Resource Efficiency

Serverless computing provides another potential mechanism for reducing idle resource consumption. In serverless environments, computing resources are dynamically assigned when functions or applications execute rather than requiring permanently active virtual machines. Jonas *et al.* (2019) identify fine-grained resource allocation and elasticity as important characteristics of serverless computing. However, cloud sustainability also faces the rebound effect: improvements in efficiency can make digital services cheaper and easier to deploy, which may increase total consumption. Efficiency improvements therefore do not automatically guarantee reductions in overall environmental impact.

5. Edge Computing and Sustainability

Edge computing complements centralized cloud computing by moving selected computation and storage closer to users and data sources. This approach is increasingly important for IoT, industrial systems, autonomous applications, and real-time services.

5.1 Reducing Data Movement

IoT systems can generate continuous streams of sensor, image, audio, and operational data. Sending all raw data to distant cloud servers can create unnecessary network traffic and latency. Edge computing enables local filtering, aggregation, compression, and analysis before information is transmitted to the cloud (Shi *et al.*, 2016).

For example, a video-monitoring system may process images locally and transmit only detected events rather than continuously uploading an entire video stream. Similarly, an environmental sensor network can aggregate routine measurements locally and communicate only summaries or abnormal readings. Reducing data movement can lower communication overhead while also improving responsiveness.

5.2 Energy-Aware Computation Offloading

Local computation is not always more energy-efficient than cloud processing. Edge devices generally possess less computational capacity and may operate under battery or power constraints. The appropriate location for a task therefore depends on its characteristics.

Computation offloading determines whether a workload should execute on the originating device, a nearby edge server, or a remote cloud platform. The decision may depend on processing requirements, available bandwidth, latency, battery level, network condition, and energy consumption.

Jiang *et al.* (2020) emphasize that energy-aware edge computing must account for the interaction among edge devices, edge servers, and cloud data centres. Sustainable task placement therefore requires a system-level perspective rather than assuming that either local or centralized computation is universally preferable.

5.3 Edge Intelligence

The integration of AI with edge computing has led to edge intelligence, in which machine-learning capabilities are deployed close to data sources. Chen and Ran (2019) describe the opportunities and challenges associated with performing deep-learning operations at the edge.

Edge intelligence can reduce repeated data transmission, support real-time decisions, and improve privacy by keeping sensitive raw data closer to its source. However, resource-constrained edge devices require compact and efficient models. Pruning, quantization, knowledge distillation, and hardware-specific optimization therefore become particularly important.

5.4 Sustainability Limitations of Edge Computing

Edge computing should not automatically be regarded as environmentally superior to cloud computing. Large data centres benefit from highly optimized infrastructure and economies of scale, while distributed edge devices may have lower utilization and shorter operational lifetimes. The deployment of billions of connected devices also increases hardware manufacturing requirements and potentially electronic waste. Sustainability therefore depends on whether the reduction in data transmission and centralized processing outweighs the environmental cost of additional distributed infrastructure.

6. Convergence of AI, Cloud, and Edge Technologies

The most important opportunities for sustainable computing emerge when AI, cloud computing, and edge technologies operate as an integrated ecosystem rather than as isolated platforms.

Modern applications increasingly operate across a device–edge–cloud continuum. Sensors and end devices generate data; edge nodes perform time-sensitive or communication-efficient processing; cloud infrastructure provides large-scale computation and storage; and AI can optimize decisions throughout the system.

A conceptual workflow can be expressed as:

Data Generation → Edge Filtering/Processing → Selective Transmission → Cloud Processing and Storage → AI-Based Optimization → Intelligent Workload Placement.

Consider a smart agricultural system. Sensors may continuously collect soil moisture, temperature, weather, and crop information. Sending every raw measurement directly to the cloud would consume unnecessary network resources. Edge nodes can preprocess data, identify conditions requiring immediate intervention, and transmit selected information. Cloud infrastructure can store historical records and perform computationally demanding analytics, while AI can predict irrigation requirements and optimize resource use.

Similar architectures are applicable to smart cities, healthcare, industrial IoT, intelligent transportation, energy systems, and environmental monitoring.

AI-based orchestration can dynamically determine where a task should execute. Latency-sensitive processing may remain at the edge, while computationally intensive but delay-tolerant tasks can be moved to the cloud. If renewable-energy availability and grid carbon intensity are incorporated into these decisions, workload placement can additionally become carbon-aware.

The key principle is that sustainability should be optimized across the entire computing system. Reducing energy use at one component may simply transfer consumption to another. An edge device may save network energy but consume more computational energy locally; conversely, centralized processing may be highly efficient but require substantial data transmission.

6.1 Federated and Distributed Intelligence

Federated learning represents an additional example of cloud–edge collaboration. Instead of transferring all raw data to a central server, participating devices can train or update models locally and exchange model information. This can provide privacy and data-locality advantages, although communication and distributed computation themselves consume energy.

Consequently, federated learning should not automatically be classified as sustainable. Its environmental performance depends on the number of participating devices, communication frequency, local computation, network characteristics, and alternative centralized architecture.

6.2 From Energy-Aware to Carbon-Aware Orchestration

An advanced sustainable architecture would integrate real-time information about workload demand, latency, hardware utilization, network conditions, energy consumption, and carbon intensity. AI could use this information to determine whether a workload should be executed locally, at an edge node, or within a particular cloud region. Such systems represent a transition from isolated energy-saving mechanisms toward intelligent sustainability-aware orchestration.

7. Challenges and Future Directions

Despite substantial progress, sustainable computing faces several technical, environmental, and organizational challenges.

7.1 Standardized Sustainability Measurement

There is no single metric capable of representing the complete environmental impact of a computing system. Energy consumption, carbon intensity, water use, hardware lifetime, embodied carbon, and electronic waste represent different dimensions.

Future research should develop transparent and standardized reporting frameworks. Computational studies could increasingly report energy use and hardware requirements alongside traditional measures such as execution time and predictive accuracy.

7.2 Growing AI Demand

Improvements in AI efficiency may be offset by rapidly increasing demand for larger models and more frequent inference. Green AI therefore needs to become a fundamental design objective rather than a secondary optimization.

Future research should emphasize efficient architectures, model compression, hardware-aware algorithms, model reuse, and appropriate model selection. A small task-specific model may sometimes provide a more sustainable solution than a very large general-purpose model.

7.3 Life-Cycle and Embodied Environmental Impact

Operational efficiency alone is insufficient. Wu *et al.* (2022) highlight the importance of both operational and manufacturing carbon in assessing AI systems. Similar considerations apply to cloud and edge infrastructure.

Longer equipment lifetimes, modular hardware, repairability, reuse, recycling, and circular-economy principles should therefore become part of sustainable computing research.

7.4 Renewable-Energy-Aware Computing

Greater integration between computing infrastructure and renewable-energy systems represents an important future direction. Workloads that are flexible in time can potentially be scheduled according to renewable-energy availability. Cloud platforms have particular advantages because geographically distributed data centres can support spatial workload shifting. Future edge systems could similarly coordinate processing according to locally available renewable energy and battery conditions.

7.5 Intelligent Cloud–Edge Orchestration

Future sustainable computing environments will require dynamic orchestration across heterogeneous infrastructure. AI-based controllers may simultaneously consider energy, carbon intensity, latency, network bandwidth, privacy, reliability, and cost.

This is fundamentally a multi-objective optimization problem. Minimizing energy alone may increase latency; reducing network traffic may increase local computation; maximizing hardware utilization may create thermal constraints. Sustainable orchestration therefore requires balancing competing objectives rather than optimizing a single parameter.

7.6 Rebound Effects and Responsible Digitalization

Perhaps the most difficult challenge is that efficiency does not necessarily reduce overall consumption. More efficient computing can lower costs and encourage greater use of digital services, potentially offsetting environmental savings. Sustainable computing therefore requires attention to both efficiency and sufficiency. Researchers and technology providers should consider whether computational resources are proportionate to the value delivered rather than assuming that greater computation is inherently desirable.

7.7 Sustainability, Accessibility, and Equity

Sustainability also has a social dimension. Advanced cloud and AI infrastructure is concentrated in a relatively small number of regions and organizations, while the environmental impacts of hardware production, resource extraction, and electronic waste may be distributed differently.

Future sustainable computing frameworks should therefore consider accessibility and equity alongside technical efficiency. Environmentally responsible computing should support broad societal benefits rather than simply transferring environmental costs between regions or communities.

Conclusion

The relationship between artificial intelligence, cloud computing, edge technologies, and sustainability is complex and multidirectional. These technologies contribute to the environmental footprint of the digital economy through electricity consumption, data transmission, hardware manufacturing, cooling requirements, and electronic waste. At the same time, they provide powerful mechanisms for reducing environmental impact through intelligent resource management, infrastructure consolidation, efficient data processing, and support for

renewable-energy integration. AI can improve workload prediction, resource allocation, energy management, and system optimization, but the increasing computational demands of modern AI models make Green AI and efficient model design increasingly important. Cloud computing can achieve significant efficiency through virtualization, resource pooling, workload consolidation, and elastic provisioning, while carbon-aware scheduling offers opportunities to align computational demand with cleaner electricity. Edge computing can reduce unnecessary data movement and support efficient real-time processing, although the environmental costs of distributed infrastructure must also be considered.

The strongest pathway toward sustainable computing lies in the coordinated convergence of AI, cloud, and edge technologies. Intelligent orchestration across the device–edge–cloud continuum can determine where, when, and how computational workloads should execute according to performance requirements, network conditions, energy availability, and carbon intensity.

Achieving this vision will require continued advances in efficient AI models, carbon-aware computing, renewable-energy integration, sustainable hardware design, and life-cycle assessment. Greater transparency and standardized sustainability metrics will also be essential. Ultimately, sustainable computing should not be viewed as a constraint on digital innovation. Rather, it provides a framework for ensuring that technological progress remains computationally effective, economically viable, and environmentally responsible as digital systems become increasingly embedded in modern society.

References

1. Ahmad, T., Zhang, D., Huang, C., Zhang, H., Dai, N., Song, Y., & Chen, H. (2021). Artificial intelligence in sustainable energy industry: Status quo, challenges and opportunities. *Journal of Cleaner Production*, 289.
2. Baliga, J., Ayre, R. W. A., Hinton, K., & Tucker, R. S. (2011). Green cloud computing: Balancing energy in processing, storage, and transport. *Proceedings of the IEEE*, 99(1), 149–167.
3. Beloglazov, A., Abawajy, J., & Buyya, R. (2012). Energy-aware resource allocation heuristics for efficient management of data centers for cloud computing. *Future Generation Computer Systems*, 28(5), 755–768.
4. Berl, A., Gelenbe, E., Di Girolamo, M., Giuliani, G., De Meer, H., Dang, M. Q., & Pentikousis, K. (2010). Energy-efficient cloud computing. *The Computer Journal*, 53(7), 1045–1051.
5. Chen, J., & Ran, X. (2019). Deep learning with edge computing: A review. *Proceedings of the IEEE*, 107(8), 1655–1674.
6. Freitag, C., Berners-Lee, M., Widdicks, K., Knowles, B., Blair, G. S., & Friday, A. (2021). The real climate and transformative impact of ICT: A critique of estimates, trends, and regulations. *Patterns*, 2(9).
7. Gao, J. (2014). *Machine learning applications for data center optimization*. Google Research.

8. Jiang, C., Fan, T., Gao, H., Shi, W., Liu, L., Cérin, C., & Wan, J. (2020). Energy aware edge computing: A survey. *Computer Communications*, 151, 556–580.
9. Jonas, E., Schleier-Smith, J., Sreekanti, V., Tsai, C.-C., Khandelwal, A., Pu, Q., Shankar, V., Carreira, J., Krauth, K., Yadwadkar, N., Gonzalez, J. E., Popa, R. A., Stoica, I., & Patterson, D. A. (2019). Cloud programming simplified: A Berkeley view on serverless computing. *arXiv preprint*, 56–80.
10. Masanet, E., Shehabi, A., Lei, N., Smith, S., & Koomey, J. (2020). Recalibrating global data center energy-use estimates. *Science*, 367(6481), 984–986.
11. Menghani, G. (2023). Efficient deep learning: A survey on making deep learning models smaller, faster, and better. *ACM Computing Surveys*, 55(12), 1–37.
12. Murugesan, S. (2008). Harnessing green IT: Principles and practices. *IT Professional*, 10(1), 24–33.
13. Mytton, D. (2021). Data centre water consumption. *npj Clean Water*, 4, 11.
14. Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L.-M., Rothchild, D., So, D., Texier, M., & Dean, J. (2021). Carbon emissions and large neural network training. *arXiv preprint*, 123–130.
15. Radovanović, A., Koningstein, R., Schneider, I., Chen, B., Duarte, A., Roy, B., Xiao, D., Haridasan, M., Hung, P., Care, N., Talukdar, S., Mullen, E., Smith, K., Cottman, M., & Cirne, W. (2023). Carbon-aware computing for datacenters. *IEEE Transactions on Power Systems*, 38(2), 1270–1280.
16. Rolnick, D., Donti, P. L., Kaack, L. H., Kochanski, K., Lacoste, A., Sankaran, K., et al. (2022). Tackling climate change with machine learning. *ACM Computing Surveys*, 55(2), 1–96.
17. Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green AI. *Communications of the ACM*, 63(12), 54–63.
18. Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637–646.
19. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3645–3650.
20. Verdecchia, R., Sallou, J., & Cruz, L. (2023). A systematic review of Green AI. *WIREs Data Mining and Knowledge Discovery*, 13(4), 242–251.
21. Wu, C.-J., Raghavendra, R., Gupta, U., Acun, B., Ardalani, N., Maeng, K., et al. (2022). Sustainable AI: Environmental implications, challenges and opportunities. *Proceedings of Machine Learning and Systems*, 4, 795–813.

SUSTAINABLE AGENTIC AI FOR RESPIRATORY TRIAGE IN LOW-RESOURCE HEALTHCARE

Jyothi K. and Pavan M.*

Department of Information Science and Engineering,
Jawaharlal Nehru New College of Engineering (JNNCE),
Affiliated to Visvesvaraya Technological University (VTU),
Belagavi, Shivamogga, Karnataka, India

*Corresponding author E-mail: pmadhvesh@jnnce.ac.in

Abstract

The equitable delivery of healthcare in low-resource settings is severely hindered by a shortage of radiologists. Highly infectious respiratory illnesses, including pneumonia and COVID-19, require rapid screening to curb transmission, yet traditional AI models remain computationally prohibitive for rural deployment. This chapter proposes a sustainable, Agentic AI framework designed as a proactive "cognitive ally" for non-specialist workers triaging chest X-rays. Unlike cloud-dependent models, our diagnostic agent leverages a parameter-efficient DenseNet-121 architecture. This enables carbon-aware, local execution directly on standard mobile devices and clinical tablets, achieving an inference latency under one second. The methodology employs targeted transfer learning and cost-sensitive loss functions to accurately identify pathological opacities. Evaluated on an independent test set, the framework achieved 94% sensitivity for pneumonia detection and an AUC of 0.9457. By providing real-time, explainable confidence score overlays, the system serves as an autonomous first-pass screener that reduces clinician workload while maintaining mandatory human-in-the-loop oversight, establishing an inclusive, mobile-ready standard of care.

Keywords: Agentic AI, Global Health Equity, DenseNet-121, Sustainable Healthcare, Edge-AI, Automated Respiratory Triage, Explainable AI (XAI), Low-Resource Settings.

Introduction

The equitable delivery of global healthcare is severely hindered by a critical shortage of radiological expertise, particularly in low-resource and rural clinical settings (Kiyasseh *et al.*, 2022). Infectious respiratory illnesses, most notably pneumonia and COVID-19, are characterized by their rapid transmission rates and remain leading causes of global morbidity. This rapid spread necessitates continuous, rapid, and accessible screening mechanisms to prevent unchecked localized outbreaks (Ranganathan *et al.*, 2025; Sidume *et al.*, 2024). As reported in the WHO Global Tuberculosis Report 2023, approximately 10.6 million people fell ill with TB in 2022 (World Health Organization, 2023). As illustrated by the regional trend data in Figure 1, the burden of respiratory infections disproportionately afflicts low- and middle-income countries, with 87% of new cases concentrated in 30 high-burden nations (World Health Organization, 2023). In these environments, chest radiography (CXR) serves as the primary

diagnostic modality; however, the lack of specialized personnel to interpret these scans frequently leads to diagnostic delays and compromised patient outcomes (Wuni *et al.*, 2021).

Estimated TB incidence rates, 2022




Figure 1: Trends in estimated Tuberculosis incidence rates (new cases per 100,000 population) by WHO region, 2010–2022. The data underscores the disproportionate burden in regions such as South-East Asia and Africa, highlighting the urgent need for localized, autonomous triage tools (World Health Organization, 2023).

In recent years, the integration of Artificial Intelligence (AI) and Deep Learning (DL) has demonstrated transformative potential in automating medical image analysis (Shaukat *et al.*, 2026). Convolutional Neural Networks (CNNs) can now achieve diagnostic accuracies comparable to expert radiologists. Despite these advancements, the practical deployment of AI in the high-burden regions highlighted in Figure 1 faces massive infrastructural bottlenecks. Traditional AI architectures are heavily reliant on continuous cloud computing, requiring high internet bandwidth rarely available in rural clinics (Ranganathan *et al.*, 2025). To illustrate the diagnostic challenge, Figure 2 presents representative samples from the dataset used in this study, showcasing the visual complexity of pathological lung opacities compared to healthy parenchyma.

Furthermore, the rapid proliferation of computationally intensive AI models has raised urgent concerns regarding their environmental sustainability and carbon footprint (Doo *et al.*, 2024). The energy consumption required to deploy heavy deep learning architectures constitutes a "double-edged sword," where clinical benefits are offset by ecological impact (Champendal *et al.*, 2026). Consequently, there is an urgent paradigm shift toward Sustainable Edge-AI—frameworks designed to execute diagnostics locally on standard, resource-constrained hardware while minimizing carbon emissions (Shaukat *et al.*, 2026).




Figure 2: Representative samples from the chest radiography dataset used for model training and validation (Kermany *et al.*, 2018). The top row illustrates healthy lung parenchyma (Normal), while the bottom row demonstrates the pathological opacities characteristic of pneumonia infections.

To bridge the gap between high-tech diagnostic capabilities and low-resource realities, this chapter proposes a sustainable, autonomous diagnostic agent for respiratory triage. Utilizing a parameter-efficient DenseNet-121 architecture, the proposed Clinical Decision Support System (CDSS) is engineered specifically for localized Edge-AI deployment. A primary motivation for this architectural choice is its exceptionally small computational footprint, which enables the solution to be executed directly on standard mobile devices and low-cost clinical tablets without internet reliance. By implementing targeted transfer learning and cost-sensitive loss functions, the model achieves 94% sensitivity for pneumonia detection with an inference latency of under one second. Crucially, the system functions as an explainable collaborative agent, overlaying confidence scores directly onto radiographs to guide, rather than replace, the frontline healthcare professional (Karunanayake, 2025).

To address the identified gaps in sustainable diagnostic deployment, the specific contributions of this chapter are threefold. First, we develop a parameter-efficient diagnostic agent utilizing a DenseNet-121 backbone, enabling continuous, carbon-aware execution directly on standard, low-resource mobile clinical devices. Second, we implement a cost-sensitive, classweighted transfer learning methodology that strictly mitigates epidemiological dataset imbalance, thereby achieving robust diagnostic sensitivity for pathological opacities without cloud reliance. Finally, we integrate dynamic, human-in-the-loop visual confidence overlays that translate black-box statistical reasoning into actionable clinical insights, ensuring safe, transparent, and accountable collaborative decision-support.

Related Work

The integration of Agentic Artificial Intelligence into clinical workflows represents a fundamental shift from passive predictive models to proactive, autonomous decision-support systems (Dietrich, 2025). However, this transition introduces complex challenges regarding medical liability, infrastructural demands, and clinical trust.

Autonomy, Liability, and Human-in-the-Loop Oversight

As Agentic AI systems evolve into sophisticated "cognitive allies," fully autonomous deployment remains legally and ethically contentious due to the "black box" nature of deep learning algorithms and the difficulty of attributing liability (Eldakak *et al.*, 2024). Recent taxonomies of healthcare agents emphasize that high-stakes medical decision-making renders standard autonomous frameworks inadequate, necessitating systems that prioritize trust and strict alignment with clinical constraints (Vatsal *et al.*, 2026). A primary vulnerability identified in contemporary evaluations is the lack of integrated safety guardrails; empirical analyses reveal that approximately 86% of medical agents fail to implement robust "Human-in-the-Loop" (HITL) mechanisms, risking unchecked autonomous execution (Vatsal *et al.*, 2026).

Despite these challenges, Agentic AI demonstrates a profound ability to augment human expertise. Studies indicate that AI systems analyzing chest radiographs for pneumonia can achieve diagnostic accuracies exceeding 96%, significantly reducing practitioner cognitive load by automating routine triage (AlZu'bi *et al.*, 2025; Alzoubi *et al.*, 2025). Consequently, the prevailing consensus in recent literature advocates for "adaptive triage" mechanisms—agents that quantify uncertainty and present diagnostic predictions as confidence-weighted recommendations rather than definitive autonomous actions (Vatsal *et al.*, 2026; Alzoubi *et al.*, 2025). By automating initial triage while requiring human verification, these systems mitigate algorithmic bias and ensure equitable care standards (Alzoubi *et al.*, 2025; Dietrich, 2025).

Infrastructural Constraints and Edge-Optimized AI

While the theoretical benefits of Agentic AI are well-documented, practical deployment in global healthcare systems is frequently hindered by severe infrastructural limitations. Contemporary architectural evaluations identify high computational overhead and heavy cloudinfrastructure dependencies as the primary weaknesses of modern agentic frameworks (Srinivasu *et al.*, 2026; Karunanayake, 2025). In rural or low-resource clinical environments, the hardware required to orchestrate massive multi-agent models is often unavailable or economically unviable, creating a pressing need for scalable, cost-effective solutions (Karunanayake, 2025).

To democratize access to intelligent diagnostics, recent studies advocate for a transition toward sustainable, lightweight architectures. For instance, AlZu'bi *et al.* (2025) demonstrated that adaptive Agentic AI frameworks can achieve high diagnostic confidence scores (96.5%) while outperforming traditional classifiers by dynamically adjusting to patient-specific imaging parameters. Addressing the global hardware gap requires aligning standard agentic operational

phases—perception, reasoning, and action—with computationally efficient models capable of localized Edge execution (Srinivasu *et al.*, 2026).

Contributions of the Proposed Framework

Building upon the identified gaps in the current literature, the methodology proposed in this chapter introduces a sustainable, collaborative Agentic AI framework built on a DenseNet-121 backbone. Unlike resource-heavy cloud models, this lightweight diagnostic agent is explicitly optimized for Edge-AI deployment on standard clinical hardware (Karunanayake, 2025; Srinivasu *et al.*, 2026).

To navigate the liability risks of full autonomy (Eldakak *et al.*, 2024), the system employs a routing Master Controller that functions strictly as a diagnostic assistant (Dietrich, 2025). By generating transparent visual confidence scores and mandating human-in-the-loop verification (Vatsal *et al.*, 2026), the proposed architecture ensures that predictive efficiency accelerates respiratory triage without ever superseding clinical safety or practitioner authority.

Proposed Agentic AI Methodology

To address the computational constraints of low-resource clinical settings, the proposed diagnostic agent is built upon a sustainable Edge-AI framework utilizing PyTorch. The methodology emphasizes computational efficiency, robust handling of imbalanced medical datasets, and human-in-the-loop explainability, avoiding the heavy infrastructural demands of continuous cloud processing (Ranganathan *et al.*, 2025; Shaukat *et al.*, 2026).

An overarching visual summary of the proposed framework, from initial data ingestion to the final human-in-the-loop clinical triage, is illustrated in Figure 3. The pipeline is strictly optimized for low-latency, localized execution.




Figure 3: System architecture of the proposed Agentic AI framework. The pipeline highlights the transition from raw radiograph ingestion through the Edge-optimized DenseNet-121 reasoning module, culminating in an explainable visual overlay for human-in-the-loop verification

Data Acquisition and Augmentation

The agentic pipeline initiates by automatically ingesting chest radiography (CXR) data. Medical imaging datasets frequently suffer from limited sample sizes, which can lead to rapid model overfitting (Sidume *et al.*, 2024). To enforce robust feature learning without requiring external synthetic data generation, the system applies dynamic spatial transformations during the data loading phase.

The training transformations include dynamic resizing to 224×224 pixels, randomized horizontal flipping, and randomized spatial rotations of up to 8 degrees. Following augmentation, the

tensors are normalized using standard ImageNet channel statistics (Means: [0.485,0.456,0.406], Standard Deviations: [0.229,0.224,0.225]). This standardization ensures the input distribution matches the feature space expected by the pre-trained neural network, optimizing the convergence rate.

Algorithmic Handling of Class Imbalance

In epidemiological datasets, normal baseline scans typically outnumber pathological cases, causing standard deep learning classifiers to develop a heavy bias toward the majority class. Rather than artificially oversampling the data, the proposed framework utilizes a mathematically rigorous cost-sensitive learning approach.

The agent calculates balanced class weights (Equation 1) inversely proportional to class frequencies:

$$w_c = \frac{N}{C \times n_c} \quad (1)$$

where w_c is the weight for class c , N is the total number of samples, C is the total number of classes, and n_c is the number of samples in class c . These weights are directly integrated into a Weighted Cross-Entropy Loss function to heavily penalize misclassifications of the minority (pathological) class.

Edge-Optimized Architecture: DenseNet-121

To satisfy the requirements of carbon-aware computing (Shaukat *et al.*, 2026), the core reasoning module is powered by the DenseNet-121 architecture. DenseNet-121 connects each layer to every other layer in a feed-forward fashion, maximizing feature reuse and significantly reducing the vanishing gradient problem. It allows the model to achieve high diagnostic accuracy with a fraction of the parameters required by traditional architectures like VGG16, making it ideal for localized, Edge-AI deployment.

Two-Stage Transfer Learning Protocol

A two-stage transfer learning protocol is employed to optimize the agent:

- **Feature Extraction:** The foundational convolutional blocks of the ImageNet-trained DenseNet-121 are frozen. This preserves the model's ability to recognize fundamental visual textures, such as edges and gradients, without exhausting local computational resources (Shaukat *et al.*, 2026).
- **Custom Classification Head:** The native classifier is replaced with a regularized, task-specific diagnostic head consisting of a 512-unit fully connected layer, ReLU activation, and a Dropout layer (0.4) to mitigate overfitting.

Agentic Inference and Explainable Action

The defining characteristic of an Agentic AI system is its ability to translate reasoning into actionable, context-aware insights (Dietrich, 2025). Upon processing an unseen radiograph, the inference module calculates a probabilistic confidence score using the Softmax function.

To fulfill the mandate for Explainable AI (XAI) and human-in-the-loop oversight (Karunanayake, 2025), the agent translates this statistical output into a visual format. It

dynamically clones the raw input image and utilizes the Python Imaging Library (PIL) to visually overlay the predicted classification and exact confidence percentage directly onto the scan. This visual feedback mechanism, illustrated empirically in Section 4, transitions the model from a "black box" algorithm to a collaborative clinical assistant, providing the reviewing radiologist with transparent triage data for rapid decision-making.

Results and Analysis

The performance of the proposed DenseNet-121 Agentic AI system was evaluated using an independent test set of 624 samples. To assess the system's viability for sustainable, lowresource triage, the evaluation prioritizes metrics that balance diagnostic safety (Recall) with operational efficiency (Inference Latency).

Dataset Description and Experimental Setup

To validate the proposed Agentic AI framework, this study utilized a rigorously curated dataset of pediatric chest X-ray (CXR) images originally published by Kermany et al. (2018). The dataset comprises 5,856 anterior-posterior chest radiographs collected from retrospective cohorts of pediatric patients, aged one to five years, treated at the Guangzhou Women and Children's Medical Center in Guangzhou, China. All imaging was acquired as part of the patients' routine clinical care. To ensure ground-truth diagnostic accuracy, every radiograph was subjected to a rigorous quality control and grading process. Initial clinical screening removed low-quality or unreadable scans. Subsequently, the diagnoses were independently graded by two expert physicians, with any evaluation discrepancies arbitrated by a third senior expert to establish a highly reliable gold standard for model training.

Regarding the main data components, the dataset is strictly categorized into two primary diagnostic classes: pneumonia and normal. The pneumonia class includes 4,273 images exhibiting pathological lung opacities, encompassing both focal lobar consolidations characteristic of bacterial infections and diffuse interstitial patterns typical of viral infections. Conversely, the normal class consists of 1,583 healthy radiographs demonstrating clear lungs with no areas of abnormal opacification. Reflecting real-world epidemiological distributions, this inherent structural imbalance explicitly necessitates the use of the class-weighted transfer learning methodology detailed in our proposed framework to prevent the model from biasing toward the majority disease class.

Implementation Details

The framework was implemented using PyTorch 2.x and executed within a CUDA-enabled Google Colab environment. Optimization utilized the Adam optimizer with a learning rate of 1×10^{-4} and a batch size of 16. The training cycle for 10 epochs was completed in approximately 25 minutes, underscoring the rapid deployment capabilities essential for Edge-AI architectures in resource-constrained settings.

Training and Validation Performance

The integration of class-weighted loss functions allowed the agent to converge rapidly despite the inherent class imbalance. As illustrated by the convergence trends in Figure 4, the model

achieved peak validation accuracy by the fifth epoch. The stability of the loss curves indicates that the two-stage transfer learning protocol successfully mitigated overfitting while preserving foundational visual features.




Figure 4: Training and Validation performance demonstrating rapid convergence and stability: (a) Model Accuracy reaching peak validation by Epoch 5; (b) Model Loss confirming the mitigation of overfitting

Feature Discrimination and Classification Accuracy

To evaluate the internal representation learning of the agent, we conducted a latent feature space analysis alongside traditional error metrics. High-dimensional feature vectors (1024 dimensions) were extracted from the final Global Average Pooling layer and visualized via t-SNE.




Figure 5: Empirical Validation of Feature Separability and Accuracy:

- (a) Clear spatial clustering of Normal vs. Pneumonia features;
- (b) Predictive distribution confirming high diagnostic recall

As illustrated in Figure 5a, the t-SNE visualization reveals two distinct, cohesive clusters with minimal overlap, proving that the DenseNet-121 backbone successfully extracted task-specific radiological markers. This clear separability in the feature space directly translates to the high diagnostic precision shown in Figure 5b, where the model achieved a Pneumonia Recall of 0.94.

This high sensitivity is critical for ensuring that pathological cases are successfully flagged for clinician review.

System Performance and Interpretability

The relationship between statistical reliability and practical clinical application is further validated through Receiver Operating Characteristic (ROC) analysis and agentic inference overlays.




Figure 6: Combined Performance Metrics: (a) ROC curve illustrating a robust AUC of 0.9457; (b) Practical application showing the real-time AI-generated confidence overlay.

As depicted in Figure 6a, the model achieved an AUC of 0.9457, indicating superior class discrimination across varying thresholds. This statistical precision is translated into an actionable tool in Figure 6b, where the agent localizes its reasoning into a transparent probability overlay (98.19%). This feedback mechanism transitions the model from a passive algorithm to a collaborative clinical assistant. The final quantitative results are summarized in Table 1.

Table 1: Summary of final diagnostic performance metrics on the independent test set

Metric	Achieved Value
Area Under Curve (AUC)	0.9457
Pneumonia Sensitivity (Recall)	0.94
Pneumonia F1-Score	0.90
Overall Test Accuracy	0.87
Inference Latency (Edge Device)	< 1.0 second

Comparative Analysis with State-of-the-Art

To contextualize the performance of the proposed Agentic AI framework, we compared our Edge-optimized DenseNet-121 model against several state-of-the-art deep learning architectures evaluated on the same pediatric chest radiography dataset. As summarized in Table 2, while complex ensemble models and heavy network architectures frequently achieve marginal gains in absolute accuracy, they do so at the severe cost of massive computational overhead, rendering them unsuitable for localized, offline clinical deployment.

For instance, while weighted ensemble architectures (Kundu *et al.*, 2021) report high accuracies exceeding 98%, they require the simultaneous execution of multiple heavy base learners (e.g., GoogLeNet, ResNet-18, and DenseNet), consuming significant memory and battery power. Similarly, early benchmark models utilizing deep Inception V3 architectures (Kermany *et al.*, 2018) rely heavily on cloud infrastructure. While custom CNN architectures built from scratch (Stephen *et al.*, 2019) offer potential Edge viability, they frequently omit comprehensive sensitivity reporting, which is a crucial metric for ensuring patient safety during clinical triage.

In contrast, our proposed DenseNet-121 framework maintains a highly competitive pneumonia sensitivity (Recall) of 94.0% and an AUC of 0.9457 using a fraction of the trainable parameters. This highly compact architecture not only minimizes memory usage but also ensures exceptional time efficiency, allowing for near-instantaneous diagnostic inference. By prioritizing parameter-efficiency and cost-sensitive loss weighting over brute-force computation, the proposed model successfully balances diagnostic rigor with the strict infrastructural constraints of Edge-AI healthcare applications.

Table 2: Comparative analysis of the proposed framework against state-of-the-art models on the pediatric pneumonia dataset

Author (Year)	Architecture	Accuracy (%)	Sensitivity (%)	Edge-Viable
Kermany <i>et al.</i> (2018) (Kermany <i>et al.</i> , 2018)	Inception V3	92.80	93.20	No
Stephen <i>et al.</i> (2019) (Stephen <i>et al.</i> , 2019)	Custom CNN	93.73	N/A	Yes
Kundu <i>et al.</i> (2021) (Kundu <i>et al.</i> , 2021)	Ensemble (GoogLeNet+ResNet+DenseNet)	98.81	98.80	No
Proposed Framework	Agentic DenseNet-121 (Class-Weighted)	87.00	94.00	Yes

Cost-Effectiveness, Sustainability, and Operational Efficiency

The practical deployment of artificial intelligence in low-resource healthcare settings is frequently bottlenecked by the prohibitive operational costs, latency issues, and high energy consumption associated with traditional cloud-centric architectures. Conventional deep learning diagnostic systems necessitate continuous, expensive subscriptions to high-performance cloud computing infrastructure and depend entirely on stable, high-bandwidth internet connectivity.

The proposed Agentic AI framework directly addresses these economic and environmental constraints by leveraging the inherent parameter efficiency of the DenseNet-121 backbone. With a highly compact architectural footprint (approximately 31 MB), the model is computationally lean enough to be deployed directly on standard commercial off-the-shelf (COTS) hardware, such as low-cost clinical tablets and mobile devices. This decentralized, Edge-AI approach drastically reduces the total cost of ownership and operation, validating the framework as a

sustainable, carbon-aware solution that makes advanced diagnostic support financially viable for underfunded healthcare facilities.

Furthermore, this localized deployment strategy significantly enhances operational efficiency by optimizing diagnostic inference time. In cloud-dependent systems, inference latency is dictated by network transmission speeds, frequently leading to unacceptable delays in regions with poor connectivity. By executing the classification locally, the proposed model completely bypasses network round-trip delays. During testing, the model demonstrated a negligible inference latency of less than one second per radiograph on local hardware. This near real-time processing capability prevents bottlenecks in high-volume patient triage scenarios, enabling frontline healthcare workers to make rapid, accountable clinical decisions directly at the point of care.

Explainable AI (XAI) and Key Clinical Findings

To bridge the trust deficit frequently associated with black-box deep learning models, this framework integrates Explainable AI (XAI) techniques as a core component of the agentic triage system. Specifically, Gradient-weighted Class Activation Mapping (Grad-CAM) is utilized to dynamically visualize the decision-making process of the DenseNet-121 model. By computing the gradients of the target diagnostic class flowing into the final convolutional layer, the system generates coarse localization heatmaps that highlight the specific spatial regions within the radiograph that most strongly influenced the model's prediction. These visual confidence overlays transform the diagnostic agent from an opaque statistical classifier into a transparent, collaborative decision-support tool.

The application of XAI within this framework yielded three key findings regarding the model's diagnostic behavior. First, the generated activation maps demonstrated strong morphological alignment with human expert analysis, consistently highlighting true pathological features such as focal lobar consolidations and diffuse interstitial opacities when predicting pneumonia cases. This is illustrated in Figure 7, where the high-confidence (99.10%) heatmap corresponds directly to relevant lung parenchyma rather than background artifacts. Second, the visual explanations served as a critical validation mechanism to ensure the model was not relying on spurious correlations or dataset biases; the heatmaps confirmed that the network successfully ignored irrelevant artifacts such as medical tubing, embedded text markers, and anatomical structures outside the lung fields. Finally, the integration of these explainable overlays proved essential for the human-in-the-loop deployment strategy. By providing visual localization alongside statistical confidence scores, the Agentic AI system enables frontline healthcare workers to quickly verify the model's reasoning, thereby fostering clinical trust and ensuring safe, accountable triage decisions in low-resource environments.

Conclusion

This chapter presented a sustainable, Edge-based Agentic AI framework for automated respiratory disease triage, addressing the critical constraints of low-resource clinical settings by eliminating prohibitive cloud infrastructure dependencies. By integrating a parameter-efficient DenseNet-121 architecture with a cost-sensitive learning paradigm, the system achieves high

diagnostic sensitivity for pathologies like pneumonia while maintaining a carbon-aware, localized computational footprint. Crucially, the framework functions as a collaborative "cognitive ally" rather than an autonomous replacement; by providing transparent, real-time confidence scores, it ensures mandatory human-in-the-loop oversight, addresses clinical liability concerns, and significantly reduces the cognitive burden on frontline practitioners. While future research will focus on expanding this modular architecture to include multi-modal data fusion—such as integrating clinical notes with imaging—the current framework already demonstrates a viable pathway for democratizing advanced radiological support. Ultimately, the transition to such explainable and lightweight diagnostic agents represents a transformative milestone in global health, ensuring that life-saving AI tools remain inclusive, economically viable, and environmentally responsible.




Figure 7: Grad-CAM visualization of a pneumonia case demonstrating high-confidence (99.10%) localization of pathological lung opacities, confirming the model's focus aligns with clinically significant anatomical regions.

References

1. Alzoubi, Y. I., Topcu, A. E., and Elbasi, E. (2025). A systematic review and evaluation of sustainable ai algorithms and techniques in healthcare. *IEEE Access*, 13:1–25.
2. AlZu'bi, S., Almiani, M., and Jararweh, Y. (2025). Agentic ai for healthcare: Solutions to intelligent patient care. In *2025 12th International Conference on Information Technology (ICIT)*, pages 33–38. IEEE.
3. Champendal, M., Lokaj, B., de Gevigney, V. D., Brulé, G., Zaghir, J., Lovis, C., Müller, H., Schmid, J., and Ribeiro, R. T. (2026). Exploring environmental sustainability of artificial intelligence in radiology: A scoping review. *European Journal of Radiology*, 194:112558.
4. Dietrich, N. (2025). Agentic ai in radiology: emerging potential and unresolved challenges. *British Journal of Radiology*, 98:1582–1584.
5. Doo, F. X., Vossenrich, J., Cook, T. S., Moy, L., Almeida, E. P. R. P., Woolen, S. A., Gichoya, J. W., Heye, T., and Hanneman, K. (2024). Environmental sustainability and ai in radiology:

7. A double-edged sword. *Radiology*, 310(1):e231968.
8. Eldakak, A., Alremeithi, A., Dahiyat, E., El-Gheriani, M., Mohamed, H., and Abdulla, M. I. A. (2024). Civil liability for the actions of autonomous ai in healthcare: an invitation to further contemplation. *Humanities and Social Sciences Communications*, 11(1):305.
9. Karunanayake, N. (2025). Next-generation agentic ai for transforming healthcare. *Informatics and Health*, 2:73–83.
10. Kermany, D. S., Goldbaum, M., Cai, W., Valentim, C. C. S., Liang, H., Baxter, S. L., McKeown, A., Yang, G., Wu, X., Yan, F., *et al.* (2018). Identifying medical diagnoses and treatable diseases by image-based deep learning. *Cell*, 172(5):1122–1131.
11. Kiyasseh, D., Zhu, T., and Clifton, D. (2022). The promise of clinical decision support systems targetting low-resource settings. *IEEE Reviews in Biomedical Engineering*, 15:354–368.
12. Kundu, R., Das, R., Geem, Z. W., Han, G.-T., and Sarkar, R. (2021). Pneumonia detection in chest x-ray images using an ensemble of deep learning models. *PLoS ONE*, 16(9):e0256630.
13. Ranganathan, R., Mythili, S., Annamalai, A., and Shuriya, B. (2025). Sustainable ai for medical imaging: A federated, edge-based approach to chest x-ray triage. In *2025 International Conference on Modern Sustainable Systems (CMSS)*. IEEE.
14. Shaukat, F., Parwez, K., Ashraf, Z., Alhabeeb, A., Alotaibi, F. A., and Alnfiai, M. M. (2026). Mobilenet-lite-health: A sustainable edge ai framework for medical image classification and carbon-aware computing. *IEEE Access*, 14:29319–29330.
15. Sidume, F., Muchuchuti, S., Ntwaetsile, K., Chengetenai, G., Ntekeng, P., and Tamukate, R. (2024). Deep learning techniques for tuberculosis detection on chest x-rays in low-resource settings: A survey of opportunities and challenges. In *2024 International Conference on Electrical and Computer Engineering Researches (ICECER)*. IEEE.
16. Srinivasu, P. N., Kumari, G. L. A., Ahmed, S., and Alhumam, A. (2026). Exploring agentic ai in healthcare: A study on its working mechanism. *Frontiers in Medicine*, 12:1753443.
17. Stephen, O., Sain, M., Maduh, U. J., and Jeong, D. U. (2019). An efficient deep learning approach to pneumonia classification in healthcare. *Journal of Healthcare Engineering*, 2019:4180949.
18. Vatsal, S., Dubey, H., and Singh, A. (2026). Agentic ai in healthcare and medicine: A sevendimensional taxonomy for empirical evaluation of llm-based agents. *IEEE Access*, 14:4840–4863.
19. World Health Organization (2023). Global tuberculosis report 2023. Technical report, World Health Organization, Geneva, Switzerland.
20. Wuni, A.-R., Botwe, B. O., and Akudjedu, T. N. (2021). Impact of artificial intelligence on clinical radiography practice: futuristic prospects in a low resource setting. *Radiography*, 27(4):1234–1240.

DESIGN AND IMPLEMENTATION OF A SMART AUTOMATIC EMERGENCY BRAKING SYSTEM FOR ELECTRIC VEHICLES USING ULTRASONIC SENSORS

C. Jayabalan

Department of Mechanical Engineering, AMET Deemed to be University, Chennai, India.

Corresponding author [E-mail: jayabalan.mech@ametuniv.ac.in](mailto:jayabalan.mech@ametuniv.ac.in)

Abstract

The increasing adoption of electric vehicles (EVs) has created a growing need for intelligent safety systems that can reduce the risk of collisions and improve passenger protection. This study presents a Smart Ultrasonic Sensor-Based Automatic Braking System designed to detect obstacles in real time and automatically activate the braking mechanism when a potential collision is identified. Ultrasonic sensors are used to continuously measure the distance between the vehicle and surrounding obstacles. When an obstacle is detected within a predefined safety range, the control unit processes the sensor information and initiates automatic braking, thereby reducing the vehicle speed and minimizing the possibility of an accident. The proposed system can provide rapid obstacle detection, improve driver response time, and enhance overall vehicle safety. The system is designed to be cost-effective, compact, and suitable for integration into low-speed electric vehicles and electric mobility platforms. Experimental evaluation can be carried out under different obstacle distances and vehicle speeds to assess detection accuracy, braking response time, and system reliability. The proposed approach demonstrates the potential of ultrasonic sensing and automated control to improve collision avoidance and contribute to safer electric vehicle operation.

Keywords: Electric Vehicle, Ultrasonic Sensor, Automatic Braking, Obstacle Detection, Collision Avoidance, Vehicle Safety, Intelligent Transportation System.

1. Introduction

The rapid development of electric vehicles (EVs) has created new opportunities for sustainable and efficient transportation. Electric vehicles offer several advantages, including reduced greenhouse gas emissions, lower operating costs, and improved energy efficiency. However, ensuring adequate vehicle safety remains an important concern, particularly with the increasing use of EVs in urban environments where vehicles frequently encounter pedestrians, other vehicles, and stationary obstacles [1].

Advanced driver assistance systems (ADAS) have become an important area of research for improving road safety. These systems use sensors and intelligent control techniques to detect potential hazards and assist the driver in avoiding collisions. Automatic emergency braking (AEB) is one such safety technology that can detect obstacles and automatically apply the brakes when a collision is considered imminent [2]. The implementation of automatic braking can reduce driver reaction time and potentially decrease the severity of collisions.

Ultrasonic sensors are widely used for short-range obstacle detection because they are relatively inexpensive, compact, and capable of measuring distance using reflected sound waves. An ultrasonic sensor transmits a high-frequency sound pulse and determines the distance to an obstacle by measuring the time taken for the reflected signal to return to the sensor [3]. These characteristics make ultrasonic sensors suitable for low-speed electric vehicles and prototype safety systems where cost and simplicity are important considerations.


Several studies have investigated sensor-based collision avoidance and automatic braking systems. Sensor fusion approaches combining ultrasonic, infrared, radar, and camera-based technologies have been explored to improve obstacle detection and vehicle safety [4]. However, complex sensing systems can increase system cost and computational requirements. A properly designed ultrasonic sensor-based system can provide a simpler and economical alternative for short-range obstacle detection, particularly in low-speed EV applications.

The effectiveness of an automatic braking system depends on accurate distance measurement, appropriate safety thresholds, rapid signal processing, and reliable braking control. When an obstacle enters the predefined critical distance, the controller must respond quickly by activating the braking mechanism. The safety distance may be determined based on vehicle speed, sensor characteristics, braking response time, and stopping distance [5]. Therefore, integrating ultrasonic sensing with an automatic braking mechanism requires careful consideration of both sensing accuracy and control response.

With the increasing demand for affordable and intelligent EV safety technologies, there is a need for compact systems capable of providing real-time obstacle detection and automatic braking. The present study therefore focuses on the design and development of a smart ultrasonic sensor-based automatic braking system for enhanced electric vehicle safety. The proposed system continuously monitors the distance between the vehicle and nearby obstacles and automatically activates the braking mechanism when an obstacle is detected within a predefined safety zone. The system aims to improve collision avoidance, reduce driver reaction dependency, and provide a cost-effective safety solution for electric vehicles.

2. Objectives of the Study

- To design and develop a smart ultrasonic sensor-based automatic braking system for electric vehicles.
- To detect obstacles in real time using ultrasonic sensors.
- To measure the distance between the electric vehicle and nearby obstacles accurately.
- To automatically activate the braking mechanism when an obstacle reaches a predefined safety distance.
- To reduce the driver's reaction time and minimize the possibility of collisions.
- To evaluate the performance of the proposed system under different obstacle distances and vehicle speeds.



3. Data and Methodology

3.1 Data Collection

The study uses experimental data obtained from the developed ultrasonic sensor-based automatic braking prototype. Experiments are conducted by placing obstacles at different distances from the vehicle and testing the system under different operating conditions.

The following parameters are recorded during the experiments:

- Distance between the vehicle and obstacle (cm)
- Vehicle speed (km/h)
- Ultrasonic sensor detection time (ms)
- Braking activation time (ms)
- Stopping distance (cm)
- Collision avoidance success/failure
- Sensor detection accuracy (%)

3.2 Hardware Components

The proposed system consists of:

- Ultrasonic sensor (HC-SR04)
- Microcontroller/Arduino
- Motor driver
- DC motor/electric vehicle drive system
- Electronic braking mechanism
- Battery/power supply
- Buzzer or warning indicator
- Obstacle/object for testing

3.3 Methodology

The working procedure of the proposed system is as follows:

Obstacle Detection → Distance Measurement → Data Processing → Safety Distance Comparison → Automatic Brake Activation → Vehicle Stopping

The ultrasonic sensor continuously transmits ultrasonic waves and receives the reflected signal from an obstacle. The controller calculates the distance using the time taken by the ultrasonic pulse to return to the sensor.

The distance is calculated using:

$$d = \frac{v \times t}{2}$$

where (d) is the distance between the vehicle and obstacle, (v) is the velocity of sound, and (t) is the round-trip time of the ultrasonic signal.

When the measured distance is greater than the predefined safety threshold, the vehicle continues to operate normally. When the measured distance becomes equal to or less than the threshold, the controller sends a signal to the braking mechanism and automatically stops or reduces the vehicle speed.

3.4 Experimental Procedure

The prototype is tested by placing obstacles at different distances, such as 20 cm, 30 cm, 40 cm, 50 cm, and 60 cm. The experiments are repeated at different vehicle speeds to evaluate the response of the automatic braking system.

The performance of the system is evaluated based on obstacle detection accuracy, braking response time, stopping distance, and collision avoidance rate.

3.5 Data Analysis

The experimental results are analysed using descriptive statistical measures such as mean, percentage, standard deviation, and error percentage. Graphs can be used to examine the relationship between vehicle speed, obstacle distance, braking response time, and stopping distance.

The overall performance of the proposed system is assessed by comparing the experimental results under different operating conditions. The methodology helps determine whether the ultrasonic sensor-based automatic braking system can provide reliable and rapid obstacle detection and improve safety in electric vehicles.

Conclusion

The study presents a Smart Ultrasonic Sensor-Based Automatic Braking System for enhancing the safety of electric vehicles. The proposed system uses ultrasonic sensors to continuously detect obstacles and measure the distance between the vehicle and the detected object. When the obstacle enters the predefined safety range, the controller automatically activates the braking mechanism, thereby reducing the possibility of collision.

The experimental methodology provides a practical approach for evaluating obstacle detection accuracy, braking response time, and stopping distance under different operating conditions. The system offers a simple, low-cost, and efficient safety solution that can be particularly useful for low-speed electric vehicles. By reducing dependence on human reaction time, the proposed system can improve collision avoidance and overall vehicle safety.

The study demonstrates the potential of integrating ultrasonic sensing, microcontroller-based processing, and automatic braking control in electric vehicles. Future improvements may include multiple ultrasonic sensors, higher-speed testing, sensor fusion with cameras or radar, adaptive braking based on vehicle speed, and integration with advanced driver-assistance systems (ADAS).

References

1. International Energy Agency (IEA). (2024). *Global EV Outlook 2024: Moving towards increased affordability*. IEA, Paris.
2. International Organization for Standardization. (2021). *ISO 22733:2021—Road vehicles: Test method to evaluate the performance of autonomous emergency braking systems*. ISO.
3. Chan, C. C. (2007). The state of the art of electric, hybrid, and fuel cell vehicles. *Proceedings of the IEEE*, 95(4), 704–718.
4. Yurtsever, E., Lambert, J., Carballo, A., & Takeda, K. (2020). A survey of autonomous driving: Common practices and emerging technologies. *IEEE Access*, 8, 58443–58469.
5. Rajamani, R. (2012). *Vehicle Dynamics and Control* (2nd ed.). Springer.

FROM DIFFERENTIAL EQUATIONS TO SMART GRIDS: A MATHEMATICAL JOURNEY

T. Jagan

United Institute of Technology, Coimbatore-641020

Corresponding author E-mail: jaganthangaraj06@gmail.com

1. Why Does a Power Grid Need Maths?

Think of the electricity grid as one giant, invisible system that has to keep everyone's lights on, phones charging, and fridges running — all at the same time, every second of the day. It sounds like a job for engineers and wires. But actually, the whole thing behaves like a living system that is constantly reacting, adjusting, and settling down after small shocks — not so different from how a person's heart rate spikes and then calms down after a surprise, or how a room's temperature drifts and gets pulled back by a thermostat.

Maths is simply the tool engineers use to describe and predict that behaviour. In this chapter, we will go through five ideas, one at a time, each with a simple everyday example first and a diagram to look at. You do not need an engineering background to follow this — the same patterns show up in psychology, biology, and daily life, and we will point those out along the way.

By the end of this chapter, you should be able to explain, in plain words:

- *Why a system “settles down” after a disturbance, and what makes it settle faster or slower.*
- *Why a system might “swing” back and forth for a while before calming down.*
- *How computers “guess” what happens next in small steps, and why some guesses are better than others.*
- *How feedback (checking and correcting) keeps a system stable.*
- *Why unpredictable things like sunshine and wind make grid planning harder.*
- *Why drawing the grid as a network of connected “people” (nodes) helps us understand it*

2. Settling Down After a Disturbance

2.1 The Phone-Charging Idea

Here is something you already know from daily life: when you plug in your phone at 20% battery, it charges fast at first, then slows down as it gets closer to 100%. It never overshoots — it just glides smoothly up to full and stays there. A voltage on a small section of the electricity grid behaves in exactly the same way after a disturbance: it rises quickly at first, then eases off as it gets close to its final, steady value.

Engineers describe this “fast then slow” rise using a number called the time constant (written as the Greek letter τ , “tau”). A small τ means the system settles down quickly — like a fast phone charger. A large τ means it takes its time — like an old, slow charger. Figure 1 shows exactly this: three curves, each settling at a different speed, all heading toward the same final value.




Figure 1: Three different “settling speeds.” All three curves reach the same final value, but the teal one gets there fastest (small τ) and the orange one is slowest (large τ) — just like a fast charger versus a slow charger

That is really all the maths in this section says: something changes quickly when it is far from where it wants to be, and more slowly as it gets close. Engineers write this as a small equation ($dv/dt = (V_s - v)/RC$), but you do not need to memorise it — the phone-charging picture in your head is doing the same job.

2.2 The Swing-Back-and-Forth Idea

Now imagine a child on a playground swing who gets pushed once and then left alone. The swing does not just glide smoothly to a stop — it goes forward, then back, then forward again, each time a little less far, until it finally comes to rest. This back-and-forth pattern, gradually shrinking, is called damped oscillation.




Figure 2: Like a playground swing after one push: without any “friction” (dashed line) it keeps swinging forever; with friction/damping (solid line) it settles down within a few seconds

A power-generator's rotor — the big spinning part inside a power plant — behaves exactly like that swing after a sudden change in the grid (say, a large factory switching on all at once). It swings back and forth in angle before settling down. If you have ever studied how the body's

arousal level spikes after a stressful event and then oscillates before returning to baseline, it is the same shape of curve. Figure 2 shows this: the dashed line is a swing with no friction (it never stops), and the solid line is a swing with some friction — or in a psychology term, some “self-regulation” — that brings it back to rest.

Why does this matter for a smart grid? Older power plants have big, heavy spinning generators, and that heaviness acts like natural friction — it resists sudden swings. Solar panels and wind turbines do not spin like that, so as more of them are added to the grid, the whole system has less natural “friction” and can swing more wildly after a disturbance. That is one of the real, current challenges engineers are solving today.

3. Guessing What Happens Next, One Small Step at a Time

A real power grid has thousands of generators and connections all interacting, so there is no neat formula that gives the exact answer instantly. Instead, computers estimate the future in tiny time-steps, a bit like predicting someone's mood tomorrow by looking at today's mood and nudging it forward hour by hour, rather than jumping straight to next week.

There are cheaper, rougher ways to do this and more careful, accurate ways. A rough method (called Euler's method) just looks at the current direction of change and takes one big step in that direction — quick, but it can drift off course. A more careful method (called Runge–Kutta, or RK4) checks the direction a few times within each step and averages them before moving — slower to compute, but far more accurate. Figure 3 shows the rough method visibly straying from the true answer, while the careful method stays right on track.




Figure 3: The grey line is the true answer. The orange “rough guess” method drifts away from it. The teal “careful guess” method (checking direction a few times per step) stays right on the curve

Grid engineers pick whichever method fits the situation: a rough-and-fast method when a system is calm and predictable, and a careful, slower method when things are changing quickly and small mistakes could snowball.

4. Feedback: How the Grid Keeps Its Balance

You already know a perfect example of feedback control: a thermostat. It checks the room temperature, compares it to what you set, and turns the heater or air conditioner on or off to close the gap. It keeps checking, again and again, correcting a little each time.

The electricity grid uses the exact same idea to keep its frequency (essentially, how fast the whole system is “ticking”) steady. A controller constantly measures the frequency, compares it to the target, and sends a correction signal to power plants to produce a little more or a little less power. This three-part correction — reacting to the current gap, reacting to how long the gap has lasted, and reacting to how fast the gap is changing — is called PID control (Proportional–Integral–Derivative). Figure 4 draws this out as a loop: measure, correct, measure again, forever.




Figure 4: The same idea as a thermostat: measure the current state, compare it to the target, send a correction, then measure again. This loop runs continuously to keep grid frequency steady

If the correction is too weak, the system drifts and never quite gets back to normal. If the correction is too strong or too delayed, the system can overcorrect and oscillate — imagine overcorrecting a swerving car and fishtailing side to side. Tuning this balance correctly is one of the main jobs of a control engineer.

5. Living With Uncertainty: Sun and Wind Don't Follow a Schedule

Everything discussed so far assumes we know exactly how much power is coming in. But solar and wind power depend on the weather, which is never perfectly predictable — a cloud passes over, the wind gusts and dies down. This is similar to how someone's mood can fluctuate day to day around a general baseline, pulled by daily ups and downs but still tending to drift back toward “average.”

Engineers model this using a random-but-anchored pattern: a value that wobbles randomly but keeps getting pulled back toward a long-term average, with the size of the random wobble and the strength of the “pull-back” as two adjustable settings. This lets engineers run thousands of simulated “what if the weather does this” scenarios and then size batteries and backup power generously enough to cover, say, 95% of the realistic bad-weather scenarios — without wildly over-building expensive equipment for the rare worst case.

6. Drawing the Grid Like a Social Network

In psychology and sociology, you may already be familiar with drawing a social network: people as dots, and lines connecting the people who know each other. A power grid — especially a small local “microgrid” with solar panels, batteries, and homes — can be drawn exactly the same

way: each solar panel, battery, or group of homes is a dot (called a node), and each cable connecting them is a line (called an edge). Figure 5 shows a small example.




Figure 5: A small local power network drawn like a social network diagram: solar and wind (green), a battery (yellow), homes/loads (pink), and a connection to the main grid (grey), all linked through two central hubs

Once the grid is drawn this way, engineers can borrow tools from network science — the same field used to study how information or disease spreads through a social network — to answer two very practical questions: how power actually flows from one point to another, and how vulnerable the network is to a single connection breaking. A network with only one path between two important points is fragile, exactly the way a friendship group with only one connecting person is fragile if that person leaves. Adding a second connection, a backup path, makes the whole network more resilient — whether it is a power grid or a social group.

7. A Simple Worked Example

Picture three points in a row: a solar-and-battery hub (Bus A), a homes hub in the middle (Bus B), and a connection to the main grid (Grid Tie). Bus A is generating more power than it needs and sending the extra out; Bus B's homes are using a little power; and the main grid absorbs whatever is left over to keep everything balanced.

Using the network idea from Section 7.6, we can calculate exactly how much power flows down each connecting cable. Doing the maths (which we will not repeat here in full) shows that the cable between Bus A and Bus B carries more power than the cable between Bus B and the Grid Tie — simply because all of Bus A's extra power has to funnel through that first cable before it can spread out anywhere else. If that first cable is only rated to carry a smaller amount safely, it becomes the weak link in the chain — something engineers would need to fix, either by upgrading that cable or by adding an alternate path.

8. Bringing it All Together

Every idea in this chapter follows a pattern you likely already understand from everyday life or from psychology: systems settle down after a disturbance (the phone charger), some systems

overshoot and oscillate before settling (the swing), computers predict the future step by step (like updating a forecast hour by hour), feedback loops correct small errors continuously (the thermostat), some inputs are unpredictable but anchored around an average (daily mood fluctuation), and networks of connected things can be mapped and studied for their weak points (a social network). The mathematics simply gives these familiar patterns precise, testable shapes so that engineers can design a grid that stays reliable even as more unpredictable, less “steady” sources like solar and wind are added to it.

References

1. Kundur, P. (1994). *Power System Stability and Control*. McGraw-Hill.
2. Boyce, W. E., & DiPrima, R. C. (2017). *Elementary Differential Equations and Boundary Value Problems* (11th ed.). Wiley.
3. Newman, M. E. J. (2010). *Networks: An Introduction*. Oxford University Press.
4. Åström, K. J., & Murray, R. M. (2008). *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press.
5. Uhlenbeck, G. E., & Ornstein, L. S. (1930). On the Theory of the Brownian Motion. *Physical Review*, 36(5), 823–841.

**Engineering Tomorrow:
Intelligent Systems, Sustainable Design and Emerging Technologies
(ISBN: 978-93-47587-83-2)**

About Editors



Neha Bagga is an Assistant Professor in the Department of Computer Science and Engineering at GNA University, Punjab, and is currently pursuing her Ph.D. in Computer Science from Guru Nanak Dev University, Amritsar. She holds an M.Tech. in Computer Science and Engineering and has over 10 years of teaching experience. She is UGC-NET and GATE qualified. Her research interests include Cybersecurity, Cyber-Physical Systems, Vehicular Cyber-Physical Systems, and Intrusion Detection Systems. She has published research papers in international journals and conferences and contributed chapters to books on cybersecurity, cryptography, blockchain, and cyber-physical systems. She has also supervised undergraduate and postgraduate research projects and actively contributes to academic mentoring and research activities. Her professional interests focus on developing innovative, secure, and reliable solutions for emerging computing and cyber-physical technologies and interdisciplinary research, while supporting students through effective teaching, guidance, and professional development.



Prof. Soniya Jain is an Assistant Professor at the Centre for Professional Courses, Gujarat University, Ahmedabad, and is pursuing a Ph.D. in Computer/IT Engineering at Gujarat Technological University. She holds an M.E. in Computer Engineering with distinction and a B.E. in Computer Engineering. With over 17 years of teaching experience, her research interests include Machine Learning, Deep Learning, Data Science, Artificial Intelligence, Computer Vision, and Graph Neural Networks. Her research focuses on EEG-based depression detection using graphical deep learning. She has published research papers, presented at national and international conferences, and received the Best Paper Presentation Award at ICMOTARSS-2026. She is UGC-NET, GSET, and professionally certified, and has contributed to faculty development, workshops, academic events, research, innovation, and technical training. She received a Gold Medal for outstanding Master's performance.



Dr. Barjinder Kaur is a Physics educator and researcher with qualifications including Ph.D., M.Sc. in Physics, and B.Ed. Her academic and research interests encompass Plasma Physics, Semiconductor Physics, Electronic Materials, Semiconductor Devices, Lasers, Optical Fibers, and Superconductivity. She has contributed to the field through publications in reputed journals and academic forums. She has several granted patents to her credit and has published numerous book chapters in her areas of expertise. Dr. Kaur has been actively involved in teaching, research, curriculum development, and academic activities. Her professional achievements include research contributions, academic mentoring, and participation in professional and educational initiatives. She has experience guiding students and developing innovative teaching-learning methodologies in Physics. Her work promotes scientific education, research, and practical applications of Physics, reflecting her continued commitment to teaching, research, publications, academic development.



Dr. Mohd Faizan Hasan is an Associate Professor in the Department of Mechanical Engineering at Integral University, Lucknow, India. He has extensive experience in teaching, research, administration, and accreditation. He holds a Ph.D. in Mechanical Engineering, an M.Tech. in Production and Industrial Engineering, and a B.Tech. in Mechanical Engineering. His research interests include Wire Electrical Discharge Machining (WEDM), advanced machining processes, composite materials, additive manufacturing, Industry 4.0, optimization, and intelligent manufacturing. He has published research papers in reputed journals and conferences and contributed to patents, books, and publications. His academic contributions include research in machining optimization, curriculum development, accreditation activities, and mentoring undergraduate and postgraduate students. His work reflects commitment to innovation, advanced manufacturing, engineering education, and interdisciplinary research.

