A Hindsight–Insight–Foresight Analysis of AI in Mission-Critical Healthcare Information Systems

Ihunanya Udodiri Ajakwe*
IT Convergence Department
Kumoh National Institute of Technology
Gumi, Gyenongsangbuk, Republic of Korea
ihunanya.ajakwe@gmail.com

Simeon Okechukwu Ajakwe†
ICT Convergence Research Centre
Kumoh National Institute of Technology
Gumi, Gyenongsangbuk, Republic of Korea
simeonajlove@gmail.com

Theodore Armand Tagne Poupi‡
Institute of Digital Anti-Aging and Healthcare
Inje University
Gimhae, Republic of Korea
poupiarmand2@gmail.com

Oluleke Babayomi§
ICT Convergence Research Centre
Kumoh National Institute of Technology
Kumoh National Institute of Technology, Gyenongsangbuk
Republic of Korea
oluleke.babayomi@ieee.org

Elvin Ugonna Eziama¶
Department of Mechanical, Automotive, and Materials Engineering
University of Windsor
Windsor, Canada
eziama@uwindsor.ca

Jae Min Lee‖
IT Convergence Engineering Department
Kumoh National Institute of Technology
Gumi, Gyenongsangbuk, Republic of Korea
ljmpaul@kumoh.ac.kr

Dong Seong Kim
IT Convergence Engineering Department
Kumoh National Institute of Technology
Gumi, Gyenongsangbuk, Republic of Korea
dskim@kumoh.ac.kr

Abstract

The convergence of Artificial Intelligence (AI) and Mission-Critical Systems (MCS) has transformed healthcare into an intelligent, interconnected, and adaptive environment. This paper presents a concise analysis of AI integration into mission-critical healthcare information systems through a threefold framework: Hindsight, Insight, and Foresight. Hindsight traces early symbolic and statistical AI systems such as MYCIN and CADIAG, which laid the foundation for clinical decision support. Insight highlights recent advances, including deep learning for perception, reinforcement learning for adaptive decision-making, explainable AI (XAI) for interpretability, and federated learning (FL) for privacy-preserving collaboration. Foresight explores the emergence of Generative and Agentic AI as enablers of real-time, distributed, and trustworthy healthcare operations within 6G-enabled environments. A compact taxonomy is proposed to classify AI approaches in healthcare mission-critical systems based on technique, system layer, and safety function. This synthesis provides a roadmap for developing robust, explainable, and privacy-preserving AI solutions aligned with computing and telecommunications convergence.

*IT Convergence Engineering Department, Kumoh National Institute of Technology, Gumi, South Korea
†ICT Convergence Research Center, Kumoh National Institute of Technology, Gumi, South Korea
‡Institute of Digital Anti-Aging and Healthcare, Inje University, South Korea
§ICT Convergence Research Center, Kumoh National Institute of Technology, Gumi, South Korea
¶Department of Mechanical, Automotive, and Materials Engineering at the University of Windsor
‖IT Convergence Engineering Department, Kumoh National Institute of Technology, Gumi, South Korea

ICACS 2025, Bangkok, Thailand

ACM ISBN 979-8-4007-1961-5/25/12
https://doi.org/10.1145/3789418.3789448

CCS Concepts

• **Information systems**; • **Computing methodologies**; • **Security and privacy**;

Keywords

Artificial Intelligence, Mission-Critical Systems, Healthcare, Explainable AI, Federated Learning, Reinforcement Learning, Generative AI, Edge Intelligence, 6G Communications.

ACM Reference Format:
Ihunanya Udodiri Ajakwe, Simeon Okechukwu Ajakwe, Theodore Armand Tagne Poupi, Oluleke Babayomi, Elvin Ugonna Eziama, Jae Min Lee, and Dong Seong Kim. 2025. A Hindsight–Insight–Foresight Analysis of AI

in Mission-Critical Healthcare Information Systems. In *The 9th International Conference on Algorithms, Computing and Systems (ICACS 2025), December 12–14, 2025, Bangkok, Thailand*. ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3789418.3789448

1 Introduction

Mission-critical healthcare systems represent the backbone of medical operations where failure is not an option, such as intensive care units (ICUs), surgical robotics, and emergency response systems [8]. These systems operate under stringent constraints of latency, reliability, and security, where even minimal disruptions can result in life-threatening consequences. The growing complexity of modern healthcare demands intelligent solutions capable of perceiving, reasoning, and acting autonomously in dynamic environments.

The intersection of healthcare with computing, management, and telecommunications has enabled the emergence of distributed and intelligent healthcare infrastructures. Advances in computing, such as cloud-edge collaboration and federated learning (FL), now enable hospitals to train AI models collaboratively without compromising data privacy [21]. Management frameworks for mission-critical operations integrate predictive analytics and reinforcement learning (RL) to optimize clinical resource allocation and workflow coordination. Meanwhile, telecommunications technologies, particularly 5G and emerging 6G networks, enable ultra-reliable low-latency communication (URLLC) between connected medical devices, remote monitoring systems, and autonomous surgical robots. Together, these pillars establish the foundation for AI-enabled Mission-Critical Healthcare Information Systems (AIMS).

The application of AI within these interconnected healthcare ecosystems extends beyond automation; it enhances decision-making, improves situational awareness, and strengthens system resilience. However, integrating AI across distributed hospital networks also introduces new challenges, including maintaining explainability, ensuring data integrity, and guaranteeing end-to-end security across telemedical communication links.

This paper adopts a structured Hindsight–Insight–Foresight (HIF) framework to examine the evolution of AI in mission-critical healthcare. The Hindsight section 3 revisits the foundations of symbolic and statistical AI systems such as MYCIN and CADIAG; the Insight section 4 explores current paradigms, including deep learning (DL), explainable AI (XAI), federated learning (FL), and reinforcement learning (RL); and the Foresight section 5 envisions future trends such as generative and agentic AI in 6G-connected healthcare environments. This perspective provides a roadmap for developing resilient, distributed, and certifiable AI-driven healthcare systems that align with the goals of computing, management, and telecommunications convergence [8, 16, 21].

2 Review Methodology & Taxonomy

This study employs a structured three-axis taxonomy, Hindsight, Insight, and Foresight (HIF), to systematically analyze the evolution, current state, and emerging trajectory of Artificial Intelligence (AI) in mission-critical healthcare information systems. The approach integrates qualitative synthesis, comparative evaluation, and conceptual modeling across computing, management, and telecommunications domains.

The Hindsight axis examines foundational models such as MYCIN and CADIAG that established interpretability and rule-based inference in medical AI. The Insight axis evaluates current paradigms, including Deep Learning (DL), Explainable AI (XAI), Federated Learning (FL), and Reinforcement Learning (RL), emphasizing how these methods enhance perception, transparency, and adaptive control. The Foresight axis forecasts the integration of generative and agentic AI paradigms within distributed and 6G-enabled healthcare ecosystems.

The relationship among these axes can be expressed as an aggregated methodological function:

$$HIF_{survey} = f(H, I, F) = \alpha H + \beta I + \gamma F, \tag{1}$$

where H, I, and F represent the knowledge contributions from each axis, and α, β, and γ are weighting coefficients capturing their relative influence on the synthesis process. This equation conceptualizes how historical analysis (H), current insight (I), and forward-looking foresight (F) jointly inform a holistic understanding of AI evolution in mission-critical healthcare systems. Table 1 presents a compact taxonomy summarizing AI paradigms across three HIF dimensions: historical evolution, present-day techniques, and future innovations.

3 Hindsight: Evolution of Early AI in Mission-Critical Healthcare

The origins of AI in healthcare trace back to the era of symbolic reasoning and expert systems, which laid the foundation for modern mission-critical decision-support tools. Early AI systems such as MYCIN and CADIAG exemplified how encoded medical expertise could be used to assist physicians in diagnosis and treatment planning [10]. These systems relied on deterministic rule sets and knowledge bases where medical conditions were represented through "if-then" logic. For instance, MYCIN applied more than 600 handcrafted rules to identify bacterial infections and recommend antibiotics, while CADIAG supported internal medicine through fuzzy inference and pattern matching. Their transparency and explainability made them valuable in controlled settings; however, they lacked scalability and adaptability when confronted with novel clinical data [14].

With the rise of statistical and machine learning (ML) techniques in the 1980s and 1990s, the limitations of rule-based reasoning gave way to data-driven inference. Logistic regression, Naïve Bayes, and Support Vector Machines (SVM) improved diagnostic accuracy by learning probabilistic relationships from patient datasets. These models supported risk prediction in intensive care units (ICUs) and early computer-aided diagnosis of cardiac and neurological disorders. For example, ML-based arrhythmia detectors demonstrated that predictive algorithms could outperform manual electrocardiogram (ECG) interpretation in speed and sensitivity [19]. The mathematical formulation of such models can be generalized as:

$$P(y|X) = \sigma(WX + b), \tag{2}$$

where $P(y|X)$ denotes the predicted probability of a clinical outcome, W and b are model parameters, and σ represents the logistic activation function.

These foundational systems highlighted two enduring lessons. First, interpretability and transparency are critical for clinical acceptance, a principle that continues to shape modern XAI systems.

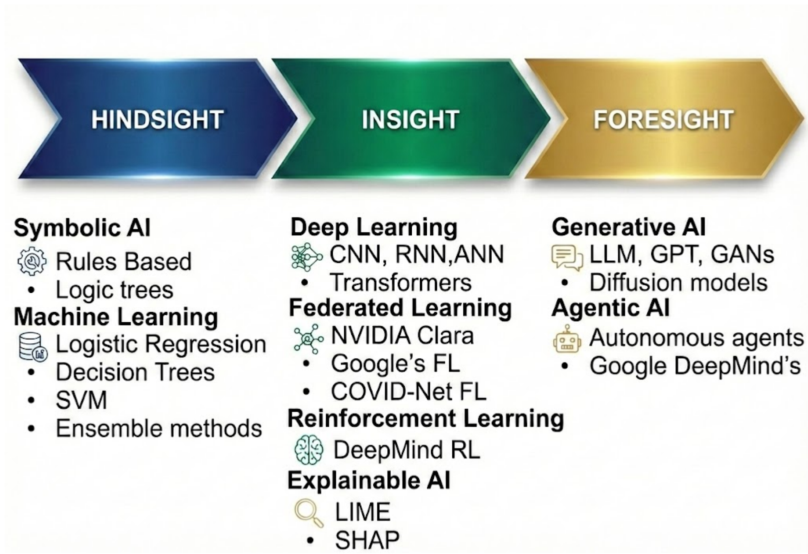

Figure 1: Chronological evolution of AI in mission-critical healthcare (Hindsight–Insight–Foresight framework).

Table 1: Taxonomy of AI Paradigms in Mission-Critical Healthcare

Phase	Technique	Healthcare Impact
Hindsight	Symbolic AI, Statistical ML	Diagnostic rule engines, ICU scoring systems
Insight	Deep Learning, XAI, FL, RL	Real-time analytics, privacy-preserving collaboration, adaptive treatment
Foresight	Generative AI, Agentic AI	Autonomous triage, multimodal decision-making, 6G-integrated healthcare

Table 2: Comparison of Early Medical Expert Systems

System	Date	Institution	Disease Type	Knowledge Base Coverage	Supported Technique	Area of Deployment
MYCIN	1970s	Stanford University	Bacteremia, Meningitis	600 rules, 100 therapies	Hand-coded rules, backward chaining logic	University Clinical Research
CADIAG-1	∼1980	University of Vienna	Internal Medicine	200 rules	Rule-based expert system	Clinical Decision Support
CADIAG-2	1980s	University of Vienna	Hepatology, Rheumatology	>300 diseases (diagnostic profiles)	Rule chaining, pattern matching	Hospital Department Use
INTERNIST-1	1974–1980s	University of Pittsburgh	General Internal Medicine	570 disease profiles	Heuristic rules, pattern recognition	Research Tool
CADUCEUS	1980s	University of Pittsburgh	Internal Medicine	1,000 disease frames	Advanced rule chaining + inference engine	Knowledge Base Expansion of INTERNIST-1

Table 3: Representative Case Studies of AI in Healthcare Systems

Case Study	Domain	Technique(s)	Institution/Developer	Impact/Use	Type
AISE – Artificial Intelligence Sepsis Expert	Sepsis (ICU)	Random Forest, Gradient Boosted Trees	University of California San Diego (UCSD)	Predicted septic shock up to 12 hrs earlier than clinicians	Academic
Epic Sepsis Model	Sepsis (EHR)	Logistic Regression + Rule-based Filters	Epic Systems (EMR vendor)	Widely deployed in hospitals; flagged sepsis risk, but high false positives	Industrial
Brain-IT Project	Neurocritical ICU	Time-series modeling + Shallow NN	University of Glasgow & EU Consortium	Predicted ICP spikes in traumatic brain injury (TBI) patients	Academic
MIT-BIH Arrhythmia Studies	Cardiology	k-NN, SVM, RNN	MIT (MIT-BIH Arrhythmia Database)	Benchmarked arrhythmia classification for clinical ECG monitoring	Academic
MIMIC-II & MIMIC-III	ICU Mortality	Logistic Regression, XGBoost	Beth Israel Deaconess / MIT (MIMIC datasets)	Predicted ICU mortality, ventilator use, and readmission with global research use	Academic

Table 4: Examples of Federated Learning Applications in Healthcare Systems

Example	Task / Domain	FL Architecture / Technique	Key Findings & Significance
FLICU (Mondrejevski et al.)	ICU mortality prediction	Cross-hospital FL workflow comparing local, centralized, and federated models using sequential deep models (LSTM, GRU, 1D-CNN)	The federated model matched or exceeded local models, demonstrating privacy-preserving viability for life-critical prediction tasks.
eICU Multi-Center Critical Care (Mehrjou et al.)	ICU survival prediction	Federated Averaging across multiple hospitals using eICU dataset	Showed how FL performance depends on client size, epoch settings, and data heterogeneity in critical care settings.
Deep Federated Learning for Early ICU Mortality	Early mortality forecasting from multivariate time series	Deep federated models combining time-series data across centers	Demonstrated that federated approach can improve predictive power while preserving data privacy in life-threatening settings.
Patient Clustering + FL for EMR Predictions	Mortality & hospital stay time using distributed EMRs	"Community-based Federated Learning" (CBFL)	Improved over basic FL under non-IID settings by grouping similar hospitals/clients to reduce heterogeneity impact.

Second, adaptability and continuous learning are essential in dynamic environments, motivating the transition from static models to data-driven and autonomous AI paradigms. The legacy of symbolic and early ML approaches thus forms the "Hindsight" layer of mission-critical AI evolution, illustrating the journey from handcrafted logic toward adaptive and intelligent healthcare systems [4, 11, 23].

4 Insight: Current AI Paradigms in Mission-Critical Healthcare

Recent advancements in Artificial Intelligence (AI) have redefined how mission-critical healthcare systems perceive, predict, and respond to clinical events in real-time environments. The integration of deep learning, explainable AI (XAI), federated learning (FL), and reinforcement learning (RL) has transformed hospitals from reactive to proactive environments capable of continuous, data-driven decision-making.

4.1 Deep Learning for Perception and Prognostics

Deep Learning (DL) architectures, including Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), underpin perception in healthcare. CNNs learn hierarchical feature representations directly from raw medical images such as X-rays or CT scans, achieving radiologist-level performance in diagnostic tasks like pneumonia detection with CheXNet [13, 16, 17]. The predictive relationship can be expressed as:

$$\hat{y} = f_{DL}(X; \theta) = \sigma(W_n g(\ldots g(W_1 X + b_1)\ldots) + b_n), \tag{3}$$

where X represents multimodal patient data, θ denotes learned parameters, and σ is the activation function producing diagnostic probabilities.

4.2 Explainable and Trustworthy AI

XAI improves trust and safety by visualizing and quantifying model reasoning. Methods such as LIME and SHAP compute local and global interpretability scores to show how specific features influence predictions [4, 5]. For instance, SHAP-based analysis in intensive care mortality prediction reveals that elevated lactate levels or low blood pressure strongly influence high-risk classifications, supporting clinician validation.

4.3 Federated Learning for Privacy Preservation

Federated Learning (FL) enables hospitals to jointly train models without sharing sensitive patient data [3] [20] [12]. Each participant computes a local update ∇_i and shares only model parameters with a central server that aggregates them as:

$$\theta^{t+1} = \sum_{i=1}^{N} w_i \theta_i^t, \tag{4}$$

where w_i is the proportional weight of each data source. This ensures compliance with GDPR and HIPAA while improving cross-institutional generalization [12, 23].

4.4 Reinforcement Learning for Adaptive Clinical Control

Reinforcement Learning (RL) optimizes sequential treatment strategies through continuous feedback. In sepsis management, an RL agent learns an optimal dosing policy $\pi^*(s)$ that maximizes the expected cumulative reward:

$$\pi^*(s) = \arg\max_{\pi} \mathbb{E}_{s_t, a_t \sim \pi} \left[\sum_t \gamma^t R_t \right], \tag{5}$$

where R_t is the clinical reward and γ is the discount factor controlling the balance between immediate and long-term patient outcomes [15]. RL provides adaptability in dynamic environments such as ICUs.

4.5 Integration Perspective

The synergy of DL, XAI, FL, and RL defines the cognitive backbone of mission-critical healthcare AI. DL enhances perception, XAI builds trust, FL ensures privacy, and RL promotes adaptability. Their integration across edge-cloud infrastructures and 6G-enabled networks represents the foundation of resilient, distributed, and explainable clinical intelligence.

5 Foresight: Emerging Directions and Future Innovations

The next frontier in mission-critical healthcare is driven by the convergence of Artificial Intelligence (AI), edge computing, and ultra-reliable low-latency communications (URLLC). s healthcare systems transition toward distributed intelligence, future AI frameworks will emphasize autonomy, trustworthiness, and interoperability within highly connected hospital ecosystems, while preserving the privacy of sensitive clinical data [6] through a tamperproof security mechanism. This section outlines the anticipated developments in Generative and Agentic AI and proposes a mathematical model that captures the trajectory toward intelligent, adaptive, and explainable healthcare systems.

5.1 Generative and Multimodal AI

Generative AI (GenAI) systems, comprising Generative Adversarial Networks (GANs), diffusion models, and Large Language Models (LLMs), are redefining data availability and multimodal reasoning. In healthcare, these models can synthesize realistic medical images for rare disease training, automate clinical documentation, and enable conversational diagnostic assistants. By generating synthetic datasets, GenAI mitigates the issue of data scarcity and supports balanced model training across institutions [4, 18]. Diffusion models enhance image quality reconstruction in low-resource environments, while domain-specific LLMs such as Med-PaLM facilitate real-time clinical Q&A for emergency decision support. However, challenges related to bias propagation, data authenticity, and model governance require standardized validation protocols.

5.2 Agentic AI and Autonomous Decision Systems

Agentic AI represents the evolution from predictive to autonomous intelligence. These systems integrate reasoning, planning, and coordination, enabling collaborative decision-making across distributed hospitals and edge nodes. Multi-Agent Reinforcement Learning (MARL) frameworks allow medical agents to share state information and optimize collective clinical outcomes. For instance, autonomous triage systems can coordinate emergency resources among connected hospitals, dynamically allocating personnel and equipment to reduce latency during critical response [9, 21]. The operational paradigm of such agents can be abstracted as:

$$\pi_i^*(s_t) = \arg\max_{\pi_i} \mathbb{E}_{s_t, a_t \sim \pi_i} [R_t + \lambda C(s_t, a_t)], \tag{6}$$

where $\pi_i^*(s_t)$ is the optimal policy for agent i, R_t is the clinical reward, and $C(s_t, a_t)$ is a collaboration term weighted by λ, encouraging cooperative decisions across the network.

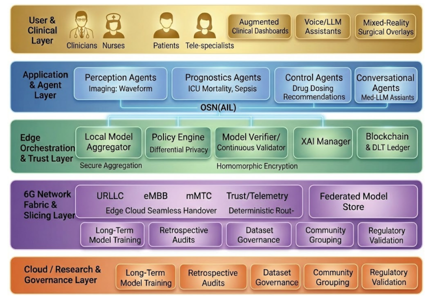

Figure 2: Proposed system architecture of the AI-Enabled Next-Generation Healthcare Information System (AIMS-6G). The framework integrates explainable and federated artificial intelligence modules with 6G edge intelligence for secure, privacy-preserving, and trustworthy healthcare delivery. Each layer—from cloud governance to user interaction—illustrates hierarchical integration of explainable AI, federated learning, blockchain-based trust, and 6G network slicing (URLLC, eMBB, mMTC) to enable ultra-reliable and low-latency medical decision support.

5.3 Integration with 6G and Edge Intelligence

Fig. 2 illustrates the layered architecture of the AI-Enabled Next-Generation Healthcare Information System (AIMS-6G), integrating explainable AI, federated learning, blockchain trust [1], and 6G edge intelligence. Each color-coded layer represents hierarchical modules enabling secure, low-latency, privacy-preserving, and trustworthy healthcare decision-making in next-generation medical systems. Thus, the future of healthcare AI (as seen in Fig. 2) will be shaped by 6G-enabled communication infrastructures supporting edge-to-cloud interoperability. Federated and distributed learning will operate over network slices with deterministic latency, ensuring rapid model updates for time-sensitive applications. Trustworthy AI pipelines will incorporate explainability (XAI), quantum-safe encryption, and continuous model verification to guarantee reliability and auditability [2, 8]. The synergy of these technologies will transform healthcare from static automation into a self-organizing, learning ecosystem that delivers tamperproof and trustworthy feedback.

5.4 Summary and Way Forward

The Foresight perspective reveals a healthcare paradigm in which AI agents collaborate across connected hospitals, leveraging generative synthesis, reinforcement learning, and edge intelligence to enable adaptive care. The strategic objective is to establish an AI-enabled Mission-Critical Healthcare System (AIMS) that optimizes performance, transparency, and resilience. This goal can be modeled as an optimization problem:

$$AIMS^* = \arg\max_{XAI, FL, RL, GenAI} [T_r + S_q - R_c], \tag{7}$$

Table 5: Healthcare Applications and AIMS Alignment

Application	Organization	Domain & Use Case	Notes	AIMS Alignment
Healthcare – Intensive Care Units (ICUs)	Hospitals, Medical Centers	Continuous monitoring and intervention for patients with life-threatening conditions	Mission-critical systems where real-time AI support ensures survival	Demonstrates AIMS role in perception, monitoring, and decision-making
Healthcare – Emergency Response	Emergency Services, Disaster Management Agencies	Rapid response during disasters, pandemics, and public health emergencies	Requires uninterrupted operations and coordinated medical delivery	AIMS ensures agility, resilience, and adaptive response
Healthcare – Diagnostic Imaging	Radiology Departments	CT, MRI, X-ray for detecting stroke or hemorrhage in real time	Delays can be fatal; AI ensures speed and accuracy	AIMS integrates perception systems for life-critical diagnostics
Surgical Robotics	Hospitals, Research Labs	High-risk operations (cardiac, neurosurgery) requiring robotic precision	Must operate flawlessly under uncertainty	AIMS enhances adaptive precision and ensures patient safety
Clinical Decision Support Systems (CDSS)	Healthcare IT, Hospitals	Real-time evidence-based recommendations using multimodal data	Supports high-stakes decision-making with explainability	AIMS integrates decision intelligence with human-in-the-loop
Security & Safety	Critical Infrastructure, National Security	Protection of mission-critical systems against cyber and physical threats	Ensures system resilience, trust, and robustness	AIMS enforces trustworthy, secure AI operations

where T_r denotes system trust and reliability, S_q represents service quality and speed, and R_c captures communication and computation risk. The solution encapsulates the multidisciplinary direction of AI in mission-critical healthcare, balancing explainability, privacy, and scalability toward a secure, connected, and adaptive medical future [7, 18, 22].

6 Conclusion

This paper provides a concise Hindsight–Insight–Foresight perspective of AI in mission-critical healthcare information systems. From symbolic and statistical reasoning to generative and agentic intelligence, AI has evolved to support real-time, trustworthy, and distributed clinical decision-making. Future research should focus on integrating explainability, privacy, and resilience within 6G-connected medical environments through standardized and certifiable AI frameworks.

Acknowledgments

This research was supported by the Priority Research Centers Program through the NRF funded by the MEST (2018R1A6A1A03024003) (50%) and by the Institute of Information & Communications Technology Planning & Evaluation (IITP)-Innovative Human Resource Development for Local Intellectualization program grant funded by the Korea government(MSIT)(IITP-2025-RS-2020-II201612) (50%).

References

[1] Ihunanya Udodiri Ajakwe, Victor Ikenna Kanu, Simeon Okechukwu Ajakwe, and Dong-Seong Kim. 2025. eBCTC: Energy-Efficient Hybrid Blockchain Architecture for Smart and Secured K-ETS. *Cleaner Engineering and Technology* (2025), 101084.
[2] Simeon Okechukwu Ajakwe, Igboanusi Ikechi Saviour, Vivian Ukamaka Ihekoronye, Odinachi U Nwankwo, Mohamed Abubakar Dini, Izuazu Urslla Uchechi, Dong-Seong Kim, and Jae Min Lee. 2024. Medical IoT Record Security and Blockchain: Systematic Review of Milieu, Milestones, and Momentum. *Big Data and Cognitive Computing* 8, 9 (2024), 121.
[3] Ons Aouedi, Alessio Sacco, Kandaraj Piamrat, and Guido Marchetto. 2022. Handling privacy-sensitive medical data with federated learning: challenges and future directions. *IEEE journal of biomedical and health informatics* 27, 2 (2022), 790–803.
[4] Arun Das and Paul Rad. 2020. Opportunities and challenges in explainable artificial intelligence (xai): A survey. *arXiv preprint arXiv:2006.11371* (2020).
[5] Arun Das and Paul Rad. 2020. Opportunities and Challenges in Explainable Artificial Intelligence (XAI): A Survey. arXiv:2006.11371 [cs.CV] https://arxiv.org/abs/2006.11371
[6] Mohamed A Dini, Simeon O Ajakwe, Igboanusi I Saviour, Vivian U Ihekoronye, Odinachi U Nwankwo, Izuazu U Uchechi, Gifar Arif Haryadi, Made Adi Paramartha Putra, Dong-Seong Kim, Taesoo Jun, et al. 2023. Patient-centric blockchain framework for secured medical record fidelity and authorization. *Proceedings of the Korean Institute of Communications and Information Sciences (KICS) Conference* (2023), 300–301.
[7] Emily Haeuser, Sam Byrne, Jason Nguyen, Catalina Raggi, Susan A McLaughlin, Catherine Bisignano, Ashley A Harris, Amanda E Smith, Paulina A Lindstedt, Georgia Smith, et al. 2025. Global, regional, and national trends in routine childhood vaccination coverage from 1980 to 2023 with forecasts to 2030: a systematic analysis for the Global Burden of Disease Study 2023. *The Lancet* (2025).
[8] Kyriakos Houliotis, Panagiotis Oikonomidis, Periklis Charchalakis, and Elias Stipidis. 2018. Mission-critical systems design framework. *Advances in Science, Technology and Engineering Systems Journal* 3, 2 (2018), 128–137.
[9] Nalan Karunanayake. 2025. Next-generation agentic AI for transforming healthcare. *Informatics and Health* 2, 2 (2025), 73–83.
[10] Casimir A Kulikowski. 1980. Artificial intelligence methods and systems for medical consultation. *IEEE Transactions on pattern analysis and Machine Intelligence* 5 (1980), 464–476.
[11] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature* 521, 7553 (2015), 436–444.
[12] G Muni Nagamani and Chanumolu Kiran Kumar. 2024. Design of an improved graph-based model for real-time anomaly detection in healthcare using hybrid CNN-LSTM and federated learning. *Heliyon* 10, 24 (2024).
[13] Loris Nanni, Stefano Ghidoni, and Sheryl Brahnam. 2017. Handcrafted vs. non-handcrafted features for computer vision classification. *Pattern recognition* 71 (2017), 158–172.
[14] Witold Pedrycz and Armando F Rocha. 2002. Fuzzy-set based models of neurons and knowledge-based networks. *IEEE Transactions on Fuzzy Systems* 1, 4 (2002), 254–266.
[15] Aniruddh Raghu, Matthieu Komorowski, Imran Ahmed, Leo Celi, Peter Szolovits, and Marzyeh Ghassemi. 2017. Deep reinforcement learning for sepsis treatment. *arXiv preprint arXiv:1711.09602* (2017).
[16] Pranav Rajpurkar, Jeremy Irvin, Kaylie Zhu, Brandon Yang, Hershel Mehta, Tony Duan, Daisy Ding, Aarti Bagul, Curtis Langlotz, Katie Shpanskaya, et al. 2017. Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning. *arXiv preprint arXiv:1711.05225* (2017).
[17] Leonardo Rundo and Carmelo Militello. 2024. Image biomarkers and explainable AI: handcrafted features versus deep learned features. *European Radiology Experimental* 8, 1 (2024), 130.
[18] Vidhi Singh, Susan Cheng, Alan C Kwan, and Joseph Ebinger. 2025. United States Food and Drug Administration regulation of clinical software in the era of artificial intelligence and machine learning. *Mayo Clinic Proceedings: Digital Health* 3, 3 (2025), 100231.
[19] KP Soman, R Loganathan, and V Ajay. 2009. *Machine learning with SVM and other kernel methods.* PHI Learning Pvt. Ltd.
[20] Simon Süwer, Md Shihab Ullah, Niklas Probul, Andreas Maier, and Jan Baumbach. 2024. Privacy-by-Design with Federated Learning will drive future Rare Disease Research. *Journal of Neuromuscular Diseases* (2024), 22143602241296276.
[21] Chris Thames and Yifan Sun. 2024. A Survey of Artificial Intelligence Approaches to Safety and Mission-Critical Systems. In *2024 Integrated Communications, Navigation and Surveillance Conference (ICNS)*. IEEE, 1–12.

[22] Tagne Poupi Theodore Armand, Md Ariful Islam Mozumder, Kouayep Sonia Carole, Opeyemi Deji-Oloruntoba, Hee-Cheol Kim, and Simeon Okechukwu Ajakwe. 2024. ELIPF: Explicit Learning Framework for Pre-Emptive Forecasting, Early Detection and Curtailment of Idiopathic Pulmonary Fibrosis Disease. *BioMedInformatics* 4, 3 (2024), 1807–1821.

[23] Neil C. Thompson, Kristjan Greenewald, Keeheon Lee, and Gabriel F. Manso. 2022. The Computational Limits of Deep Learning. arXiv:2007.05558 [cs.LG] https://arxiv.org/abs/2007.05558