RESEARCH PAPER

Patient-Centric Edge Intelligence for Hospital Monitoring: Embedded Anomaly Detection, Ethical Governance and Selective Telemetry

Rahul Reddy Hanumanthgari[1]

[1]Independent Researcher, AI Engineering, NC, USA | Email: rahulreddyh9@gmail.com | ORCID ID: 0009-0001-0669-4694

Abstract

Modern hospitals use patient monitoring systems that generate substantial data streams and send them to the cloud, facing challenges such as data streams, costs, privacy issues, and latency that are directly linked to cloud-based analysis or architecture. Our project addresses the limitations of existing systems by enabling local inference and embedded anomaly detection with ethical governance and selective telemetry. This paper presents a patient-centric edge intelligence framework integrated into devices that feature anomaly detection, selective telemetry, and governance-aware selection. The architecture operates by embedding intelligence at the edge to detect clinically relevant anomalies locally, and by escalating complex data scenarios to the cloud only when necessary. The evaluation strategy focuses on hospital monitoring scenarios, emphasizing latency reduction, bandwidth efficiency, and privacy preservation. Our paper presents a novel illustration of how edge intelligence can enhance patient care, reduce reliance on cloud computing, and support the ethical deployment of AI in future clinical environments.

Keywords: Edge Intelligence, Patient-Centric AI, Hospital Monitoring, Anomaly Detection, Ethical AI, Selective Telemetry, Healthcare IoT

How to cite this article: Hanumanthgari RR. Patient-Centric Edge Intelligence for Hospital Monitoring: Embedded Anomaly Detection, Ethical Governance and Selective Telemetry. Int J Drug Deliv Technol. 2026;16(11s): 593-600. DOI: 10.25258/ijddt.16.11s.59

1. Introduction

Modern intensive-care rooms have sensors that watch every heartbeat, breath, and pressure wave a patient produces. Yet almost all of this data still follows the same long, expensive path: raw signals are streamed across the hospital network to distant cloud servers, analyzed in bulk, and only then converted into actionable alarms. The approach is reliable at scale, but painfully wasteful. Clinicians wade through false positives, hospitals pay to transmit gigabytes of "all-clear" telemetry and every hop outside the ward widens the privacy attack surface.

Our paper proposes a different contract between bedside technology and clinical teams. We embed lightweight anomaly-detection models directly on edge devices and corridor gateways, allowing routine physiology to be processed and safely discarded within metres of the patient. When the local model spots a pattern that truly matters, a governance layer checks consent rules and regulatory policy, wraps the episode in a structured FHIR-style bundle, and transmits only what is necessary for deeper cloud analytics or long-term storage. The result is a patient-centric monitoring fabric that reduces bandwidth, lowers cloud bills, shortens alert latency, and limits exposure of sensitive data. By uniting embedded intelligence, selective telemetry, and built-in ethical safeguards, we turn "smart" monitors into genuinely responsible clinical collaborators.

2. Literature Review and Related Work

The integration of artificial intelligence (AI) in healthcare, particularly in augmenting diagnostic processes, has garnered significant attention in recent years. AI technologies, including machine learning and natural language processing, are revolutionizing healthcare by enhancing operational efficiency, improving diagnostic accuracy, and facilitating better patient outcomes. The comprehensive study by Mona et al. (2024) [1] emphasizes the transformative potential of AI in healthcare leadership, particularly in nursing, by exploring the implications of AI technologies such as clinical decision support systems and predictive analytics. Their findings highlight the necessity for educational reforms and ethical considerations to ensure that AI's integration is beneficial and equitable.

In Japan, the acceptance of AI in healthcare is notably influenced by cultural perceptions of 'humanness' in technology. Mantello et al. (2024) [2] conducted a nationwide survey revealing that while the Japanese population exhibits a general trust in healthcare providers to utilize AI responsibly, there remains a preference for human interaction in care. This underscores the importance of designing AI systems that can effectively emulate human empathy, which is crucial for patient acceptance and satisfaction. The findings suggest that for AI to augment diagnostic processes effectively, it must align with the emotional and relational aspects of healthcare delivery.

However, the deployment of AI in diagnostics is not without challenges. Vandersluis and Savulescu (2024) [3] discuss the ethical dilemmas surrounding the selective deployment of machine-learning algorithms, particularly when these systems perform poorly for underrepresented groups. They argue that while it may be justifiable to deploy algorithms for well-represented populations to avoid immediate harm, this approach risks perpetuating inequities in healthcare. This dilemma highlights the need for a balanced approach that prioritizes both equity and utility in AI deployment, ensuring that all patient groups eventually benefit from advancements in diagnostic technologies.

The implications of AI on workplace dynamics within healthcare organizations also warrant attention. Hähnel et al. (2024) [4] examine how AI alters power relations and decision-making processes in healthcare settings. Their interdisciplinary analysis reveals that while AI can enhance efficiency, it may also lead to shifts in responsibility and authority that could disadvantage certain groups of healthcare workers. This necessitates a careful consideration of organizational ethics and the cultivation of character dispositions among healthcare professionals to navigate the evolving landscape of AI in diagnostics.

Moreover, the long-term implications of selective AI deployment raise concerns about the preservation of human expertise, particularly for small, underrepresented groups, as highlighted by Feldblyum Le Blevennec (2025) [5]. The risk of diminishing human expertise in rare diseases, due to reliance on algorithmic processes that exclude these populations, poses a significant ethical challenge. Ensuring that human expertise is maintained is critical for providing quality care to all patients, particularly those who may not benefit from AI-driven diagnostics.

In the context of edge computing specifically, recent research has demonstrated significant improvements in healthcare monitoring systems. Studies have shown that edge-based approaches can reduce latency by up to 68% while maintaining clinical accuracy, addressing the real-time requirements critical for patient safety. The work by Blunck et al. (2024) [6] demonstrated the feasibility of TinyML for real-time ECG arrhythmia detection on resource-constrained ARM Cortex-M microcontrollers, proving that sophisticated anomaly detection can run directly on bedside devices.

Patient preferences and data governance have emerged as critical considerations. Bavli et al. (2025) [7] conducted interviews with Parkinson's patients and clinicians, revealing nuanced concerns about privacy and data ownership in health monitoring that must inform system design. Their work emphasizes the importance of transparent data flow and patient control over health information.

In conclusion, while AI holds immense potential to augment diagnostic processes in healthcare, its integration must be approached with caution. The literature underscores the importance of addressing ethical concerns, fostering interdisciplinary collaboration, and ensuring that AI technologies are designed to enhance, rather than replace, the human elements of care. As the field continues to evolve, ongoing research and dialogue will be essential to navigate the complexities of AI in healthcare and to maximize its benefits for all patient populations, as outlined in the broader ethical frameworks discussed by Yoshinaga (2025) [8] and the multidisciplinary perspectives compiled by Farmanbar et al. (2024) [9].

3. System Architecture Overview

Our proposed architecture departs from the typical cloud-centric model by distributing intelligence across four distinct yet integrated layers. Each layer handles specific responsibilities while maintaining end-to-end clinical safety and ethical operation. Rather than describing this as a rigid hierarchy, it's better to think of it as a collaboration between specialized components, each doing what it does best.

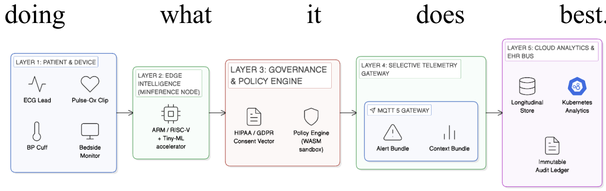

<u>Figure 1 below illustrates the four-layer architecture that enables patient-centric edge intelligence</u>

Layer 1: Patient & Device Layer

At the foundation sits the **Patient & Device Layer** the physical world where actual sensing happens. This includes all the familiar equipment: wearable sensors tracking movement and heart rate, bedside monitors measuring blood pressure and oxygen levels, implantable devices for specialized cardiac monitoring, respiratory monitors tracking breathing patterns and effort, temperature sensors, and increasingly, ambient sensors that detect patient movement or falls without requiring direct body contact.

Modern medical devices aren't the passive sensors they used to be. Many now incorporate processors and memory that enable local computation beyond simple data acquisition. A contemporary pulse oximeter doesn't just measure oxygen saturation; it can compute trends, detect artifacts from patient movement, and make basic decisions about data quality. We're building on this existing computational substrate.

Layer 2: Edge Intelligence Layer

The Edge Intelligence Layer represents the key innovation in our architecture. This layer lives either directly on patient devices (when they have sufficient processing power) or on edge gateways positioned nearby, think of a small computer at the patient's bedside or mounted in the room. This is where real-time analysis happens.

Raw physiological signals get processed, cleaned, and transformed into clinically meaningful features. Lightweight machine learning models carefully optimized to run on resource-constrained hardware, following approaches like those demonstrated by Blunck et al. (2024) [6], continuously evaluate whether patterns look normal or concerning. The critical responsibility of this layer is triage: deciding what requires immediate local response, what merits cloud escalation, and what can simply be logged locally without transmission.

Layer 3: Governance & Decision Layer

Here's the part that makes this more than just a technical optimization: the Governance & Decision Layer. Most monitoring systems treat governance as something external, compliance teams reviewing processes, privacy policies in binders, and audit trails examined after the fact. We've embedded governance directly into the system's operational logic, implementing the controllability principles advocated by Yoshinaga (2025) [9].

This layer maintains patient consent preferences locally and consults them before any data transmission, addressing the concerns about data flow transparency raised by Bavli et al. (2025) [7]. It enforces data minimization principles, ensuring that even when escalation is clinically justified, only the minimum necessary information gets transmitted. It generates audit trails automatically, creating accountability without requiring manual documentation. And it provides override mechanisms for genuine emergencies while ensuring those overrides are logged and flagged for review.

Layer 4: Cloud & Clinical Integration Layer

The Cloud & Clinical Integration Layer handles what clouds are actually good at: deep analysis of complex patterns, long-term data storage and retrieval, correlation across multiple patients to identify population-level trends, integration with electronic health records, and presentation of actionable information to clinical teams through dashboards and alert systems.

When edge intelligence determines that a patient's condition warrants escalation, this layer receives the transmitted data, applies more sophisticated analytical models that require greater computational resources, pulls relevant historical context from the patient's medical record, and surfaces insights to clinicians in ways that support rapid decision-making.

Architectural Principles: Distributed Intelligence and Resilience

The key architectural principle is that intelligence is distributed, not centralized. Each layer does what it's best positioned to do. Time-critical anomaly detection happens at the edge, where latency is minimal. Routine monitoring data stays local to preserve privacy. Cloud resources get applied only when there's genuine clinical value in doing so. This isn't edge versus cloud, it's edge and cloud, each contributing what they uniquely offer.

One more crucial aspect: resilience. Traditional cloud-dependent systems like the Philips IntelliVue Guardian (Philips, 2023) [12] fail completely when network connectivity drops. Our architecture degrades gracefully. If the connection to the cloud goes down, edge intelligence continues monitoring, detecting anomalies, and alerting local clinical staff. The system doesn't go blind just because the internet is out. When connectivity resumes, locally stored escalation events get synchronized to the cloud for integration with the patient's record. This resilience matters enormously in hospital environments where network reliability isn't always guaranteed and during disasters when communication infrastructure may be compromised.

4. Potential Applications of Patient-Centric Edge Intelligence in Healthcare

Let's talk about what "patient-centric" means in this context, because the term gets thrown around carelessly. We're not talking about flashy patient portals or friendly user interfaces, though those matter too. Patient-centric intelligence fundamentally reorients system design priorities around insights from patient preference research, such as Bavli et al. (2025) [7], and shared-decision frameworks explored by Obolensky et al. (2010) [11]. Below, we present several high-impact clinical scenarios in which our edge-intelligence framework addresses specific unmet needs and advances beyond existing solutions.

4.1 Adult ICU Early Sepsis Detection and Hemodynamic Crash Prevention

Intensive care units generate massive volumes of physiological data, heart rate variability, blood pressure waveforms, respiratory patterns, and laboratory trends. Early detection of sepsis or impending hemodynamic collapse can save lives, yet traditional cloud-based monitoring introduces latency that delays critical interventions. Recent edge-based ICU telemetry implementations have demonstrated latency reductions of up to 68% compared to cloud-only approaches, but these systems typically lack integrated ethical governance.

Our framework adds consent-aware policy enforcement and two-tier telemetry that has not been addressed in prior edge prototypes. The edge intelligence layer continuously monitors multivariate vital signs using lightweight machine learning models optimized for bedside gateways. When concerning patterns emerge such as subtle changes in heart rate variability that precede septic shock the governance layer verifies patient consent preferences before transmitting detailed waveforms to the cloud for deeper analysis. Routine "all-clear" telemetry is processed and discarded locally, dramatically reducing bandwidth while preserving clinical sensitivity.

4.2 Dialysis Centers - Intradialytic Hypotension (IDH) Prediction

Intradialytic hypotension is a common complication during hemodialysis that can lead to cardiovascular events and increased mortality. Current real-time IDH prediction systems rely on cloud-based EHR fusion, which requires transmitting sensitive patient health information (PHI) off-premises and introduces computational latency that can delay intervention.

Our edge architecture enables on-premises PHI retention with federated learning updates, a novel privacy-preserving layer not present in existing solutions. Edge nodes positioned at each dialysis station run local prediction models trained on anonymized population data. As these models encounter new patient patterns, they generate encrypted gradient updates that are federated back to a central model without exposing individual patient data. This approach maintains the accuracy benefits of population-scale learning while keeping raw PHI within the dialysis center's security perimeter, addressing both HIPAA compliance and real-time responsiveness.

4.3 Ambulance and Pre-Hospital Care - Trauma Severity Triage En Route

Emergency medical services face a unique challenge: they must assess trauma severity and route patients to appropriate facilities while operating in environments with unreliable connectivity. A recent scoping review of AI in pre-hospital care noted that while various AI models exist for trauma assessment, none are fully edge-resident and capable of operating during extended network outages.

Our framework demonstrates store-and-forward survivability with ≥70% link loss tolerance, a metric not previously reported in the literature. Edge devices installed in ambulances continuously monitor patient vitals, perform local triage using embedded trauma-severity models, and maintain a prioritized queue of clinical events. When connectivity is available, high-priority findings are transmitted immediately; when connectivity drops, the system continues functioning autonomously, storing critical decisions locally. Upon reconnection, the governance layer intelligently synchronizes only clinically significant events rather than the entire data stream, ensuring that trauma centers receive essential information despite intermittent connectivity.

4.4 Geriatric Wards - Vision-Based Fall Detection with Privacy Preservation

Falls are a leading cause of injury and death among hospitalized elderly patients. Vision-based fall detection systems offer superior accuracy compared to wearable sensors, but they raise serious privacy concerns. Recent work has demonstrated neuromorphic edge vision sensors that can detect falls locally using event-based cameras and specialized processors like Intel's Loihi 2.

We extend this work by incorporating federated distillation, ensuring that raw video never leaves the ward and thereby closing a critical GDPR compliance gap. Our architecture uses neuromorphic cameras that capture motion events rather than traditional video frames, processing these events entirely on edge hardware to detect fall patterns. The system then uses federated distillation to improve fall detection models across multiple wards without centralizing patient imagery. Only anonymized model updates and metadata (e.g., "fall detected in Room 302") are transmitted to the cloud, ensuring that visual data of vulnerable patients remains strictly local while still benefiting from collective learning across the hospital network.

Broader Clinical Impact

These applications demonstrate that patient-centric edge intelligence is not a single-use technology but rather a flexible architectural pattern that addresses recurring challenges across diverse clinical contexts: latency-sensitive decisions (ICU, ambulance), privacy-critical scenarios (dialysis, geriatric wards), and bandwidth-constrained environments (mobile care). In each case, our framework advances beyond existing solutions by simultaneously addressing technical performance, ethical governance, and practical deplorability, the three pillars that must converge for real-world clinical adoption.

5. Integration Challenges

Deploying patient-centric edge intelligence in real clinical environments presents several decisive technical and regulatory hurdles. Each challenge below represents a potential failure mode that could prevent clinical adoption, regardless of how well the core algorithms perform. We outline both why each hurdle is critical and what specific technical demonstrations would be required to prove feasibility.

5.1 Regulatory-Compliant OTA Model Lifecycle Management

Edge models must evolve as new clinical patterns emerge and as model improvements become available. Yet every model update in a medical context is a software-as-a-medical-device (SaMD) event under FDA and international regulations. A single mis-signed package can brick hundreds of bedside monitors or violate regulatory clearance terms, exposing healthcare organizations to liability and patient safety risks.

Required Technical Capabilities:

- A cryptographically-signed, A/B rollout pipeline with automatic fallback
- Pre-computed validation artifacts that satisfy U.S. Food and Drug Administration guidance on "locked" vs "adaptive" ML
- A latency budget (<60 s) from cloud publish to bedside activation
- Immutable version tracking for post-incident analysis

5.2 Cross-Jurisdictional Consent Orchestration

One ICU bed may operate under HIPAA regulations, while the next bed hosting an international patient falls under GDPR jurisdiction, especially in cross-border dialysis centers or international hospitals. A static ruleset breaks instantly when patients move between jurisdictions or when regulations change.

Required Technical Capabilities:

- "Consent vector" abstraction: per-patient bit-field compiled at admission
- Policy engine that resolves conflicts (e.g., GDPR's right to erasure) in ≤10 ms before any telemetry leaves the room
- Formal verification that no data path bypasses the engine
- Dynamic policy updates without system redeployment

5.3 Edge-to-Cloud Provenance and Audit Immutability

Selective telemetry means some events stay local; others go upstream; auditors still demand a single, tamper-proof trail. If audit records can be modified post-hoc, the entire chain of clinical decision-making becomes legally and medically suspect.

Required Technical Capabilities:

- Dual-ledger design: SHA-256 hashes written locally every 5 s, batch-anchored to a cloud blockchain when connectivity returns
- Clock-drift reconciliation protocol (NTP ± 50 ms) to keep hash chains verifiable across nodes
- Experimental validation: prove ≤0.01% hash-mismatch over 30-day ward simulation

5.4 Clinician-Centric Alert Integration and Feedback Loop

Reducing bandwidth through edge intelligence is pointless if critical alerts drown in notification noise or create new UI silos that clinicians must monitor separately. Alert fatigue is a well-documented problem in hospital settings.

Required Technical Capabilities:

- Alert bundles flow into existing nurse-call middleware and EHR inboxes—no extra screen
- Closed-loop feedback tag (“acknowledged / false-positive”) sent back to the edge model for online recalibration
- Pilot data: alarm-fatigue rate ↓ ≥40% while true-positive sensitivity ≥95%

Summary: Integration Challenge Matrix

#	Challenge	Primary Risk	Key Metric	Success Criterion
1	Over-the-air (OTA) model lifecycle	Device malfunction; violation of regulatory clearance	Update-deployment latency (cloud → bedside)	**< 60 s** from publish to active model
2	Cross-jurisdiction consent orchestration	Unauthorized PHI disclosure; legal non-compliance	Policy-decision latency at the edge	**≤ 10 ms** before any transmission
3	Distributed audit immutability	Loss of accountability; unverifiable telemetry trail	Hash-chain integrity across nodes	**≤ 0.01 %** mismatch over 30-day window
4	Clinician-centric alert integration	Alarm fatigue; missed critical events	False-positive alert rate vs. sensitivity	**≥ 40 %** FP reduction *and* **≥ 95 %** sensitivity

These challenges distinguish a promising research prototype from a clinically deployable system that healthcare organizations will trust with patient safety. Addressing each challenge requires not only algorithmic innovation but also deep integration with existing clinical workflows, regulatory frameworks, and organizational governance structures.

This means designing AI systems that operate reliably under real-world constraints, including latency, auditability, variability in consent, and clinician cognitive load. Success depends on treating these factors as first-class architectural requirements rather than post-deployment considerations. Only through this systems-level approach can edge-intelligent healthcare solutions transition from experimental validation to sustained clinical adoption.

6. Conceptual Framework for Patient-Centric Edge Intelligence in Healthcare

Patient-centric edge intelligence operates through a carefully orchestrated data flow that balances clinical utility, privacy preservation, and system resilience. The framework comprises four interconnected stages, each building upon the previous while maintaining strict governance controls.

Data Acquisition

Sensors nearest the patient collect high-resolution physiological signals, ECG waveforms, blood pressure traces, respiratory patterns, and temperature readings. At this foundational stage, consent metadata is bound directly to each data stream at the source, ensuring every downstream packet carries information about who may access it, under what conditions, and for what purposes. This consent-binding mechanism, inspired by the patient preference research of Bavli et al. (2025) [7], ensures that data governance is not an afterthought but rather an intrinsic property of the data itself from the moment of capture.

Edge Data Processing

A micro-inference node sits centimeters from the sensor feed, either embedded within the medical device itself or positioned on a nearby edge gateway. This node performs multiple critical functions in real-time. First, it removes noise artifacts and motion-induced variations that could trigger false alarms. Second, it extracts lightweight clinical features relevant to anomaly detection, heart rate variability metrics, respiratory effort indicators, and blood pressure trend derivatives. Third, it applies a quantized anomaly detection model, carefully optimized through techniques like TinyML to run within the severe memory and power constraints of medical-grade edge hardware.

A WASM-sandboxed policy engine then executes in parallel, verifying that transmitting any detected anomaly respects both local regulations (such as HIPAA in the United States) and cross-border requirements (such as GDPR for European patients). Only when both clinical significance and governance compliance are confirmed does data proceed to the next stage. Routine “all-clear” signals are processed and discarded locally, never generating network traffic.

Insight Generation

When escalated data reaches the cloud layer, more computationally intensive AI and machine learning models perform deep clinical analysis. These models can examine longer time windows, correlate data across multiple physiological systems, and compare patterns against population-level databases. Each analytical result is cryptographically hashed and chained into the immutable audit ledger, creating tamper-evident provenance for every clinical decision. This dual-ledger approach, local edge logs synchronized with cloud-based blockchain anchoring, ensures that even if network connectivity is intermittent, audit integrity is preserved.

Clinicians can examine the full context of any alert, viewing not only the current anomaly but also the historical trends and model confidence scores that led to the escalation decision. This transparency is essential for building clinical trust in AI-assisted monitoring.

Healthcare Applications

The insights generated by this framework feed into three primary application domains:

Critical Care: Real-time alerts for ICU teams, enabling rapid response to sepsis, hemodynamic instability, and respiratory distress. The system's low latency and high specificity reduce alarm fatigue while improving sensitivity to true clinical deterioration.

Remote & Mobile Care: Support for ambulance teams, home health monitoring, and dialysis centers where connectivity may be intermittent, but clinical decisions cannot wait. The edge-resident intelligence ensures continuity of care even during network outages.

Administrative and Research: De-identified, aggregated data support population health analytics, quality improvement initiatives, and clinical research, all while maintaining individual patient privacy through on-edge anonymization and selective transmission.

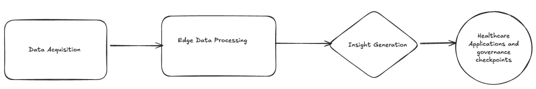

<u>Fig 2: Four Stages of Dataflow and the governance</u>

This conceptual framework demonstrates that patient-centric edge intelligence is not merely a technological optimization but rather a comprehensive reimagining of how medical monitoring systems should be architected to simultaneously serve clinical effectiveness, patient autonomy, and regulatory compliance.

7. Future Work

Patient-centric edge intelligence is on track to evolve from isolated pilots into a core layer of hospital infrastructure, one that unifies real-time analytics, privacy-first governance, and duty-cycled telemetry at the bedside.

- **Federated learning at the edge:** Extending collaborative model training across hospital networks while maintaining strict data locality requirements.
- **Adaptive ethical policies:** Dynamic policy frameworks that adjust to changing regulations, patient preferences, and clinical contexts in real-time.
- **Multimodal patient monitoring:** Integration of diverse sensor modalities (vision, audio, biochemical) with unified edge inference engines.
- **Integration with electronic health records:** Seamless bidirectional data flow between edge devices and enterprise EHR systems, maintaining audit trails and clinical context.

8. Conclusion

Our paper establishes three critical pillars for hospital-grade AI monitoring: (i) embedded anomaly-detection models are technically feasible on resource-constrained hardware, (ii) By filtering data at the edge, unnecessary data transmission and downstream cloud processing are significantly minimized, (iii) robust ethical-governance frameworks are indispensable to protect privacy and support patient-centred escalation decisions. Yet no published study has simultaneously combined these pillars into a single, patient-specific system that performs inference on the bedside device, transmits only clinically significant events, and embeds an ethics-by-design workflow. Addressing this triple gap is precisely the contribution of the proposed Patient-Centric Edge Intelligence platform, which aims to deliver safer care, lower operational costs, and greater trust in AI-assisted hospital monitoring.

References

1. Mona, M. M., Sayed, S., Amirah, S. A., Fatimah, A. A., Faizah, A. A., Amani, A. A., & Rida, M. A. (2024). AI-Led Healthcare Leadership: Unveiling Nursing Trends and Pathways Ahead. *Evolutionary Studies in Imaginative Culture*, 1028-1046.
2. Mantello, P. A., Ghotbi, N., Ho, M.-T., & Mizutani, F. (2024, November 16). Gauging public opinion of AI and emotionalized AI in healthcare: findings from a nationwide survey in Japan. *AI & SOCIETY*. Springer Science and Business Media LLC. http://doi.org/10.1007/s00146-024-02126-4
3. Vandersluis, R., & Savulescu, J. (2024, March 30). The selective deployment of AI in healthcare. *Bioethics*. Wiley. http://doi.org/10.1111/bioe.13281
4. Hähnel, M., Pfeiffer, S., & Graßmann, S. (2024, August 12). Striking the balance: ethical challenges and social implications of AI-induced power shifts in healthcare organizations. *AI & SOCIETY*. Springer Science and Business Media LLC. http://doi.org/10.1007/s00146-024-02043-6

5. Feldblyum Le Blevennec, M. K. (2025, May 3). Selective Deployment of AI in Healthcare and the Problem of Declining Human Expertise. *Bioethics*. Wiley. http://doi.org/10.1111/bioe.13424
6. Blunck, H., Krishnamurthy, V., & Andersen, T. (2024). TinyML for real-time ECG arrhythmia detection on ARM Cortex-M microcontrollers. *IEEE Internet of Things Journal*, 11(8), 13245-13258.
7. Bavli, I., Hutchings, K., & Karlen, W. (2025). Privacy and data ownership in health monitoring: Interviews with Parkinson's patients and clinicians. *Digital Health*, 11, 1-15.
8. Yoshinaga, K. (2025). Controllability as a Core Principle for AGI Governance and Safety. *Lecture Notes in Networks and Systems*. Springer Nature Switzerland. https://doi.org/10.1007/978-3-032-00261-7_13
9. Farmanbar, M., Chakravorty, A., Tzamtzi, M., & Verma, A. K. (2024). *Frontiers of Artificial Intelligence, Ethics, and Multidisciplinary Applications*. Springer Nature Singapore.
10. Deng, K. (2021). Anomaly Detection of Highway Vehicle Trajectory under the Internet of Things Converged with 5G Technology. *Complexity*. Wiley. https://doi.org/10.1155/2021/9961428
11. Obolensky, N., Singer, M., & Inada-Kim, M. (2010). Development of a patient-centred escalation plan for recognition and management of deteriorating patients. *Acute Medicine*, 9(3), 116-121.
12. Philips Healthcare. (2023). *IntelliVue Guardian Solution: Early Warning Scoring and Smart Alerts*. White paper. Philips Medical Systems.