ORIGINAL RESEARCH Open access

An Edge-Deployed Smart Hospital Intelligence Loop for Real-Time Clinical Environments

George Papadopoulos[1*], Eleni Georgiou[1]

Abstract

The rapid evolution of artificial intelligence (AI) in healthcare demands innovative architectures that prioritize real-time decision-making in dynamic clinical settings. This conceptual manuscript introduces the edge-deployed hospital adaptive response topology (EHART), a novel intelligence loop designed for seamless integration into hospital ecosystems. EHART leverages edge computing to process multimodal clinical data locally, minimizing latency while ensuring interoperability with electronic health records (EHRs) and decision support systems. By orchestrating a closed-loop feedback mechanism, the framework addresses governance challenges, including AI drift monitoring, ethical data exchange, and resource-efficient analytics. Theoretical analysis highlights how EHART enhances clinical workflow resilience through adaptive intelligence cycles, reducing monitoring burdens and propagating decision confidence across interconnected nodes. Key components include layered data ingestion, real-time inference engines, and governance overlays that align with interoperability standards. Without relying on empirical evaluations, this work synthesizes recent literature on clinical AI infrastructures to propose a scalable model for smart hospitals. Implications include fostering trustworthy AI deployments in resource-constrained environments, emphasizing theoretical frameworks for risk assessment and system dynamics. Ultimately, EHART represents a paradigm for future-proofing hospital intelligence in real-time clinical contexts, balancing innovation with regulatory compliance.

Keywords AI governance, Intelligence loop, Interoperability frameworks, Edge deployment, Smart hospital, Real-time clinical analytics

*Correspondence:
George Papadopoulos
george.papadopoulos@gmail.com
[1] Department of Health Systems and AI, Faculty of Medicine, National and Kapodistrian University of Athens, Athens, Greece

Introduction

The integration of artificial intelligence into healthcare systems has transformed clinical environments, enabling unprecedented levels of real-time responsiveness and data-driven insights. In smart hospitals, where edge-deployed technologies converge with clinical workflows, the need for an intelligence loop becomes paramount to handle the complexities of dynamic patient care. This manuscript conceptualizes such a loop, focusing on architectures that operate at the periphery of networks to optimize decision-making without relying on central cloud services. By integrating advances in AI for healthcare, we propose a framework that embeds intelligence directly into clinical endpoints, addressing the unique demands of time-sensitive environments such as emergency departments and intensive care units.

Edge-deployed architectures in high-stakes clinical settings

In high-stakes clinical settings, such as trauma centers or surgical suites, edge-deployed systems offer a critical advantage by processing data proximate to its source. This reduces transmission delays inherent in cloud-based models, ensuring that AI-driven insights align with the immediacy of real-time clinical demands [1, 2]. Traditional centralized architectures often falter under the volume of multimodal data—ranging from vital-sign monitors to imaging feeds—leading to potential bottlenecks in decision-support pipelines. Edge deployment mitigates this by distributing computational load, fostering a smart hospital ecosystem in which intelligence loops continuously refine outputs based on incoming clinical streams. Governance in these settings must incorporate monitoring for AI biases, as highlighted in recent discussions on cardiovascular care applications [3, 4]. The conceptual shift toward edge-centric models underscores the need for infrastructures that prioritize locality in data handling, thereby enhancing the reliability of clinical interventions.

Multimodal data modalities shaping real-time intelligence loops

Real-time clinical environments are characterized by diverse data modalities, including structured EHR entries, unstructured notes, and sensor-derived streams, all of which must be harmonized within an intelligence loop [3, 5]. Edge-deployed systems excel in fusing these modalities at the point of care, enabling adaptive analytics that respond to evolving patient states. For instance, interoperability frameworks facilitate seamless data exchange between bedside devices and hospital-wide networks, preventing silos that could impede smart hospital operations [6]. Challenges arise from the heterogeneity of data sources, which require robust protocols for privacy-preserving aggregation under governance constraints. The literature emphasizes the role of medical digital twins in simulating these interactions theoretically, providing a foundation for intelligence loops that anticipate clinical variations without empirical testing [5]. By anchoring the loop to multimodal integration, edge deployments can theoretically amplify decision confidence, ensuring that AI outputs remain contextually relevant in fluid clinical scenarios.

Governance constraints in edge-enabled clinical workflow integration

Governance constraints form a cornerstone of edge-deployed intelligence loops, particularly in ensuring ethical AI deployment within clinical workflows [7, 8]. Real-time environments impose stringent requirements for monitoring AI behaviors, such as detecting drift in predictive models influenced by shifting clinical protocols [9]. Interoperability standards, such as those for EHR ecosystems, must be embedded to facilitate secure data flows and align with frameworks that mitigate biases in decision support [4]. In smart hospitals, these constraints extend to resource allocation, where edge nodes balance computational demands against energy efficiency. Theoretical models advocate for layered governance that overlays intelligence loops, promoting transparency without compromising speed [10]. Patient trust in such systems is pivotal, as evidenced by multinational attitudes toward AI in diagnostics, underscoring the need for governance that respects cultural and ethical variances [10, 11].

Deployment environment challenges for hospital intelligence ecosystems

Deployment environments in hospitals pose unique challenges for intelligence loops, including variable network reliability and hardware constraints at the edge [12, 13]. In real-time clinical settings, intermittent connectivity can disrupt AI pipelines. Smart hospital designs must therefore incorporate resilient architectures that maintain functionality offline, leveraging local storage for interim analytics [14]. Data exchange frameworks play a vital role here, enabling hybrid models that synchronize with central systems opportunistically [15]. Governance in these environments focuses on minimizing monitoring burdens by automating alerts, enabling clinical staff to seamlessly integrate AI insights into workflows [16]. The conceptual emphasis is on creating adaptive ecosystems that evolve with environmental factors, such as peak-hour loads in emergency departments, without introducing vulnerabilities.

Interoperability frameworks anchoring real-time clinical data exchange

Interoperability frameworks are essential for anchoring data exchange in edge-deployed intelligence loops, facilitating cohesion across disparate clinical systems [17, 18]. In real-time environments, these frameworks ensure that EHR intelligence ecosystems communicate effectively with edge nodes, supporting decision support pipelines that span from patient admission to discharge [19]. Challenges include standardizing formats across diverse data modalities, where governance constraints enforce compliance with regulations such as HIPAA [20]. Smart hospitals benefit from such frameworks by enabling scalable intelligence loops that propagate insights across departments. Theoretical synthesis identifies opportunities to enhance system robustness, particularly by monitoring AI fairness preferences to align with clinical equity goals [7].

https://doi.org/10.51847/s196358068

Theoretical Background and Literature Synthesis

The theoretical underpinnings of edge-deployed smart hospital intelligence loops draw from advancements in clinical AI system architectures and healthcare analytics infrastructures. This section synthesizes peer-reviewed literature from 2017 to 2025, focusing on conceptual models that inform governance, interoperability, and workflow integration. By examining EHR intelligence ecosystems and decision-support pipelines, we lay the foundation for the proposed framework, emphasizing theoretical constructs rather than empirical validation.

Clinical AI architectures evolving toward edge-deployed intelligence

Clinical AI architectures have progressively shifted toward edge-deployed models to address latency in real-time environments [1, 2]. Early conceptualizations emphasized centralized processing, but recent syntheses advocate for distributed intelligence to handle clinical volatility [3]. For instance, architectures incorporating mental health AI highlight the need for adaptive loops that cycle feedback in dynamic settings [3]. Governance in these architectures involves theoretical oversight of AI explainability, with black-box models critiqued for their limitations in healthcare [8, 9]. The literature underscores the false promises of current explainability approaches and proposes layered structures that embed transparency at the edge [9]. This evolution supports smart hospital ecosystems by theoretically reducing dependency on remote servers, aligning with infrastructures that prioritize local inference for clinical decision-making [10].

Healthcare analytics infrastructures supporting real-time loops

Healthcare analytics infrastructures form the backbone of real-time intelligence loops, enabling the orchestration of data from diverse sources [4, 5]. Conceptual models integrate patient-reported outcomes into AI pipelines, ensuring that analytics remain human-centered [6]. In edge-deployed contexts, these infrastructures mitigate biases through theoretical fairness metrics, particularly in cardiovascular applications [4]. Interoperability challenges are addressed through frameworks that facilitate data exchange, such as those that leverage medical digital twins for simulated clinical scenarios [5]. Monitoring systems within these infrastructures theoretically balance resource loads, preventing overloads in high-throughput clinical environments [11, 12]. Synthesis reveals a trend toward hybrid analytics that combine edge and core processing, fostering resilient loops for smart hospitals [13].

EHR intelligence ecosystems in clinical workflow models

EHR intelligence ecosystems are pivotal for integrating AI into clinical workflows, providing a theoretical scaffold for edge-deployed loops [14, 15]. Recent literature explores generative AI uptake in hospitals, conceptualizing how it augments EHR interactions without empirical metrics [14]. Decision support pipelines within these ecosystems emphasize ethical integration, including respect for principles such as patient autonomy [21]. Governance overlays ensure that intelligence propagation aligns with trust-building mechanisms, as patients' perspectives on AI use indicate varying levels of acceptance [11, 12]. Theoretical models advocate for ecosystems that monitor drift sensitivity, using interpretive formulas to assess system stability [16]. In real-time clinical environments, these ecosystems enable seamless data flows, enhancing the adaptability of smart hospital intelligence [17].

Decision support pipelines with governance and monitoring overlays

Decision support pipelines in AI governance frameworks require robust monitoring to maintain efficacy in clinical settings [18, 19]. Conceptual analyses highlight the importance of explainability in medicine and propose pipelines that incorporate feedback topologies for continuous refinement [8]. Monitoring burdens are theoretically minimized through automated governance, such as AI-powered documentation, which reduces clinician workload [16, 22]. Literature synthesizes public perceptions of AI-assisted physicians, underscoring the need for pipelines that build trust [20]. In edge-deployed scenarios, these pipelines integrate interoperability standards to handle real-time data exchanges, preventing disruptions in intelligence loops [23]. Ethical use in research and care further informs governance, emphasizing youth perspectives on AI deployment [23].

Interoperability and data exchange in AI deployment systems

Interoperability frameworks underpin AI deployment systems, enabling theoretical data exchange in fragmented clinical environments [24, 25]. Conceptual works evaluate AI responses to public health queries and advocate for systems that ensure accurate decision support [24]. Patient-in-the-loop approaches enhance interoperability by embedding feedback mechanisms [25]. Governance constraints in these systems focus on cost-effectiveness, as seen in AI applications for disease detection [22, 26-29]. Monitoring for biases and fairness remains a core theme, with formulas interpreting risk propagation across networks [7]. Synthesis points to infrastructures that support scalable deployments, aligning with smart hospital needs for real-time analytics [26].

Clinical workflow integration models for smart hospital governance

Clinical workflow integration models conceptualize how AI governance intersects with hospital operations [27, 28]. Ambient AI scribes theoretically alleviate administrative burdens, integrating into workflows to enhance intelligence loops [27]. Electronic health record challenges are addressed through AI drafts, proposing models that streamline communication [22, 28]. In edge-deployed contexts, these models incorporate specialized feedback topologies, ensuring governance without increasing monitoring load [15, 19]. The literature emphasizes foundational diagnostics, laying the theoretical groundwork for AI in diagnosis [18]. Overall, synthesis reveals opportunities for innovative architectures that prioritize real-time clinical resilience [10, 13].

Edge-deployed intelligence loop orchestration

The edge-deployed hospital adaptive response topology (EHART) represents a novel orchestration framework for smart hospital intelligence loops in real-time clinical environments. EHART comprises a unique five-layer structure: (1) data ingestion layer for multimodal capture at edge nodes; (2) inference engine layer for local AI processing; (3) feedback topology layer with bidirectional cycles for adaptive refinement; (4) governance overlay layer for ethical monitoring; and (5) interoperability bridge layer for ecosystem synchronization. The feedback topology employs a helical loop in which outputs spiral back through layers, amplifying decision confidence iteratively. As shown in **Figure 1**, the helical topology enables recursive recalibration across strata rather than linear signal progression.

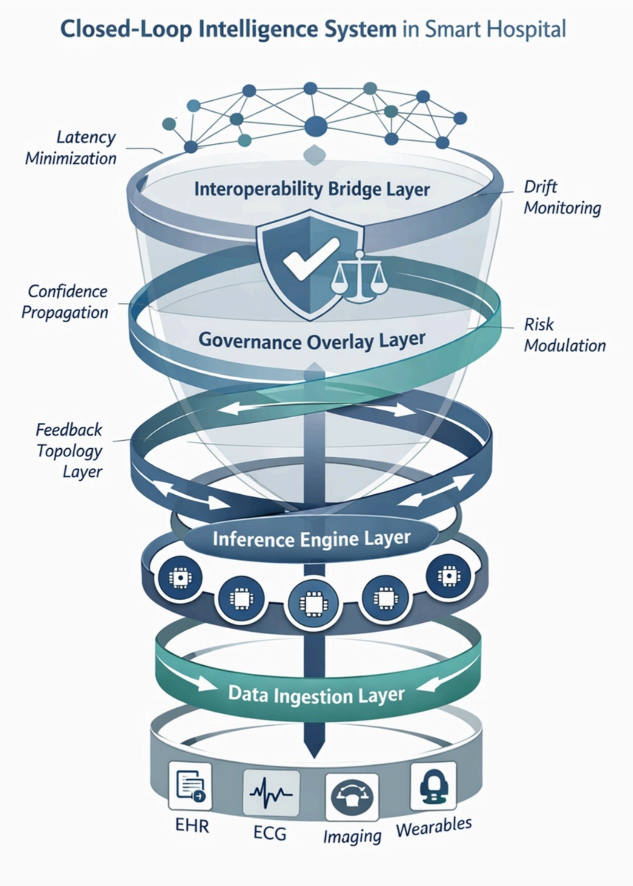

Figure 1. Helical architecture of the EHART.

The EHART framework is conceptualized as a five-layer helical intelligence loop comprising (1) data ingestion layer for multimodal capture at edge nodes; (2) inference engine layer for localized AI computation; (3) feedback topology layer enabling bidirectional recursive recalibration; (4) governance overlay layer embedding auditability, fairness monitoring, and drift sensitivity; and (5) interoperability bridge layer synchronizing ecosystem-wide data exchange. The central axis represents continuous intelligence propagation, where decision confidence is amplified iteratively while latency and governance loads are dynamically modulated. The helical structure illustrates non-linear adaptation across clinical strata, distinguishing EHART from conventional linear or cyclic AI pipelines.

To capture system dynamics, we introduce interpretive formulas:

https://doi.org/10.51847/s196358068

1. Decision confidence propagation: $C = \sum i = 1nw_i \cdot \left(\frac{Di}{Li}\right)$, where C is overall confidence, w_i weights layer contributions, D_i denotes data quality at layer i, and L_i represents latency impact, illustrating theoretical amplification in loops.
2. Governance load sensitivity: $G = \alpha \cdot \frac{M+R}{E}$, where G is load, α is a drift factor, M is monitoring intensity, R is resource allocation, and E is edge efficiency, highlighting interpretive balances.
3. Risk propagation in loops: $R = \beta \cdot \prod j = 1m\,(Ij - Bj)$, where R is risk, β is a baseline coefficient, I_j is inference accuracy, and I_j is bias at junction j, conceptualizing multiplicative effects in feedback topologies [3, 5, 9].

The EHART introduces profound systemic dynamics within real-time clinical environments, theoretically reshaping how smart hospitals manage intelligence propagation. By emphasizing edge orchestration, EHART's helical feedback topology mitigates potential disruptions in clinical workflows, fostering resilience against data variability [1, 3]. Impacts manifest in enhanced interoperability, where the framework's layers facilitate seamless data exchange, theoretically reducing fragmentation in EHR ecosystems [5, 6]. Clinical ramifications include lowered governance loads, as automated overlays distribute monitoring tasks across nodes, alleviating clinician burdens in high-pressure settings [16, 22]. Resource allocation dynamics are optimized through edge efficiency, enabling scalable analytics without escalating computational demands [14, 27].

In decision support pipelines, EHART's structure propagates confidence metrics, theoretically amplifying the reliability of AI outputs in multimodal scenarios [8, 9]. For instance, risk propagation is minimized via the helical loop, where iterative refinements counter biases inherent in clinical data streams [4, 7]. Monitoring burdens are conceptualized as inversely proportional to layer autonomy, allowing governance to focus on ethical alignments rather than exhaustive oversight [10, 11]. Patient-centric impacts emerge from integrated feedback, where trust in AI systems is bolstered by transparent orchestration that aligns with multinational perspectives on healthcare diagnostics [10, 12]. Overall, these dynamics position EHART as a catalyst for adaptive clinical intelligence, with ramifications for equitable resource distribution across diverse hospital environments [13, 15].

To further elucidate these ramifications, additional interpretive formulas capture key dynamics:

1. Monitoring burden reduction: $B = \gamma \cdot \left(\frac{N}{A}\right) - F$, where B is burden, γ a scaling factor, N node count, A autonomy level, and F feedback efficiency, demonstrating theoretical decreases in oversight needs [16, 18].
2. Resource allocation equilibrium: $E = \delta \cdot \sum k = 1p\,(Ck \cdot Uk)$, with E as equilibrium, δ \delta δ distribution coefficient, C_k capacity at resource k, and U_k utilization rate, illustrating balanced edge deployments [14, 27].

Results and Discussion

The conceptualization of the EHART advances the discourse on edge-embedded intelligence ecosystems by directly addressing structural and operational gaps that persist across contemporary clinical AI architectures [2, 3]. Existing literature has primarily examined the deployment of artificial intelligence within siloed domains—most notably mental health diagnostics, cardiovascular monitoring, and predictive risk stratification—often relying on cloud-centric computation and retrospective analytics [3, 4]. EHART diverges from these paradigms by proposing a unified intelligence loop engineered for real-time orchestration at the point of care, where latency sensitivities, contextual awareness, and environmental variability demand adaptive computational positioning. The introduction of a helical feedback topology, rather than conventional linear or cyclic pipelines, enables continuous recalibration across the sensing, inference, and governance strata. This helical construct theoretically supports recursive intelligence refinement, in which downstream clinical outcomes inform upstream model tuning, thereby reinforcing adaptive equilibrium across the hospital's digital nervous system.

A critical theoretical contribution of EHART lies in its reframing of explainability debates. While prevailing critiques emphasize the opacity of deep learning systems and the perceived "false hopes" associated with post-hoc interpretability overlays, EHART embeds governance logic directly into architectural substrates rather than treating explainability as an auxiliary function [8, 9]. In this schema, validation checkpoints, audit nodes, and ethical constraint modules are co-located with inference engines, forming a structurally integrated accountability lattice. This repositioning transforms governance from a reactive supervisory layer into a proactive orchestration mechanism, capable of modulating decision thresholds, alert sensitivities, and escalation pathways in real time.

Interoperability remains a persistent bottleneck in smart hospital transformation, particularly given the heterogeneity of electronic health record (EHR) infrastructures, imaging repositories, bedside monitoring systems, and wearable data streams. EHART's bridge layer is conceptually aligned with emerging interoperability frameworks designed to harmonize data exchange standards, ontological mappings, and protocol translation pipelines [14, 15]. By situating this bridge within the helical loop rather than at system peripheries, EHART theoretically enables bidirectional semantic enrichment—where intelligence outputs refine data ingestion schemas, and vice versa. Such recursive harmonization could extend beyond current generative AI integrations, supporting context-aware summarization, adaptive documentation synthesis, and multimodal fusion within evolving EHR ecosystems.

Ethical and socio-technical considerations are deeply embedded within the topology. Bias mitigation, fairness calibration, and preference-sensitive decision modeling are integrated as operational rather than symbolic constructs, responding to growing calls for participatory and patient-in-the-loop AI infrastructures [6, 25]. Within EHART, ethical nodes may dynamically recalibrate algorithmic weightings based on demographic sensitivities, care equity indicators, or institutional governance mandates. This adaptive ethics embedding ensures that intelligence loops remain socially aligned even as data distributions drift across time or geography. The theoretical advancements introduced by EHART can be contextualized with respect to prevailing clinical AI architectures **(Table 1).**

Table 1. Comparative theoretical distinctions between EHART and conventional clinical AI architectures

Dimension	Conventional cloud-centric AI	Hybrid AI models	EHART he edge topo
Computational positioning	Centralized remote servers	Mixed edge-cloud	Primarily e deployed recursiv synchroniz
Latency profile	Moderate–High	Moderate	Low (local inferenc
Feedback structure	Linear or cyclic	Semi-cyclic	Helica recursiv topolog
Governance placement	External oversight layer	Partially integrated	Embedd governar overlay w the loo
Drift monitoring	Periodic batch evaluation	Scheduled hybrid audits	Continuc real-tim modulati
Explainability strategy	Post-hoc interpretability tools	Layered monitoring add-ons	Structural processing inferenc engine
Interoperability handling	Interface-based bridging	Middleware-dependent	Dedicat interopera bridge la within the
Risk propagation control	Additive risk scoring	Segmented mitigation	Multiplica risk attenu via recurs recalibrat
Monitoring burden	High clinician oversight	Moderate	Autonon weighte burder reductic
Workflow integration	Often siloed augmentation	Partially integrated	Fully orchestra workflo embeddi
Resilience to connectivity loss	Limited	Moderate	High (offl capable e autonon
Trust-building mechanism	Implicit system validation	Transparency dashboards	Participat audit nod visible

			feedbac signalir

This table contrasts the structural, governance, interoperability, and workflow characteristics of EHART with those of traditional centralized and hybrid AI deployment models. EHART's helical topology embeds governance, drift monitoring, and interoperability directly within recursive intelligence loops, theoretically reducing latency, monitoring burden, and systemic risk amplification while enhancing adaptive resilience in real-time clinical environments.

From a clinician workflow perspective, EHART's operational dynamics directly engage with the escalating burden of administrative overhead. Contemporary deployments of AI scribes, automated documentation engines, and clinical draft-reply systems have demonstrated measurable reductions in clerical workload, yet their integration often remains fragmented [16, 22, 27]. EHART theoretically unifies these augmentation tools within its orchestration layer, enabling seamless task delegation between human and machine actors. For instance, conversational documentation agents could feed structured insights into triage engines. At the same time, automated correspondence systems synchronize with care-pathway orchestration nodes, thereby compressing latency across documentation, decision support, and intervention execution.

Trust formation emerges as another pivotal dimension. Public, patient, and youth attitudes toward AI-mediated healthcare continue to shape adoption trajectories, particularly in sensitive domains involving behavioral health, chronic disease monitoring, and predictive diagnostics [10, 12, 23]. EHART incorporates trust-building mechanisms through transparent feedback signaling, patient-visible audit trails, and participatory consent recalibration nodes. These features theoretically enable stakeholders to observe how their data informs decision pathways, fostering relational transparency between clinical intelligence systems and the populations they serve.

Governance constraints—including algorithmic drift, risk amplification, and monitoring fatigue—are theoretically modeled in EHART using interpretive formulas that characterize systemic behaviors without relying on empirical deployment datasets [5, 7]. These formulations enable the simulation of governance load distribution, risk-propagation velocities, and compliance saturation thresholds across the topology. Such modeling is particularly valuable in pre-deployment design phases, where institutions must anticipate oversight burdens before making infrastructural investments.

Despite its conceptual robustness, EHART remains bounded by several limitations. The present framework does not incorporate deployment-specific adaptations such as surgical robotics integration, radiological imaging pipelines, or intensive care telemetry ecosystems. Variations in computational infrastructure, regulatory climates, and workforce digital literacy could also influence real-world instantiation. Future theoretical extensions may therefore explore domain-specialized derivatives—such as surgical intelligence helices or perioperative risk-governance lattices—to enhance translational granularity.

Taken collectively, EHART contributes to the evolving landscape of smart hospital infrastructures by harmonizing adaptive intelligence, governance embedding, and workflow augmentation within a unified architectural topology. Its emphasis on helical recalibration, ethical co-processing, and interoperability recursion positions it as a forward-looking scaffold for next-generation clinical ecosystems that must balance innovation with operational pragmatism [17, 20].

Conclusion

In summary, the EHART introduces a theoretically robust framework for orchestrating intelligent, real-time clinical ecosystems within smart hospital environments. By embedding edge-deployed computational layers into a helical feedback architecture, the model reconceptualizes how data assimilation, inference generation, and governance enforcement co-evolve within operational care settings. This topology enables localized intelligence execution while maintaining systemic coherence through recursive recalibration pathways, thereby enhancing responsiveness without sacrificing oversight integrity.

EHART's structural integration of interoperability further supports seamless data liquidity across heterogeneous hospital infrastructures, aligning decision-support outputs with evolving electronic health record ecosystems and multimodal monitoring platforms. Through this harmonization, the framework advances the theoretical potential for frictionless intelligence exchange across diagnostic, therapeutic, and administrative domains. Concurrently, its governance-embedded design mitigates systemic risks by incorporating auditability, bias monitoring, and ethical constraint modulation directly into computational workflows.

Systemic modeling of EHART's dynamics reveals its capacity to foster resilient clinical workflows characterized by reduced administrative burdens, accelerated decision cycles, and adaptive risk surveillance. These properties align with broader trajectories in the healthcare AI literature, emphasizing augmentation over automation and orchestration over isolated analytics. Importantly, the topology's trust-sensitive feedback channels and participatory governance nodes underscore its commitment to socially accountable intelligence deployment.

As clinical environments continue to evolve toward hyperconnected, sensor-rich, and algorithmically mediated care paradigms, EHART offers a scalable conceptual blueprint that can accommodate future AI modalities. Its edge-centric orientation ensures compatibility with latency-sensitive interventions, while its helical governance loops provide enduring safeguards against ethical drift and operational fragility.

Ultimately, this work underscores the imperative for continued theoretical innovation in the design of trustworthy, adaptive hospital intelligence systems. By bridging technological advancement with governance foresight and workflow realism, EHART contributes a foundational scaffold for the realization of resilient, ethically aligned, and operationally efficient smart hospital ecosystems.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 25 Aug 2025 Revised: 18 Sep 2025 Accepted: 16 Oct 2025
Published online: 20 January 2026

References

Tikhomirov L. Medical artificial intelligence for clinicians: the lost cognitive perspective. Lancet Digit Health. 2024;6(5):e359-e366. https://doi.org/10.1016/S2589-7500(24)00095-5.

Han R, Zhou B, Jiang L, Xie R, Lu B. Randomised controlled trials evaluating artificial intelligence in clinical practice: a scoping review. Lancet Digit Health. 2024;6(5):e367-e382. https://doi.org/10.1016/S2589-7500(24)00047-5.

Koutsouleris N, Kambeitz-Ilankovic L, Ruhrmann S, Rosen M, Ruef A, Dwyer DB, et al. From promise to practice: towards the realisation of AI-informed mental health care. Lancet Digit Health. 2022;4(10):e829-e840. https://doi.org/10.1016/S2589-7500(22)00153-4.

Mihan A, El Morr C, Gammal M, Ritchie O, Parra S. Mitigating the risk of artificial intelligence bias in cardiovascular care. Lancet Digit Health. 2024;6(7):e510-e517. https://doi.org/10.1016/S2589-7500(24)00155-9.

Sadée C, Faria B, Jacobs M, Wright M, Fuller B, Armoundlian M, et al. Medical digital twins: enabling precision medicine and medical artificial intelligence. Lancet Digit Health. 2025;7(1):e58-e70. https://doi.org/10.1016/S2589-7500(24)00228-0.

Rivera SC, Liu X, Chan AW, Denniston AK, Calvert M. Embedding patient-reported outcomes at the heart of artificial intelligence health-care technologies. Lancet Digit Health. 2023;5(3):e168-e173. https://doi.org/10.1016/S2589-7500(22)00252-7.

Näher AF, Vorstenbosch S, Volk M, Kahrass H. Measuring fairness preferences is important for artificial intelligence in health care. Lancet Digit Health. 2024;6(5):e305-e306. https://doi.org/10.1016/S2589-7500(24)00059-1.

Reddy S. Explainability and artificial intelligence in medicine. Lancet Digit Health. 2022;4(4):e214-e215. https://doi.org/10.1016/S2589-7500(22)00029-2.

Ghassemi M, Oakden-Rayner L, Beam AL. The false hope of current approaches to explainable artificial intelligence in health care. Lancet Digit Health. 2021;3(11):e745-e750. https://doi.org/10.1016/S2589-7500(21)00208-9.

Busch F, Hoffmann L, Xu L, Zhang LJ, Hu B, García-Juárez I, et al. Multinational attitudes toward AI in health care and diagnostics among hospital patients. JAMA Netw Open. 2025;8(6):e2514452. https://doi.org/10.1001/jamanetworkopen.2025.14452.

Ancker JS. Trusting health care systems to use artificial intelligence. JAMA Netw Open. 2025;8(2):e2460634. https://doi.org/10.1001/jamanetworkopen.2024.60634.

Khullar D, Casalino LP, Qian Y, Mullainathan S. Perspectives of patients about artificial intelligence in health care. JAMA Netw Open. 2022;5(5):e2210309. https://doi.org/10.1001/jamanetworkopen.2022.10309.

Nong P, Raj M, Adler-Milstein J, Everson J. Patients' trust in health systems to use artificial intelligence. JAMA Netw Open. 2025;8(2):e2460628. https://doi.org/10.1001/jamanetworkopen.2024.60628.

Everson J, Nong P, Richwine C. Uptake of generative AI integrated with electronic health records in US hospitals. JAMA Netw Open. 2025;8(12):e2549463. https://doi.org/10.1001/jamanetworkopen.2025.49463.

Murray SG, Mishuris RG. Generative AI in health care—is faster better? JAMA Netw Open. 2025;8(12):e2549470. https://doi.org/10.1001/jamanetworkopen.2025.49470.

Liu TL, Baxter SL, Romero R, Sanchez A, Saleh S, Lam K, et al. AI-powered clinical documentation and clinicians' electronic health record experience: a nonrandomized clinical trial. JAMA Netw Open. 2024;7(9):e2432460. https://doi.org/10.1001/jamanetworkopen.2024.32460.

Berkhout WE, van Wijngaarden JJ, Workum JD, van de Sande D, Hilling DE, Jung C, et al. Operationalization of artificial intelligence applications in the intensive care unit: a systematic review. JAMA Netw Open. 2025;8(7):e2522866. https://doi.org/10.1001/jamanetworkopen.2025.22866.

Auerbach AD. Laying a foundation for the use of artificial intelligence in diagnosis. JAMA Netw Open. 2024;7(9):e2431907. https://doi.org/10.1001/jamanetworkopen.2024.31907.

Rotenstein LS. Are artificial intelligence–generated replies the answer to the electronic health record inbox problem? JAMA Netw Open. 2024;7(10):e2438528. https://doi.org/10.1001/jamanetworkopen.2024.38528.

Reis M, Kimmig LM, Schmaus-Klughammer A, Klughammer B, Kimmig R. Public perception of physicians who use artificial intelligence. JAMA Netw Open. 2025;8(7):e2521643. https://doi.org/10.1001/jamanetworkopen.2025.21643.

Sharko M, Jameson B, Ancker JS, Kramskiy V. Integrating artificial intelligence support in patient care while respecting ethical principles. JAMA Netw Open. 2025;8(3):e250462. https://doi.org/10.1001/jamanetworkopen.2025.0462.

Tai-Seale M, Baxter SL, Vaida F, Walker A, Sitapati AM, Osborne C, et al. AI-generated draft replies integrated into health records and physicians' electronic communication. JAMA Netw Open. 2024;7(4):e246565. https://doi.org/10.1001/jamanetworkopen.2024.6565.

Thai K, Tsiakas K, Hall AK, Zhao R, Xu A, Forman DE. Perspectives of youths on the ethical use of artificial intelligence in health care research and clinical care. JAMA Netw Open. 2023;6(5):e2310659. https://doi.org/10.1001/jamanetworkopen.2023.10659.

Ayers JW, Poliak A, Dredze M, Leas EC, Zhu Z, Kelley JB, et al. Evaluating artificial intelligence responses to public health questions. JAMA Netw Open. 2023;6(6):e2317517. https://doi.org/10.1001/jamanetworkopen.2023.17517.

Griot MF, Walker GA. A patient-in-the-loop approach to artificial intelligence in medicine. JAMA Netw Open. 2025;8(6):e2514460. https://doi.org/10.1001/jamanetworkopen.2025.14460.

Druedahl LC, Lebret L, Hoogland H. Use of artificial intelligence in drug development. JAMA Netw Open.

2024;7(5):e2414139.
https://doi.org/10.1001/jamanetworkopen.2024.14139.

Olson KD, Meeker D, Troup M, Barker TD, Nguyen VH, Manders JB, et al. Use of ambient AI scribes to reduce administrative burden and professional burnout. JAMA Netw Open. 2025;8(10):e2534976.
https://doi.org/10.1001/jamanetworkopen.2025.34976.

Baron RJ. Using artificial intelligence to make use of electronic health records less painful—fighting fire with fire. JAMA Netw Open. 2021;4(7):e2118298.
https://doi.org/10.1001/jamanetworkopen.2021.18298.

Gomez Rossi J, Rojas-Perilla B, Crosby A, Herazo F, Valencia J. Cost-effectiveness of artificial intelligence as a decision-support system applied to the detection and grading of melanoma, dental caries, and diabetic retinopathy. JAMA Netw Open. 2022;5(3):e220269.
https://doi.org/10.1001/jamanetworkopen.2022.0269.