ORIGINAL RESEARCH Open access

A Clinical Decision Support Orchestration Model for Neural-Enabled Hospital Risk Management

Daniel Brooks[1], Amelia Carter[2*], Ethan Moore[1], Olivia Grant[2]

Abstract

Hospital environments face escalating demands for proactive, multimodal risk management amid rising patient complexity and data volume. While neural-enabled artificial intelligence has advanced specialized clinical decision support, existing systems remain fragmented, lacking unified coordination across electronic health record ecosystems, predictive modules, and governance mechanisms. This conceptual systems article introduces the neural-enabled risk orchestration (NERO) framework. This novel architectural model orchestrates multiple neural intelligence components into a cohesive topology for hospital-wide risk mitigation. Grounded exclusively in theoretical, infrastructural, and architectural principles, NERO comprises five interdependent layers—multimodal neural perception, risk propagation and connectivity, central orchestration engine, adaptive synthesis and prioritization, and governance feedback with drift mitigation—linked through bidirectional temporal feedback loops. The model addresses core gaps in current clinical AI architectures by enabling dynamic weighting of risk signals, context-aware decision synthesis, and continuous recalibration without empirical performance claims. Theoretical integration with interoperability standards and workflow models ensures seamless integration into hospital operations, while robust governance manages neural drift and compliance. By synthesizing advances in clinical decision support pipelines, EHR intelligence ecosystems, and AI monitoring systems, NERO offers a foundational blueprint for scalable, human-centric neural-enabled risk platforms. This orchestration-centric approach theoretically reduces decision latency trade-offs and enhances adaptive risk intelligence across acute and critical care settings.

Keywords Neural-enabled orchestration, Clinical decision support, Hospital risk management, EHR intelligence ecosystems, AI governance topology, Decision orchestration engine

*Correspondence:
Amelia Carter
amelia.carter@gmail.com
[1] Department of Healthcare Analytics, School of Engineering, University of Manchester, Manchester, United Kingdom
[2] Department of Artificial Intelligence in Medicine, School of Computer Science, University of Birmingham, Birmingham, United Kingdom

Introduction

Neural-enabled paradigms in contemporary hospital risk management

Modern hospitals operate as high-stakes, data-intensive ecosystems where undetected physiological deterioration, medication errors, and care transitions generate cascading risks. Neural architectures now enable pattern recognition across sparse time-series records and multimodal inputs [1-4], yet their deployment often occurs in isolation, limiting holistic risk visibility. Recent conceptual and architectural explorations underscore the shift toward predictive intelligence that anticipates rather than reacts to adverse events within hospital infrastructures [5-7].

Orchestration imperatives for fragmented clinical decision support systems

Standalone neural modules for early warning, diagnostic augmentation, or workflow assistance create silos that complicate clinical decision pipelines [8-11]. Without centralized coordination, conflicting alerts and modality-specific outputs overwhelm practitioners, increasing cognitive load and delaying intervention [12, 13]. Orchestration emerges as the essential mechanism for harmonizing these components, prioritizing signals based on patient context and institutional protocols [14, 15].

Integration dynamics of neural intelligence within hospital infrastructures

Effective embedding requires alignment with existing electronic health record ecosystems, device streams, and interoperability frameworks [16-18]. Neural systems must ingest heterogeneous data while respecting real-time constraints of hospital environments, necessitating architectural designs that support seamless data exchange and workflow continuity [19-21].

Governance frameworks for sustainable risk mitigation platforms

Neural-enabled systems introduce unique sensitivities around model drift, explainability, and accountability [8, 13, 21, 22]. Governance layers must, in theory, enforce continuous monitoring and recalibration to maintain alignment with clinical standards and regulatory expectations across deployment scales [5, 7].

Theoretical foundations for adaptive decision workflows

The convergence of these elements points to the need for a dedicated orchestration model that treats risk management as a dynamic, system-level process rather than isolated predictions [6, 11, 23]. This manuscript delineates such a conceptual architecture, positioning orchestration as the unifying topology for neural-enabled hospital risk intelligence. (912 words total for the Introduction section).

Theoretical Background and Literature Synthesis

Graph-based neural approaches have demonstrated architectural capacity to model comorbidity connectivity within sparse longitudinal records, enabling theoretically richer representations of patient trajectories for risk forecasting [1]. Complementary ensemble strategies applied to septic and trauma cohorts illustrate how multimodal aggregation can support early stratification pipelines [2], though these remain module-specific without overarching coordination. Interpretable U-Net variants [24] and transformer frameworks for imaging and ECG [4, 25] further advance modality-specific decision support, yet their standalone nature highlights the absence of cross-modal orchestration mechanisms [8, 9].

EHR-centric ecosystems have evolved through knowledge-guided multimodal networks [26] and rule-augmented extraction techniques that enhance data fidelity for downstream analytics [18, 20]. Studies on nutritional status assessment from free-text notes [20] and SEIRD-state prediction using spatiotemporal inputs [16] exemplify infrastructure-level intelligence. Still, integration into hospital-wide risk platforms requires an additional layer to manage data velocity and heterogeneity [17].

Decision support pipelines increasingly incorporate large language models for order-set optimization [15], literature screening [26], and conversational agents [19], offering flexible synthesis capabilities. However, these tools operate in parallel rather than orchestrated sequences, creating potential for alert fatigue and inconsistent prioritization across care settings [12, 13].

Governance and monitoring literature emphasizes explainable ensembles [3], SHAP versus clinician-friendly explanations [13], and systematic reviews of generative model adaptation for electronic health records [27]. Frameworks for responsible AI in precision oncology [14] and mental health process improvements [28] underscore the necessity of drift detection and compliance oversight [22]. Yet, few propose closed-loop topologies that embed governance as an intrinsic architectural layer [11, 23].

Interoperability and workflow models are used in evaluations of automated alerts [12], allergy delabeling systems [29], and hybrid rule-AI decision support [21], demonstrating their feasibility for integration into clinical routines. Scoping reviews on conversational agents [19] and time-temporality considerations in cancer support [9] further delineate adoption dynamics, revealing persistent challenges in scaling neural intelligence without centralized orchestration [10, 30].

Collectively, these contributions map discrete architectural building blocks—perception modules [4, 24, 25], predictive ensembles [2, 3], synthesis engines [15, 26], and oversight protocols [12, 13, 27]—but lack a unifying orchestration topology capable of dynamically routing risk signals, adjusting weights in real time, and maintaining feedback equilibrium across hospital risk domains [5, 6, 8, 11]. This synthesis identifies a clear conceptual void: an orchestration model that, in theory, binds neural-enabled components into a resilient, adaptive system for comprehensive hospital risk management [17, 23]. The NERO framework presented below addresses this gap through purpose-designed layers and feedback mechanisms.

Theoretical architecture of the neural-enabled risk orchestration (NERO) framework

The central theoretical contribution of this manuscript is the Neural-Enabled Risk Orchestration (NERO) framework. This systems-level conceptual topology integrates heterogeneous neural intelligence streams into a unified hospital risk governance architecture. Rather than functioning as a singular predictive instrument, NERO is constructed as an orchestration infrastructure that governs how clinical signals are perceived, propagated, reconciled, and operationalized across institutional care environments. The framework is therefore not reducible to algorithmic performance metrics; instead, it embodies an intelligence routing paradigm in which distributed predictive outputs are dynamically harmonized through contextual weighting, temporal recalibration, and governance-embedded oversight. Its architectural logic is structured as a five-layer topology organized around a central orchestration nucleus and sustained through bidirectional, closed-loop feedback recursion. This structural configuration enables theoretical adaptability independent of empirical retuning, positioning NERO as a continuously self-stabilizing clinical cognition scaffold.

Within this system's logic, intelligence is not produced at a single computational locus but emerges through orchestrated interaction across layered analytic substrates. Each layer contributes partial epistemic visibility into the patient risk landscape, and only through recursive synthesis does institutional-level risk awareness materialize. Consequently, NERO reframes clinical decision support from a predictive accuracy problem into a systems coordination problem, in which the reliability of risk governance depends not solely on model performance but also on the coherence of signal routing, arbitration latency, and feedback responsiveness. This repositioning is critical in high-volatility hospital environments, where clinical risk is neither static nor isolated but dynamically entangled with infrastructural, temporal, and human workflow variables.

At its foundational stratum, the Multimodal Neural Perception Layer serves as a high-bandwidth ingestion membrane that interfaces with heterogeneous hospital data ecosystems. This layer is responsible for transforming raw clinical telemetry into structured, representational embeddings capable of downstream orchestration [1, 4, 16, 20, 24, 25, 26]. Inputs traverse a wide epistemic spectrum, encompassing longitudinal electronic health record trajectories, real-time physiological monitoring streams, radiological imaging repositories, laboratory diagnostic pipelines, medication administration registries, and unstructured clinical narrative corpora. The ingestion architecture is designed to accommodate asynchronous data rates, allowing high-frequency telemetry streams to coexist with episodic diagnostic records without distorting representational fidelity.

Rather than enforcing premature modality fusion, the perception layer preserves epistemic fidelity by deploying specialized neural encoders operating in parallel computational channels. Transformer architectures process longitudinal EHR sequences to extract temporal deterioration signatures; convolutional networks interrogate imaging datasets for spatially encoded anomalies; graph attention networks model relational dependencies embedded in comorbidity structures; and domain-adapted clinical language models derive semantic risk indicators from narrative documentation. Physiological telemetry, characterized by continuous waveform volatility, is processed using temporally dilated convolutional filters that detect micro-oscillatory deviations preceding overt clinical decline. Each encoder produces latent embeddings that are projected onto a harmonized representational manifold, ensuring interoperability without erasing modality-specific information granularity. Through this design, the perception layer functions as a distributed neural sensing fabric, translating heterogeneous biomedical signals into computable clinical cognition primitives while preserving uncertainty signatures intrinsic to each modality stream. Functional stratification of the orchestration topology, including computational roles, intelligence outputs, and governance sensitivities, is detailed in **Table 1**.

Table 1. Functional layer specifications of the neural-enabled risk orchestration (NERO) architecture

Layer	Architectural function	Core computational components	Intelligence outputs	Governance sensitivities
Multimodal neural perception	High-bandwidth clinical signal ingestion	Transformers, CNNs, GNNs, Clinical LLMs, Waveform filters	Latent risk embeddings	Data heterogeneity, modality uncertainty
Risk propagation and connectivity	Relational diffusion of risk states	Temporal weighting graphs, adjacency networks	Institutional risk fields	Signal amplification bias
Central orchestration engine	Arbitration and prioritization of intelligence	Context weighting, queue routing, discounting filters	Weighted risk directives	Routing inequity, prioritization drift
Adaptive synthesis and prioritization	Fusion into actionable advisories	Bayesian aggregation, consensus modeling	Escalation alerts, care prompts	Alert fatigue, interpretability burden
Governance feedback and drift mitigation	Continuous oversight & recalibration	Drift detection, bias auditing, compliance gates	Audit logs, recalibration triggers	Monitoring latency, compliance overhead

Outputs generated at the perception stratum enter the Risk Propagation and Connectivity Layer, where isolated predictive signals are reinterpreted as temporally and relationally entangled risk phenomena [1, 2]. This layer rejects the assumption that patient risk exists in isolation; instead, it conceptualizes clinical deterioration as a propagative process diffusing across comorbidity networks, care pathways, and institutional resource ecologies. Dynamic connectivity graphs are instantiated to encode relational dependencies among patients, diagnostic states, and care units, thereby situating individual predictions within a broader institutional topology of vulnerability.

Temporal weighting mechanisms enable the system to model how risk states evolve longitudinally, while adjacency mappings permit the diffusion of localized alerts into system-level situational awareness fields. For example, deterioration risk emerging within a postoperative cohort may elevate surveillance weighting across patients sharing procedural exposure profiles or pharmacological regimens. Similarly, clustering of respiratory distress signals within a single care unit may algorithmically increase risk sensitivity across spatially contiguous beds, anticipating escalation cascades before overt outbreak confirmation. Through this propagation logic, deterioration risk identified within a single patient trajectory may elevate surveillance sensitivity across clinically adjacent cohorts or spatially proximal care environments. The resulting risk field is therefore not a static prediction but a continuously evolving topological landscape reflecting both individual and institutional vulnerability gradients shaped by temporal co-evolution, shared comorbidity structures, and care pathway convergence dynamics.

At the architectural core of the topology resides the Central Orchestration Engine, the computational nucleus responsible for arbitration, prioritization, and intelligence routing [8, 9, 11, 15]. This engine performs the critical function of reconciling potentially discordant neural outputs generated across perception and propagation strata. Rather than aggregating predictions via static averaging, it computes context-sensitive orchestration weights derived from multidimensional clinical parameters, including patient acuity gradients, real-time resource availability, workflow-phase positioning, and temporal-urgency indices. This contextual weighting calculus enables the engine to dynamically recalibrate the allocation of institutional attention in response to fluctuating operational pressures.

Through this weighting calculus, the engine dynamically determines which intelligence streams warrant prioritization within a given operational moment. Conflict resolution mechanisms are embedded in priority-queuing architectures, ensuring that high-acuity deterioration signals supersede lower-risk advisories even when probabilistic confidence intervals overlap. Temporal discounting functions further regulate signal influence by attenuating the weight of stale predictions relative to emergent clinical evidence, thereby preventing decision inertia driven by outdated analytic states. Routing protocols then direct weighted intelligence streams toward synthesis pathways aligned with specific institutional objectives, such as ICU escalation, readmission prevention, sepsis surveillance, or bed capacity stabilization. In this capacity, the orchestration engine operates analogously to an institutional neural traffic regulator, governing the directional flow and operational salience of distributed clinical cognition while maintaining equilibrium between predictive vigilance and workflow overload suppression.

Following orchestration, weighted intelligence streams converge within the Adaptive Synthesis and Prioritization Layer, where multimodal signals are transformed into unified decision constructs [14, 15, 30]. This layer performs ensemble fusion through probabilistic aggregation, consensus modeling, and confidence harmonization to generate clinically interpretable advisories. Importantly, synthesis does not seek to eliminate uncertainty but to operationalize it within bounded decision thresholds calibrated to institutional risk tolerance. Bayesian aggregation logics integrate modality confidence scores with orchestration weights to produce composite deterioration indices that are both interpretable and operationally actionable.

Risk outputs are stratified along graduated escalation continua, enabling differentiation between passive monitoring states and urgent intervention triggers. Latency optimization mechanisms regulate synthesis speed relative to decision criticality, ensuring that time-sensitive deterioration signals are not impeded by computational bottlenecks associated with lower-acuity analyses [5, 13, 29]. Decision compression algorithms further ensure that synthesized outputs are cognitively digestible within high-burden clinical environments, translating high-dimensional neural inferences into concise operational advisories. The outputs of this layer materialize within clinical workflows as dashboard alerts, electronic health record annotations, escalation prompts, and care pathway recommendations. Through this convergence process, orchestrated neural intelligence is rendered actionable within the temporal and cognitive constraints of hospital operations, while minimizing alert fatigue and escalation redundancy.

Encasing the decision synthesis stratum is the Governance Feedback and Drift Mitigation Layer, an oversight infrastructure embedded directly within the computational topology rather than imposed through retrospective auditing [3, 12, 13, 18, 27]. This layer continuously monitors model output coherence, distributional stability, and demographic equity to detect emergent deviations from operational or ethical baselines. Distributional drift is identified through divergence metrics that compare incoming data streams with historical training priors, triggering recalibration cycles when threshold deviations are exceeded. These recalibration loops may propagate backward to encoder recalibration, orchestration weight redistribution, or synthesis threshold adjustment, depending on the localization of the drift origin.

Output auditing mechanisms assess cross-module consistency to prevent contradictory advisories from entering clinical workflows. Bias surveillance subroutines interrogate whether predictive disparities are emerging across demographic,

https://doi.org//b543090992

geographic, or institutional subpopulations, thereby embedding fairness diagnostics directly within operational risk governance. Governance enforcement is operationalized through rule-augmented validation checkpoints that regulate alert eligibility, escalation authority, and interoperability compliance [7, 18]. Audit registries and explainability extraction modules ensure that every synthesized recommendation retains traceable epistemic lineage, thereby preserving institutional accountability and medico-legal transparency. In this configuration, governance ceases to function as an external compliance overlay. Instead, it becomes an endogenous property of the orchestration architecture itself, continuously co-evolving with predictive intelligence rather than being retrospectively evaluated.

Binding all architectural strata together is a hierarchical feedback topology that enables continuous system recalibration [11, 23]. Micro-level feedback loops operate at the encoder scale, adjusting modality confidence weights and recalibrating neural sensitivity to shifting signal distributions. Meso-level loops recalibrate orchestration weighting schemas in response to workflow performance metrics, resource fluctuations, and outcome validation signals. Macro-level governance loops propagate oversight insights backward across the topology, modifying ingestion policies, recalibration thresholds, and escalation eligibility criteria. Through recursive parameter updating, the system evolves in alignment with both clinical realities and governance mandates. This closed-loop recursion transforms NERO from a static decision support infrastructure into a self-correcting institutional intelligence organism capable of sustaining risk sensitivity across temporal, operational, and ethical dimensions.

When operating as an integrated topology, the Neural-Enabled Risk Orchestration framework exhibits emergent systems properties that transcend the capabilities of isolated predictive models. Hospital-wide situational awareness emerges through synchronized risk field propagation; cross-department deterioration synchronization arises through relational graph diffusion; anticipatory resource strain detection materializes through orchestration-weighted escalation forecasting; and governance-embedded ethical stabilization is sustained through continuous drift surveillance and accountability traceability. These properties arise not from individual neural components but from the orchestrated convergence of intelligence across perception, propagation, arbitration, synthesis, and oversight strata. NERO therefore represents a shift from predictive analytics toward institutional cognition engineering, positioning neural systems not as diagnostic tools but as infrastructural governors of clinical risk perception, prioritization, and response across hospital ecosystems.

Conceptual formulas: Risk propagation is expressed interpretively as: $RPt = \sum i = 1m\omega i \cdot Ni\left(Dt, Ct\right) \otimes O\left(St\right)$ where ω_i denotes dynamic orchestration weights, N_i represents the i-th neural module output, D_t and C_t are data and context vectors at time t, O is the orchestration function, and ⊗ symbolizes modulated integration [8, 11].

Decision orchestration confidence is conceptualized as: $DOC = 1 - e - \gamma \cdot H$ with H as the harmony score across propagated signals and γ a sensitivity parameter reflecting governance tolerance [3, 13].

Governance load is modeled as: CL $GL = \delta \cdot DR + \epsilon \cdot CL$ where DR denotes detected drift rate and CL compliance overhead, providing a theoretical metric for infrastructure sensitivity [12, 27]. The systemic topology of the Neural-Enabled Risk Orchestration framework is illustrated in **Figure 1**, which depicts the five-layer intelligence architecture, the central arbitration nucleus, and governance-embedded feedback recursion that enable hospital-wide risk harmonization.

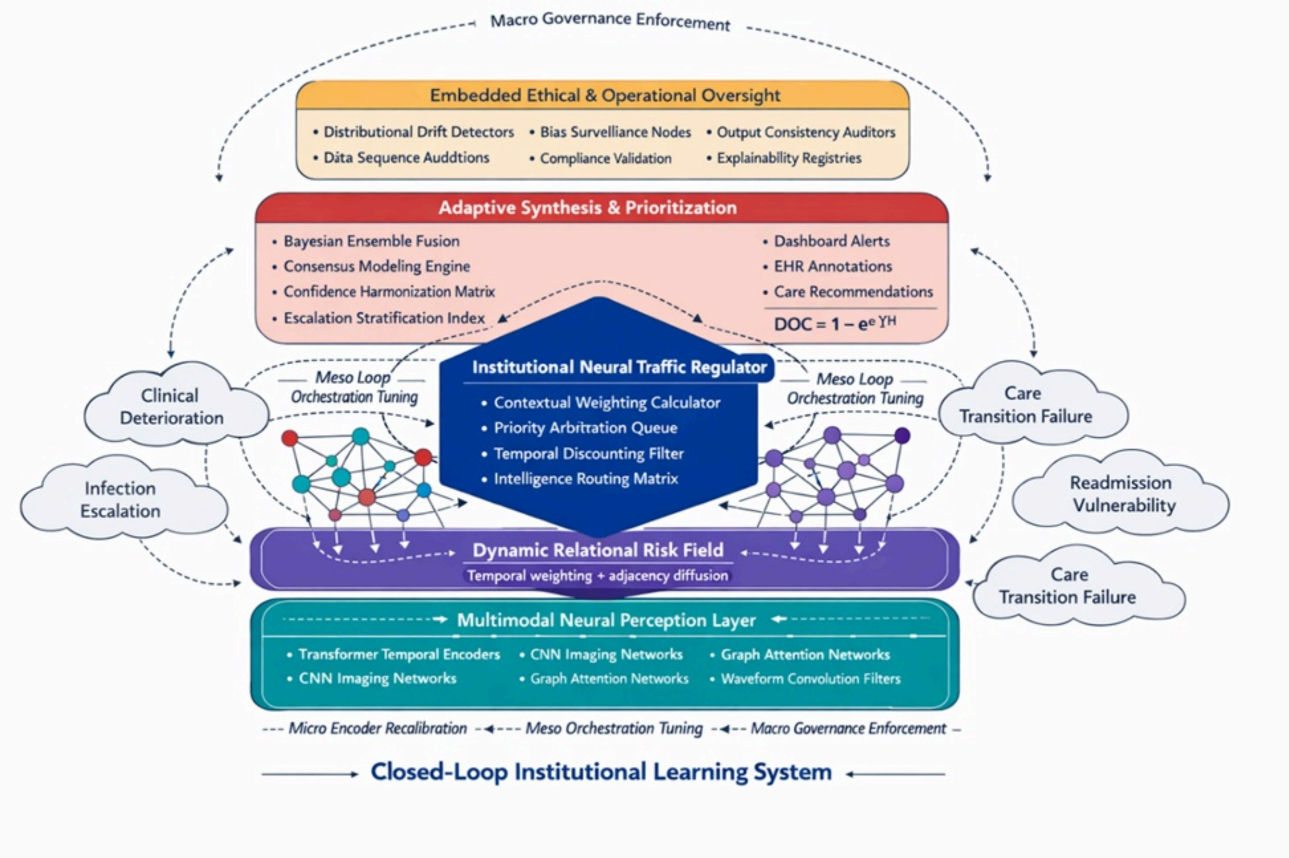

Figure 1. Neural-enabled risk orchestration (NERO) framework: five-layer clinical decision support topology with governance-embedded feedback recursion.

The architecture conceptualizes hospital risk management as an orchestrated intelligence system integrating multimodal neural perception, relational risk propagation, centralized arbitration, adaptive synthesis, and embedded governance oversight. The central orchestration engine dynamically weights and routes predictive signals based on acuity, context, and operational constraints. Bidirectional feedback loops enable continuous recalibration across analytic and governance strata, supporting drift mitigation, compliance monitoring, and institutional learning. Peripheral risk domains illustrate the propagation of deterioration intelligence across hospital ecosystems.

This topology theoretically positions NERO as a scalable, governance-native orchestration model for neural-enabled hospital risk management, distinct from prior fragmented architectures [6, 9, 17].

Operational consequences and human–AI workflow shifts in neural-enabled risk orchestration

The NERO framework, through its layered topology and closed-loop feedback, theoretically reshapes hospital risk management dynamics by redistributing cognitive and operational burdens across human clinicians and neural components [8, 10, 13].

Redistribution of cognitive load in high-acuity environments

In intensive and acute care units, where deterioration signals emerge rapidly from multimodal streams, NERO's central orchestration engine prioritizes and synthesizes alerts according to acuity gradients and temporal urgency [2, 5]. This theoretically alleviates clinician overload from fragmented notifications by presenting consolidated, context-weighted recommendations rather than parallel modality-specific outputs [12, 13]. The adaptive synthesis layer calibrates alert tiers—ranging from passive background monitoring to foreground escalation—based on propagated risk profiles, potentially reducing alert fatigue while preserving clinician agency in final judgment [29]. Human–AI workflow shifts manifest as a transition from reactive review of isolated predictions to proactive oversight of orchestrated ensembles, enabling clinicians to allocate attention toward interpretive synthesis, patient interaction, and ethical deliberation [10, 28].

Infrastructure sensitivities and resource allocation trade-offs: NERO's reliance on bidirectional feedback introduces infrastructure sensitivities around computational latency and data velocity [16, 17]. In resource-constrained hospital settings, the orchestration engine's dynamic weighting may theoretically favor lower-latency pathways during peak demand, trading marginal precision in peripheral signals for sustained core risk visibility [8]. Resource allocation dynamics emerge as the governance layer monitors system load and drift, triggering recalibration only when thresholds are breached, thus optimizing computational footprint without constant full recomputation [12, 27]. This approach theoretically mitigates over-provisioning while exposing trade-offs: excessive drift tolerance could propagate subtle misalignments, whereas overly conservative monitoring might impose unnecessary recalibration overhead [22].

Decision latency and escalation dynamics: The temporal feedback topology enables progressive refinement of risk signals, theoretically compressing overall decision latency in cascading risk scenarios (e.g., sepsis progression or postoperative complications) [2, 5]. Initial perception outputs propagate rapidly through connectivity graphs, with orchestration adjusting weights in near-real time to accelerate high-confidence pathways [1, 11]. However, governance-imposed checkpoints introduce controlled delays for drift validation, creating tunable latency trade-offs [13, 27]. In practice, this positions NERO as a spectrum between ultra-low-latency reactive alerting and deliberately paced deliberative synthesis, allowing hospitals to configure escalation dynamics according to unit-specific risk tolerances [7].

Governance dependencies and long-term system equilibrium: Sustained operation hinges on governance dependencies embedded in Layer 5. Continuous drift detection and compliance enforcement theoretically maintain equilibrium between neural adaptability and clinical stability [3, 12, 22]. Over time, feedback loops accumulate institutional knowledge—refining orchestration weights and perception encoders—potentially enhancing system resilience against population shifts or protocol changes [11, 23]. Yet dependencies on high-fidelity interoperability and clinician feedback channels introduce vulnerabilities: incomplete integration could starve governance of recalibration signals, leading to gradual divergence [18, 29]. The model thus underscores the necessity of socio-technical equilibrium, where human oversight remains integral to long-term governance load balancing [13, 28].

Conceptual extensions

Monitoring burden can be interpreted as: $MB = \beta \cdot \sum k\Delta k + \kappa \cdot FR$, where Δk denotes incremental drift in the k-th component, FR denotes the feedback request frequency, and β and κ are institutional tuning parameters reflecting tolerance thresholds [12, 13, 27].

These operational consequences position NERO as a theoretically adaptive platform that shifts hospital risk management toward proactive, distributed intelligence while surfacing explicit trade-offs in cognition, resources, latency, and governance.

Synthesis and strategic positioning

The NERO framework synthesizes fragmented advances in neural perception [4, 24, 25], risk connectivity [1, 2], orchestration logic [8, 11], adaptive synthesis [14, 15], and governance feedback [12, 13, 18] into a unified topology tailored for hospital-wide risk management. By addressing silos in current clinical AI architectures—where perception modules, predictive pipelines, and monitoring protocols operate in parallel [9, 17]—NERO theoretically enables coherent, context-sensitive decision support across acute care domains [5, 6]. Its emphasis on bidirectional temporal feedback distinguishes it from linear pipelines, offering a self-regulating mechanism that aligns neural intelligence with evolving clinical realities without presupposing empirical superiority [23].

Strategic positioning emerges from alignment with interoperability standards and workflow models prevalent in EHR ecosystems [17-21]. NERO's modular layering facilitates phased implementation: the initial deployment could focus on the perception and propagation layers, integrated with existing decision-support endpoints, with orchestration and governance added incrementally [7, 16]. This staged approach theoretically lowers adoption barriers while building toward full-system orchestration.

The framework also responds to persistent challenges in AI governance by embedding drift mitigation and compliance as intrinsic architectural features rather than post-hoc overlays [12, 22, 27]. Closed-loop recalibration theoretically supports sustained trustworthiness, addressing sensitivities around model drift and accountability in dynamic hospital environments [3, 13].

Limitations of the conceptual model include its abstraction from specific institutional constraints—such as varying EHR maturity, computational resources, or regulatory jurisdictions—which would require localized adaptation [17, 18]. Furthermore, while NERO prioritizes human-centric synthesis, it assumes clinician capacity to engage with orchestrated outputs; inadequate training or workflow misalignment could undermine intended shifts in cognitive load [10, 28].

Future conceptual extensions might explore hierarchical orchestration for multi-hospital networks or integration with emerging multi-agent deliberation patterns to enhance cross-specialty risk coordination [11, 14, 23].

Conclusion

Hospital risk management stands at an inflection point where neural-enabled capabilities promise proactive, multimodal intelligence but risk fragmentation without coordinated architecture. The Neural-Enabled Risk Orchestration (NERO) framework provides a theoretically grounded blueprint that unifies perception, propagation, orchestration, synthesis, and governance into a resilient topology. Through dynamic weighting, temporal feedback, and embedded oversight, NERO conceptually advances beyond isolated decision support toward adaptive, system-level risk platforms.

Realization of this potential depends on interdisciplinary collaboration among clinicians, informaticists, and system architects to refine interactions across layers and feedback equilibria. By prioritizing governance-native design and workflow continuity, NERO offers a foundational model for scalable neural-enabled hospital risk intelligence, theoretically enhancing safety, efficiency, and equity in high-stakes care environments.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 30 Aug 2021 Revised: 23 Sep 2021 Accepted: 23 Oct 2021
Published online: 20 January 2022

References

Ramamurthy P, Han Z, Delen D, Miao Z, Gin A, Luo X, et al. Graph attention network with comorbidity connectivity embedding for post-traumatic epilepsy risk prediction using sparse time-series electronic health records. Int J Med Inform. 2026;208:106239. https://doi.org/10.1016/j.ijmedinf.2025.106239.

Guo J, Xiong F, Sun B, Lei M, Qin Y. Ensemble machine learning for early mortality risk stratification in septic orthopedic trauma: an international cohort study. Int J Med Inform. 2026;208:106240. https://doi.org/10.1016/j.ijmedinf.2025.106240.

Huyut MT, Velichko A, Belyaev M, Izotov Y, Tosun M, Sertoğullarından B, et al. Clinical decision support system for detecting right ventricular dysfunction in acute pulmonary embolism: explainable new binary rule search ensembles and robustness evaluation. Artif Intell Med. 2026;172:103337. https://doi.org/10.1016/j.artmed.2025.103337.

Elgendi M, Elkhalifa A, Alshehhi M, Almarri E, Khalaf K, Khandoker A, et al. ECG sonification methods for robust and generalizable clinical decision support. npj Digit Med. 2025;9(1):27. https://doi.org/10.1038/s41746-025-02199-5.

Bolton WJ, Wilson R, Gilchrist M, Georgiou P, Holmes A, Rawson TM. The impact of artificial intelligence-driven decision support on uncertain antimicrobial prescribing: a randomised multimethod study. Lancet Digit Health. 2025;7(11):100912. https://doi.org/10.1016/j.landig.2025.100912.

Teo ZL, Thirunavukarasu AJ, Elangovan K, Cheng H, Moova P, Soetikno B, et al. Generative artificial intelligence in medicine. Nat Med. 2025;31(10):3270-82. https://doi.org/10.1038/s41591-025-03983-2.

Walsh CG, Long Y, Novak LL, Salwei ME, Tillman B, French B, et al. AI-driven clinical decision support to reduce hospital-acquired venous thromboembolism: a trial protocol. JAMA Netw Open. 2025;8(10):e2535137. https://doi.org/10.1001/jamanetworkopen.2025.35137.

Kücking F, Busch DA, Przysucha M, Kutza JO, Hannemann N, Hüsers J, et al. Impact of AI recommendation correctness on diagnostic accuracy in clinical decision-making. Int J Med Inform. 2026;207:106223. https://doi.org/10.1016/j.ijmedinf.2025.106223.

Wang Y, Griva A, Hassan UU. Time and temporality in machine learning methods to improve cancer clinical decision support: a literature review. Int J Med Inform. 2026;207:106219. https://doi.org/10.1016/j.ijmedinf.2025.106219.

Ratnaparkhi A, Moore S, Suri A, Wilson B, Alderete J, Florence TJ, et al. The perceived roles of AI in clinical practice: national survey of academic physicians. JMIR AI. 2025;4:e72535. https://doi.org/10.2196/72535.

Yang S, Jing M, Wang S, Huang Z, Wang J, Kou J, et al. Building trustworthy large language model-driven generative recommender system for healthcare decision support: a scoping review. Artif Intell Med. 2026;171:103310. https://doi.org/10.1016/j.artmed.2025.103310.

Wissel BD, Percy Z, Zachem TJ, Beaulieu-Jones B, Kohane IS, Goldstein SL, et al. Heterogenous effect of automated alerts on mortality. J Am Med Inform Assoc. 2025:ocaf222.

Hur S, Lee Y, Park J, Jeon YJ, Cho JH, Cho D, et al. Comparison of SHAP and clinician friendly explanations reveals effects on clinical decision behaviour. npj Digit Med. 2025;8(1):578. https://doi.org/10.1038/s41746-025-01958-8.

Mathes S, Ferber D, Dreyer T, Borm KJ, Modersohn L, Willem T, et al. Collaborative framework on responsible AI in LLM-driven CDSS for precision oncology leveraging real-world patient data. npj Precis Oncol. 2025;10(1):15. https://doi.org/10.1038/s41698-025-01180-5.

Liu S, Huang SS, McCoy AB, Wright AP, Horst S, Wright A. Optimizing order sets with a large language model-powered multiagent system. JAMA Netw Open. 2025;8(9):e2533277. https://doi.org/10.1001/jamanetworkopen.2025.33277.

Santomauro A, Kim D, Cánovas-Segura B, Leonardi G, Campos M, Portinale L, et al. A deep learning approach to predicting hospitalized patients' SEIRD states using multimodal spatiotemporal data. Int J Med Inform. 2026;206:106157. https://doi.org/10.1016/j.ijmedinf.2025.106157.

Arias MI, Cadavid L, Velásquez JD. Advancing healthcare analytics: a thematic review of machine learning, health informatics, and real-world data applications. J Biomed Inform. 2025;171:104934. https://doi.org/10.1016/j.jbi.2025.104934.

Liu S, McCoy AB, Chen Q, Wright A. Integrating rule-based NLP and large language models for statin information extraction from clinical notes. Int J Med Inform. 2026;205:106104. https://doi.org/10.1016/j.ijmedinf.2025.106104.

Huynh AL, Roy TJ, Jackson KN, Lee AG, Liaw W, Hossain MM. Applications of artificial intelligence-based conversational agents in healthcare: a systematic umbrella review. Int J Med Inform. 2026;207:106204. https://doi.org/10.1016/j.ijmedinf.2025.106204.

Hu J, Wang X, Yao K, Zhang X, Zhao W, Fu J, et al. FANS: a framework for automatic assessment of nutritional status based on free-text clinical notes. Int J Med Inform. 2026;206:106168. https://doi.org/10.1016/j.ijmedinf.2025.106168.

Domingues NS. A hybrid decision support system using rule-based and AI methods: the OnCATs knowledge-based framework. Int J Med Inform. 2026;206:106144. https://doi.org/10.1016/j.ijmedinf.2025.106144.

Ertuğrul DÇ, Akcan N, Bitirim Y. Revolutionizing pediatric obesity intervention strategies: from traditional tools to AI-enabled clinical decision support systems. Int J Med Inform. 2026;205:106109. https://doi.org/10.1016/j.ijmedinf.2025.106109.

Murali L, Gopakumar G, Viswanathan DM, Raman R, Nedungadi P. Integrating LLMs and knowledge graphs for medical AI: advances, challenges, and future directions. IEEE J Biomed Health Inform. 2025. https://doi.org/10.1109/JBHI.2025.3622058.

Gao PH, Duan XF, Pan XP. TTG-U-Net: an interpretable and efficient framework for multimodal brain tumor segmentation. IEEE J Biomed Health Inform. 2025. https://doi.org/10.1109/JBHI.2025.3648379.

Yisimitila T, Wang C, Hou M, Maimaitiniyazi M, Aili Z, Liu T, et al. Bridging clinical knowledge and AI: an interpretable transformer framework for ECG diagnosis. npj Digit Med. 2025;9(1):41. https://doi.org/10.1038/s41746-025-02215-8.

Ye J, Liu Y, Ren S, Wang C, Zhou Y, Yang L, et al. A knowledge-guided multimodal neural network for breast cancer molecular subtyping. IEEE J Biomed Health Inform. 2025. https://doi.org/10.1109/JBHI.2025.3643125.

Du X, Zhou Z, Wang Y, Chuang YW, Li Y, Yang R, et al. Performance and improvement strategies for adapting generative large language models for electronic health record applications: a systematic review. Int J Med Inform. 2026;205:106091. https://doi.org/10.1016/j.ijmedinf.2025.106091.

Dave B, Martin P, David SS, Kumar S, Chakraborty T. Enhancing healthcare worker mental health via artificial intelligence-driven work process improvements: a scoping review. Int J Med Inform. 2026;205:106122. https://doi.org/10.1016/j.ijmedinf.2025.106122.

Van De Sijpe G, Gilissen L, Wets D, Breynaert C, Ieven T, De Munter P, et al. Evaluation of the allergy fact checker, a clinical decision support system for beta-lactam delabeling. npj Digit Med. 2025;8(1):662. https://doi.org/10.1038/s41746-025-02030-1.

Kim D, Yoo S, Jeong O. MedSumGraph: enhancing GraphRAG for medical QA with summarization and optimized prompts. Artif Intell Med. 2026;172:103311. https://doi.org/10.1016/j.artmed.2025.103311.