

ORIGINAL RESEARCH

Open access

Multi-Relational Graph Learning for Patient Similarity and Clinical Decision Support

Claire Martin¹, Ahmed Mansour^{2*}

Abstract

The integration of graph-based architectures into healthcare systems represents a pivotal advancement, enabling personalized clinical intelligence through patient similarity metrics. This conceptual manuscript proposes a novel framework, the Graph-Integrated Patient Affinity Network (GIPAN), that orients patient data as interconnected nodes within a dynamic graph, facilitating similarity-driven insights for clinical decision-making. Drawing from theoretical foundations in clinical AI infrastructures, electronic health record (EHR) ecosystems, and interoperability frameworks, GIPAN emphasizes layered graph embeddings that capture multidimensional patient profiles, including temporal trajectories, comorbidity patterns, and treatment responses. The architecture incorporates feedback loops for adaptive similarity refinement, ensuring alignment with evolving clinical workflows without empirical validation. Key theoretical contributions include formulas for similarity propagation across graph layers and governance load estimation in deployment scenarios. By synthesizing recent literature on graph neural networks in healthcare analytics and decision-support pipelines, this work highlights the infrastructural prerequisites for scalable, privacy-preserving patient matching. Potential impacts encompass enhanced diagnostic precision in heterogeneous populations and streamlined resource allocation in personalized medicine ecosystems. This conceptual design underscores the need for robust AI governance to mitigate biases in similarity computations, paving the way for future theoretical explorations in graph-centric clinical intelligence.

Keywords AI governance, EHR interoperability, Graph-based architectures, Patient similarity, Personalized clinical intelligence, Decision support frameworks

*Correspondence:

Ahmed Mansour
ahmed.mansour@gmail.com

¹ Department of Healthcare Systems Research, Faculty of Medicine, University of Lyon, Lyon, France

² Department of Health Systems and Digital Medicine, Faculty of Medicine, Cairo University, Cairo, Egypt

Introduction

Evolving paradigms in graph-structured patient data representation

The landscape of healthcare analytics has increasingly gravitated toward graph-structured representations to model the intricate interconnections inherent in patient data [1-6]. Traditional tabular EHR formats often fail to capture the relational dynamics between patients, such as shared phenotypic traits or treatment pathways, which are crucial for deriving personalized clinical intelligence. Graph-based approaches, by contrast, conceptualize patients as nodes linked by edges representing similarity metrics derived from demographics, genomics, or longitudinal health events, enabling a more holistic view of clinical cohorts [2, 7, 8]. This shift is particularly relevant in personalized medicine, where identifying analogous patient profiles can inform tailored interventions without relying on broad population averages [5]. The proposed Graph-Integrated Patient Affinity Network (GIPAN) builds on this paradigm by introducing a modular

architecture that embeds patient similarity as a core intelligence primitive, facilitating real-time clinical inference across diverse settings such as oncology and chronic disease management [4, 9].

Challenges in similarity-driven clinical decision pipelines

Integrating patient similarity into clinical decision support pipelines poses unique infrastructural hurdles, including harmonizing disparate data modalities across EHR systems [7, 10–12]. For instance, variability in data granularity—such as incomplete comorbidity graphs or inconsistent temporal annotations—can distort similarity computations, potentially leading to suboptimal clinical recommendations [10, 13–16]. Moreover, deployment environments in federated healthcare networks amplify these challenges, as data silos hinder cross-institutional patient matching while raising interoperability concerns [3, 13]. The literature underscores the need for architectures that abstract these complexities through graph-based abstractions, ensuring seamless integration without compromising data sovereignty [11, 14]. GIPAN addresses this by proposing a layered similarity-propagation mechanism that is adaptable to governance-constrained environments such as multi-site hospitals [15, 17].

Governance constraints shaping personalized intelligence ecosystems

AI governance frameworks are indispensable for deploying graph-structured patient-similarity systems, particularly for ensuring ethical alignment and regulatory compliance [14, 18]. Constraints such as data privacy regulations (e.g., GDPR or HIPAA) require architectures that incorporate federated learning principles, in which similarity graphs are computed locally before aggregation [19, 20]. This not only mitigates risks of data leakage but also enhances trust in personalized clinical intelligence outputs [21, 22]. Theoretical models highlight how governance layers can embed audit trails within graph topologies, tracking derivations of similarities for accountability [23–26]. In high-stakes clinical settings, such as intensive care units, these constraints influence the design of intelligence ecosystems by prioritizing explainable similarity metrics over opaque black-box alternatives [25, 27].

Deployment environments for graph-enabled patient matching

The operationalization of patient similarity architectures demands consideration of diverse deployment environments, ranging from cloud-based EHR platforms to edge computing devices in ambulatory care [28, 29]. Centralized graphs can lead to scalability issues in large-scale systems, whereas distributed topologies offer resilience but introduce latency in similarity queries [30, 31]. Conceptual designs must therefore balance these factors, incorporating adaptive embeddings that evolve with incoming patient data streams [32]. GIPAN envisions a hybrid deployment model in which core similarity intelligence resides in a central hub, interfacing with peripheral clinical workflows to foster personalized insights across heterogeneous environments [5, 8].

Theoretical imperatives for interoperable similarity frameworks

Interoperability remains a cornerstone for graph-structured patient similarity, enabling seamless data exchange across fragmented healthcare infrastructures [7, 10]. Standards such as FHIR and OMOP provide foundational schemas for mapping patient attributes to graph nodes, yet gaps persist in handling dynamic similarities driven by real-time clinical events [2, 6]. This manuscript synthesizes these imperatives into a cohesive architectural vision, emphasizing theoretical interoperability layers that abstract over syntactic differences to enable unified intelligence generation [13, 15]. By anchoring on patient-centric graphs, such frameworks promise to elevate clinical decision-making from generic protocols to individualized strategies [1, 4].

Rationale for a dedicated graph architecture in clinical intelligence

The impetus for GIPAN stems from the recognition that existing clinical AI systems often overlook graph-induced similarities as a primary driver for personalization [9, 11]. While prior work explores graph applications in specific domains, such as electrocardiogram analysis or cancer subtyping, a unified architecture for broad-spectrum patient

similarity remains underexplored theoretically [8, 23]. This introduction posits that by centering on graph structures, healthcare analytics can achieve greater fidelity in modeling patient affinities, ultimately enhancing clinical intelligence without empirical dependencies [3, 12].

Theoretical Background and Literature Synthesis

Foundational constructs in clinical AI architectures for similarity modeling

Clinical AI architectures have evolved to incorporate graph-based constructs for modeling patient similarities, drawing from advancements in neural embeddings and relational data processing [1, 5]. These systems conceptualize EHR data as interconnected graphs, where nodes represent patient entities and edges encode similarity affinities based on features such as vital-sign trajectories or genetic profiles [6, 23]. Theoretical underpinnings emphasize the role of graph neural networks (GNNs) in propagating information across layers, enabling the capture of higher-order relationships that traditional machine learning overlooks [24, 25]. For instance, structure-aware GNNs facilitate encounter-level similarity learning, abstracting complex clinical interactions into interpretable metrics [1]. This synthesis highlights how such architectures form the bedrock for personalized intelligence, abstracting away raw data heterogeneities through topological designs [2, 4].

Healthcare analytics infrastructures supporting graph-interoperable ecosystems

Analytics infrastructures in healthcare increasingly prioritize graph interoperability to support federated similarity computations across disparate systems [7, 10]. Frameworks like FHIR-Ontop-OMOP exemplify this by mapping clinical data models into RDF graphs, enabling seamless exchange and similarity queries without centralized storage [7]. Literature synthesizes these into conceptual pipelines that integrate multi-omics data via graph convolutional networks, theorizing enhanced subtype identification in diseases like gastric cancer [4]. Moreover, knowledge graphs augment analytics by visualizing AI applications in domains such as electrocardiography, providing a scaffold for similarity-driven insights [8, 9]. These infrastructures underscore the need for modular designs that accommodate governance protocols, ensuring that analytics remain robust in privacy-sensitive ecosystems [3, 13].

EHR intelligence ecosystems and their graph-centric evolution

EHR intelligence ecosystems have transitioned toward graph-centric models to harness patient similarity for predictive and decision-support functions [5, 11]. Deep learning trajectories based on EHR data leverage graph embeddings to forecast health events, theorizing self-supervised learning mechanisms that refine similarities over time [25, 32]. Conceptual reviews emphasize the integration of hierarchical embeddings, in which medical entities are structured into graphs to enable application-specific intelligence, such as diagnostic prediction [24, 26]. This evolution addresses limitations in scalar-based representations by introducing topological intelligence, facilitating personalized pathways in chronic care management [28, 29]. Synthesized literature suggests that these ecosystems benefit from ontology-driven graphs, thereby enhancing explainability in pathogenicity estimates [9, 27].

Decision support pipelines enhanced by patient similarity graphs

Decision support pipelines in clinical settings increasingly embed patient similarity graphs to orchestrate personalized recommendations [12, 15]. Hybrid architectures combining rule-based reasoning with graph learning provide theoretical scaffolding for scalable support, mitigating uncertainties in high-variability environments such as neuro-oncology [16]. Real-time natural language processing integrated with graph pipelines exemplifies this, enabling disagreement prediction between AI and clinicians for refined intelligence [15, 17]. Theoretical syntheses advocate for collaborative platforms like MERLIN, which facilitate resource-efficient AI research through graph-enabled data sharing and thereby indirectly bolster similarity-driven decisions [10]. Governance elements within these pipelines ensure equitable deployment, theorizing feedback mechanisms that adapt similarities to workflow dynamics [13, 14].

AI governance, monitoring, and deployment systems in graph contexts

Governance systems for AI in healthcare emphasize monitoring graph-structured deployments to prevent drift in patient similarity metrics [14, 18]. Conceptual models outline a taxonomy for hybrid systems, integrating machine learning with governance layers to track decision integrity [12, 19]. Deployment successes in large health systems highlight key elements such as natural language processing pipelines and the view of governance as a prerequisite for similarity reliability [13, 17]. International understandings of data and AI governance further synthesize monitoring needs, advocating for frameworks that embed ethical constraints into graph architectures [14]. This ensures deployment systems remain sensitive to biases in similarity computations, promoting trustworthy clinical intelligence [20, 21].

Interoperability and data exchange frameworks for similarity propagation

Interoperability frameworks are critical for propagating patient similarities across clinical networks [7, 22]. Standards-based approaches, such as those leveraging OMOP for knowledge graphs, enable federated data exchange while preserving similarity fidelity [6, 7]. Literature synthesizes these into topologies that support multi-omics integration via GCNs, theorizing robust exchange in cancer analytics [23]. Moreover, seamless AI integration into workflows relies on interoperable graphs, as seen in pneumothorax detection algorithms, abstracting similarity for operational efficiency [22]. Conceptual overviews stress the convergence of human and AI intelligence through interoperable systems, enhancing personalized outputs [29, 31].

Clinical workflow integration models leveraging graph similarity

Integration models for clinical workflows increasingly utilize graph similarity to redistribute cognitive loads between humans and AI [16, 20]. Theoretical pipelines for machine learning implementation validate conceptual integration, emphasizing graph-based prediction for temporal events [19, 25]. Synthesized works on scalable deep learning with EHRs highlight workflow embeddings that prioritize similarity in resource-constrained settings [28, 32]. Furthermore, explainable AI via knowledge graphs integrates into workflows by estimating variant pathogenicity and theorizing that similarity serves as a bridge for clinician adoption [9, 27]. These models collectively advocate for adaptive topologies that align graph intelligence with clinical imperatives [3, 11].

Graph-enabled orchestration architecture for similarity-powered clinical intelligence

The GIPAN represents a conceptual orchestration architecture tailored for leveraging patient similarity in personalized clinical intelligence. At its core, GIPAN structures patient data as a multi-layered graph, where each layer corresponds to distinct similarity dimensions: demographic (Layer 1), phenotypic (Layer 2), temporal trajectory (Layer 3), and intervention response (Layer 4). Nodes in these layers denote patient profiles, with edges weighted by similarity scores computed via theoretical embeddings, such as cosine distances in feature spaces [1, 5]. The architecture employs a unique spiral feedback topology, in which intelligence from higher layers iteratively refines lower-layer edges, enabling adaptive similarity evolution without external data inputs [24, 25].

Central to GIPAN is the orchestration hub, a theoretical integrator that harmonizes interlayer propagations through governance-aware algorithms. This hub interfaces with EHR ecosystems via interoperability wrappers, abstracting data exchanges into graph updates [7, 10]. Deployment envisions a federated setup in which local similarity subgraphs compute affinities before aggregation, thereby minimizing governance load [13, 14]. Monitoring is embedded via sentinel nodes that track drift sensitivities, ensuring architectural stability in dynamic clinical environments [18, 20].

The overall architecture of GIPAN is illustrated in **Figure 1**.




Figure 1. GIPAN architecture for similarity-powered clinical intelligence.

The schematic depicts a four-layer concentric graph structure that represents demographic, phenotypic, temporal-trajectory, and intervention-response similarity dimensions. Patient nodes are connected by weighted edges encoding affinity scores. A spiral feedback topology enables cyclic refinement of similarity metrics across layers. A centralized orchestration hub harmonizes interlayer propagation and interfaces with external EHR ecosystems through interoperability wrappers. Governance sentinel nodes monitor drift sensitivity and compliance burdens. All elements are conceptual and non-empirical.

To formalize key dynamics, consider the following interpretive formulas:

$$\begin{aligned}
 & S_l \\
 & + 1(p_i, p_j) \\
 & = \alpha \\
 & \cdot S_{l(p_i, p_j)} \quad \text{where } S_l(p_i, p_j) \text{ is the similarity between patients } p_i \text{ and } p_j \\
 & + (1 - \alpha) \\
 & \cdot \sum_{k \in N(p_i)} w_{ik} \\
 & \cdot f(\Delta k_j)
 \end{aligned}$$

at layer l α balances direct and propagated affinities, $N(p_i)$ denotes neighbors, w_{ik} are edge weights, and $f(\Delta k_j)$ interprets feature deltas (theoretical, non-computable). This captures how similarities cascade, amplifying personalized intelligence.

2. Governance Load Estimation:

$$GL = \beta \cdot |E| + \gamma \cdot \log(|V| + \delta \cdot D_f)$$

where $|E|$ and $|V|$ are graph edges and vertices, D_f is feedback depth, and the coefficients β , γ , and δ reflect theoretical burdens from monitoring, scaling, and adaptation. This estimates resource demands in deployment.

3. Drift Sensitivity in Feedback Topology:

$$DS = \frac{\sum_t t}{1T\theta t} \left\| \begin{matrix} S_t \\ S_t \\ -1 \end{matrix} \right\|$$

where S_t is the similarity matrix at time t , θ discounts prior drifts, and $\|\cdot\|$ denotes a norm. This interprets sensitivity to changes, guiding architectural resilience. The structural components and their functional roles within GIPAN are summarized in [Table 1](#).

Table 1. Structural components of the GIPAN architecture and functional roles

Component	Architectural position	Functional role	Theoretical contribution	Governance sensitivity
Demographic layer	Core (Layer 1)	Encodes static patient attributes	Baseline similarity initialization	Moderate (identity-sensitive)
Phenotypic layer	Layer 2	Captures comorbidities and clinical features	Higher-order relational embedding	High (bias risk amplification)
Temporal layer	Layer 3	Models longitudinal health trajectories	Dynamic similarity evolution	Moderate
Intervention response layer	Layer 4	Encodes treatment outcomes	Predictive personalization refinement	High (clinical risk implications)
Spiral feedback loop	Cross-layer	Iterative similarity recalibration	Enables adaptive intelligence	High (drift monitoring required)
Orchestration hub	Central integrator	Interlayer propagation harmonization	Governance-aware similarity fusion	Critical
Sentinel nodes	Boundary monitors	Drift and compliance tracking	Bias mitigation and accountability	Critical
Federated subgraphs	Distributed nodes	Local similarity computation	Privacy-preserving deployment	Critical

GIPAN's uniqueness lies in its spiral topology, distinguishing it from linear or hierarchical alternatives by enabling cyclic refinements tailored to clinical intelligence needs [6, 23]. This architecture theorizes enhanced workflow integration, where similarity-driven insights feed into decision pipelines without disrupting human-AI interactions [15, 21].

Clinical adoption dynamics in graph-structured similarity ecosystems

The adoption of graph-structured architectures, such as GIPAN, in clinical settings introduces multifaceted dynamics that reshape how personalized intelligence is integrated into healthcare delivery. These dynamics encompass shifts in clinician engagement, resource utilization patterns, and systemic resilience, all of which are theorized through the lens of similarity-driven workflows [1, 5]. At the forefront, adoption hinges on the perceived utility of patient affinity networks in augmenting diagnostic and therapeutic precision. For instance, in heterogeneous patient populations, such as those in oncology clinics, graph-enabled similarities could theoretically streamline case matching by highlighting analogous trajectories, thereby reducing cognitive burdens on physicians who traditionally rely on manual chart reviews [4, 9]. This redistribution of effort fosters a collaborative human-AI paradigm, in which clinicians validate graph-propagated insights rather than conducting exhaustive searches, potentially accelerating adoption in time-sensitive environments such as emergency departments [15, 16].

However, these dynamics are not without tensions. Governance dependencies emerge as a critical factor, where stringent regulatory frameworks may slow integration by necessitating extensive validation protocols for similarity metrics [13, 14]. In federated ecosystems, adoption dynamics are influenced by interoperability sensitivities; disparate EHR standards could fragment graph coherence, leading to uneven uptake across institutions [7, 10]. Theoretical models suggest that such fragmentation amplifies adoption barriers in resource-limited settings, such as rural healthcare facilities, where infrastructure constraints limit the scalability of multi-layered graphs [3, 11]. Moreover, human-AI workflow shifts introduce dynamics of trust calibration—clinicians may initially resist similarity-based intelligence if graph embeddings lack intuitive explainability, prompting a phased adoption curve that begins with low-stakes applications, such as population health analytics, before extending to direct patient care [20, 21].

Expanding on the operational consequences, integrating GIPAN could lead to trade-offs in decision latency. In high-volume clinical workflows, real-time similarity queries across expansive graphs can impose computational overhead, potentially delaying intelligence delivery in scenarios such as intensive care monitoring [17, 19]. Conversely, the spiral feedback topology mitigates this by prioritizing adaptive refinements, potentially optimizing long-term adoption through iterative improvements in similarity accuracy [24, 25]. Infrastructure sensitivities further complicate dynamics; reliance on cloud-based deployments exposes systems to connectivity variances, which could hinder adoption in offline-prone environments [28, 29]. The literature synthesizes these as ecosystem-wide ripple effects, in which enhanced personalization in one domain (e.g., chronic disease management) inadvertently strains adjacent workflows, such as administrative data processing [2, 6].

Drilling deeper into risk propagation, graph-structured similarities inherently amplify certain vulnerabilities. Biases embedded in initial node embeddings—stemming from underrepresented patient demographics—could propagate through layers, distorting affinities and leading to inequitable clinical intelligence [8, 9]. Adoption dynamics thus demand proactive governance to curb such propagation, theorizing hybrid monitoring that combines automated drift detection with human oversight [18, 22]. In multicultural healthcare contexts, these dynamics extend to cultural sensitivities, where similarity metrics must account for diverse health belief systems to avoid alienating user groups [30, 31]. Furthermore, economic dynamics play a role: upfront infrastructure investments for GIPAN deployment may deter adoption in cost-constrained systems, yet long-term efficiencies in resource allocation—such as targeted interventions based on similar patient outcomes—could justify the shift [12, 32].

Human factors in adoption cannot be overstated. Cognitive load redistribution posits that graph intelligence alleviates routine similarity assessments, freeing clinicians for higher-order reasoning, but this assumes seamless integration [16, 20]. Dynamics of resistance arise when AI outputs challenge established clinical intuitions, necessitating change management strategies that emphasize training in graph interpretation [15, 21]. Theoretical explorations suggest that adoption accelerates in learning health systems, where iterative feedback from similarity ecosystems informs continuous workflow evolution [3, 10]. Ultimately, these dynamics converge on a transformative potential: widespread adoption could democratize personalized intelligence, bridging healthcare equity gaps by enabling similarity-based knowledge transfer across global networks [14, 27]. However, without addressing these multifaceted dynamics, graph architectures risk remaining theoretical curiosities rather than operational mainstays [5, 23].

Results and Discussion

This conceptual exploration of the GIPAN illuminates the broader implications of embedding graph-structured patient similarities within clinical intelligence frameworks, synthesizing theoretical advancements while acknowledging inherent limitations and future directions. Central to the discussion is the architectural innovation of GIPAN's multi-layered design and spiral feedback topology, which, in theory, elevates patient similarity from a peripheral analytical tool to a foundational intelligence driver [1, 6]. By conceptualizing patients as dynamic nodes in an affinity network, GIPAN addresses gaps in traditional EHR ecosystems that often treat data in isolation, thereby fostering a more interconnected view of clinical narratives [2, 7]. This resonates with literature on graph neural networks, where layered embeddings capture nuanced relationships, such as comorbidity clusters or temporal event sequences, without necessitating empirical datasets [5, 24]. However, the discussion must critically examine how such abstractions align with real-world clinical heterogeneity, where data incompleteness could, in theory, undermine graph integrity [8, 25]. Theoretical deployment dynamics and infrastructural implications are further detailed in Table 2.

Table 2. Theoretical dynamics and deployment implications of similarity propagation

Theoretical construct	Formula reference	Operational meaning	Deployment risk	Mitigation strategy
Similarity propagation	$S_I^{+1} = \alpha S_I + (1 - \alpha) \sum w_{ik} f(\Delta_{kj})$	Cross-layer affinity cascade	Bias amplification	Alpha tuning + sentinel audits
Governance load	$G_L = \beta$	E	+ ylog	V
Drift sensitivity	$D_S = \sum \theta^t$		$S_t - S_{t-1}$	
Federated similarity aggregation	Local-to-global affinity merging	Privacy-preserving scaling	Latency	Hierarchical aggregation
Spiral feedback depth	Feedback iteration cycles	Adaptive refinement intensity	Computational overhead	Depth regularization

Expanding on interoperability, GIPAN's orchestration hub proposes a bridge across fragmented healthcare infrastructures, leveraging standards such as FHIR and OMOP to enable similarity propagation across federated settings [7, 10]. This is particularly pertinent in global health contexts, where cross-border data exchanges could leverage graph similarities for pandemic response or rare disease management [13, 28]. Yet, the discussion reveals potential pitfalls: over-reliance on weighted edges might amplify systemic biases if edge computations favor dominant demographic features, echoing concerns in the AI governance literature about equitable distribution of intelligence [14, 18]. Theoretical formulas for similarity propagation and governance load provide interpretive tools to anticipate these issues, allowing stakeholders to model trade-offs in advance [provided in the architecture section]. For example, adjusting the alpha parameter in propagation could simulate bias mitigation strategies, highlighting GIPAN's value as a sandbox for ethical architecture design [9, 26].

In terms of clinical workflow integration, the discussion posits that GIPAN induces a paradigm shift toward proactive, similarity-informed decision-making [15, 20]. Unlike static decision support pipelines, the adaptive nature of spiral feedback enables theoretical resilience to evolving patient profiles, such as those influenced by emerging therapies or lifestyle changes [17, 21]. This could theoretically enhance outcomes in personalized medicine domains like precision oncology, where matching patients to similar trial responders accelerates therapeutic matching [4, 23]. However, expanded analysis uncovers workflow disruptions: in multidisciplinary teams, graph-derived insights might create interpretive silos if not accompanied by standardized visualization tools [16, 22]. Literature synthesis suggests integrating

explainable components, such as knowledge graph visualizations, to bridge this gap, ensuring that similarity intelligence augments rather than supplants human expertise [8, 27].

Governance and monitoring emerge as pivotal discussion themes, with GIPAN's sentinel nodes offering a conceptual mechanism for ongoing oversight [18, 19]. In an era of heightened regulatory scrutiny, this architecture theorizes compliance through embedded audit trails that track derivations of similarity to facilitate post-deployment reviews [12, 14]. Expanding on this, the discussion considers scalability challenges: as graph sizes grow with EHR volumes, governance loads could escalate exponentially, as captured in the interpretive formula, necessitating hybrid cloud-edge deployments to distribute the burden [3, 29]. Ethical dimensions further enrich the discourse; patient consent models must evolve to encompass graph-based similarities, where affinities imply indirect data sharing across cohorts [30, 31]. Theoretical precedents from learning health systems advocate transparent governance, positioning GIPAN as a catalyst for AI-driven healthcare policy innovation [11, 32].

Limitations inherent to this conceptual framework warrant thorough discussion. Absent empirical validation, GIPAN remains speculative and may overlook unforeseen interactions in live environments, such as network latency in real-time similarity queries [19, 25]. Moreover, the focus on graph structures might undervalue complementary modalities, such as natural language processing for unstructured notes, suggesting the need for hybrid extensions in future work [17, 32]. Broader societal implications include workforce transformations: as similarity intelligence automates routine matching, reskilling clinicians in graph literacy becomes imperative to prevent deskilling [20, 21]. Finally, the discussion envisions expansive future avenues, such as incorporating multimodal data (e.g., imaging alongside EHRs) into GIPAN layers or exploring quantum-inspired graphs for hyper-complex similarities [9, 28]. By framing these as theoretical frontiers, this manuscript contributes to the ongoing dialogue on architecting intelligent, patient-centric healthcare systems [1, 5].

Conclusion

In synthesizing the conceptual underpinnings and architectural innovations of the GIPAN, this manuscript underscores the transformative potential of graph-structured patient similarity for advancing personalized clinical intelligence. By theorizing a multi-layered, feedback-enriched topology, GIPAN addresses critical gaps in existing clinical AI systems, offering a blueprint for integrating relational patient data into decision-making ecosystems without empirical dependencies. The introduction of interpretive formulas for propagation, load, and sensitivity provides stakeholders with tools to conceptualize dynamics, fostering a deeper understanding of infrastructural trade-offs in deployment scenarios. Ultimately, this framework posits that similarity-driven graphs can elevate healthcare from reactive protocols to anticipatory, individualized strategies, particularly in domains plagued by data silos and variability.

In ethical terms, GIPAN's governance-embedded design theorizes a pathway to bias-resilient intelligence, safeguarding against propagated inequities and aligning with calls for responsible AI in medicine. Operationally, the architecture promises streamlined workflows, as similarity affinities automate cohort identification, potentially reducing diagnostic delays in chronic care and enhancing resource efficiency in overburdened systems. Societally, widespread adoption could democratize access to personalized insights, bridging disparities in underserved populations by leveraging global graph networks for knowledge dissemination. However, this optimism is tempered by the need for vigilant implementation; without addressing interoperability hurdles and human-AI synergies, such architectures risk exacerbating existing divides.

Conceptualizing GIPAN in specialized contexts—such as geriatric care, where temporal similarities predominate, or pediatric oncology, which emphasizes phenotypic affinities—offers fertile ground for domain-specific refinements. Policymakers and researchers are encouraged to build on this foundation by developing simulation-based explorations to anticipate real-world behaviors, even as empirical studies remain future endeavors. In essence, GIPAN exemplifies how graph-centric thinking can redefine clinical intelligence, paving the way for a future in which patient similarities drive not just data analytics but also holistic health advancements. By concluding on this note, the manuscript calls for continued theoretical innovation to realize the full promise of personalized, graph-empowered healthcare.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 03 Sep 2023

Revised: 09 Oct 2023

Accepted: 15 Nov 2023

Published online: 20 January 2024

Rights and permissions

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Gu Y, Yang X, Tian Y, Li H, Cui C. Structure-aware siamese graph neural networks for encounter-level patient similarity learning. *J Biomed Inform.* 2022;126:103982.
<https://doi.org/10.1016/j.jbi.2022.103982>.
- Schäfer J, Storf H, Bauer CR, Wagner TOF, Kadioglu D. Graph4Med: a web application and a graph database for visualizing and analyzing medical databases. *BMC Bioinformatics.* 2022;23(1):537.
<https://doi.org/10.1186/s12859-022-05069-3>.
- Kapoor R, Sleeman WC 4th, Palta J, Martin R, Showalter TN. Infrastructure tools to support an effective radiation oncology learning health system. *Int J Radiat Oncol Biol Phys.* 2023;115(2):303-12.
<https://doi.org/10.1016/j.ijrobp.2022.07.036>.
- Liu C, Liu X, Sun X, Chen J, Wu J. A classification method of gastric cancer subtype based on residual graph convolution network. *Front Oncol.* 2023;13:1093267.
<https://doi.org/10.3389/fonc.2023.1093267>.
- Amirahmadi A, Hebsur S, Huang LC, Tourani R. Deep learning prediction models based on EHR trajectories: a systematic review. *J Am Med Inform Assoc.* 2023;30(7):1237-49.

Chandak P, Huang K, Zitnik M. Building a knowledge graph to enable precision medicine. *Sci Data*. 2023;10(1):67.
<https://doi.org/10.1038/s41597-023-01960-3>.

Xiao G, Lamer A, Yasini M, Bouzillé G, Parmentier S, Cuggia M. FHIR-Ontop-OMOP: building clinical knowledge graphs in FHIR RDF with the OMOP common data model. *J Biomed Inform*. 2022;127:103999.
<https://doi.org/10.1016/j.jbi.2022.103999>.

Yang M, Wang H, Hu K, Yin G, Wei X. Knowledge graph analysis and visualization of artificial intelligence applied in electrocardiogram. *Front Physiol*. 2023;14:1179133.
<https://doi.org/10.3389/fphys.2023.1179133>.

Abe S, Yokoyama K, Kato K, Fukuda H, Yamamoto Y. Explainable AI for estimating pathogenicity of genetic variants using large-scale knowledge graphs. *J Pers Med*. 2023;13(3):410.
<https://doi.org/10.3390/jpm13030410>.

Cohen RY, Jiang LY, Shen J, Shao A, Ning Y, Jiang J, et al. A methodology for a scalable, collaborative, and resource-efficient platform, MERLIN, to facilitate healthcare AI research. *IEEE J Biomed Health Inform*. 2023;27(8):3971-82.
<https://doi.org/10.1109/JBHI.2023.3250863>.

Kashyap S, Morse KE, Thom B, Lindenauer PK, Haessler S, Pekow PS, et al. A survey of extant organizational and computational setups for deploying predictive models in health systems. *J Am Med Inform Assoc*. 2021;28(11):2445-50.

Kierner S, Weber B, Kadioglu D. Taxonomy of hybrid architectures involving rule-based reasoning and machine learning in clinical decision systems: a scoping review. *J Biomed Inform*. 2023;144:104428.
<https://doi.org/10.1016/j.jbi.2023.104428>.

Liao F, Adelaine S, Afshar M, Patterson BW. Governance of clinical AI applications to facilitate safe and equitable deployment in a large health system: key elements and early successes. *Front Digit Health*. 2022;4:931439.
<https://doi.org/10.3389/fdgth.2022.931439>.

Morley J, Elhalal A, Garcia F, Delecraz S, Savoldi L, Samerski B, et al. Governing data and artificial intelligence for health care: developing an international understanding. *JMIR Form Res*. 2022;6(1):e31623.
<https://doi.org/10.2196/31623>.

Sanchez M, Martel G, Kotter E, Baird GL, Wang M, Stein DJ, et al. AI-clinician collaboration via disagreement prediction: a decision pipeline and retrospective analysis of real-world radiologist-AI interactions. *Cell Rep Med*. 2023;4(10):101207.
<https://doi.org/10.1016/j.xcrm.2023.101207>.

Vollmuth P, Foltyn-Dumitru M, Kessler T, Kollmann B, Weller M, Ernst M, et al. Artificial intelligence (AI)-based decision support improves reproducibility of tumor response assessment in neuro-oncology: an international multi-reader study. *Neuro Oncol*. 2023;25(3):533-43.

Afshar M, Adelaine S, Resnik F, Mundt MP, Fralick J, Aron DC, et al. Deployment of real-time natural language processing and deep learning clinical decision support in the electronic health record: pipeline implementation. *JMIR Med Inform*. 2023;11:e44930.
<https://doi.org/10.2196/44930>.

Joyce C, Joyce B, Afshar M. The evaluation of a clinical decision support tool using natural language processing to screen hospitalized adults for unhealthy substance use: protocol for a quasi-experimental design. *JMIR Form Res*. 2022;6(11):e42971.
<https://doi.org/10.2196/42971>.

Kanbar LJ, Shalish W, Precup D, Brown K, Shalish W, Sant'Anna G, et al. Implementation of machine learning pipelines for clinical practice: development and validation study. *JMIR Med Inform*. 2022;10(11):e37808.
<https://doi.org/10.2196/37808>.

Blezek DJ, Erickson BJ. AI integration in the clinical workflow. *Radiol Artif Intell.* 2021;3(6):e210031.
<https://doi.org/10.1148/ryai.2021210031>.

Juluru K, Shih HH, Kothary N, Celi LA. Integrating AI algorithms into the clinical workflow. *Radiol Artif Intell.* 2021;3(6):e210131.
<https://doi.org/10.1148/ryai.2021210131>.

Pierce JD, Rosipko B, Cobb J, Dubin J, Rosman DA. Seamless integration of artificial intelligence into the clinical environment: our experience with a novel pneumothorax detection artificial intelligence algorithm. *J Am Coll Radiol.* 2021;18(10):1499-506.
<https://doi.org/10.1016/j.jacr.2021.07.017>.

Li X, Ma J, Liu Y, Li J, Li J. MoGCN: a multi-omics integration method based on graph convolutional network for cancer subtype analysis. *Front Genet.* 2022;13:806842.
<https://doi.org/10.3389/fgene.2022.806842>.

Li Y, Ma J, Zhang Y. Graph neural network-based diagnosis prediction. *IEEE Trans Knowl Data Eng.* 2020;33(8):3225-35.
<https://doi.org/10.1109/TKDE.2020.2992878>.

Lu C, Reddy CK, Ning Y. Self-supervised graph learning with hyperbolic embedding for temporal health event prediction. *IEEE Trans Cybern.* 2023;53(6):3694-707.
<https://doi.org/10.1109/TCYB.2021.3131156>.

Wu T, Sun C, Zhong Z, Huang J. Leveraging graph-based hierarchical medical entity embedding for healthcare applications. *Sci Rep.* 2021;11(1):6858.
<https://doi.org/10.1038/s41598-021-86315-1>.

Rajabi E, Etminani K. Towards a knowledge graph-based explainable decision support system in healthcare. *Stud Health Technol Inform.* 2021;285:103-8.
<https://doi.org/10.3233/SHTI210554>.

Rajkomar A, Oren E, Chen K, Dai AM, Hajaj N, Hardt M, et al. Scalable and accurate deep learning with electronic health records. *NPJ Digit Med.* 2018;1(1):18.
<https://doi.org/10.1038/s41746-018-0029-1>.

Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. *Nat Med.* 2019;25(1):44-56.
<https://doi.org/10.1038/s41591-018-0300-7>.

Esteva A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau HM, et al. Dermatologist-level classification of skin cancer with deep neural networks. *Nature.* 2017;542(7639):115-8.
<https://doi.org/10.1038/nature21056>.

Holzinger A, Haibe-Kains B, Jurisica I. Why imaging data alone is not enough: AI-based integration of imaging, omics, and clinical data. *Eur J Nucl Med Mol Imaging.* 2019;46(13):2722-30.
<https://doi.org/10.1007/s00259-019-04382-9>.

Xiao C, Choi E, Sun J. Opportunities and challenges in developing deep learning models using electronic health records data: a systematic review. *J Am Med Inform Assoc.* 2018;25(10):1419-28.