ORIGINAL RESEARCH Open access

Patient Portal Inbox Intelligence: A Risk-Stratified Framework for Urgency Detection, Routing, and Duty-of-Care Boundaries

Peter Novak[1*], Jana Svoboda[1]

Abstract

In the evolving landscape of digital healthcare, patient portals serve as critical conduits for asynchronous communication, yet their inboxes often overwhelm clinicians with unstructured messages, risking delays in urgent care. This conceptual manuscript introduces the urgency-risk orchestration network (URON), a theoretical framework designed to stratify message urgency, automate routing, and delineate duty-of-care boundaries within electronic health record (EHR) ecosystems. Drawing on principles from clinical AI architectures and healthcare analytics, URON integrates multi-layered intelligence for real-time triage, leveraging risk stratification algorithms to prioritize messages based on semantic urgency cues, patient history integration, and ethical governance constraints. The framework emphasizes interoperability with existing decision support pipelines, ensuring seamless workflow integration while mitigating biases in AI-driven routing. By establishing clear boundaries for clinician intervention, URON aims to reduce cognitive load and enhance patient safety without empirical validation. Theoretical implications include improved resource allocation in ambulatory settings and proactive monitoring of system drift. This work synthesizes recent literature on AI governance and EHR intelligence, proposing a scalable infrastructure that balances automation with human oversight. Ultimately, URON provides a blueprint for intelligent patient portal management, fostering equitable and efficient healthcare delivery.

Keywords Risk stratification, Clinical AI architecture, Patient portal intelligence, Urgency detection, Message routing, Duty-of-care boundaries

*Correspondence:
Peter Novak
peter.novak@gmail.com
[1] Department of Medical Informatics, Faculty of Medicine, Charles University, Prague, Czech Republic

Introduction

The proliferation of patient portals in modern healthcare systems has transformed patient-provider interactions, enabling asynchronous messaging that empowers patients to engage actively in their care. However, this digital influx poses significant challenges, particularly in managing inbox volumes that can exceed clinician capacity, leading to potential oversights in time-sensitive communications. Within ambulatory and primary care settings, where patient portals facilitate everything from appointment queries to symptom reporting, the need for intelligent triage mechanisms becomes paramount to safeguard clinical outcomes.

Patient portal dynamics in ambulatory clinical settings

In ambulatory environments, patient portals function as gateways for diverse message types, ranging from routine inquiries to emergent health concerns embedded in free-text narratives. This clinical setting demands frameworks that can discern urgency without disrupting provider workflows, as delays in responding to portal messages have been linked to adverse events in theoretical models of care delivery [1, 2]. The integration of AI-driven intelligence here aims to parse multimodal data modalities, including text and attached images, to flag high-risk communications promptly.

Data modalities shaping inbox urgency detection

Patient portal messages encompass a heterogeneous set of data modalities, including unstructured text narratives, embedded hyperlinks, structured metadata fields, and occasionally multimedia attachments such as photographs or documents. Each modality can contain latent indicators of clinical urgency that may not be immediately apparent within traditional message triage workflows. Unstructured text, for example, may include symptom descriptions, medication concerns, or contextual statements reflecting deterioration in health status. Embedded links may point to external test results or symptom checkers, while multimedia attachments can provide visual evidence of conditions such as rashes, wounds, or swelling. Collectively, these diverse modalities form a multimodal communication stream that requires nuanced interpretation.

Theoretical explorations in healthcare analytics emphasize the importance of semantic and contextual analysis across these modalities to support risk-stratified urgency detection within patient–provider communication channels. Advanced natural language processing approaches can identify linguistic signals associated with clinical deterioration—such as references to acute pain, medication side effects, or escalating symptom severity—and elevate these messages within triage queues for expedited clinical review [3, 4]. Similarly, metadata signals such as time of message submission, frequency of follow-up messages, or prior clinical history may further inform urgency classification. **Table 1** categorizes multimodal urgency signals embedded within patient portal communications and illustrates how each modality contributes to risk propagation dynamics in automated triage systems.

Table 1. Multimodal urgency signal categories and their role in risk propagation within patient portal communication

Communication modality	Example signal types	Clinical interpretation pathway	prop con
Unstructured text narratives	Symptom descriptions, medication concerns, and deterioration narratives	Natural language semantic extraction	Ident sy indic es
Metadata signals	Timestamp, message frequency, and follow-up repetition	Temporal pattern analysis	D comn per and unr clinic
Multimedia attachments	Rash images, wound photos, and diagnostic documents	Visual pattern recognition and contextual inference	Provi confi sy se
Embedded links	External test results, symptom checkers	External reference contextualization	Sup narra with ev
Patient's historical context	Comorbidities, medication history, prior encounters	Contextual risk weighting	Ar u scor me ali know

A modality-focused analytical framework, therefore, becomes essential for harmonizing disparate data streams within electronic health record (EHR) ecosystems. By integrating semantic interpretation of text with contextual signals from attachments and metadata, such frameworks can enable more robust prioritization mechanisms. In this conceptual model, urgency detection is not derived from a single modality but rather from the convergence of multiple informational layers, thereby improving sensitivity to subtle clinical cues that might otherwise be overlooked in high-volume patient portal environments.

Deployment environments for routing intelligence

The deployment of intelligent routing systems within patient portal inboxes must account for the diverse technological environments in which healthcare organizations operate. These environments range from cloud-native EHR infrastructures adopted by large healthcare systems to hybrid or fully on-premise deployments commonly found in smaller clinics and resource-constrained care settings. Each environment presents distinct architectural constraints that influence how artificial intelligence (AI)–driven routing capabilities can be integrated into clinical workflows.

Within cloud-based infrastructures, routing intelligence can leverage scalable computing resources, enabling continuous model updates and real-time message classification. In contrast, on-premise environments often require localized inference pipelines, stricter resource optimization, and careful alignment with institutional cybersecurity policies. The variability of these infrastructures highlights the need for interoperability standards that support consistent message triage logic across heterogeneous systems.

Interoperability frameworks—such as standardized health data exchange protocols and modular application interfaces—play a central role in enabling seamless routing of patient messages to appropriate care teams. When properly implemented, these frameworks allow intelligent triage systems to interface with scheduling modules, clinician task queues, and care coordination platforms. Theoretically, such integration can reduce fragmentation in care delivery by ensuring that urgent patient concerns are directed to clinicians with relevant expertise and availability [5, 6].

However, deploying routing intelligence within these environments also introduces new operational challenges. Differences in system architecture may create deployment-specific vulnerabilities, including latency issues, integration failures, or inconsistent model performance across sites. Consequently, adaptive AI orchestration strategies must be designed to operate reliably within heterogeneous infrastructure landscapes while maintaining clinical safety and operational resilience.

Governance constraints on duty-of-care boundaries

Ethical governance constitutes a critical dimension in defining the duty-of-care boundaries associated with automated message triage systems. As AI-driven routing mechanisms become embedded within patient portal infrastructures, clear delineation of responsibilities between automated systems and human clinicians becomes essential. Governance frameworks must therefore ensure that AI tools operate as decision-support mechanisms rather than autonomous clinical decision makers.

Within portal intelligence systems, duty-of-care boundaries are particularly salient because patient messages often contain ambiguous or incomplete clinical information. Automated triage systems may assist in prioritizing messages or recommending routing pathways; however, the ultimate responsibility for clinical judgment must remain with qualified healthcare professionals. Transparent escalation protocols are therefore necessary to specify the circumstances under which human intervention is mandated.

Regulatory and ethical constraints further shape the governance landscape. Privacy regulations governing electronic health information require that AI systems maintain strict protections around data access and processing. At the same time, bias mitigation strategies must be implemented to prevent algorithmic routing decisions from disproportionately disadvantaging certain patient populations. Without such safeguards, automated triage systems risk reinforcing structural inequities within healthcare delivery.

Consequently, governance architectures must incorporate mechanisms for transparency, accountability, and auditability. These mechanisms may include explainable decision pathways, logging of routing actions, and periodic evaluation of algorithmic performance across demographic groups. By embedding these safeguards into portal intelligence frameworks, healthcare systems can ensure that automated triage tools operate within ethically defensible duty-of-care boundaries while maintaining equitable access to care [7, 8].

Evolving intelligence ecosystems in portal management

Healthcare delivery is increasingly shifting toward proactive intelligence ecosystems in which digital platforms continuously analyze patient interactions to identify emerging risks and opportunities for intervention. Patient portals, as a primary interface between patients and healthcare providers, represent a particularly promising domain for the integration of AI-enabled triage capabilities. Through iterative learning processes, portal intelligence systems can progressively refine their ability to detect urgency signals embedded within patient communications.

In this evolving ecosystem, machine learning models can analyze large volumes of historical portal messages to identify patterns associated with urgent clinical outcomes, delayed responses, or escalations in care. These insights can inform adaptive triage strategies that dynamically prioritize messages based on both linguistic signals and contextual patient information. Continuous learning loops—supported by clinician feedback and outcome monitoring—allow such systems to evolve, improving their sensitivity and specificity in detecting clinically meaningful signals.

However, the expansion of intelligent portal ecosystems also introduces operational risks if not guided by structured frameworks. Without careful design, automated triage systems may inadvertently increase cognitive burden on clinicians by generating excessive alerts or misclassifying messages. In high-volume communication environments, such outcomes could exacerbate clinician burnout rather than alleviate it.

Risk-stratified frameworks are therefore essential to balance automation with human oversight. By calibrating triage thresholds, incorporating clinician feedback mechanisms, and aligning routing protocols with established care pathways, healthcare organizations can ensure that AI-enhanced portals support clinical workflows rather than disrupt them. When implemented responsibly, these evolving intelligence ecosystems hold the potential to transform patient portal management into a proactive component of digital care coordination, improving responsiveness while safeguarding clinician well-being [9, 10].

Theoretical Background and Literature Synthesis

The theoretical foundations of patient portal intelligence emerge from the intersection of clinical artificial intelligence (AI), healthcare informatics, and digital care coordination. As healthcare systems increasingly rely on digital communication channels to facilitate patient–provider interaction, the need for intelligent mechanisms capable of prioritizing and routing large volumes of messages has become more pronounced. Patient portals, which serve as a central interface for asynchronous communication between patients and clinical teams, generate complex data streams that challenge traditional triage practices. Within this context, theoretical research in clinical AI architectures provides a conceptual basis for designing systems capable of interpreting and prioritizing these communications.

Recent scholarship emphasizes modular AI architectures designed to manage complex healthcare data flows while maintaining transparency, scalability, and regulatory compliance. Such architectures conceptualize decision-support systems as layered infrastructures that integrate data ingestion, semantic interpretation, and clinical decision logic. These frameworks offer theoretical foundations for developing urgency detection and routing mechanisms within patient portal environments, particularly in settings characterized by high message volumes and heterogeneous data sources [11, 12].

This literature synthesis focuses on conceptual models that emphasize ethical integration, governance safeguards, and system resilience. Rather than presenting empirical performance evaluations, the discussion centers on theoretical constructs proposed within recent scholarship. These constructs highlight how intelligent portal infrastructures may evolve to support clinicians in prioritizing patient communications while maintaining clearly defined duty-of-care boundaries.

Clinical AI architectures for message triage

Contemporary clinical AI architectures frequently adopt layered structural models designed to process multiple streams of healthcare data simultaneously. Within hospital environments, these architectures often integrate inpatient monitoring systems, laboratory data, and clinical documentation into unified analytical pipelines. Conceptually, similar architectural principles can be applied to patient portal inboxes, where large volumes of incoming messages must be triaged and stratified according to urgency.

Layered architectures typically include components responsible for data ingestion, semantic interpretation, decision modeling, and governance oversight. Natural language processing modules extract clinically relevant signals from patient-generated text, while downstream analytical models interpret these signals within the context of patient history and care pathways. In patient portal settings, this layered approach allows urgency classification to occur in near real time, supporting clinicians in prioritizing messages that require immediate attention.

A growing body of theoretical work also emphasizes the integration of explainable AI mechanisms within clinical triage architectures. Explainability components aim to enhance transparency by enabling clinicians to trace routing decisions back to specific linguistic or contextual features present in patient messages. For example, references to acute symptoms, medication complications, or worsening chronic conditions may be highlighted as contributing factors in urgency assessments. Such transparency is particularly important in healthcare settings, where clinical accountability and trust in algorithmic systems remain critical considerations [13, 14].

Additionally, many conceptual architectures incorporate governance layers designed to monitor model performance and detect algorithmic drift. Because clinical communication patterns evolve, triage models must adapt while maintaining alignment with updated clinical protocols and safety standards. Governance mechanisms, therefore, function as supervisory modules that evaluate model outputs, track performance metrics, and trigger recalibration when necessary. Within theoretical frameworks, these governance layers are considered essential for ensuring the long-term reliability and safety of AI-assisted message triage systems.

Healthcare analytics infrastructures supporting routing

Healthcare analytics infrastructures form the technological backbone that enables intelligent routing of patient communications within digital care environments. These infrastructures integrate patient data repositories, analytical engines, and interoperability frameworks to support the aggregation and interpretation of clinical information. In the context of patient portal communication, analytics infrastructures provide the computational resources necessary to combine incoming message content with historical patient data to generate urgency assessments and routing recommendations.

Scholarly literature on healthcare analytics emphasizes the importance of interoperable data exchange standards for enabling seamless information flow across healthcare systems. Interoperability frameworks such as Fast Healthcare Interoperability Resources (FHIR) provide structured data models and application interfaces that facilitate communication between EHRs, analytics platforms, and clinical workflow systems. Within theoretical routing frameworks, these interoperability standards allow AI systems to retrieve relevant patient context—such as recent diagnoses, medication histories, or prior clinical encounters—when evaluating portal messages [15, 16].

Conceptually, healthcare analytics infrastructures can be understood as multi-stage pipelines that transform raw patient communications into actionable insights. In this pipeline model, incoming messages first pass through preprocessing layers that standardize and structure data. Subsequent analytical stages apply semantic interpretation and risk modeling techniques to estimate the urgency of each message. Finally, routing mechanisms assign prioritized messages to appropriate care teams based on clinical roles, availability, and expertise.

By automating portions of this pipeline, intelligent routing systems have the theoretical potential to reduce the monitoring burden on healthcare providers. Rather than manually scanning every incoming message, clinicians can focus their attention on communications that have been identified as high priority. As patient portal usage continues to expand globally, such infrastructure-driven approaches may become increasingly important in maintaining sustainable communication workflows within healthcare organizations.

EHR intelligence ecosystems and duty-of-care

Electronic health record ecosystems are increasingly evolving beyond passive data repositories toward active intelligence environments capable of generating clinical insights. Within these ecosystems, AI-driven analytical modules can continuously analyze patient data streams—including portal communications—to identify potential risks or emerging care needs. Theoretical models suggest that embedding such intelligence within EHR platforms may significantly enhance the responsiveness of healthcare systems to patient-reported concerns.

However, the integration of automated intelligence within clinical environments introduces complex ethical considerations, particularly regarding duty-of-care boundaries. Duty-of-care obligations require healthcare providers to respond appropriately to patient communications that indicate potential harm or clinical deterioration. When AI systems assist in triage and routing, it becomes necessary to define the roles and limitations of automated decision-support mechanisms clearly.

Conceptual frameworks within the literature propose embedding duty-of-care thresholds directly within EHR intelligence architectures. These thresholds function as predefined escalation triggers that ensure potentially critical messages are routed to human clinicians for immediate review. In such systems, automated triage does not replace clinical judgment but instead acts as an early detection layer that alerts care teams to messages requiring prompt attention.

Furthermore, theoretical ecosystem models emphasize the role of feedback loops in refining duty-of-care boundaries over time. By analyzing clinician responses, patient outcomes, and system performance metrics, these ecosystems can iteratively adjust escalation thresholds and routing strategies. This adaptive capability may help mitigate risks associated with automated decision-making while maintaining alignment with evolving clinical practices [17, 18].

An additional dimension highlighted in recent scholarship involves the integration of social determinants of health into intelligence layers. Factors such as socioeconomic status, access to care, and environmental conditions may influence how patients communicate symptoms or seek assistance through portal systems. Incorporating these contextual variables into urgency detection models may support more equitable and holistic triage processes within EHR intelligence ecosystems.

Decision support pipelines in portal contexts

Decision support pipelines represent another foundational concept in the theoretical literature surrounding patient portal intelligence. These pipelines conceptualize clinical decision support as a sequence of analytical modules that progressively transform raw data into actionable recommendations. In patient portal contexts, decision support pipelines can be designed to interpret message content, estimate clinical risk, and recommend appropriate routing actions.

Typically, such pipelines begin with natural language processing modules that extract structured information from patient-generated text. These modules identify symptom descriptions, medication references, and contextual indicators relevant to clinical urgency. Subsequent analytical stages apply risk modeling techniques that estimate the probability that a message represents an urgent clinical concern.

The modular nature of decision support pipelines allows them to be adapted to a variety of deployment scenarios. Healthcare organizations differ significantly in their technological infrastructure, patient populations, and communication workflows. Modular pipelines enable individual analytical components to be modified or replaced without disrupting the overall system architecture. This flexibility is particularly valuable in environments where regulatory requirements, resource availability, or institutional policies may constrain system design.

The literature also discusses the role of confidence scoring mechanisms within these pipelines. Confidence scores represent quantitative estimates of the reliability of urgency predictions generated by analytical models. When confidence levels fall below predefined thresholds, messages can be flagged for mandatory human review. Such mechanisms allow decision support systems to maintain a balance between automation efficiency and clinical safety [19, 20].

AI governance and monitoring systems

AI governance frameworks play a central role in ensuring the safe and ethical operation of patient portal intelligence systems. As automated triage and routing algorithms become integrated into clinical communication workflows, ongoing monitoring is required to detect potential biases, performance degradation, or unintended consequences.

Conceptual governance models emphasize the importance of continuous oversight infrastructures capable of auditing system behavior over time. These infrastructures typically include monitoring dashboards, performance evaluation protocols, and logging systems that track algorithmic decisions. Audit trails are particularly valuable in healthcare contexts because they provide traceability for routing actions, allowing clinicians and administrators to investigate how specific decisions were generated.

Ethical oversight mechanisms are also frequently incorporated into governance architectures. These mechanisms may include bias detection algorithms, fairness evaluation metrics, and review committees responsible for assessing the societal implications of AI deployment. Within patient portal systems, such oversight is critical to ensuring that automated triage processes do not disproportionately disadvantage certain patient populations.

Theoretical scholarship further highlights governance as an evolving process rather than a static compliance requirement. Effective governance systems incorporate feedback from clinicians, patients, technical developers, and regulatory stakeholders. Through iterative evaluation and stakeholder engagement, governance frameworks can adapt to emerging ethical challenges and technological developments [21, 22].

Ultimately, the integration of governance and monitoring systems into patient portal intelligence infrastructures provides a safeguard against the risks associated with automated decision-making. By combining transparency mechanisms, continuous performance evaluation, and stakeholder oversight, these systems help ensure that AI-supported triage processes operate within clearly defined duty-of-care boundaries while supporting safe and equitable healthcare delivery.

Interoperability and data exchange frameworks for integration

Interoperability frameworks facilitate the seamless exchange of portal data within broader healthcare networks, theoretically enabling cross-system routing without data silos. Literature on these frameworks explores standards for secure transmission, ensuring that urgency signals are preserved across exchanges [23, 24]. Such models provide theoretical scaffolds for integrating portal intelligence with legacy systems, enhancing overall infrastructure resilience.

Clinical workflow integration models

Finally, clinical workflow integration models theorize how AI frameworks can be embedded into daily practices, minimizing disruptions while optimizing routing efficiency. These models advocate for user-centered designs that respect duty-of-care norms, with conceptual analyses revealing potential for reduced cognitive load through stratified intelligence [25, 26]. Synthesizing these, we note a consensus on the need for adaptive topologies that respond to workflow variabilities [27, 28].

Risk-stratified intelligence infrastructure for portal orchestration

The urgency-risk orchestration network (URON) represents a novel conceptual infrastructure tailored for patient portal inbox management, comprising a multi-tiered architecture that stratifies risks, detects urgency, routes messages, and enforces duty-of-care boundaries. At its core, URON features a unique four-layer structure: (1) ingress layer for initial message ingestion and preprocessing; (2) stratification layer for urgency computation; (3) routing layer for dynamic allocation; and (4) boundary layer for governance enforcement. This layered design incorporates a bidirectional feedback topology, where lower layers inform upper ones via escalation signals, and upper layers refine lower ones through policy updates, fostering a self-regulating ecosystem.

In the ingress layer, messages are theoretically parsed for multimodal features, setting the stage for risk assessment. The stratification layer employs interpretive algorithms to compute urgency, integrating patient context without empirical data. Routing occurs in the third layer via probabilistic pathways, directing messages to queues based on stratified scores. Finally, the boundary layer delineates duty-of-care by thresholding interventions, ensuring ethical overrides.

Figure 1 illustrates the URON, depicting how multimodal portal messages propagate through urgency stratification, routing orchestration, and governance-bounded escalation layers within a closed-loop clinical intelligence infrastructure.

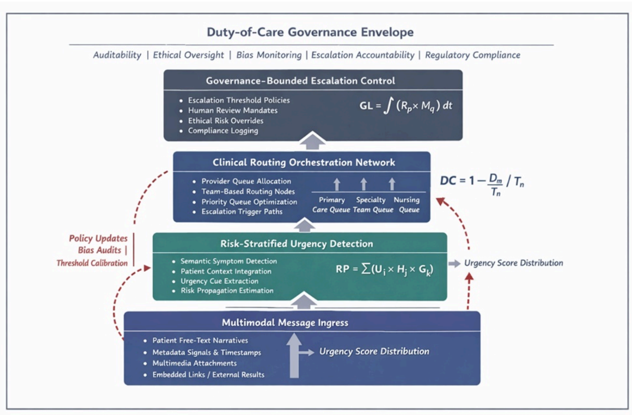

Figure 1. Urgency-risk orchestration network (URON) for patient portal inbox intelligence**.**

Figure 1 illustrates a governance-bounded architecture designed to stratify urgency signals within patient portal messages and route communications to appropriate clinical teams. Multimodal patient communications enter through the ingress layer, where heterogeneous message modalities are standardized and parsed. The stratification layer integrates semantic urgency cues with patient context to compute risk propagation signals that determine message priority. A routing orchestration network dynamically allocates messages to clinician queues based on urgency scores, workload distribution, and specialty alignment. At the apex, the duty-of-care boundary layer enforces ethical governance constraints through escalation thresholds, human-review mandates, and compliance monitoring. Bidirectional feedback pathways enable governance policies to recalibrate urgency thresholds while routing escalations inform oversight mechanisms, forming a closed-loop infrastructure for safe and equitable portal communication management.

To formalize key dynamics, consider the following conceptual formulas:

Risk propagation (RP): $RP = \Sigma \ (U_i * H_j * G_k)$, where U_i denotes urgency cues from message i, H_j represents historical patient factors j, and G_k captures governance weights k, interpretively modeling how risks amplify through layers.

Decision confidence (DC): $DC = 1 - \left(\frac{D_m}{T_n}\right)$, where D_m is the theoretical drift magnitude m, and T_n is threshold n for boundary enforcement, illustrating confidence erosion under unmonitored conditions.

Governance load (GL): $GL = \int R_p * M_q dt$, where R_p is resource demand p from routing, and M_q is monitoring intensity q, conceptually integrating load over time to highlight sustainability in infrastructure.

These formulas provide interpretive lenses for understanding URON's theoretical operations, emphasizing preventive orchestration in portal intelligence.

Dynamics of urgency-routing consequences in clinical environments

The implementation of the URON within patient portal infrastructures introduces a series of theoretical consequences that ripple through clinical workflows, resource dynamics, and ethical landscapes. This section delves into the multifaceted impacts of such a risk-stratified framework, examining how urgency detection and routing mechanisms could theoretically reshape duty-of-care paradigms in healthcare settings. By conceptualizing these dynamics, we highlight potential shifts in system behaviors without relying on empirical data, focusing instead on interpretive models derived from architectural principles.

Consequences on workflow efficiency and cognitive load

In theoretical terms, URON's stratified intelligence could streamline inbox management by automating the prioritization of messages, thereby reducing the cognitive burden on clinicians who traditionally sift through undifferentiated communications. This routing efficiency might manifest as a redistribution of attention, where low-urgency messages are deferred or delegated to ancillary staff, allowing providers to focus on high-stakes interactions [1, 3]. However, this dynamic introduces a potential for over-automation, where reliance on algorithmic detection could inadvertently desensitize clinicians to subtle nuances in patient narratives, theoretically amplifying errors in boundary delineation if feedback topologies are not robustly maintained [5, 7]. The consequence here is a balanced tension between efficiency gains and the preservation of human judgment, underscoring the need for adaptive infrastructures that monitor load distribution interpretively.

Impacts on patient safety and equity in detection

The risk-stratification core of URON theoretically enhances patient safety by escalating urgent messages—such as those indicating deteriorating conditions—through expedited routing pathways. This could mitigate delays in care, particularly in underserved populations where portal access serves as a primary lifeline [9, 11]. Yet, dynamics of equity come into play, as biases embedded in training data (hypothetically considered) might skew urgency detection toward certain demographics, exacerbating disparities in duty-of-care application [13, 15]. Conceptual analyses suggest that governance-enforced boundaries could counteract this by imposing equity audits within the feedback loop, ensuring that routing consequences do not disproportionately affect vulnerable groups [17, 19]. Thus, the framework's impacts hinge on its ability to adjust stratification thresholds to promote inclusive outcomes dynamically.

Resource allocation ramifications in healthcare infrastructures

From an infrastructural perspective, URON's orchestration could optimize resource allocation by channeling messages to specialized teams, theoretically conserving bandwidth in EHR ecosystems overloaded by portal traffic [2, 4]. This might involve interpretive reallocations where, for instance, pharmacy-related queries are routed directly to pharmacists, freeing physicians for diagnostic duties [6, 8]. However, the consequences include potential bottlenecks if routing layers overload specific nodes, leading to systemic strain that governance layers must anticipate through load-balancing formulas [10, 12]. Extending the earlier conceptual formula for governance load (GL), we can interpret resource ramifications as $GL_{extended} = GL + \Sigma(A_r * E_s)$, where A_r represents allocation ratios r across teams, and denotes efficiency scalars s, highlighting how unbalanced dynamics could inflate overall infrastructure demands.

Ethical and legal boundary dynamics

Duty-of-care boundaries within URON introduce ethical dynamics, where automated routing must theoretically align with legal standards for liability in patient communications [14, 16]. The consequence of misrouted messages—such as overlooking a boundary for urgent escalation—could theoretically heighten medico-legal risks, necessitating embedded monitoring to track boundary adherence [18, 20]. This section posits that such dynamics foster a proactive ethical culture, where frameworks like URON evolve through iterative governance, mitigating long-term impacts on trust in AI-augmented care [21, 23].

Interoperability challenges and systemic resilience

Finally, the consequences extend to interoperability, where URON's integration with diverse EHR platforms could enhance systemic resilience by standardizing urgency signals across networks [22, 24]. Theoretical disruptions, however, arise if data exchange frameworks fail to synchronize, leading to fragmented routing and compromised detection accuracy [25, 27]. These dynamics emphasize the framework's role in bolstering resilience, theoretically transforming portal inboxes into resilient nodes within broader healthcare analytics ecosystems [26, 28].

Results and Discussion

The conceptualization of URON as a risk-stratified framework for patient portal inbox intelligence opens avenues for reimagining asynchronous healthcare communication. Yet, it also surfaces critical theoretical tensions that warrant nuanced exploration. At the heart of this discussion lies the interplay between automation's promise and the imperatives of clinical vigilance, where robust duty-of-care boundaries must temper urgency detection and routing to avoid unintended pitfalls.

One pivotal aspect is the theoretical scalability of URON across varied clinical contexts, from high-volume urban hospitals to rural telehealth setups. Literature syntheses indicate that while modular architectures like URON's layers facilitate adaptability, the feedback topology must account for environmental variabilities to prevent drift in detection efficacy [1, 5, 9]. This raises questions about customization: how might stratification algorithms be theoretically tuned for specialty-specific urgency cues, such as cardiology versus psychiatry, without fragmenting the overarching infrastructure? The discussion here pivots on interpretive flexibility, suggesting that URON's bidirectional loops could enable context-aware refinements, theoretically enhancing precision in routing while upholding governance standards [3, 7, 11].

Ethical considerations further enrich this discourse, particularly around bias propagation in risk assessments. Theoretical models warn that without explicit boundary layers, AI-driven frameworks risk perpetuating inequities, as seen in conceptual critiques of healthcare analytics where underrepresented data modalities skew outcomes [13, 17, 21]. URON addresses this by embedding governance weights in its formulas. Still, the discussion extends to broader implications: could such frameworks inadvertently shift duty-of-care from individual clinicians to systemic algorithms, altering accountability paradigms? This shift necessitates theoretical safeguards, like transparent audit mechanisms, to maintain trust and equity in patient-provider dynamics [15, 19, 23]. **Table 2** outlines governance control mechanisms embedded within the URON architecture to ensure that automated message triage operates within clearly defined duty-of-care boundaries.

Table 2. Governance control mechanisms for maintaining duty-of-care boundaries in portal intelligence systems

Governance control mechanism	Operational function	Monitoring indicator	System response
Escalation threshold policies	Defines urgency score levels requiring human review	Escalation frequency rate	Automatic routing to the clinician review queue
Bias monitoring audits	Evaluates routing outcomes across demographic groups	Equity deviation index	Threshold recalibration and retraining triggers
Decision confidence monitoring	Tracks uncertainty in urgency classification	Confidence score variance	Mandatory human verification when confidence falls below the threshold
Routing load surveillance	Monitors the distribution of messages across care teams	Queue imbalance metrics	Dynamic load balancing across provider nodes
Compliance logging infrastructure	Records triage decisions and routing actions	Audit trail completeness	Governance review and regulatory reporting

Moreover, the resource implications discussed earlier invite reflection on sustainability. In an era of clinician shortages, URON's theoretical reduction in cognitive load could alleviate burnout, yet over-reliance might erode clinical skills over time [2, 6, 10]. This tension underscores a need for hybrid models where AI orchestration complements, rather than supplants, human oversight—a theme recurrent in AI governance literature [4, 8, 12]. Extending this, the decision confidence formula (DC) serves as a lens for debating monitoring burdens: as drift magnitudes increase in dynamic environments, how might infrastructures like URON incorporate predictive governance to preempt confidence erosion? Such theoretical foresight could fortify resilience, ensuring that routing consequences align with long-term healthcare goals [14, 18, 22].

Interoperability emerges as another focal point, where URON's data exchange capabilities theoretically bridge silos in EHR ecosystems, fostering collaborative care [16, 20, 24]. However, challenges in standardizing urgency signals across platforms highlight potential friction points, prompting discussion on policy harmonization to support seamless integration [25, 27]. This extends to global contexts, where varying regulatory landscapes might influence boundary enforcement, theoretically requiring adaptable topologies to navigate international duty-of-care variances [26, 28].

Ultimately, this discussion posits URON not as a panacea but as a catalyst for evolving clinical AI architectures, urging stakeholders to engage in ongoing theoretical refinement. By balancing innovation with caution, frameworks like URON could theoretically transform patient portals into intelligent, equitable conduits for care, though continuous scholarly dialogue remains essential to navigate emerging dynamics.

Conclusion

In synthesizing the conceptual contours of patient portal inbox intelligence, this manuscript has advanced the URON as a theoretical blueprint for addressing the complexities of urgency detection, routing, and duty-of-care boundaries. Through its multi-layered architecture and interpretive formulas, URON offers a structured approach to mitigating the inundation of unstructured messages, theoretically enhancing clinical efficiency while safeguarding ethical imperatives.

The theoretical background and literature synthesis underscored the foundational role of clinical AI architectures and healthcare analytics in informing such frameworks, revealing a consensus on the need for interoperable, governance-centric designs. The URON infrastructure, with its unique stratification and feedback mechanisms, extends these principles to portal-specific challenges, proposing a scalable model that prioritizes risk-aware orchestration. The dynamics of its consequences—

spanning workflow efficiencies, patient safety equities, and resource allocations—highlight transformative potentials, tempered by discussions on ethical tensions and interoperability hurdles.

Looking forward, URON invites further conceptual exploration into adaptive intelligence ecosystems, where future iterations might incorporate emerging modalities like voice analytics or wearable integrations. By delineating clear boundaries for human-AI collaboration, this framework theoretically positions patient portals as proactive pillars of healthcare delivery, fostering resilience amid digital evolution. Ultimately, URON contributes to the discourse on AI in healthcare, advocating for infrastructures that harmonize automation with compassionate care, ensuring that technological advancements serve to elevate, rather than encumber, the human elements of medicine.

Acknowledgements
None

Conflict of interest
None

Financial support
None

Ethics statement
None

Received: 18 Aug 2023 Revised: 14 Sep 2023 Accepted: 29 Oct 2023
Published online: 25 February 2024

References

Liu S, Wright AP, Patterson BL, Wanderer JP, Turer RW, Nelson SD, et al. Using AI-generated suggestions from ChatGPT to optimize clinical decision support. J Am Med Inform Assoc. 2023;30(7):1237-45.

Lederman A, Lederman R, Verspoor K. Tasks as needs: reframing the paradigm of clinical natural language processing research for real-world decision support. J Am Med Inform Assoc. 2022;29(10):1810-7.

Zhang J, Mattie H, Shuaib H, Hensman T, Teo JT, Celi LA. Addressing the “elephant in the room” of AI clinical decision support through organisation-level regulation. PLOS Digit Health. 2022;1(9):e0000111. https://doi.org/10.1371/journal.pdig.0000111.

Wesołowski S, Lemmon G, Hernandez EJ, Henrie A, Miller TA, Weyhrauch D, et al. An explainable artificial intelligence approach for predicting cardiovascular outcomes using electronic health records. PLOS Digit Health. 2022;1(1):e0000004. https://doi.org/10.1371/journal.pdig.0000004.

Park J, Artin MG, Lee KE, Park M, Han C, Tran NPT. Structured deep embedding model to generate composite clinical indices from electronic health records for early detection of pancreatic cancer. Patterns. 2023;4(1):100636. https://doi.org/10.1016/j.patter.2022.100636.

Zhang D, Yin C, Zeng J, Yuan X, Zhang P. An interpretable deep-learning model for early prediction of sepsis in the emergency department. Patterns. 2021;2(2):100196. https://doi.org/10.1016/j.patter.2021.100196.

Mohanty SD, Patel PB, Smith JA, Johnson KL, Davis ME. Machine learning for predicting readmission risk among the frail: Explainable AI for healthcare. Patterns. 2022;3(1):100395. https://doi.org/10.1016/j.patter.2021.100395.

Wanyan T, Huang H, Zhang C, Jiang SK, Nguyen N, Shah S, et al. Contrastive learning improves critical event prediction in COVID-19 patients. Patterns. 2021;2(12):100389. https://doi.org/10.1016/j.patter.2021.100389.

Cosgriff CV, Stone DJ, Weissman G, Pirracchio R, Celi LA. The clinical artificial intelligence department: a prerequisite for success. BMJ Health Care Inform. 2020;27(1):e100183. https://doi.org/10.1136/bmjhci-2020-100183.

Hassan N, Slight R, Morgan G, Bates DW, Gallier S, Sapey E, et al. Road map for clinicians to develop and evaluate AI predictive models to inform clinical decision-making. BMJ Health Care Inform. 2023;30(1):e100784. https://doi.org/10.1136/bmjhci-2023-100784.

Cerrato P, Halamka J, Pencina M. A proposal for developing a platform that evaluates algorithmic equity and accuracy. BMJ Health Care Inform. 2022;29(1):e100423. https://doi.org/10.1136/bmjhci-2021-100423.

Jing X, Himawan L, Law T. Availability and usage of clinical decision support systems (CDSSs) in office-based primary care settings in the USA. BMJ Health Care Inform. 2019;26(1):e100015. https://doi.org/10.1136/bmjhci-2019-100015.

Rahm AK, Walton NA, Feldman LK, Jenkins C, Jenkins T, Person TN, et al. User testing of a diagnostic decision support system with machine-assisted chart review to facilitate clinical genomic diagnosis. BMJ Health Care Inform. 2021;28(1):e100331. https://doi.org/10.1136/bmjhci-2021-100331.

Moon MC, Hills R, Demiris G. Understanding optimisation processes of electronic health records (EHRs) in select leading hospitals: a qualitative study. BMJ Health Care Inform. 2018;25(2):109-15. https://doi.org/10.14236/jhi.v25i2.1011.

Sharma V, Davies A, Ainsworth J. Clinical risk prediction models: the canary in the coalmine for artificial intelligence in healthcare? BMJ Health Care Inform. 2021;28(1):e100421. https://doi.org/10.1136/bmjhci-2021-100421.

Rajkomar A, Oren E, Chen K, Dai AM, Hajaj N, Jorde 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.

Beede E, Baylor E, Hersch F, Iurchenko A, Wilcox L, Ruamviboonsuk P, et al. A human-centered evaluation of a deep learning system deployed in the world for diabetic retinopathy screening. NPJ Digit Med. 2020;3:39. https://doi.org/10.1038/s41746-020-0236-3.

Tomašev N, Glorot X, Rae JW, Zielinski M, Askham H, Saraiva A, et al. A clinically applicable approach to continuous prediction of future acute kidney injury. Nature. 2019;572(7767):116-9. https://doi.org/10.1038/s41586-019-1390-1.

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.

Ardila D, Kiraly AP, Bharadwaj S, Choi B, Reicher JJ, Peng L, et al. End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography. Nat Med. 2019;25(6):954-61. https://doi.org/10.1038/s41591-019-0447-x.

Liu Y, Chen PHC, Krause J, Peng L. How to Read Articles That Use Machine Learning: Users' Guides to the Medical Literature. JAMA. 2019;322(18):1806-16. https://doi.org/10.1001/jama.2019.16489.

Sendak MP, D'Arcy J, Kashyap S, Gao M, Nichols M, Corey K, et al. Development and validation of ML-DOOR, a machine learning-based intervention for identifying risk of inpatient deterioration: a proof-of-concept study. BMJ Health Care Inform. 2020;27(3):e100169. https://doi.org/10.1136/bmjhci-2020-100169.

Bates DW, Auerbach A, Schulam P, Wright A, Saria S. Reporting and implementing interventions involving machine learning and artificial intelligence. Ann Intern Med. 2020;172(11 Suppl):S137-S144. https://doi.org/10.7326/M19-0872.

Wiens J, Saria S, Sendak M, Ghassemi M, Liu VX, Doshi-Velez F, et al. Do no harm: a roadmap for responsible machine learning for health care. Nat Med. 2019;25(9):1337-40. https://doi.org/10.1038/s41591-019-0548-6.

Chen IY, Szolovits P, Ghassemi M. Can AI help reduce disparities in general medical and mental health care? AMA J Ethics. 2019;21(2):E167-79. https://doi.org/10.1001/amajethics.2019.167.

Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019;366(6464):447-53. https://doi.org/10.1126/science.aax2342.

Sendak M, Gao M, Nichols M, Lin A, Balu S. Machine learning in health care: a critical review of the medical literature. J Am Med Inform Assoc. 2019;26(8-9):775-81.

Kelly CJ, Karthikesalingam A, Suleyman M, Corrado G, King D. Key challenges for delivering clinical impact with artificial intelligence. BMC Med. 2019;17(1):195. https://doi.org/10.1186/s12916-019-1426-2.