El-Kholy *et al. Journal of Health Informatics and Digital Systems* *(2026) 1:15*
https://doi.org/10.51847/m612899625

ORIGINAL RESEARCH **Open access**

Clinical Decision Latency as a Safety Variable: A Temporal Accountability Framework for Detecting Harmful Delays in Care

Ahmed El-Kholy[1*], Nour Abdelrahman[1], Karim Hassan[2]

Abstract

Clinical decision latency, defined as the temporal interval from the moment actionable clinical data becomes available to the initiation of a corresponding therapeutic or diagnostic action, constitutes an under-recognized yet critical safety variable in contemporary healthcare delivery. Prevailing patient safety paradigms predominantly concentrate on categorical errors of commission or omission while largely treating time as an exogenous operational factor rather than an intrinsic propagative risk element capable of independently driving harm. This conceptual systems article reframes clinical decision latency as a primary, quantifiable, and governable safety variable. It proposes the clinical latency oversight lattice (CLOL)—an original infrastructural framework specifically designed to detect, quantify, assign accountability for, and interrupt harmful temporal delays across care pathways. Drawing on a targeted synthesis of literature that collectively addresses clinical decision support limitations, diagnostic uncertainty propagation, consequences of treatment delays, health IT-induced temporal vulnerabilities, and AI integration challenges, the manuscript argues that latency functions not as mere logistical inefficiency but as a dynamic, modality-sensitive, context-dependent risk multiplier. The CLOL architecture organizes temporal accountability into four interdependent lattice layers linked by a bidirectional feedback topology that enables real-time drift monitoring, explicit actor/system responsibility mapping, safety-variable score propagation, and orchestrated mitigation responses. Three interpretive mathematical expressions capture core dynamics: risk propagation across pathways, exponential decay of decision confidence under accumulating latency, and cumulative governance/monitoring burden. By institutionalizing latency as a traceable safety variable within a closed-loop accountability structure, CLOL offers healthcare analytics and AI system designers a theoretical and architectural foundation for shifting from retrospective error analysis toward prospective temporal harm prevention in high-stakes clinical environments.

Keywords Clinical decision latency, Temporal accountability, Safety variable, Harmful delays, Clinical decision support systems, Artificial intelligence in healthcare

*Correspondence:
Ahmed El-Kholy
ahmed.elkholy@outlook.com
[1] Department of Health Information Systems, Faculty of Medicine, Alexandria University, Alexandria, Egypt
[2] Department of Clinical Informatics, Faculty of Medicine, Ain Shams University, Cairo, Egypt

Introduction

Emergence of clinical decision latency as a primary safety variable in acute care environments

The accelerating integration of continuous physiologic monitoring, automated laboratory reporting, point-of-care imaging, and predictive AI models into acute care settings has dramatically increased the volume and velocity of clinically relevant data available to decision-makers [1-3]. Paradoxically, this data-rich environment has not proportionally reduced preventable harm; instead, numerous studies document that critical therapeutic windows continue to close due to delays between data availability and clinician action [4-6]. In high-acuity domains such as intensive care, trauma resuscitation, sepsis management, and intraoperative decision-making, even relatively short latency periods—measured in minutes to hours—can initiate irreversible physiologic cascades [7-11]. These observations compel a paradigm shift: clinical decision latency should no longer be conceptualized as background process friction or a quality-of-care efficiency indicator, but rather elevated to the status of a primary safety variable whose magnitude, direction, and propagation directly modulate patient risk independent of the correctness of the eventual decision [1, 6]. Treating latency in this manner aligns with emerging systems-thinking approaches to patient safety that seek to identify latent preconditions of harm before adverse events materialize [3, 12-17].

Temporal dimensions of harmful delays across heterogeneous data modalities

Modern electronic health record ecosystems aggregate information from fundamentally dissimilar sources: high-frequency streaming data from bedside monitors (seconds-level resolution), intermittent laboratory results (minutes to hours), radiology reports and images (tens of minutes to hours), and unstructured clinician documentation (variable) [18-25]. Each modality carries its own intrinsic temporal semantics and clinically acceptable latency envelope. Yet most existing decision-support and alerting infrastructures apply uniform or arbitrary time thresholds—if they consider time at all—creating a mismatch that systematically under-detects harmful delays in slower modalities while generating excessive noise in faster ones [5, 12]. For instance, a 45-minute delay in troponin result acknowledgment may represent catastrophic latency in suspected acute coronary syndrome, whereas the same interval is inconsequential for routine thyroid-function testing [4, 14]. A robust temporal accountability system must therefore incorporate modality-stratified, context-aware thresholds that reflect physiologic urgency, evidence-based intervention windows, and severity-adjusted risk gradients rather than calendar time alone [13, 14, 26-29]. Failure to do so renders downstream AI tools—however accurate—vulnerable to propagating already-delayed decisions [10, 24].

Deployment constraints in real-time monitoring infrastructures

Implementing continuous latency surveillance faces formidable infrastructural barriers. Most hospital information systems were architected for episodic documentation and billing rather than real-time temporal auditing [17]. Interoperability standards (e.g., HL7 FHIR, older CDA) rarely mandate sub-minute timestamp precision or latency metadata propagation across components [2, 8]. Alert engines prioritize content relevance over timeliness-of-response tracking. When they do monitor response intervals, they frequently do so only after alerts have fired—missing pre-alert latency entirely [2, 7]. Cloud-based analytics platforms introduce additional latency through data ingestion queues, model inference times, and result delivery. Yet, few incorporate mechanisms to expose or govern these pipeline delays as safety-relevant variables [17, 27]. These architectural legacies collectively render harmful decision delays invisible until downstream harm manifests, at which point root-cause analysis is retrospective and punitive rather than prospective and preventive [18, 20]. A purpose-built temporal accountability layer must therefore operate orthogonally to existing infrastructures, intercepting timestamped events at ingestion points and enforcing accountability without requiring wholesale system replacement.

Governance constraints and accountability imperatives in AI-augmented decision pathways

Current regulatory, accreditation, and institutional governance structures lack explicit constructs for temporal accountability. Adverse-event reporting taxonomies capture “delay in diagnosis/treatment” as a secondary descriptor rather than a root causal pathway [14, 16, 18]. Ethical guidelines for AI in healthcare emphasize explainability, fairness, and bias mitigation but rarely address timeliness as a dimension of trustworthiness [6, 9]. When AI augments human decision-making, latency introduced by model computation, result review, or overridden alerts creates shared-responsibility zones where accountability diffuses [1, 9]. Without a formal governance topology that traces latency accumulation to specific nodes (human, algorithmic, infrastructural), these zones become liability black holes. The proposed framework directly confronts this deficit by instituting a closed-loop accountability lattice that records, scores, and escalates temporal deviations in real time while supporting retrospective forensic auditing and prospective threshold refinement [3, 6]. **Table 1** defines how harmful clinical decision latency differs across data modalities by aligning temporal signal structure, failure mode, accountability locus, and governance requirement within the CLOL framework.

Table 1. Temporal accountability matrix across clinical data modalities, urgency contexts, and latency failure modes

Clinical data modality	Temporal signal characteristics	Typical decision dependency	s
Continuous physiologic monitoring	High-frequency and second-level streaming data	Immediate recognition of physiologic instability	V de
Discrete laboratory results	Intermittent, batch-like release with timestamped completion	Confirmation or exclusion of emergent diagnostic/treatment decisions	H s
Imaging/report outputs	Acquisition-to-interpretation-to-action sequence	Escalation of treatment, triage change, or procedural intervention	M v d o
Unstructured clinical documentation	Variable, asynchronous authoring and retrieval	Context consolidation, handoff continuity, and deferred decision support	I cu i
AI/CDSS output streams	Event-triggered inference with processing and review delay	Prioritization, prediction-assisted action, or override support	H wi
Referral / cross-service communication	Episodic, organizationally mediated transfer events	Specialist evaluation, procedural authorization, or cross-team decision progression	fr p

These interlocking dimensions—acute-care urgency, modality heterogeneity, infrastructural opacity, and governance vacuum—collectively justify the development of a dedicated temporal accountability architecture. The following sections synthesize the evidentiary foundation and present the clinical latency oversight lattice (CLOL) as a coherent systems-level response.

Theoretical Background and Literature Synthesis

Limitations of conventional clinical decision support in capturing temporal risk

Clinical decision support systems (CDSS) have matured considerably since the mid-2010s, with substantial literature addressing alert fatigue, override rates, and content appropriateness [2, 7]. However, temporal performance remains a glaring blind spot. Multiple studies document that even well-tuned alerts frequently fail to alter clinician behavior when the time between alert delivery and required action exceeds narrow safety windows [2]. In anesthesia and perioperative domains, intraoperative CDSS tools detect physiologic deviations promptly but suffer from silent accumulation of response latency until critical events occur [11]. Surgical AI-augmented decision-making similarly reveals that model outputs, while accurate, arrive too late in fast-moving scenarios to influence outcomes meaningfully [1]. These patterns indicate that conventional CDSS architectures treat time as an external constraint rather than an internal variable requiring active governance—leaving latency-induced harm unaddressed even when content-level support functions optimally [6, 9].

Diagnostic uncertainty and the propagation of temporal errors

Diagnostic error remains one of the most prevalent and costly forms of medical harm, with delay consistently identified as a leading subtype [14, 16]. Emergency department reviews show that uncertainty compounds nonlinearly when temporal windows for confirmatory testing (imaging, labs) are missed, converting ambiguous presentations into definitive adverse events [4]. In pediatrics and ambulatory care, fragmented referral loops—where latency between initial suspicion and specialist evaluation exceeds governance limits—produce cascading diagnostic failures that evade standard safety classification schemes [16, 18]. Perhaps most compellingly, large-scale meta-analyses of antifibrinolytic administration in acute severe bleeding establish a clear temporal dose-response curve: mortality reduction degrades continuously and measurably with each incremental minute of delay [13]. Together, these bodies of work elevate delay from an incidental factor to a mechanistically central driver of harm propagation, warranting its reconceptualization as a standalone safety variable amenable to targeted infrastructural intervention [14, 18].

Artificial intelligence applications and the amplification of latency drift

Recent AI applications in healthcare showcase impressive discriminative performance across domains: continuous cardiac-signal classification, multiparametric prostate MRI interpretation, longitudinal EHR trajectory reconstruction, real-time antibiotic dosing, and acute-kidney-injury early warning [24–27, 29]. Yet virtually every implementation inherits the temporal vulnerabilities of the underlying data pipelines and human–machine interfaces. When inference latency, result delivery delay, or clinician review time push total decision latency beyond clinical thresholds, superior model accuracy provides little protective value [10, 24]. Moreover, AI can inadvertently exacerbate drift: automated alerts may desensitize users (increasing response latency), while black-box outputs may prolong interpretive time compared with simpler rules-based systems [9]. These dynamics underscore an urgent architectural requirement: any AI augmentation layer must be nested within a latency-aware governance superstructure capable of monitoring, scoring, and mitigating temporal degradation end-to-end [6, 10, 25].

Health information technology and systemic contributors to undetected delays

Health IT-related adverse events frequently involve temporal dimensions that current reporting systems under-specify [17]. Configuration errors, interface incompatibilities, alert suppression logic, and data-synchronization failures all inflate end-to-end decision latency without triggering conventional safety alarms [17, 20]. Patient-safety culture assessments link inadequate temporal monitoring practices to elevated diagnostic error rates across settings [20]. Conversely, diagnostic-stewardship programs that enforce strict timing rules for high-risk tests (e.g., Clostridioides difficile assays) achieve significant reductions in unnecessary testing and downstream harm, indirectly demonstrating the protective potential of enforced temporal boundaries [22, 23]. Spanning critical care, radiology error analysis, nephrology prediction tools, infectious-disease stewardship, and oncology diagnostic pathways, this diverse evidence base converges on a single structural conclusion: undetected or ungoverned clinical decision latency functions as a systemic precondition for preventable harm. It therefore demands dedicated architectural attention [1, 3, 6, 9, 17].

Orchestrating temporal accountability: the clinical latency oversight lattice for harm prevention

The clinical latency oversight lattice (CLOL) constitutes a purpose-designed, infrastructural response to the temporal accountability deficit identified above. Unlike conventional linear pipelines or hierarchical CDSS architectures, CLOL organizes temporal governance as a four-layer lattice whose nodes and edges explicitly represent decision points, latency vectors, accountability mappings, and bidirectional feedback flows.

Layer 1 – Latency sensing and multimodal ingestion

All incoming clinical data streams—continuous vital signs, discrete laboratory values, imaging metadata with acquisition timestamps, free-text notes with dictation times—are captured at source with high-fidelity timestamps [25, 27]. Dedicated ingestion adapters normalize heterogeneous update frequencies into a common temporal-event format, producing per-decision-node latency vectors that carry forward provenance information (source modality, generation time, transmission delays). This layer establishes the foundational data substrate for all subsequent accountability computations without altering original clinical workflows.

Layer 2 – Temporal drift analysis and accountability mapping

A continuously operating drift engine evaluates each latency vector against a library of context-conditioned thresholds. Thresholds are stratified by clinical urgency tier (e.g., life-threatening, serious but stable, routine), modality velocity class, and institutional governance policies [8, 10]. Deviations trigger immediate accountability mapping: responsibility is attributed to the most proximal accountable entity—human clinician, automated agent, interface component, or infrastructure layer—using metadata trails preserved from layer 1. Mapping occurs at the earliest detectable point to minimize diffusion of responsibility.

Layer 3 – Safety variable scoring and risk propagation

Latency is elevated to an explicit, dynamic safety variable through quantitative scoring. Risk propagation across sequential or parallel care-pathway segments is formalized as:

$$R_{p(t)} = \beta \sum w_i \cdot \Delta L_i(t) \cdot S_i \cdot \prod (1-\gamma_j)$$

where $R_p(t)$ is cumulative propagated risk at time t, β \beta β is a tunable domain scaling factor, w_i modality-specific weights, $\Delta L_i(t)$ incremental latency on segment i, S_i clinical-severity multiplier, and γ_j mitigation-efficacy terms from prior interventions (allowing dampening of propagation when early mitigations succeed) [13, 29].

Decision-confidence decay under accumulating latency follows an exponential model motivated by psychophysical and physiologic time-sensitivity literature:

$$C(t) = C0 \exp(-\alpha \cdot Lcum(t)) \cdot f(context) \quad (1)$$

with C(t) residual confidence, C_0 baseline, α decay rate (higher for time-critical states), $L_{cum}(t)$ cumulative latency, and f(context) a context-sensitivity adjustment [14].

Layer 4 – Governance response and feedback orchestration

When aggregated risk $R_p(t)$ or confidence C(t) breaches predefined escalation thresholds, layer 4 instantiates tiered responses: passive logging → visible escalation alerts → forced handoff → resource reallocation prompts → executive dashboard flagging for oversight committees. The governance burden imposed by continuous monitoring is modeled as:

$$M_b = \int \rho(\tau) \cdot (1 + \kappa \cdot dR_p d\tau)\, d\tau$$

where ρ(τ) is baseline resource demand per unit time, and κ amplifies burden during periods of rapid risk escalation. **Figure 1** illustrates the CLOL as a bidirectional temporal-accountability architecture that converts multimodal delay signals into attributable risk propagation, confidence decay, and tiered mitigation across sequential decision nodes.

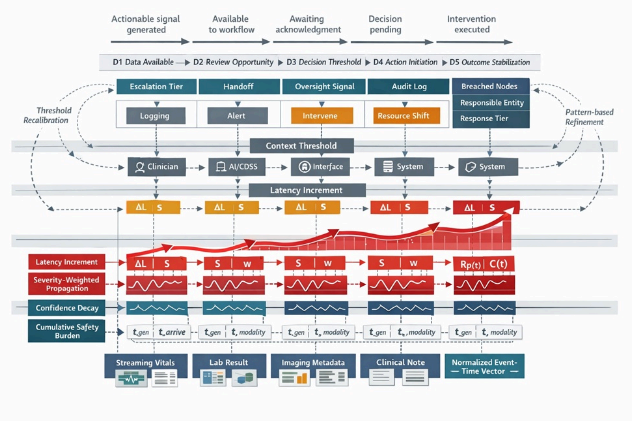

Figure 1. Clinical latency oversight lattice (CLOL) for temporal accountability in care pathways**.** The figure depicts the manuscript's proposed four-layer bidirectional architecture for governing clinical decision latency as an explicit safety variable. Multimodal time-stamped clinical inputs are first normalized into latency-aware event vectors, then evaluated against context-conditioned temporal thresholds to identify harmful drift and assign accountability to the most proximal responsible actor or system component. Detected delays are transformed into dynamic safety scores through propagated risk accumulation and decision-confidence decay across sequential care-pathway nodes. When threshold breaches occur, tiered governance responses are triggered, ranging from passive logging to escalation, handoff, resource reallocation, and oversight review. Dashed backward feedback arcs indicate prospective threshold refinement and retrospective

learning, distinguishing CLOL from conventional feed-forward decision-support designs. **Table 2** analytically differentiates the four CLOL layers by specifying the distinct inputs, operations, outputs, accountability roles, and failure-containment functions that together enable prospective temporal harm governance.

Table 2. Conceptual differentiation of the four CLOL layers: inputs, core operations, outputs, accountability role, and failure containment function

CLOL layer	**Primary analytical input**	**Core operation**	**Immediate output**
Layer 1: Latency sensing and multimodal ingestion	Source-level timestamps, modality provenance, transmission metadata, event-generation times	Normalize heterogeneous clinical events into comparable latency-aware vectors	Structured temporal event strea with preserved provenanc
Layer 2: Temporal drift analysis and accountability mapping	Normalized latency vectors plus urgency-, modality-, and policy-conditioned threshold library	Detect deviation from acceptable temporal envelopes and assign responsibility to the most proximal accountable actor or component	Drift flags, breach states, anc the responsibilit map
Layer 3: Safety variable scoring and risk propagation	Drift events, latency increments, severity multipliers, modality weights, mitigation terms, and cumulative elapsed time	Compute propagated risk and confidence erosion across sequential care nodes	Dynamic ris trajectory and confidence decay profile, pathway-level temporal burden signal
Layer 4: Governance response and feedback orchestration	Risk thresholds, confidence thresholds, responsibility map, audit context, and resource burden indicators	Trigger tiered mitigation and feed lessons backward for recalibration and policy refinement	Escalation action, handoff, resource prompt, oversight signal, anc learning loo

The bidirectional topology fundamentally differentiates CLOL from feed-forward designs: forward flows enable immediate harm interruption, while backward flows support continuous model refinement, cross-institutional pattern sharing, and governance policy evolution—all without requiring empirical training or performance benchmarking at this conceptual stage [1, 6, 9].

System impact dynamics of temporal accountability integration in healthcare ecosystems [1, 6, 17]

The CLOL introduces a fundamental reconfiguration of how temporal dimensions influence safety, workflow, resource utilization, and institutional governance within healthcare delivery systems [3, 6, 9]. By designating clinical decision latency as an explicit safety variable and embedding it within a layered, bidirectional accountability architecture, CLOL generates cascading effects across multiple subsystems—clinical, operational, technological, economic, and regulatory—without presupposing specific implementation outcomes or quantitative performance gains [1, 17, 20]. This section explores these dynamics in depth, examining how the lattice's layers and feedback topology interact with existing healthcare ecosystems to foster temporal harm prevention while addressing potential frictions in adoption and scalability.

Clinical workflow realignment and human-AI symbiosis enhancement

Integration of CLOL shifts clinician attention from reactive error recovery toward proactive temporal risk management [1, 6, 11]. In high-velocity environments such as emergency departments or intensive care units, where data streams arrive asynchronously from physiologic monitors, laboratory interfaces, and imaging systems, the lattice's layer 1 ingestion and layer 2 drift analysis provide continuous visibility into accumulating latency vectors [25, 27]. This visibility enables clinicians to prioritize actions based on propagated risk scores rather than solely on clinical urgency cues or interruptive alerts, potentially transforming disjointed workflows into temporally orchestrated sequences [2, 7, 10]. For example, in surgical decision-making contexts, where intraoperative delays can compound silently, CLOL's real-time drift mapping ensures that accountability is assigned at the point of deviation, encouraging preemptive adjustments before critical thresholds are breached [1, 11]. The bidirectional feedback topology further supports symbiosis between human judgment and infrastructural monitoring: forward propagation interrupts imminent harm through escalation protocols, while backward loops allow clinicians to annotate latency events, refining threshold models over time without requiring external recalibration [3, 6, 8]. Such closed-loop interaction mitigates automation bias by keeping temporal accountability transparent and contestable, ensuring that AI-augmented pathways remain human-centered rather than algorithmically dominant [9, 25]. In diagnostic pathways characterized by uncertainty—such as those in emergency or pediatric settings—the exponential confidence-decay model formalized in layer 3 quantifies how prolonged latency erodes interpretive reliability, prompting earlier escalation or parallel testing strategies and thereby reducing the propagation of diagnostic errors into treatment phases [4, 14, 16]. This realignment not only enhances individual clinician efficacy but also promotes team-level coordination, as shared latency dashboards could facilitate handover communications, minimizing information loss during shift changes or multidisciplinary rounds [18, 20].

Resource allocation optimization and monitoring burden distribution

CLOL's governance-load formula $M_b = \int \rho\left(\tau\right) \cdot \left(1 + \kappa \cdot dR_p d\tau\right) d\tau$ conceptualizes monitoring as a dynamic cost that escalates nonlinearly during periods of rapid risk accumulation, providing a theoretical lens for balancing surveillance intensity against resource constraints [17, 20]. This formulation highlights a key tension: continuous latency surveillance imposes baseline resource demand (ρ(τ)), amplified by risk velocity ($\kappa \cdot dR_p d\tau$), which could strain computational infrastructure, clinician attention, or administrative overhead in resource-limited settings [2, 7, 17]. However, by concentrating oversight on high-severity, high-latency pathways—such as acute kidney injury prediction or antibiotic dosing in critical care—the lattice enables stratified resource deployment [10, 24, 29]. This stratification has the potential to redistribute institutional monitoring burden away from blanket alert fatigue toward targeted temporal governance, freeing cognitive and computational resources for value-adding activities such as complex case discussion, patient education, or preventive care planning [2, 7, 22]. Operationally, the architecture supports predictive resource forecasting: aggregated $R_p\left(t\right)$ trajectories across patient cohorts can inform staffing models, operating-room scheduling, diagnostic-capacity planning, and even supply-chain adjustments for time-sensitive interventions, transforming latency from an unpredictable drain into a governable planning variable [13, 20, 23]. In stewardship programs for infectious diseases or diagnostic testing, where enforced timing rules have demonstrated harm reduction, CLOL's feedback loops could automate burden-minimizing adaptations, such as dynamic threshold loosening during low-risk periods to conserve monitoring capacity [22, 23]. Ultimately, this optimization dynamic positions temporal accountability as a lever for sustainable resource use, aligning with broader healthcare imperatives to maximize efficiency amid growing data volumes and AI integration [1, 9, 25].

Economic and reimbursement implications in value-based care models

In environments transitioning toward value-based reimbursement, where penalties for readmissions, prolonged lengths of stay, and preventable complications are increasingly tied to temporal performance metrics, CLOL offers a theoretical scaffold for linking latency governance to financial accountability [15, 20]. Propagated risk scores serve as proxies for downstream harm probability, enabling institutions to demonstrate proactive mitigation of costly delays (e.g., in sepsis bundles, acute myocardial infarction reperfusion, or postoperative deterioration) through auditable temporal logs [10, 13, 16]. Backward feedback loops facilitate institutional learning that refines latency thresholds based on observed harm patterns, supporting continuous quality improvement programs eligible for shared-savings arrangements under models like accountable care organizations [15, 20]. From a payer perspective, standardized latency metadata propagated through the lattice could underpin new reimbursement modifiers—rewarding systems that maintain low cumulative $L_{cum}\left(t\right)$ in high-risk cohorts or penalizing persistent drift beyond governance-defined envelopes, thereby incentivizing infrastructural upgrades [15]. This alignment creates economic incentives for investment in temporal accountability, shifting cost structures from reactive harm remediation (e.g., litigation, extended hospitalizations) toward preventive temporal optimization [16, 18, 21]. In specialized domains like oncology or rare genetic diseases, where diagnostic delays incur substantial economic burdens through missed therapeutic windows, CLOL's modality-specific scoring could justify tiered reimbursement for accelerated pathways, fostering equity in resource distribution across patient populations [21, 26]. Moreover, by quantifying governance load, the framework aids in cost-benefit analyses for AI adoption, highlighting scenarios where latency reduction offsets implementation expenses [1, 9, 24].

Regulatory and ethical accountability evolution

CLOL addresses a core deficit in existing governance frameworks: the absence of mechanisms to trace and assign responsibility for temporal drift in AI-augmented pathways [3, 6, 9]. By mapping deviations to specific nodes (human, algorithmic, infrastructural) at the earliest detectable stage, the lattice establishes auditable temporal accountability chains that support forensic review following adverse events, aligning with patient-safety reporting requirements [14, 17, 18]. This traceability aligns with emerging ethical imperatives for AI in healthcare, where explainability must extend beyond diagnostic reasoning to include timeliness of reasoning, as delays can exacerbate inequities in access or outcomes [6, 9, 21]. Regulatory bodies could leverage CLOL-derived metadata to define latency-related performance boundaries for non-device clinical decision support software, particularly in time-sensitive contexts where independent clinician review is constrained by workflow pressures [1, 3, 8]. The bidirectional topology further enables prospective governance evolution: learned patterns from institutional deployments inform threshold harmonization across organizations, fostering standardized temporal safety benchmarks without mandating centralized control or compromising institutional autonomy [6, 20]. Ethically, explicit scoring of decision-confidence decay under latency reinforces the principle of non-maleficence by quantifying how delay undermines beneficent action, providing a conceptual basis for justifying escalation protocols even when immediate clinical indicators remain equivocal [14, 16]. In pandemic-era diagnostics or telemedicine scenarios, where temporal errors amplified harm, CLOL's architecture could integrate with regulatory audits to ensure equitable temporal governance across diverse deployment environments [19, 28].

Cross-institutional pattern propagation and scalability considerations

The lattice's backward feedback arcs enable propagation of temporal patterns beyond single institutions—aggregated anonymized latency-event data could inform cross-system threshold models, accelerating identification of systemic vulnerabilities (e.g., interoperability-induced delays in regional health information exchanges or supply-chain latencies in global pandemics) [17, 19, 27]. This scalability distinguishes CLOL from localized quality initiatives: forward prevention remains institution-specific, while backward learning supports ecosystem-wide resilience against shared temporal risks [6, 9, 20]. However, scalability introduces governance challenges—data-sharing protocols must preserve patient privacy while enabling pattern mining, and institutional autonomy must be balanced against collective safety gains, potentially requiring federated learning paradigms [25, 27, 28]. The architecture accommodates such approaches, where local lattices exchange threshold refinements without centralizing raw latency vectors, preserving data sovereignty while advancing temporal accountability at scale [27]. In multi-site research consortia or national health systems, this could facilitate benchmarking of latency norms across demographics, revealing disparities (e.g., higher delays in underserved areas) and guiding policy interventions [21, 28]. Overall, these dynamics position CLOL as a catalyst for scalable, interconnected temporal governance, evolving healthcare ecosystems toward greater temporal resilience [1, 3, 17].

Results and Discussion

The repositioning of clinical decision latency as a primary safety variable through the CLOL framework confronts a persistent structural blind spot in healthcare systems and AI augmentation strategies [1, 3, 6]. Traditional patient-safety paradigms—rooted in taxonomies of error commission, omission, or process deviation—treat time as an implicit background variable, rarely subjecting it to the same rigorous governance applied to diagnostic accuracy or medication correctness [14, 16, 17]. Yet the synthesized evidence underscores that latency operates as a propagative multiplier: modest delays in data acknowledgment or action initiation cascade into disproportionate harm, particularly in time-sensitive domains where physiologic windows close irreversibly, as seen in hemorrhage management, acute kidney injury, or intraoperative crises [10, 13, 29].

By formalizing latency within an explicit architectural lattice, CLOL bridges this conceptual gap, providing infrastructure that makes temporal risk visible, measurable, and accountable across modalities and actors [3, 6, 9]. The four-layer design—ingestion for multimodal normalization, drift analysis for deviation mapping, scoring/propagation for risk quantification, and orchestrated response for mitigation—combined with bidirectional feedback, creates a closed temporal governance loop absent in linear CDSS pipelines [2, 7, 8]. The interpretive formulas advance theoretical understanding: risk propagation captures inter-pathway amplification with modality weights and severity multipliers, confidence decay models psychophysiologic erosion of judgment reliability under cumulative latency, and governance-load quantification highlights the resource trade-offs inherent in continuous surveillance, offering tools for theoretical simulation of system behaviors [13, 14, 29].

Several limitations must be acknowledged at the conceptual level. The framework assumes availability of high-fidelity timestamps and metadata provenance across heterogeneous systems—an assumption challenged by legacy infrastructures and interoperability failures documented in health IT adverse-event analyses [17]. Threshold calibration remains context-dependent and institution-specific, requiring iterative refinement that the backward loops support but do not automate, potentially leading to initial inconsistencies in multi-site deployments [8, 20]. Ethical tensions arise in balancing proactive escalation (potentially increasing interruptive burden and alert fatigue) against watchful waiting, particularly when confidence decay models predict erosion without overt clinical deterioration, as in diagnostic uncertainty scenarios [2, 4, 14]. Economic modeling of latency governance requires further elaboration to quantify return on infrastructural investment in diverse reimbursement environments, where value-based care incentives may vary by jurisdiction [15]. Additionally, while CLOL mitigates AI-amplified latency drift, it does not inherently address upstream data-quality issues (e.g., sensor inaccuracies or EHR entry errors) that could confound latency vectors [12, 25, 27].

Despite these boundaries, CLOL offers a coherent path forward for healthcare analytics and AI system architects [1, 6, 9]. It reframes temporal performance from an operational nuisance to a core safety domain, aligning with broader movements toward proactive, systems-level harm prevention in critical care, radiology, and stewardship programs [3, 11, 22]. Future conceptual extensions could incorporate multi-agent coordination (e.g., across interprofessional care teams or federated institutions), integration with emerging real-time predictive models for anticipatory latency adjustment, or expansion to non-acute settings like chronic disease management, always preserving the foundational commitment to explicit accountability [24, 25, 28]. By synthesizing gaps in diagnostic error, treatment delay, and IT vulnerabilities, the framework underscores the need for temporal-centric redesigns in AI-driven healthcare [14, 16, 17].

Conclusion

Clinical decision latency, long treated as an invisible operational artifact, warrants recognition as a primary, propagative safety variable whose governance is essential for safe, equitable, and effective modern healthcare. The CLOL provides a novel architectural response: a four-layer, bidirectional lattice that senses, analyzes, scores, and orchestrates temporal accountability across modalities, actors, and institutions.

Through modality-stratified ingestion normalizing heterogeneous data velocities, explicit drift mapping assigning responsibility at deviation points, dynamic risk/confidence propagation quantifying harm potential, and tiered governance responses enabling mitigation, CLOL institutionalizes proactive temporal harm interruption while supporting continuous system learning via feedback loops. The interpretive formalisms—risk propagation with weighted summation for pathway interdependencies, confidence decay exponentially modeling judgment erosion, and governance-load integration capturing monitoring costs—supply theoretical rigor to an area previously lacking structured conceptualization, facilitating deeper analysis of temporal dynamics in diverse clinical contexts.

By elevating latency to an explicit safety variable embedded in infrastructural accountability, this framework advances AI-augmented healthcare from reactive error mitigation toward preventive delay governance, addressing vulnerabilities in CDSS, diagnostic pathways, and health IT ecosystems. Implementation pathways remain open for empirical validation in settings like critical care or oncology, but the conceptual foundation establishes temporal accountability as indispensable for next-generation clinical systems. Ultimately, recognizing and governing decision latency as a core safety construct promises to narrow preventable-harm windows, optimize resource deployment across economic models, reinforce ethical traceability in AI integration, and enhance scalability through pattern propagation—fostering resilient healthcare ecosystems capable of minimizing temporal-induced harm.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 22 Dec 2021 Revised: 29 Jan 2022 Accepted: 29 Mar 2022

Rights and permissions

References

Loftus TJ, Tighe PJ, Filiberto AC, Efron PA, Brakenridge SC, Mohr AM, et al. Artificial intelligence and surgical decision-making. JAMA Surg. 2020;155(2):148-58. https://doi.org/10.1001/jamasurg.2019.4917.

Olakotan O, Mohd Yusof M, Ezat Wan Puteh S. A systematic review on CDSS alert appropriateness. Stud Health Technol Inform. 2020;270:906-10. https://doi.org/10.3233/SHTI200293.

Belard A, Buchman T, Forsberg J, Potter BK, Dente CJ, Kirk A, et al. Precision diagnosis: a view of the clinical decision support systems (CDSS) landscape through the lens of critical care. J Clin Monit Comput. 2017;31(2):261-71. https://doi.org/10.1007/s10877-016-9849-1.

Medford-Davis LN, Singh H, Mahajan P. Diagnostic decision-making in the emergency department. Pediatr Clin North Am. 2018;65(6):1097-105. https://doi.org/10.1016/j.pcl.2018.07.003.

Piva E, Sciacovelli L, Pelloso M, Plebani M. Performance specifications of critical results management. Clin Biochem. 2017;50(10-11):617-21. https://doi.org/10.1016/j.clinbiochem.2017.05.010.

El-Kareh R, Sittig DF. Enhancing diagnosis through technology: decision support, artificial intelligence, and beyond. Crit Care Clin. 2022;38(1):129-39. https://doi.org/10.1016/j.ccc.2021.08.004.

Bittmann JA, Haefeli WE, Seidling HM. Modulators influencing medication alert acceptance: an explorative review. Appl Clin Inform. 2022;13(2):468-85. https://doi.org/10.1055/s-0042-1748146.

Dziorny AC, Heneghan JA, Bhat MA, Karavite DJ, Sanchez-Pinto LN, McArthur J, et al. Clinical decision support in the PICU: implications for design and evaluation. Pediatr Crit Care Med. 2022;23(8):e392-6. https://doi.org/10.1097/PCC.0000000000002973.

Jenders RA. Advances in clinical decision support: highlights of practice and the literature 2015-2016. Yearb Med Inform.

2017;26(1):125-32.
https://doi.org/10.15265/IY-2017-012.

Wilson FP, Greenberg JH. Acute kidney injury in real time: prediction, alerts, and clinical decision support. Nephron. 2018;140(2):116-9.
https://doi.org/10.1159/000492064.

Nair BG, Gabel E, Hofer I, Schwid HA, Cannesson M. Intraoperative clinical decision support for anesthesia: a narrative review of available systems. Anesth Analg. 2017;124(2):603-17.
https://doi.org/10.1213/ANE.0000000000001636.

Waite S, Scott JM, Legasto A, Kolla S, Gale B, Krupinski EA. Systemic error in radiology. AJR Am J Roentgenol. 2017;209(3):629-39.
https://doi.org/10.2214/AJR.16.17719.

Gayet-Ageron A, Prieto-Merino D, Ker K, Shakur H, Ageron FX, Roberts I, et al. Effect of treatment delay on the effectiveness and safety of antifibrinolytics in acute severe haemorrhage: a meta-analysis of individual patient-level data from 40 138 bleeding patients. Lancet. 2018;391(10116):125-32.
https://doi.org/10.1016/S0140-6736(17)32455-8.

Bhise V, Rajan SS, Sittig DF, Morgan RO, Chaudhary P, Singh H. Defining and measuring diagnostic uncertainty in medicine: a systematic review. J Gen Intern Med. 2018;33(1):103-15.
https://doi.org/10.1007/s11606-017-4164-1.

Berenson R, Singh H. Payment innovations to improve diagnostic accuracy and reduce diagnostic error. Health Aff (Millwood). 2018;37(11):1828-35.
https://doi.org/10.1377/hlthaff.2018.0714.

Rinke ML, Singh H, Heo M, Adelman JS, O’Donnell HC, Choi SJ, et al. Diagnostic errors in primary care pediatrics: Project RedDE. Acad Pediatr. 2018;18(2):220-7.
https://doi.org/10.1016/j.acap.2017.08.005.

Kim MO, Coiera E, Magrabi F. Problems with health information technology and their effects on care delivery and patient outcomes: a systematic review. J Am Med Inform Assoc. 2017;24(2):246-50.

Nehls N, Yap TS, Salant T, Aronson M, Schiff G, Olbricht S, et al. Systems engineering analysis of diagnostic referral closed-loop processes. BMJ Open Qual. 2021;10(4):e001603.
https://doi.org/10.1136/bmjoq-2021-001603.

Shen L, Levie A, Singh H, Murray K, Desai S. Harnessing event report data to identify diagnostic error during the COVID-19 pandemic. Jt Comm J Qual Patient Saf. 2022;48(2):71-80.
https://doi.org/10.1016/j.jcjq.2021.10.002.

Campione JR, Mardon RE, McDonald KM. Patient safety culture, health information technology implementation, and medical office problems that could lead to diagnostic error. J Patient Saf. 2019;15(4):267-73.
https://doi.org/10.1097/PTS.0000000000000531.

Merker VL, Slobogean B, Jordan JT, Langmead S, Meterko M, Charns MP, et al. Understanding barriers to diagnosis in a rare, genetic disease: delays and errors in diagnosing schwannomatosis. Am J Med Genet A. 2022;188(9):2672-83.
https://doi.org/10.1002/ajmg.a.62860.

Madden GR, Enfield KB, Sifri CD. Patient outcomes with prevented vs negative Clostridioides difficile tests using a computerized clinical decision support tool. Open Forum Infect Dis. 2020;7(4):ofaa094.

Madden GR, Sifri CD. Reduced Clostridioides difficile tests among solid organ transplant recipients through a diagnostic stewardship bundled intervention. Ann Transplant. 2019;24:304-11.
https://doi.org/10.12659/AOT.915168.

Roggeveen LF, Guo T, Driessen RH, Fleuren LM, Thoral P, van der Voort PHJ, et al. Right dose, right now: development of AutoKinetics for real time model informed precision antibiotic dosing decision support at the bedside of critically ill patients. Front Pharmacol. 2020;11:646.
https://doi.org/10.3389/fphar.2020.00646.

Kiyasseh D, Zhu T, Clifton D. A clinical deep learning framework for continually learning from cardiac signals across diseases, time, modalities, and institutions. Nat Commun. 2021;12(1):4221.
https://doi.org/10.1038/s41467-021-24483-0.

Mehta P, Antonelli M, Ahmed HU, Emberton M, Punwani S, Ourselin S. Computer-aided diagnosis of prostate cancer using multiparametric MRI and clinical features: a patient-level classification framework. Med Image Anal. 2021;73:102153.
https://doi.org/10.1016/j.media.2021.102153.

Najafabadipour M, Zanin M, Rodríguez-González A, Torrente M, Nuñez García B, Cruz Bermudez JL, et al. Reconstructing the patient’s natural history from electronic health records. Artif Intell Med. 2020;105:101860.
https://doi.org/10.1016/j.artmed.2020.101860.

Zikos D, DeLellis N. CDSS-RM: a clinical decision support system reference model. BMC Med Res Methodol. 2018;18(1):137.
https://doi.org/10.1186/s12874-018-0587-6.