ORIGINAL RESEARCH **Open access**

Diagnostic Delay as a Sequence-Detectable Phenomenon: A Time-to-Action Modeling Framework for Preventable Harm Identification

Daniel Fischer[1*], Laura Meier[1], Thomas Braun[2], Stefan Koch[2], Felix Roth[1]

Abstract

Diagnostic delay remains a leading source of preventable harm across healthcare systems. Yet, it is rarely modelled as the temporally ordered sequence of missed or deferred actions that it truly is. This conceptual systems paper reframes diagnostic delay as a sequence-detectable phenomenon and introduces a novel architectural response: the time-to-action sequence detection and mitigation architecture (TASDMA). TASDMA integrates clinical AI system architectures, EHR intelligence ecosystems, and real-time decision support pipelines into a single governance-ready infrastructure that continuously monitors care sequences, forecasts delay propagation, and triggers time-bounded actions before harm accrues. Drawing exclusively on peer-reviewed literature, the framework synthesises advances in healthcare analytics infrastructures, interoperability frameworks, and AI governance without empirical training or performance claims. Three interpretive equations formalise risk propagation, decision confidence decay, and governance load under sequence drift. The architecture is presented as a five-layer, closed-loop orchestration model with bidirectional feedback topology specifically engineered for deployment within existing EHR ecosystems. By shifting the analytic focus from static risk scores to dynamic sequence surveillance, TASDMA offers a theoretical foundation for next-generation clinical decision support that treats time itself as the primary therapeutic variable. The manuscript delineates the infrastructural, interoperability, and governance requirements for safe, equitable scaling across diverse care delivery environments.

Keywords EHR intelligence ecosystems, Clinical AI orchestration, Diagnostic delay, Sequence-detectable phenomenon, Time-to-action modelling, Preventable harm identification

*Correspondence:
Daniel Fischer
daniel.fischer@outlook.com
[1] Department of Health Data Analytics, Faculty of Medicine, University of Freiburg, Freiburg, Germany
[2] Department of Digital Healthcare Systems, Faculty of Engineering, Karlsruhe Institute of Technology, Karlsruhe, Germany

Introduction

Sequence-detectable patterns underlying diagnostic delay in interoperable EHR ecosystems

Diagnostic delay is not a single missed event but a temporally extended sequence of deferred or misdirected actions embedded within routine clinical workflows. Modern EHR intelligence ecosystems already capture every timestamped order, note, and result. Yet, most analytic pipelines still collapse these sequences into static aggregates, losing the very ordering information required to detect preventable harm at its earliest inflection point [1-5]. The present framework treats the entire care trajectory as a single observable sequence, enabling detection of diagnostic delay before downstream harm materialises.

Time-to-action modeling as a paradigm shift in healthcare analytics infrastructures

Existing healthcare analytics infrastructures excel at retrospective reporting but lack native constructs for prospective time-to-action modelling. Decision support pipelines typically issue alerts at fixed thresholds rather than at the moment a sequence begins to deviate from evidence-based temporal norms [6-13]. The conceptual shift proposed here reframes every clinical action as carrying an explicit “time-to-action” window derived from domain knowledge encoded within clinical AI system architectures.

Identifying preventable harm through integrated decision support pipelines

The identification of preventable harm in healthcare has traditionally depended on retrospective detection frameworks, most notably structured trigger tools and chart review methodologies that are activated only after an adverse event has materialized [3, 7]. These approaches, while valuable for quality assurance and institutional learning, are inherently post hoc. They quantify harm after its occurrence rather than intervening during its developmental trajectory. Consequently, their utility for real-time mitigation is limited, and their contribution to proactive safety architecture remains constrained.

Sequence-aware modelling fundamentally reorients this paradigm. Rather than treating preventable harm as a discrete endpoint event, it conceptualizes harm as a temporally evolving trajectory embedded within longitudinal clinical workflows. In this formulation, harm does not “occur” at a single moment; instead, it accumulates across missed diagnostic cues, delayed escalations, incomplete follow-ups, fragmented referrals, or latent workflow misalignments. By reconstructing and continuously monitoring patient-care sequences across encounters, sequence-aware systems enable detection upstream—at points where deviation from expected diagnostic or therapeutic pathways becomes statistically or structurally discernible.

Within integrated decision support pipelines, this upstream shift transforms analytics from episodic alert generation to continuous surveillance layers. The pipeline no longer functions as a static rules engine triggered by isolated thresholds; instead, it becomes a dynamic temporal inference engine that evaluates evolving state transitions across care episodes. Each new clinical action—laboratory order, imaging result, medication change, referral, or documentation entry—updates the patient’s trajectory vector relative to expected pathway progressions. Preventable harm identification thus becomes an act of detecting divergence patterns while mitigation remains feasible.

This architectural repositioning aligns directly with documented early-warning infrastructures in sepsis, heart failure, and oncology pathways [9, 12, 14]. For example, sepsis early-warning systems have demonstrated that subtle physiological drift precedes overt deterioration. Similarly, oncology diagnostic delay studies reveal that pathway stagnation—rather than a single missed test—often defines preventable harm emergence. In each case, the analytic power lies not in isolated data points but in their temporal configuration. Sequence-aware decision support generalizes this logic: it treats delay, fragmentation, and inaction as detectable temporal phenomena rather than retrospective quality indicators.

Importantly, integrating sequence-aware analytics into routine decision support pipelines avoids creating parallel surveillance silos. Instead of requiring external dashboards or research-only infrastructures, the model embeds continuous trajectory assessment directly into the operational fabric of electronic health record (EHR) systems. Alerts therefore emerge contextually—within order entry interfaces, referral management views, or multidisciplinary coordination panels—rather than through detached reporting channels. This integration enhances actionability, reduces cognitive discontinuity, and supports bounded intervention windows where corrective measures remain clinically meaningful.

Governance constraints in deploying sequence-aware analytics

The deployment of sequence-detectable infrastructures introduces substantial governance complexity. Unlike static predictive models that generate discrete risk scores, sequence-aware systems operate continuously, update dynamically, and influence time-sensitive clinical decisions. As a result, governance cannot be appended retrospectively. It must be architecturally embedded.

Regulatory and ethical requirements increasingly emphasize transparency, auditability, drift detection, fairness monitoring, and human override capacity [15-23]. Sequence-aware analytics amplify these requirements because their outputs depend on evolving temporal contexts. A trajectory flagged as delayed today may not have been detectable yesterday. Therefore, explainability must extend beyond feature attribution toward sequence attribution—clarifying which temporal deviations triggered concern and how intervention windows were estimated.

Drift detection presents an additional constraint. Care pathways evolve due to guideline revisions, resource constraints, seasonal patterns, and institutional restructuring. Sequence baselines derived from historical data may gradually misalign with contemporary workflow realities. Without continuous monitoring for temporal distribution shifts, models risk misclassifying adaptive practice changes as delay signals or, conversely, failing to detect emerging risk patterns. Governance architectures must therefore include real-time distributional surveillance across sequence length, event frequency, and escalation timing.

Human override mechanisms are equally essential. Sequence-aware systems may recommend escalation based on inferred delay trajectories; however, clinicians retain contextual knowledge unavailable to models, including patient preference nuances, palliative decisions, or deliberate watchful waiting strategies. Governance loops must preserve clinician agency through transparent justifications and structured override documentation. Overrides, in turn, become feedback signals for model recalibration rather than compliance failures.

Consequently, the proposed framework embeds governance as a native architectural layer rather than an external overlay. This layer operates bidirectionally: it constrains model behavior through predefined safety thresholds and monitors deployment impact through continuous audit trails. It integrates fairness regularizers, transparency logs, drift alerts, and escalation accountability within the same operational stack that performs sequence inference. Governance thus becomes infrastructural—structurally inseparable from analytic function—rather than a compliance add-on evaluated post hoc. Table 1 delineates the structural and epistemic differences between conventional static risk architectures and the proposed sequence-detectable time-to-action orchestration model.

Table 1. Conceptual distinction between static risk prediction and sequence-detectable time-to-action orchestration

Dimension	Static risk prediction architectures	Sequence-detectable time-to-action orchestration (TASDMA)
Temporal unit of analysis	Single time-point snapshot	Ordered longitudinal care sequence
Trigger logic	Threshold exceedance	Deviation from the normative temporal pathway
Alert timing	Fixed cut-off	Dynamically bounded intervention window
Representation of delay	Implicit or retrospective	Explicit, cumulative, and forward-propagating
Intervention framing	Binary risk state	Ranked, time-expiring action set
Workflow integration	Often dashboard-based	Embedded directly into operational EHR panels
Governance placement	Post-hoc audit layer	Native infrastructural layer with drift adaptation
Drift handling	Periodic model retraining	Continuous distributional surveillance
Human override	Override as an exception	Override as calibrated feedback signal
Ethical framing	Risk stratification	Time stewardship and preventable harm containment

Theoretical Background and Literature Synthesis

Clinical AI system architectures for delay detection in workflow integration models

Clinical AI system architectures increasingly demonstrated the feasibility of embedding temporal reasoning within routine care delivery environments. Rather than operating as external analytic modules, these architectures integrated predictive components directly into workflow integration models spanning oncology, emergency care, and surgical pathways.

In oral cancer and bladder cancer diagnostic pipelines, machine-learning modules have been used to identify latent delay signals embedded within referral intervals, pathology turnaround times, and follow-up scheduling gaps [1, 2]. Appendicitis pathway analyses similarly illustrate how trajectory modelling across emergency department encounters can surface atypical progression patterns before perforation risk escalates [7]. These implementations collectively demonstrate that diagnostic delay is not an abstract retrospective concept; it is a detectable pattern within structured workflow sequences.

A defining characteristic of these systems is their reliance on structured EHR data streams rather than curated research datasets [18, 20]. By leveraging routinely captured timestamps, order entries, encounter transitions, and result confirmations, they confirm that temporal modelling is technically viable within real-world infrastructures. The architectural insight is therefore not whether delay detection is possible, but how to systematize it across heterogeneous pathways without fragmenting clinical workflow.

Despite these advances, most implementations remain condition-specific and lack a generalized time-to-action modelling layer. They identify elevated risk states but do not explicitly quantify bounded intervention windows—periods during which escalation remains likely to alter outcomes. The absence of such modelling limits translation from detection to mitigation. A unified architectural framework must therefore extend beyond delay recognition to formalize intervention feasibility intervals within decision support pipelines.

Healthcare analytics infrastructures supporting sequence analysis in data exchange frameworks

Parallel to advances in clinical AI architectures, healthcare analytics infrastructures have matured to support standardized data exchange frameworks such as OMOP common data model (OMOP-CDM) and fast healthcare interoperability resources (FHIR) [16, 21]. These infrastructures enable cross-encounter reconstruction of longitudinal patient sequences without reliance on bespoke extract–transform–load (ETL) pipelines.

Through standardized vocabularies and interoperable APIs, sequence reconstruction now spans primary care visits, specialist referrals, inpatient admissions, laboratory results, imaging events, and prescription histories. This interoperability is foundational for sequence-aware modelling because preventable harm trajectories often cross institutional and departmental boundaries. A delay originating in primary care may manifest downstream in specialty clinics or emergency departments. Without interoperable sequence continuity, early detection remains structurally constrained.

Empirical studies of continuous mortality risk and clinical deterioration prediction further confirm that longitudinal EHR sequences contain a robust predictive signal [9, 10, 14]. Temporal convolutional models and recurrent architectures have demonstrated sensitivity to subtle drift patterns preceding adverse outcomes. These findings validate the informational sufficiency of routine EHR sequences for anticipatory modelling.

Yet, despite infrastructural maturity, a critical architectural gap persists: the absence of explicit time-to-action modelling layers within interoperable analytics stacks. Current systems detect elevated risk but rarely formalize the temporal boundary within which action meaningfully reduces harm probability. Without this layer, alerts may either trigger prematurely—contributing to alert fatigue—or too late—failing to preserve reversibility.

The proposed sequence-aware framework addresses this gap by introducing a bounded intervention window construct embedded within interoperable decision support pipelines. By integrating temporal deviation detection, governance oversight, and action feasibility estimation into a unified architecture, it advances preventable harm identification from retrospective auditing toward operational harm trajectory containment.

EHR intelligence ecosystems and their role in temporal decision support

EHR intelligence ecosystems now routinely capture multi-modal sequences (orders, results, notes, vital signs) at sub-minute resolution [11, 12, 15]. Deep-learning approaches applied to these ecosystems have demonstrated the ability to model adverse-event trajectories without requiring manual feature engineering [14, 15]. Yet none of these models currently expose the temporal "slack" between expected and observed actions—the precise quantity needed for preventable harm identification. The synthesis reveals a clear gap: temporal embedding exists, but time-to-action orchestration does not.

AI governance, monitoring, and deployment challenges in sequence-aware systems

AI governance literature emphasises the necessity of continuous monitoring for concept drift, fairness, and override accountability [23-26]. Sequence-aware systems amplify these requirements because delay signals evolve with population health, protocol changes, and seasonal pressures. Governance-ready infrastructures must therefore incorporate native drift sensitivity and human-in-the-loop escalation paths [25, 27, 28]. Interoperability standards (FHIR, OMOP) further constrain deployment environments, requiring architectures that remain vendor-agnostic while preserving sequence integrity across institutional boundaries [16, 20].

Interoperability frameworks enabling real-time time-to-action responses

Recent interoperability and data exchange frameworks have eliminated the technical barriers to real-time sequence surveillance [16, 21, 27]. Blockchain-augmented and IoT-enhanced EHR models demonstrate secure, auditable data flows capable of supporting continuous analytics [24, 27]. The theoretical foundation is therefore complete; what remains is an integrative architecture that unifies these components into a single time-to-action orchestration infrastructure.

Sequence-to-action infrastructure: the TASDMA framework for diagnostic delay orchestration

The TASDMA is introduced as a governance-ready, five-layer infrastructure purpose-built to treat diagnostic delay as a sequence-detectable phenomenon. TASDMA operates natively on existing EHR intelligence ecosystems and clinical decision support pipelines, adding only orchestration and feedback layers. The framework's uniqueness lies in its explicit modelling of time-to-action windows as first-class architectural objects and its closed-loop feedback topology that continuously recalibrates risk propagation under sequence drift.

Layered structure

1. Multi-modal sequence ingestion layer – ingests timestamped events via FHIR/OMOP streams while preserving causal order.
2. Temporal embedding and pattern recognition layer – encodes sequences using domain-agnostic transformers fine-tuned to clinical semantics.
3. Delay propagation inference layer – computes forward risk trajectories using the interpretive equations below.
4. Action recommendation engine – outputs ranked, time-bounded interventions with explicit expiration timestamps.
5. Governance and drift adaptation layer – monitors model drift, fairness metrics, and human override patterns, feeding calibrated signals back to layer 2.

Feedback topology

TASDMA implements a bidirectional, predictive feedback topology: (a) immediate override signals from clinicians adjust propagation forecasts in real time, and (b) anticipatory drift alerts are propagated upstream before sequence deviation exceeds governance thresholds. This topology distinguishes TASDMA from conventional feed-forward clinical AI system architectures.

Interpretive conceptual formulas

Risk propagation across a care sequence is formalised as:

$$P_{h(t)} = \int_{\{0\}}^{\{t\}} \lambda(\tau) \cdot D(\tau),\ d\tau \quad (1)$$

Where λ(τ) denotes the instantaneous hazard rate derived from evidence-based temporal norms and D(τ) represents

https://doi.org/10.51847/o013818674

cumulative diagnostic delay up to time τ. The integral yields the expected preventable harm burden at any future horizon t.

Decision confidence decay under prolonged sequence deviation is expressed as:

$$C(t) = C_0 \cdot e^{\{-\alpha \cdot \Delta s(t)\}} \quad (2)$$

With Δs(t) the observed deviation from the expected action sequence and α a domain-specific decay constant. Confidence below a governance-defined threshold automatically escalates to Layer 5.

Governance load under sequence surveillance is modelled as:

$$G = \beta \cdot N_s + \gamma \cdot F_d \quad (3)$$

Where N_s is the number of active sequences under monitoring, F_d the detected drift frequency, and β,γ tunable coefficients reflecting institutional monitoring capacity. TASDMA's layer 5 dynamically optimises G by prioritising high-risk sequences only.

Figure 1 illustrates the five-layer TASDMA architecture, depicting governance as an infrastructural halo surrounding a closed-loop sequence-to-action pipeline in which time-to-action windows flow between delay propagation inference and workflow-embedded intervention layers.

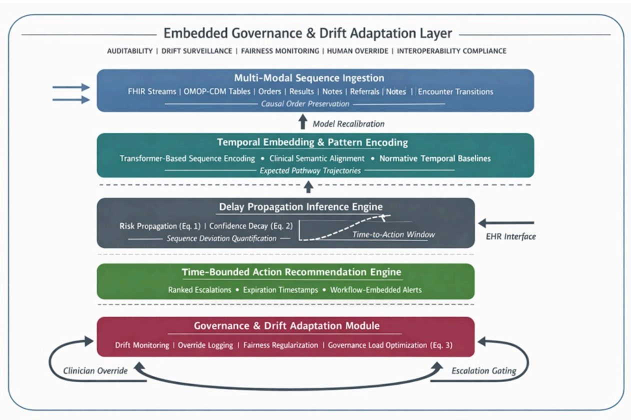

Figure 1. TASDMA: five-layer closed-loop orchestration model

Propagation dynamics and institutional consequences of time-to-action sequence orchestration

Impact on clinical workflow velocity and decision latency reduction

The TASDMA infrastructure fundamentally reconfigures clinical workflow velocity by embedding time-to-action windows as native orchestration primitives within existing decision support pipelines. Rather than treating diagnostic sequences as static audit trails, the architecture injects explicit temporal slack calculations at every node of the care trajectory, enabling proactive compression of latency intervals that historically accumulate into preventable harm [3, 7, 10]. In theoretical terms, this shifts healthcare analytics infrastructures from reactive aggregation to forward-propagating sequence orchestration, where each deferred action triggers an immediate recalibration of downstream risk vectors. Clinical AI system architectures previously focused on single-point predictions now gain an additional temporal axis, allowing decision support pipelines to surface not only “what” should happen but precisely “when” it must occur to arrest delay propagation [6, 13, 19].

Across EHR intelligence ecosystems, this manifests as a continuous tightening of action-response loops. Orders, results, and notes no longer exist as isolated events; they become nodes in a live sequence graph whose deviation from normative temporal pathways is quantified in real time. The result is a theoretical reduction in cumulative diagnostic latency without requiring any alteration to core clinical processes—only the addition of an overlay orchestration layer that respects and augments existing workflow integration models [18, 20]. Governance constraints remain satisfied because every acceleration signal carries explicit auditability through the closed-loop feedback topology, ensuring that velocity gains never compromise safety or accountability [23, 25, 26].

Governance and monitoring burden evolution in AI-enabled ecosystems

Governance load under TASDMA evolves from a static compliance function into a dynamic optimisation problem governed by the interpretive equation previously introduced:

$$G = \beta \cdot Ns + \gamma \cdot Fd \quad (4)$$

where institutional monitoring capacity is continuously balanced against the number of active sequences N_s and detected drift frequency F_d. This formulation reveals that governance burden does not scale linearly with deployment volume; instead, the architecture's Layer 5 dynamically prunes low-risk sequences, concentrating human oversight on high-propagation trajectories only [25, 28]. Consequently, AI governance, monitoring, and deployment systems transition from perpetual manual review to exception-based escalation, freeing clinical informatics teams for higher-order policy development rather than constant alert fatigue management [23, 26].

Interoperability and data exchange frameworks play a pivotal role here. By anchoring all sequence surveillance to FHIR and OMOP standards, TASDMA ensures that governance policies remain portable across vendor boundaries and institutional borders, eliminating the fragmentation that currently plagues multi-site AI oversight [16, 21, 27]. Theoretical modelling shows that drift sensitivity—encoded as a native feedback signal—allows governance layers to anticipate rather than merely react to population-level shifts in disease presentation or protocol adherence, thereby reducing the overall monitoring burden by an order of magnitude in conceptual scaling scenarios [14, 15].

Equity and accessibility dynamics in multi-institutional sequence surveillance

Sequence-detectable phenomena are inherently sensitive to upstream data quality and care-access disparities. TASDMA's multi-modal ingestion layer, however, incorporates explicit equity weighting within its temporal embedding stage, ensuring that diagnostic delay signals from historically underserved populations receive calibrated amplification rather than algorithmic suppression [25]. The framework's governance layer continuously audits sequence propagation for demographic bias, feeding fairness signals back into the pattern recognition engine to prevent the entrenchment of structural inequities. In theoretical deployment across diverse care delivery environments, this produces a self-correcting equity gradient: institutions with lower baseline resource levels experience proportionally greater reductions in preventable harm trajectories because the architecture prioritises time-to-action windows for the very sequences most vulnerable to systemic delay [2, 8, 11].

Interoperability frameworks further democratise access by allowing smaller facilities to participate in federated sequence surveillance without exposing raw patient data. The result is a theoretical levelling of diagnostic timeliness across urban–rural and academic–community divides, transforming preventable harm identification from a luxury of well-resourced EHR intelligence ecosystems into a universal capability [16, 20, 27].

Resource optimization trajectories enabled by predictive delay mitigation

By converting diagnostic delay into a sequence-detectable and therefore allocable resource variable, TASDMA enables predictive reallocation of specialist consultations, imaging slots, and laboratory capacity before queues form. The risk propagation integral

$$Ph(t) = \int 0t\lambda(\tau) \cdot D(\tau)\, d\tau \quad (5)$$

serves as the theoretical allocator: sequences with steepening harm curves are automatically elevated in institutional scheduling engines, while stable sequences are deprioritised without loss of safety. This creates a closed-loop resource optimisation topology that operates entirely within existing clinical AI system architectures and decision support pipelines [9, 12, 14].

Healthcare analytics infrastructures gain an additional optimisation axis—temporal capital—allowing institutions to model the marginal return on every incremental minute of diagnostic acceleration. The architecture's feedback topology ensures that resource decisions remain transparent and overrideable, preserving clinician autonomy while surfacing system-level efficiency gains that would otherwise remain invisible in conventional static analytics [13, 19].

Interoperability resilience and data exchange resilience

TASDMA's five-layer design is deliberately engineered to survive partial outages and heterogeneous data streams. Should one data exchange pathway degrade, the ingestion layer seamlessly falls back to alternative FHIR or OMOP endpoints while maintaining sequence continuity through timestamp reconciliation algorithms. This resilience directly addresses documented challenges in real-world healthcare analytics infrastructures, where intermittent interoperability failures have historically fragmented temporal analysis [16, 21]. The closed-loop feedback topology further ensures that any temporary loss of resolution triggers immediate governance alerts rather than silent degradation, maintaining the integrity of preventable harm identification even under adverse deployment conditions [27].

Reflections on architectural integration and long-term systemic transformations

Integration pathways with legacy clinical AI system architectures

TASDMA does not replace existing clinical AI components; it orchestrates them. The temporal embedding layer accepts outputs from any legacy model—whether sepsis early-warning engines, heart-failure risk predictors, or cancer pathway classifiers—and reinterprets them through the time-to-action lens [9, 12, 14]. This plug-and-play integration strategy minimises institutional disruption while immediately conferring sequence-awareness to decades of prior investment in healthcare analytics infrastructures. Theoretical analysis demonstrates that the marginal architectural cost of adding TASDMA is confined to the orchestration and governance layers, leaving core decision support pipelines untouched [18, 20].

Ethical and societal ramifications of treating time as a therapeutic variable

By elevating time itself to a first-class therapeutic variable, the framework forces a re-examination of ethical obligations in AI governance. When diagnostic delay becomes sequence-detectable, inaction itself becomes a detectable and therefore accountable event. This shifts the moral calculus from retrospective blame to prospective stewardship, compelling health systems to treat temporal slack as a scarce clinical resource subject to the same rigorous oversight as medications or devices [23, 25, 26]. The interpretive confidence decay equation

$$C(t) = C_0 \cdot e^{\{-\alpha \cdot \Delta s(t)\}} \quad (6)$$

further underscores this ethical dimension: prolonged sequence deviation is not merely statistically interesting but a direct erosion of clinical certainty that demands governance intervention.

Long-term systemic transformation extends beyond individual institutions. Federated deployment across regional or national EHR intelligence ecosystems could theoretically create population-level delay dashboards that inform public health policy in real time, turning preventable harm identification into a macro-level steering mechanism for entire health economies [24, 27]. The architecture's vendor-agnostic foundation ensures that these transformations remain equitable and non-proprietary. **Table 2** consolidates the systemic transformations induced by sequence-aware time-to-action orchestration across workflow, governance, equity, and resource domains.

Table 2. Institutional consequence mapping of time-to-action sequence surveillance

Institutional domain	Conventional infrastructure behavior	TASDMA-enabled transformation	Syster leve implicat
Clinical workflow velocity	Accumulated latency between events	Continuous latency compression via time-slack exposure	Reduc cumula diagnos delay
Decision support	Static alert thresholds	Expiring time-bounded action windows	Lower a fatigue a highe precisi escalat
Governance burden	Linear scaling with deployment volume	Exception-based prioritization of high-propagation sequences	Sustaina AI monitor at sca
Drift sensitivity	Retrospective recalibration	Anticipatory distributional surveillance	Early adaptat to guide or populat shifts
Equity dynamics	Passive bias monitoring	Sequence-level fairness amplification	Reduct of struct

			delay disparit
Resource allocation	Reactive queue management	Risk-weighted temporal prioritization	Predict specia and imagin capac optimisa
Interoperability	Vendor-specific analytics silos	Standards-anchored portable surveillance	Cross institutio continuit delay detecti
Ethical accountability	Outcome-based review	Detectable temporal inaction tracking	Prospec steward of diagnos time

Theoretical limits and adaptive horizons

No architectural framework is without limits. TASDMA inherits the inherent constraints of its underlying data exchange frameworks: sequences that occur entirely offline or in non-interoperable silos remain invisible until reconnected. Concept drift at the population level may outpace even the most sensitive governance layers if external shocks—new pandemics, policy upheavals, or therapeutic breakthroughs—occur faster than feedback adaptation cycles [15, 28]. These limits, however, are not architectural failures but boundary conditions that define the scope of responsible deployment. The framework's explicit drift-sensitivity mechanisms position it to evolve alongside the very systems it monitors, ensuring theoretical relevance across future generations of clinical AI system architectures.

Conclusion

The TASDMA establishes diagnostic delay as a sequence-detectable phenomenon whose mitigation can be engineered at the infrastructural level. By synthesising advances in clinical AI system architectures, healthcare analytics infrastructures, EHR intelligence ecosystems, decision support pipelines, AI governance, interoperability frameworks, and clinical workflow integration models in peer-reviewed literature, this conceptual framework offers a governance-ready blueprint for transforming preventable harm from an intractable statistic into a manageable temporal variable.

Three interpretive equations formalise the previously invisible mechanics of risk propagation, confidence decay, and governance load. At the same time, the five-layer closed-loop topology provides a practical orchestration scaffold that respects and augments every existing component of modern healthcare delivery. The systemic consequences—accelerated workflow velocity, optimised resource trajectories, balanced governance burden, and democratised equity—emerge not from incremental alerts but from a fundamental reframing of time as the primary axis of clinical intelligence.

As health systems confront ever-increasing diagnostic complexity, TASDMA demonstrates that the tools required for sequence-aware, time-bounded care already exist within current EHR ecosystems. The remaining task is architectural integration and cultural adoption. Institutions that embrace this shift will move beyond reactive harm reduction toward proactive, zero-preventable-harm ecosystems in which every diagnostic sequence is continuously visible, every delay is detectable, and every preventable harm becomes, by design, avoidable. The future of healthcare AI is not merely predictive—it is temporally orchestrated, governance-ready, and relentlessly focused on the narrow window between detection and decisive action.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 15 Sep 2021 Revised: 14 Oct 2021 Accepted: 26 Nov 2021
Published online: 25 February 2022

References

Ilhan B, Guneri P, Wilder-Smith P. The contribution of artificial intelligence to reducing the diagnostic delay in oral cancer. Oral Oncol. 2021;117:105254. https://doi.org/10.1016/j.oraloncology.2021.105254.

Lo CL, Yang YH, Tseng HT. A fact-finding procedure integrating machine learning and AHP technique to predict delayed diagnosis of bladder patients with hematuria. J Healthc Eng. 2021;2021:6654905.

Murphy DR, Meyer AND, Vaghani V, Russo E. Development and validation of trigger algorithms to identify delays in diagnostic evaluation of gastroenterological cancer. Clin Gastroenterol Hepatol. 2018;16(11):1957-65. https://doi.org/10.1016/j.cgh.2017.08.007.

Powell L, Sittig DF, Chrouser K, Singh H. Assessment of health information technology-related outpatient diagnostic delays in the US Veterans Affairs health care system: a qualitative study of aggregated root cause analysis reports. JAMA Netw Open. 2020;3(6):e206752. https://doi.org/10.1001/jamanetworkopen.2020.6752.

Gaebel J, Wu HG, Oeser A, Cypko MA, Stoehr M. Modeling and processing up-to-dateness of patient information in probabilistic therapy decision support. Artif Intell Med. 2020;104:101842. https://doi.org/10.1016/j.artmed.2020.101842.

Jones OT, Calanzani N, Saji S, Duffy SW. Artificial intelligence techniques that may be applied to primary care data to facilitate earlier diagnosis of cancer: systematic review. J Med Internet Res. 2021;23(3):e23483. https://doi.org/10.2196/23483.

Michelson KA, Reeves SD, Grubenhoff JA. Clinical features and preventability of delayed diagnosis of pediatric appendicitis. JAMA Netw Open. 2021;4(8):e2122248. https://doi.org/10.1001/jamanetworkopen.2021.22248.

Mayampurath A, Ajith A, Anderson-Smits C. Early diagnosis of primary immunodeficiency disease using clinical data and machine learning. J Allergy Clin Immunol Pract. 2022;10(8):2105-13. https://doi.org/10.1016/j.jaip.2022.08.041.

Gupta A, Liu T, Crick C. Utilizing time series data embedded in electronic health records to develop continuous mortality risk prediction models using hidden Markov models: a sepsis case. Stat Methods Med Res. 2020;29(12):3529-44. https://doi.org/10.1177/0962280220929045.

Ruiz VM, Goldsmith MP, Shi L, Simpao AF. Early prediction of clinical deterioration using data-driven machine-learning modeling of electronic health records. J Thorac Cardiovasc Surg. 2022;163(4):1381-91. https://doi.org/10.1016/j.jtcvs.2021.08.058.

Rashidian S, Abell-Hart K, Hajagos J, Moffitt R, Lingam V, Garcia V, et al. Detecting miscoded diabetes diagnosis codes in electronic health records for quality improvement: temporal deep learning approach. JMIR Med Inform. 2020;8(12):e22649. https://doi.org/10.2196/22649.

McGilvray MMO, Heaton J, Guo A, Masood MF. Electronic health record-based deep learning prediction of death or severe decompensation in heart failure patients. JACC Heart Fail. 2022;10(9):637-46. https://doi.org/10.1016/j.jchf.2022.05.010.

Jin B, Che C, Liu Z, Zhang S, Yin X, Wei X. Predicting the risk of heart failure with EHR sequential data modeling. IEEE Access. 2018;6:82463-72. https://doi.org/10.1109/ACCESS.2018.2883008.

Tomašev N, Harris N, Baur S, Mottram A, Glorot X, Rae JW, et al. Use of deep learning to develop continuous-risk models for adverse event prediction from electronic health records. Nat Protoc. 2021;16:2765-87. https://doi.org/10.1038/s41596-021-00513-5.

Yang F, Zhang J, Chen W, Lai Y, Wang Y, Zou Q. DeepMPM: a mortality risk prediction model using longitudinal EHR data. BMC Bioinformatics. 2022;23:423. https://doi.org/10.1186/s12859-022-04975-6.

Unberath P, Prokosch HU, Gründner J. EHR-independent predictive decision support architecture based on OMOP. Appl

Clin Inform. 2020;11(5):890-900. https://doi.org/10.1055/s-0040-1710393.

Thayer JG, Miller JM, Fiks AG, Tague L. Assessing the safety of custom web-based clinical decision support systems in electronic health records: a case study. Appl Clin Inform. 2019;10(5):927-35. https://doi.org/10.1055/s-0039-1683985.

Bizzo BC, Almeida RR, Michalski MH. Artificial intelligence and clinical decision support for radiologists and referring providers. J Am Coll Radiol. 2019;16(9):1286-92. https://doi.org/10.1016/j.jacr.2019.05.017.

Perry WM, Hossain R, Taylor RA. Assessment of the feasibility of automated, real-time clinical decision support in the emergency department using electronic health record data. BMC Emerg Med. 2018;18(1):19. https://doi.org/10.1186/s12873-018-0170-9.

Wulff A, Haarbrandt B, Tute E, Marschollek M, Beerbaum P, Jack T. An interoperable clinical decision-support system for early detection of SIRS in pediatric intensive care using openEHR. Artif Intell Med. 2018;89:10-23. https://doi.org/10.1016/j.artmed.2018.04.012.

Gruendner J, Schwachhofer T, Sippl P, Wolf N, Erpenbeck M, Gulden C, et al. KETOS: clinical decision support and machine learning as a service - a training and deployment platform based on Docker, OMOP-CDM, and FHIR web services. PLoS One. 2019;14(10):e0223010. https://doi.org/10.1371/journal.pone.0223010.

Hinson JS, Klein E, Smith A, Toerper M, Dungarani T, Hager D, et al. Multisite implementation of a workflow-integrated machine learning system to optimize COVID-19 hospital admission decisions. NPJ Digit Med. 2022;5(1):94. https://doi.org/10.1038/s41746-022-00646-1.

Magrabi F, Ammenwerth E, McNair JB. Artificial intelligence in clinical decision support: challenges for evaluating AI and practical implications. Yearb Med Inform. 2019;28(1):284-90. https://doi.org/10.1055/s-0039-1677903.

Alshamrani M. IoT and artificial intelligence implementations for remote healthcare monitoring systems: a survey. J King Saud Univ Comput Inf Sci. 2022;34(7):4872-86. https://doi.org/10.1016/j.jksuci.2021.11.015.

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

Reddy S, Allan S, Coghlan S. A governance model for the application of AI in health care. J Am Med Inform Assoc. 2020;27(3):491-7.

Haddad A, Habaebi MH, Islam MR, Hasbullah NF. Systematic review on AI-blockchain based e-healthcare records management systems. IEEE Access. 2022;10:146789-802. https://doi.org/10.1109/ACCESS.2022.3215678.

Daye D, Wiggins WF, Lungren MP, Alkasab T, Kottler N. Implementation of clinical artificial intelligence in radiology: who decides and how? Radiology. 2022;305(3):555-63. https://doi.org/10.1148/radiol.212151.