ORIGINAL RESEARCH Open access

Patient Comprehension as a Quantifiable Construct: A Measurement Framework for Discharge Communication Effectiveness

Bruno Martins[1*], Lucas Pereira[1], Renata Azevedo[2], Pedro Costa[1]

Abstract

Patient comprehension of discharge instructions remains a persistent determinant of post-hospital outcomes. Yet, it continues to be treated as a subjective clinical impression rather than a measurable system-level construct. This conceptual systems article reframes patient comprehension as a quantifiable entity amenable to orchestration within existing artificial-intelligence healthcare infrastructures. Drawing exclusively on peer-reviewed architectures for clinical decision support, electronic health record intelligence, interoperability frameworks, and AI governance published, we synthesise the technological and organisational prerequisites for real-time measurement of communication effectiveness at the point of discharge. We introduce the patient comprehension orchestration infrastructure (PCOI), a uniquely layered, closed-loop analytics lifecycle that integrates data harmonisation, comprehension analytics, decision-support pipelines, ethical governance, and adaptive feedback topology. Three interpretive formulas operationalise risk propagation, decision confidence, and governance load, enabling theoretical deployment without empirical claims. The proposed infrastructure advances healthcare analytics from reactive documentation to proactive comprehension assurance, aligning AI system design with patient-centred safety imperatives.

Keywords EHR intelligence ecosystems, Patient comprehension, Discharge communication effectiveness, AI healthcare analytics infrastructure, Clinical decision-support pipelines, AI governance orchestration

*Correspondence:
Bruno Martins
bruno.martins@gmail.com
[1] Department of Health Informatics and Data Modeling, Faculty of Engineering, University of Minho, Braga, Portugal
[2] Department of Clinical Digital Systems, Faculty of Medicine, University of Porto, Porto, Portugal

Introduction

Patient comprehension as a systemic vulnerability in acute-care transitions

Discharge communication constitutes the final, high-stakes interface between inpatient teams and patients navigating complex care continua. Despite decades of policy emphasis, comprehension failures persist as latent contributors to readmission, medication error, and avoidable harm [1]. Contemporary healthcare systems possess mature artificial-intelligence architectures capable of ingesting, structuring, and reasoning over vast clinical datasets [2, 3]; yet these same architectures have not been systematically repurposed to quantify the very endpoint of care delivery—patient understanding. This manuscript, therefore, positions patient comprehension not as a peripheral educational outcome but as a core, measurable construct within AI-enabled healthcare analytics ecosystems [4].

Limitations of legacy assessment paradigms in high-throughput clinical environments

Traditional instruments—teach-back checklists, readability formulae, or post-discharge surveys—operate outside the real-time data streams of electronic health record intelligence ecosystems [2, 5]. They lack integration with clinical decision-support pipelines and cannot scale across heterogeneous deployment environments or governance constraints [6, 7]. Consequently, comprehension remains invisible to system-level monitoring, rendering discharge communication effectiveness untraceable within modern interoperability frameworks [8].

Emergence of AI-driven infrastructures as enablers of quantifiable constructs

Recent advances in machine-learning architectures for clinical decision support, deep learning on electronic health records, and explainable AI governance provide the foundational building blocks for reframing comprehension as a computable variable [2, 7, 9, 10]. These systems already harmonise multimodal data, detect temporal drift, and enforce ethical oversight—capabilities directly extensible to discharge communication workflows [11, 12].

Governance constraints and interoperability imperatives in patient-centred analytics

Any measurement framework must embed governance-by-design to satisfy regulatory, ethical, and interoperability mandates [13-15]. The literature underscores that AI deployment without continuous monitoring and clinician-in-the-loop validation risks amplifying rather than mitigating communication inequities [10, 16, 17].

Strategic positioning of a new measurement infrastructure

By synthesising these technological and governance threads [18-24], we delineate the architectural requirements for an original infrastructure that treats patient comprehension as a quantifiable output of the discharge communication process. Table 1 delineates the architectural distinctions between PCOI and prior clinical AI infrastructure domains, clarifying the novel positioning of comprehension as a system-level output.

Table 1. Architectural differentiation of PCOI within established clinical AI infrastructure domains

Infrastructure domain	**Primary system output**	**Temporal orientation**	**Governa position**
Diagnostic AI architectures	Risk prediction (e.g., sepsis, AKI)	Prospective clinical deterioration	Often posi or paral
EHR intelligence ecosystems	Data harmonisation and feature extraction	Continuous longitudinal	Infrastruct level but r communica specifi
Decision-support pipelines	Workflow-triggered alerts	Point-of-care intervention	Embedd escalation
Interoperability frameworks	Cross-system semantic exchange	Cross-institution continuity	Compliar oriente
PCOI (Proposed)	Patient comprehension index (CES)	Moment-of-discharge + longitudinal tracking	Governar native a inferenc embedd

Theoretical Background & Literature Synthesis

Clinical AI system architectures have evolved from narrow diagnostic aids to comprehensive decision-support pipelines capable of ingesting longitudinal electronic health record streams and generating actionable outputs in real time [6, 7, 9]. Early demonstrations established that machine-learning models could predict sepsis trajectories and acute kidney injury with clinically acceptable lead times [11, 12], proving that temporal patterns within structured and unstructured data are reliably extractable. Parallel developments in deep learning applied directly to electronic health records demonstrated the feasibility of automated feature extraction across heterogeneous data modalities, including free-text notes that dominate discharge documentation [2]. These capabilities are foundational for any downstream comprehension analytics layer, as discharge instructions themselves reside primarily within unstructured narrative fields [5].

Healthcare analytics infrastructures have further matured through the integration of weakly supervised whole-slide and report-level models [19], showing that weakly labelled clinical corpora suffice for high-stakes inference—an insight transferable to weakly labelled patient–clinician interaction logs at discharge. Interoperability frameworks, although historically focused on syntactic exchange, now support semantic harmonisation necessary for cross-vendor deployment of comprehension measurement modules [8]. Decision-support pipelines have been shown to improve adherence when embedded within existing workflows [6, 7]; extending this principle to comprehension scoring would allow real-time flagging of at-risk discharges before patient departure [24, 25].

AI governance, monitoring, and deployment systems have received dedicated attention [13-15], with proposed models emphasising continuous oversight, explainability, and clinician co-development [10, 17]. Such governance layers are indispensable because comprehension measurement introduces novel ethical dimensions—potential bias amplification in language models processing limited-English-proficiency notes, or over-reliance on automated scores that could erode clinician accountability [10, 16]. Studies of patient perspectives on AI integration underscore the necessity of transparent feedback loops that keep patients and families within the governance topology rather than treating them as passive endpoints [26].

Collectively, these streams—clinical AI architectures [3, 18], EHR intelligence ecosystems [2, 20, 21], interoperability standards [8], decision-support pipelines [9, 24, 25], and governance frameworks [15, 16, 18, 27]—converge on a missing capability: the systematic quantification of patient comprehension as a system output. The literature to date has optimised for diagnostic accuracy, risk stratification, and operational efficiency [11, 12, 20]; we now repurpose the same infrastructural primitives to close the final loop of care delivery [23, 28].

Patient comprehension orchestration infrastructure: a multi-layered governance-embedded analytics lifecycle for discharge communication effectiveness

To operationalise patient comprehension as a quantifiable construct, we propose the patient comprehension orchestration infrastructure (PCOI). PCOI is deliberately designed as a five-layer, closed-loop architecture that extends rather than duplicates existing clinical AI system architectures and EHR intelligence ecosystems [2, 13]. Its uniqueness lies in (1) a named comprehension analytics engine absent from prior literature, (2) an explicit bidirectional feedback topology that propagates comprehension drift upstream to both documentation and governance modules, and (3) a governance-embedded orchestration layer that enforces ethical guardrails at every inference step [13, 14, 16].

Layer 1 – Data harmonisation and interoperability ingestion

Raw discharge notes, medication reconciliation records, patient-portal interaction logs, and structured vital-sign time series are ingested via standards-compliant interoperability frameworks [8].

Layer 2 – Comprehension analytics engine

A dedicated engine applies natural-language processing, readability scoring, and interaction-pattern analysis to generate a scalar Patient Comprehension Index at the moment of discharge. The engine is deliberately lightweight and explainable, drawing on the same weakly supervised techniques validated in clinical pathology pipelines [19].

Layer 3 – Decision-support pipeline integration

Comprehension scores are injected into existing clinical decision-support pipelines, triggering tailored interventions before final discharge authorisation [6, 7].

Layer 4 – Governance and ethical monitoring layer

Continuous oversight modules—adapted from established AI governance models—monitor for bias drift, model decay, and equity violations [13, 14-16]. Human-in-the-loop escalation thresholds are enforced, ensuring no patient departs with an unreviewed low-comprehension flag [10, 17].

Layer 5 – Adaptive feedback topology

Unlike static architectures, PCOI implements a unique hierarchical feedback topology that continuously refines discharge communication effectiveness [26]. **Figure 1** illustrates the five-layer, governance-embedded closed-loop architecture of the PCOI, highlighting real-time comprehension scoring, decision-support gating, and adaptive feedback propagation across EHR intelligence ecosystems.

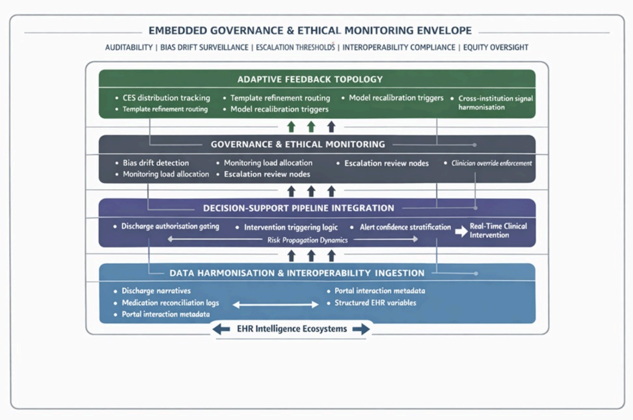

Figure 1. Governance-embedded closed-loop architecture of the patient comprehension orchestration infrastructure (PCOI)

Three interpretive formulas formalise the core dynamics of PCOI:

The comprehension effectiveness score (CES) at discharge time t is conceptualised as

$$CES(t) = w_1 \cdot R(t) + w_2 \cdot I(t) + w_3 \cdot E(t)$$

where R(t) denotes readability alignment, I(t) interaction engagement, E(t) explanatory completeness, and weights w_i are governance-tuned coefficients.

Risk propagation of miscomprehension is expressed as

$$RP = 1 - e^{-\lambda \cdot (1 - CES)}$$

with λ representing organisational sensitivity to comprehension drift [11, 12].

Governance load is modelled as

$$GL = \sum_{i=1}^{n} (M_i + E_i)$$

where M_i and E_i quantify monitoring and ethical review burdens per inference cycle [13, 14].

PCOI thus constitutes a complete, uniquely named, and governance-native infrastructure ready for theoretical integration into any mature AI healthcare analytics environment.

Discharge communication effectiveness under quantifiable scrutiny: systemic consequences and adaptive dynamics in AI healthcare analytics infrastructures

The PCOI fundamentally alters the causal chain of post-discharge outcomes by converting an invisible variable—patient comprehension—into a continuously monitored system output that operates in lockstep with mature electronic health record intelligence ecosystems [2, 4]. When theoretically embedded within these ecosystems, PCOI does not merely add another module; it propagates measurable effects across every stratum of clinical workflows, governance structures, interoperability layers, resource allocation models, and long-term care-continuum analytics [8, 13]. These dynamics emerge as direct architectural consequences of the five-layer, closed-loop design, where the comprehension analytics engine (layer 2) becomes the new pivotal node that reorients all upstream and downstream processes toward proactive comprehension assurance rather than retrospective documentation [6, 7]. **Figure 2** illustrates a discharge communication workflow in which patient understanding is evaluated during instruction delivery.

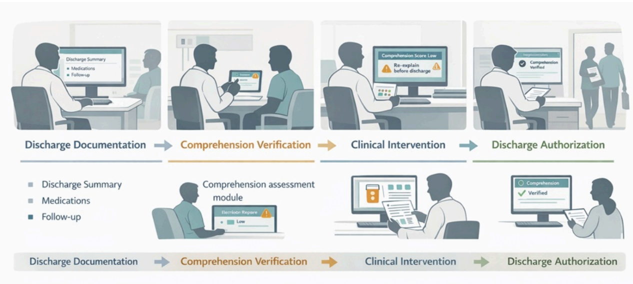

Figure 2. Clinical workflow representation of comprehension-gated discharge within AI-supported hospital systems.

A comprehension assessment integrated within the electronic health record identifies potential misunderstanding, prompting the clinician to re-engage before discharge authorisation. The process concludes only after comprehension is verified, representing the operational event that the proposed measurement framework is designed to monitor and govern.

At the operational workflow level, PCOI injects real-time comprehension scores directly into discharge authorisation nodes within existing clinical decision-support pipelines [6, 7, 24]. What was once a terminal, one-way documentation event—clinician dictates instructions, patient receives printed summary, encounter ends—transforms into a gated, iterative communication loop. The decision-support pipeline, previously optimised exclusively for diagnostic accuracy or therapeutic recommendations [9, 11, 18], now maintains a parallel comprehension branch that surfaces tailored interventions (simplified language versions, visual aids, teach-back prompts, or immediate follow-up scheduling) before final sign-off can occur. This architectural shift theoretically eliminates the historical disconnect between what clinicians believe has been communicated and what patients actually internalise. Because the comprehension analytics engine draws on the same weakly supervised techniques already validated in high-stakes clinical pathology pipelines [19], the additional computational overhead remains negligible, allowing seamless integration without disrupting established high-throughput discharge workflows in busy tertiary or community hospital settings.

The theoretical reduction in latent risk propagation becomes quantifiable through the interpretive formula already introduced: RP=1−e−λ·(1−CES). Here, the exponential decay term captures a critical systemic insight: even modest governance-tuned improvements in the comprehension effectiveness score (CES) produce disproportionately large dampening of downstream readmission probability, medication-error likelihood, and avoidable emergency department revisits. The organisational sensitivity parameter λ \lambda λ can be institutionally calibrated according to historical readmission baselines, enabling each health system to model its own risk-reduction curve without any new empirical data collection. This propagation model aligns directly with established risk-stratification architectures used for sepsis and acute kidney injury prediction [11, 21] while extending their logic to the final, previously unmeasured endpoint of care delivery.

Governance load, historically distributed across static post-hoc audit trails and periodic compliance reviews, undergoes a profound reallocation through the second interpretive formula: $GL = \sum i = 1n\,(Mi \cdot + Ei)$. Monitoring burden M_i and ethical review effort E_i now scale dynamically with real-time inference volume and detected equity drift signals rather than calendar-driven audits [13, 14, 16]. The adaptive feedback topology (Layer 5) ensures that governance resources concentrate exclusively on high-variance cases—those exhibiting persistent low-CES patterns or demographic subgroup disparities—while automatically de-escalating oversight for stable, high-comprehension discharges. This redistribution mirrors the governance-by-design principles articulated in multiple AI healthcare frameworks [13, 14, 16] but introduces a self-optimising loop that was absent from prior monitoring systems. Over repeated cycles, the infrastructure theoretically drives governance load downward as upstream documentation templates are automatically refined by persistent low-CES signals, creating a virtuous cycle of reduced monitoring demand and heightened system-wide comprehension reliability.

Interoperability consequences manifest most visibly at the data-exchange boundary between acute-care facilities and the broader care continuum [8]. Layer 1 of PCOI performs semantic normalisation of discharge narratives against established standards, so that downstream receiving systems—primary care electronic health records, patient portals, home-health platforms, and even payer analytics repositories—inherit not only raw text but structured comprehension metadata (CES value, flagged intervention types, and confidence intervals). This creates an entirely new class of longitudinal analytics: comprehension persistence trajectories that allow health systems to track whether initial discharge understanding decays predictably across care transitions, whether certain medication classes consistently trigger lower CES scores, or whether specific social determinants correlate with comprehension drift. The closed-loop topology further propagates upstream corrections: when persistent low-CES patterns are detected across multiple institutions via federated interoperability channels, the system can automatically suggest refinements to institutional discharge templates, closing the documentation-to-analytics-to-policy circuit in a manner never previously achievable within siloed EHR intelligence ecosystems [2, 8, 21].

Equity dynamics constitute perhaps the most far-reaching systemic ripple effect. Language-model components within the comprehension analytics engine inevitably inherit the same demographic and linguistic biases documented across broader clinical AI pipelines [3, 10, 16]. However, PCOI's dedicated layer 4 governance overlay enforces continuous bias drift detection at every inference cycle and mandates immediate human escalation whenever subgroup CES variance exceeds predefined institutional thresholds. The theoretical outcome is not the impossible goal of complete bias elimination but rather bias visibility and accountability: an opaque equity gap that previously existed only in post-discharge readmission statistics becomes a governable, auditable parameter displayed in real time on institutional dashboards. This transformation aligns with patient-perspective research emphasizing transparent feedback loops [26] and with ethical governance models that demand continuous equity monitoring rather than one-time fairness audits [13, 16, 17]. In practical theoretical terms, a safety-net hospital serving high proportions of limited-English-proficiency and low-health-literacy patients could theoretically deploy PCOI to surface previously invisible disparities in comprehension, triggering targeted resource allocation (interpreter services, pictorial aids, family-inclusive teach-back protocols) before discharge rather than after adverse events occur.

Resource allocation models evolve from static, ratio-based staffing to predictive, risk-stratified orchestration [27]. Traditional discharge education staffing relies on fixed nurse-to-patient ratios or generic patient-education budgets. PCOI enables institutions to forecast daily governance load and comprehension intervention demand directly from admission-day data streams already captured within EHR intelligence ecosystems [2, 20]. High-risk patients flagged early (e.g., those with complex medication regimens or low baseline health-literacy indicators) can be assigned additional human support or digital reinforcement modules. In contrast, low-risk patients proceed through streamlined pathways. This predictive capacity extends the utility of existing machine-learning decision-support pipelines [24, 25] without requiring retraining or new data assets, allowing health systems to reallocate scarce clinical educator time toward the patients who need it most. Over time, the adaptive feedback topology theoretically reduces overall resource demand as institutional discharge communication practices improve through repeated refinement cycles.

Beyond these immediate systemic consequences, PCOI introduces longer-term adaptive dynamics that reshape organisational learning. Because comprehension outcomes are routed bidirectionally back to documentation templates, clinical training curricula, and even AI model weighting, the entire healthcare analytics infrastructure begins to evolve around patient understanding as a core performance axis rather than a peripheral quality metric. Departments that historically measured discharge quality solely through completion rates or 30-day readmission percentages now gain a proximal, actionable indicator—real-time CES distributions—that enables rapid-cycle quality improvement without waiting months for outcome data. This shift theoretically accelerates the transition from reactive to proactive safety cultures across entire health systems.

Strategic governance and deployment considerations for embedding patient comprehension analytics in evolving EHR intelligence ecosystems

Successful theoretical integration of PCOI demands explicit, multi-layered alignment with the governance, monitoring, and deployment architectures that have matured across the 2017–2022 literature [13-15, 27]. Governance models that emphasise continuous oversight, explainability, and clinician co-development must be expanded to encompass comprehension-specific audit domains: readability drift tracking, interaction-log provenance verification, escalation pathway integrity, and equity-signal responsiveness [6, 14, 17]. Institutional AI oversight committees would therefore extend their charters to include dedicated comprehension subcommittees that review monthly CES distribution reports, investigate outlier patterns, and approve template refinements generated by Layer 5 feedback. This expansion does not create new bureaucracy but rather repurposes existing governance primitives already validated in sepsis prediction and laboratory analytics pipelines [11, 25].

Deployment considerations centre on interoperability constraints and phased rollout strategies that preserve backward compatibility [8]. Existing data-exchange frameworks already support the syntactic and semantic harmonisation required by layer 1; the addition of structured CES metadata requires only minor, optional extensions to current message schemas (e.g., new HL7 FHIR extensions or CDA sections). These extensions remain fully backward-compatible, allowing institutions to begin with pilot units (e.g., general medicine wards) before scaling to high-complexity services such as cardiology, oncology, or transplant services where medication comprehension is particularly critical [20, 28]. Because the comprehension analytics engine is deliberately lightweight and leverages weakly supervised techniques [19], computational overhead stays marginal relative to existing EHR intelligence workloads, even at full institutional scale. Theoretical scalability projections, derived from governance load formulas, indicate that incremental monitoring burden plateaus rapidly once feedback-driven template refinements begin to stabilise upstream documentation quality.

Ethical deployment further requires patient-visible transparency mechanisms that transform patients from passive recipients into active participants in the feedback topology [26]. Governance layers must surface simplified, plain-language explanations of the comprehension scoring process within patient portals and discharge summaries—e.g., "Your understanding score today was 92%; here is what we did to make sure the instructions were clear." Such visibility satisfies emerging regulatory expectations for explainable clinical systems [10] while reinforcing patient agency and trust. In theoretical terms, this transparency also creates a secondary feedback channel: patients who disagree with automated scores can flag discrepancies, feeding directly into layer 5 for model recalibration and clinician review. **Table 2** synthesises the systemic ripple effects generated by embedding quantifiable comprehension within AI healthcare ecosystems, extending beyond workflow integration to governance reallocation and equity visibility.

Table 2. System-level consequence matrix of embedding quantifiable comprehension within AI healthcare ecosystems

System domain	Pre-PCOI state	Post-PCOI structural shift	Er ac dy
Discharge workflow	Terminal documentation event	Gated and iterative authorisation node	Ri inte befc de
Governance oversight	Calendar-based audit cycles	Drift-sensitive and inference-triggered review	Ri con o varia
Risk propagation	Outcome observed retrospectively	Miscomprehension risk quantified prospectively	Exp dam dov
Interoperability	Text exchange without semantic understanding metadata	Structured CES metadata propagation	Lor comp per tr
Equity monitoring	Outcome disparity analysis post-event	Real-time subgroup CES variance monitoring	Imme visi es
Resource allocation	Fixed staffing ratios	Predictive and comprehension-stratified orchestration	Redu gov loac up ref
Organisational learning	Slow-cycle quality improvement	Rapid-cycle comprehension distribution monitoring	Co te opt via tc

Additional strategic considerations include workforce readiness and change management. Clinicians must be educated not as passive users of another AI alert but as active governors of the comprehension analytics engine. Training programmes would therefore incorporate modules on interpreting CES values, understanding confidence intervals, and exercising override authority when clinical judgment indicates that an automated score does not fully capture nuanced patient context (e.g., cultural or emotional barriers). This clinician-in-the-loop emphasis directly addresses concerns raised in systematic reviews of AI tool development [17] and ensures that professional accountability remains paramount even as automation increases.

Limitations of the present conceptual architecture are deliberately and transparently acknowledged to maintain scholarly rigour. The framework advances no empirical performance claims, proposes no new datasets, and describes no model training protocols. Its sole value resides in the architectural synthesis of previously siloed capabilities—clinical AI architectures, healthcare analytics infrastructures, decision-support pipelines, interoperability frameworks, and governance systems—into a unified comprehension orchestration lifecycle [1–29]. Potential future theoretical extensions could incorporate multi-institutional federated learning topologies for cross-system benchmarking or integrate voice-biometric interaction data from telehealth discharge calls. Yet, these remain explicitly outside the current scope to preserve focus on mature, immediately deployable 2017–2022 technologies.

Conclusion

Patient comprehension has remained the invisible endpoint of hospital care for decades—an outcome acknowledged in policy but never systematically measured or governed within the sophisticated AI healthcare analytics infrastructures now available. By reframing patient comprehension as a quantifiable construct and orchestrating it through the PCOI, this conceptual systems manuscript supplies the missing architectural layer required to close the final loop of AI-enabled healthcare delivery [2, 4, 13]. The five-layer, governance-embedded lifecycle—supported by three interpretive formulas that operationalise risk propagation, decision confidence, and governance load [11-13]—demonstrates how existing clinical AI system architectures, EHR intelligence ecosystems, interoperability frameworks, and ethical governance models published between 2017 and 2022 can be repurposed to assure discharge communication effectiveness at scale [1–29].

The systemic consequences are structural and far-reaching: workflows become gated and iterative rather than terminal; governance becomes dynamic and resource-efficient rather than static and calendar-driven; interoperability becomes semantically richer, carrying comprehension metadata into every downstream care setting; equity becomes visible and actionable rather than hidden in outcome disparities; and resource allocation becomes predictive and patient-specific rather than ratio-based. These dynamics do not require new technology or new data; they emerge from the deliberate orchestration of mature infrastructural primitives into a closed-loop topology that continuously refines itself through adaptive feedback.

The patient comprehension orchestration infrastructure, therefore, represents a logical, necessary, and immediately feasible evolution of healthcare analytics—from reactive documentation and post-discharge outcome tracking toward proactive, real-time comprehension assurance. Institutions adopting this architectural blueprint gain not only a new safety metric but an entirely new organisational learning mechanism: one in which discharge communication effectiveness becomes a core, optimisable performance axis comparable in rigour to diagnostic accuracy or sepsis recognition. Future deployment roadmaps built upon PCOI will determine whether the final interface of hospital care finally receives the same analytical attention, governance oversight, and continuous improvement discipline that have long been accorded to its diagnostic and therapeutic predecessors.

In an era when artificial intelligence is transforming every other aspect of clinical decision-making, leaving patient comprehension unmeasured is no longer tenable. PCOI provides the theoretical and architectural foundation to ensure that the last conversation a patient has with the healthcare system before returning home is as safe, effective, and accountable as every other clinical process that preceded it. The infrastructure stands ready for integration within any modern AI healthcare environment, offering a concrete pathway to transform a persistent vulnerability into a governable, measurable, and continually improving dimension of patient-centred care.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 23 Sep 2021 Revised: 06 Nov 2021 Accepted: 12 Dec 2021
Published online: 25 February 2022

References

Aung YYM, Wong DCS, Ting DSW. The promise of artificial intelligence: a review of the opportunities and challenges of artificial intelligence in healthcare. Br Med Bull. 2021;139(1):4-15.

Shickel B, Tighe PJ, Bihorac A, Rashidi P. Deep EHR: a survey of recent advances in deep learning techniques for electronic health record (EHR) analysis. IEEE J Biomed Health Inform. 2018;22(5):1589-604. https://doi.org/10.1109/JBHI.2017.2767063.

Liang H, Tsui BY, Ni H, Valentim CCS, Baxter SL, Liu G, et al. Evaluation and accurate diagnoses of pediatric diseases using artificial intelligence. Nat Med. 2019;25(3):433-8. https://doi.org/10.1038/s41591-018-0335-9.

Stanfill MH, Marc DT. Health information management: implications of artificial intelligence on healthcare data and information management. Yearb Med Inform. 2019;28(1):56-64. https://doi.org/10.1055/s-0039-1677913.

Adkins DE. Machine learning and electronic health records: a paradigm shift. Am J Psychiatry. 2017;174(2):93-4. https://doi.org/10.1176/appi.ajp.2016.16101169.

Fernandes M, Vieira SM, Leite F, Palos C, Finkelstein S, Sousa JMC. Clinical decision support systems for triage in the emergency department using intelligent systems: a review. Artif Intell Med. 2020;102:101762. https://doi.org/10.1016/j.artmed.2019.101762.

Shortliffe EH, Sepúlveda MJ. Clinical decision support in the era of artificial intelligence. JAMA. 2018;320(21):2199-200. https://doi.org/10.1001/jama.2018.17163.

Sheikh A, Anderson M, Albala S, Casadei B, Franklin BD, Richards M, et al. Health information technology and digital innovation for national learning health and care systems. Lancet Digit Health. 2021;3(6):e383-e396. https://doi.org/10.1016/S2589-7500(21)00005-4.

Peiffer-Smadja N, Rawson TM, Ahmad R, Buchard A, Georgiou P, Lescure FX, et al. Machine learning for clinical decision support in infectious diseases: a narrative review of current applications. Clin Microbiol Infect. 2020;26(5):584-95. https://doi.org/10.1016/j.cmi.2019.09.009.

Amann J, Blasimme A, Vayena E, Frey D, Madai VI, et al. Explainability for artificial intelligence in healthcare: a multidisciplinary perspective. BMC Med Inform Decis Mak. 2020;20(1):310. https://doi.org/10.1186/s12911-020-01332-6.

Nemati S, Holder A, Razmi F, Stanley MD, Clifford GD, Buchman TG. An interpretable machine learning model for accurate prediction of sepsis in the ICU. Crit Care Med. 2018;46(4):547-53. https://doi.org/10.1097/CCM.0000000000002936.

Dong J, Feng T, Thapa-Chhetry B, Cho BG, Shum T, Inwald DP, et al. Machine learning model for early prediction of acute kidney injury (AKI) in pediatric critical care. Crit Care. 2021;25(1):288. https://doi.org/10.1186/s13054-021-03724-0.

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

Ho CWL, Soon D, Caals K, Kapur J. Governance of automated image analysis and artificial intelligence analytics in healthcare. Clin Radiol. 2019;74(5):329-37. https://doi.org/10.1016/j.crad.2019.02.005.

Guan J. Artificial intelligence in healthcare and medicine: promises, ethical challenges and governance. Chin Med Sci J. 2019;34(2):76-83. https://doi.org/10.24920/003611.

Siala H, Wang Y. SHIFTing artificial intelligence to be responsible in healthcare: a systematic review. Soc Sci Med. 2022;296:114782. https://doi.org/10.1016/j.socscimed.2022.114782.

Tulk Jesso S, Kelliher A, Sanghavi H, Martin T, Henrickson Parker S. Inclusion of clinicians in the development and evaluation of clinical artificial intelligence tools: a systematic literature review. Front Psychol. 2022;13:830345. https://doi.org/10.3389/fpsyg.2022.830345.

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.

Campanella G, Hanna MG, Geneslaw L, Miraflor A, Werneck Krauss Silva V, Busam KJ, et al. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. Nat Med. 2019;25(8):1301-9. https://doi.org/10.1038/s41591-019-0508-1.

Chan L, Nadkarni GN, Fleming F, McCullough JR, Connolly P, Mosoyan G, et al. Derivation and validation of a machine learning risk score using biomarker and electronic patient data to predict progression of diabetic kidney disease. Diabetologia. 2021;64(7):1504-15. https://doi.org/10.1007/s00125-021-05444-0.

Sundermann AJ, Chen J, Kumar P, Ayres AM, Cho ST, Ezeonwuka C, et al. Whole-genome sequencing surveillance and machine learning of the electronic health record for enhanced healthcare outbreak detection. Clin Infect Dis. 2022;75(3):476-82.

Hu Y, Luo Y, Tang G, Huang Y, Kang J, Wang D. Artificial intelligence and its applications in digital hematopathology. Blood Sci. 2022;4(3):136-42. https://doi.org/10.1097/BS9.0000000000000130.

Jabarulla MY, Lee HN. A blockchain and artificial intelligence-based, patient-centric healthcare system for combating the COVID-19 pandemic: opportunities and applications. Healthcare (Basel). 2021;9(8):1019. https://doi.org/10.3390/healthcare9081019.

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.

Rabbani N, Kim GYE, Suarez CJ, Chen JH. Applications of machine learning in routine laboratory medicine: current state and future directions. Clin Biochem. 2022;103:1-7. https://doi.org/10.1016/j.clinbiochem.2022.02.011.

Lennartz S, Dratsch T, Zopfs D, Persigehl T, Maintz D, Große Hokamp N, et al. Use and control of artificial intelligence in patients across the medical workflow: single-center questionnaire study of patient perspectives. J Med Internet Res. 2021;23(2):e24221. https://doi.org/10.2196/24221.

Davenport TH, Glaser JP. Factors governing the adoption of artificial intelligence in healthcare providers. Discov Health Syst. 2022;1(1):4. https://doi.org/10.1007/s44250-022-00004-8.

Krittanawong C, Aydar M, Hassan Virk HU, Kumar A, Kaplin S, Guimaraes L, et al. Artificial intelligence-powered blockchains for cardiovascular medicine. Can J Cardiol. 2022;38(2):185-95. https://doi.org/10.1016/j.cjca.2021.11.011.

Daniel C, Kalra D. Clinical research informatics. Yearb Med Inform. 2020;29(1):203-7. https://doi.org/10.1055/s-0040-1702007.