

ORIGINAL RESEARCH

Open access

# A Systems-Level Architecture for AI-Enabled Hospital Readmission Risk Governance

Lucas Meyer<sup>1\*</sup>, Anna Schmid<sup>2</sup>, Stefan Braun<sup>1</sup>

## Abstract

Hospital readmission rates are a critical metric in healthcare systems, reflecting operational inefficiencies, patient safety risks, and resource-allocation challenges within clinical environments. AI-enabled analytics have emerged as tools for predicting and mitigating these risks. Yet their integration into hospital workflows demands robust governance architectures to address privacy, interoperability, and accountability for decision-making. This conceptual manuscript identifies a gap in systems-level frameworks that holistically govern AI-driven readmission risk models from data ingestion through clinical deployment. We propose the readmission risk oversight scaffold (RROS), a novel architecture comprising layered components for data harmonization, model monitoring, workflow integration, and governance feedback loops. RROS emphasizes interoperability with electronic health records (EHRs), privacy-preserving analytics pipelines, and clinician-AI collaboration to enhance risk governance. Implications include improved hospital resource management, reduced bias in predictive analytics, and scalable oversight mechanisms for AI in healthcare informatics. By framing readmission risk as a governed systems process, RROS offers interpretive insights into balancing technological capabilities with clinical imperatives, potentially informing future informatics infrastructures without empirical validation. This work underscores the need for architectural designs that prioritize safety and equity in AI-enabled hospital settings.

**Keywords** AI-enabled governance, Hospital readmission risk, Healthcare systems architecture, Clinical analytics pipelines, EHR interoperability, Decision support monitoring

\*Correspondence:

Lucas Meyer

[lucas.meyer@gmail.com](mailto:lucas.meyer@gmail.com)

<sup>1</sup> Department of Healthcare Data Science, Faculty of Engineering, ETH Zurich, Zurich, Switzerland

<sup>2</sup> Department of Clinical Information Systems, Faculty of Medicine, University of Bern, Bern, Switzerland

## Introduction

### Evolving paradigms in AI-enabled hospital readmission risk assessment

The integration of artificial intelligence into hospital readmission risk governance represents a significant paradigm shift in clinical informatics. Historically, readmission analysis served primarily as a retrospective quality-improvement metric, used to evaluate institutional performance after adverse utilization outcomes had occurred. Contemporary AI-enabled systems, however, reframe readmission risk as a prospective, continuously

monitored phenomenon embedded within clinical workflows. This transition from retrospective evaluation to proactive, systems-level intervention reflects broader transformations in digital health ecosystems, where predictive analytics increasingly shape operational and clinical decision-making.

Hospital readmission risk—particularly in the context of chronic conditions such as heart failure, chronic obstructive pulmonary disease, and post-surgical complications—presents complex, multifactorial determinants that extend beyond isolated clinical variables [1-5]. Socioeconomic conditions, care transitions, medication adherence, and

discharge planning all contribute to readmission probability. Consequently, governance architectures must move beyond single-model optimization toward integrated frameworks that contextualize AI-generated predictions within dynamic hospital environments.

AI-enabled tools that analyze electronic health records (EHRs), longitudinal patient trajectories, and transition-of-care data can generate sophisticated risk stratifications. Yet predictive accuracy alone does not constitute effective governance. Systems-level architectures must incorporate mechanisms for interpretability, accountability, workflow embedding, and regulatory compliance to ensure that predictive outputs translate into ethically and operationally sound interventions [3, 4, 6]. This evolution underscores the necessity of governance frameworks that treat readmission risk not merely as an algorithmic output but as a managed pipeline from data ingestion to clinical action.

## Challenges in interoperability for AI-enabled readmission risk governance

Interoperability remains a foundational challenge in deploying AI-enabled architectures for hospital readmission risk governance. Modern healthcare institutions operate across fragmented data landscapes, including EHR platforms, claims databases, laboratory systems, pharmacy records, and increasingly, wearable and remote monitoring devices. These heterogeneous data sources often lack standardized schemas, temporal synchronization, and unified access protocols, creating friction within analytics pipelines.

AI models designed to predict readmission risk must traverse these silos to produce actionable insights. However, without robust governance structures, fragmented data integration may propagate inconsistencies, amplify documentation biases, and reduce prediction reliability [7-9]. Errors introduced during harmonization—such as misaligned timestamps or inconsistent coding standards—can cascade through predictive systems, undermining clinician trust and the validity of decision support.

Systems-level designs should therefore prioritize standardized interfaces, interoperable APIs, and governance-aware integration protocols that enable seamless AI embedding within hospital workflows. Interoperability is not solely a technical objective; it functions as a governance enabler, facilitating

transparency, traceability, and coordinated oversight. In the context of inpatient-outpatient transitions—where readmission risks dynamically evolve—interoperability becomes central to sustaining continuity in predictive monitoring [10, 11]. Effective governance architectures must conceptualize interoperability as a structural pillar of AI-enabled readmission management rather than a peripheral technical feature.

## Privacy imperatives in systems-level hospital readmission risk architectures

Privacy considerations occupy a central position in AI-enabled hospital readmission risk governance, given the sensitive nature of health data processed by predictive infrastructures. Systems-level architectures must embed privacy-aware mechanisms that safeguard patient information across the entire data lifecycle—from ingestion and model training to deployment and monitoring—while preserving analytic utility [12-14].

The expansion of data modalities and cross-institutional analytics introduces heightened exposure risks. Federated learning approaches, distributed model training environments, and data-sharing agreements between hospitals or public health agencies can enhance predictive performance, yet they also introduce governance complexities. Without structured oversight, these mechanisms may expose vulnerabilities in readmission risk models, including unintended data leakage or insufficient auditability.

Effective architectures, therefore, incorporate privacy-enhancing technologies such as differential privacy, secure multi-party computation, role-based access controls, and immutable audit trails [15-17]. These safeguards extend governance beyond model outputs to encompass infrastructural resilience and institutional accountability. In this context, privacy is not treated as a compliance afterthought but as an embedded architectural principle that sustains stakeholder trust and regulatory alignment within AI-enabled readmission ecosystems.

## Bias mitigation strategies for AI-enabled readmission risk governance frameworks

Bias in AI models presents a significant governance concern in hospital readmission risk assessment. Predictive

systems trained on historical healthcare data may inadvertently encode longstanding disparities related to socioeconomic status, race, access to care, or institutional resource allocation. Without structured oversight, these embedded inequities may manifest as disproportionate risk stratification outcomes across patient cohorts [18–20].

Systems-level governance architectures must therefore integrate bias detection and mitigation protocols as core components rather than optional evaluative add-ons. This includes fairness audits, subgroup performance analyses, and monitoring of downstream clinical interventions influenced by risk predictions. Architectural designs should ensure that bias-auditing mechanisms are embedded in clinical workflows, enabling transparency into how readmission risks are assigned and interpreted.

Importantly, bias mitigation is not limited to statistical parity adjustments; it involves socio-technical engagement with how predictive insights are operationalized. Governance layers that facilitate interdisciplinary review and contextual recalibration of models reinforce institutional accountability. Such strategies underscore the interpretive role of governance in refining AI-enabled systems for equitable hospital deployment [21, 22].

## Positioning the need for a novel systems-level architecture

Collectively, evolving predictive paradigms, interoperability constraints, privacy imperatives, and bias mitigation strategies reveal a conceptual gap in current scholarship. Existing literature extensively addresses model development, algorithmic fairness, and ethical guidelines; however, it rarely synthesizes these components into a cohesive, systems-level governance architecture tailored specifically to AI-enabled hospital readmission risk management [23, 24].

This manuscript positions the readmission risk oversight scaffold (RROS) as an original conceptual contribution designed to address this void. Rather than focusing solely on predictive performance or isolated governance principles, RROS offers a layered framework that integrates data harmonization, model orchestration, workflow embedding, bias monitoring, privacy safeguards, and feedback mechanisms within a unified architectural construct. By conceptualizing readmission risk as a governed socio-technical pipeline, the framework advances interpretive clarity regarding how AI systems may be

responsibly integrated into hospital informatics environments.

## Theoretical Background and Literature Synthesis

### Data modalities in EHR-driven hospital readmission risk analytics

Electronic health records (EHRs) serve as foundational data modalities in AI-enabled hospital readmission risk governance. They provide both structured inputs—such as laboratory values, medication histories, diagnostic codes, and vital signs—and unstructured narratives, including clinical notes and discharge summaries. The literature emphasizes the importance of EHR integration in constructing systems-level architectures capable of processing longitudinal patient data for risk assessment [1, 5, 25].

In specialized hospital settings, such as cardiology wards, EHR-derived features, including heart rate variability, medication adherence patterns, and procedural histories, inform predictive modeling strategies. However, governance considerations arise due to variability in documentation practices, coding inconsistencies, and differential data completeness across departments. Without standardized data flows and harmonization protocols, EHR-driven analytics may generate fragmented or inconsistent risk insights [2, 6, 26].

Systems-level architectures must therefore formalize data normalization pipelines, interoperability standards, and monitoring checkpoints to ensure analytic reliability. Governance in this domain extends beyond data ingestion to encompass traceability, quality assurance, and alignment with clinical decision-support mechanisms.

### Imaging and waveform data integration for readmission risk governance in specialized units

Beyond traditional EHR data, imaging and waveform modalities increasingly contribute to AI-enabled architectures for hospital readmission risk governance, particularly in radiology departments and intensive care units (ICUs). Convolutional neural networks applied to imaging data—such as chest radiographs or computed

tomography scans—can detect latent physiological markers associated with elevated readmission probability [6, 7, 15].

Similarly, waveform analytics derived from electrocardiogram (ECG) signals or continuous vital sign monitoring provide real-time insights into patient stability during post-acute care transitions. These modalities enhance predictive granularity but introduce interoperability and governance complexities when merged with conventional clinical workflows. Differences in data formats, storage infrastructure, and real-time processing requirements complicate the development of unified analytics pipelines.

Systems-level synthesis highlights that without structured governance, multimodal integrations risk producing siloed analytics in high-stakes hospital environments. Governance architectures must address cross-modality harmonization, privacy safeguards during model training, and workflow-compatible visualization mechanisms to prevent fragmentation of risk insights [3, 9, 27]. By embedding these considerations into architectural design, AI-enabled readmission governance can leverage advanced data modalities while preserving institutional coherence and clinical safety.

Natural language processing (NLP) of clinical notes in AI-enabled readmission risk systems: NLP of clinical notes represents a key data modality in hospital readmission risk governance, extracting narrative insights from EHRs to refine AI-driven risk models. Studies synthesize NLP applications in oncology and primary care settings, where unstructured text reveals social determinants that influence readmission risk [4, 10, 11]. Systems-level architectures require governance to mitigate biases in NLP outputs, ensuring that extracted features align with privacy standards and informatics infrastructures. Literature underscores the interpretive value of NLP in workflow integration, yet calls for oversight to maintain decision-making reliability [7, 12, 18].

Genomics and wearables in public health-oriented readmission risk architectures: Genomics data modalities, combined with wearables, extend AI-enabled hospital readmission risk governance to public health and telemedicine contexts, enabling personalized risk profiling. Conceptual syntheses explore how genomic variants and continuous monitoring data inform predictive analytics, with governance emphasizing data security in distributed systems [13, 14, 16]. In primary care transitions, wearables

provide real-time metrics that feed into readmission models, but architectures must govern interoperability to avoid data drift. The background literature highlights trade-offs in scalability and advocates systems-level designs that balance innovation with clinical safety [17, 19, 20].

Claims data analytics for workflow integration in hospital readmission: Governance claims data, as administrative modalities, support AI-enabled architectures by linking financial and clinical outcomes in hospital readmission risk governance. Synthesis from informatics journals illustrates how claims-integrated pipelines enhance decision support in emergency departments (EDs), yet governance is critical to address privacy in cross-institutional data sharing [8, 21, 22]. Systems-level frameworks must incorporate audit mechanisms for claims-based risk models, ensuring alignment with hospital workflows and reducing liability in governance processes [23, 24].

Monitoring infrastructures in AI-driven readmission risk ecosystems: Monitoring infrastructures form the backbone of theoretical governance in AI-enabled hospital readmission risk systems, encompassing drift detection and performance oversight. Literature synthesizes approaches to real-time analytics in healthcare settings, where monitoring loops provide feedback for model recalibration without empirical testing [6, 25, 27]. In telemedicine extensions, these infrastructures govern interoperability across data modalities, offering conceptual insights into sustaining AI reliability in dynamic hospital environments [7, 27].

## Proposed conceptual framework

The readmission risk oversight scaffold (RROS) represents an original systems-level architecture designed to govern AI-enabled hospital readmission risk through integrated layers of data harmonization, model orchestration, workflow embedding, and governance feedback. RROS conceptualizes readmission risk as a cyclical process within healthcare informatics, starting from multifaceted data inputs (EHRs, imaging, NLP notes, genomics, wearables, claims) that feed into AI analytics pipelines. The architecture comprises four interconnected layers: (1) Data Harmonization Layer, which standardizes inputs via interoperability protocols to mitigate silos; (2) Model Orchestration Layer, where AI algorithms process risks with embedded bias checks; (3) Workflow Embedding Layer, integrating predictions into clinical decision support for seamless hospital actions; and (4) Governance Feedback

Layer, incorporating monitoring loops for drift, privacy audits, and clinician overrides. Feedback mechanisms ensure iterative refinement, such as automated alerts for data drift or governance reports on decision accountability, fostering a closed-loop system that balances technological efficacy with clinical safety.


Layer-specific governance functions and operational outputs are systematically summarized in Table 1.

**Table 1.** Functional layer responsibilities within the readmission risk oversight scaffold (RROS)

| Architectural layer       | Core functions                                                        | Governance mechanisms                                                      | Operation outputs                     |
|---------------------------|-----------------------------------------------------------------------|----------------------------------------------------------------------------|---------------------------------------|
| Data harmonization Layer  | Multimodal data ingestion, normalization, interoperability alignment  | Data provenance tracking, privacy filtering, schema validation             | Structured analytics-ready datasets   |
| Model orchestration layer | Predictive risk modeling, uncertainty scoring, bias detection         | Fairness audits, federated learning governance, and explainability modules | Probabilistic readmission risk score  |
| Workflow embedding layer  | Integration into clinical systems, discharge planning support, alerts | Human-AI collaboration protocols, override controls                        | Actionable decision support outputs   |
| Governance feedback layer | Monitoring, audit logging, compliance oversight, and drift detection  | Accountability tracing, regulatory reporting, and recalibration triggers   | Governance reports & model refinement |

Central to RROS is the data → model → decision → clinical action pipeline logic, where data ingestion informs probabilistic risk modeling, leading to governed decisions that trigger actions like targeted interventions or resource

allocations. This pipeline includes safeguards for privacy-preserving computations, such as federated learning proxies, and interoperability standards aligned with health informatics norms. The layered governance dynamics and closed-loop monitoring structure of the readmission risk oversight scaffold (RROS) are illustrated in Figure 1.



**Figure 1.** Systems-level governance architecture for AI-enabled hospital readmission risk: the readmission risk oversight scaffold (RROS).

The scaffold conceptualizes readmission risk governance as a vertically integrated socio-technical pipeline progressing from multimodal healthcare data ingestion to governed clinical action. The Data harmonization layer standardizes heterogeneous inputs, which are then processed by AI-driven risk modeling in the model orchestration layer. Predictions are operationalized via clinical decision infrastructures in the workflow embedding layer, while the governance feedback layer provides continuous oversight through privacy auditing, fairness monitoring, drift detection, and accountability logging. Horizontal feedback loops illustrate iterative recalibration mechanisms that ensure the safe, equitable, and interoperable deployment of predictive analytics within hospital environments.

To formalize key dynamics, the interaction between governance cost and risk mitigation in RROS can be

conceptualized as 
$$G(R) = \alpha \cdot D + \beta \cdot M + \gamma \cdot W - \delta \cdot F$$
, where  $G(R)$  denotes

governed readmission risk utility,  $D$  is data harmonization effort,  $M$  is model monitoring intensity,  $W$  is workflow integration workload,  $F$  is feedback loop efficiency, and

coefficients  $\alpha$ ,  $\beta$ ,  $\gamma$ ,  $\delta$  capture interpretive trade-offs in resource allocation. This expression captures the interaction between operational investments and governance gains, highlighting how enhanced feedback (higher  $\delta$ ) may offset costs in high-risk hospital scenarios.

Furthermore, uncertainty propagation in the pipeline may be expressed as  $U_p = \sum_{i=1}^n (u_i \cdot w_i)$ , where  $U_p$  is pipeline uncertainty,  $u_i$  represents uncertainty from each layer (data, model, etc.), and  $w_i$  are weights reflecting governance priorities—e.g., higher  $w$  for privacy-sensitive layers. This formula illustrates systems-level insights into how unmitigated uncertainties cascade, advocating for monitoring to minimize  $U_p$  in clinical deployments.

Finally, bias-fairness trade-offs in decision loops can be conceptualized as  $B_f = \left(\frac{B_m}{F_g}\right) * (1 - e^{-k \cdot A})$ , where  $B_f$  is a balanced fairness metric,  $B_m$  is model bias,  $F_g$  is the governance filtering strength,  $A$  is the audit frequency, and  $k$  is the decay constant; this captures the interpretive diminution of bias through sustained oversight.

Analytical implications

**Operational enhancements in AI-enabled hospital readmission risk workflows:** The readmission risk oversight scaffold (RROS) offers analytical opportunities to optimize hospital workflows through AI-enabled governance, with systems-level architectures that streamline readmission risk assessments from admission to discharge. By embedding interoperability-focused layers, RROS facilitates seamless data flows between EHR systems and clinical decision support, potentially reducing administrative burdens and enhancing resource allocation in high-volume hospital settings [1, 2, 5]. Interpretive insights suggest that such governance could harmonize predictive analytics with operational realities, allowing clinicians to prioritize interventions based on governed risk outputs rather than isolated predictions. This workflow integration implies a shift toward proactive risk management, where AI-driven insights inform staffing and bed management without overriding clinical judgment [3, 6, 9]. Key governance risks and their corresponding mitigation infrastructures across the RROS pipeline are synthesized in Table 2.

**Table 2.** Systems-level governance risk mitigation matrix for AI-enabled readmission analytics

| Governance domain | Risk vector                     | Monitoring mechanism                     | Mitigation strategy                    |
|-------------------|---------------------------------|------------------------------------------|----------------------------------------|
| Interoperability  | Data silos, schema misalignment | API trace logs, harmonization audits     | Standardized integration protocols     |
| Privacy           | Data leakage, re-identification | Differential privacy checks, access logs | Federated learning & encryption layers |
| Bias and fairness | Disparity amplification         | Subgroup performance dashboards          | Fairness recalibration pipelines       |
| Model reliability | Predictive drift                | Real-time performance monitoring         | Automated retraining triggers          |
| Clinical safety   | Over-automation                 | Clinician override tracking              | Human-AI review checkpoints            |
| Accountability    | Decision opacity                | Immutable audit trails                   | Governance reporting systems           |

Governance trade-offs in privacy-aware readmission risk architectures, privacy governance within RROS highlights analytical trade-offs in AI-enabled hospital readmission risk systems, balancing data utility with compliance imperatives. The architecture's feedback loops enable interpretive evaluations of privacy impacts on analytics pipelines, where differential privacy mechanisms may attenuate model sensitivity but preserve overall risk governance efficacy [12-14]. Systems-level implications include mitigated liability through audit-integrated decision loops that foster trust in interoperable infrastructures across hospital networks. This trade-off analysis underscores how governance layers can modulate privacy costs against clinical benefits, conceptualizing scenarios in which enhanced monitoring offsets potential data-exposure risks in readmission analytics [15-17].

Bias and fairness dynamics in clinical action loops for readmission risk: RROS provides systems-level insights into bias dynamics within AI-enabled hospital readmission risk governance, where monitoring scaffolds interpretation and addresses disparities in clinical action pathways. Analytical implications reveal how layered architectures can

detect bias propagation from data to decisions, promoting equitable risk stratification in diverse patient populations [18–20]. By incorporating fairness-aware feedback, the framework is expected to improve clinical outcomes through adjusted governance, such as recalibrating models to better reflect underrepresented cohorts in hospital datasets. This interpretive lens on bias-fairness interactions suggests operational safeguards that align AI deployments with ethical informatics standards, enhancing overall system resilience [21, 22].

Monitoring feedback for interoperability in readmission risk pipelines: Interoperability implications in RROS emphasize the analytical role of monitoring in sustaining AI-enabled readmission risk governance across fragmented healthcare systems. The architecture's pipeline logic interprets data-model-decision-action flows as interdependent, with feedback loops enabling conceptual validations of system coherence [7, 8, 23]. In hospital contexts, this implies scalable oversight for integrating modalities such as EHRs and wearables, thereby reducing drift in risk predictions over time. Systems-level trade-offs here involve governance costs versus interoperability gains, where enhanced monitoring could streamline clinical workflows in telemedicine extensions of readmission management [24–26].

Safety and accountability insights in AI-governed hospital readmission systems: safety-oriented implications of RROS; focus on accountability in AI-enabled hospital readmission risk architectures; interpreting human-AI collaborations as governed processes. The framework's override mechanisms imply analytical pathways for clinician intervention, ensuring safety in decision loops amid uncertain predictions [4, 10, 11]. This systems-level perspective highlights trade-offs between automation efficiency and liability, conceptualizing governance as a buffer against adverse events in clinical deployments. Furthermore, drift monitoring dynamics may be expressed

as 
$$D_t = \frac{(\sum_{j=1}^m d_j \cdot p_j)}{(1 + \lambda \cdot G_f)},$$

where  $D_t$  is temporal drift,  $d_j$  are data shifts per modality,  $p_j$  are pipeline priorities,  $G_f$  is governance frequency, and  $\lambda$  is a stabilization factor; this captures the interaction between evolving data environments and oversight intensity, illustrating how proactive governance minimizes safety risks in hospital settings [6, 7, 27].

## Results and Discussion

### Conceptual limitations in systems-level readmission risk governance architectures

While RROS advances interpretive frameworks for AI-enabled hospital readmission risk governance, its conceptual limitations stem from its theoretical grounding, potentially overlooking nuanced interactions within real-world informatics infrastructures. Systems-level architectures such as RROS assume structured interoperability, harmonized data exchange, and coordinated oversight layers. However, in practice, legacy EHR systems and vendor-specific constraints may limit data pipeline efficiency, model retraining cycles, and the capacity for workflow embedding, underscoring the need for adaptive governance interpretations rather than static architectural assumptions [1, 5, 12].

The abstraction inherent in systems-level design may obscure socio-technical frictions, including documentation variability, cross-departmental data silos, and inconsistent model monitoring infrastructures. These infrastructural asymmetries suggest that governance architectures must be interpreted as flexible scaffolds rather than prescriptive blueprints. In heterogeneous hospital environments, regulatory heterogeneity and institutional compliance cultures further complicate uniform deployment, reinforcing the importance of interpretive adaptability within governance frameworks.

Privacy trade-offs, although formally conceptualized within RROS, remain inherently context-dependent. Data minimization, consent interpretation, secondary use governance, and audit traceability differ across hospital regulatory landscapes and jurisdictions. Thus, privacy governance cannot be treated as a fixed module; instead, it requires reflexive alignment with institutional, national, and transnational policy environments [13, 14]. This discussion highlights the need for conceptual refinements that acknowledge variability without advancing empirical claims, maintaining RROS as a theoretical construct responsive to infrastructural realities.

### Future directions for AI-enabled workflow integration in hospital risk systems

Future conceptual explorations may extend RROS to incorporate emerging data modalities, including real-time genomics, remote patient monitoring streams, and telemedicine-integrated biomarkers. The integration of such data types into readmission governance architectures introduces interpretive complexity in harmonization pipelines, latency management, and cross-platform interoperability [15–17]. Expanding the scaffold to support multimodal analytics could enhance comprehensive oversight of readmission risk while preserving governance transparency.

Scaling governance architectures across public health networks presents another avenue for theoretical development. Workflow integration across distributed hospital systems implies collaborative AI deployments, shared audit protocols, and federated data governance models. Such scaling challenges conventional institution-centric oversight, instead suggesting systems-level insights into inter-organizational AI governance structures [7, 8, 23].

Moreover, dynamic feedback loops capable of recalibrating predictive thresholds in response to evolving epidemiological, operational, or policy conditions represent a critical interpretive direction. Adaptive oversight mechanisms could allow RROS to function not as a static architecture but as a responsive governance ecosystem. These future pathways emphasize balancing innovation with foundational clinical informatics principles, ensuring transparency, safety, and interpretability remain central to AI-enabled workflow integration.

## Interpretive challenges in bias monitoring for readmission risk frameworks

Bias monitoring within RROS presents interpretive challenges in AI-enabled hospital governance, as systems-level frameworks must navigate inherent data inequities without prescriptive solutions. Historical disparities in healthcare documentation, resource allocation, and access to care may become embedded within predictive pipelines, amplifying inequities through algorithmic outputs if not continuously scrutinized [18–20].

Conceptually, fairness governance can be strengthened by adding additional audit layers that evaluate both model outputs and downstream clinical actions. Rather than relying solely on statistical parity metrics, interpretive

safeguards may include longitudinal monitoring, contextual threshold adjustments, and interdisciplinary oversight committees. Such enhancements imply conceptual pathways for mitigating amplified disparities in clinical decision-making processes.

This challenge reflects broader informatics discourses on accountability and socio-technical systems governance, where fairness is not treated as a one-time compliance objective but as an evolving institutional responsibility [21, 22]. Within this perspective, RROS must be understood as an adaptable governance scaffold capable of integrating emerging fairness methodologies while maintaining theoretical coherence.

## Deployment considerations for privacy and safety in readmission governance

Deployment-oriented discussion of RROS emphasizes conceptual considerations for privacy and safety in AI-enabled readmission risk systems. Interoperability functions as a governance enabler, facilitating traceability, audit transparency, and coordinated oversight across clinical workflows. However, increased monitoring intensity may introduce operational burdens, particularly in resource-constrained hospitals, where computational infrastructure and workforce capacity remain limited [3, 6, 9].

Balancing oversight robustness with workflow feasibility becomes a central interpretive challenge. Excessive alerting mechanisms or monitoring layers may contribute to clinician fatigue, whereas insufficient oversight risks undermining safety accountability. Sustainable architectures, therefore, require calibrated governance intensity aligned with institutional capabilities.

This synthesis advocates for conceptual integrations that prioritize human–AI synergies, ensuring predictive systems augment rather than supplant clinical judgment. Embedding transparency, explainability, and shared decision-making principles within governance layers fosters resilient informatics ecosystems capable of adapting to infrastructural variability [24–26].

## Conclusion

The readmission risk oversight scaffold (RROS) introduces a novel systems-level architecture for governing AI-enabled hospital readmission risks, unifying data harmonization,

model orchestration, workflow embedding, bias monitoring, privacy safeguards, interoperability oversight, and feedback mechanisms within healthcare informatics. By conceptualizing readmission as a governed pipeline from data to clinical action, RROS offers interpretive insights into operational enhancements, privacy trade-offs, bias dynamics, interoperability monitoring, and safety accountability, without empirical claims.

Implications extend to improved resource management, strengthened clinician trust in AI deployments, and more structured accountability in hospital governance environments. By addressing conceptual gaps in current literature through layered oversight integration, the framework contributes to evolving discourses on responsible AI in healthcare.

Ultimately, RROS underscores the interpretive value of architectural governance designs in balancing technological advancement with clinical imperatives. It advances conceptual pathways for equitable, transparent, and safe AI integration in hospital systems while maintaining sensitivity to socio-technical complexities and regulatory diversity.

## Acknowledgements

None

## Conflict of interest

None

## Financial support

None

## Ethics statement

None

Received: 29 Jan 2022   Revised: 04 Mar 2022   Accepted: 03 Apr 2022  
Published online: 20 July 2022

## Rights and permissions

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

## References

- Johnson KW, Yahi A, Miotto R, Li L, Ricks D, Jebakaran J, et al. Predictive modeling of hospital readmission rates using electronic medical record-wide machine learning: a case-study using Mount Sinai heart failure cohort. *Pac Symp Biocomput.* 2017;22:276-87.  
[https://doi.org/10.1142/9789813207813\\_0027](https://doi.org/10.1142/9789813207813_0027).
- Marafino BJ, Schuler A, Liu VX, Escobar GJ, Baiocchi M. Predicting preventable hospital readmissions with causal machine learning. *Health Serv Res.* 2020;55(6):993-1002.  
<https://doi.org/10.1111/1475-6773.13586>.
- Kalagara S, Nakhla J, Chang V, Skovrlj B, Caridi JM, Hillbrand AS, et al. Machine learning modeling for predicting hospital readmission following lumbar laminectomy. *J Neurosurg Spine.* 2019;30(3):344-52.  
<https://doi.org/10.3171/2018.8.SPINE1869>.
- Alzeer AH, Alshaygi A. Using machine learning to reduce unnecessary rehospitalization of cardiovascular patients in Saudi Arabia. *Int J Med Inform.* 2021;154:104565.  
<https://doi.org/10.1016/j.ijmedinf.2021.104565>.
- Golas SB, Shibahara T, Agboola S, Otaki H, Sato J, Nakae T, et al. A machine learning model to predict the risk of 30-day readmissions in patients with heart failure: a retrospective analysis of electronic medical records data. *BMC Med Inform Decis Mak.* 2018;18(1):44.  
<https://doi.org/10.1186/s12911-018-0620-z>.
- Morgan DJ, Bame B, Zimberg P, Pimienta J, Scully K, Desai R, et al. Assessment of machine learning vs standard prediction rules for predicting hospital readmissions. *JAMA Netw Open.*

2019;2(3):e190348.

<https://doi.org/10.1001/jamanetworkopen.2019.0348>.

Theek B, Roeder I, Grimm L, Schmitz F, Schrader J, Steuw M, et al. Automation of data analysis in molecular cancer imaging and its potential impact on future clinical practice. *Methods*. 2021;188:30-6.

<https://doi.org/10.1016/j.ymeth.2020.06.019>.

Amritphale A, Gurm HS, Gul F, Amritphale N. Predictors of 30-day unplanned readmission after carotid artery stenting using artificial intelligence. *Adv Ther*. 2021;38(6):2954-72.

<https://doi.org/10.1007/s12325-021-01709-7>.

Rojas JC, Carey KA, Edelson DP, Venable LR, Howell MD, Churpek MM, et al. Predicting intensive care unit readmission with machine learning using electronic health record data. *Ann Am Thorac Soc*. 2018;15(7):846-53.

<https://doi.org/10.1513/AnnalsATS.201710-787OC>.

Hung M, Li W, He Y, Li C, Hon ES, Lipsky MS, et al. Prediction of 30-day hospital readmissions for all-cause dental conditions using machine learning. *BMC Oral Health*. 2020;20(1):307.

<https://doi.org/10.1186/s12903-020-01302-2>.

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.

Morley J, Machado CCV, Burr C, Cows J, Joshi I, Taddeo M, et al. The ethics of AI in health care: a mapping review. *Soc Sci Med*. 2020;260:113172.

<https://doi.org/10.1016/j.socscimed.2020.113172>.

Price WN II, Cohen IG. Privacy in the age of medical big data. *Nat Med*. 2019;25(1):37-43.

<https://doi.org/10.1038/s41591-018-0272-7>.

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>.

Currie G, Hawk KE, Rohren E, Vial A, Klein R. Machine learning and deep learning in medical imaging: intelligent imaging. *J Med Imaging Radiat Sci*. 2019;50(4):477-87.

<https://doi.org/10.1016/j.jmir.2019.09.005>.

Jia H. Yi Zeng: promoting good governance of artificial intelligence. *Natl Sci Rev*. 2020;7(12):1954-6.

Alruwaili FF. Artificial intelligence and multi agent based distributed ledger system for better privacy and security of

electronic healthcare records. *PeerJ Comput Sci*. 2020;6:e323.

<https://doi.org/10.7717/peerj-cs.323>.

Stewart C, Cristofaro S, Schneider P, Bauer C, Cortellesi P. Mapping ethico-legal principles for the use of artificial intelligence in gastroenterology. *J Gastroenterol Hepatol*. 2021;36(9):2413-9.

<https://doi.org/10.1111/jgh.15494>.

Mišić VV, Gabel E, Hofer I, Raji MA, Mahajan A. Machine learning prediction of postoperative emergency department hospital readmission. *Anesthesiology*. 2020;132(4):663-77.

<https://doi.org/10.1097/ALN.0000000000003140>.

Huang Y, Talwar A, Chatterjee S, Aparicio HJ. Application of machine learning in predicting hospital readmissions: a scoping review of the literature. *BMC Med Res Methodol*. 2021;21(1):96.

<https://doi.org/10.1186/s12874-021-01285-z>.

Li Q, Yao X, Eilerman BS, Xu Y. How good is machine learning in predicting all-cause 30-day hospital readmission? Evidence from administrative data. *Inquiry*. 2020;57:46958020962970.

<https://doi.org/10.1177/0046958020962970>.

Morel D, Yu F, Kalathil S, Ferrand E, Boucekine M, Falissard B, et al. Predicting hospital readmission in patients with mental or substance use disorders: a machine learning approach. *Int J Med Inform*. 2020;139:104136.

<https://doi.org/10.1016/j.ijmedinf.2020.104136>.

Artetxe A, Beristain A, Graña M. Predictive models for hospital readmission risk: a systematic review of methods. *Comput Methods Programs Biomed*. 2018;164:49-64.

<https://doi.org/10.1016/j.cmpb.2018.06.006>.

Romero-Brufau S, Huddleston JM, Escobar GJ, Liebow M. Implementation of artificial intelligence-based clinical decision support to reduce hospital readmissions at a regional hospital. *Appl Clin Inform*. 2020;11(4):570-7.

<https://doi.org/10.1055/s-0040-1715826>.

Baig MM, GholamHosseini H, Moqem AA, Mirza F, Lindén M. A machine learning model for predicting risk of hospital readmission within 30 days of discharge: validated with LACE index and patient at risk of hospital readmission (PARR) model. *Med Biol Eng Comput*. 2020;58(3):545-56.

<https://doi.org/10.1007/s11517-019-02105-4>.

Hung M, Li W, Hon ES, Su S, Zhou G, Lipsky MS, et al. Machine learning approach to predict risk of 90-day hospital readmissions in patients with atrial fibrillation: implications for quality improvement in healthcare. *Heart Mind*. 2020;4(3):59-65.

[https://doi.org/10.4103/hm.hm\\_30\\_20](https://doi.org/10.4103/hm.hm_30_20).

Gruendner J, Schwachhofer T, Sippl P, Gulden C, Kapsner LA, Zierk J, et al. Integrating genomics and clinical data for statistical analysis by using GENome MINIng (GEMINI) and fast healthcare interoperability resources (FHIR): system

design and implementation. *J Med Internet Res.* 2020;22(10):e19830.  
<https://doi.org/10.2196/19830>.