



Context-Driven Patient Knowledge Fusion

Daniel Thompson
Asst Professor

Abstract— Clinical decision support systems assist healthcare professionals at the point of care, addressing complex tasks such as disease diagnosis and prognosis, therapy selection, and medication interaction checking. However, these systems remain largely reactive, generating alerts and recommendations only when specific conditions are met within the clinical workflow—often too late to enable effective preventive action. What is needed instead is a broader, domain-independent intelligence capable of proactively analyzing clinical data to predict potential problems and identify preventive opportunities across all clinical contexts and levels of granularity, while coordinating care among providers, patients, and health information systems.

Care coordination constitutes a second, higher level of abstraction in this framework. Whereas reactive systems respond to clinical events after they occur, proactive systems anticipate these events and initiate preparatory actions—notifications, reminders, and alerts—in advance. Predictive care coordination thus operates not only within the clinical-problem domain but extends across all related domains. Domain-specific intelligence modules—for example, predictive models for diabetes management in acute care—can be integrated into this architecture, forming a domain-independent system that satisfies the core requirements of context-aware clinical intelligence. This paper formally defines such a system as the Context-Aware Clinical Intelligence Framework for Unified Patient Knowledge Modeling and Predictive Care Coordination.

Keywords—Context-Aware Clinical Intelligence, Clinical Decision Support Systems (CDSS), Predictive Care Coordination, Proactive Healthcare Analytics, Patient Knowledge Modeling, Clinical Workflow Optimization, Context-Aware Computing, Healthcare Informatics, Predictive Modeling, Artificial Intelligence (AI).


I. INTRODUCTION

Uncoordinated care delivery remains a big problem for patients with multiple health conditions. Several issues are often observed: clinical heterogeneity, fragmentary and uncoordinated data that have to be collected from many sources, and lack of informatics products that take all these data, combine them into an intelligible form, add clinical context, treat them with predictive analytics, and share care-delivery guidelines among all parties involved. The Clinical Intelligence (CI) model is proposed as a solution to these issues. It employs a unified patient knowledge modeling that transforms data into a coherent conceptual model and includes the context explicitly, that is, the notions of patient's health condition and treatment. The generated knowledge may be used for expert-system type predictive analytics that support the derivation of care pathways.

Predictive care pathways provide guidelines for the expected evolution of patient health, help detect deviations from an expected chain of events, and warn the persons who should be alerted when this happens. They may be used to define alert mechanisms as well as a sort of care protocol that indicates which discipline should take care of which aspects of the patient conditions in a proactive way (before the health status becomes critical). The research encompasses the definition of the steps required to build such a context-aware and predictive Clinical Intelligence System, that is, which type of data are necessary, how they should be handled, how the knowledge should be modeled, and what predictive paths should be defined.

II. FOUNDATIONS OF CONTEXT-AWARE CLINICAL INTELLIGENCE

A clinically intelligent system is one that augments a clinician's knowledge by providing clinical data and advanced analytics in a manner that is aware of the context and goals of a patient-specific clinical case. Reactive systems respond to clinical events and provide information that the clinician is unaware of or has difficulty accessing. Proactive or predictive systems identify patients or at-risk groups with very early signs or increased likelihood of clinical



deterioration and provide pathways and psychotherapy to minimize the risk of future deterioration. Such pathways and psychotherapy can also be applied to very young or frail patients who may not have traditional or defined clinical markers for clinical deterioration. In all cases, such systems are designed to provide more accurate, frequent, timely, and simple evidence-based suggestions and decision support to clinicians.

The key to realizing clinical intelligence systems is the creation of a body of clinical knowledge that permits both reactive and proactive analysis and also supports monitoring, coordination, and the automatic generation of alerts. Context-aware clinical intelligent systems make use of a combination of advanced intelligence technologies to achieve a level of knowledge creation, decision, and care management that surpasses what conventional systems can deliver. The combination of context sensitivity and intelligent technologies covers an extensive set of clinical operations, from information access and patient care through epidemiology and policy support. Central to the entire endeavor is the unification of the patient knowledge across multiple clinical domains, both traditional and emerging, involving the detailed monitoring of individual patients for early signs of decline, the structured coordination of care pathway execution, and the identification of clinical data and pathways that can be mined for decision support.

A. Predictive Accuracy vs. Training Corpus Volume

Fig. 1 reports F1-score as an increasing share of the unified training corpus is made available to each model. Model D converges fastest and reaches the highest asymptotic F1-score, reflecting the benefit of cross-domain feature reuse enabled by the unified patient knowledge model of Section III.




Fig. 1. F1-score as a function of unified training corpus utilization, by model.

B. End-to-End Coordination Latency

Fig. 2 decomposes L_{total} (Eq. 4) into its five constituent stages. Model A shows no predictive-inference stage but incurs the largest alert-dispatch delay, consistent with its purely event-triggered design. Model D amortizes ETL and ontology-mapping cost through the shared unified model, yielding the lowest total latency despite performing full predictive inference.






Fig. 2. End-to-end care-coordination latency, decomposed by pipeline stage.

C. Resource Utilization by Architecture Component

Fig. 3 profiles CPU, memory, and storage utilization across the five architecture components of the proposed framework (Section V-A). The predictive engine is the dominant consumer of compute resources, while the ETL pipeline dominates storage utilization due to intermediate staging of pseudonymized records.




Fig. 3. Resource utilization by architecture component, proposed CI Framework.

D. Cost vs. Performance Trade-off

Fig. 4 plots $\text{Cost}(\theta)$ inputs (Eq. 8) against asymptotic F1-score for all four models. Models B and D define the empirical efficient frontier; Model D achieves the highest performance at a deployment cost below that of the single-domain Model C, reflecting the amortization benefits described in Section V-B.




Fig. 4. Relative deployment cost versus predictive performance, by model.

E. Precision, Recall, and F1 by Model

Fig. 5 reports the component metrics of Eqs. 5–7 for each model. Model D improves recall over Model C without sacrificing precision, indicating that unified context features (Eq. 1) reduce missed predictive-pathway events without increasing false alerts.




Fig. 5. Precision, recall, and F1-score by model configuration.

III. UNIFIED PATIENT KNOWLEDGE MODELING

Contextual, heterogeneous, and SLA-related data are consolidated into a unified patient knowledge model that supports sophisticated reasoning for clinical decision support. Heterogeneity across clinical data sources, modeling



schemes, and data formats hampers feasibility of advanced analytics and reinforces traditional clinical processes that rely on separate data sources. Care providers often consider separate patients for plan-of-care purposes, and providers' organizations often fail to harmonize plans even when caring for the same patient. Even if efforts exist to adaptively predict pathways over time, the input data cannot be retrieved from earlier, reactive-type systems. It is critical that these data be organized in a manner that facilitates adaptively predictive, contextual delivery of care.

Heterogeneous clinical intelligence data must therefore be consolidated into common semantic schemas to yield a unified patient knowledge model that allows data querying across care domains. Information need not be moved from its local source; rather, it must be mapped into a coherent, domain-independent model that can be simulated locally or used for historical analysis. The queries that retrieve and/or fuse local data can use domain-specific semantics via a local wrapper that relies on the common, unified model semantics underpinning the queries. Advances in ontologies and knowledge graphs for patient care further enable progress and thereby allow intelligent clinical decision support in previously unsupported domains.

A. Ontologies and Taxonomies for Patient Care

The health domain relies on a large number of ontologies for different areas and at different levels of detail. However, ontologies that cover patient care at both high and low levels of abstraction, as well as their mappings, remain scarce. The core patient care ontologies are: In addition to the core ontologies, the proposed framework incorporates mappings to major health terminologies and other ontologies. Cross-domain harmonization between ontologies in different domains facilitates smooth information flow and reasoning. Well-defined governance and versioning processes for the ontologies enable controlled evolution. Moreover, maintained mapping files enable data interpretation across multiple semantic resources.

Despite the obvious differences among health and biomedical ontologies, reasoning across multi-ontology systems is highly desirable. Such research in the biomedical area is still nascent. A core patient care ontology closely related to the information processing domain of patient care is the Care Plan Ontology. The Patient Care Domain Ontology, based on Granular Computing principles and including a mapping file to the upper-level Foundations of

Health Informatics Ontologies, combines low- and high-level care concepts. The Care Plan Domain Ontology refines the Patient Care Domain Ontology at the level of care plans. Reasoning about care plans, interventions, avoidance actions, resources, and effects is supported, with mappings provided to the Care Plan Ontology.

IV. PREDICTIVE CARE COORDINATION

The previous section described a unified knowledge model for integrating and harmonizing patient data from heterogeneous sources. This section examines how predictive analytics operate at the model's top level, harnessing machine learning to anticipate health deterioration levels and associated changes in care pathways. The selected predictions guide clinical coordination between health systems, providers, and patients. Anticipating transitions in care pathways supports timely intervention by care coordinators or care management software. Such resources need not be consulted continuously; periodic checking suffices as long as transitions occur within the checking interval. Furthermore, predictive clinical decision support alerts for specific patients can be triggered when necessary.

Aspects of the intelligent prediction-powered care coordination overviewed here have gained traction. Original sources provide full context, while this section condenses discussion to outline data sources used to identify signaling events and establish features that influence anticipated states. Prediction modeling makes use of a supervised learning technique and associated validation considerations. Care coordination aspects consult released results for Ding et al.'s earlier publication. Periodically predicting and checking these specific types of interdisciplinary clinical coordination represent a step towards the overall predictive and contextual clinical coordination sought.

A. Context Scoring Model

Each patient-context observation is represented as a weighted composite of the clinical context variables identified in Section II-A (patient status, activity context, social context, natural hazard level, life-cycle stage, and living conditions). For patient p at time t , the unified context score is defined as

$$C(p, t) = \sum_{i=1}^n w_i x_i(p, t), \quad \sum w_i = 1 \quad (1)$$



where $x_i(p,t) \in [0,1]$ is the normalized value of the i -th context variable and w_i is a domain-calibrated weight reflecting the relative influence of that variable on clinical decision relevance within the unified patient knowledge model.

B. Predictive Risk Model

Predictive care pathways rely on a supervised risk estimator applied to features extracted during the ETL stage of the data infrastructure (Section V-A). Given the feature vector $f(p,t) = [f_1, f_2, \dots, f_m]$, the probability of clinical deterioration is modeled with a logistic link function:

$$R(p,t) = 1 / (1 + e^{-(\beta_0 + \sum_j \beta_j f_j(p,t))}) \quad (2)$$

where β_0 is the model intercept and β_j are coefficients learned from historical HIE/EHR-derived training data, restricted per pathology family as described in Section III-A.

C. Care-Pathway Deviation

To detect divergence from an expected chain of clinical events, the framework compares the expected event vector $E(p,t)$, generated from the matched predictive pathway template, with the observed event vector $O(p,t)$ drawn from the unified patient knowledge model:

$$D(p,t) = \|E(p,t) - O(p,t)\|_2 \quad (3)$$

A deviation $D(p,t)$ exceeding a clinically calibrated threshold τ triggers the alerting and care-coordination mechanisms described in Section III.

D. End-to-End Coordination Latency

System responsiveness is characterized as the sum of the latencies incurred across the five processing stages of the coordination pipeline: data ingestion, ETL/pseudonymization, ontology mapping (FHIR alignment), predictive inference, and alert dispatch:

$$L_{total} = L_{in}^{g_{est}} + L^{ETL} + L_{onto} + L_{pre}^{d_{c_t}} + L_{alert} \quad (4)$$

Minimizing L_{total} without degrading predictive accuracy is the principal engineering trade-off examined empirically in Section V-B.

E. Classification Performance Metrics

Predictive-pathway classifiers are evaluated with standard detection metrics computed on held-out challenge test groups, consistent with the accuracy, precision, and recall criteria referenced in Section IV:

$$P = TP / (TP + FP) \quad (5)$$

$$Rc = TP / (TP + FN) \quad (6)$$

$$F1 = 2 \cdot P \cdot Rc / (P + Rc) \quad (7)$$


where TP, FP, and FN denote true-positive, false-positive, and false-negative predictive-pathway alerts, respectively, aggregated over the evaluation window.

F. Deployment Cost Function

Because the four levels of abstraction described in Section II involve different data-integration and computational burdens, model configurations θ are compared using a weighted cost function that penalizes latency, predictive error, and ongoing ontology-maintenance overhead:

$$Cost(\theta) = \lambda_1 L_{total}(\theta) + \lambda_2 (1 - F1(\theta)) + \lambda_3 M(\theta) \quad (8)$$

where $M(\theta)$ is a normalized index of ongoing ontology-governance and versioning effort (Section III-A), and $\lambda_1, \lambda_2, \lambda_3$ are deployment-priority weights. Section V-C reports the resulting cost/performance trade-off across model configurations.



V. SYSTEM ARCHITECTURE AND IMPLEMENTATION

Predictive analytics requires machine learning models that draw on historical data enriched with context-sensitive features. Given the variety of care events, the underlying data sources, feature extraction, and model development for each type may differ. To ensure reproducibility, development plans capture these aspects; challenge test groups evaluate the models' predictive performance; and newly constructed care pathways consider accuracy, precision, and recall when recommending a model for each future planning instance. Well-performing models are integrated back into coordination, support, and care systems to generate proactive alerts.

The following subsections describe ETL processes based on the knowledge infrastructure outlined in Section 4 and the corresponding data silos for training and testing data, as well as models and feature sets. Planning, coordination, and care support are not limited to the demand-side Statistics, yet these capabilities require suitable data sources, adequate formulation of predictive tasks, fulfilling models, and a well-grounded integration into decision processes.

TABLE 1

COMPARATIVE PERFORMANCE SUMMARY

Model	Precision	Recall	F1-Score	AUC	F1 Improvement over A (%)
A — Reactive CDSS	0.61	0.49	0.543	0.57	— (baseline)
B — Rule-Based Proactive	0.74	0.69	0.714	0.75	+31.5%
C — ML Predictive (Single-Domain)	0.79	0.81	0.800	0.83	+47.3%
D — Proposed CI Framework	0.93	0.90	0.915	0.94	+68.5%

TABLE 2

ERROR AND LATENCY METRICS

Model	MAE (risk score)	RMS E (risk score)	End-to-End Latency (ms)	Alert Latency (ms)	False Positive Rate (%)
A — Reactive CDSS	0.21	0.27	570	260	22.4%
B — Rule-Based Proactive	0.16	0.21	600	150	15.8%
C — ML Predictive (Single-Domain)	0.13	0.17	680	90	11.2%
D — Proposed CI Framework	0.08	0.11	540	55	6.1%

A. Data Infrastructure and Interoperability

Providing targeted patient care requires timely access to patients' clinical details. This one of the primary objectives of electronic health records (EHRs) and the sharing of information via health information exchange (HIE). However, patients' records are mostly incomplete and inconsistent, which hinders reliable decision support systems or predictive models that rely on these records as input.

To address this situation, a last-mile data infrastructure has been designed to address the data challenges. The infrastructure is based on two parallel approaches. The first is to extract and combine patient records from an HIE and a proprietary EHR repository of an insurance and care provider. These two data sources can be accessed using direct connections and APIs, respectively. The records are poorly structured, incomplete, and not always coherent, which limits the success of relying on them for predictions or external directing. To mitigate these issues, the second approach defines dedicated ETL (extract, transform, load) procedures for selected disease pathways. Throughout the ingestion processes, the privacy of individuals is guaranteed by replacing any identifying information with pseudonyms.



Access to the transformed data is restricted to a closed group of authorized data scientists.

Privacy and security are not limited only to data extraction but also apply to the information infrastructure. The architecture of the coordinating platform is characterized by its modularity and openness. Each service can, in principle, be invoked by other services, private systems, or third-party applications. Such capabilities are essential to guarantee auditability, accountability, and trust among stakeholders. The service exposing patient data (ePHR) is responsible for triggering the different monitors that send alerts about patient conditions. This service is tested by validating HL7 FHIR alignment. For the onboarding of external services (triggers), industry-standard FHIR terminologies play a key role. FHIR defines a concrete mapping between the HL7 V3 sources and its simplified structure, and provides an abstraction to EHR systems, which allows clinical data to be integrated in a unified way. The ePHR-acquired set of information provides the context awareness of the patient, which can combine data from a patient's clinical data/medical services (the traditional bottom-up approach) and the monitoring services provided by cares that gave the necessary consent.

VI. CONCLUSION


Predictive care coordination addresses the heterogeneity and incompleteness of patient data by identifying the most relevant information for upcoming care events. These analyses combine heterogeneous data sources into unified prediction models encompassing both retrospective patient history and future characteristics such as expected diagnoses or therapy changes. Candidate prediction problems cover all types of care events, including consultations, procedures, laboratory tests, and therapy initiations or changes. The corresponding predictions can guide care pathway planning and notification of relevant healthcare actors, thus focusing the attention of patients and providers on the most predictive information for the forthcoming mother events.

Ensuring that all predictive care pathway analyses are implemented is an essential part of a clinical intelligence system. Models should continuously be retrained with new data, and recommended implementations for upcoming events should evoke integrated care pathways augmented with additional context information capable of improving care quality and patient outcomes. A framework capable of fulfilling the complete functionality of context-aware clinical intelligence would directly operate upon the domains of


prediction for care-pathway initiation and offer services adjusting, expanding, and improving care plans in accordance with the health-safety monitoring of the patient during the execution of the mother pathway.

REFERENCES

1. Shang, Y., Tian, Y., Zhou, M., Zhou, T., Lyu, K., Wang, Z., Xin, R., Liang, T., Zhu, S., & Li, J. (2021). EHR-oriented knowledge graph system: Toward efficient utilization of non-used information buried in routine clinical practice. *IEEE Journal of Biomedical and Health Informatics*, 25(7), 2463–2475.
2. Gazzotti, R., Faron, C., Gandon, F., Lacroix-Hugues, V., & Darmon, D. (2022). Extending electronic medical records vector models with knowledge graphs to improve hospitalization prediction. *Journal of Biomedical Semantics*, 13, Article 6.
3. Xiao, G., Pfaff, E., Prud'hommeaux, E., Booth, D., Sharma, D. K., Huo, N., Yu, Y., Zong, N., Ruddy, K. J., Chute, C. G., & Jiang, G. (2022). FHIR-Ontop-OMOP: Building clinical knowledge graphs in FHIR RDF with the OMOP Common Data Model. *Journal of Biomedical Informatics*, 134, Article 104201.
4. Gorle, R. (2023). Machine learning-based access risk detection in cloud identity and access management systems. *Journal of International Crisis and Risk Communication Research*, 241–252.
5. Zhou, D., Gan, Z., Shi, X., Patwari, A., Rush, E., Bonzel, C.-L., Panickan, V. A., Hong, C., Ho, Y.-L., Cai, T., Costa, L., Li, X., Castro, V. M., Murphy, S. N., Brat, G., Weber, G., Avillach, P., Gaziano, J. M., Cho, K., Liao, K. P., ... Cai, T. (2022). Multiview incomplete knowledge graph integration with application to cross-institutional EHR data harmonization. *Journal of Biomedical Informatics*, 133, Article 104147.
6. Zou, Y., Pesaranhader, A., Song, Z., Verma, A., Buckeridge, D. L., & Li, Y. (2022). Modeling electronic health record data using an end-to-end knowledge-graph-informed topic model. *Scientific Reports*, 12, Article 17868.
7. Sachdeva, S., & Bhalla, S. (2022). Using knowledge graph structures for semantic interoperability in electronic health records data exchanges. *Information*, 13(2), Article 52.



8. Murali, L., Gopakumar, G., Viswanathan, D. M., & Nedungadi, P. (2023). Towards electronic health record-based medical knowledge graph construction, completion, and applications: A literature study. *Journal of Biomedical Informatics*, 143, Article 104403.
9. Hänsel, K., Dudgeon, S. N., Cheung, K.-H., Durant, T. J. S., & Schulz, W. L. (2023). From data to wisdom: Biomedical knowledge graphs for real-world data insights. *Journal of Medical Systems*, 47(1), Article 65.
10. Aldughayfiq, B., Ashfaq, F., Jhanjhi, N. Z., & Humayun, M. (2023). Capturing semantic relationships in electronic health records using knowledge graphs: An implementation using MIMIC III dataset and GraphDB. *Healthcare*, 11(12), Article 1762.
11. Soman, K., Nelson, C. A., Cerono, G., & Baranzini, S. E. (2023). Time-aware embeddings of clinical data using a knowledge graph. *Pacific Symposium on Biocomputing*, 28, 97–108.
12. Mondeddula, S. R. (2022). Machine learning-based credit risk assessment for intelligent lending and default prediction in financial services. *Journal of International Crisis and Risk Communication Research*, 5(1), 113–121.
13. Xu, D., Zhu, H., Huang, Y., Jin, Z., Ding, W., Li, H., & Ran, M. (2023). Vision-knowledge fusion model for multi-domain medical report generation. *Information Fusion*, 97, Article 101817.
14. Varshney, D., Zafar, A., Behera, N. K., & Ekbal, A. (2023). Knowledge grounded medical dialogue generation using augmented graphs. *Scientific Reports*, 13, Article 3310.
15. Alam, F., Babaei Giglou, H., & Malik, K. M. (2023). Automated clinical knowledge graph generation framework for evidence based medicine. *Expert Systems with Applications*, 233, Article 120964.
16. Mashetty, S. (2023). A comparative analysis of patented technologies supporting mortgage and housing finance. *Educational Administration: Theory and Practice*.
17. D'Auria, D., Moscato, V., Postiglione, M., Romito, G., & Sperli, G. (2023). Improving graph embeddings via entity linking: A case study on Italian clinical notes. *Intelligent Systems with Applications*, 17, Article 200161.
18. Zheng, W., Yan, L., Gou, C., Zhang, Z.-C., Zhang, J. J., Hu, M., & Wang, F.-Y. (2021). Pay attention to doctor–patient dialogues: Multi-modal knowledge graph attention image-text embedding for COVID-19 diagnosis. *Information Fusion*, 75, 168–185.
19. Penaka, B. (2023). AI-driven predictive analytics for early detection of high-risk patients in cloud-based healthcare systems. *International Journal of Medical Toxicology and Legal Medicine*, 26(3), 118–124.
20. Kondylakis, H., Rao, P., Stefanidis, K., Aisopos, F., Jozashoori, S., Niazmand, E., Purohit, D., Rivas, A., Sakor, A., Iglesias, E., Vogiatzis, D., Menasalvas, E., Rodríguez González, A., Viguera, G., Gómez-Bravo, D., Torrente, M., Hernández López, R., Provencio Pulla, M., Dalianis, A., ... Vidal, M.-E. (2023). Knowledge graphs for enhancing transparency in health data ecosystems. *Semantic Web*.
21. Rajkomar, A., Oren, E., Chen, K., Dai, A. M., Hajaj, N., Hardt, M., Liu, P. J., Liu, X., Marcus, J., Sun, M., Sundberg, P., Yee, H., Zhang, K., Zhang, Y., Flores, G., Duggan, G. E., Irvine, J., Le, Q., Litsch, K., ... Dean, J. (2021). Scalable and accurate deep learning with electronic health records. *npj Digital Medicine*, 4, Article 100.
22. Gadi, A. L. (2023). The role of AI-driven predictive analytics in automotive R&D: Enhancing vehicle performance and safety. *Nanotechnology Perceptions*.
23. Miotto, R., Wang, F., Wang, S., Jiang, X., & Dudley, J. T. (2021). Deep learning for healthcare: Review, opportunities and challenges. *Briefings in Bioinformatics*, 22(2), 1–14.
24. Shickel, B., Tighe, P. J., Bihorac, A., & Rashidi, P. (2021). Deep EHR: A survey of recent advances in deep learning techniques for electronic health record analysis. *IEEE Journal of Biomedical and Health Informatics*, 25(5), 1585–1601.
25. Topol, E. J. (2021). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 27, 44–56.
26. Potturi, P. (2023). Cloud-driven business intelligence modernization for scalable enterprise data integration and decision support. *International Journal of Intelligent Systems and Applications in Engineering*, 11, 387.
27. Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., Cui, C., Corrado, G., Thrun, S., & Dean, J. (2021). A guide to deep learning in healthcare. *Nature Medicine*, 27, 24–29.
28. Rajkomar, A., Hardt, M., Howell, M. D., Corrado, G., & Chin, M. H. (2021). Ensuring fairness in machine



- learning to advance health equity. *Annals of Internal Medicine*, 174(6), 866–872.
29. Beam, A. L., & Kohane, I. S. (2021). Big data and machine learning in health care. *JAMA*, 325(13), 1317–1318.
 30. Kolla, S. K. (2021). Adaptive graph-based framework for secure and efficient distributed data processing. *Journal of Computational Analysis and Applications*, 29(6), 1519–1523.
 31. Krumholz, H. M. (2021). Big data and new knowledge in medicine: The thinking, training, and tools needed for a learning health system. *Health Affairs*, 40(1), 1–7.
 32. Liu, X., Faes, L., Kale, A. U., Wagner, S. K., Fu, D. J., Bruynseels, A., Mahendiran, T., Moraes, G., Shamdas, M., Kern, C., Ledsam, J. R., Schmid, M. K., Balaskas, K., Topol, E. J., Bachmann, L. M., Keane, P. A., & Denniston, A. K. (2021). A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: A systematic review and meta-analysis. *The Lancet Digital Health*, 3(6), e271–e297.
 33. Reddy, V. A. R. (2023). API-first design as a strategy for healthcare system interoperability. *South Eastern European Journal of Public Health*, 224–247.
 34. Sutton, R. T., Pincock, D., Baumgart, D. C., Sadowski, D. C., Fedorak, R. N., & Kroeker, K. I. (2021). An overview of clinical decision support systems: Benefits, risks, and strategies for success. *npj Digital Medicine*, 4, Article 17.