Interoperability in Healthcare: A Primer and New Frontiers

Andra Ionescu
KTH Royal Institute of Technology
Stockholm, Sweden
adio@kth.se

Paris Carbone
KTH Royal Institute of Technology
Stockholm, Sweden
parisc@kth.se

Sebastiaan Meijer
KTH Royal Institute of Technology
Stockholm, Sweden
smeijer@kth.se

Christos Koutras
New York University
New York City, USA
christos.koutras@nyu.edu

Juliana Freire
New York University
New York City, USA
juliana.freire@nyu.edu

Katja Hose
TU Wien
Vienna, Austria
katja.hose@tuwien.ac.at

Jayanth Raghothama
KTH Royal Institute of Technology
Stockholm, Sweden
jayanthr@kth.se

ABSTRACT

Interoperability in healthcare remains difficult despite decades of standards and infrastructure. Although messaging protocols and standards enable systems to exchange data, they do not guarantee that clinical meaning is preserved across organisations, workflows, and cultural contexts. This persistent gap between where semantics are used in care and how they are exchanged continues to drive fragmentation, bespoke mappings, and costly manual integration.

This tutorial offers a primer and a forward-looking agenda by framing interoperability as three generations. We first revisit the pre-interoperability era of proprietary systems and point-to-point integrations. We then cover the standards-based second generation, highlighting what it solved (e.g., connectivity) and what it left unresolved (e.g., semantic alignment). Finally, we outline an emerging third generation towards true interoperability, and conclude with research directions for building interoperability in the AI era.

PVLDB Reference Format:
Andra Ionescu, Paris Carbone, Sebastiaan Meijer, Christos Koutras, Juliana Freire, Katja Hose, and Jayanth Raghothama. Interoperability in Healthcare: A Primer and New Frontiers. PVLDB, 19(12): 4912 - 4917, 2026.
doi:10.14778/3827998.3828149

1 INTRODUCTION

Interoperability in healthcare is a veritable Pandora's box of problems, ranging from legacy systems to inadequately defined and described information, data silos, poor adherence to common data models and standards, etc. Beyond implementation issues, there are fundamental underlying issues, such as a lack of agreement on the semantics of terms, differences in understanding and interpreting standards, and a genuine diversity and ambiguity in how medical and clinical terms are used and applied in care settings [3]. While poor implementation has often been blamed, a deeper issue is the persistent mismatch between *where* semantics are defined and used (i.e., within clinical workflows, organisational conventions and cultural contexts) and *how* semantics are exchanged (i.e., via rigid, standardised structures).

Data exchange in healthcare has benefited from the introduction of standards, which brought "*secure telephone lines between providers*". However, standards flatten diversity, ignoring important contextual nuance that clinicians rely upon. Despite the standards, Electronic Health Record (EHR) and healthtech vendors still defined their own information models to describe the same concepts, making a stable connection insufficient to solve the bigger problem: more standards brought more fragmentation and variability [2].

The past decades of interoperability work have been shaped largely by messaging protocols (e.g., HL7 [20]), system APIs (e.g., FHIR [6]), content standards (e.g., openEHR [38]), and terminologies (e.g., ICD [31], SNOMED-CT [15], LOINC [25]), which provided the essential "*plumbing*" for exchanging health information across systems. However, these technical connections are insufficient to ensure that the receiving side interprets the data as intended in its local clinical context. As a result, despite decades of progress with standards definition and technological improvement, we still face the same central challenge: *true* semantic interoperability, which requires that the precise meaning of data is understood identically by all parties [16, 52].

Recent advances in AI-driven methods and knowledge graphs offer new opportunities to capture and align meaning across heterogeneous data sources, motivating a re-examination of interoperability beyond standardisation. In this tutorial, we examine the technical foundations of semantic interoperability in healthcare, the influence of data systems on this problem, and guide researchers towards emerging unique research directions made possible by the latest developments in AI, such as large language models (LLMs).

Tutorial Structure. The tutorial will last 90 minutes and is structured into three parts lasting 30 minutes each. In the first part, we introduce the motivation and background (Section 2), and discuss the first generation of interoperability (Section 3). In the second

Proceedings of the VLDB Endowment, Vol. 19, No. 12 ISSN 2150-8097.
doi:10.14778/3827998.3828149

part, we examine the second generation of standards-based interoperability and explain why semantic alignment remains difficult in practice (Section 4). In the third part, we discuss emerging third-generation approaches towards full interoperability, and conclude with a forward-looking agenda (Section 5). The tutorial material will be available on the tutorial website before the conference.

Target Audience. The target audience for our tutorial are researchers and practitioners with an interest in the applications of data management systems in healthcare, aiming to solve interoperability. The tutorial is self-contained and assumes no prior knowledge of healthcare informatics or interoperability.

Prior Tutorials. Interoperability in healthcare has been presented in the form of a tutorial at VLDB 2012 [19]. The focus then was on standards, the adoption of EHR across institutions, and the pipeline of clinical studies. With this tutorial, we aim to bridge the database systems and health informatics perspectives by looking at common developments towards interoperability, such as data harmonisation, semantic alignment, as well as the emerging approaches around LLMs, knowledge graphs, conformal predictions, multimodality, and agentic AI.

2 BACKGROUND

Clinical Data. EHRs capture patient information across clinical settings (e.g., diagnoses, procedures, laboratory results, medications, etc). Beyond structured EHR data, clinical information includes unstructured narratives (e.g., clinical notes), medical imaging, physiological signals, patient-generated and genomic data, introducing heterogeneity in both format and semantics. Moreover, clinical meaning is context-dependent, relying on time, setting, and intent of recording, and similar concepts may be encoded using different schemata, coding systems, and levels of detail across institutions. As such, harmonising clinical data requires aligning structure, values, and preserving the context in which data is generated.

Interoperability. Interoperability can be broadly defined as the ability of two or more systems to exchange information *and* to use the information that has been exchanged without extensive human intervention or processing [47]. The Healthcare Information and Management Systems Society (HIMSS) defines four levels of interoperability[1] [35].

Technical interoperability – represents exchanging data between systems. In practice, this is the "plumbing", where standardised protocols allow data transfer, regardless of whether the receiving side can interpret the payload.

Syntactic interoperability – concerns the format and structure in which data is encoded and exchanged (e.g., XML, JSON, etc). This layer is shaped by standards bodies and implementation guides.

Semantic interoperability – is the ability to share meaning across systems, a common language that remains understandable to both humans and machines. Many efforts have concentrated on exchange mechanisms and messaging protocols, assuming that meaning would "*follow*" the message. However, the problem must be addressed at the semantic layer(s) by distinguishing between ontic and epistemic dimensions of health data, clarifying both the concept model (i.e., terminologies) and the information model.

[1]Throughout the tutorial, we will focus on the challenges and opportunities of *semantic interoperability*.

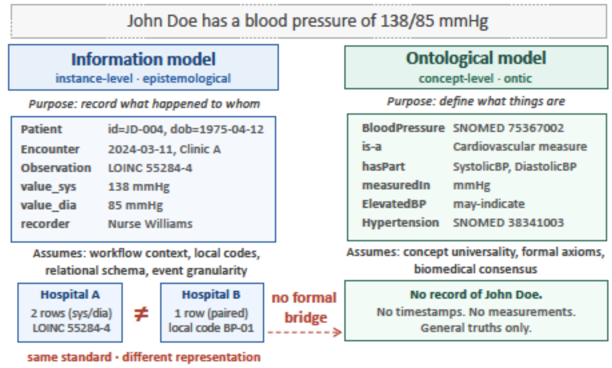

Figure 1: The difference between the information model and the ontological model.

Example. Take a canonical example, illustrated in Figure 1: the information model that captures blood pressure measurements for John is not the same as the ontological description of blood pressure. In the information model, the focus is on information/content, reflecting an epistemological view: knowledge gathered about individuals (e.g., encounters, events) in concrete settings. In the ontological model, the focus is on terminology/concepts, reflecting an ontic view: formal descriptions of general truths and relations (e.g., diseases and their biomedical relationships).

Organisational interoperability – is achieved when exchanging data across health IT systems aligns workflows, governance, and incentives, not only technical interfaces. Progress depends on policy and regulation that incentivise and sometimes enforce interoperable exchange. Achieving organisational interoperability remains the future goal and an open challenge until semantic interoperability is solved.

3 FIRST GENERATION: PRE-INTEROPERABILITY

The first generation of systems (≈1960 – 1995), what we refer to as the pre-interoperability era (Figure 2), aimed at securing proprietary health record management systems. The focus of clinical information systems during this time was primarily on digitising records rather than enabling data exchange across settings [17]. Therefore, standards surfaced to cover the emerging need for hospital databases and informatics systems to "talk to each other".

Challenges. The pre-2000 health information landscape lacked all three foundations necessary for interoperability: a shared information model, a shared terminology, and a shared exchange mechanism [23]. The fragmentation of the period created a long-lasting technical debt: once hospitals built bespoke EHR infrastructures, retrofitting interoperability onto them became increasingly difficult [37, 62]. These limitations shaped the agenda of the second generation, where the community shifted from digitisation toward systems interoperability, emphasising messaging and vocabulary standards, and syntactic compatibility.

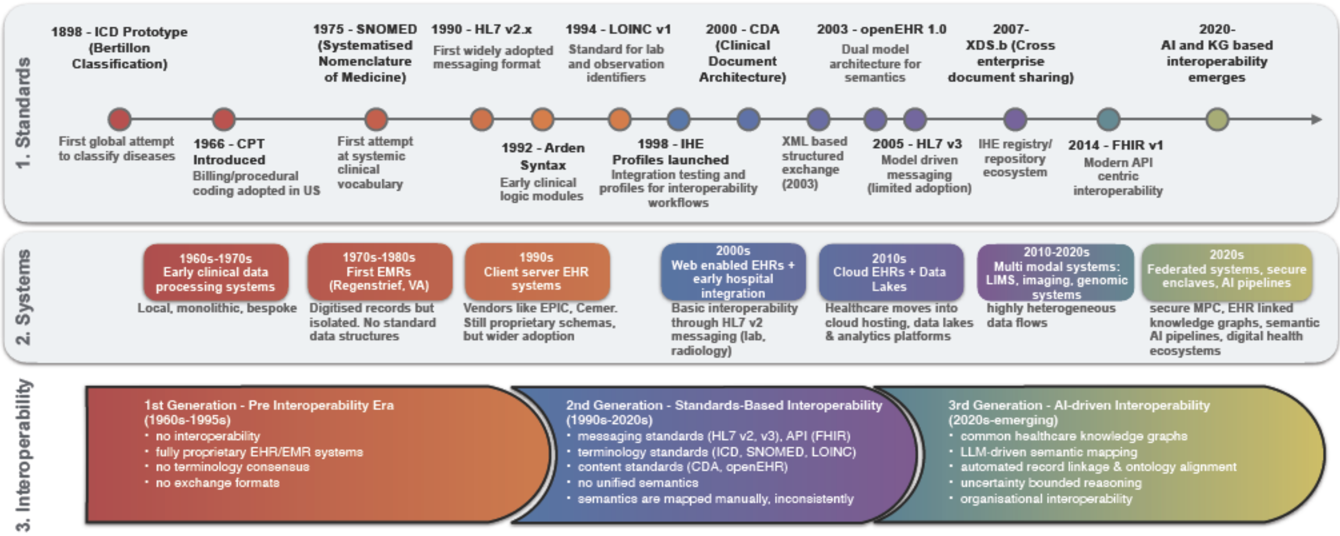

Figure 2: A Timeline of Health Data Interoperability

4 SECOND GENERATION: STANDARDS-BASED INTEROPERABILITY

The standards-based second generation of healthcare interoperability (≈1990–2020) developed bottom-up, layer by layer. Starting with exchange messaging, HL7 v2 gave hospitals a shared communication channel for structured clinical events. Terminology systems (e.g., ICD, SNOMED-CT) followed as an independent effort, providing controlled vocabularies for coding clinical concepts. Information modelling was developed next, with HL7 v3's Reference Information Model (RIM) and OpenEHR gaining traction across health systems. Finally, the API-based exchange layer FHIR enabled cross-organisational data flows [2].

Although the second generation is perceived as a success story, its defining limitation is structural. Because each layer was standardised independently, the information model was never explicitly defined, but emerged implicitly from the intersection of conventions developed in parallel. For example, hospitals adopting the same standards still produce semantically incompatible datasets [16, 52], mirroring a well-known problem in data management: structural alignment does not imply semantic alignment [28, 34]. Next, we present three core challenges, common across data management and health informatics communities, that delay interoperability.

Data Discovery Challenges. Mappings between schemata, terminologies, and data elements are incomplete, institution-specific, and poorly maintained, mirroring data lake challenges where finding related tables requires dataset discovery, join discovery, and schema matching [10, 11, 24, 36, 43, 55]. In classical data integration, a mediated schema provides a stable global view of heterogeneous sources [18], while in healthcare, such schemata are rarely available.

Data Harmonisation Challenges. Even when institutions adopt the same standards, clinical concepts are represented differently across systems: one hospital stores blood pressure as two separate observations, another as a paired record under a local code as also shown in Figure 1. This mirrors the core challenge of entity resolution, as structurally compatible records referring to the same real-world entity cannot be trivially unified [12, 13]. Standard approaches (e.g., blocking, similarity matching, active learning) [27, 40] face additional constraints in healthcare, as privacy limits available signals, clinical values are context-sensitive, and errors carry patient safety consequences.

Semantic Alignment Challenges. The healthcare ecosystem contains a large number of overlapping standards and terminologies grounded in different conceptual assumptions [1, 57], introducing multiple representations of the same clinical concept and persistent ambiguity across systems. The dataspace vision [29] reframed this problem: rather than demanding full semantic alignment upfront, integration can tighten incrementally on a pay-as-you-go basis. LLM-based ontology matching and entity linking [21, 59, 60] offer a promising avenue for reducing manual effort, though clinical terminology remains challenging due to its specificity and the high cost of misalignment.

5 TOWARDS A THIRD GENERATION: NEW FRONTIERS

In the second generation, the information model came last, and every downstream integration had to compensate for the incompatibilities. In the third generation, a top-down structure is proposed. Rather than standardising exchange and hoping semantic agreement follows, it begins by defining a shared information model and derives storage, terminology, and exchange from it. Initiatives such as the OMOP Common Data Model [51] represent early implementations of the top-down approach. However, interoperability is still driven by manual ETL pipelines. Thus, we propose several opportunities for advancing the next generation of interoperability.

Closing the Semantic Gap using LLMs. A fundamental challenge in semantic interoperability is bridging the gap between informal clinical language and formal machine-readable representations [64].

In the third generation, semantic alignment becomes a first-class pipeline component rather than a preprocessing step. LLMs act as flexible semantic matchers that interpret schema elements, attribute descriptions, and clinical context to generate correspondence candidates [8, 26, 48, 53, 54, 58]. Critically, LLM-based mapping must not operate in isolation, but can rely on clinical ontologies to constrain the space of valid interpretations, suppress hallucinations, and bound the output [5, 44].

Interactive Tools for Data Harmonisation. While LLM-based mapping methods can automate parts of the semantic alignment pipeline, healthcare practitioners ultimately need to inspect, validate, and refine their outputs before they can be trusted in clinical settings. Therefore, modern healthcare interoperability systems should provide interactive, human-in-the-loop tools that allow practitioners to curate AI-driven mapping results. Towards this direction, visualisation environments that let experts explore and correct LLM-suggested schema correspondences have begun to emerge in the biomedical domain [63], as have end-to-end harmonisation toolkits [49]. In addition, conversational interfaces powered by LLM agents can lower the barrier for non-technical users to carry out harmonisation tasks, though open questions around trust, reproducibility and reliability remain [56]. Together, these developments suggest that while effective semantic interoperability requires AI-based tools to handle the bulk of alignment and disambiguation, practitioners should retain control over the final results.

Graph-Driven Semantic Interoperability. The fragmentation of the healthcare standards ecosystem is an architectural property of the second generation that no single standard can undo. The third generation requires a standard-agnostic semantic layer that decouples data representation from any particular schema [16]. Knowledge graphs provide the right abstraction, enabling semantic alignment at the level of subgraphs and interconnected concepts rather than isolated attributes [1, 61, 65]. Graph neural networks extend this by capturing multi-hop reasoning paths that reflect causal and taxonomic relationships in clinical ontologies [22, 41, 45], improving robustness to noise and sparsity endemic to real-world EHR data, and supporting principled missing data imputation [9].

Multi-Modal Interoperability. Current frameworks implicitly assume a single data modality as the primary source for interoperability. In practice, beyond structured EHR, clinical data spans heterogeneous sources: free-text notes, medical imaging, genomic profiles, wearable sensor streams, administrative data, etc [30, 42]. Attempts to sidestep this problem and reduce the need for concept harmonisation across systems show how persistent and challenging the problem is [46]. The third generation calls for multi-modal systems capable of ingesting and aligning data across these diverse formats while preserving modality-specific semantics.

Callibrated Prediction Accuracies. While generative AI base models can be semantically well calibrated, predictive or reasoning processes, such as instruction tuning and chain of thought, can significantly distort model confidence, creating overconfident or underconfident predictions without changing accuracy [50]. In healthcare, this directly undermines trust, since clinicians must rely on how sure an AI system is, not only on its average correctness. For example, given a link prediction task that infers missing diagnoses, a predictive model may propose one or several likely diagnoses. By wrapping this model with conformal prediction [4], the system can test whether the uncertainty exceeds a user-defined risk cap. Within an interoperability pipeline, conformalised link prediction can offer a mechanism to quantify and bound uncertainties, and also recalibrate predictions so that reported confidence matches observed correctness [32, 33].

Interoperability in the Agentic AI Era. Agentic, neuro-symbolic approaches offer a path toward adaptive and context-aware interoperability by combining LLMs with structured knowledge from ontologies and knowledge graphs. This enables dynamic semantic alignment across heterogeneous data and reduces reliance on rigid, manually curated mappings. However, challenges remain in ensuring reliability, consistency, and reproducibility in safety-critical settings. Robust evaluation methods, principled integration of symbolic and learned representations, and issues of scalability, cost, and governance (e.g., transparency and auditability) must be addressed before such systems can be deployed in practice [7, 14, 39].

6 PRESENTERS

Andra Ionescu is a Postdoctoral Researcher at KTH Royal Institute of Technology, in the Department of Biomedical Engineering and Health Systems. Her current work focuses on graph-based clinical data harmonisation and semantic interoperability challenges.
Paris Carbone is Associate Professor at KTH Royal Institute of Technology and director of the Data Systems Lab[2]. He is working on scalable data processing systems, generative computing, data streaming, and graph databases. He is a co-recipient of the ACM SIGMOD Systems Award 2023 (Apache Flink).
Sebastiaan Meijer is Professor of Health Care Logistics at KTH Royal Institute of Technology. He is working on healthcare, health prevention and promotion systems, and is the coordinator of the European Digital Innovation Hub Health Data Sweden.
Christos Koutras is a Postdoctoral Research Associate at the VIDA Center of NYU Tandon School of Engineering. His research addresses core data integration challenges, while contributing open-source tools accessible to both domain experts and developers.
Juliana Freire is an Institute Professor at NYU's Tandon School of Engineering and co-directs the VIDA Center. Her research develops methods and systems that enable users to obtain trustworthy insights from data, spanning large-scale data analysis and integration, visualization, machine learning, provenance management, and web information discovery.
Katja Hose is Professor at TU Wien, leading the Data Management and Knowledge-Driven AI Lab[3]. She is working on intelligent data management, knowledge graphs, and applied machine learning with interdisciplinary collaboration in bioscience and health.
Jayanth Raghothama is Associate Professor at KTH Royal Institute of Technology, in the Department of Biomedical Engineering and Health Systems. His research interests include complex systems, health informatics, policy design, and applied ML.

ACKNOWLEDGMENTS

This work was partially funded by Digital Futures.

[2]https://datasystems.nu/
[3]https://dmki-tuwien.github.io/

REFERENCES

[1] Francisco Abad-Navarro and Catalina Martinez-Costa. 2024. A knowledge graph-based data harmonization framework for secondary data reuse. *Computer Methods and Programs in Biomedicine* 243 (2024), 107918.

[2] Kola Adegoke, Abimbola Adegoke, Deborah Dawodu, Akorede Adekoya, Ayoola Bayowa, Temitope Kayode, and Mallika Singh. 2025. Interoperability as a Catalyst for Digital Health and Therapeutics: A Scoping Review of Emerging Technologies and Standards (2015–2025). *International Journal of Environmental Research and Public Health* 22, 10 (2025), 1535.

[3] Fouzia Amar, Alain April, and Alain Abran. 2024. Electronic health record and semantic issues using fast healthcare interoperability resources: systematic mapping review. *Journal of medical Internet research* 26 (2024), e45209.

[4] Anastasios N Angelopoulos and Stephen Bates. 2023. Conformal prediction: A gentle introduction. *Foundations and Trends in Machine Learning* 16, 4 (2023), 494–591.

[5] Paolo Atzeni, Teodoro Baldazzi, Luigi Bellomarini, Eleonora Laurenza, and Emanuel Sallinger. 2025. Semantic-aware query answering with Large Language Models. *Data & Knowledge Engineering* (2025), 102494.

[6] Duane Bender and Kamran Sartipi. 2013. HL7 FHIR: An Agile and RESTful approach to healthcare information exchange. In *Proceedings of the 26th IEEE international symposium on computer-based medical systems*. IEEE, 326–331.

[7] Bikram Pratim Bhuyan, Amar Ramdane-Cherif, Ravi Tomar, and TP Singh. 2024. Neuro-symbolic artificial intelligence: a survey. *Neural Computing and Applications* 36, 21 (2024), 12809–12844.

[8] Christopher Buss, Mahdis Safari, Arash Termehchy, David Maier, and Stefan Lee. 2025. Towards Scalable Schema Mapping using Large Language Models. In *Proceedings of the 1st workshop connecting academia and industry on Modern Integrated Database and AI Systems*. 12–15.

[9] Riccardo Cappuzzo, Saravanan Thirumuruganathan, and Paolo Papotti. 2024. Relational data imputation with graph neural networks. In *EDBT/ICDT 2024, 27th International Conference on Extending Database Technology*.

[10] Raúl Castro Fernandez, Ziawasch Abedjan, Felix Koko, Guoliang Yuan, Samuel Madden, and Michael Stonebraker. 2018. Aurum: A Data Discovery System. In *2018 IEEE 34th International Conference on Data Engineering (ICDE)*. IEEE, 1001–1012.

[11] Martin Pekár Christensen, Aristotelis Leventidis, Matteo Lissandrini, Laura Di Rocco, Renée J Miller, and Katja Hose. 2025. Fantastic Tables and Where to Find Them: Table Search in Semantic Data Lakes.. In *EDBT*. 397–410.

[12] Martin Pekár Christensen, Matteo Lissandrini, and Katja Hose. 2026. An End-to-End Re-Evaluation of Table Entity-Linking Systems. In *2026 IEEE 42th International Conference on Data Engineering (ICDE)*. to appear.

[13] Vassilis Christophides, Vasilis Efthymiou, Themis Palpanas, George Papadakis, and Kostas Stefanidis. 2020. An overview of end-to-end entity resolution for big data. *ACM Computing Surveys (CSUR)* 53, 6 (2020), 1–42.

[14] Bernardo G Collaco, Syed Ali Haider, Srinivasagam Prabha, Cesar A Gomez-Cabello, Ariana Genovese, Nadia G Wood, Sanjay P Bagaria, Narayanan Gopala, Cui Tao, and Antonio Jorge Forte. 2026. The role of agentic artificial intelligence in healthcare: a scoping review. *npj Digital Medicine* (2026).

[15] Roger A Cote and Stanley Robboy. 1980. Progress in medical information management: Systematized Nomenclature of Medicine (SNOMED). *Jama* 243, 8 (1980), 756–762.

[16] Blanda Helena De Mello, Sandro José Rigo, Cristiano André Da Costa, Rodrigo da Rosa Righi, Bruna Donida, Marta Rosecler Bez, and Luana Carina Schunke. 2022. Semantic interoperability in health records standards: a systematic literature review. *Health and technology* 12, 2 (2022), 255–272.

[17] Don E Detmer, Elaine B Steen, and Richard S Dick. 1997. The computer-based patient record: an essential technology for health care. (1997).

[18] AnHai Doan, Alon Halevy, and Zachary Ives. 2012. *Principles of data integration*. Elsevier.

[19] Asuman Dogac. 2012. Interoperability in eHealth systems. *Proceedings of the VLDB Endowment* 5, 12 (2012), 2026–2027.

[20] Robert H Dolin, Liora Alschuler, Calvin Beebe, Paul V Biron, Sandra Lee Boyer, Daniel Essin, Elliot Kimber, Tom Lincoln, and John E Mattison. 2001. The HL7 clinical document architecture. *Journal of the American Medical Informatics Association* 8, 6 (2001), 552–569.

[21] Xin Dong, Alon Halevy, and Jayant Madhavan. 2005. Reference reconciliation in complex information spaces. In *Proceedings of the 2005 ACM SIGMOD international conference on Management of data*. 85–96.

[22] Zhengxiao Du, Chang Zhou, Jiangchao Yao, Teng Tu, Letian Cheng, Hongxia Yang, Jingren Zhou, and Jie Tang. 2021. CogKR: Cognitive graph for multi-hop knowledge reasoning. *IEEE Transactions on Knowledge and Data Engineering* 35, 2 (2021), 1283–1295.

[23] Marco Eichelberg, Thomas Aden, Jörg Riesmeier, Asuman Dogac, and Gokce B Laleci. 2005. A survey and analysis of electronic healthcare record standards. *Acm Computing Surveys (Csur)* 37, 4 (2005), 277–315.

[24] Grace Fan, Jin Wang, Yuliang Li, and Renée J Miller. 2023. Table discovery in data lakes: State-of-the-art and future directions. In *Companion of the 2023 International Conference on Management of Data*. 69–75.

[25] Arden W Forrey, Clement J Mcdonald, Georges DeMoor, Stanley M Huff, Dennis Leavelle, Diane Leland, Tom Fiers, Linda Charles, Brian Griffin, Frank Stalling, et al. 1996. Logical observation identifier names and codes (LOINC) database: a public use set of codes and names for electronic reporting of clinical laboratory test results. *Clinical chemistry* 42, 1 (1996), 81–90.

[26] Juliana Freire, Grace Fan, Benjamin Feuer, Christos Koutras, Yurong Liu, Eduardo Peña, Aécio Santos, Cláudio T Silva, and Eden Wu. 2025. Large language models for data discovery and integration: Challenges and opportunities. *IEEE Data Engineering Bulletin* (2025).

[27] Evan French and Bridget T McInnes. 2023. An overview of biomedical entity linking throughout the years. *Journal of biomedical informatics* 137 (2023), 104252.

[28] David George. 2005. Understanding structural and semantic heterogeneity in the context of database schema integration. *Journal of the Department of Computing, UCLAN* 4, 1 (2005), 29–44.

[29] Alon Halevy, Michael Franklin, and David Maier. 2006. Principles of dataspace systems. In *Proceedings of the twenty-fifth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. 1–9.

[30] Yan Hao, Chao Cheng, Juanjuan Li, Hongwen Li, Xingsi Di, Xiaoxia Zeng, Shoumei Jin, Xiaodong Han, Chongsong Liu, Qianqian Wang, et al. 2025. Multimodal integration in health care: development with applications in disease management. *Journal of medical Internet research* 27 (2025), e76557.

[31] JA Hirsch, G Nicola, G McGinty, RW Liu, RM Barr, MD Chittle, and L Manchikanti. 2016. ICD-10: history and context. *American Journal of Neuroradiology* 37, 4 (2016), 596–599.

[32] Sonia Horchidan and Paris Carbone. 2023. ORB: Empowering Graph Queries through Inference.. In *ESWC Workshops*.

[33] Sonia Horchidan, Fabian Zeiher, Henrik Boström, and Paris Carbone. 2026. CoNANN: Conformal Approximate Nearest Neighbor Search. *Proceedings of the VLDB Endowment* 19 (2026).

[34] Richard Hull. 1997. Managing semantic heterogeneity in databases: a theoretical prospective. In *Proceedings of the sixteenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*. 51–61.

[35] The Healthcare Information and Management Systems Society (HIMSS). 2026. *Interoperability in Healthcare*. https://nationalcapitalarea.himss.org/resources/interoperability-healthcare

[36] Andra Ionescu, Kiril Vasilev, Florena Buse, Rihan Hai, and Asterios Katsifodimos. 2024. Autofeat: Transitive feature discovery over join paths. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*. IEEE, 1861–1873.

[37] Ashish K Jha, Catherine M DesRoches, Eric G Campbell, Karen Donelan, Sowmya R Rao, Timothy G Ferris, Alexandra Shields, Sara Rosenbaum, and David Blumenthal. 2009. Use of electronic health records in US hospitals. *New England Journal of Medicine* 360, 16 (2009), 1628–1638.

[38] Dipak Kalra, Thomas Beale, and Sam Heard. 2005. The openEHR foundation. *Studies in health technology and informatics* 115 (2005), 153–173.

[39] Bardia Khosravi, Pouria Rouzrokh, Tugba Akinci D'Antonoli, Mana Moassefi, Shahriar Faghani, Aawez Mansuri, Keno Bressem, Ali Tejani, and Judy Gichoya. 2026. Agentic AI in Radiology: Evolution from Large Language Models to Future Clinical Integration. *Radiology: Artificial Intelligence* 8, 2 (2026), e250651.

[40] Enas Khwaileh and Yannis Velegrakis. 2025. Dataset Discovery using Semantic Matching. In *28th International Conference on Extending Database Technology, EDBT 2025*. OpenProceedings. org, 649–660.

[41] Yunsoo Kim, Yusuf Abdulle, and Honghan Wu. 2025. Biohopr: A benchmark for multi-hop, multi-answer reasoning in biomedical domain. In *Findings of the Association for Computational Linguistics: ACL 2025*. 12894–12908.

[42] Adrienne Kline, Hanyin Wang, Yikuan Li, Saya Dennis, Meghan Hutch, Zhenxing Xu, Fei Wang, Feixiong Cheng, and Yuan Luo. 2022. Multimodal machine learning in precision health: A scoping review. *NPJ digital medicine* 5, 1 (2022), 171.

[43] Christos Koutras, George Siachamis, Andra Ionescu, Kyriakos Psarakis, Jerry Brons, Marios Fragkoulis, Christoph Lofi, Angela Bonifati, and Asterios Katsifodimos. 2021. Valentine: Evaluating matching techniques for dataset discovery. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*. IEEE, 468–479.

[44] Ernests Lavrinovics, Russa Biswas, Johannes Bjerva, and Katja Hose. 2025. Knowledge graphs, large language models, and hallucinations: An nlp perspective. *Journal of Web Semantics* 85 (2025), 100844.

[45] Ernests Lavrinovics, Russa Biswas, Katja Hose, and Johannes Bjerva. 2025. Multihal: Multilingual dataset for knowledge-graph grounded evaluation of LLM hallucinations. *arXiv preprint arXiv:2505.14101* (2025).

[46] Simon A Lee, Sujay Jain, Alex Chen, Kyoka Ono, Arabdha Biswas, Ákos Rudas, Jennifer Fang, and Jeffrey N Chiang. 2025. Clinical decision support using pseudo-notes from multiple streams of EHR data. *npj Digital Medicine* 8, 1 (2025), 394.

[47] Moritz Lehne, Julian Sass, Andrea Essenwanger, Josef Schepers, and Sylvia Thun. 2019. Why digital medicine depends on interoperability. *NPJ digital medicine* 2, 1 (2019), 79.

[48] Yurong Liu, Eduardo HM Pena, Aécio Santos, Eden Wu, and Juliana Freire. 2025. Magneto: Combining Small and Large Language Models for Schema Matching. *Proceedings of the VLDB Endowment* 18, 8 (2025), 2681–2694.

[49] Roque Lopez, Aécio Santos, Christos Koutras, and Juliana Freire. [n. d.]. BDI-Kit: An AI-powered toolkit for biomedical data harmonization. *Patterns* ([n. d.]).

[50] Preetum Nakkiran, Arwen Bradley, Adam Goliński, Eugene Ndiaye, Michael Kirchhof, and Sinead Williamson. 2025. Trained on Tokens, Calibrated on Concepts: The Emergence of Semantic Calibration in LLMs. *arXiv preprint arXiv:2511.04869* (2025).

[51] OHDSI Community. 2024. Standardized Data: The OMOP Common Data Model. https://www.ohdsi.org/data-standardization/. Accessed 2025-12-03.

[52] Sari Palojoki, Lasse Lehtonen, and Riikka Vuokko. 2024. Semantic interoperability of electronic health records: Systematic review of alternative approaches for enhancing patient information availability. *JMIR medical informatics* 12, 1 (2024), e53535.

[53] Marcel Parciak, Brecht Vandevoort, Frank Neven, Liesbet M Peeters, and Stijn Vansummeren. [n. d.]. Schema Matching with Large Language Models: an Experimental Study. *Proceedings of the VLDB Endowment. ISSN* 2150 ([n. d.]), 8097.

[54] Marcel Parciak, Brecht Vandevoort, Frank Neven, Liesbet M Peeters, and Stijn Vansummeren. 2025. Llm-matcher: A name-based schema matching tool using large language models. In *Companion of the 2025 International Conference on Management of Data*. 203–206.

[55] Erhard Rahm and Philip A Bernstein. 2001. A survey of approaches to automatic schema matching. *the VLDB Journal* 10, 4 (2001), 334–350.

[56] Aécio Santos, Eduardo HM Pena, Roque Lopez, and Juliana Freire. [n. d.]. Interactive Data Harmonization with LLM Agents: Opportunities and Challenges. In *Novel Optimizations for Visionary AI Systems Workshop at SIGMOD 2025*.

[57] Stefan Schulz and Martin Boeker. 2013. BioTopLite: An upper level ontology for the life sciences. Evolution, design and application. In *INFORMATIK 2013–Informatik angepasst an Mensch, Organisation und Umwelt*. Gesellschaft für Informatik eV, 1889–1899.

[58] Nabeel Seedat and Mihaela van der Schaar. [n. d.]. Matchmaker: Self-Improving Large Language Model Programs for Schema Matching. In *GenAI for Health: Potential, Trust and Policy Compliance*.

[59] Wei Shen, Jianyong Wang, and Jiawei Han. 2014. Entity linking with a knowledge base: Issues, techniques, and solutions. *IEEE Transactions on Knowledge and Data Engineering* 27, 2 (2014), 443–460.

[60] Marta Contreiras Silva, Daniel Faria, and Catia Pesquita. 2022. Matching multiple ontologies to build a knowledge graph for personalized medicine. In *European Semantic Web Conference*. Springer, 461–477.

[61] Dominik Tomaszuk, Arslan Smajevic, Tomer Sagi, and Katja Hose. 2025. FHIR Lens: A Graph-Based Approach to Semantic EHR Exploration. In *2025 IEEE 38th International Symposium on Computer-Based Medical Systems (CBMS)*. IEEE, 1–6.

[62] Jan Walker, Eric Pan, Douglas Johnston, Julia Adler-Milstein, David W Bates, and Blackford Middleton. 2005. The Value of Health Care Information Exchange and Interoperability: There is a business case to be made for spending money on a fully standardized nationwide system. *Health affairs* 24, Suppl1 (2005), W5–10.

[63] Eden Wu, Dishita G Turakhia, Guande Wu, Christos Koutras, Sarah Keegan, Wenke Liu, Beata Szeitz, David Fenyo, Cláudio T Silva, and Juliana Freire. 2025. BDIViz: An Interactive Visualization System for Biomedical Schema Matching with LLM-Powered Validation. *IEEE Transactions on Visualization and Computer Graphics* (2025).

[64] Dukyong Yoon, Changho Han, Dong Won Kim, Songsoo Kim, SungA Bae, Jee An Ryu, and Yujin Choi. 2024. Redefining health care data interoperability: empirical exploration of large language models in information exchange. *Journal of medical Internet research* 26 (2024), e56614.

[65] Doudou Zhou, Ziming Gan, Xu Shi, Alina Patwari, Everett Rush, Clara-Lea Bonzel, Vidul A Panickan, Chuan Hong, Yuk-Lam Ho, Tianrun Cai, et al. 2022. Multiview Incomplete Knowledge Graph Integration with application to cross-institutional EHR data harmonization. *Journal of Biomedical Informatics* 133 (2022), 104147.