Toward High-Fidelity Healthcare Digital Twins: Integrating Real-Time Processing, Data Mesh, and MDM

Somnath Banerjee
IEEE Senior Member, Forbes Technology Council Member, NJ, USA

doi: https://doi.org/10.37745/ejcsit.2013/vol13n51114 Published July 10, 2025

Citation: Banerjee S. (2025) Toward High-Fidelity Healthcare Digital Twins: Integrating Real-Time Processing, Data Mesh, and MDM, *European Journal of Computer Science and Information Technology*, 13(51),1-14

Abstract: *Healthcare digital twins are emerging as powerful tools for simulating patient conditions and operational workflows in real time. This paper explores the architectural and technical foundation necessary for building high-fidelity digital twins—those capable of accurate, synchronized, and responsive modeling. It identifies key challenges, including fragmented data, latency, poor semantic alignment, and identity inconsistencies. To overcome these, the study proposes a five-layer architecture integrating real-time data processing, data mesh principles, and master data management (MDM). Through case studies involving heart failure monitoring and hospital operations, the research demonstrates improvements in fidelity, latency, and interoperability. The study concludes with strategic guidance for healthcare organizations and outlines future research topics, including automated twin generation and federated implementations. By aligning infrastructure with intelligence, the proposed model advances the promise of high-fidelity digital twins from concept to clinical reality.*

Keywords: digital twin, healthcare, real-time processing, data mesh, master data management (MDM)

INTRODUCTION

Digital twin (DT) technology is rapidly emerging as a transformative innovation in healthcare, enabling real-time, data-driven simulations of physical entities—whether patients, medical devices, or entire hospital systems. Defined as a dynamic, digital representation of a physical system continuously updated with real-world data, the digital twin concept enables advanced monitoring, predictive analytics, and simulation capabilities that were previously unattainable [1]. In the context of healthcare, digital twins can be broadly classified into two categories: clinical digital twins (also known as human-body digital twins), which replicate individual patient states to support personalized medicine [2], and operational digital twins (also known as medical digital twins), which simulate the functioning of healthcare systems to optimize workflow, resource allocation, and care delivery [3, 4].

The potential of digital twins in healthcare has been recognized globally. The technology was highlighted by Gartner as a top strategic trend in 2020, and market projections estimate the global digital twin healthcare market will surpass $183 billion by 2031 [5]. Use cases now range from simulating chronic disease progression and predicting treatment outcomes to managing hospital bed availability in real time and coordinating interdepartmental operations [6, 7]. However, the implementation of high-fidelity digital twins—those capable of replicating physical systems with high accuracy and minimal latency—faces several challenges.

High-fidelity digital twins rely on the seamless integration of vast volumes of real-time, heterogeneous data. In traditional healthcare settings, this is hindered by legacy IT systems that are often siloed, operate in batch-processing modes, and lack a unified framework for patient identity resolution [8]. These limitations impair the twin's ability to continuously synchronize with its real-world counterpart, leading to discrepancies, delays, and loss of actionable insights. Furthermore, data quality issues, such as missing values, inconsistent coding standards, and duplicated records, undermine the integrity of the digital twin's simulations [9, 10].

This paper reviews challenges in implementing high-fidelity DTs, examines how real-time processing, data mesh, and Master Data Management (MDM) can address these challenges, and proposes a modular, layered architecture that integrates these technologies. Real-time processing ensures continuous data flow with minimal latency, enabling the digital twin to maintain synchronous behaviour with physical systems [11]. The data mesh paradigm breaks organizational data silos by structuring data into domain-specific products governed by standardized policies, facilitating scalable and federated access [12, 13]. MDM provides the backbone of consistency and reliability, reconciling disparate data sources into a single, authoritative view of core healthcare entities such as patients, providers, and devices [8, 14].

Through an in-depth synthesis of literature, theoretical frameworks, and applied system design, the paper offers both a conceptual foundation and a practical roadmap for digital twin deployment in healthcare enterprises.

LITERATURE REVIEW & THEORETICAL UNDERPINNING

Evolution of Digital Twins in Healthcare

The concept of digital twins (DTs) originated in a book in the 1990s, followed by their introduction in manufacturing in the 2000s, and subsequently expanded to the aerospace sector before entering healthcare [1]. One of the first notable applications was the Archimedes project, which simulated diabetic patient behaviour for policy and clinical decisions [15]. The European DISCIPULUS initiative later promoted patient-specific digital avatars across Europe, setting the stage for the clinical and administrative adoption of DTs [16]. More recently, the U.S. Food and Drug Administration (FDA) has supported virtual patients as a component of in silico trials [3].

In healthcare, DTs are generally categorized into clinical digital twins, which simulate physiological or disease-specific conditions, and operational digital twins, which model healthcare system dynamics like bed management and staffing [1, 4]. While clinical DTs require high-resolution biometric, genomic, or behavioural data for predictive modelling, operational DTs rely on administrative data such as EHR timestamps, patient acuity scores, and throughput statistics [6, 7].

Figure 1 illustrates common digital twin implementations in healthcare.

European Journal of Computer Science and Information Technology, 13(51),1-14, 2025
Print ISSN: 2054-0957 (Print)
Online ISSN: 2054-0965 (Online)
Website: https://www.eajournals.org/
Publication of the European Centre for Research Training and Development -UK

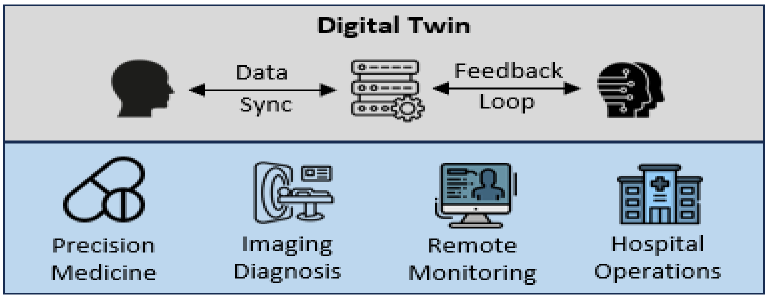

Figure 1: Digital Twins

Table 1 compares the various aspects of clinical and operational DTs.

Table 1: Clinical vs. Operational Healthcare Digital Twins

Dimension	**Clinical Digital Twins**	**Operational Digital Twins**
Primary Focus	Individual patient modelling (physiology, disease progression, treatment response) [6, 17]	System-level modelling (workflow, capacity, staffing, logistics) [7, 18]
Data Sources	Wearables, vitals, genomics, lab results, EHR data	Bed occupancy, staff schedules, admission-discharge-transfer data, telemetry
Use Cases	Predictive alerts, virtual trials, treatment optimization	Capacity planning, ED/ICU flow, resource utilization forecasting
Temporal Resolution	High-frequency, near real-time (sub-minute to hourly)	Near real-time to daily updates, depending on system maturity
Complexity of Modelling	High: requires personalization and physiological models	Medium: driven by historical and real-time operational data
Key Dependencies	Accurate biometric capture, timely data streaming, strong identity resolution [8].	Seamless integration across operational systems, cross-domain metadata governance [13]
Role of MDM	Crucial for entity resolution across devices, EHRs, and biometric feeds	Vital for linking beds, staff, locations, and patients across silos
Governance Model	Clinical governance teams (cardiology, oncology, primary care)	Administrative leadership (operations, scheduling, facilities)
Privacy Considerations	High sensitivity due to personal health and genomic data [14]	Moderate; still subject to HIPAA/GDPR but less risk from indirect identifiers
Outcome Impact	Reduces preventable admissions, improves chronic care, enables precision medicine [3]	Increases efficiency, reduces bottlenecks, enables surge planning [4, 18]

Challenges to High-Fidelity Digital Twins

Achieving high fidelity in Digital Twins—defined as the ability to maintain synchronized, accurate, and real-time representations of physical systems—faces several technical and organizational barriers.

First, data silos and interoperability challenges remain pervasive. Hospital data often resides in EHRs, lab systems, radiology archives, and third-party platforms, each with proprietary formats or inconsistent standards [4]. This fragmentation hinders integration efforts and results in data latency and loss of semantic clarity.

Second, batch processing models still dominate many healthcare IT systems. In scenarios such as ICU telemetry or ED patient intake, updates are often delayed by hours or days—making such systems incompatible with the responsiveness required by high-fidelity DTs [9].

Third, data quality concerns—such as duplicate records, mismatched identifiers, and missing attributes—undermine the reliability of simulations. Without robust identity resolution mechanisms, digital twins may be operating on fragmented or incorrect data, thereby reducing their utility [8].

Fourth, computational complexity arises in multi-scale modelling. Whether simulating cardiac output in a clinical DT or emergency room throughput in an operational DT, real-time modelling requires both algorithmic precision and optimization efficiency [17, 19].

Finally, privacy, security, and governance constraints—governed by regulations like HIPAA or GDPR—restrict data sharing and integration. Privacy-preserving computation and de-identified pipelines are necessary but often under-implemented [14].

Technological Enablers: Real-Time Processing, Data Mesh, and MDM

There are multiple solutions across data engineering that offer powerful tools to mitigate the above challenges:

- **Real-Time Processing**: Tools such as Apache Kafka, Apache Flink, and AWS Kinesis enable real-time ingestion and transformation of data from medical devices, EHRs, and ambient sensors. These systems reduce data latency and allow immediate insights. A recent study showed that real-time telemetry pipelines in ICUs significantly improved early warning scores.
- **Data Mesh**: Coined by Dehghani (2019) [12], the data mesh architecture decentralizes data stewardship, assigning accountability to domain-specific teams. In healthcare, this means radiology, cardiology, and pharmacy units manage their own semantically rich, discoverable data products [13]. Metadata and governance layers ensure compliance and interoperability.
- **Master Data Management (MDM):** MDM systems consolidate patient, provider, and device identities, offering referential integrity and entity resolution. By ensuring that a patient’s identity is consistently tracked across systems, MDM strengthens trust and data usability.

Figure 2 shows a typical data mesh implementation leveraging MDM.

Print ISSN: 2054-0957 (Print)
Online ISSN: 2054-0965 (Online)
Website: https://www.eajournals.org/
Publication of the European Centre for Research Training and Development -UK

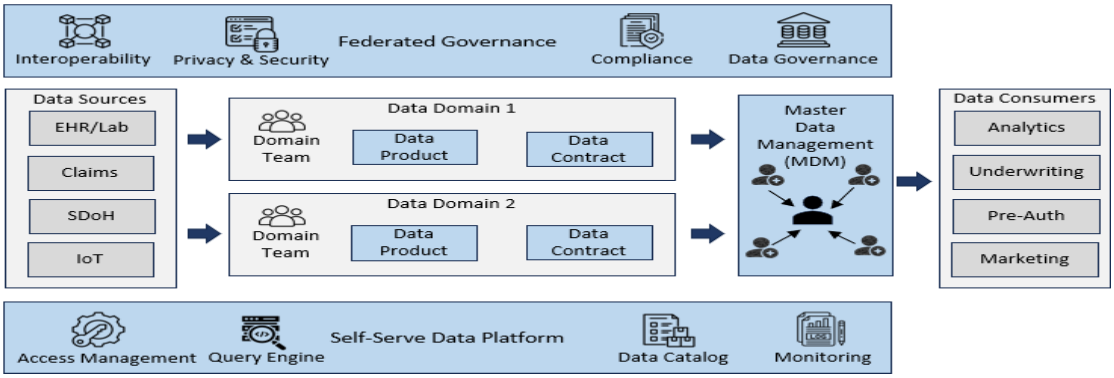

Figure 2: Typical Data Mesh Implementation with MDM

Comparative Role of Real-Time Processing and Data Mesh

Table 2 illustrates the synergistic roles of real-time data streams and data mesh principles. Real-time processing ensures temporal fidelity, i.e., data arrives fast enough to reflect current physical states. Data mesh, by contrast, ensures semantic and organizational fidelity—that data is trusted, well-described, and easily discoverable by consumers, including digital twins [13, 20].

Table 2: Comparison – Data Mesh vs. Real-Time Data Processing

Dimension	Data Mesh	Real-Time Data Processing
Core Purpose	Decentralized data ownership and federated governance across domains [12, 13]	Fast ingestion and transformation of streaming data from physical systems [11].
Target Problem	Semantic fragmentation, lack of ownership, and siloed data across departments	Latency in decision-making caused by batch-oriented systems
Primary Enabler	Domain-driven design and metadata-rich, reusable data products	Event stream processing platforms like Apache Kafka, Flink, AWS Kinesis
Granularity of Data	Managed data products (e.g., lab orders, cardiology images, discharge notes)	Fine-grained event data (e.g., heart rate spikes, room entry, telemetry packets)
Interoperability Role	Encourages standards via shared schemas, metadata layers, and versioning [20]	Ensures interoperability by aligning real-time feeds with downstream APIs and analytic services
Impact on Digital Twins	Enhances semantic fidelity through high-quality, domain-contextualized data [13]	Supports temporal fidelity by enabling low-latency feedback loops from physical to digital and back
Example in Healthcare	A “Radiology Orders” product stewarded by imaging team with standardized codes and lineage [13]	Real-time vitals pipeline detecting bradycardia and triggering alert within 10 seconds [18]
Complementary Use	Works best when paired with real-time systems that feed into data products [12]	Maximized when events are exposed as products within a governed mesh ecosystem [20]

The combination of data mesh and real-time data processing provides a foundation in which streaming data can be exposed as governed, version-controlled data products, enabling digital twins to consume high-quality, low-latency inputs across domains.

METHODOLOGY

This research was conducted as desktop-based exploratory research combining literature review, case study analysis, and architectural design. A systematic review was performed of academic and industry literature (2018–2025) on healthcare digital twins, data integration architectures (real-time streaming, data mesh), and MDM in healthcare. Key insights were drawn from recent systematic reviews, industry whitepapers, and technology standards.

The study followed an iterative methodology: derive requirements from the literature, propose an architecture, and evaluate it against case scenarios. By using emblematic case scenarios, theory was bridged with plausible practice, ensuring the proposed design addresses concrete challenges. Relevant frameworks were consulted, like the Digital Twin Consortium's reference architecture (which highlights data ingestion, semantic layer, simulation, etc.) [21] to align the design with emerging standards.

Theoretical Basis for the Proposed Architecture

The proposed layered architecture draws from four theoretical domains:

- **Systems Control Theory**: Feedback latency impacts system stability. Shorter delays in sensing and acting yield higher fidelity and control accuracy [3].
- **Socio-Technical Systems Theory**: The data mesh's decentralization aligns with socio-technical design by promoting ownership, transparency, and collaborative governance across clinical and operational teams [13].
- **Master Data Theory**: MDM facilitates trust in shared data environments by resolving identity conflicts and enabling consistent semantics, particularly in federated health networks [8].
- **Layered Architecture Models**: Reference models like those proposed by Noeikham et al. (2024) [10] argue for modular, loosely coupled layers in healthcare DT systems, isolating data acquisition, processing, storage, identity resolution, and simulation into distinct architectural functions.

Together, these theories provide the conceptual and design foundation for the five-layer framework proposed in this study.

Architectural Design Approach

Drawing on modular systems design and layered architectural principles [10], the proposed framework is composed of five key layers:

1. **Physical and Data Acquisition Layer** – Captures real-world signals from IoT devices, EHR interfaces, and medical equipment via HL7/FHIR protocols.

European Journal of Computer Science and Information Technology, 13(51),1-14, 2025
Print ISSN: 2054-0957 (Print)
Online ISSN: 2054-0965 (Online)
Website: https://www.eajournals.org/
Publication of the European Centre for Research Training and Development -UK

2. **Real-Time Processing Layer** – Uses tools like Apache Kafka, Flink, and AWS Kinesis to ingest, filter, and transform data streams for immediate analysis.
3. **Data Mesh Layer** – Applies federated governance and domain-driven data ownership, supported by technologies like Lake Formation and Azure Purview [20].
4. **MDM Layer** – Resolves patient and provider identities across disparate systems using probabilistic and referential matching [8].
5. **Digital Twin Model Layer** – Hosts simulation and predictive engines, enabling insights and control signals to be generated and routed back to the physical layer [19].

Each layer should be loosely coupled, independently scalable, and compliant with privacy mandates such as HIPAA and GDPR [14].

Figure 3 illustrates the core components of a healthcare digital twin architecture. The physical twin (patients, devices, environment) continuously transmits data via sensors/IT systems to the digital realm, where a real-time data pipeline processes incoming events. A data mesh layer organizes domain data and ensures shareability across the system, while an MDM layer provides a unified reference for key entities (e.g., patient identity). The digital twin model layer consumes this data to simulate, analyze, and predict the state of the physical twin. A feedback loop can send insights or control actions back to the physical environment.

European Journal of Computer Science and Information Technology, 13(51),1-14, 2025
Print ISSN: 2054-0957 (Print)
Online ISSN: 2054-0965 (Online)
Website: https://www.eajournals.org/
Publication of the European Centre for Research Training and Development -UK

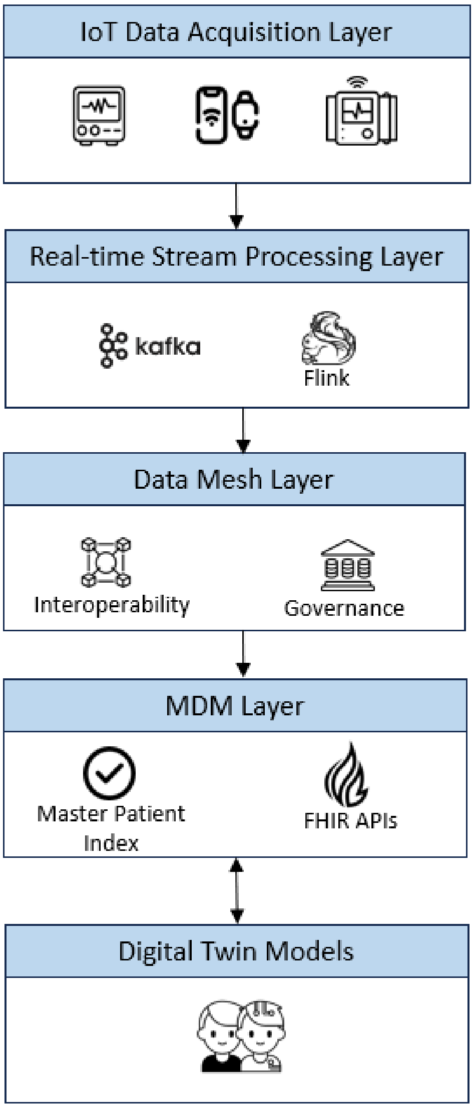

Figure 3: Digital Twin architecture with IoT and Data Mesh

FINDINGS

Analysing the proposed architecture through two emblematic case studies – a heart failure patient twin (clinical) and a hospital operations twin (operational) – demonstrates notable improvements in fidelity, latency, and interoperability compared to traditional setups. Performance improvements were assessed by referencing outcomes documented in prior studies and pilot deployments. The findings are presented in a summary table (Table 3 below) of key performance improvements.

Case Study 1: Heart Failure Patient Digital Twin

Scenario Overview

A 68-year-old male patient with chronic heart failure is outfitted with wearable sensors (continuous ECG patch, blood pressure cuff, weight scale) after hospital discharge. The goal is to use a digital twin to monitor his condition at home, predict decompensation (fluid overload, arrhythmias), and proactively adjust treatment or trigger interventions (e.g. telemedicine visit or medication change) to prevent readmission [6].

Baseline Conditions (Legacy System)

Without the proposed architecture, monitoring relies on sporadic data: the patient might log daily weight manually, and clinicians review data during periodic clinic visits or if the patient calls in symptoms. Data from different devices are not integrated in real-time (perhaps checked separately via each device's app), and there is no unified model predicting heart status. Care adjustments are reactive and based on lagging indicators. Patient deterioration may be noticed only after symptomatic escalation [7].

Proposed Architecture Implementation

The proposed architecture integrates device telemetry into a continuous data stream. The real-time processing layer detects abnormal trends, such as rapid weight gain. The data mesh layer links this data to EHR history and medications [13]. MDM ensures accurate identity resolution across devices and records [8]. The digital twin model predicts a high likelihood of pulmonary edema within a few hours, triggering care team alerts.

Case Study 2: Hospital Operations Digital Twin

Scenario Overview

A large urban hospital implements an operational digital twin to manage patient flow in the emergency department (ED) and intensive care unit (ICU). The twin's goal is to predict bottlenecks (like ED boarding due to no ICU beds) and optimize resource allocation (staffing, bed assignments) in real time, improving throughput and patient outcomes (e.g. reduced wait times).

Baseline Conditions (Legacy System)

Operational decisions are made by hospital staff based on daily bed management meetings, manual data aggregation (bed boards, phone calls), and static predictive tools (if any). Data like current bed occupancy, expected discharges, ED incoming ambulances, etc., are spread across separate systems or spreadsheets. Surge events are reacted to late, increasing boarding time and patient risk [9].

Proposed Architecture Implementation

The physical layer ingests live data from occupancy sensors and ambulatory feeds. The real-time processing layer identifies spikes in ICU demand [11]. The data mesh integrates OR schedules and staffing data products [20]. MDM connects beds, patients, and staff to unified identifiers [8]. The DT simulates ICU overflow in advance and recommends actions.

Table 3 below summarizes key performance improvements.

Table 3: Improvements in Fidelity, Latency, and Interoperability with Proposed Architecture

Metric	Legacy Systems	Proposed Architecture	Supporting References
Fidelity	Partial synchronization; delays and inconsistent semantics reduce model accuracy	High-fidelity simulation enabled via real-time data, identity resolution, and semantic alignment	[3, 8, 13]
Latency	Batch updates every few hours or days; delayed response to clinical or operational events	Sub-minute ingestion and alert generation using stream processing tools	[11]
Data Integration	Manual, siloed ETL workflows; limited EHR-device synchronization	Real-time ingestion across devices, EHRs, and sensors using streaming infrastructure	[7, 17]
Semantic Consistency	Inconsistent terminologies across domains; no unified schema	Domain-governed data products with shared metadata standards and lineage tracking	[13, 20]
Interoperability	Difficult cross-system linkages; proprietary formats	Standards-based APIs (e.g., FHIR, HL7); mesh-enabled governance for multi-system alignment	[12, 18]
Simulation Accuracy	Forecasts based on historical trends; limited personalization	Real-time model updates enable personalized predictions and actionable insights	[6, 19]
Responsiveness	Reactive interventions after adverse events	Proactive response to predicted deterioration or bottlenecks	[17, 18]

The findings demonstrate that architecting for high fidelity – through real-time pipelines, integrated data mesh, and strong MDM – yields tangible benefits: more accurate twins (reflecting reality and predicting future states), more timely insights (enabling proactive action), and more holistic integration of information (breaking down silos). These translate to improved patient care and operational efficiency, which is the ultimate goal of healthcare digital twins.

DISCUSSION & IMPLICATIONS

Theoretical Contributions

This research offers several theoretical advancements to the understanding and design of high-fidelity digital twins (DTs) in healthcare. First, it reconceptualizes fidelity as a system-level outcome rather than a property of the simulation model alone. In the proposed architecture, fidelity emerges from the timely flow of real-world data, its semantic harmonization through data mesh, and accurate entity resolution via Master Data Management (MDM) [3, 8]. This redefinition shifts attention from model complexity to data architecture and system orchestration.

Second, the paper unifies previously fragmented research strands—real-time telemetry [11], distributed data governance [12], and identity management [14]—into a cohesive architectural vision grounded in healthcare informatics. While past efforts have explored these components individually, few have demonstrated their combined potential to enhance real-time feedback, data interoperability, and simulation responsiveness in healthcare DTs [1, 7].

Third, the framework advances socio-technical systems theory by showing how federated governance, enabled through data mesh, and trustable identity, enforced via MDM, create the conditions for resilient, data-driven decision environments. This aligns with the principles of learning health systems that continuously adapt and improve based on operational and clinical feedback [6, 13].

Practical Implications for Healthcare Organizations

For healthcare providers and IT professionals, the proposed architecture provides a roadmap for implementing digital twins at scale. Some practical implications and recommendations include:

- **Invest in Data Infrastructure:** To enable digital twins, hospitals must invest in technologies like streaming platforms, FHIR APIs, and MDM systems. These foundational upgrades not only support twin development but also modernize the overall data infrastructure, improving clinical and operational outcomes.
- **Cross-Functional Collaboration:** Successful digital twins require clinical, IT, and administrative teams to co-develop data products and governance models. This collaboration fosters a culture of data sharing and creates benefits across departments—from smoother ED operations to faster patient care.
- **Regulatory Compliance & Security by Design:** Privacy, role-based access, and de-identification must be integrated into the architecture from the start to meet HIPAA and GDPR requirements. Doing so simplifies audits, builds institutional trust, and allows for expanded use cases like research or FDA-aligned simulations.
- **System Optimization and Cost-Benefit:** Digital twins help hospitals test "what-if" scenarios before implementation, leading to smarter decisions and reduced risk. The resulting improvements in throughput, readmissions, and staffing can produce strong ROI despite upfront tech investment.
- **Workforce Implications:** Adopting digital twins changes staff workflows, requiring training and trust-building, but it can also ease administrative burdens. With phased implementation and proven results, resistance can be overcome and job satisfaction improved through better decision-making tools.

Limitations

This research, while comprehensive in scope, is not without limitations:

- **Model Complexity and Generalizability**: Building sophisticated physiological models is time-consuming and may not generalize across populations. Domain-specific validation is required [16].

- **Cost of Implementation**: Real-time infrastructure and high-grade MDM systems require significant upfront investment, potentially excluding smaller healthcare systems [10].
- **Data Overload**: Streaming data from thousands of endpoints risks overwhelming processing pipelines. Event-driven architectures and edge filtering may be necessary.
- **Regional Constraints**: The framework is optimized for high-resource settings with centralized data governance. Adaptations are required for distributed or low-resource health environments [4].

CONCLUSION

The purpose of this research was to chart a path towards architecting high-fidelity digital twins that truly live up to their promise in healthcare. It was shown that by weaving together real-time data flow, mesh-based data sharing, and rigorous master data management, one can construct a digital twin that is comprehensive, current, and correct – the hallmarks of high fidelity. Such a twin effectively becomes a powerful decision-support ally: for clinicians, a second set of (digital) eyes on each patient, continuously vigilant; for administrators, a real-time systems engineer optimizing the hospital's inner workings. The implications – better health outcomes, more efficient care delivery, and informed strategic planning – are aligned with the pressing needs of healthcare today. The contributions of this work thus lie not only in theory but in practical blueprinting of systems that can realize those benefits. Future research and deployment will no doubt refine these ideas, but the core message is clear: building high-fidelity healthcare digital twins is an achievable goal – and with the architecture presented here, healthcare organizations have a guiding template to begin that journey.

FUTURE RESEARCH

The study opens several pathways for continued exploration:

- **Standardized Fidelity Metrics**: Developing benchmarks for DT performance, including latency, semantic accuracy, and prediction robustness, will aid in cross-system comparison and certification [3].
- **Automated Twin Generation**: AI-driven construction of digital twins from EHRs and sensor data holds promise, especially through transfer learning and graph-based modelling [19].
- **Cross-Institutional DTs**: Federated digital twins spanning multiple hospitals could enable regional surge planning and public health simulation, requiring enhanced identity coordination and governance [1].
- **Ethical and Legal Frameworks**: As DTs become more autonomous, frameworks are needed to clarify liability, consent, and oversight boundaries [14].
- **Generative AI for Twin Data Augmentation**
  Synthetic data can augment digital twin training datasets, support rare disease modeling, and mitigate data scarcity [22].

- **Human-in-the-Loop Explainability & Trust**
  Clinicians need intuitive interfaces and narratives to trust digital twin outputs [23].

REFERENCES

1. Kamel Boulos, M. N., & Zhang, P. (2021). Digital twins: From personalised medicine to precision public health. *Journal of Personalized Medicine, 11*(8), 745. Available: https://doi.org/10.3390/jpm11080745
2. Teng Li, et al. (2024), The status quo and future prospects of digital twins for healthcare, *EngMedicine, Volume 1, Issue 3*, Available: *https://doi.org/10.1016/j.engmed.2024.100042*)
3. Laubenbacher, R., Mehrad, B., Shmulevich, I., & Trayanova, N. (2024). Digital twins in medicine. *Nature Computational Science, 4*(3), 184–191. Available: https://doi.org/10.1038/s43588-024-00607-6
4. Vallée, A. (2023). Digital twin for healthcare systems. *Frontiers in Digital Health, 5*, 1253050. Available: https://doi.org/10.3389/fdgth.2023.1253050
5. Scott, J. (2024). What are digital twins and how can they be used in healthcare? *HealthTech Magazine*. Available: https://healthtechmagazine.net/article/2024/01/what-are-digital-twins-and-how-can-they-be-used-healthcare
6. Armeni, P., Polat, I., De Rossi, L. M., Diaferia, L., Meregalli, S., & Gatti, A. (2022). Digital twins in healthcare: Is it the beginning of a new era of evidence-based medicine? *Journal of Personalized Medicine, 12*(8), 1255. Available: https://doi.org/10.3390/jpm12081255
7. Ringeval M, Etindele Sosso F, Cousineau M, Paré G (2025). Advancing Health Care With Digital Twins: Meta-Review of Applications and Implementation Challenges. J Med Internet Res 2025;27:e69544. DOI: 10.2196/69544. Available: https://www.jmir.org/2025/1/e69544.
8. Verato. (n.d.). Reimagining master data management for healthcare organizations. *Verato Blog*. Available: https://verato.com/blog/data-management-for-healthcare/
9. DePaolis, J. (2024). Journey toward better decision-making using digital twins. *American Planning Association Blog*. Available: https://www.planning.org/blog/9290026/journey-toward-better-decision-making-using-digital-twins/
10. Noeikham, P., Buakum, D., & Sirivongpaisal, N. (2024). Architecture designing of digital twin in a healthcare unit. *Health Informatics Journal, 30*(4), 14604582241296792. Available: https://doi.org/10.1177/14604582241296792
11. Erol, T., Mendi, A. F., & Doğan, D. (2020). The digital twin revolution in healthcare. In *2020 4th International Symposium on Multidisciplinary Studies and Innovative Technologies (ISMSIT)* (pp. 1–7). IEEE. Available: https://doi.org/10.1109/ISMSIT50672.2020.9255249
12. Dehghani, Z. (2019). How to move beyond a monolithic data lake to a distributed data mesh. *martinfowler.com*. Available: https://martinfowler.com/articles/data-monolith-to-mesh.html
13. Nice, B., Hare, J., Montes Nogueira, A. A., & Soundar, J. (2023). Three capabilities that can help healthcare organizations adopt data mesh. *Thoughtworks Insights*. Available: https://www.thoughtworks.com/insights/articles/3-capabilities-healthcare-organizations-adopt-data-mesh
14. Bruynseels, K., de Sio, F. S., & van den Hoven, J. (2018). Digital twins in health care: Ethical implications of an emerging engineering paradigm. Frontiers in Genetics, 9, 31. Available: https://doi.org/10.3389/fgene.2018.00031

15. Eddy, D. M., & Schlessinger, L. (2003). Archimedes: a trial-validated model of diabetes. Diabetes Care, 26(11), 3093–3101. Available: https://doi.org/10.2337/diacare.26.11.3093
16. Elkefi, S., & Asan, O. (2022). Digital twins for managing health care systems: Rapid literature review. Journal of Medical Internet Research, 24(8), e37641. Available: https://doi.org/10.2196/37641
17. Corral-Acero, J., Margara, F., Marciniak, M., et al. (2020). The 'Digital Twin' to enable the vision of precision cardiology: European Heart Journal, 41(48), 4556–4564. Available: https://doi.org/10.1093/eurheartj/ehaa159
18. HealthManagement.org. (2024). Implementing digital twin architecture in healthcare units. *EXEC News*. Available: https://healthmanagement.org/c/hospital/News/implementing-digital-twin-architecture-in-healthcare-units
19. Barbiero, P., Viñas Torné, R., & Lió, P. (2021). Graph representation forecasting of patient's medical conditions: toward a digital twin. *Frontiers in Genetics*, *12*, 652907. Available: https://doi.org/10.3389/fgene.2021.652907
20. Denti, V., & Grüebler, A. (2022). Design a data mesh with event streaming for real-time recommendations on AWS. *AWS Big Data Blog*. Available: https://aws.amazon.com/blogs/big-data/design-a-data-mesh-with-event-streaming-for-real-time-recommendations-on-aws/
21. McKee, D (2023), Platform Stack Architectural Framework: An Introductory Guide. Digital Twin Consortium. Available: https://www.digitaltwinconsortium.org/wp-content/uploads/sites/3/2023/07/Platform-Stack-Architectural-Framework-2023-07-11.pdf
22. Frid-Adar, M., Klang, E., Amitai, M., Goldberger, J., & Greenspan, H. (2018). Synthetic data augmentation using GAN for improved liver lesion classification. Neurocomputing, 321, 321–331. Available: https://doi.org/10.1016/j.neucom.2018.09.013
23. Holzinger, A., Malle, B., Saranti, A., & Pfeifer, B. (2021). Towards multi-modal causability with Graph Neural Networks enabling information fusion for explainable AI. *Information Fusion*, *71*, 28–37. Available: https://doi.org/10.1016/j.inffus.2021.01.008