DIGITAL TWINS IN HEALTHCARE: SIMULATION, DATA, AND THE FUTURE OF PERSONALIZED MEDICINE

Lyuben Zyumbilski[1]
e-mail: lzyumbilski@unwe.bg; l.zyumbilski@gmail.com

Абстракт

Дигиталните близнаци – виртуални представяния на физически системи, които се актуализират непрекъснато с реални данни – се превръщат в трансформираща технология в здравеопазването. Чрез интеграция на физиологични модели, медицински данни и информация за пациенти в реално време, дигиталните близнаци позволяват симулации, прогнозиране и оптимизация на здравни състояния и резултати от лечението. В статията се разглеждат архитектурата, интеграцията на данни, ролята на изкуствения интелект, етичните и правни аспекти, както и бъдещите тенденции.

Abstract

Digital twins-virtual representations of physical systems continuously updated with real-world data-are emerging as a transformative technology in healthcare. By integrating physiological models, medical data, and real-time patient information, digital twins enable simulation, prediction, and optimization of individual health states and treatment outcomes. This paper explores the concept and architecture of digital twins in medicine, focusing on data integration, modeling techniques, and the role of artificial intelligence in maintaining accurate and adaptive simulations. Ethical and privacy challenges are also discussed, along with the potential of digital twin ecosystems for clinical research, preventive medicine, and personalized treatment planning. The study concludes that digital twins represent a paradigm shift from reactive to proactive healthcare, driven by data, computation, and intelligent modeling.

Keywords: Digital Twin, Healthcare, Artificial Intelligence, Data Integration, Simulation

JEL: I10, O33, C88

Introduction

Healthcare systems are increasingly data-intensive, with electronic health records (EHRs), imaging archives, laboratory information systems, and connected medical devices generating high-velocity, high-volume streams of heterogeneous information. Yet clinical decision-making remains largely reactive-responding to events after they occur rather than anticipating and preventing them. Digital twins (DTs) propose an alternative paradigm: continuously updated computational representations that mirror the evolving state of a patient, organ, or clinical process. These representations enable simulation, prediction, and optimization, thereby supporting proactive, personalized care.

While DTs originated in manufacturing and aerospace, where they are used to monitor equipment, optimize maintenance, and enhance safety, their translation to healthcare introduces unique challenges. Biological systems are nonlinear, multi-scale, and subject to substantial inter-individual variability; moreover, healthcare data are fragmented across institutional, regulatory, and technical boundaries. This paper analyzes the foundational components needed to realize healthcare DTs, synthesizes

[1] Докторант към катедра ИТК, УНСС, email: lzyumbilski@unwe.bg; l.zyumbilski@gmail.com

implementation patterns, and outlines the policy and ethical safeguards that must accompany deployment.

The contribution of this work is threefold. First, it proposes a clear conceptual architecture for healthcare digital twins that integrates data pipelines, computational models, and intelligent services. Second, it identifies data and algorithmic requirements for multi-scale modeling and continuous learning. Third, it discusses implementation barriers and governance mechanisms necessary to ensure safety, privacy, and equity. The analysis is scoped to clinical and operational use-cases and focuses on practical pathways from pilot to scale.

Concept and Architecture of Digital Twins

A healthcare digital twin can be conceptualized as a layered system that couples a physical entity (the patient or clinical process) with a virtual entity maintained by data-driven and mechanistic models. The coupling is bi-directional: telemetry, measurements, and clinical events update the virtual state, while simulated trajectories and counterfactuals inform decisions in the physical domain. Four architectural layers are commonly distinguished: the physical layer, the data acquisition and integration layer, the modeling and simulation layer, and the intelligence and interaction layer.

The physical layer encompasses sensors and devices (bedside monitors, implantables, wearables), as well as clinical workflows that produce discrete observations (laboratory results, imaging, clinician notes). The integration layer standardizes inputs using interoperability profiles (e.g., HL7 FHIR resources for clinical data and DICOM objects for imaging) and performs data quality management, identity resolution, and semantic harmonization. Data are persisted in secure repositories with auditable access control. Streaming pipelines enable near real-time ingestion for time-sensitive monitoring tasks.

The modeling and simulation layer combines mechanistic physiology (e.g., compartmental or finite-element models) with data-driven components (e.g., state-space models, recurrent neural networks) to form hybrid digital representations. Model calibration uses patient-specific parameters inferred from historical and current data. The intelligence and interaction layer exposes services-risk scoring, anomaly detection, what-if simulations-through decision-support interfaces. Importantly, this layer maintains explainability artifacts (saliency maps, feature attributions, uncertainty estimates) to support clinical validation and trust.

Data Integration and Multi-Scale Modeling

Constructing clinically useful digital twins requires reconciling heterogeneous data with varying cadence, fidelity, and semantics. Time-aligned integration is essential: physiological sensor streams (millisecond to minute resolution) must be synchronized with episodic EHR updates and periodic imaging. Robust timestamping, clock synchronization, and data versioning are necessary to reconstruct a faithful longitudinal state. Missingness patterns-common in routine care-should be modeled explicitly using probabilistic frameworks rather than imputed naively.

From a representational standpoint, graph-based data structures are advantageous. Patients, encounters, observations, and interventions can be modeled as typed nodes and edges, allowing the twin to capture temporal dependencies and causal pathways. Graph neural networks (GNNs) extend this representation with learnable message passing, supporting individualized predictions. At the same time, mechanistic models encode biophysical constraints that stabilize learning and improve extrapolation.

Multi-scale modeling links molecular-level dynamics (e.g., pharmacokinetics/pharmacodynamics), tissue or organ-level behavior (e.g., electrophysiology), and system-level responses (e.g., hemodynamics, metabolism). Coupling across scales can be achieved through surrogate models-reduced-order approximations trained on high-fidelity simulations-that provide tractable components for real-time use. Bayesian calibration harmonizes priors from population studies with patient-specific data to update parameter distributions and quantify uncertainty.

Interoperability remains a practical bottleneck. Adoption of open standards such as HL7 FHIR for resources (Patient, Observation, MedicationStatement), DICOM for imaging, and OMOP CDM for observational research enables reproducible pipelines and cross-institution validation. Privacy-preserving computation-federated learning, secure enclaves, and differential privacy-allows multi-center modeling without centralizing sensitive data.

Artificial Intelligence and Simulation

AI augments digital twins along three axes: perception, prediction, and control. Perception algorithms transform raw signals into semantically rich features-arrhythmia detection from ECG, motion artifacts from accelerometers, phenotype extraction from clinical notes using natural language processing (NLP). Prediction models estimate risk trajectories, treatment response, and counterfactual outcomes under alternative interventions. Control leverages reinforcement learning and model-predictive control to recommend actions subject to safety constraints.

A critical requirement is continual learning. As clinical practice evolves and populations shift, static models decay. Online learning paradigms-bounded by safeguards such as shadow deployment, drift detection, and human-in-the-loop review-enable twins to adapt without compromising safety. Uncertainty quantification (e.g., ensembles, Bayesian neural networks) is essential for calibrated decision-making, with thresholds and alerts tuned to clinical utility rather than raw accuracy metrics.

Explainability must align with the task. Clinicians often prefer concise, mechanistically plausible rationales over opaque saliency maps. Hybrid twins offer a practical route: they fuse mechanistic structures-where cause–effect relations are explicit-with machine-learned residuals that capture idiosyncratic patterns. This design yields interpretable outputs while maintaining predictive performance.

Practical Applications and Case Studies

Patient monitoring twins integrate bedside telemetry and wearable data to detect early physiological deterioration in general wards. By modeling baseline variability and individual set points, the twin suppresses false alarms and elevates clinically meaningful deviations. In perioperative care, twin-guided hemodynamic simulations inform fluid management strategies tailored to patient-specific cardiovascular dynamics.

In chronic disease management, long-term twins synthesize medication adherence, activity profiles, and lab trends to anticipate exacerbations and recommend adjustments. Hospital operations benefit from process twins that simulate patient flow and resource allocation, enabling proactive staffing and bed management during demand surges. Pharmaceutical R&D explores cohort-level twins-population models calibrated to trial data-to optimize protocol design and identify responders.

Early industrial deployments demonstrate feasibility: vendors have piloted DT-enabled monitoring platforms that integrate EHR streams with telemetry for real-time risk stratification, while academic–

industry consortia have produced high-fidelity organ models (e.g., cardiac electrophysiology) validated against clinical outcomes. Although many pilots remain pre-market, they provide a template for translating DTs from research to clinical practice.

Implementation Barriers and Policy Perspectives

Scaling digital twins beyond pilots faces organizational, technical, and regulatory barriers. Data custodianship is fragmented; hospitals, device manufacturers, and software vendors maintain siloed systems with divergent incentives. Robust data-sharing agreements and standardized APIs are prerequisites for end-to-end integration. Economically, sustainable business models must demonstrate value through reduced adverse events, shorter length of stay, or optimized resource utilization.

From a safety standpoint, DTs implicate software as a medical device (SaMD) regulations. Lifecycle management requires documented verification and validation, real-world performance monitoring, and post-market surveillance. Change control for continuously learning systems must be formalized: updates should pass through staged evaluation with rollback capability and audit trails. Clinical governance needs clear delineation of responsibility between decision-support recommendations and clinician judgment.

Equity considerations demand rigorous bias assessment and inclusivity in training data. Federated studies should ensure representation across demographic groups and care settings to avoid performance gaps. Policymakers can encourage safe innovation through regulatory sandboxes, reference implementations, and public datasets curated to benchmark fairness, robustness, and privacy-preserving efficacy.

Future Trends and Challenges

Three technical trends will shape the next decade of healthcare DTs. First, high-performance and cloud–edge hybrid computing will enable real-time multiscale simulations at the point of care. Second, standardized provenance and model cards will make DT components composable and auditable, supporting reuse across institutions. Third, privacy-enhancing technologies-federated analytics with secure aggregation, homomorphic encryption for selected operations-will allow richer collaboration without exposing raw data.

On the methodological front, digital physiology libraries will expand, providing validated mechanistic building blocks (e.g., circulatory models, pulmonary mechanics) that can be parameterized for individuals. Surrogate modeling and physics-informed neural networks will accelerate complex simulations. Human factors engineering will gain prominence: DT interfaces must integrate seamlessly into clinical workflows, minimize alert fatigue, and present recommendations with appropriate context and uncertainty.

A realistic trajectory is progressive deployment: start with narrow, high-value tasks (e.g., deterioration detection on a specific ward), establish data and governance foundations, and iteratively broaden scope. Success metrics should extend beyond AUC to encompass calibration, clinical impact, equity, and cost-effectiveness. Ultimately, DTs can help shift health systems from episodic care to continuous, preventive, and personalized management.

Conclusion

Digital twins offer a coherent framework that unifies data integration, modeling, and AI-driven decision support for personalized and proactive healthcare. By mirroring patient states and simulating

interventions, DTs can improve timeliness of care, reduce avoidable harm, and optimize resource use. However, realizing this potential requires rigorous engineering, careful governance, and sustained collaboration across clinical, technical, and policy domains.

This paper articulated a reference architecture, clarified data and modeling requirements, and examined practical applications alongside barriers to scale. Future work should prioritize prospective evaluations in diverse health systems, transparent reporting of real-world performance, and open, standards-based ecosystems that enable reproducibility and trust.

References

1. Björnsson, B., Borrebaeck, C., Elander, N. et al. (2020). Digital twins to personalize medicine. Genome Medicine, 12(1), 4.
2. 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.
3. Fuller, A., Fan, Z., Day, C., & Barlow, C. (2020). Digital Twin: Enabling Technologies, Challenges and Open Research. IEEE Access, 8, 108952–108971.
4. Kapteyn, M. G., et al. (2021). Physics-informed machine learning: modeling and simulation. Annual Review of Fluid Mechanics, 53, 491–518.
5. Viceconti, M., & Hunter, P. (2016). The Virtual Physiological Human – A Review of Integrative Biomedical Simulation. Interface Focus, 6(2), 20150092.
6. Bruynseels, K., Santoni de Sio, F., & van den Hoven, J. (2018). Digital Twins in Health Care: Ethical Implications of an Emerging Engineering Paradigm. Frontiers in Genetics, 9, 31.
7. Wong, K. C., & Goh, W. L. (2022). Edge Computing for Medical Digital Twins. Journal of Biomedical Informatics, 130, 104086.
8. Topol, E. (2019). High-performance medicine: the convergence of human and artificial intelligence. Nature Medicine, 25, 44–56.
9. Rumsfeld, J. S., Joynt, K. E., & Maddox, T. M. (2016). Big Data Analytics to Improve Cardiovascular Care: Promise and Challenges. Journal of the American College of Cardiology (JACC), 69(10), 1187–1199.
10. Rajpurkar, P., Chen, E., Banerjee, O., & Topol, E. J. (2022). AI in health and medicine. Nature Medicine, 28, 31–38.
11. Björnsson, B., et al. (2021). Personalized medicine and digital twins: challenges and opportunities. Frontiers in Genetics, 12, 626685.
12. European Commission. (2023). DigiTwins Project – Digital Twins for Preventive Health. Horizon Europe Programme. https://cordis.europa.eu/project/id/777435
13. WHO (2023). Ethics and Governance of Artificial Intelligence for Health. Geneva: World Health Organization. https://www.who.int/publications/i/item/9789240029200
14. Hunter, P., Viceconti, M., & Coveney, P. (2022). Computational biomedicine and the digital twin: The road ahead. Medical Engineering & Physics, 104, 103812.
15. Tolk, A., Diallo, S. Y., & Turnitsa, C. (2020). Modeling and Simulation Support for System of Systems Engineering Applications. John Wiley & Sons.