IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, November 2026
Impact Factor: 8.2

A Patient-Specific Digital Twin Framework for Personalized Healthcare using Hybrid Artificial Intelligence and Physiological Simulation

Bhakti Haridas Maykar[1] and Prof. Divya Ashok Rahane[2]
[1]Student, Department of Artificial Intelligence and Data Science
[2]Assistant Professor, Department of Artificial Intelligence and Data Science
Sahyadri Valley College of Engineering and Technology, Rajuri, India

Abstract : *Personalized healthcare requires clinical decisions that account for differences between individuals rather than relying only on population averages. A patient-specific digital twin (PSDT) is a computational representation of an individual that can be updated using clinical observations and used to explore plausible health trajectories. This paper proposes a conceptual hybrid framework that combines artificial intelligence (AI) with physiological simulation to support patient-level monitoring, risk forecasting, and treatment "what-if" analysis. The framework integrates electronic health records, laboratory results, medical imaging, physiological signals, and wearable data through a quality-control and temporal-alignment layer. Machine-learning models extract patterns from heterogeneous data, while mechanistic or reduced-order simulators represent selected physiological processes. A personalization loop calibrates twin parameters as new observations arrive and communicates uncertainty alongside predictions. The output is designed as clinician-supervised decision support, not autonomous diagnosis or treatment. The paper describes the architecture, workflow, candidate evaluation protocol, application scenarios, and governance requirements. As a conceptual study, it does not report patient experiments or measured clinical performance; prospective validation is identified as necessary before clinical use.*

Keywords: patient-specific digital twin, personalized healthcare, hybrid AI, physiological simulation, clinical decision support, uncertainty quantification, privacy-preserving learning.

I. INTRODUCTION

Healthcare is moving from episodic and reactive care toward continuous, predictive, and individualized management. Digital twins are virtual counterparts of physical systems that are informed by data from their real-world counterparts. In medicine, a patient-specific twin may combine an individual's clinical history, measurements, anatomy, and estimated physiological parameters to represent aspects of health over time [1], [2]. Such models may help clinicians compare possible trajectories and understand how a patient could respond under alternative interventions. The term "digital twin" should be used carefully: a static risk score or an isolated predictive model is not, by itself, a continuously updated patient twin. A meaningful PSDT requires an identifiable patient-specific state, a data-update mechanism, a model of relevant dynamics, and a defined clinical purpose.

Purely data-driven AI can learn complex patterns from large datasets, but it may be difficult to interpret outside the distribution represented in training data. Physiological simulations encode domain knowledge and causal structure, yet may be computationally expensive and require parameters that are not directly measurable. A hybrid approach can combine these strengths: AI estimates latent states and patient-specific parameters, while simulation checks plausible dynamics and explores scenarios. Existing research describes the potential of digital twins in precision cardiology and broader healthcare, while also emphasizing substantial technical and translational challenges [1], [3].

Copyright to IJARSCT
www.ijarsct.co.in
DOI: 10.48175/IJARSCT-41156

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, November 2026
Impact Factor: 8.2

This paper presents a conceptual, modular framework rather than a completed clinical system. Its intended contribution is a practical workflow for integrating multimodal patient data, AI-based inference, physiology-informed simulation, personalization, uncertainty reporting, and clinician review. The framework is designed to be adaptable to a selected disease pathway instead of claiming to model the entire human body.

II. LITERATURE REVIEW

Digital twin concepts were first developed in engineering and systems design, where a virtual model is connected to its physical counterpart throughout operation. Healthcare applications extend this idea to biological systems, where the "physical" counterpart is a patient and the twin is necessarily partial, probabilistic, and bounded by available measurements. Corral-Acero et al. discuss the promise of digital twins for precision cardiology, including the integration of patient data and computational models for individualized cardiovascular care [1].

Reviews of digital-twin technologies in healthcare identify opportunities in monitoring, diagnosis, treatment planning, and predictive care, but note that data integration, model credibility, interoperability, and validation remain barriers [2], [3]. A related development is the use of AI to analyze clinical data and support personalized decisions. Topol highlights the potential of deep learning and human-centered medicine while emphasizing the importance of clinicians and trustworthy implementation [4].

Ethical work on healthcare digital twins cautions that a highly detailed model can create false confidence if its limits are not communicated. Patient consent, privacy, responsibility for decisions, and the possibility of unequal performance must be addressed from design onward [5]. Federated learning offers one possible approach to collaborative model development without centralizing all raw patient data, although it does not eliminate privacy risks and requires governance [6].

The present framework builds on these directions by explicitly separating data quality, AI inference, physiological simulation, personalization, uncertainty communication, and clinical review. It treats the twin as a decision-support artifact whose scope and validation evidence must be specified for each intended use.

III.PROPOSED METHODOLOGY

The proposed PSDT framework contains five connected stages (Fig. 1). The implementation should begin with one clinically meaningful use case—for example, monitoring a chronic condition or estimating response to a defined therapy—rather than attempting to model every organ system.

A. Patient Data Acquisition and Governance

Potential inputs include structured EHR variables (diagnoses, medication, demographics, laboratory values), imaging-derived measurements, physiological signals (such as ECG), and wearable observations (heart rate, activity, or sleep). Data are collected only for the approved purpose, with appropriate consent or lawful authorization. Data provenance, timestamps, units, and missingness are retained.

B. Data Quality and Temporal Alignment

The pipeline checks validity ranges, duplicates, inconsistent units, sensor artifacts, and missing values. Records are aligned to a common timeline; clinically meaningful sampling intervals are preserved. Missingness is not automatically treated as normal. Imputation methods are selected according to data type and evaluated for bias. A data-quality score accompanies the patient state so that poor inputs can reduce confidence or trigger abstention.

C. Hybrid AI and Simulation Core

Separate encoders process different modalities: tabular models for structured clinical variables, temporal models for signals, and imaging models when relevant. The model produces a latent patient state and estimates parameters for a mechanistic physiological simulator. The simulator may use differential equations, compartment models, or a validated reduced-order model for the selected organ or disease. AI predictions and simulation outputs are combined through a transparent fusion layer. The architecture does not assume that one model type is universally best.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, November 2026
Impact Factor: 8.2

D. Patient Personalization and Updating

At initialization, population-level parameters are adapted using available patient observations. As new measurements arrive, the twin updates its state and parameter estimates using a suitable calibration method, such as Bayesian updating or constrained optimization. Updates are bounded by physiological plausibility and are logged. The twin should not silently change its assumptions or extrapolate beyond its validated operating range.

E. Clinical Output and Human Review

Outputs may include a forecast interval, risk category, trend visualization, and comparison of simulated scenarios. The system should display the evidence used, time of last update, uncertainty, and known limitations. A qualified clinician reviews the output alongside the patient's context and makes the final decision. The framework does not autonomously prescribe, alter medication, or replace emergency care.

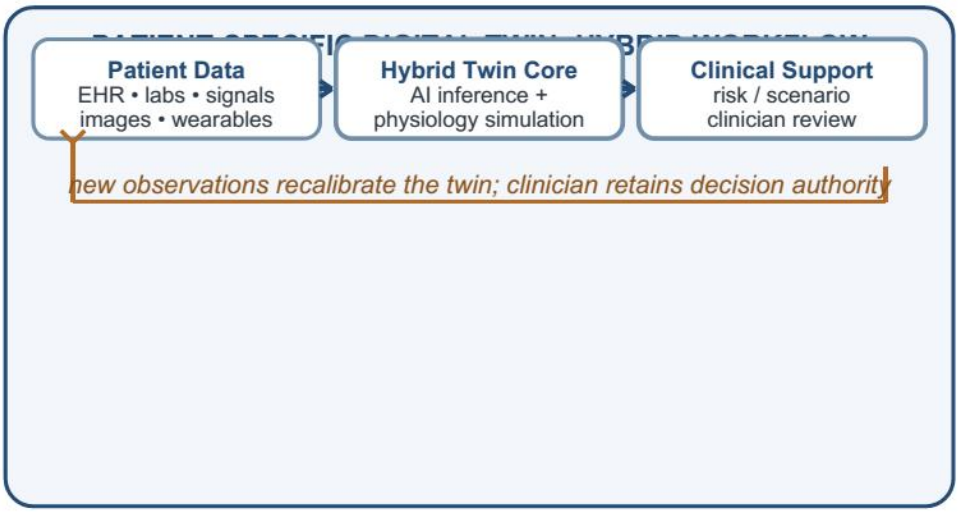

Fig. 1. Conceptual workflow of the proposed patient-specific digital twin. Data updates support personalization; clinical recommendations remain under clinician supervision.

F. Conceptual Model

Let the patient-specific physiological x(t) denote the estimated patient state, parameters, u(t) an intervention or scenario, and y(t) observed measurements. A simplified hybrid model can be expressed as:

$$\dot{x}(t) = f_{phys}(x(t), u(t), \theta) + g_{AI}(z(t))$$
$$y(t) = h(x(t)) + \varepsilon(t)$$
$$\theta_{t+1} = \text{Update}(\theta_t, y_t, \dot{x}_t)$$

Here, fphys represents the physiological simulator, gAI captures data-driven corrections or latent effects, z(t) is the encoded multimodal
input, represents h maps the internal state to measurable quantities, and measurement noise. These equations define a conceptual structure; the specific functions and parameters must be selected and validated for each clinical use case.

IV. ETHICAL PRINCIPLES AND CHALLENGES

A patient twin may combine highly sensitive data across time and sources. Privacy protection therefore requires data minimization, role-based access, encryption, secure audit logs, and clear retention policies. Consent should explain which data are used, what the model can and cannot predict, and whether data may be used for secondary research. Federated or privacy-preserving learning may reduce unnecessary data movement, but model updates and outputs still require security assessment [6].

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, November 2026
Impact Factor: 8.2

Bias and representativeness: Training data may underrepresent age groups, genders, ethnic or socioeconomic groups, comorbidities, or care settings. Performance should be reported across relevant subgroups, and monitoring should look for unequal error rates.

Model credibility: A plausible-looking simulation is not proof of clinical validity. Verification checks whether the software implements its equations correctly; validation assesses whether the model represents the intended physiology for the intended use. Uncertainty and failure conditions must be documented.

Human factors and responsibility: Clinicians need understandable outputs and a way to reject or override recommendations. The interface should avoid automation bias, make stale data visible, and identify the responsible clinical team. Ethical and governance requirements should be incorporated throughout development, consistent with responsible AI guidance [5], [7].

V. APPLICATIONS, RISKS AND EVALUATION

Potential applications include (i) chronic disease monitoring, where a twin summarizes trends and flags changes; (ii) cardiovascular care, where a calibrated model may explore patient-specific hemodynamic scenarios;
(iii) oncology or rehabilitation planning, where simulations may compare defined treatment or recovery pathways; and (iv) remote monitoring, where wearable data update the patient state. These are prospective use cases, not claims of demonstrated benefit.

The main risks are poor-quality inputs, drift, overconfident predictions, computational delay, interoperability failures, privacy breaches, and inappropriate clinical reliance. A twin may also be incomplete: it can represent only the physiology and context for which it was designed.

Proposed evaluation protocol: The framework should be assessed in stages. First, verify software components and data pipelines. Second, conduct retrospective evaluation using a clearly defined cohort and patient-level separation between training and test data. Third, assess calibration, discrimination, forecast error, uncertainty coverage, subgroup performance, and robustness to missing or noisy inputs. Fourth, compare AI-only, simulation-only, and hybrid variants to determine whether the hybrid design adds value. Finally, perform prospective silent-mode evaluation before any clinician-facing deployment, followed by a carefully governed clinical study if justified.

Suggested metrics include AUROC/AUPRC for classification, MAE or RMSE for continuous prediction, calibration error or Brier score for probabilistic outputs, and prediction-interval coverage for uncertainty. Clinical utility should be evaluated with decision-curve analysis and workflow measures, not accuracy alone. No performance values are reported in this conceptual paper because no experiments or patient dataset were conducted.

VI. CONCLUSION

This paper proposed a conceptual patient-specific digital twin framework that combines multimodal clinical data, AI-based inference, physiological simulation, patient-level updating, uncertainty communication, and clinician review. The hybrid design aims to balance the pattern-learning capability of AI with the mechanistic structure of simulation. Its value depends on a narrow and well-defined clinical purpose, reliable data, credible models, transparent uncertainty, and rigorous evaluation. The framework is not a validated medical product and should not be used for diagnosis or treatment without appropriate research, clinical validation, and regulatory review. Future work should implement the architecture for a single disease pathway, evaluate it using an ethically approved dataset, and determine whether hybrid modeling provides measurable benefit over simpler baselines.

ACKNOWLEDGMENT

The authors acknowledge the academic guidance and institutional support of the Department of Artificial Intelligence and Data Science, Sahyadri Valley College of Engineering and Technology, Rajuri. Faculty co-authorship and the final manuscript should be reviewed and approved by all listed authors before submission.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, November 2026
Impact Factor: 8.2

REFERENCES

[1] A. Corral-Acero et al., "The 'digital twin' to enable the vision of precision cardiology," European Heart Journal, vol. 41, no. 48, pp. 4556–4564, 2020, doi: 10.1093/eurheartj/ehaa159.
[2] A. K. Kamel Boulos and P. Zhang, "Digital twins: From personalised medicine to precision public health," Journal of Personalized Medicine, vol. 11, no. 8, art. 745, 2021, doi: 10.3390/jpm11080745.
[3] M. Nadeem, S. Kostic, M. Dornhöfer, C. Weber, and M. Fathi, "A comprehensive review of digital twin in healthcare in the scope of simulative health-monitoring," Digital Health, vol. 11, 2025, doi: 10.1177/20552076241304078.
[4] E. J. Topol, "High-performance medicine: The convergence of human and artificial intelligence," Nature Medicine, vol. 25, pp. 44–56, 2019, doi: 10.1038/s41591-018-0300-7.
[5] A. Bruynseels, F. Santoni de Sio, and J. van den Hoven, "Digital twins in health care: Ethical implications of an emerging engineering paradigm," Frontiers in Genetics, vol. 9, art. 31, 2018, doi: 10.3389/fgene.2018.00031.
[6] N. Rieke et al., "The future of digital health with federated learning," NPJ Digital Medicine, vol. 3, art. 119, 2020, doi: 10.1038/s41746-020-00323-1.
[7] World Health Organization, Ethics and Governance of Artificial Intelligence for Health: WHO Guidance. Geneva: WHO, 2021.

AUTHOR PROFILE



Bhakti Haridas Maykar
Student, AI & Data Science, Sahyadri Valley College of Engineering and Technology, Rajuri, India.
Interests:AI, machine learning, personalized healthcare, and digital health.



Prof. Divya Ashok Rahane
Faculty, AI & Data Science, Sahyadri Valley College of Engineering and Technology, Rajuri, India.
Interests:AI, machine learning, data analysis, and healthcare technology

Submission note: This manuscript is a conceptual framework paper. Please verify all author details, references, institutional requirements, and co-author approval before uploading to a journal or conference website.