BioTwin-Q: A Causal Multimodal Quantum Digital Twin for Personalized Drug Discovery and Precision Healthcare

Maria Michael Visuwasam L
Lincoln University College, 47301, Petaling Jaya, Selangor Darul Ehsan, Malaysia.
pdf.maria@lincoln.edu.my
Dept of CSE (Artificial Intelligence), RMK College of Engineering and Technology,
mariamichaelads@rmkcet.ac.in

Pawan Kumar Chaurasia
Babasaheb Bhimrao Ambedkar Central University Lucknow Uttar Pradesh India.
pkc.gkp@gmail.com

Abstract: Precision healthcare requires models that can integrate heterogeneous patient data, reason about treatment effects rather than correlations alone, and evaluate therapeutic options before clinical deployment. Existing digital-twin approaches provide patient-specific simulation, while causal machine learning supports individualized treatment-effect estimation and quantum machine learning offers emerging tools for molecular representation and optimization. This paper proposes BioTwin-Q, a causal multimodal quantum digital twin framework that combines electronic health records, omics, medical imaging, wearable signals, clinical laboratory data, drug molecular representations and biomedical knowledge in a continuously updated patient twin. A structural causal model forms the reasoning layer for intervention and counterfactual analysis, while a hybrid quantum-classical module is selectively applied to molecular property prediction and candidate optimization. The framework identifies three design-level findings: multimodal fusion is required to reduce single-source bias, causal constraints are required to distinguish actionable treatment effects from predictive associations, and quantum components should be used as task-specific accelerators rather than replacements for classical models. BioTwin-Q is intended for drug prioritization, treatment-response simulation, adverse-event risk assessment and clinician-in-the-loop precision decision support. The paper also defines an evaluation protocol covering predictive accuracy, calibration, individualized treatment effects, molecular validity, uncertainty and safety.

Keywords: causal AI; multimodal learning; quantum machine learning; drug discovery; precision healthcare;

Introduction

Healthcare data are increasingly multimodal: structured electronic health records coexist with longitudinal laboratory measurements, medical images, genomic and transcriptomic profiles, wearable streams, medication histories and external biomedical knowledge. However, these sources are commonly analyzed in separate pipelines, making it difficult to maintain a coherent patient-specific computational representation. Medical digital twins address this gap by creating computational models that represent an individual and can be updated as new observations arrive. Recent reviews identify personalized therapy, risk prediction, drug discovery and clinical-trial simulation as important digital-twin applications, while also emphasizing data integration, privacy and validation challenges [1], [2].

At the same time, artificial intelligence is changing both molecular discovery and clinical decision support. Deep learning has enabled high-quality protein-structure prediction [3] and machine-learning systems can generate and optimize molecular candidates, although the chemical search space and real-world validation remain difficult [4]. Causal machine learning adds a complementary capability: instead of asking only whether an outcome can be predicted, it estimates how an outcome may change under a treatment or intervention for a particular patient [5]. This distinction is important because a high-performing predictive model does not automatically identify an effective intervention.

Quantum machine learning (QML) introduces a further computational direction. Recent reviews describe variational quantum circuits, quantum data encoding and hybrid quantum-classical models for molecular property prediction and generation, but also emphasize the immaturity of current hardware and the absence of established quantum advantage in practical healthcare workloads [6], [7]. BioTwin-Q therefore does not assume that quantum computing is universally superior. Instead, it treats quantum models as optional, task-specific components within a broader clinically governed digital-twin architecture.

The objective of this paper is to present BioTwin-Q as an integrated framework connecting multimodal patient representation, causal counterfactual reasoning, and quantum-assisted drug discovery. The framework addresses the question of how a patient-specific digital twin can combine heterogeneous biological and clinical evidence with causal intervention modeling and emerging quantum computation while preserving uncertainty, traceability, and human oversight.

Related Work

Digital twins in medicine have progressed from physiological simulation toward dynamic, data-driven representations of individual patients. Laubenbacher et al. describe medical digital twins as a promising route to personalized medicine while highlighting integration and privacy as major barriers [1]. Katsoulakis et al. further identify applications ranging from individualized therapy and clinical trials to biomarker and drug discovery [2]. A recent review found that healthcare digital-twin studies span diagnosis, therapy optimization, monitoring and system-level modeling, but validation and implementation remain open problems [8].

Causal AI addresses a different limitation. Feuerriegel et al. show that causal machine learning can estimate individualized treatment effects from clinical-trial and real-world data, while warning that biased observational data can produce misleading conclusions [5]. A 2026 causal-digital-twin framework explicitly combines structural causal models, potential outcomes and reinforcement learning to move from prediction toward adaptive intervention [9]. BioTwin-Q extends this direction by coupling the causal layer to multimodal patient states and drug representations.

In drug discovery, AlphaFold demonstrated that deep learning can substantially improve computational protein-structure prediction [3], while recent molecular-design research surveys generative and optimization methods for candidate creation [4]. QML research has explored molecular property prediction and molecular generation using hybrid quantum-classical approaches [6]. However, a systematic review of QML for digital health reported limited evidence for empirical quantum advantage and substantial concerns about realistic data encoding and hardware conditions [7]. These findings motivate a conservative hybrid design rather than a quantum-first architecture.

Table 1. Comparison of BioTwin-Q with representative research directions

Work	Multimodal	Causal	Quantum	Patient twin	Drug discovery
Digital twins [1], [2]	Yes	No	No	Yes	Yes
Causal ML [5], [9]	Partial	Yes	No	Partial	Partial
QML drug discovery [6]	Partial	No	Yes	No	Yes
Generative molecular AI [4]	Partial	No	No	No	Yes
BioTwin-Q	Yes	Yes	Yes	Yes	Yes

Dataset and Experimental Setup

BioTwin-Q follows a cross-domain multimodal dataset strategy rather than claiming direct patient-level linkage between unrelated public repositories. MIMIC-IV v3.1 provides the clinical/EHR stream, including 364,627 unique individuals, 546,028 hospitalizations and 94,458 ICU stays, with demographics, diagnoses, procedures, laboratory measurements,

medications and other hospital data [11]. TCGA provides the molecular layer, with more than 20,000 primary cancer and matched-normal samples across 33 cancer types and genomic, epigenomic, transcriptomic and proteomic measurements [12]. TCIA provides disease-oriented medical imaging collections, including CT, MRI, PET and digital pathology, with supporting outcomes, treatment, genomics or image-analysis data where available [13]. ChEMBL 37 supplies the drug-discovery layer and, at its current release, contains 18,552 targets, 2,921,148 distinct compounds and 24,527,044 bioactivity records [14].

The datasets are aligned through disease concepts, treatment/target terminology, molecular identifiers, temporal context and a shared biomedical knowledge layer. Thus, MIMIC-IV is used to construct and evaluate the patient-state and treatment-response components; TCGA and TCIA provide complementary molecular and imaging evidence for disease-aware representation learning; and ChEMBL supports compound representation, target association and molecular-property prediction. The framework does not merge these repositories as if they described the same individuals. Instead, modality-specific encoders are trained or evaluated within their respective data domains and their latent representations are connected through disease- and treatment-level concepts. This design avoids unsupported patient identity matching while preserving the intended multimodal digital-twin workflow.

Table 2. Datasets and their role in the BioTwin-Q experimental design

Dataset	Primary modality	Scale / coverage	Role in BioTwin-Q
MIMIC-IV v3.1	EHR / clinical	364,627 individuals; 546,028 hospitalizations; 94,458 ICU stays	Patient state, treatment/outcome modeling
TCGA	Omics / molecular	20,000+ samples; 33 cancer types	Disease and molecular representation
TCIA	Imaging	Disease-oriented imaging collections	Imaging representation and multimodal complement
ChEMBL 37	Drug / bioactivity	18,552 targets; 2.92M compounds; 24.53M activities	Drug representation and candidate ranking

Key Contribution

BioTwin-Q contributes an integrated architecture rather than a single predictive model. First, it defines a persistent patient state representation that is updated from heterogeneous observations. Second, it places a causal model between representation learning and decision support so that treatment options can be evaluated through interventions and counterfactuals. Third, it introduces a quantum layer only for selected subproblems where compact representations and hybrid optimization are appropriate. Finally, it treats uncertainty, provenance, privacy and clinician review as first-class components rather than post-processing steps.

The framework is organized around five principles: (i) multimodal evidence should be aligned to a common patient timeline; (ii) causal assumptions should be explicit and auditable; (iii) quantum components should be benchmarked against strong classical baselines; (iv) predictions should include calibrated uncertainty and provenance; and (v) the twin should support, not autonomously replace, clinical judgment.

Method, Experiments and Results

Figure 1 presents the proposed BioTwin-Q pipeline. The physical patient is represented through longitudinal clinical and biological observations. Data ingestion performs temporal alignment, missingness characterization, normalization and modality-specific quality checks.

Encoders then convert each modality into latent representations. Cross-attention or gated fusion produces a patient state vector z_t that summarizes current and historical evidence.

The digital-twin state can be expressed as z_t = F(z_(t-1), x_t, k_t), where x_t denotes newly observed patient data and k_t denotes contextual biomedical knowledge. The causal layer represents treatment, disease and outcome relationships using a structural causal model. For a candidate intervention a, the twin estimates a counterfactual outcome Y(a) and an individualized treatment effect $\tau(x,a)=E[Y(a)-Y(a_0)|X=x]$. This enables comparison of candidate therapies for a specific patient instead of ranking treatments only by population-level prediction.

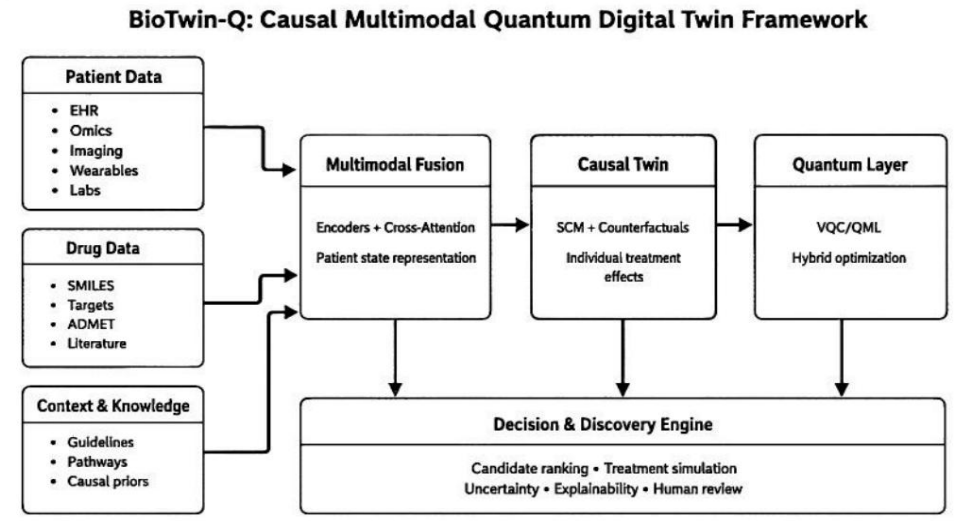

Figure 1. BioTwin-Q architecture integrating multimodal fusion, causal simulation and a hybrid quantum discovery layer.

The drug-discovery branch represents compounds using molecular graphs, SMILES-derived embeddings, physicochemical descriptors and target information. Classical encoders first reduce the representation to a compact feature space. A variational quantum circuit can then encode selected features and optimize a parameterized circuit for tasks such as molecular property classification, ranking or constrained optimization. The resulting score is fused with classical predictions for efficacy-related properties, toxicity, pharmacokinetics and manufacturability. No quantum advantage is assumed; every quantum experiment is paired with a classical baseline under comparable data and evaluation conditions.

Because BioTwin-Q is a framework proposal, this version defines the experimental protocol and dataset mapping without fabricating numerical performance results. The planned experiments use temporally separated MIMIC-IV cohorts for patient-side prediction and treatment-response analysis, disease-stratified TCGA/TCIA subsets for multimodal representation learning, and ChEMBL subsets for molecular property prediction and candidate ranking. All quantum experiments will be compared with matched classical baselines.

The principal design-level result is that the three layers solve different failure modes: multimodal fusion improves evidence coverage, causal reasoning changes the output from association to intervention-oriented estimates, and the quantum module provides a controlled experimental pathway for selected optimization tasks. A second result is methodological: the architecture should be evaluated with uncertainty and calibration rather than accuracy alone. A third result is practical: quantum components are most defensible when the task is compact enough for present hardware and when a clear classical comparator is available [6], [7].

Discussions

BioTwin-Q is intended to bridge three research communities that often use different objectives and evaluation criteria. Digital-twin research emphasizes continuous patient representation and simulation; causal AI emphasizes valid intervention reasoning; and QML emphasizes computational representations and optimization on quantum hardware. Their integration creates opportunities but also exposes dependencies. A causal model built on biased multimodal data cannot become reliable simply because it is embedded in a digital twin. Likewise, a quantum circuit does not make a molecular prediction causal or clinically valid.

An important design choice is the separation between prediction and intervention. Predictive models can identify patients likely to respond to a drug, but treatment selection requires estimating what would happen if the patient received one intervention rather than another. The causal graph should therefore encode clinically defensible relationships and temporal ordering, with expert review of confounders, mediators and treatment pathways. Where randomized data are available, they should anchor the causal estimates; observational data should be accompanied by sensitivity analysis [5], [9].

Privacy and governance are equally important. The twin should use data minimization, access control, provenance tracking and secure model-training practices. Federated or privacy-preserving learning can be considered where data cannot be centralized. For drug development, model documentation should cover dataset origin, intended use, performance limits, monitoring and change control. Current regulatory work on AI in drug development emphasizes human-led governance, data quality, model validation and ongoing monitoring [10]. These requirements align naturally with BioTwin-Q's clinician-in-the-loop design.

Quantum computing remains a limitation rather than a guaranteed advantage. Current QML studies frequently operate under restrictive assumptions about hardware, data encoding and circuit depth [7]. Consequently, BioTwin-Q uses a hybrid architecture in which classical models perform data-intensive representation learning while quantum circuits are reserved for compact, experimentally testable subproblems. This also makes the framework future-proof: if quantum advantage is not demonstrated for a task, the classical pathway remains operational.

Conclusions

BioTwin-Q addresses the fragmentation of multimodal patient data and the gap between predictive AI and intervention-oriented precision healthcare by combining multimodal representation learning, a structural causal digital twin, counterfactual treatment simulation and a selectively applied hybrid quantum-classical drug-discovery module. The framework indicates that multimodal fusion, explicit causal assumptions and rigorous classical-versus-quantum benchmarking are complementary requirements, while the proposed evaluation protocol emphasizes calibration, individualized treatment effects, molecular validity, uncertainty and subgroup robustness rather than accuracy alone. However, BioTwin-Q is currently a framework proposal and does not establish clinical benefit or quantum advantage. Future work should implement the twin using longitudinal multi-omics and EHR datasets, validate causal assumptions with clinical experts, test quantum circuits on noisy hardware, evaluate fairness and privacy, and conduct prospective clinical and pharmaceutical validation before deployment.

References

[1] R. Laubenbacher, B. Mehrad, I. Shmulevich, and N. Trayanova, "Digital twins in medicine," Nature Computational Science, vol. 4, pp. 184–191, 2024, doi: 10.1038/s43588-024-00607-6.

[2] E. Katsoulakis et al., "Digital twins for health: a scoping review," npj Digital Medicine, vol. 7, art. no. 77, 2024, doi: 10.1038/s41746-024-01073-0.

[3] J. Jumper et al., "Highly accurate protein structure prediction with AlphaFold," Nature, vol. 596, pp. 583–589, 2021, doi: 10.1038/s41586-021-03819-2.

[4] "Machine learning-aided generative molecular design," Nature Machine Intelligence, 2024, doi: 10.1038/s42256-024-00843-5.

[5] S. Feuerriegel et al., "Causal machine learning for predicting treatment outcomes," Nature Medicine, vol. 30, pp. 958–968, 2024, doi: 10.1038/s41591-024-02902-1.

[6] A. M. Smaldone et al., "Quantum machine learning in drug discovery: applications in academia and pharmaceutical industries," Chemical Reviews, 2025, doi: 10.1021/acs.chemrev.4c00678.

[7] R. S. Gupta et al., "Quantum machine learning for digital health? A systematic review," npj Digital Medicine, 2025, doi: 10.1038/s41746-025-01597-z.

[8] K. Zhang et al., "Concepts and applications of digital twins in healthcare and medicine," Patterns, vol. 5, no. 8, art. no. 101028, 2024, doi: 10.1016/j.patter.2024.101028.

[9] A. Vallée, "From prediction to intervention: causal digital twins for personalized clinical decision support," Journal of Translational Medicine, vol. 24, art. no. 441, 2026, doi: 10.1186/s12967-026-07895-8.

[10] U.S. Food and Drug Administration, "Using Artificial Intelligence & Machine Learning in the Development of Drug & Biological Products: Discussion Paper and Request for Feedback," revised Feb. 2025.