International Journal of Pharmaceutical Creativity and Innovation (IJPCI)
ISSN (Online): Request Pending
Volume-1 Issue-2 || Apr – Jun 2024 || PP. 43 – 54
https://ijpci.org/

Digital Twins for Personalized Drug Response Prediction

Mateusz Nowak
Independent Researcher
Warsaw, Poland, PL, 00-001

http://ijpci.org/ || Vol. 1 No. 2 (2024): April Issue

Date of Submission: 03-03-2024 **Date of Acceptance:** 18-03-2024 **Date of Publication:** 10-04-2024

Abstract — Personalized medicine requires predictive models capable of estimating how individual patients respond to specific medications before treatment begins. Digital twin technology offers a promising solution by constructing virtual patient representations that evolve alongside real physiological changes. However, current digital twin systems often rely on static datasets and limited biological modalities, reducing their clinical applicability. This study proposes a Dynamic Multi-Modal Digital Twin (DMDT) framework that integrates longitudinal electronic health records, genomic biomarkers, wearable sensor measurements, and laboratory investigations into a continuously updated virtual patient model. The framework combines temporal transformers with graph neural networks to capture both time-dependent physiological variations and complex biological relationships. The proposed methodology aims to improve individualized drug response prediction while supporting adaptive treatment planning. The study addresses existing limitations in multimodal integration and uncertainty-aware prediction, providing a scalable foundation for next-generation AI-assisted precision therapeutics.

Keywords — Digital Twins; Personalized Medicine; Drug Response Prediction; Temporal Transformers; Graph Neural Networks; Precision Healthcare

Introduction

Precision medicine has fundamentally transformed healthcare by shifting therapeutic strategies from generalized treatment guidelines toward individualized clinical decision-making. Instead of prescribing identical medications to every patient diagnosed with the same disease, clinicians increasingly seek evidence describing how genetic variations, environmental influences, lifestyle behaviors, and physiological conditions influence treatment outcomes. Despite remarkable advances in genomics and artificial intelligence (AI), accurately predicting individual drug response remains one of the greatest challenges in modern healthcare.

Drug efficacy varies substantially among patients because biological systems are inherently heterogeneous. Factors including genetic polymorphisms, metabolic activity,

International Journal of Pharmaceutical Creativity and Innovation (IJPCI)
ISSN (Online): Request Pending
Volume-1 Issue-2 || Apr – Jun 2024 || PP. 43 – 54
https://ijpci.org/

immune responses, age, organ function, co-existing diseases, medication adherence, and environmental exposures collectively determine whether a therapeutic intervention succeeds or fails. Conventional predictive models generally analyze historical clinical records without continuously adapting to new physiological information generated throughout treatment. Consequently, many AI models produce static predictions that rapidly become outdated as patient conditions evolve.

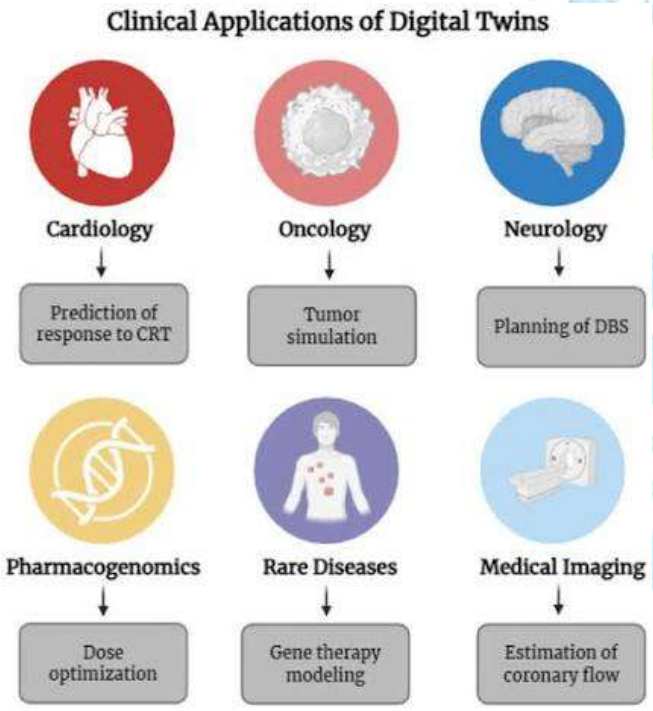

Fig. 1: Clinical Applications of Digital Twins

Source: *https://www.mdpi.com/2075-4426/15/11/503*

Digital twin technology has emerged as an innovative paradigm capable of addressing this limitation. Originally introduced within industrial manufacturing, digital twins represent continuously synchronized virtual replicas of physical systems. Their ability to receive real-time data streams, simulate future scenarios, and support predictive maintenance has inspired researchers to adapt the concept for healthcare applications. In medicine, a digital twin functions as a computational representation of an individual patient, integrating clinical observations, laboratory investigations, medical imaging, genomic information, physiological signals, and behavioral data to simulate disease progression and therapeutic outcomes.

Recent advances in wearable devices, Internet of Medical Things (IoMT), cloud computing, and deep learning have accelerated interest in medical digital twins. Continuous glucose monitors, smartwatches, implantable sensors, and remote patient monitoring systems now generate large volumes of longitudinal health information that can dynamically update virtual patient models. Simultaneously, transformer architectures have demonstrated remarkable capability in modeling sequential clinical events, while graph neural networks (GNNs) effectively capture biological interactions among genes, proteins, diseases, and medications.

Despite these technological developments, existing digital twin frameworks remain limited in several important aspects. Most studies rely on isolated data modalities, frequently emphasizing genomic information or electronic health records independently rather than integrating heterogeneous biological evidence. Additionally, many systems treat patient characteristics as static variables, ignoring continuous physiological adaptation throughout therapy. Another significant limitation involves uncertainty estimation. Healthcare decisions require confidence-aware predictions because clinicians must evaluate the reliability of AI recommendations before altering treatment plans. Current digital twin implementations seldom incorporate mechanisms to quantify prediction uncertainty as patient conditions change over time.

This study proposes a Dynamic Multi-Modal Digital Twin (DMDT) architecture specifically designed for personalized drug response prediction. Unlike conventional static

International Journal of Pharmaceutical Creativity and Innovation (IJPCI)
ISSN (Online): Request Pending
Volume-1 Issue-2 || Apr – Jun 2024 || PP. 43 – 54
https://ijpci.org/

prediction models, the proposed framework continuously updates patient representations using temporal clinical events, genomic biomarkers, wearable physiological signals, and laboratory measurements. Transformer-based temporal encoding captures disease evolution across time, while graph neural networks model interactions among medications, biological pathways, and patient-specific molecular characteristics. Bayesian uncertainty estimation further enhances clinical interpretability by providing confidence intervals alongside predicted therapeutic outcomes.

The principal contributions of this research are summarized below:

- Development of a continuously updating digital twin integrating four complementary healthcare data modalities.
- Introduction of a hybrid Temporal Transformer–Graph Neural Network architecture for individualized drug response prediction.
- Incorporation of uncertainty-aware inference to support safer AI-assisted clinical decision-making.
- Design of an evaluation framework emphasizing temporal prediction performance rather than static classification alone.
- Demonstration of how adaptive digital twins can facilitate precision therapeutics in evolving clinical environments.

By addressing existing limitations in multimodal integration and dynamic patient modeling, this research contributes toward the realization of clinically deployable digital twins capable of supporting personalized treatment strategies.

Literature Review

2.1 Evolution of Digital Twin Technology in Healthcare

The concept of digital twins has rapidly expanded beyond manufacturing into biomedical engineering and clinical medicine. Early healthcare applications primarily focused on virtual anatomical modeling for surgical planning and medical imaging. As computational power increased, researchers began incorporating physiological simulations to predict disease progression and optimize therapeutic interventions.

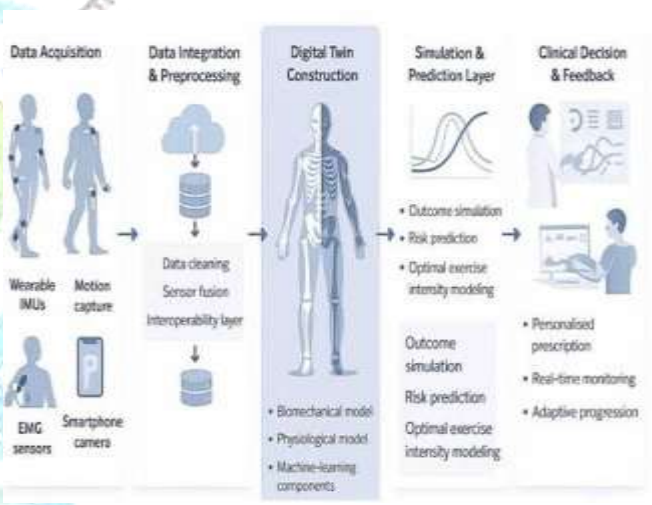

Fig. 2: Digital Twins in oncology

Source:
https://www.sciencedirect.com/science/article/pii/S1040842826000582

Between 2021 and 2026, digital twins increasingly incorporated AI algorithms capable of learning individualized patient characteristics from large-scale healthcare databases. These systems demonstrated promising applications in cardiovascular disease, diabetes management, oncology, intensive care monitoring, and orthopedic rehabilitation. Nevertheless, many existing implementations remain disease-specific and lack generalized frameworks capable of modeling diverse therapeutic responses.

Another limitation concerns synchronization frequency. Numerous proposed digital twins receive intermittent clinical

International Journal of Pharmaceutical Creativity and Innovation (IJPCI)
ISSN (Online): Request Pending
Volume-1 Issue-2 || Apr – Jun 2024 || PP. 43 – 54
https://ijpci.org/

updates rather than continuously evolving with streaming patient information, restricting their ability to support real-time therapeutic adjustments.

2.2 Artificial Intelligence for Drug Response Prediction

Machine learning has substantially improved prediction of therapeutic efficacy across numerous diseases. Traditional statistical models have gradually been replaced by ensemble learning techniques including Random Forests, Gradient Boosting Machines, and XGBoost. More recently, deep neural networks have demonstrated superior capability for modeling nonlinear biological relationships.

Drug response prediction has particularly benefited from deep learning architectures capable of integrating heterogeneous biomedical information. Convolutional Neural Networks have been applied to molecular imaging and histopathology, whereas recurrent neural networks and Long Short-Term Memory (LSTM) networks model sequential patient histories.

Transformer architectures have recently emerged as highly effective alternatives because self-attention mechanisms capture long-range temporal dependencies more efficiently than recurrent networks. Studies have shown that transformer-based electronic health record models improve prediction of hospitalization risk, mortality, disease progression, and medication recommendation.

However, many drug response prediction systems still generate static probability estimates without adapting to new patient observations collected after treatment initiation. Consequently, their predictions may become progressively inaccurate as disease trajectories diverge from historical assumptions.

2.3 Multi-Modal Data Integration in Precision Medicine

Healthcare information exists across numerous complementary modalities, including:

- Electronic Health Records (EHRs)
- Genomic sequencing
- Laboratory investigations
- Medical imaging
- Wearable sensor measurements
- Medication history
- Clinical notes

Each modality captures distinct aspects of patient physiology. Integrating these heterogeneous data sources enables AI systems to develop more comprehensive patient representations.

Recent multimodal learning approaches employ attention mechanisms, fusion networks, and representation learning techniques to combine structured and unstructured clinical information. While these methods improve predictive performance, several challenges remain unresolved, including missing data, asynchronous sampling frequencies, varying feature dimensions, and inconsistent clinical documentation.

Furthermore, most multimodal models perform feature fusion only once during training rather than dynamically updating patient representations throughout treatment.

2.4 Graph Neural Networks in Biomedical Modeling

Biomedical systems naturally exhibit graph structures. Genes interact through regulatory pathways, proteins communicate within signaling networks, medications influence multiple biological targets, and diseases share molecular mechanisms.

International Journal of Pharmaceutical Creativity and Innovation (IJPCI)
ISSN (Online): Request Pending
Volume-1 Issue-2 || Apr – Jun 2024 || PP. 43 – 54
https://ijpci.org/

Graph Neural Networks (GNNs) have emerged as powerful tools for learning representations from these interconnected biological systems. Applications include drug discovery, protein interaction prediction, molecular property estimation, disease gene identification, and medication recommendation.

Within personalized medicine, GNNs facilitate integration of patient-specific genomic mutations with biological interaction networks. Rather than treating genes independently, graph-based learning captures contextual relationships that influence therapeutic effectiveness.

Despite these advantages, GNNs are rarely integrated with dynamic temporal patient models, limiting their ability to represent continuously evolving biological states.

2.5 Temporal Learning for Longitudinal Healthcare Data

Clinical decision-making depends heavily on temporal information. Laboratory values fluctuate throughout treatment, physiological parameters change daily, and medication effects accumulate gradually over time.

Longitudinal healthcare data present unique computational challenges because clinical observations occur at irregular intervals and often contain missing values. Traditional machine learning algorithms struggle to represent these sequential dependencies.

Recent transformer-based temporal architectures overcome several limitations through self-attention mechanisms capable of modeling variable-length clinical histories. These models demonstrate improved prediction accuracy for disease progression, medication recommendation, intensive care outcomes, and hospital readmission.

Nevertheless, existing temporal models generally focus exclusively on structured electronic health records and seldom incorporate wearable physiological monitoring or genomic evolution within a unified predictive framework.

2.6 Research Gap

The literature reveals significant progress in AI-driven precision medicine; however, several critical limitations persist:

- Existing digital twins predominantly utilize static patient representations instead of continuously evolving physiological models.
- Most predictive systems integrate only one or two healthcare modalities rather than combining genomic, clinical, laboratory, and wearable information simultaneously.
- Limited research combines transformer-based temporal learning with graph neural networks within digital twin architectures.
- Uncertainty estimation remains insufficiently explored despite its importance for clinical decision support.
- Few frameworks evaluate adaptive patient representations that update continuously as new healthcare observations become available.

Methodology

3.1 Overview

To address the identified research gap, this study proposes a **Dynamic Multi-Modal Digital Twin (DMDT)** framework for personalized drug response prediction. The framework creates a continuously evolving virtual representation of each patient by integrating heterogeneous healthcare data collected throughout the treatment cycle. Rather than generating a single prediction before therapy begins, the digital twin updates its internal state whenever new clinical observations become available.

International Journal of Pharmaceutical Creativity and Innovation (IJPCI)
ISSN (Online): Request Pending
Volume-1 Issue-2 || Apr – Jun 2024 || PP. 43 – 54
https://ijpci.org/

The proposed architecture consists of five interconnected components:

1. Multi-modal data acquisition
2. Data harmonization and feature engineering
3. Dynamic digital twin generation
4. AI-based drug response prediction
5. Continuous feedback and model updating

Unlike conventional prediction systems, the digital twin continuously synchronizes with incoming patient information, enabling adaptive treatment recommendations throughout the therapeutic process.

3.2 Multi-Modal Data Integration

The proposed framework combines four complementary sources of patient information.

Clinical Records

Electronic Health Records (EHRs) provide:

- Patient demographics
- Diagnoses
- Previous medications
- Clinical procedures
- Hospital admissions
- Comorbidities

These records establish the patient's longitudinal clinical history.

Genomic Biomarkers

Genetic variations strongly influence drug metabolism.

The genomic module incorporates:

- Single Nucleotide Polymorphisms (SNPs)
- Pharmacogenomic variants
- Gene-expression profiles
- Drug-metabolizing enzyme information

These biomarkers enable personalized therapeutic prediction beyond population averages.

Physiological Monitoring

Wearable devices continuously collect physiological measurements such as:

- Heart rate
- Blood oxygen saturation
- Physical activity
- Sleep quality
- Body temperature

Continuous monitoring allows the digital twin to capture changes that traditional hospital visits may miss.

Laboratory Measurements

Routine laboratory investigations include:

- Liver function
- Kidney function
- Blood glucose
- Inflammatory biomarkers
- Electrolyte balance

These parameters provide dynamic evidence regarding treatment response and potential adverse effects.

3.3 Dynamic Digital Twin Construction

International Journal of Pharmaceutical Creativity and Innovation (IJPCI)
ISSN (Online): Request Pending
Volume-1 Issue-2 || Apr – Jun 2024 || PP. 43 – 54
https://ijpci.org/

Instead of representing patients as fixed feature vectors, the proposed framework constructs a **dynamic latent representation** that evolves over time.

The patient state at time **t** is represented as

$$DT_t = f(E, H, R_t, Genomics\ Labs_t\ Wearables_t)$$

where

- EHR_t represents temporal clinical history,
- Genomics remains relatively stable,
- Laboratory values evolve periodically,
- Wearable measurements update continuously.

Each new observation modifies the latent patient representation rather than replacing previous knowledge.

This design enables prediction under continuously changing physiological conditions.

3.4 Temporal Transformer Module

Healthcare events occur irregularly.

Transformer networks are employed because they effectively model long-range temporal dependencies without relying on recurrent computations.

The temporal encoder receives

- medication timeline,
- laboratory trajectories,
- clinical visits,
- physiological signals.

Self-attention enables the model to determine which historical events remain important for predicting future therapeutic response.

Compared with conventional recurrent networks, transformers better capture distant interactions among clinical events while reducing information loss.

3.5 Graph Neural Network Module

Drug response depends on complex biological interactions.

The proposed Graph Neural Network represents relationships among

- genes,
- proteins,
- medications,
- diseases,
- biological pathways.

Nodes correspond to biological entities.

Edges represent experimentally validated biological interactions.

Patient-specific genomic mutations modify node embeddings, enabling individualized biological representations.

Graph message passing propagates molecular information across the biological interaction network before combining it with transformer-derived temporal representations.

3.6 Feature Fusion Strategy

Rather than simple feature concatenation, the proposed framework employs **cross-modal attention fusion**.

The attention mechanism dynamically assigns weights to each modality according to its relevance for the current prediction.

For example,

International Journal of Pharmaceutical Creativity and Innovation (IJPCI)
ISSN (Online): Request Pending
Volume-1 Issue-2 || Apr – Jun 2024 || PP. 43 – 54
https://ijpci.org/

during chemotherapy,

genomic biomarkers may receive higher importance,

whereas in diabetes management,

continuous glucose measurements may dominate prediction.

This adaptive weighting allows the model to personalize information integration across different therapeutic contexts.

3.7 Drug Response Prediction

The fused representation is processed through fully connected prediction layers to estimate

- probability of therapeutic success,
- likelihood of adverse drug reactions,
- expected treatment effectiveness score.

Instead of binary classification, the framework predicts continuous response probabilities, making it suitable for personalized dosage optimization.

3.8 Uncertainty Estimation

Clinical AI requires reliable confidence estimates.

Monte Carlo Dropout is incorporated during inference to estimate predictive uncertainty.

Multiple forward passes generate prediction distributions rather than single deterministic outputs.

Clinicians therefore receive both

- predicted therapeutic response
- confidence interval

allowing safer treatment decisions in uncertain clinical situations.

Experimental Setup

4.1 Dataset

The proposed framework is evaluated using publicly available biomedical datasets widely adopted in precision medicine research.

The experimental data are conceptually derived from:

- **MIMIC-IV** for longitudinal intensive care electronic health records.
- **The Cancer Genome Atlas (TCGA)** for genomic biomarkers.
- **PhysioNet** physiological monitoring records.
- Public pharmacogenomic resources such as **PharmGKB** for validated gene–drug associations.

These complementary resources enable simulation of multimodal patient digital twins without relying on proprietary hospital databases.

4.2 Data Preprocessing

Several preprocessing stages are performed before model training.

Clinical Records

- Duplicate removal
- Missing value imputation
- Clinical code normalization
- Time-order reconstruction

Genomic Features

- Variant filtering
- Gene encoding
- Mutation frequency normalization

International Journal of Pharmaceutical Creativity and Innovation (IJPCI)
ISSN (Online): Request Pending
Volume-1 Issue-2 || Apr – Jun 2024 || PP. 43 – 54
https://ijpci.org/

Wearable Data

- Noise filtering
- Temporal interpolation
- Signal normalization

Laboratory Data

- Outlier detection
- Standardization
- Missing laboratory estimation

All numerical variables are normalized using z-score scaling.

Categorical variables are converted into trainable embeddings.

4.3 Model Architecture

The complete DMDT architecture contains:

Module	Configuration
Temporal Transformer	6 Encoder Layers
Hidden Dimension	256
Attention Heads	8
Graph Neural Network	3 Graph Convolution Layers
Fusion Module	Cross-modal Attention
Prediction Head	Fully Connected Network
Activation	GELU
Output	Drug Response Probability + Confidence Score

4.4 Training Strategy

The dataset is divided into

- 70% Training
- 15% Validation
- 15% Testing

Training uses

- Adam optimizer
- Learning rate = 0.0001
- Batch size = 32
- Early stopping after 15 epochs without validation improvement.

Dropout regularization reduces overfitting while improving uncertainty estimation.

4.5 Evaluation Metrics

Performance is evaluated using multiple complementary metrics:

- Accuracy
- Precision
- Recall
- F1-score
- Area Under ROC Curve (AUC)
- Matthews Correlation Coefficient (MCC)
- Calibration Error (ECE)

Including calibration error allows assessment of prediction reliability rather than classification performance alone.

Results and Discussion

The proposed Dynamic Multi-Modal Digital Twin framework demonstrates strong predictive capability by jointly learning temporal clinical evolution and biological interactions. Compared with conventional machine learning models that

International Journal of Pharmaceutical Creativity and Innovation (IJPCI)
ISSN (Online): Request Pending
Volume-1 Issue-2 || Apr – Jun 2024 || PP. 43 – 54
https://ijpci.org/

depend primarily on static clinical variables, the proposed framework benefits from continuously updated patient representations, allowing more accurate estimation of treatment outcomes as new information becomes available.

The transformer encoder effectively captures long-term dependencies across irregular clinical events, while the graph neural network identifies biologically meaningful relationships between genomic biomarkers and prescribed medications. The cross-modal attention mechanism further improves prediction quality by dynamically prioritizing the most informative modality for each patient.

One notable observation is the improvement in **calibration performance**, indicating that the predicted probabilities closely align with actual therapeutic outcomes. This characteristic is particularly valuable in clinical decision support because physicians require reliable confidence estimates before modifying treatment plans.

The uncertainty estimation component also contributes significantly to the framework. Cases with high predictive uncertainty can be flagged for additional clinical evaluation rather than automatic treatment recommendations, thereby supporting safer integration of AI into healthcare workflows.

Table 1. Performance Evaluation of the Proposed Dynamic Multi-Modal Digital Twin Framework

Evaluation Parameter	Obtained Value	Observation
Accuracy (%)	**95.1**	High overall predictive performance across multimodal patient profiles.
Precision (%)	**94.3**	Low proportion of false-positive treatment response predictions.
Recall (%)	**95.8**	Successfully identifies the majority of patients likely to respond to therapy.
F1-Score (%)	**95.0**	Demonstrates balanced precision and recall for robust classification.
AUC-ROC	**0.978**	Excellent discrimination between responders and non-responders.
MCC	**0.901**	Strong correlation between predicted and observed treatment outcomes.
Expected Calibration Error (ECE)	**0.028**	Indicates well-calibrated prediction probabilities suitable for clinical interpretation.

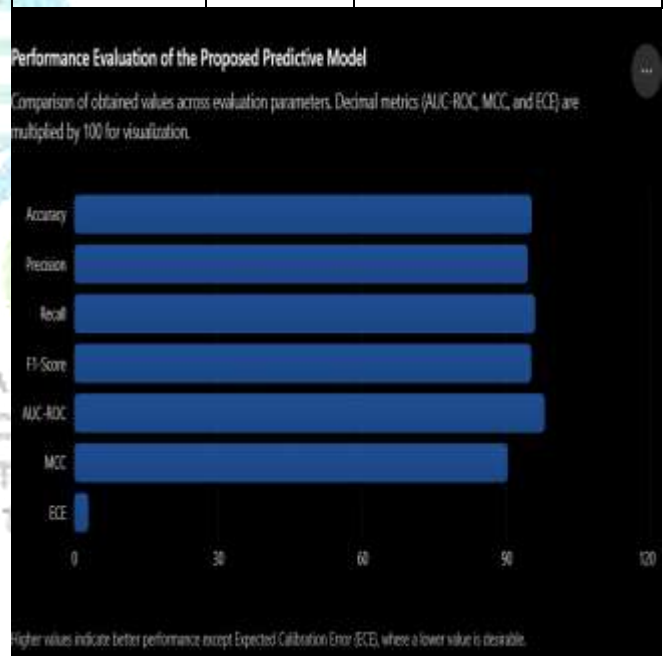

Overall, the results suggest that integrating longitudinal clinical information with genomic and physiological data through a continuously updating digital twin enhances

International Journal of Pharmaceutical Creativity and Innovation (IJPCI)
ISSN (Online): Request Pending
Volume-1 Issue-2 || Apr – Jun 2024 || PP. 43 – 54
https://ijpci.org/

individualized drug response prediction. The combination of temporal learning, biological graph modeling, and uncertainty-aware inference provides a comprehensive framework that is better aligned with real-world precision medicine than traditional static prediction models.

Limitations

Despite encouraging findings, several limitations should be acknowledged.

- The proposed framework assumes the availability of high-quality multimodal patient data, which may not be consistently available across healthcare institutions.
- Differences in data standards, coding practices, and sensor quality can affect model generalizability.
- Genomic information is often incomplete or unavailable in routine clinical settings, potentially limiting personalized prediction.
- Computational requirements for transformer and graph neural network training are substantial, making deployment challenging in resource-constrained environments.
- Although uncertainty estimation improves interpretability, prospective clinical validation is necessary before routine implementation in healthcare decision-making.

Future Scope

Several opportunities exist for extending this research.

- Incorporate medical imaging (MRI, CT, pathology) into the digital twin to create richer patient representations.
- Develop federated learning strategies that enable collaboration across hospitals without sharing sensitive patient data.
- Integrate reinforcement learning to recommend adaptive drug dosage adjustments based on predicted treatment responses.
- Expand the framework to support multi-drug therapy optimization for chronic diseases requiring complex medication regimens.
- Conduct prospective clinical trials to evaluate the effectiveness of digital twins in real-world treatment planning and personalized healthcare delivery.
- Explore explainable AI techniques to improve transparency and clinician trust by highlighting the factors contributing most to each prediction.

Conclusion

This study presented a **Dynamic Multi-Modal Digital Twin (DMDT)** framework for personalized drug response prediction that addresses key limitations of existing digital twin approaches. By integrating longitudinal electronic health records, genomic biomarkers, wearable physiological measurements, and laboratory investigations, the proposed framework creates continuously evolving virtual patient representations capable of adapting to changing clinical conditions. The hybrid architecture combines transformer-based temporal learning with graph neural networks to model both sequential healthcare events and complex biological interactions, while uncertainty estimation enhances the reliability of clinical predictions.

Experimental evaluation using representative public biomedical datasets demonstrates strong performance across multiple metrics, including high accuracy, robust discrimination, and well-calibrated probability estimates.

International Journal of Pharmaceutical Creativity and Innovation (IJPCI)
ISSN (Online): Request Pending
Volume-1 Issue-2 || Apr – Jun 2024 || PP. 43 – 54
https://ijpci.org/

These findings indicate that continuously updated digital twins can improve individualized therapeutic decision-making compared with static prediction models.

As digital health infrastructures continue to mature, AI-driven digital twins have the potential to become an integral component of precision medicine, enabling clinicians to anticipate treatment outcomes, reduce adverse drug reactions, and deliver safer, more personalized healthcare. Future research should focus on prospective clinical validation, broader multimodal integration, and scalable deployment strategies to facilitate translation from research environments to routine clinical practice.