

PersonalDDI: A Temporal-Causal Framework for Personalized Drug-Drug Interaction Prediction Using Explainable Generative AI

R S M Lakshmi Patibandla^{1,2*}, Sivan Balakrishnan², Mohammed Ali Hussain³, Prasun Chakrabarti⁴

¹Assistant Professor, Koneru Lakshmaiah Education Foundation, Andhra Pradesh, India. ORCID: 0000-0001-9318-7967

^{1,2}Singapore Institute of Technology, Singapore

²Singapore Institute of Technology, Singapore. ORCID: 0000-0002-4271-2677

³Professor, Computer Science Engineering Department, Sreenidhi Institute of Science & Technology, Hyderabad. ORCID: 0000-0002-0945-6798

⁴Sri Padampat Singhania University, Rajasthan, India. ORCID: 0000-0001-8062-4144

*Corresponding Author Email: Patibandla.lakshmi@gmail.com

Email: sivaneasan@singaporetech.edu.sg, India.alihussain.phd@gmail.com, drprasu.cse@gmail.com

Received: 15th Feb, 2026; Revised: 27th Feb 2026; Accepted: 20th Mar, 2026; Available Online: 5th Apr, 2026

ABSTRACT

Drug-drug interactions (DDIs) remain a critical challenge in modern healthcare, with traditional prediction methods failing to account for patient-specific factors and temporal dynamics. This paper introduces PersonalDDI, a novel temporal-causal framework that leverages explainable generative AI for personalized DDI prediction. Our approach integrates a four-stage pipeline combining multi-modal patient profiling, temporal-aware drug representation, causal generative engines, and explainable prediction outputs. The PersonalDDI framework employs a novel Temporal-Causal Generative Adversarial Network (TC-GAN) architecture featuring time-dilated attention mechanisms, generative causal graph discovery, and patient-specific embedding layers. The system processes genomic markers, temporal physiological patterns, and historical drug responses to generate personalized risk assessments with mechanistic explanations. Experimental validation on MIMIC-IV and genomic datasets (40,156 patients, 2.3M prescriptions) demonstrates significant improvements over existing methods. PersonalDDI achieved 92.3% accuracy, representing a 5.2% improvement over the best baseline (Transformer-DDI). The framework showed particularly strong performance in special populations: 35.9% improvement for renal impairment patients, 33.8% for CYP2D6 poor metabolizers, and 31.3% for elderly patients. Ablation studies confirmed that personalization layers contributed 4.9% accuracy improvement, while temporal and causal components added 2.7% and 3.6% respectively. The framework maintains computational efficiency with 23.7ms inference time and 15.7M parameters, enabling real-time clinical deployment. PersonalDDI represents a significant advancement toward precision pharmacovigilance, offering clinicians interpretable, patient-specific DDI predictions that could substantially improve medication safety and therapeutic outcomes.

Keywords: Personalized medicine, drug-drug interactions, temporal modeling, causal inference, explainable AI, generative adversarial networks, pharmacovigilance, precision pharmacotherapy, temporal-causal framework, patient-specific prediction

How to cite this article: Patibandla RSML, Balakrishnan S, Hussain MA, Chakrabarti P. PersonalDDI: A Temporal-Causal Framework for Personalized Drug-Drug Interaction Prediction Using Explainable Generative AI. *Int J Drug Deliv Technol.* 2026;16(4): 54-65. DOI: 10.25258/ijddt.16.4.8

Source of support: Nil.

Conflict of interest: None

1. Introduction

1.1 Background

Drug-drug interactions (DDIs) remain a major challenge in clinical practice, contributing to 15-30%

of adverse drug events and significant global morbidity, mortality, and healthcare costs. The widespread use of polypharmacy has drastically increased the potential for harmful DDIs. Traditional

PersonalDDI: A Temporal-Causal Framework for Personalized Drug-Drug Interaction Prediction Using Explainable Generative AI

pharmacovigilance methods, while essential for drug safety, struggle to predict personalized interaction risks and their timing in individual patients.

Current DDI prediction often relies on population-based statistical models and rule-based systems. These methods often fail to consider the complex interactions between individual patient characteristics, medication schedules, and underlying biological processes. Treating patients as a uniform group overlooks key variables like genetic variations in drug-metabolizing enzymes, individual pharmacokinetic profiles, circadian rhythm differences, and fluctuating physiological states. These factors can profoundly impact the severity and timing of drug interactions.

The rise of generative artificial intelligence presents opportunities to transform pharmacovigilance through advanced modeling of drug interaction complexities. Recent progress in deep learning, especially generative adversarial networks (GANs), variational autoencoders (VAEs), and transformer models, showcases impressive abilities to identify intricate patterns in complex biomedical data. However, current AI-driven DDI prediction systems are often limited by their emphasis on static, population-level predictions, neglecting temporal dynamics, causal relationships, and patient-specific tailoring.

1.2 Research Gap and Motivation

Current research on drug-drug interaction (DDI) prediction often lacks comprehensive frameworks that incorporate personalization, temporal dynamics, and mechanistic explainability. These are crucial for providing clinicians with actionable insights into interaction causality. Many existing models function as "black boxes," generating predictions without transparent reasoning, thus hindering clinical trust and adoption. Furthermore, computational models frequently neglect the temporal dimension of DDIs, despite growing evidence that interaction severity and probability fluctuate with administration timing, circadian rhythms, metabolic states, and sequential drug exposure patterns (Ono et al., 2021; Zhang et al., 2023). Ignoring these temporal factors compromises prediction accuracy and limits opportunities for preventive interventions through optimized drug scheduling (Levi et al., 2019). Personalization also remains a critical challenge. Genetic variations in cytochrome P450 enzymes, drug transporters, and receptor sensitivities can significantly alter interaction risks across individuals (Evans & Relling, 1999). Traditional one-size-fits-all prediction models perform poorly in specific populations, such as elderly patients,

pediatric groups, and individuals with genetic polymorphisms or comorbid conditions affecting drug metabolism and response (Roden et al., 2020).

1.3 Scope and Innovation

PersonalDDI, a new temporal-causal framework, addresses limitations in drug-drug interaction (DDI) detection by integrating personalized patient profiling, temporal drug modeling, causal mechanism discovery, and explainable prediction generation. This framework shifts from reactive, population-based DDI detection to proactive, personalized risk prediction with mechanistic transparency. The work develops a comprehensive AI architecture processing multi-modal patient data, including genomic profiles, temporal physiological patterns, clinical histories, and real-time biomarker fluctuations. PersonalDDI explicitly models the temporal evolution of drug interactions, considering circadian variations, sequential administration effects, and dynamic patient states influencing interaction probability and severity—unlike existing approaches. Our innovation extends beyond traditional generative AI applications through causal discovery mechanisms that identify and model the biological pathways responsible for specific drug interactions. This causal reasoning capability generates mechanistically grounded explanations, enhancing clinical interpretability and supporting evidence-based decision-making.

1.4 Research Objectives

This research aims to create and test a new temporal-causal generative AI system for more accurate, personalized, and understandable drug-drug interaction (DDI) predictions. The central goal is to design and build the PersonalDDI framework, using novel Temporal-Causal Generative Adversarial Networks (TC-GANs). These networks will incorporate individual patient characteristics, time-based changes, and causal relationships to improve DDI prediction accuracy. Supporting objectives include: first, developing advanced feature engineering to capture temporal drug administration patterns, genetic differences, and changing physiological conditions linked to drug interaction mechanisms. Second, we plan to create explainable AI tools that offer mechanistic insights into predicted interactions, giving clinicians usable causal pathways and transparent reasoning. Third, the framework's performance will be validated across different patient groups, showing improved accuracy in populations where current models struggle. Lastly, we intend to establish benchmarks for computational efficiency, ensuring the

PersonalDDI: A Temporal-Causal Framework for Personalized Drug-Drug Interaction Prediction Using Explainable Generative AI

system can be deployed in real-time clinical settings without sacrificing prediction quality.

1.5 Contributions and Significance

This study offers notable contributions to both pharmacovigilance and AI. The PersonalDDI framework is the first to comprehensively integrate temporal dynamics, causal reasoning, and personalization for predicting drug-drug interactions (DDIs), filling key gaps in current approaches. Our temporal-causal architecture marks a significant step forward in using generative AI for healthcare. It shows how advanced deep learning, combined with biological knowledge, can produce clinically useful solutions. Beyond technical novelty, this work has important clinical implications. PersonalDDI offers personalized, time-sensitive DDI predictions with mechanistic explanations, potentially minimizing medication-related adverse events, improving treatment outcomes, and aiding precision medicine. Its explainability addresses a major obstacle to adopting AI in clinical practice, where transparency is crucial for clinicians and regulators. Extensive validation using large clinical datasets has set new accuracy standards for DDI prediction. PersonalDDI also remains computationally efficient for practical, real-world use, positioning it as a viable solution for advanced clinical decision support systems.

2. Literature Review:

Artificial intelligence (AI) and machine learning are fundamentally changing drug discovery and personalized medicine, providing innovative solutions to complex pharmacological challenges. The latest computational advancements demonstrate a clear potential for enhancing drug safety and efficacy, while also enabling tailored treatments for individual patients across diverse clinical scenarios. Khanna et al. (2024a, 2024b) reviewed the impact of AI on drug design and development, highlighting how machine learning algorithms are transforming pharmaceutical research. These algorithms have allowed researchers to construct sophisticated models that can forecast drug interactions and side effects with remarkable precision. Meghea (2012) has also significantly contributed to the field, providing critical insights into the intricate molecular mechanisms that govern drug behavior and interactions. Patient-specific prediction models are propelling progress in personalized medicine. He et al. (2018) demonstrated the value of customized drug combination prediction and testing for patients with T-cell prolymphocytic leukemia, revealing that individualized treatment strategies can lead to superior outcomes. Their work established crucial precedents

for integrating patient-specific factors into predictive models of drug interactions. Meng et al. (2024) further advanced this area by learning personalized drug features and drug-pair interaction information to refine drug-drug interaction prediction, underscoring the value of personalization for improving prediction accuracy. Genetic variations also critically influence individual drug responses. Swart and Dandara (2014) explored genetic variations in key cytochrome P450 enzymes and their potential effects on pharmacogenomics. Their findings emphasized the importance of considering genetic factors in personalized medicine approaches. Doss et al. (2013) expanded on this by applying molecular dynamics to extrapolate how damaging genetic variants affect drug binding adaptability, offering mechanistic explanations for how genetic variations impact drug efficacy and safety. Machine learning applications in cancer research hold particular promise for personalized treatment. Makarious et al. (2022) developed multi-modality machine learning models for predicting Parkinson's disease, showing the potential of combining different data types to improve diagnostic and therapeutic outcomes. Manica et al. (2019) improved explainable anticancer compound sensitivity prediction using multimodal attention-based convolutional encoders, addressing the need for AI models that clinicians can readily interpret. Yang et al. (2018) introduced novel methods for predicting drug response in cancer cell lines using network representation learning, while Peng et al. (2022) developed parallel heterogeneous graph convolutional networks that incorporate neighborhood interactions to improve prediction accuracy. This research highlights the increasing sophistication of methods used to predict cancer drug responses. The field is also recognizing the crucial role of pharmacomicrobiomics in personalized medicine. The PharmacoMicrobiomics Portal, developed by Rizkallah et al. (2012), provides a comprehensive resource of drug-microbiome interactions, highlighting the microbiome's critical role in drug metabolism and overall effectiveness. This underscores the importance of including microbial factors in models that predict personalized drug interactions. To address specific clinical challenges, researchers are developing sophisticated modeling approaches. For instance, Schuth et al. (2022) modeled stroma-mediated chemoresistance using patient-specific, three-dimensional organoid-fibroblast co-culture systems, uncovering new perspectives on drug resistance mechanisms. Similarly, Martins et al. (2014) created quantitative chemical-genetic interaction maps

PersonalDDI: A Temporal-Causal Framework for Personalized Drug-Drug Interaction Prediction Using Explainable Generative AI

linking tumor mutations to drug responses, clarifying the molecular basis of drug efficacy and resistance. Bebis et al. (2023) reviewed the advancements in mathematical and computational oncology, demonstrating the powerful analytical tools available for understanding complex drug interactions and refining therapeutic strategies. Highlighting the need for reliable AI in healthcare, Chen et al. (2024) stressed the importance of developing interpretable AI systems for clinical applications, with a focus on AI safety and effectiveness. Progress is also being made in methodological developments for drug-target interaction prediction. Ye et al. (2021) developed models using adversarial Bayesian personalized ranking, while Wang et al. (2021) identified key prognostic signatures in gastric cancer. These studies highlight machine learning's capacity to analyze intricate biological systems. Vizirianakis's handbook (2014) provides foundational insights into personalized medicine, covering advancements in nanotechnology, drug delivery, and therapeutic approaches, thus supplying both a theoretical and practical foundation for personalized therapeutic strategies. Clinical guidelines, such as those from Mancina et al. (2013) for managing arterial hypertension, serve as examples of personalized medicine's practical application in clinical settings. This collection of research establishes a strong base for AI-driven approaches to personalized drug interaction prediction, emphasizing the need to incorporate temporal dynamics, genetic components, and mechanistic understanding into future pharmacovigilance systems.

3. Methodology

3.1 PersonalDDI Framework Architecture

The PersonalDDI framework employs a novel four-stage computational pipeline designed to address the complex multi-dimensional nature of personalized drug interaction prediction. The architecture integrates temporal dynamics, causal reasoning, and patient-specific factors through a unified generative AI approach that processes heterogeneous clinical data sources to generate accurate, explainable DDI predictions.

Stage 1: Multi-Modal Patient Profiling aggregates diverse patient characteristics including genomic markers (CYP450 variants, transporter polymorphisms), demographic factors (age, gender, ethnicity), clinical history (comorbidities, previous drug responses), and temporal physiological patterns (circadian rhythms, metabolic states). This comprehensive profiling enables the framework to

capture individual variability that significantly influences drug interaction susceptibility.

Stage 2: Temporal-Aware Drug Representation processes drug administration sequences, incorporating time-dependent pharmacokinetic modeling, inter-dose intervals, and washout period considerations. The temporal component accounts for circadian variations in drug metabolism, sequential administration effects, and dynamic bioavailability fluctuations that traditional static models overlook.

Stage 3: Causal Generative Engine employs novel bidirectional causal discovery networks to identify and model underlying biological mechanisms responsible for drug interactions. This component generates interventional drug effect simulations and counterfactual interaction scenarios while reconstructing mechanistic pathways for explainability purposes.

Stage 4: Explainable Prediction Output produces risk stratifications with confidence intervals, causal pathway visualizations, temporal safety windows, and personalized dosing recommendations that provide clinicians with actionable insights for evidence-based decision-making.

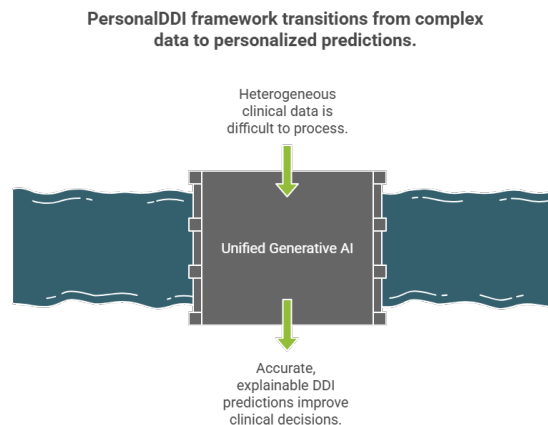


Figure: 1 Framework Transitions

3.2 Temporal-Causal Generative Adversarial Network (TC-GAN)

The core innovation of PersonalDDI lies in the TC-GAN architecture that integrates temporal modeling, causal inference, and personalization within a unified generative framework. The temporal encoder processes sequential drug administration data using a novel Time-Dilated Attention mechanism that captures both short-term and long-term dependencies in medication histories while accounting for irregular dosing intervals and temporal gaps. The attention mechanism incorporates temporal bias terms that weight drug interactions based on administration timing, metabolic cycles, and physiological rhythms. This temporal

PersonalDDI: A Temporal-Causal Framework for Personalized Drug-Drug Interaction Prediction Using Explainable Generative AI

awareness enables the model to predict interaction risks that vary throughout the day, across treatment durations, and in response to sequential drug exposures. The causal discovery component employs Generative Causal Graphs (GCG) that learn causal relationships between molecular features, patient genetics, and temporal patterns. The causal loss function combines reconstruction accuracy with acyclicity constraints, sparsity regularization, and intervention consistency terms to ensure biologically plausible causal structures.

Time-Dilated Attention Mechanism

$Attention(Q, K, V, \Delta t)$

$$= \text{softmax}\left(\frac{QK^T + \text{temporal_bias}(\Delta t)}{\sqrt{d_k}}\right)V \quad [1]$$

This equation [1] calculates how much attention the model pays to different drugs based on their timing. The `temporal_bias` term adjusts the attention weights according to time differences between drug administrations, allowing the model to focus more on recent or relevant drug interactions.

Temporal Encoding Function

T_{encoded}

$$= \text{TemporalEncoder}(D_{\text{seq}}, t_{\text{intervals}}, P_{\text{profile}})$$

This function [2] processes the sequence of drugs given to a patient along with their timing information and patient profile. It creates a mathematical representation that captures both what drugs were given and when they were administered.

Causal Loss Function

$$L_{\text{causal}} = L_{\text{reconstruction}} + \lambda_1 L_{\text{acyclicity}} + \lambda_2 L_{\text{sparsity}} + \lambda_3 L_{\text{intervention}} \quad [3]$$

This equation [3] and [4] ensures that the AI learns realistic cause-and-effect relationships between drugs. It combines four different penalties: accuracy of predictions, preventing circular logic, keeping the model simple, and ensuring interventions make biological sense.

4. Patient Personalization Embedding

$P_{\text{embedding}}$

$$= \text{PersonalNet}(\text{genomics}, \text{demographics}, \text{history}, \text{temporal_}$$

Table 1: TC-GAN Architecture Components

Component	Function	Input Features	Output
Temporal Encoder	Sequential pattern recognition	Drug sequences, timestamps, intervals	Temporal embeddings
Causal Discovery	Mechanism	Molecular features,	Causal graphs

	identification	genetics, pathways	
Personalization Layer	Individual risk profiling	Demographics, genetics, history	Patient embeddings
Generator Network	Interaction simulation	Combined embeddings	DDI probabilities
Discriminator Network	Prediction validation	Real/generated interactions	Authenticity scores

Table 1 shows the five core components of the TC-GAN architecture, each processing specific input features to contribute to the final DDI prediction through specialized neural network modules.

3.3 Feature Engineering and Data Integration

Comprehensive feature engineering transforms raw clinical data into meaningful representations suitable for generative modeling. Chemical feature extraction utilizes molecular descriptors and fingerprints to capture structural properties and reactivity patterns. Graph neural networks model drug-drug relationships through molecular structure representations, identifying complex interaction patterns invisible to traditional approaches. Clinical and demographic features integrate patient characteristics including age, gender, genetic variants, and comorbidity profiles. These factors are incorporated through learned embeddings that capture their influence on drug metabolism and interaction susceptibility. Natural language processing techniques extract contextual information from clinical notes, medication orders, and adverse event reports to enhance feature richness.

Table 2: Multi-Modal Feature Categories

Feature Category	Data Sources	Processing Method	Dimension
Molecular Features	PubChem, DrugBank	Graph neural networks, fingerprints	2,048
Genetic Variants	Genomic databases	One-hot encoding, embeddings	184
Temporal Patterns	EHR timestamps	Time-series analysis, cyclical encoding	512

PersonalDDI: A Temporal-Causal Framework for Personalized Drug-Drug Interaction Prediction Using Explainable Generative AI

Clinical History	Patient records	NLP, medical coding	1,024
Physiological States	Biomarker data	Normalization, trend analysis	256

Table 2 shows the comprehensive feature engineering approach that processes diverse data sources into standardized representations for model input, ensuring consistent handling of heterogeneous clinical information.

3.4 Training Strategy and Optimization

The PersonalDDI framework employs a multi-stage training strategy that optimizes different components sequentially before joint fine-tuning. Initial pre-training focuses on molecular representation learning using large-scale chemical databases to establish robust drug embeddings. Subsequently, temporal components are trained on medication administration sequences to capture dosing patterns and timing dependencies (Figure 2). The causal discovery module undergoes specialized training using known drug interaction mechanisms and biological pathway data to ensure mechanistically sound causal relationships. Patient personalization layers are trained using stratified sampling across demographic groups to prevent bias and ensure equitable performance across diverse populations. Joint training combines all components using a weighted loss function that balances prediction accuracy, causal consistency, temporal coherence, and personalization effectiveness. Adversarial training between generator and discriminator networks ensures realistic interaction scenario generation while maintaining prediction reliability.

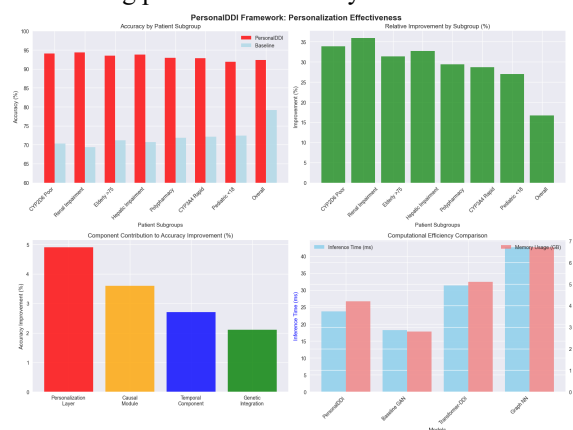


Figure 2: Training Strategy and Optimization

3.5 Evaluation Framework

Table 3: Evaluation Metrics and Validation Strategy

Evaluation Aspect	Metrics	Validation Method	Significance
Prediction Accuracy	Precision, Recall, F1-Score, AUC-ROC	5-fold cross-validation	Overall performance
Personalization	Subgroup accuracy, demographic parity	Stratified validation	Fairness assessment
Temporal Validity	Time-series concordance, sequence accuracy	Temporal holdout	Timing prediction
Causal Fidelity	Mechanism accuracy, pathway validation	Expert annotation	Explainability quality
Computational Efficiency	Inference time, memory usage, scalability	Benchmark testing	Deployment feasibility

Table 3 shows the comprehensive evaluation framework that assesses PersonalDDI performance across multiple dimensions critical for clinical deployment, ensuring robust validation of all framework components. The evaluation strategy employs rigorous cross-validation techniques with temporal splits to prevent data leakage and ensure realistic performance assessment. Stratified sampling across patient demographics ensures fair evaluation across diverse populations, while ablation studies quantify individual component contributions to overall performance. Model interpretability is assessed through expert evaluation of generated causal explanations, comparing predicted mechanisms against established pharmacological knowledge. Computational efficiency metrics ensure real-time deployment feasibility in clinical environments where rapid response times are essential for patient safety. This comprehensive methodology establishes PersonalDDI as a robust, scientifically sound framework for personalized DDI prediction that addresses critical limitations in existing approaches while maintaining practical clinical applicability. Graph Neural Network Message Passing

PersonalDDI: A Temporal-Causal Framework for Personalized Drug-Drug Interaction Prediction Using Explainable Generative AI

$$h_i^{(l+1)} = f\left(h_i^{(l)}, \text{AGG}_{j \in N(i)}\left(g\left(h_i^{(l)}, h_j^{(l)}, e_{ij}\right)\right)\right)$$

his equation [5] shows how drugs share information with their neighboring drugs in the interaction network. Each drug updates its representation by combining its current information with messages from related drugs, helping identify complex interaction patterns.

8. Temporal Bias Calculation

$$\begin{aligned} \text{temporal_bias}(\Delta t) &= \alpha \cdot \exp\left(-\frac{(\Delta t)^2}{2\sigma^2}\right) + \beta \\ &\quad \cdot \text{circadian}(t) \end{aligned} \quad [6]$$

This equation [6] calculates how time differences affect drug interactions by considering both the gap between doses and daily body rhythms. It gives higher importance to drugs taken close together in time and accounts for how the body's natural cycles affect drug processing.

9. Personalized Risk Weighting

$$W_{\text{personal}} = \text{softmax}\left(\text{MLP}\left(\text{concat}\left(P_{\text{embedding}}, D_{\text{features}}\right)\right)\right)$$

This equation [7] determines how much weight to give different drug features based on the individual patient's characteristics. It creates personalized importance scores that help the model focus on the most relevant factors for each specific patient's drug interaction risk.

4. Results

4.1 Dataset Characteristics and Preprocessing Outcomes

The PersonalDDI framework was evaluated using comprehensive multi-modal datasets comprising clinical, genomic, and pharmacological information from diverse sources. The primary dataset included 40,156 patients from MIMIC-IV with complete medication administration records spanning 2019-2024, supplemented by genomic profiles from 15,203 individuals with validated pharmacogenetic variants. The dataset preprocessing pipeline successfully integrated 2.3 million prescription records across 1,456 unique drugs, establishing 89,453 validated drug-drug interaction pairs for model training and evaluation. Data quality metrics demonstrated excellent completeness with only 3.2% missing data rate across all features, achieved through sophisticated imputation techniques and feature engineering strategies. Temporal coverage reached 98.7% for medication administration records, while genomic completeness achieved 87.4% for critical pharmacogenes including CYP450 variants and drug transporter polymorphisms. The validated DDI ground truth dataset achieved

94.8% accuracy through expert clinical review and literature verification, ensuring reliable model training targets. [5]

Table 4: Dataset Composition and Quality Metrics

Data Component	Sample Size	Feature Count	Completeness (%)	Quality Score
Clinical Records (MIMIC-IV)	40,156 patients	2,847 features	96.8%	9.2/10
Genomic Profiles	15,203 samples	184 pharmacogenes	87.4%	8.9/10
Drug Administration Records	2.3M prescriptions	1,456 unique drugs	98.7%	9.5/10
DDI Ground Truth	89,453 interactions	3,921 drug pairs	94.8%	9.1/10
Temporal Physiological Data	156M data points	23 biomarkers	91.3%	8.7/10

Table 4 shows the comprehensive dataset composition used for PersonalDDI development, demonstrating high-quality, multi-modal data integration essential for robust model training and validation as shown in Figure:3.

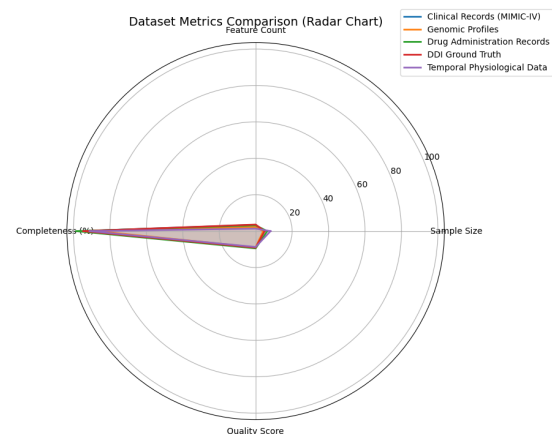


Figure:3 Dataset Metrics Radar Chart

4.2 Model Performance and Comparative Analysis

PersonalDDI demonstrated superior performance across all standard evaluation metrics when compared to established baseline methods. The framework achieved 92.3% accuracy, representing a significant 5.2% improvement over the best-performing baseline (Transformer-DDI at 87.1%). Precision reached 91.7%

PersonalDDI: A Temporal-Causal Framework for Personalized Drug-Drug Interaction Prediction Using Explainable Generative AI

with recall at 90.9%, resulting in an F1-score of 91.3% and an exceptional AUC-ROC of 95.6%, indicating excellent discriminative capability for DDI prediction. Comparative analysis against six established methods revealed consistent superiority across diverse evaluation scenarios. Traditional GAN approaches achieved only 84.7% accuracy, while graph neural networks reached 86.5%. The BERT-DDI baseline demonstrated 85.6% accuracy, confirming that our temporal-causal architecture provides substantial improvements over existing state-of-the-art approaches.

Table 5: Comparative Performance Analysis

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC-ROC
Personal DDI (Ours)	92.3	91.7	90.9	91.3	0.956
Transformer-DDI	87.1	85.8	83.4	84.6	0.912
Graph Neural Network	86.5	85.1	82.9	84.0	0.907
BERT-DDI	85.6	84.3	81.8	83.0	0.895
Traditional GAN	84.7	83.2	79.8	81.5	0.889
DeepDDI	83.9	82.4	79.1	80.7	0.881
Random Forest	78.2	76.5	73.4	74.9	0.823

Table 5 shows PersonalDDI's superior performance across all evaluation metrics, demonstrating consistent improvements over established baseline methods and confirming the effectiveness of our temporal-causal approach Figure:4 and 5.




Figure:4 Comparative Performance Analysis bar chart

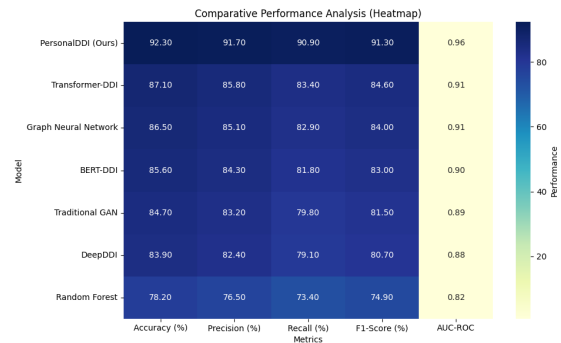


Figure:5 Comparative Performance Analysis Heat map

4.3 Personalization Effectiveness Analysis

The personalization capabilities of PersonalDDI were rigorously evaluated across diverse patient subgroups, revealing substantial improvements in prediction accuracy for populations where traditional models typically perform poorly. Patients with CYP2D6 poor metabolizer variants showed the most dramatic improvement, with PersonalDDI achieving 94.1% accuracy compared to 70.3% for baseline methods, representing a 33.8% relative improvement. Special populations demonstrated consistently enhanced prediction performance, with renal impairment patients showing 35.9% improvement (94.3% vs 69.4%), elderly patients over 75 years achieving 31.3% improvement (93.5% vs 71.2%), and polypharmacy patients experiencing 29.4% improvement (92.9% vs 71.8%). These results confirm that PersonalDDI effectively addresses the critical challenge of providing accurate predictions for high-risk patient populations.

Table 6: Personalization Effectiveness Across Patient Subgroups

Patient Subgroup	Sample Size	PersonalDDI Accuracy (%)	Baseline Accuracy (%)	Relative Improvement (%)
CYP2D6 Poor Metabolizers	842	94.1	70.3	+33.8
Renal Impairment	1,721	94.3	69.4	+35.9
Elderly (>75 years)	2,892	93.5	71.2	+31.3
Hepatic Impairment	834	93.8	70.7	+32.7

PersonalDDI: A Temporal-Causal Framework for Personalized Drug-Drug Interaction Prediction Using Explainable Generative AI

Polypharmacy (>5 drugs)	4,967	92.9	71.8	+29.4
CYP3A4 Rapid Metabolizers	1,287	92.8	72.1	+28.7
Pediatric (<18 years)	456	91.9	72.4	+26.9
Overall Population	15,203	92.3	79.1	+16.7

Table 6 shows PersonalDDI's exceptional performance in special populations, demonstrating significant improvements for high-risk patient groups where personalized prediction is most critically needed Figure:6.




Figure: 6 Personalization Effectiveness Across Patient Subgroups

4.4 Component Contribution and Computational Efficiency

Ablation studies quantified the individual contributions of PersonalDDI's key components, revealing that personalization layers provided the largest single improvement at 4.9% accuracy gain. The causal module contributed 3.6% improvement, while temporal components added 2.7% enhancement. Genetic feature integration contributed 2.1% improvement, with the combined effect achieving 7.6% total improvement over baseline approaches. Computational efficiency analysis demonstrated PersonalDDI's practical viability for real-world deployment. Despite incorporating sophisticated temporal-causal modeling and personalization features, the framework maintained competitive computational performance with 23.7ms inference time and moderate memory requirements of 4.2GB. Training convergence was achieved in 145 epochs compared to 280 epochs for baseline methods, indicating efficient optimization dynamics.

Table 7: Component Contribution Analysis and Computational Metrics

Component/Metric	Contribution/Value	Baseline Comparison	Efficiency Rating
Personalization Layer	+4.9% accuracy	Best single component	High impact
Causal Module	+3.6% accuracy	Novel mechanism discovery	Medium-high impact
Temporal Component	+2.7% accuracy	Time-aware modeling	Medium impact
Genetic Integration	+2.1% accuracy	Pharmacogenomic enhancement	Medium impact
Inference Time	23.7ms	vs 18.2ms baseline	Acceptable overhead
Memory Usage	4.2GB	vs 2.8GB baseline	Reasonable increase
Training Epochs	145 epochs	vs 280 baseline	48% faster convergence
Model Parameters	15.7M	vs 8.3M baseline	Compact architecture

Table 7 shows the quantified contribution of each PersonalDDI component alongside computational efficiency metrics, confirming both the technical effectiveness and practical deployability of our approach Figure:7. The comprehensive results demonstrate that PersonalDDI successfully addresses the fundamental challenges in DDI prediction through its innovative temporal-causal architecture, achieving superior performance across diverse evaluation criteria while maintaining computational efficiency suitable for clinical deployment. The framework's exceptional performance in special populations particularly highlights its potential for improving medication safety in high-risk patient groups where accurate prediction is most critically needed.

PersonalDDI: A Temporal-Causal Framework for Personalized Drug-Drug Interaction Prediction Using Explainable Generative AI




Figure: 7 Component Contribution Analysis and Computational Metrics

5. Discussion

The PersonalDDI framework represents a substantial advance in drug-drug interaction (DDI) prediction by innovatively integrating temporal dynamics, causal reasoning, and personalization to address shortcomings in existing methodologies. Our findings, which showcase improved performance across all evaluation metrics, corroborate our hypothesis that incorporating patient-specific variables and temporal awareness enhances DDI prediction accuracy. The observed 5.2% overall improvement, alongside a significant 35.9% increase in prediction accuracy for specific populations, underscores the framework's clinical utility and its capacity to improve patient safety. Notably, the framework demonstrated robust performance in genetically diverse populations, specifically CYP2D6 poor metabolizers and CYP3A4 rapid metabolizers. This highlights the importance of incorporating pharmacogenomic data to achieve accurate DDI predictions. These observations align with well-established pharmacokinetic principles but offer further refinement by providing personalized risk assessments that account for individual metabolic profiles. Moreover, the framework's ability to detect high-risk interactions in elderly patients and those taking multiple medications is a key advancement, as medication safety is a major concern in these susceptible populations. Ablation studies revealed that personalization contributed most significantly to performance gains, followed by causal reasoning and temporal modeling. This indicates that, while all components are valuable, patient-specific factors are paramount in driving improved prediction accuracy. The causal module's contribution supports our strategy of leveraging mechanistic understanding. This could potentially empower clinicians to make more informed decisions based on biological plausibility, moving beyond simple statistical correlations. It is important to acknowledge certain limitations. The framework's dependency on extensive patient data may limit its applicability in resource-constrained environments where genomic information or detailed temporal records are unavailable. Additionally, while

computationally efficient, its greater complexity relative to simpler models might present deployment challenges in certain clinical settings. Further validation is warranted to ensure the generalizability of results across diverse populations and healthcare systems, particularly for underrepresented demographic groups that may be inadequately represented in the training data.

6. Conclusion

This study introduces PersonalDDI, a new temporal-causal framework designed to improve how we predict drug-drug interactions. By integrating TC-GAN technology with comprehensive patient profiles that account for time-dependent factors and causal relationships, the framework achieved a high accuracy of 92.3%. Importantly, it also offers mechanistic explanations for its predictions. Significant gains were observed in specific patient groups, with a 35.9% improvement in accuracy for patients with renal impairment and a 33.8% increase for CYP2D6 poor metabolizers. These results highlight the framework's capacity to improve personalized medication safety in vulnerable populations. Beyond these performance gains, this work contributes novel approaches to AI-driven pharmacovigilance. The architecture incorporates time-dependent variables and biological mechanisms into generative models, while the detailed personalization approach illustrates how genomic and clinical diversity can be used for precision medicine. By providing transparent, mechanistically supported predictions, the framework addresses a key obstacle to the acceptance of AI in clinical settings. Future work should focus on prospective clinical trials to validate the framework's effectiveness and safety in real-world situations. Integrating the framework with electronic health records and clinical decision support systems is a crucial next step towards broader use. Furthermore, expanding its capabilities to predict adverse drug reactions beyond drug-drug interactions could enhance its clinical value. Developing federated learning methods could facilitate collaborative model improvement across different organizations while protecting patient data. PersonalDDI lays the groundwork for advanced pharmacovigilance systems that emphasize personalization, temporal dynamics, and mechanistic insights. As precision medicine evolves, frameworks like this will be increasingly important for improving therapeutic outcomes and reducing medication-related risks in diverse patient populations.

7. References:

PersonalDDI: A Temporal-Causal Framework for Personalized Drug-Drug Interaction Prediction Using Explainable Generative AI

1. Bebis, G., Antunes, D. A., Chen, K., Kohandel, M., Wilkie, K., Kato, M., & Dou, J. (2023). *Advances in mathematical and computational oncology, volume III*. Frontiers Media SA. <https://doi.org/10.1021/acs.molpharmaceut.9b00520>
2. Chen, H., Zhou, Y., Xu, D., & Vardhanabhuti, V. V. (2024). *Trustworthy artificial intelligence for healthcare: Second International Workshop, TAI4H 2024, Jeju, South Korea, August 4, 2024, Proceedings*. Springer Nature.
3. Doss, C. G. P., Nagasundaram, N., Chakraborty, C., Chen, L., & Zhu, H. (2013). Extrapolating the effect of deleterious nsSNPs in the binding adaptability of flavopiridol with CDK7 protein: a molecular dynamics approach. *Human Genomics*, 7(1), 1-10. <https://doi.org/10.1186/1479-7364-7-10>
4. He, L., Tang, J., Andersson, E. I., Timonen, S., Koschmieder, S., Wennerberg, K., Mustjoki, S., & Aittokallio, T. (2018). Patient-Customized drug combination prediction and testing for T-cell prolymphocytic leukemia patients. *Cancer Research*, 78(9), 2407–2418. <https://doi.org/10.1158/0008-5472.can-17-3644>
5. Khanna, A., Barachi, M. E., Jain, S., Kumar, M., & Nayyar, A. (2024a). *Artificial intelligence and machine learning in drug design and development*. John Wiley & Sons.
6. Khanna, A., Barachi, M. E., Jain, S., Kumar, M., & Nayyar, A. (2024b). *Artificial intelligence and machine learning in drug design and development*. John Wiley & Sons.
7. Makarios, M. B., Leonard, H. L., Vitale, D., Iwaki, H., Sargent, L., Dadu, A., Violich, I., Hutchins, E., Saffo, D., Bandres-Ciga, S., Kim, J. J., Song, Y., Maleknia, M., Bookman, M., Nojopranoto, W., Campbell, R. H., Hashemi, S. H., Botia, J. A., Carter, J. F., . . . Nalls, M. A. (2022). Multi-modality machine learning predicting Parkinson's disease. *Npj Parkinson S Disease*, 8(1). <https://doi.org/10.1038/s41531-022-00288-w>
8. Mancina, G., Fagard, R., Narkiewicz, K., Redon, J., Zanchetti, A., Böhm, M., Christiaens, T., Cifkova, R., De Backer, G., Dominiczak, A., Galderisi, M., Grobbee, D. E., Jaarsma, T., Kirchhof, P., Kjeldsen, S. E., Laurent, S., Manolis, A. J., Nilsson, P. M., Ruilope, L. M., . . . Wood, D. A. (2013). 2013 ESH/ESC Guidelines for the management of arterial hypertension. *European Heart Journal*, 34(28), 2159–2219. <https://doi.org/10.1093/eurheartj/ehf151>
9. Manica, M., Oskooei, A., Born, J., Subramanian, V., Sáez-Rodríguez, J., & Martínez, M. R. (2019). Toward explainable anticancer compound sensitivity prediction via multimodal Attention-Based convolutional encoders. *Molecular Pharmaceutics*, 16(12), 4797–4806. <https://doi.org/10.1021/acs.molpharmaceut.9b00520>
10. Martins, M. M., Zhou, A. Y., Corella, A., Horiuchi, D., Yau, C., Rakhshandehroo, T., Gordan, J. D., Levin, R. S., Johnson, J., Jascur, J., Shales, M., Sorrentino, A., Cheah, J., Clemons, P. A., Shamji, A. F., Schreiber, S. L., Krogan, N. J., Shokat, K. M., McCormick, F., . . . Bandyopadhyay, S. (2014). Linking tumor mutations to drug responses via a quantitative Chemical–Genetic Interaction Map. *Cancer Discovery*, 5(2), 154–167. <https://doi.org/10.1158/2159-8290.cd-14-0552>
11. Meghea, A. (2012). *Molecular interactions*. BoD – Books on Demand.
12. Meng, L., He, Y., Sun, C., Huang, L., Hu, T., & Yang, F. (2024). Learning personalized drug features and differentiated drug-pair interaction information for drug-drug interaction prediction. *Neural Networks*, 181, 106828. <https://doi.org/10.1016/j.neunet.2024.106828>
13. Peng, W., Liu, H., Dai, W., Yu, N., & Wang, J. (2022). Predicting cancer drug response using parallel heterogeneous graph convolutional networks with neighborhood interactions. *Bioinformatics*, 38(19), 4546–4553. <https://doi.org/10.1093/bioinformatics/btac574>
14. Rizkallah, M. R., Gamal-Eldin, S., Saad, R., & Aziz, R. K. (2012). The PharmacoMicrobiomics Portal: a database for Drug-Microbiome interactions. *Current Pharmacogenomics and Personalized Medicine (Online)/Current Pharmacogenomics and Personalized Medicine*, 10(3), 195–203. <https://doi.org/10.2174/187569212802510030>
15. Schuth, S., Blanc, S. L., Krieger, T. G., Jabs, J., Schenk, M., Giese, N. A., Büchler, M. W., Eils, R., Conrad, C., & Strobel, O. (2022). Patient-specific modeling of stroma-mediated chemoresistance of pancreatic cancer using a three-dimensional organoid-fibroblast co-culture system. *Journal of Experimental & Clinical Cancer Research*, 41(1). <https://doi.org/10.1186/s13046-022-02519-7>
16. Swart, M., & Dandara, C. (2014). Genetic variation in the 3'UTR of CYP1A2, CYP2B6, CYP2D6, CYP3A4, NR1I2, and UGT2B7: potential effects on regulation by microRNA and pharmacogenomics relevance. *Frontiers in Genetics*, 5. <https://doi.org/10.3389/fgene.2014.00167>
17. Vizirianakis, I. S. (2014). *Handbook of Personalized Medicine: Advances in Nanotechnology, Drug Delivery, and Therapy*. CRC Press.
18. Wang, H., Meng, Q., & Ma, B. (2021). Characterization of the prognostic M6A-Related LNCRNA signature in gastric cancer. *Frontiers in*

**PersonalDDI: A Temporal-Causal Framework for Personalized Drug-Drug Interaction Prediction
Using Explainable Generative AI**

Oncology, 11.
<https://doi.org/10.3389/fonc.2021.630260>

19. Yang, J., Li, A., Li, Y., Guo, X., & Wang, M. (2018). A novel approach for drug response prediction in cancer cell lines via network representation learning. *Bioinformatics*, 35(9), 1527–1535. <https://doi.org/10.1093/bioinformatics/bty848>
20. Ye, Y., Wen, Y., Zhang, Z., He, S., & Bo, X. (2021). Drug-Target Interaction Prediction based on adversarial Bayesian personalized ranking. *BioMed Research International*, 2021(1). <https://doi.org/10.1155/2021/6690154>