Patient-State–Conditioned Pharmacodynamics Modeling With Counterfactual Inference for ADR Risk in Peritoneal Dialysis

Yudi Zhang[1], Xinqiu Li[2], Kai Niu[1], Jie Dong[2,*], Zhiqiang He[1,*]
[1]Key Laboratory of Universal Wireless Communications, Ministry of Education,
Beijing University of Posts and Telecommunications, Beijing, China
[2]Renal Division, Department of Medicine, Peking University First Hospital, Beijing, China
Email: zed@bupt.edu.cn, 2011210231@stu.pku.edu.cn, niukai@bupt.edu.cn,
jie.dong@bjmu.edu.cn, hezq@bupt.edu.cn

***Abstract*—Patients with chronic kidney disease (CKD) undergoing peritoneal dialysis (PD) commonly experience multimorbidity and polypharmacy, making adverse drug reaction (ADR) risk highly patient-specific. Conventional trials and drug-centric computational models often lack sufficient patient context and longitudinal validity to capture context-specific safety signals in real-world PD follow-ups. We propose PS-PDNet, a patient-state–conditioned deep learning framework that integrates clinical indicators and longitudinal prescriptions to predict ADR risk in PD. By conditioning pharmacodynamics modeling on clinical state and exposure history, PS-PDNet supports individualized safety assessment. Coupled with counterfactual inference, it highlights high-risk drugs and combinations and summarizes them into category-level exposure patterns for ADR surveillance and medication optimization in real-world PD care.**

***Index Terms*—ADR Prediction, Pharmacodynamics Modeling, Counterfactual Inference, Peritoneal Dialysis**

I. Introduction

Patients with chronic kidney disease (CKD) often live with multimorbidity and long-term medication exposure, making polypharmacy commonplace [1]–[3]. In this setting, adverse drug reactions (ADRs) and drug–drug interactions (DDIs) are not isolated events but safety signals that may evolve with changing patient physiology, comorbidity profiles, and concomitant medications [4]–[6]. However, conventional drug trials are often underpowered for rare or context-specific ADRs, and their restricted follow-up and homogeneous enrollment limit generalizability to real-world PD care [7], [8].

Existing ADR/DDI prediction studies largely rely on drug-centric learning from curated knowledge bases or pharmacovigilance resources, providing mechanistic priors but limited patient context [9]–[13]. In contrast, electronic health records (EHRs) provide a rich, patient-grounded source for detecting adverse drug events and medication-related relations in real-world care, enabling pharmacovigilance to move beyond drug-only signals toward context-aware safety assessment [14], [15]. EHR-based medication modeling has been explored for recommendation and outcome prediction, yet many methods treat longitudinal history as aggregated covariates rather than explicitly modeling state-conditioned, time-evolving pharmacodynamics [16]–[19]. For PD surveillance, this gap is clinically consequential: the same drug or combination may be tolerable in one patient but precipitate electrolyte or hepatic abnormalities in another, depending on baseline reserve, longitudinal laboratory trends, and prior ADR trajectories [20], [21].

To address these challenges, we frame ADR trajectory prediction in PD follow-ups as patient-state–conditioned pharmacodynamics modeling. We propose PS-PDNet, which couples prescription records with a previous-visit patient state, so that the same polypharmacy regimen is mapped to different risk representations under different clinical states, enabling individualized medication-effect estimation aligned with PD decision-making. Importantly, PS-PDNet is counterfactual-ready by design: its tokenized medication representations allow systematic removal-based perturbations to quantify PD-sensitive drugs and co-medication patterns. With an interpretation workflow, it translates model-derived signals into clinically readable category-level findings and exposure patterns for actionable ADR surveillance under real-world polypharmacy. The primary contributions of this research are summarized as follows:

- We propose PS-PDNet, a deep learning network for ADR prediction in PD under multimorbidity and polypharmacy, and couple it with counterfactual inference to prioritize high-risk drugs and drug combinations. An interpretation workflow is then applied to contextualize these findings and assess their plausibility.
- PS-PDNet introduces a two-stage attention architecture with IPIT for intra-prescription interaction modeling and IPTT for inter-prescription temporal modeling. Both modules are conditioned on the previous-visit patient state, enabling patient-specific pharmacodynamic representations.
- We incorporate switchable pairwise and triplet feature extraction modules to construct explicit interaction tokens.

*Corresponding authors.

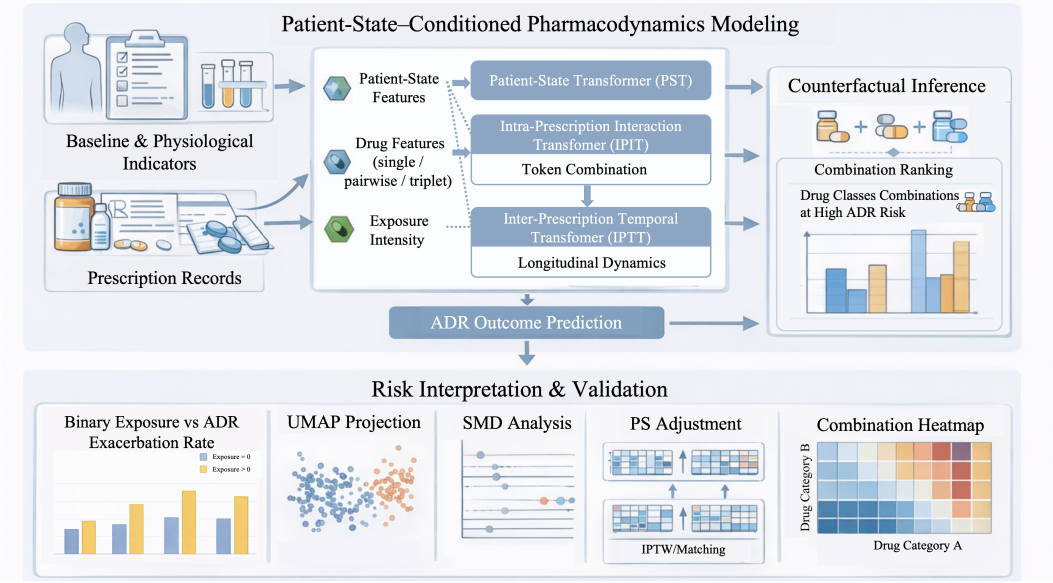

Fig. 1. Overview of PS-PDNet and the interpretation pipeline.

Using counterfactual perturbations, we quantify each drug and combination's marginal contribution to a target ADR endpoint, supporting post-hoc interpretation.

II. Data Preparation

The dataset was obtained from the Peritoneal Dialysis Telemedicine-based Management Platform (PDTAP), an internet-enabled clinical management system for PD follow-up. This retrospective analysis was conducted on a single-center prospective cohort approved by the Biomedical Research Ethics Committee of Peking University First Hospital (Approval No. 2018 Scientific Research-100) and registered on ClinicalTrials.gov (NCT03571451; 2018-06-20). Written informed consent was obtained from all participants.

We collected longitudinal follow-up records spanning 2016–2025, including repeated laboratory measurements (e.g., serum biochemistry and complete blood counts) and time-stamped prescription records. The cohort comprised 29,038 visits from 1,098 patients; after excluding visits without a prior follow-up within 3 months, 26,348 visits were retained as eligible samples. On average, each patient contributed 25 samples, each sample contained 2.5 prescriptions, and each prescription included 8.5 distinct medications across approximately 500 medications.

For each follow-up instance, laboratory results were aligned to the visit time, and medication exposure was reconstructed from the prescription sequence preceding that visit. After eligibility screening and variable harmonization, we retained 66 baseline and physiological indicators as inputs. The prediction targets comprised six clinically relevant ADR endpoints: alanine aminotransferase increased, aspartate aminotransferase increased, blood bilirubin increased, γ-glutamyltransferase increased, hyperkalemia, and hypokalemia. Each endpoint was formulated as a three-class outcome to reflect clinically meaningful trajectories across consecutive visits, where class (2) indicates exacerbation, class (1) indicates no change, and class (0) indicates remission. Exacerbation—our primary endpoint of interest—accounted for 4.73% of eligible samples. To improve clinical interpretability and mechanistic coherence under polypharmacy, each medication was mapped to a three-level category system based on clinical use, pharmacological mechanism, and therapeutic specialty (excluding traditional Chinese medicines). Baseline and physiological variables were grouped by type (continuous, binary, ordinal, and nominal) and processed with type-specific normalization and encoding.

III. Proposed Methodology

A. Model Overview

We propose PS-PDNet, a patient-state–conditioned deep learning framework for ADR risk prediction in PD follow-ups (Fig. 2). Each visit is modeled as a joint input of current physiological indicators, a previous-visit patient state, and a longitudinal prescription record.

Prescription records are mapped to drug features via hierarchical drug semantic embeddings initialized from multi-level drug category text using a Sentence Transformer. To represent higher-order polypharmacy, PS-PDNet constructs single-drug, pairwise, and triplet tokens. IPIT encodes within-prescription interactions under previous-visit state conditioning, and IPTT models temporal dynamics across prescriptions with a causal mask. The resulting medication representation is fused with physiological features to predict the multi-class ADR outcome.

B. Feature Vectorization Encoder

A PD follow-up sample is indexed by $i \in \{1, \ldots, M\}$, where M is the cohort size. We represent each sample as the tuple $\mathcal{D}_i = \left(\mathbf{x}_{\text{pre}}^{(i)}, \mathbf{x}_{\text{cur}}^{(i)}, \{\mathcal{P}_t^{(i)}\}_{t=1}^{T_i}\right)$, where $\mathbf{x}_{\text{cur}}^{(i)} \in \mathbb{R}^{d_{\text{base}}}$ denotes current-visit baseline and physiological indicators, and

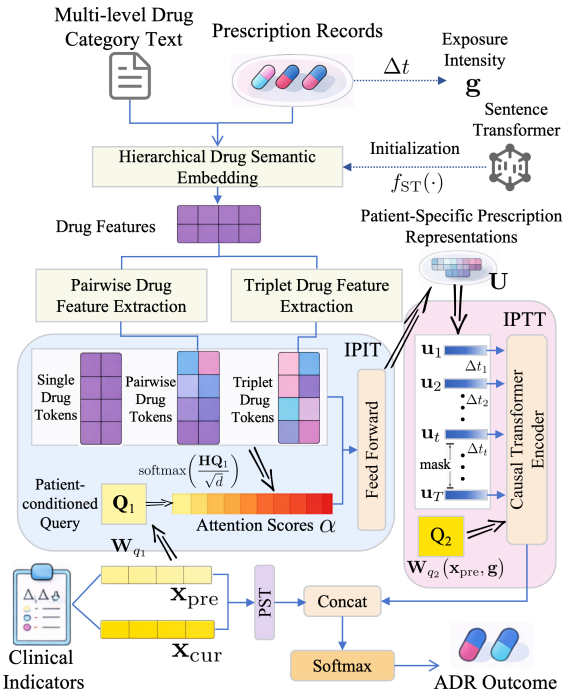

Fig. 2. Architecture of PS-PDNet

$\mathbf{x}_{\text{pre}}^{(i)} \in \mathbb{R}^{d_{\text{base}}+LC}$ denotes the previous-visit patient-state features augmented with prior ADR status. The prior ADR state is encoded over L ADR endpoints, each with C classes.

Prescription records are represented as a time-ordered sequence $\{\mathcal{P}_t^{(i)}\}_{t=1}^{T_i}$, where each $\mathcal{P}_t^{(i)}$ is a multiset of medications. Let $\mathcal{V}$ be the drug vocabulary and $v \in \mathcal{V}$ denote a drug with hierarchical category text $\mathcal{T}(v)$. A Sentence Transformer encoder $f_{\text{ST}}(\cdot)$ maps $\mathcal{T}(v)$ to a dense embedding $\mathbf{e}_v \in \mathbb{R}^d$. These text-derived vectors initialize trainable drug embeddings and improve transferability by encoding mechanism- and class-level semantics. For prescription t of sample i, we denote single-drug tokens as $\{\mathbf{s}_{t,n}^{(i)}\}_{n=1}^{N_{1,t}^{(i)}}$ with $\mathbf{s}_{t,n}^{(i)} \in \mathbb{R}^d$, where $N_{1,t}^{(i)}$ is the number of unique medication instances. To construct higher-order medication tokens, we use learnable extractors $\phi_2 : \mathbb{R}^{2d} \to \mathbb{R}^d$ and $\phi_3 : \mathbb{R}^{3d} \to \mathbb{R}^d$. For selected pairs (n, m) and triplets (n, m, k) within prescription t, we form $\mathbf{p}_{t,(n,m)}^{(i)} = \phi_2\big(\text{concat}(\mathbf{s}_{t,n}^{(i)}, \mathbf{s}_{t,m}^{(i)})\big) \in \mathbb{R}^d$, and $\mathbf{r}_{t,(n,m,k)}^{(i)} = \phi_3\big(\text{concat}(\mathbf{s}_{t,n}^{(i)}, \mathbf{s}_{t,m}^{(i)}, \mathbf{s}_{t,k}^{(i)})\big) \in \mathbb{R}^d$. Pairwise and triplet tokens were retained only after informatics-based screening of clinically plausible and sufficiently supported co-medication patterns. Denote the counts of pairwise and triplet tokens as $N_{2,t}^{(i)}$ and $N_{3,t}^{(i)}$, respectively, and define $N_t^{(i)} = N_{1,t}^{(i)} + N_{2,t}^{(i)} + N_{3,t}^{(i)}$. Stacking all tokens yields $\mathbf{H}_t^{(i)} \in \mathbb{R}^{N_t^{(i)} \times d}$, which serves as the input token matrix for intra-prescription modeling.

C. IPIT and IPTT Module

Given the token matrix $\mathbf{H}_t^{(i)}$, IPIT performs within-prescription pharmacodynamic aggregation conditioned on the PD patient's recent clinical state. Specifically, the previous-visit patient state $\mathbf{x}_{\text{pre}}^{(i)}$ is used as a conditioning signal in attention allocation, so an identical drug set can induce different risk-relevant summaries under different physiological reserves and prior ADR trajectories. We project $\mathbf{x}_{\text{pre}}^{(i)}$ into the token space via $\mathbf{Q}_{t,1}^{(i)} = \mathbf{W}_{q_1}\mathbf{x}_{\text{pre}}^{(i)} + \mathbf{b}_{q_1} \in \mathbb{R}^d$, compute patient-conditioned attention scores $\boldsymbol{\alpha}_t^{(i)} = \text{softmax}\left(\frac{\mathbf{H}_t^{(i)}\mathbf{Q}_{t,1}^{(i)}}{\sqrt{d}}\right) \in \mathbb{R}^{N_t^{(i)}}$ and obtain the prescription representation by a weighted pooling, $\mathbf{u}_t^{(i)} = (\boldsymbol{\alpha}_t^{(i)})^\top \mathbf{H}_t^{(i)} \in \mathbb{R}^d$. This design couples prior patient tolerance, including previous ADR status, to the contribution of single-, pair-, and triplet-level medication tokens, which matches PD practice where exacerbation risk is often driven by state-dependent susceptibility rather than drug identity alone.

IPTT further models across-prescription temporal dynamics on the time-ordered sequence $\{\mathbf{u}_t^{(i)}\}_{t=1}^{T_i}$. To encode causal exposure evolution, IPTT uses a unidirectional attention mask so that the representation at time t only attends to $\{1, \ldots, t\}$. Let $\mathbf{U}^{(i)} = \text{stack}(\mathbf{u}_1^{(i)}, \ldots, \mathbf{u}_{T_i}^{(i)}) \in \mathbb{R}^{T_i \times d}$ and $\mathbf{M}^{(i)} \in \{-\infty, 0\}^{T_i \times T_i}$ be the causal mask with $\mathbf{M}_{t,\tau}^{(i)} = -\infty$ for $\tau > t$. IPTT forms a state-and-time conditioned query so that prescription salience depends on both patient history and exposure timing. Let $\mathbf{g}_t^{(i)} \in \mathbb{R}^d$ denote a learned embedding of temporal distances that summarizes the gap from prescription t to the current visit and to the next prescription, and define $\mathbf{Q}_{t,2}^{(i)} = \mathbf{W}_{q_2}\,\text{concat}\big(\mathbf{x}_{\text{pre}}^{(i)}, \mathbf{g}_t^{(i)}\big) + \mathbf{b}_{q_2} \in \mathbb{R}^d$. Using standard scaled dot-product attention with masked future positions, one layer can be written as $\text{Attn}\Big(\mathbf{Q}_2^{(i)}, \mathbf{K}^{(i)}, \mathbf{V}^{(i)}\Big) = \text{softmax}\left(\frac{\mathbf{Q}_2^{(i)}(\mathbf{K}^{(i)})^\top}{\sqrt{d}} + \mathbf{M}^{(i)}\right)\mathbf{V}^{(i)}$, where $\mathbf{Q}_2^{(i)} \in \mathbb{R}^{T_i \times d}$ stacks $\mathbf{Q}_{t,2}^{(i)}$ over time and $\mathbf{K}^{(i)}$, $\mathbf{V}^{(i)} \in \mathbb{R}^{T_i \times d}$ are linear projections of $\mathbf{U}^{(i)}$. Stacking such causal layers yields $\tilde{\mathbf{U}}^{(i)} \in \mathbb{R}^{T_i \times d}$, which captures how earlier prescriptions propagate forward under state-dependent pharmacodynamics. This is important because ADR exacerbation depends on accumulated exposure history, while future prescriptions must not leak into earlier risk estimation.

D. ADR Prediction and Masked Multi-Task Optimization

Given the IPTT-derived visit-level medication representation $\mathbf{h}_{\text{rx}}^{(i)} \in \mathbb{R}^d$ and the physiology representation $\mathbf{h}_{\text{phys}}^{(i)} \in \mathbb{R}^d$ ($\mathbf{h}_{\text{phys}}^{(i)}$ is produced by the Patient-State Transformer (PST) from $\mathbf{x}_{\text{pre}}^{(i)}$ and $\mathbf{x}_{\text{cur}}^{(i)}$), where d is the latent feature dimension, we fuse them by concatenation $\mathbf{h}^{(i)} = \text{concat}\big(\mathbf{h}_{\text{rx}}^{(i)}, \mathbf{h}_{\text{phys}}^{(i)}\big) \in \mathbb{R}^{2d}$. A prediction head $g_\theta(\cdot)$ maps $\mathbf{h}^{(i)}$ to multi-endpoint logits $\mathbf{o}^{(i)} = g_\theta\big(\mathbf{h}^{(i)}\big) \in \mathbb{R}^{L \times C}$, where L and C denote the numbers of endpoints and classes, respectively. For each endpoint ℓ, the class probability is $\mathbf{p}_\ell^{(i)} = \text{softmax}\big(\mathbf{o}_{\ell,:}^{(i)}\big) \in [0,1]^C$, where $y_\ell^{(i)} \in \{0, 1, \ldots, C-1\}$ denotes the observed ADR-change class between two consecutive follow-ups.

To handle missing endpoint annotations, we introduce an endpoint-wise validity mask $\mathbf{m}^{(i)} \in \{0,1\}^L$ and optimize only

TABLE I
COMPARISON OF PREDICTIVE PERFORMANCE ACROSS BASELINE MODELS AND PS-PDNET

Methods	Accuracy	F1-macro	AUROC-macro	AUPRC-macro
LSTM	0.881	0.654	0.898	0.649
DeepSets	0.897	0.653	0.899	0.649
RETAIN	0.895	0.648	0.889	0.638
ChronoFormer	0.906	0.654	0.901	0.649
SHGT	0.892	0.647	0.894	0.640
ICP-Transformer	0.894	0.648	0.895	0.634
PS-PDNet	0.897	0.656	0.905	0.643

evaluable endpoints. Let $p_\ell^{(i)} = \mathbf{p}_\ell^{(i)}\left[y_\ell^{(i)}\right] \in (0,1]$ be the probability assigned to the observed class. The masked multi-class focal objective for sample i is $\mathcal{L}_{\text{Focal}}^{(i)} = \frac{1}{\sum_{\ell=1}^{L} m_\ell^{(i)}} \sum_{\ell=1}^{L} m_\ell^{(i)} \cdot \alpha_{\ell,y_\ell^{(i)}} (1-p_\ell^{(i)})^\gamma \left(-\log p_\ell^{(i)}\right)$, where $\gamma > 0$ is the focusing exponent and $\alpha_{\ell,y} \geq 0$ is a class weight. The cohort-level objective is $\mathcal{L} = \frac{1}{M}\sum_{i=1}^{M} \mathcal{L}_{\text{Focal}}^{(i)}$. This mask-aware formulation preserves supervision under intermittently missing laboratory endpoints, while the focal factor emphasizes rare but clinically critical trajectory changes.

IV. EXPERIMENTS

A. Experimental Protocol and Baseline Comparison

We evaluated ADR risk prediction for PD follow-up visits under a unified protocol designed to reflect real-world longitudinal care and polypharmacy. The dataset was split into training and validation sets at a 9:1 ratio. Importantly, the split was performed at the patient level to prevent information leakage across repeated visits from the same individual. We compared PS-PDNet against six representative baselines spanning temporal sequence learning (LSTM), permutation-invariant set modeling (DeepSets [22]), clinically motivated visit-level attention (RETAIN [17]), time-aware Transformer architectures for irregular clinical timelines (ChronoFormer [23]), structure-aware hypergraph Transformers for higher-order medication relations (SHGT [24]), and an interaction-conditioned prescription Transformer (ICP-Transformer [25]). Collectively, these baselines cover common modeling assumptions in medication–outcome studies, ranging from order-invariant pooling to sequential dependence, explicit time modeling, and high-order interaction reasoning. For a fair comparison, all baselines were trained under the same data preprocessing pipeline, using comparable model depth/capacity, and a consistent training protocol; all models converged on the validation set.

Performance was summarized using Accuracy, F1-macro, AUROC-macro, and AUPRC-macro (Table I). All baselines were adapted to our dataset and task formulation, so the results may not reflect their original upper bounds. Notably, macro-level F1 remains modest across methods, largely due to severe class imbalance (the dominant "No change" class accounts for $\sim 90.5\%$ of samples) and stochastic, non-medication factors that shape ADR trajectories. Nevertheless, PS-PDNet achieves the best overall performance under the same protocol (Accuracy = 0.897, F1-macro = 0.656, AUROC-macro = 0.905, and AUPRC-macro = 0.643), and its state-conditioned, temporally causal design better aligns with the medication–effect pathway, providing a principled basis for the counterfactual analyses that follow.

TABLE II
ABLATION EXPERIMENTS OF PS-PDNET WITH DIFFERENT COMPONENTS

Component	Choice				
Hierarchical Embedding		✓	✓	✓	✓
Combo Drug Tokens	✓		✓	✓	✓
Causal IPTT	✓	✓		✓	✓
Patient-State Conditioning	✓	✓	✓		✓
F1-macro	0.645	0.646	0.655	0.643	**0.656**

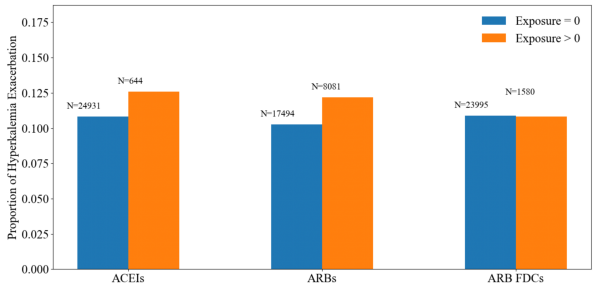

Fig. 3. Proportion of hyperkalemia exacerbation stratified by exposure status (exposure = 0 vs exposure $>$ 0) for three representative highest-Δ drug categories.

B. Ablation Study

In ablation studies, we report F1-macro for class-balanced evaluation. Table II shows that all components contribute to PS-PDNet (F1-macro = 0.656). Replacing hierarchical text-initialized embeddings with a single ID embedding reduces F1-macro to 0.645, indicating that category-level semantics improves generalization beyond drug identifiers. Removing pair/triple combination tokens reduces F1-macro to 0.646, suggesting that higher-order co-medication tokens capture risk signatures beyond single-drug tokens. Replacing causal IPTT with a non-causal variant yields a small decrease (0.655), but the causal mask remains clinically necessary to enforce temporal validity and prevent future-prescription leakage when modeling exposure accumulation over follow-ups. Finally, removing patient-state conditioning causes the largest drop to 0.643, corroborating that ADR exacerbation in PD is strongly state-dependent: the same medication set may be tolerated or become hazardous depending on baseline reserve and prior ADR trajectory; conditioning attention on the previous-visit state therefore constitutes the primary driver of patient-specific pharmacodynamic aggregation.

C. Counterfactual Inference

1) Counterfactual Discovery and Category-level Summarization: We performed counterfactual inference on the trained PS-PDNet using hyperkalemia exacerbation as an exemplar endpoint. Attribution was confined to a clinically coherent high-risk subset $\mathcal{S} = \{i \mid y^{(i)} = 2,\ \hat{y}^{(i)} = 2\}$. For

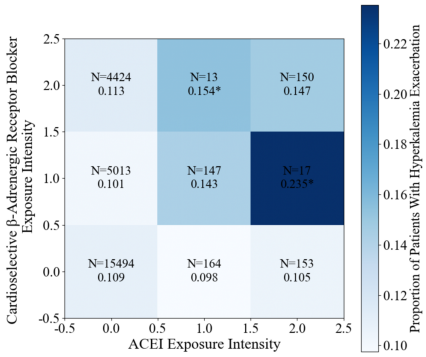

Fig. 4. Two-dimensional exposure–response heatmap for the ACEI and cardioselective β-adrenergic receptor blocker combination, where color intensity reflects the proportion of hyperkalemia exacerbation.

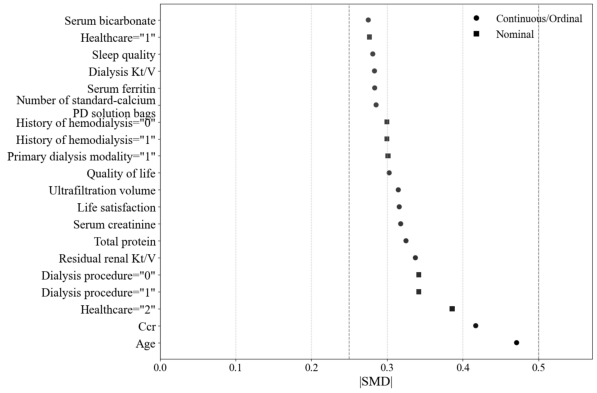

Fig. 5. Signed Love plot of SMDs comparing ACEI–exposed versus unexposed groups across the top 20 baseline and physiological covariates.

each $i \in \mathcal{S}$, we cached the baseline class-2 score $s_2^{(i)}$, then generated token-level counterfactuals by removing one medication token τ (single/pair/triplet) from the prescription token set prior to IPIT pooling (equivalently, masking τ and excluding it from attention normalization), while keeping all timestamps and remaining tokens unchanged. A forward pass yields the perturbed score $s_2^{(i,-\tau)}$, and the counterfactual shift is $\Delta^{(i)}(\tau) = s_2^{(i)} - s_2^{(i,-\tau)}$. We aggregate token effects by summing $\Delta^{(i)}(\tau)$ over $i \in \mathcal{S}$ only when $\Delta^{(i)}(\tau) > 0$, and normalize by the number of such positive occurrences to obtain $\Delta(\tau)$. Finally, token-level attributions are mapped to level-3 therapeutic categories (and category tuples for combination tokens) and aggregated to yield category-level impact scores.

We next examined the counterfactual findings from the perspective of the empirical sample distribution in the original cohort. A per-sample category exposure vector $\mathbf{X}_{\text{mech}}$ based on level-3 categories was constructed, with each element encoding the presence (and strength) of exposure along the prescription sequence. We dichotomized exposure for three representative high-Δ categories into "exposure = 0" versus "exposure > 0" and plotted bar charts of the observed proportion of hyperkalemia exacerbation in the two groups.

2) Interpretation of Counterfactual Findings: Angiotensin-converting enzyme inhibitors (ACEIs) and angiotensin II receptor blockers (ARBs) can aggravate hyperkalemia by suppressing renin–angiotensin–aldosterone activity, reducing distal tubular potassium excretion and predisposing susceptible CKD patients to potassium accumulation [26], [27]. In Fig. 3, the ARB-FDC–exposed group does not show a clearly higher proportion of exacerbation because our dominant angiotensin II receptor blocker-based fixed-dose combinations (ARB FDCs) are predominantly thiazide-containing regimens, whose kaliuretic effect may counterbalance RAAS-related potassium retention [28], [29]. In Fig. 3, the first two categories show only small differences in the outcome proportion between exposed and unexposed groups. Moreover, when projecting samples to a two-dimensional manifold via UMAP based on $\mathbf{X}_{\text{mech}}$, no clear exposure-dependent gradient in hyperkalemia exacerbation is visually discernible. Together, these observations suggest that medication signals are diluted by confounding and heterogeneity. Fig. 4 presents a two-dimensional exposure–response heatmap for the ACEI plus cardioselective β-adrenergic receptor blocker combination, a representative high-Δ pairwise co-medication category, where color intensity denotes the proportion of hyperkalemia exacerbation. Given the plausible mechanisms by which this combination can amplify potassium risk [30], [31], higher joint exposure aligns with a higher hyperkalemia exacerbation proportion.

Using ACEI as an example, exposed and unexposed groups exhibited substantial baseline imbalance (many covariates with $|\text{SMD}| > 0.3$; Fig. 5), indicating that exposure status corresponds to distinct patient profiles. We therefore applied propensity-score adjustment (IPTW and matching) to test whether the observed hyperkalemia difference is attributable to ACEI itself rather than confounding. After adjustment, effect estimates attenuated and became unstable, with confidence intervals crossing the null, precluding a reliable causal conclusion from observational data alone. In this context, PS-PDNet's counterfactual perturbation provides an orthogonal, model-based line of evidence by directly quantifying the predicted risk shift under token removal, serving as an objective corroboration signal for prioritizing candidate drug effects under real-world confounding.

V. CONCLUSION

In this paper, we present PS-PDNet for ADR trajectory prediction in PD follow-ups under real-world polypharmacy. PS-PDNet operationalizes patient-specific pharmacodynamics by conditioning both intra-prescription interaction aggregation (IPIT) and inter-prescription temporal modeling (IPTT) on the prior-visit patient state, while enforcing temporal validity through causal modeling of longitudinal prescriptions. Empirically, PS-PDNet achieves the strongest overall performance among representative baselines, and ablation results show that patient-state conditioning is a primary driver of the gain, with additional benefits from hierarchical semantics and explicit combination tokens.

Beyond prediction, we couple PS-PDNet with counterfactual inference to prioritize high-impact drugs and drug

combinations for a target ADR endpoint (illustrated with hyperkalemia exacerbation), and we summarize token-level attributions into clinically readable category-level signals and exposure patterns. This end-to-end design supports actionable PD medication safety surveillance by linking trajectory-aware risk estimation to interpretable candidates for review and monitoring. Currently, our pharmacodynamic modeling focuses on safety via ADR prediction; in future work, we will extend the framework to jointly model medication benefits and safety to better support individualized treatment optimization.

ACKNOWLEDGMENT

This study was partly supported by the National Key Research and Development Program of China (2025YFC2511800), the Excellent Plan for Medicine Innovation and Translation Project (YC202401QX0710), the National High Level Hospital Clinical Research Funding (High Quality Clinical Research Project of Peking University First Hospital 2022CR82), the National High Level Hospital Clinical Research Funding (Scientific and Technological Achievements Transformation Incubation Guidance Fund Project of Peking University First Hospital 2022CX09), and the Beijing Special Funding Program for the Cultivation of Innovative Medicines and Platforms (Z241100009024052).

REFERENCES

[1] I. J. Oosting *et al.*, "Polypharmacy in Patients with CKD: A Systematic Review and Meta-Analysis," *Kidney360*, vol. 5, no. 6, pp. 841–850, Jun. 2024.

[2] N. Panajotovikj, A. Fernández-Torras, and J. Mestres, "Drug safety in healthy aging," *npj Aging*, vol. 11, no. 1, p. 28, 2025.

[3] A. Mouheb, M. Metzger, N. Alencar de Pinho, *et al.*, "Hyperpolypharmacy in patients with chronic kidney disease and its impact on clinical outcomes," *Scientific Reports*, vol. 16, no. 1, p. 2653, 2026.

[4] A. O. Gerard, R. Lombardi, D. Merino, *et al.*, "A new chapter in pharmacology: Artificial intelligence's expanding role in pharmacokinetics, pharmacodynamics, and pharmacovigilance," *Therapies*, 2025.

[5] R. Mehta, "Postmarket drug safety monitoring," in *Translational Gastroenterology*. Academic Press, 2025, pp. 395–397.

[6] S. Fauziyah, M. Radji, and R. Andrajati, "Polypharmacy in elderly patients and their problems," *Asian J. Pharm. Clin. Res.*, vol. 10, no. 7, pp. 44–49, 2017.

[7] D. Galeano, S. Li, M. Gerstein, *et al.*, "Predicting the frequencies of drug side effects," *Nature Communications*, vol. 11, no. 1, p. 4575, 2020.

[8] B. Vora, *Integration of Benchwork, Clinical Trials, and Real World Data to Investigate Drug Interactions*. University of California, San Francisco, 2021.

[9] J. Y. Ryu, H. U. Kim, and S. Y. Lee, "Deep learning improves prediction of drug–drug and drug–food interactions," *Proc. Natl. Acad. Sci. U.S.A.*, vol. 115, no. 18, pp. E4304–E4311, 2018.

[10] M. Zitnik, M. Agrawal, and J. Leskovec, "Modeling polypharmacy side effects with graph convolutional networks," *Bioinformatics*, vol. 34, no. 13, pp. i457–i466, 2018.

[11] S. Lin, Y. Wang, L. Zhang, *et al.*, "MDF-SA-DDI: predicting drug–drug interaction events based on multi-source drug fusion, multi-source feature fusion and transformer self-attention mechanism," *Briefings in Bioinformatics*, vol. 23, no. 1, p. bbab421, 2022.

[12] Y. Yuan, J. Yue, R. Zhang, *et al.*, "PHGL-DDI: a pre-training based hierarchical graph learning framework for drug–drug interaction prediction," *Expert Systems with Applications*, vol. 270, p. 126408, 2025.

[13] D. Li, Y. Yang, Z. Cui, *et al.*, "LLM-DDI: Leveraging Large Language Models for Drug-Drug Interaction Prediction on Biomedical Knowledge Graph," *IEEE J. Biomed. Health Inform.*, 2025.

[14] F. Christopoulou, T. T. Tran, S. K. Sahu, *et al.*, "Adverse drug events and medication relation extraction in electronic health records with ensemble deep learning methods," *J. Am. Med. Inform. Assoc.*, vol. 27, no. 1, pp. 39–46, 2020.

[15] J. Yang and F. Li, "The Impact of Technological Progress on Pharmacovigilance," in *Pharmacovigilance–Facts, Challenges, Limitations and Opportunities: Facts, Challenges, Limitations and Opportunities*, 2025, p. 43.

[16] C. Scherkl, T. Dierkes, M. Metzner, *et al.*, "Towards a prescribing monitoring system for medication safety evaluation within electronic health records: a scoping review," *BMC Med. Inform. Decis. Mak.*, vol. 25, p. 244, 2025.

[17] E. Choi, M. T. Bahadori, J. Sun, *et al.*, "RETAIN: An interpretable predictive model for healthcare using reverse time attention mechanism," in *Advances in Neural Information Processing Systems*, vol. 29, 2016.

[18] C. Yang, C. Xiao, F. Ma, *et al.*, "SafeDrug: Dual molecular graph encoders for recommending effective and safe drug combinations," *arXiv preprint* arXiv:2105.02711, 2021.

[19] C. Gao, S. Yin, H. Wang, *et al.*, "Medical-knowledge-based graph neural network for medication combination prediction," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 10, pp. 13246–13257, 2023.

[20] A. O. Basile, A. Yahi, and N. P. Tatonetti, "Artificial intelligence for drug toxicity and safety," *Trends Pharmacol. Sci.*, vol. 40, no. 9, pp. 624–635, 2019.

[21] P. Mathur, A. K. Khanna, J. B. Cywinski, *et al.*, "2019 YEAR IN REVIEW: MACHINE LEARNING IN HEALTHCARE," *Team BrainX, BrainX Community*, 2019.

[22] M. Zaheer, S. Kottur, S. Ravanbakhsh, *et al.*, "Deep sets," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.

[23] Y. Zhang and S. Li, "ChronoFormer: Time-Aware Transformer Architectures for Structured Clinical Event Modeling," *arXiv preprint* arXiv:2504.07373, 2025.

[24] H. Wang and Y. Yuan, "Structure-aware Hypergraph Transformer for Diagnosis Prediction in Electronic Health Records," *arXiv preprint* arXiv:2508.20500, 2025.

[25] Y. Lu, W. Chang, D. Yang, *et al.*, "A Three-Stage Prediction Model Based on Transformer for Chronic Kidney Disease," *Concurrency and Computation: Practice and Experience*, vol. 37, no. 25–26, e70322, 2025.

[26] M. A. Raebel, "Hyperkalemia associated with use of angiotensin-converting enzyme inhibitors and angiotensin receptor blockers," *Cardiovascular Therapeutics*, vol. 30, no. 3, pp. e156–e166, 2012.

[27] H. Kim, S. Park, J. H. Jhee, *et al.*, "Urinary angiotensinogen level is associated with potassium homeostasis and clinical outcome in patients with polycystic kidney disease: a prospective cohort study," *BMC Nephrology*, vol. 20, no. 1, p. 104, 2019.

[28] R. J. McNally, B. Farukh, P. J. Chowienczyk, *et al.*, "Effect of diuretics on plasma aldosterone and potassium in primary hypertension: A systematic review and meta-analysis," *British Journal of Clinical Pharmacology*, vol. 88, no. 5, pp. 1964–1977, 2022.

[29] D. Fuwa, M. Fukuda, Y. Ogiyama, *et al.*, "Addition of hydrochlorothiazide to angiotensin receptor blocker therapy can achieve a lower sodium balance with no acceleration of intrarenal renin angiotensin system in patients with chronic kidney disease," *Journal of the Renin–Angiotensin–Aldosterone System*, vol. 17, no. 2, p. 1470320316652032, 2016.

[30] T. W. B. Gehr and D. A. Sica, "Hyperkalemia in congestive heart failure," *Congestive Heart Failure*, vol. 7, no. 2, pp. 97–100, 2001.

[31] P. A. McCullough, T. M. Beaver, E. Bennett-Guerrero, *et al.*, "Acute and chronic cardiovascular effects of hyperkalemia: new insights into prevention and clinical management," *Reviews in Cardiovascular Medicine*, vol. 15, no. 1, pp. 11–23, 2014.