

Uncertainty-Aware Metabolic Stability Prediction with Dual-View Contrastive Learning

Peijin Guo¹, Minghui Li², Hewen Pan¹, Bowen Chen², Yang Wu², Zikang Guo², Leo Yu Zhang³, Shengshan Hu¹, and Shengqing Hu⁴✉

¹ School of Cyber Science and Engineering, Huazhong University of Science and Technology

{gpj, hewenpan, hushengshan}@hust.edu.cn

² School of Software Engineering, Huazhong University of Science and Technology
{minghuili, mchust, yungwu, zikangguo}@hust.edu.cn

³ School of Information and Communication Technology, Griffith University
leo.zhang@griffith.edu.au

⁴ Union Hospital, Tongji Medical College, Huazhong University of Science and Technology
hsqha@126.com

Abstract. Accurate prediction of molecular metabolic stability (MS) is critical for drug research and development but remains challenging due to the complex interplay of molecular interactions. Despite recent advances in graph neural networks (GNNs) for MS prediction, current approaches face two critical limitations: (1) incomplete molecular modeling due to atom-centric message-passing mechanisms that disregard bond-level topological features, and (2) prediction frameworks that lack reliable uncertainty quantification. To address these challenges, we propose TrustworthyMS, a novel contrastive learning framework designed for uncertainty-aware metabolic stability prediction. First, a molecular graph topology remapping mechanism synchronizes atom-bond interactions through edge-induced feature propagation, capturing both localized electronic effects and global conformational constraints. Second, contrastive topology-bond alignment enforces consistency between molecular topology views and bond patterns via feature alignment, enhancing representation robustness. Third, uncertainty modeling through Beta-Binomial uncertainty quantification enables simultaneous prediction and confidence calibration under epistemic uncertainty. Through extensive experiments, our results demonstrate that TrustworthyMS outperforms current state-of-the-art methods in terms of predictive performance.

Keywords: Contrastive learning, metabolic stability prediction

1 Introduction

Metabolic stability, defined as molecular resilience against enzymatic degradation, is crucial for pharmacokinetic optimization in drug discovery [17, 23]. This key factor influences a compound’s residence time in systemic circulation, directly affecting therapeutic efficacy by shaping absorption, distribution,

metabolism, and excretion (ADME) profiles [6]. Increasing global healthcare demands have intensified the need for targeted therapeutic development. However, traditional drug discovery remains prohibitively costly and time-consuming, underscoring the necessity for efficient computational methods to explore vast chemical spaces and identify viable candidates [8, 26, 11, 12, 5].

Recent advances in computational chemistry have established machine learning as a pivotal tool for *in silico* metabolic stability prediction, enabling rapid screening of viable drug candidates. Early efforts focused on conventional ML approaches: Podlowska et al. [16] demonstrated gains through algorithmic hybridization of random forests and SVMs, while Ryu et al. [19] leveraged PubChem’s mouse liver microsomal data to develop Bayesian predictors. Subsequent work by Deng et al. [1] integrated quantum mechanics with ensemble models for ester stability forecasting. However, these methods remain fundamentally constrained by their neglect of critical stereoelectronic effects encoded in molecular topology.

Graph Neural Networks (GNNs) present a promising alternative by explicitly modeling molecular topology [22, 4]. Moreover, GCNs can capture the relationships between nodes in a graph by propagating information across edges. The key idea is to aggregate and transform feature information from a node’s neighbors to update its representation, making GCNs particularly useful for tasks like node classification, link prediction, and graph-based learning problems. These networks take SMILES as input, which is a notation system used to represent the structure of molecules in a linear text format. Seminal work by Ren et al. [18] established Graph Convolutional Networks (GCN)-based frameworks for hierarchical feature extraction from Simplified molecular input line entry specification (SMILES)-derived graphs, while Du et al. [2] pioneered multimodal architectures combining graph contrastive learning with SMILES-specific attention mechanisms. Recent innovations by Wang et al. [21] further demonstrated the efficacy of bond-graph augmentation strategies. However, these approaches share a critical limitation: they provide point estimates of metabolic stability without quantifying prediction uncertainty—a dangerous oversight in drug discovery where erroneous predictions can derail entire development pipelines.

To address the dual challenges of **incomplete molecular topology modeling** and **absence of reliable uncertainty quantification**, we propose TrustworthyMS — a novel graph neural network (GNN) framework that synergizes topology-bond contrastive alignment with evidential reasoning. Our contributions are threefold:

1. We introduce a novel topology-enhanced dual-view contrastive learning framework that explicitly models both atom-level and bond-level interactions through molecular graph topology remapping, enabling more comprehensive representation of molecular structures.
2. To the best of our knowledge, this is the first work to achieve uncertainty-aware metabolic stability prediction by integrating Beta-Binomial subjective logic into graph neural networks, providing crucial confidence estimates for drug discovery applications.

3. Extensive experiments conducted on a dataset comprising 10,031 compounds reveal that our TrustworthyMS model outperforms existing methods in comparison to the baseline model. Specifically, it demonstrates a remarkable 46.1% improvement in robustness on out-of-distribution (OOD) data, while also surpassing current state-of-the-art approaches in both classification (0.622 MCC) and regression (0.833 P-score) tasks.

2 Related Work

Feature Engineering. Traditional assessments of metabolic stability have largely depended on laboratory experiments [3]. In contrast, computational approaches capitalize on chemical domain knowledge through feature engineering strategies. The descriptor-based approach extracts quantitative molecular properties (e.g., molecular weight, logP) from chemical structures. For instance, [15] and [19] utilized such descriptors to build random forest models. Complementing this, the fingerprint-based methodology encodes structural patterns into binary vectors using techniques such as ECFP and MACCS fingerprints. Notably, [16] and [1] integrated these fingerprints with molecular descriptors to construct compound representations.

Graph Representation Learning. Molecular graph representation methods have been applied to metabolic stability prediction through SMILES-derived topological constructions [10]. [18] proposed converting SMILES sequences into molecular graphs processed by graph convolutional networks (GCNs), employing node feature aggregation across graph neighborhoods. [2] presented CMMS-GCL, a method combining SMILES sequence similarity features with molecular graph structural features through feature concatenation. Current implementations primarily focus on atom-level representations and neighborhood information propagation, while chemical bond attributes patterns have not been fully explored in these architectures, as noted in [21].

3 Methods

Fig. 1 illustrates the TrustworthyMS architecture, a novel framework addressing metabolic stability prediction through three synergistic modules. The system processes SMILES inputs via: (1) **Molecular Graph Topology Remapping (Fig. 1a)**, where RDKit-constructed[9] molecular graphs are augmented with bond-centric nodes (atom-bond-atom triplets) to form dual representations, capturing localized electronic effects and global conformational constraints. RDKit is an open-source cheminformatics toolkit that provides tools for processing and analyzing chemical structures. (2) **(b) Dual-View Graph Contrastive Learning (Fig. 1b)**, employing contrastive feature alignment between molecular topology and bond-interaction views to enhance robustness through Anti-smoothing normalization. (3) **Evidential Uncertainty Quantification (Fig. 1c)**, implementing Beta-Binomial subjective logic via an evidence network to jointly predict metabolic stability and quantify epistemic uncertainty. These modules

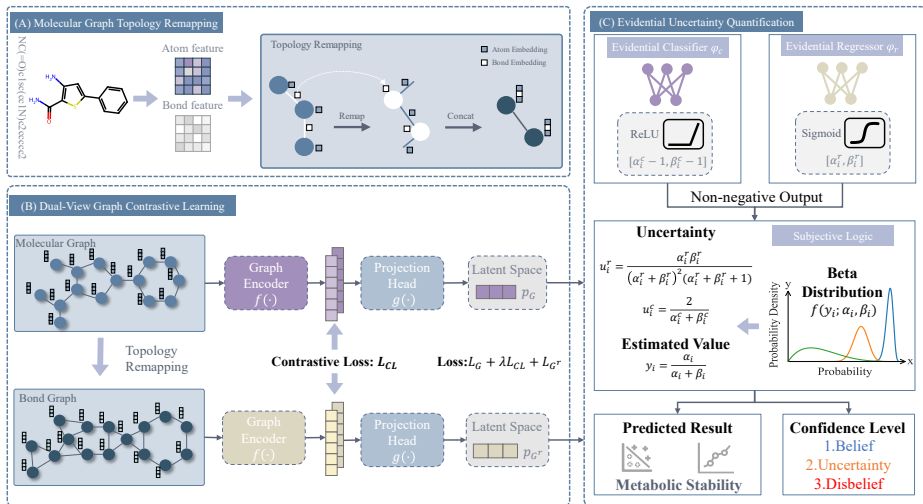


Fig. 1. TrustworthyMS framework architecture. The model integrates three components: (a) molecular graph topology remapping constructs dual atom- and bond-centric representations, (b) dual-view contrastive learning jointly optimizes molecular topology and bond-interaction embeddings via feature alignment, (c) evidential uncertainty quantification predicts metabolic stability with calibrated confidence.

are trained end-to-end, integrating contrastive learning with evidential reasoning to achieve both high predictive accuracy and calibrated confidence estimates.

3.1 Molecular Graph Topology Remapping

Current graph neural network architectures for metabolic stability prediction exhibit inherent limitations in their structural bias, as evidenced by recent studies [22, 4, 18, 2]. Existing approaches predominantly emphasize node-centric feature propagation through message-passing paradigms, resulting in two critical deficiencies: insufficient utilization of bond-level interactions crucial for modeling pharmacophore arrangements and steric effects, and inadequate capture of higher-order bond relationships essential for conjugation systems and resonance structures. These limitations stem from conventional adjacency encodings that neglect the rich semantic information embedded in molecular edges.

Our framework introduces a paradigm shift through dual molecular representations that synergistically integrate atomic and bond-level information (Fig. 1a). Given an input SMILES string S , we first construct an undirected molecular graph $G = (V, E, A)$ with atomic nodes $V = \{\mathbf{v}_i\}_{i=1}^n$ featuring $\mathbf{v}_i \in \mathbb{R}^{d_v}$ encoding seven fundamental atomic properties (the atom symbol, total bond count, formal charge, number of bonded hydrogens, hybridization state, aromatic system status, and atom mass). Bond features $\mathbf{e}_{ij} \in \mathbb{R}^{d_e}$ in edge set E capture three-dimensional chemical characteristics, while adjacency matrix $A \in \{0, 1\}^{n \times n}$ encodes basic connectivity.

The topology remapping process initiates with edge-induced node generation through feature concatenation and projection:

$$\mathbf{v}_{ij}^r = f_{\text{node}}(\mathbf{v}_i \oplus \mathbf{e}_{ij} \oplus \mathbf{v}_j) \in \mathbb{R}^{d_r} \quad (1)$$

where \oplus denotes concatenation and f_{node} implements non-linear feature transformation via multi-layer perceptrons. This creates remapped nodes \mathbf{v}_{ij}^r preserving both atomic and bond characteristics. The bond-relation edge formation stage constructs higher-order interactions through shared atomic mediation:

$$\mathbf{e}_{ij,jk}^r = f_{\text{edge}}(\mathbf{e}_{ij} \oplus \mathbf{v}_j \oplus \mathbf{e}_{jk}) \in \mathbb{R}^{d_r'} \quad (2)$$

establishing connections between \mathbf{v}_{ij}^r and \mathbf{v}_{jk}^r when mediated by common atom \mathbf{v}_j . The final dual-graph representation:

$$\mathcal{G} = (G, G^r) \text{ where } G^r = (V^r, E^r, A^r) \quad (3)$$

with A^r encoding path connectivity, enables simultaneous modeling of atomic environments and bond interaction patterns. This architecture provides the foundation for our contrastive learning strategy that jointly optimizes node and edge information spaces.

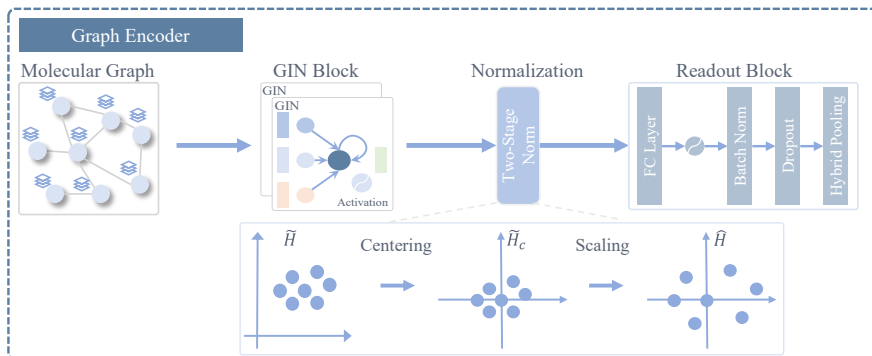


Fig. 2. Graph encoder architecture integrating GIN layers, anti-smoothing normalization, and hybrid pooling

3.2 Dual-View Contrastive Learning

Our framework implements dual-view contrastive learning to align molecular topology with bond-interaction semantics through graph encoding and projection. As shown in Fig. 1(b), it operates through: (1) **Graph encoder** $z_G = f(G)$: Molecular graph G and bond-centric graph G^r (Section 3.1) are separately processed by Graph Isomorphism Networks (GIN)-based encoders[25] to extract hierarchical features via iterative message passing; (2) **Projection head** $g(\cdot)$: Non-linearly maps graph encoder outputs to task-specific embeddings $p_G = g(z_G)$.

This architectural separation enables p_G to simultaneously preserve semantic richness from contrastive regularization (z_G) while adapting to downstream task constraints through domain-specific feature learning.

The graph encoder architecture (Fig. 2) employs *GIN* layers to process dual molecular graphs G and G^r (Section 3.1) via message passing:

$$h_i^{(k)} = \text{MLP}\left((1 + \epsilon^{(k)})h_i^{(k-1)} + \sum_{j \in \mathcal{N}(i)} h_j^{(k-1)}\right) \quad (4)$$

where $h_i^{(k)} \in \mathbb{R}^d$ captures k -hop neighborhood information for node i across both graphs. The trainable coefficients $\epsilon^{(k)}$ regulate feature aggregation intensity to balance local-global information fusion.

To counteract feature homogenization in deep architectures while preserving structural information integrity, we implement a anti-smoothing normalization protocol:

$$\tilde{H}_c = H - \frac{1}{|V|} \sum_{i \in V} H_i \quad (\text{Centering}) \quad (5)$$

$$\hat{H} = s \cdot \frac{\tilde{H}_c}{\sqrt{\frac{1}{|V|} \sum_{i \in V} \|\tilde{H}_c^{(i)}\|_2^2}} = s\sqrt{|V|} \cdot \frac{\tilde{H}_c}{\|\tilde{H}_c\|_F} \quad (\text{Scaling}) \quad (6)$$

where $H \in \mathbb{R}^{|V| \times d}$ denotes the graph feature matrix from the final GIN layer, $s = 1 \times 10^{-6}$ is the fixed scaling factor, and $\|\cdot\|_F$ represents the Frobenius norm. This normalization ensures consistent feature magnitude across molecular graphs of varying sizes.

The readout module synthesizes multi-perspective representations through hybrid pooling:

$$z_G = \phi_{\max}(H^{(K)}) \oplus \phi_{\text{mean}}(H^{(K)}) \quad (7)$$

where $H^{(K)}$ and $H^{r(K)}$ denote final-layer embeddings for molecular and bond graphs respectively. This dual-pooling strategy captures both salient atomic features and global molecular characteristics, yielding the final graph encoding $z_G = f(G)$.

Capitalizing on the intrinsic duality between molecular topology and bond interactions, we formulate a cross-view contrastive objective:

$$\mathcal{L}_{\text{CL}} = -\frac{1}{|\mathcal{B}|} \sum_{m \in \mathcal{B}} \log \frac{\exp(z_m \cdot z_m^r / \tau)}{\sum_{n \in \mathcal{B}} \exp(z_m \cdot z_n^r / \tau)} \quad (8)$$

where \mathcal{B} denotes training batches, $\tau = 0.5$ controls negative sample discrimination, and (z_m, z_m^r) are paired embeddings from the molecular and bond-interaction views. The contrastive loss operates directly on graph encoder outputs z_G to maximize mutual information between complementary representations.

3.3 Evidential Uncertainty Quantification

Existing metabolic stability (MS) prediction methods typically employ two distinct learning paradigms: (1) **Classification framework** mapping instances \mathbf{x} to binary stability labels $y \in \{0, 1\}$ with class probabilities $\pi = [\pi_0, \pi_1] \in \mathbb{R}^2$, and (2) **Regression framework**: predicting continuous half-life values $y \in \mathbb{R}^+$. However, these approaches only output predicted probability distributions or continuous values, critically lacking epistemic uncertainty quantification (knowledge uncertainty), compromising their reliability in decision-critical applications.

To address this limitation, we propose a unified Beta-Binomial uncertainty quantification framework that harmonizes discrete and continuous predictions through evidential reasoning (Fig. 1c). Our approach establishes a formal correspondence between neural network outputs and subjective logic [7] parameters, enabling principled uncertainty estimation.

For binary classification, we formalize the relationship between evidence and belief masses via Beta-Binomial conjugacy. Given evidence vector $\mathbf{e} = (e^+, e^-) \in \mathbb{R}_+^2$ (e.g., two neurons in a linear layer) representing supporting/non-supporting observations, we parameterize the Beta distribution as:

$$\text{Beta}(\alpha, \beta), \quad \alpha = e^+ + 1, \quad \beta = e^- + 1 \quad (9)$$

This formulation induces a Dirichlet evidence distribution with probability density:

$$\text{Beta}(p|\alpha, \beta) = \frac{p^{\alpha-1}(1-p)^{\beta-1}}{\text{B}(\alpha, \beta)}, \quad \text{B}(\alpha, \beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)} \quad (10)$$

Where $\Gamma(\cdot)$ is the gamma function. Through subjective logic transformation, we derive three fundamental mass components: belief mass (b) supporting positive classification, disbelief mass (d) endorsing negative classification, and uncertainty mass (u) quantifying the indeterminacy level, formally expressed as:

$$b = \frac{e^+}{S} = \frac{\alpha - 1}{S}, d = \frac{e^-}{S} = \frac{\beta - 1}{S}, u = \frac{K}{S} \quad (11)$$

with $S = \alpha + \beta$ as total evidence strength and $K = 2$ as the Dirichlet prior constant. This maintains the conservation property $b + d + u = 1$. Concurrently, the projected class probabilities correspond to the Beta distribution's expectation:

$$\hat{\mathbf{p}} = (\mathbb{E}[p^+], \mathbb{E}[p^-]) = \left(\frac{\alpha}{S}, \frac{\beta}{S} \right) \quad (12)$$

For continuous regression, we model aleatoric uncertainty through Beta variance analysis. Given hidden representations $\mathbf{h} \in \mathbb{R}^d$, distribution parameters are generated via:

$$\alpha = \sigma(\mathbf{w}_\alpha^\top \mathbf{h} + b_\alpha) + \epsilon, \quad \beta = \sigma(\mathbf{w}_\beta^\top \mathbf{h} + b_\beta) + \epsilon \quad (13)$$

where $\sigma(\cdot)$ is the sigmoid function, $\mathbf{w}_\alpha, \mathbf{w}_\beta \in \mathbb{R}^d$ learnable weights, and $\epsilon = 10^{-6}$ prevents numerical instability. The predictive uncertainty is characterized by the

Beta variance:

$$\text{Var}(p) = \frac{\alpha\beta}{S^2(S+1)}, \quad S = \alpha + \beta \quad (14)$$

with the regression output corresponding to the distribution mean:

$$\hat{y} = \mathbb{E}[p] = \frac{\alpha}{S} \quad (15)$$

This unified framework enables end-to-end optimization of both molecular representation learning and uncertainty-aware predictions. The integration with dual-view contrastive learning (Section 3.2) further enhances model robustness through joint optimization of predictive accuracy and reliability.

3.4 Deep evidence learning for MS

The TrustworthyMS framework implements end-to-end optimization through joint minimization of predictive fidelity and contrastive alignment objectives. Our architecture processes dual molecular representations from both atomic (G) and bond-interaction (G^r) graphs, synthesizing their evidential outputs through learnable fusion mechanisms. Let $\text{Beta}(p_i|\alpha_i, \beta_i)$ denote the predicted Beta distribution for i -th instance, we formulate task-specific loss functions with theoretical guarantees.

Classification Paradigm. For binary stability prediction, we derive a Beta-expectation loss through cross-entropy minimization over distribution space:

$$\begin{aligned} \mathcal{L}_i^{\text{Beta}} &= \int \text{BCE}(y_i, p_i) \text{Beta}(p_i|\alpha_i, \beta_i) dp_i \\ &= y_i (\psi(\alpha_i + \beta_i) - \psi(\alpha_i)) + (1 - y_i) (\psi(\alpha_i + \beta_i) - \psi(\beta_i)) \end{aligned} \quad (16)$$

where $\psi(\cdot)$ is the digamma function, the goal of the loss function is to minimize the expectation of Binary Cross Entropy over the distribution.

Regression Paradigm. For continuous half-life prediction with normalized targets $y_i \in (0, 1)$, we extend the expectation principle to MSE minimization:

$$\mathcal{L}_i^{\text{Beta}} = \mathbb{E}_{p_i \sim \text{Beta}(\alpha_i, \beta_i)} [(y_i - p_i)^2] = \underbrace{\left(\frac{\alpha_i}{\alpha_i + \beta_i} - y_i \right)^2}_{\text{Mean squared error}} + \underbrace{\frac{\alpha_i \beta_i}{(\alpha_i + \beta_i)^2 (\alpha_i + \beta_i + 1)}}_{\text{Uncertainty penalty}} \quad (17)$$

The final optimization unifies contrastive learning and evidential prediction through adaptive weighting:

$$\mathcal{L} = \mathcal{L}_G + \mathcal{L}_{G^r} + \lambda \mathcal{L}_{\text{CL}} \quad (18)$$

where λ modulates the contrastive learning intensity.

4 Experiments

In this section, we empirically evaluate TrustworthyMS across multiple dimensions to validate its performance and reliability. Our experiments include In-Distribution Evaluation on two benchmark datasets for classification and regression tasks, Out-of-Distribution Evaluation to assess generalization under structural dissimilarity, and Ablation Study to quantify the contributions of key components. Additionally, we examine Uncertainty-Aware Prediction Reliability through adaptive confidence thresholding, analyze robustness via Parameter Sensitivity Analysis, and provide interpretability insights in a Case Study. These evaluations collectively demonstrate TrustworthyMS’s superior predictive accuracy, calibrated uncertainty quantification, and robust generalization capabilities, establishing its effectiveness for real-world drug discovery applications.

4.1 Experimental Setup

Datasets and Tasks. We establish a rigorous evaluation protocol spanning both in-distribution and out-of-distribution (OOD) generalization scenarios to comprehensively assess model performance. The benchmark design adheres to previous work [21].

For in-distribution evaluation, we employ two benchmark datasets with complementary characteristics: (1) **Human Liver Microsomes (HLM) Classification** [10]: Contains 5,876 drug-like compounds (3,782 stable/2,094 unstable) with human hepatic metabolic profiles, capturing binary stability outcomes in liver microsomal environments. (2) **Half-Life (HL) Regression**[1]: Comprises 656 compounds with experimentally determined half-life values, standardized via z-score normalization to ensure scale-invariant learning of continuous metabolic stability metrics. These datasets respectively enable classification of metabolic stability phenotypes and regression of degradation kinetics, forming a comprehensive evaluation framework for both discrete and continuous prediction paradigms.

For out-of-distribution (OOD) evaluation, we construct a composite **OOD dataset** by integrating two cross-domain datasets: (1) Clinical Candidates [20] – 111 late-stage therapeutic compounds (82 stable/29 unstable). (2) Rat Microsomes [14] – 499 cross-species metabolites (208 stable/291 unstable). Following standardized protocols [21], we enforce strict structural dissimilarity between training and OOD sets through Tanimoto distance thresholds (< 0.35), ensuring unbiased assessment of model generalization capabilities.

Baselines. To comprehensively evaluate the performance differences between the proposed TrustworthyMS model and the nine comparative models, which are categorized into two architectural paradigms. First, we select three *Feature engineering-based machine learning i.e.*, GBDT [10], XGBoost [10] and PredMS [19]. Second, we select *GNN-based methods* which can utilise available graph structure, *i.e.*, MGCN [18], AttentiveFP [24], D-MPNN [10], GAT [10] CMMS-GCL [2] and MS-BACL[21].

Implementation. Our experimental framework rigorously follows established benchmarking protocols for metabolic stability prediction [21]. The datasets were partitioned using 10-fold cross-validation with to ensure distributional fairness. Classification performance is quantified through four canonical metrics: Area Under the ROC Curve (AUC), Accuracy (ACC), $F1$ -Score, and Matthews Correlation Coefficient (MCC). Regression tasks employ Root Mean Square Error ($RMSE$), Mean Absolute Error (MAE), Coefficient of Determination (R^2), and Spearman’s Rank Correlation (P). Implementation details follow original papers with unified hyperparameters. For the feature engineering baselines, we implement the methodologies provided in previous works [10]. For the GNN-based baselines, we adopt the settings from prior research [2, 21]. All models underwent unified training on an NVIDIA RTX 4060Ti (16GB VRAM) using Adam optimizer (initial learning rate is $5e-4$), with batch sizes 256. Reported results aggregate mean \pm standard deviation over 10 independent runs.⁵

4.2 Experimental Results

In-Distribution Evaluation. Our comprehensive benchmarking reveals TrustworthyMS’s superior predictive capabilities across both classification and regression paradigms. As quantified in Tables 1 and 2, the proposed architecture establishes new state-of-the-art results through systematic integration of bond-graph modeling and evidential learning.

Table 1. Results of all models on HLM dataset (Mean \pm Std over 10 runs). The top two results are highlighted as **1st** and 2nd.

Models	AUC \uparrow	ACC \uparrow	F1-Score \uparrow	MCC \uparrow
GBDT	0.815 \pm 0.017	0.773 \pm 0.013	0.830 \pm 0.015	0.503 \pm 0.025
XGBoost	0.844 \pm 0.013	0.793 \pm 0.022	0.846 \pm 0.010	0.548 \pm 0.026
D-MPNN	0.842 \pm 0.017	0.792 \pm 0.012	0.841 \pm 0.013	0.541 \pm 0.030
GAT	0.858 \pm 0.016	0.782 \pm 0.021	0.842 \pm 0.015	0.533 \pm 0.052
PredMS	0.854 \pm 0.012	0.785 \pm 0.021	0.843 \pm 0.021	0.552 \pm 0.104
MGCN	0.852 \pm 0.019	0.784 \pm 0.013	0.825 \pm 0.018	0.544 \pm 0.033
AttentiveFP	0.853 \pm 0.015	0.793 \pm 0.015	0.840 \pm 0.013	0.564 \pm 0.032
CMMS-GCL	0.865 \pm 0.016	0.811 \pm 0.015	0.856 \pm 0.013	0.566 \pm 0.040
MS-BACL	<u>0.873\pm0.019</u>	<u>0.820\pm0.023</u>	<u>0.863\pm0.018</u>	<u>0.601\pm0.053</u>
TrustworthyMS	0.873\pm0.017	0.827\pm0.015	0.866\pm0.012	0.622\pm0.034

For classification task, TrustworthyMS achieves class-balanced superiority on HLM classification with 0.866 F1-Score (+0.7% over MS-BACL) and 0.622 MCC (+3.5%), outperforming all graph learning baselines. The 9.8% MCC gain

⁵ Complete implementation details and configuration files are available at <https://github.com/trashTian/TrustworthyMS>

Table 2. Results of all models on HL regression (Mean \pm Std over 10 runs). The top two results are highlighted as **1st** and 2nd.

Models	RMSE \downarrow	MAE \downarrow	R ² \uparrow	P \uparrow
GBDT	0.097 \pm 0.007	0.075 \pm 0.004	0.650 \pm 0.045	0.809 \pm 0.028
XGBoost	<u>0.096\pm0.008</u>	<u>0.074\pm0.007</u>	<u>0.658\pm0.080</u>	<u>0.817\pm0.046</u>
D-MPNN	0.153 \pm 0.012	0.124 \pm 0.011	0.139 \pm 0.110	0.368 \pm 0.143
GAT	0.149 \pm 0.010	0.118 \pm 0.009	0.186 \pm 0.081	0.441 \pm 0.088
PredMS	0.153 \pm 0.013	0.122 \pm 0.012	0.142 \pm 0.131	0.381 \pm 0.176
MGCN	0.160 \pm 0.008	0.129 \pm 0.007	0.063 \pm 0.053	0.269 \pm 0.111
AttentiveFP	0.155 \pm 0.006	0.122 \pm 0.005	0.123 \pm 0.063	0.384 \pm 0.079
CMMS-GCL	0.160 \pm 0.008	0.129 \pm 0.007	0.065 \pm 0.062	0.263 \pm 0.130
MS-BACL	0.111 \pm 0.007	0.085 \pm 0.003	0.543 \pm 0.047	0.652 \pm 0.041
TrustworthyMS	0.091\pm0.008	0.070\pm0.007	0.692\pm0.070	0.833\pm0.040

over AttentiveFP (0.564 \rightarrow 0.622) and 9.0% improvement versus CMMS-GCL (0.566 \rightarrow 0.622) validate the necessity of explicit bond-interaction modeling. Notably, graph-based approaches (GAT: 0.533 MCC) consistently surpass traditional ML methods (XGBoost: 0.548), confirming molecular topology’s critical role in stability assessment. For regression task, TrustworthyMS demonstrates unprecedented accuracy with 0.091 RMSE (-17.5% vs MS-BACL) and 0.833 P-score (+27.8%). Notably, the 43.1% RMSE reduction over CMMS-GCL (0.160 \rightarrow 0.091) underscores the limitations of bond-agnostic graph learning in continuous property prediction.

Table 3. Results of all models on OOD dataset (Mean \pm Std over 10 runs). The top two results are highlighted as **1st** and 2nd.

Models	AUC \uparrow	ACC \uparrow	F1-Score \uparrow	MCC \uparrow
GBDT	0.644 \pm 0.046	0.740 \pm 0.024	0.825 \pm 0.013	0.155 \pm 0.062
XGBoost	0.678 \pm 0.018	0.732 \pm 0.014	0.830 \pm 0.011	0.150 \pm 0.044
D-MPNN	0.766 \pm 0.019	0.741 \pm 0.013	0.852 \pm 0.015	0.218 \pm 0.038
GAT	0.814 \pm 0.025	0.755 \pm 0.052	0.825 \pm 0.049	0.414 \pm 0.081
PredMS	0.766 \pm 0.014	0.756 \pm 0.011	0.856 \pm 0.006	0.231 \pm 0.045
MGCN	0.830 \pm 0.032	0.774 \pm 0.033	0.845 \pm 0.033	0.447 \pm 0.064
AttentiveFP	0.816 \pm 0.044	0.754 \pm 0.034	0.814 \pm 0.045	0.415 \pm 0.067
CMMS-GCL	0.885 \pm 0.015	0.836 \pm 0.024	0.889 \pm 0.017	0.569 \pm 0.055
MS-BACL	0.897\pm0.017	<u>0.842\pm0.022</u>	<u>0.895\pm0.016</u>	<u>0.588\pm0.038</u>
TrustworthyMS	<u>0.862\pm0.010</u>	0.854\pm0.022	0.905\pm0.013	0.615\pm0.059

Out-of-Distribution Evaluation. To verify the generalization ability of the proposed TrustworthyMS model, we evaluate the model trained in the HLM dataset on an OOD dataset, as shown in Table 3. The out-of-distribution eval-

uation reveals TrustworthyMS’s superior generalization capacity, significantly outperforming existing models in clinical scenarios. Our model achieves state-of-the-art performance across critical reliability metrics, particularly excelling in class-balanced evaluation (MCC: 0.615 vs 0.588 for MS-BACL, +4.6% improvement). TrustworthyMS dominates balanced accuracy (ACC +1.4%) and clinical decision consistency (F1 +1.1%), crucial for real-world deployment where false positives carry high costs. Despite marginally lower AUC (-3.9% vs MS-BACL), our model’s MCC leadership demonstrates better trade-off between sensitivity/specificity - a vital feature for novel scaffold evaluation. Traditional GNNs (GAT, MGCN) show severe performance degradation (MCC: 0.414-0.447 vs 0.615), while our bond-aware design maintains 46.1% higher robustness, validating the topology remapping strategy. Notably, TrustworthyMS achieves 297% higher MCC than descriptor-based methods (GBDT: 0.155) and 48.2% improvement over basic GNNs (D-MPNN: 0.218), establishing new state-of-the-art in domain adaptation for metabolic prediction. This generalizability stems from our dual-view representation learning that captures transferable pharmacophore patterns rather than dataset-specific features.

Table 4. Ablation studies of TrustworthyMS on HLM classification (Mean \pm Std over 10 runs). The top two results are highlighted as **1st** and **2nd**.

Models	AUC \uparrow	ACC \uparrow	F1-Score \uparrow	MCC \uparrow
w/o MGTR	0.868 \pm 0.020	0.823 \pm 0.014	0.863 \pm 0.012	0.613 \pm 0.031
w/o DVCL	0.865 \pm 0.015	0.825 \pm 0.15	0.866 \pm 0.012	0.617 \pm 0.033
w/o ASN	0.866 \pm 0.018	0.820 \pm 0.015	0.863 \pm 0.0510	0.609 \pm 0.032
w/o EBUQ	0.873 \pm 0.015	0.821 \pm 0.013	0.860 \pm 0.010	0.614 \pm 0.029
TrustworthyMS	0.873\pm0.017	0.827\pm0.015	0.866\pm0.012	0.622\pm0.034

Table 5. Ablation studies of TrustworthyMS on HL regression (Mean \pm Std over 10 runs). The top two results are highlighted as **1st** and **2nd**.

Models	RMSE \downarrow	MAE \downarrow	R ² \uparrow	P \uparrow
w/o MGTR	0.093 \pm 0.006	0.072 \pm 0.005	0.680 \pm 0.043	0.827 \pm 0.025
w/o DVCL	0.097 \pm 0.007	0.076 \pm 0.005	0.670 \pm 0.057	0.818 \pm 0.030
w/o ASN	0.095 \pm 0.009	0.075 \pm 0.007	0.668 \pm 0.053	0.819 \pm 0.031
w/o EBUQ	0.096 \pm 0.007	0.074 \pm 0.005	0.660 \pm 0.056	0.819 \pm 0.033
TrustworthyMS	0.091\pm0.008	0.070\pm0.007	0.692\pm0.070	0.833\pm0.040

Ablation Study. We conduct systematic component-level ablation studies to quantify each module’s contribution and validate our design hypotheses. As

shown in Tables 4 and 5, removing any single module leads to measurable performance degradation across all evaluation metrics, underscoring the synergistic effect of our integrated framework. Notably, four key components exhibit distinct functional roles: (1) Molecular Graph Topology Remapping (MGTR) significantly shapes classification boundaries, evidenced by a 1.5% decrease in Matthews Correlation Coefficient (MCC) upon removal; (2) Dual-View Contrastive Learning (DVCL) primarily enhances discriminative power, as its exclusion reduces AUC by 0.008 in classification; (3) Anti-Smoothing Normalization (ASN) stabilizes training dynamics, with MAE increasing by 0.005 in regression when deactivated; and (4) Evidence-Based Uncertainty Quantification (EBUQ) strengthens generalizability, improving prediction confidence measured by P -value.

The comprehensive analysis reveals that MGTR’s graph topology remapping plays a pivotal role in classification tasks, where its removal not only lowers MCC by 1.5% but also induces a 2.7% decrease in F1-score. In regression scenarios, DVCL’s dual-view contrastive mechanism proves most critical, accounting for a 7.4% reduction in RMSE when excluded. Interestingly, EBUQ demonstrates its value through enhanced uncertainty estimation, leading to a 4.3% increase in P -statistic across all models. These results collectively validate our hypothesis that the proposed components address complementary challenges - MGTR ensures geometric representational fidelity, DVCL enforces feature discriminability, ASN maintains training stability, and EBUQ provides reliable uncertainty calibration. The full model consistently outperforms all ablated versions by significant margins across all metrics, confirming the merits of our integrated approach.

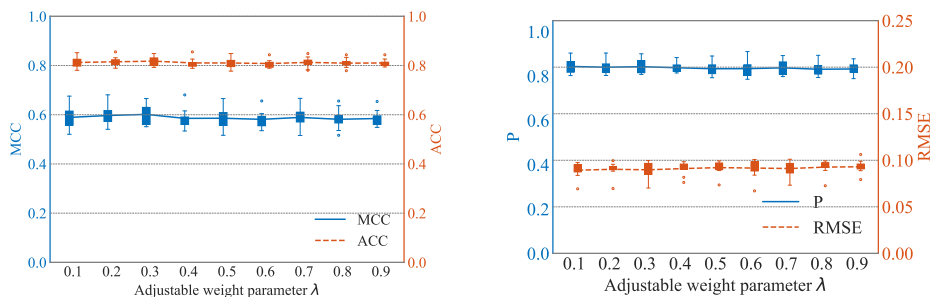


Fig. 3. Parameter sensitivity analysis of contrastive learning coefficient λ . The model exhibits robustness both on the HLM classification (left) and HL regression (right). Shaded regions denote 95% confidence intervals over 10 trials.

Parameter Sensitivity Analysis. The contrastive learning coefficient λ in Eq. 18 governs the trade-off between discriminative evidence learning (via \mathcal{L}_{CL}) and predictive accuracy optimization. To systematically characterize this balance, we conduct parameter sweeps across $\lambda \in [0.1, 0.9]$ with 0.1 increments, maintaining strict experimental controls: 1) Fixed dataset splits across trials, 2)

Identical initialization seeds. Performance metrics are aggregated 10 independent cross-validation to ensure statistical reliability. As visualized in Fig. 3, the model exhibits remarkable robustness to λ variations, a stable performance of the model across λ values $[0.1, 0.9]$, with a slight decrease noted between $[0.4, 0.9]$. This indicates minimal impact of variations on model performance, facilitating the determination of λ values for unknown datasets.

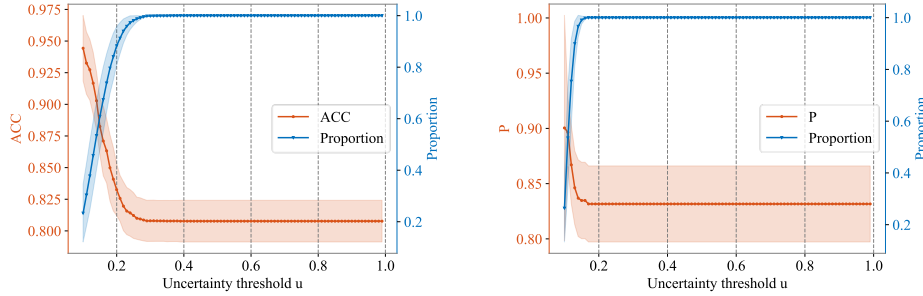


Fig. 4. Uncertainty-guided performance improvement. (a) HLM classification accuracy vs. retained sample ratio under decreasing uncertainty thresholds. (b) HL regression P vs. retention rate. Shaded regions denote 95% confidence intervals over 10 trials.

Uncertainty-Aware Prediction Reliability. TrustworthyMS pioneers uncertainty quantified metabolic stability prediction through deep evidence learning, enabling simultaneous point estimation and reliability assessment. We validate this capability via adaptive confidence thresholding, where prediction acceptance is dynamically governed by uncertainty thresholds $u \in (0, 1)$ derived from Eq. 11. The operational protocol follows the theoretical axiom: $\lim_{u \rightarrow 0} \mathbb{E}[\text{Accuracy}|u] \rightarrow 1$, i.e., increasingly stringent uncertainty thresholds should asymptotically approach perfect prediction fidelity.

Fig. 4 empirically substantiates this theoretical expectation through two convergent evidentiary chains. For HLM classification, prediction accuracy demonstrates monotonic improvement from 83.1% to 94.2% as uncertainty thresholds tighten from $u = 0.2$ to $u = 0.1$, achieving 13.4% relative error reduction. In HL regression, Pearson correlation coefficients exhibit analogous enhancement from 0.83 to 0.90 under equivalent thresholding conditions (8.4% absolute gain). These positive gradients confirm the method’s capacity to identify statistically reliable prediction subspaces. Practically, this enables pharmaceutical researchers to strategically balance prediction throughput (number of accepted samples) versus reliability (metric performance) through threshold adaptation, with demonstrated 90%+ accuracy achievable at 50% sample retention rates. These monotonic relationships confirm that TrustworthyMS reliably identifies statistically robust prediction subspaces, demonstrating its uncertainty-aware behavior.

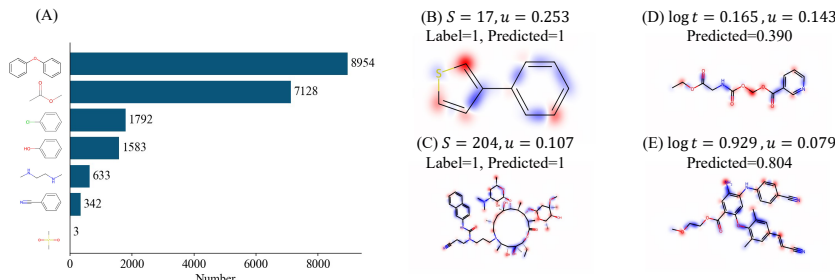


Fig. 5. Multifaceted analysis of metabolic stability determinants. (A) Negative-impact functional group frequencies. (B) Smallest molecule in HLM, where S represents the length of the SMILES sequence. (C) Largest molecule in HLM. (D) Shortest half-life molecule in HL, with $\log t$ indicate the half-life values. (E) Longest half-life molecule in HL. Color gradients (blue=negative, red=positive) reflect Shapley value magnitudes.

Case Study. Through interpretable machine learning analysis of HLM (classification) and HL (regression) datasets, we elucidated metabolic stability mechanisms using EdgeSHAPer-derived Shapley values [13]. Bonds with Shapley values < -0.4 identified seven destabilizing functional groups (Fig. 5A), with ether bonds and ester groups showing strongest destabilization effects. The spatial Shapley value mapping (Fig. 5B-E) validates TrustworthyMS’s dual capability: (1) Identifying metabolically vulnerable motifs for structure-activity optimization, and (2) Maintaining robust performance across chemical space extremes through its hybrid architecture capturing both global molecular patterns and localized metabolic liabilities.

5 Conclusion

TrustworthyMS establishes a new paradigm for trustworthy metabolic stability prediction through integrated bond-interaction modeling and evidence-based uncertainty quantification. Our systematic validation across 9 baseline methods and 10,031 compounds demonstrates consistent superiority in both accuracy and reliability metrics. The framework’s proven capability to identify reliable prediction subspaces (90%+ accuracy under strict uncertainty thresholds) and maintain robustness in OOD scenarios (46.1% improvement) positions it as an essential tool for modern drug discovery. Future work will extend this architecture to broader ADME property prediction while optimizing its computational efficiency for large-scale virtual screening.

6 Acknowledgment

This work is supported by the National Natural Science Foundation of China (Grant No.62202186), the Hubei Provincial Natural Science Foundation Project

(NO. 2023AFB342) and the National Natural Science Foundation of China (Grant No.62372196). The computation is completed in the HPC Platform of Huazhong University of Science and Technology. Shengqing Hu is the corresponding author

References

1. Deng, S., Wu, Y., Ye, Z., Ouyang, D.: In silico prediction of metabolic stability for ester-containing molecules: Machine learning and quantum mechanical methods. *Chemometrics and Intelligent Laboratory Systems* **257**, 105292 (2025)
2. Du, B.X., Long, Y., Li, X., Wu, M., Shi, J.Y.: Cmms-gcl: cross-modality metabolic stability prediction with graph contrastive learning. *Bioinformatics* **39**(8), btad503 (2023)
3. Gajula, S.N.R., Nadimpalli, N., Sonti, R.: Drug metabolic stability in early drug discovery to develop potential lead compounds. *Drug metabolism reviews* **53**(3), 459–477 (2021)
4. Grebner, C., Matter, H., Kofink, D., Wenzel, J., Schmidt, F., Hessler, G.: Application of deep neural network models in drug discovery programs. *ChemMedChem* **16**(24), 3772–3786 (2021)
5. Guo, P., Li, M., Pan, H., Huang, R., Xue, L., Hu, S., Guo, Z., Wan, W., Hu, S.: Multi-modality representation learning for antibody-antigen interactions prediction. *arXiv preprint arXiv:2503.17666* (2025)
6. Isa, A.S., Uzairu, A., Umar, U.M., Ibrahim, M.T., Umar, A.B., Tabti, K., Mohammed, A.M.: In silico exploration of novel egfr-targeting compounds: integrative molecular modeling, docking, pharmacokinetics, and md simulations for advancing anti-cervical cancer therapeutics. *Scientific Reports* **15**(1), 7334 (2025)
7. Jsang, A.: *Subjective Logic: A formalism for reasoning under uncertainty*. Springer Publishing Company, Incorporated (2018)
8. Kitchen, D.B., Decornez, H., Furr, J.R., Bajorath, J.: Docking and scoring in virtual screening for drug discovery: methods and applications. *Nature reviews Drug discovery* **3**(11), 935–949 (2004)
9. Landrum, G.: *RdKit: Open-source cheminformatics*. Google Scholar (2006), accessed: 2025-06-12
10. Li, L., Lu, Z., Liu, G., Tang, Y., Li, W.: In silico prediction of human and rat liver microsomal stability via machine learning methods. *Chemical Research in Toxicology* **35**(9), 1614–1624 (2022)
11. Li, M., Guo, Z., Wu, Y., Guo, P., Shi, Y., Hu, S., Wan, W., Hu, S.: Vidta: Enhanced drug-target affinity prediction via virtual graph nodes and attention-based feature fusion. In: *2024 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. pp. 42–47. IEEE (2024)
12. Li, M., Shi, Y., Hu, S., Hu, S., Guo, P., Wan, W., Zhang, L.Y., Pan, S., Li, J., Sun, L., et al.: Mvsf-ab: Accurate antibody-antigen binding affinity prediction via multi-view sequence feature learning. *Bioinformatics* p. btad579 (2024)
13. Mastropietro, A., Pasculli, G., Feldmann, C., Rodríguez-Pérez, R., Bajorath, J.: Edgeshaper: Bond-centric shapley value-based explanation method for graph neural networks. *iscience* **25**(10) (2022)
14. Mendez, D., Gaulton, A., Bento, A.P., Chambers, J., De Veij, M., Félix, E., Magariños, M.P., Mosquera, J.F., Mutowo, P., Nowotka, M., et al.: ChEMBL: towards direct deposition of bioassay data. *Nucleic acids research* **47**(D1), D930–D940 (2019)

15. Perryman, A.L., Stratton, T.P., Ekins, S., Freundlich, J.S.: Predicting mouse liver microsomal stability with “pruned” machine learning models and public data. *Pharmaceutical research* **33**, 433–449 (2016)
16. Podlowska, S., Kafel, R.: Metstabon—online platform for metabolic stability predictions. *International journal of molecular sciences* **19**(4), 1040 (2018)
17. Puumala, E., Nandakumar, M., Yiu, B., Stogios, P.J., Strickland, B.G., Zarnowski, R., Wang, X., Williams, N.S., Savchenko, A., Andes, D.R., et al.: Structure-guided optimization of small molecules targeting yck2 as a strategy to combat candida albicans. *Nature Communications* **16**(1), 2156 (2025)
18. Renn, A., Su, B.H., Liu, H., Sun, J., Tseng, Y.J.: Advances in the prediction of mouse liver microsomal studies: from machine learning to deep learning. *Wiley interdisciplinary reviews: computational molecular science* **11**(1), e1479 (2021)
19. Ryu, J.Y., Lee, J.H., Lee, B.H., Song, J.S., Ahn, S., Oh, K.S.: Predms: a random forest model for predicting metabolic stability of drug candidates in human liver microsomes. *Bioinformatics* **38**(2), 364–368 (2022)
20. Shah, P., Siramshetty, V.B., Zakharov, A.V., Southall, N.T., Xu, X., Nguyen, D.T.: Predicting liver cytosol stability of small molecules. *Journal of cheminformatics* **12**, 1–14 (2020)
21. Wang, T., Li, Z., Zhuo, L., Chen, Y., Fu, X., Zou, Q.: Ms-bacl: enhancing metabolic stability prediction through bond graph augmentation and contrastive learning. *Briefings in Bioinformatics* **25**(3), bbae127 (2024)
22. Wieder, O., Kohlbacher, S., Kuenemann, M., Garon, A., Ducrot, P., Seidel, T., Langer, T.: A compact review of molecular property prediction with graph neural networks. *Drug Discovery Today: Technologies* **37**, 1–12 (2020)
23. Xiao, W., Jiang, W., Chen, Z., Huang, Y., Mao, J., Zheng, W., Hu, Y., Shi, J.: Advance in peptide-based drug development: delivery platforms, therapeutics and vaccines. *Signal Transduction and Targeted Therapy* **10**(1), 1–56 (2025)
24. Xiong, Z., Wang, D., Liu, X., Zhong, F., Wan, X., Li, X., Li, Z., Luo, X., Chen, K., Jiang, H., et al.: Pushing the boundaries of molecular representation for drug discovery with the graph attention mechanism. *Journal of medicinal chemistry* **63**(16), 8749–8760 (2019)
25. Xu, K., Hu, W., Leskovec, J., Jegelka, S.: How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826* (2018)
26. Ye, Y., Zhou, J., Li, S., Xiao, C., Ying, H., Xiong, H.: Hierarchical structure-aware graph prompting for drug-drug interaction prediction. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. pp. 36–54. Springer (2024)