CSIL: A CHIRALITY-STRUCTURAL INTERACTION LEARNING FRAMEWORK FOR CHIRAL DRUG-DRUG INTERACTION PREDICTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Accurate prediction of drug-drug interactions (DDIs) is essential for drug safety and healthcare research. Many drugs are chiral, and their molecules can exist in non-superimposable mirror-image forms. These forms share the same molecular formula and atomic connectivity but differ in chiral configuration. When two chiral drugs interact, different configuration combinations may alter biological recognition and lead to different interaction outcomes. This motivates the study of Chiral Drug-Drug Interaction (CDDI) prediction. However, many existing DDI methods focus on molecular connectivity and substructures without explicitly considering chirality, which may limit their ability to capture differences caused by chiral configuration and accurately predict CDDIs. Moreover, their substructure-based explanations do not explicitly account for chirality-related information, making it difficult to explain how chiral configuration affects CDDIs. To overcome these limitations, we propose a Chirality-Structural Interaction Learning (CSIL) framework for CDDI prediction. CSIL integrates hierarchical motif representations with configuration-sensitive chirality representations and employs chirality-motif collaborative cross-attention to model their cross-drug dependencies. It further enables attention-based analysis of interaction-relevant motifs and configuration-associated differences in motif prioritization. Extensive experiments on a CDDI dataset curated from DrugBank demonstrate that CSIL achieves the best overall rank against representative baselines across four settings. The source code and datasets are submitted as Supplementary Materials for Reproducibility.

1 INTRODUCTION

Accurate prediction of drug-drug interactions (DDIs) is essential for drug safety and healthcare research Lu et al. (2026); Hashir et al. (2026). Many drugs are chiral, and their molecules have mirror-image forms that cannot overlap exactly Karayavuz & Kirmizioglu (2026). These forms share the same molecular formula and atomic connectivity but differ in chiral configuration Chawla (2025); Okunlola & Aleku (2026). Such configurational differences can alter biological recognition and processing, potentially resulting in different interaction outcomes when two chiral drugs interact Karayavuz & Kirmizioglu (2026); Okunlola & Aleku (2026). As illustrated in Figure 1, R-warfarin and S-warfarin share the same molecular connectivity but differ in chiral configuration. When paired with the same chiral drug, Abiraterone, the two drug pairs exhibit different interaction annotations Knox et al. (2024). However, molecular connectivity alone cannot capture configuration-associated interaction differences. This limitation motivates the study of Chiral Drug-Drug Interaction (CDDI) prediction, which focuses on DDI prediction between chiral drug pairs.

Molecular representation is fundamental to structure-based DDI prediction. Early approaches, such as DeepDDI, used molecular structural similarity profiles to predict interaction types Ryu et al. (2018). Subsequent studies adopted finer-grained representations to model interactions between molecular components. SSI-DDI models pairwise interactions between drug substructures Nyamabo et al. (2021). More recent methods, including MeTDDI and MS-DDI, employ motif-based and multi-scale interaction modeling, respectively, while also supporting structural interpretation of predicted DDIs Zhong et al. (2024); Zou et al. (2026). In parallel, LSA-DDI combines two-dimensional topological features with three-dimensional coordinate, distance, and angle encodings

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

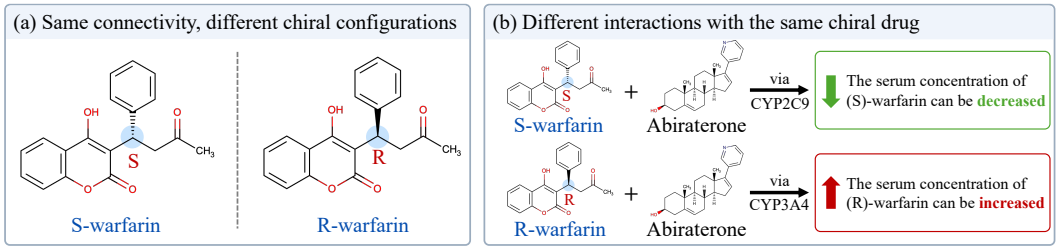


Figure 1: **Chiral configuration may alter drug-drug interactions**

to incorporate spatial information into DDI prediction Wang et al. (2025b). More recently, DSMV-DDI incorporates stereochemical molecular visual representations derived from three-dimensional conformers for DDI prediction Chen et al. (2026). These developments enrich DDI representations with increasingly detailed structural, spatial, and stereochemical information.

However, richer molecular representations do not necessarily capture differences caused by chiral configuration. Many existing DDI methods still primarily rely on molecular connectivity and substructure information without explicitly considering chirality Wang et al. (2025a); Shi et al. (2026). Since molecular connectivity and substructure composition can be identical across different chiral configurations, such representations may fail to distinguish configuration-associated differences and thus limit accurate CDDI prediction Gaiński et al. (2023); Papusha & Leonhard (2026). Moreover, existing substructure-based explanations mainly highlight molecular components associated with predicted interactions Zhong et al. (2024); Zou et al. (2026). Because the same connectivity-defined substructures can be shared across different chiral configurations Adams et al. (2022), these explanations do not explicitly reveal how chiral configuration affects CDDIs Yan et al. (2025). These limitations highlight the need for CDDI models to capture configuration-associated interaction differences while providing structural and chirality-related explanations.

In this study, we propose a Chirality-Structural Interaction Learning (CSIL) framework for CDDI prediction. CSIL jointly learns hierarchical motif representations and configuration-sensitive chirality representations from molecular conformers, and hierarchically aggregates them across conformers and candidate configurations to obtain drug-level representations. A chirality-motif collaborative cross-attention mechanism further models cross-drug dependencies between structural and chirality-related features, producing pair-aware representations for CDDI prediction. Beyond prediction, CSIL leverages the learned interaction attention to highlight interaction-relevant motifs and characterize configuration-associated differences in motif prioritization.

The primary contributions are summarized as follows:

- **CDDI prediction.** We formulate the Chiral Drug-Drug Interaction (CDDI) prediction problem for chiral drug pairs, explicitly incorporating chiral configuration information, and construct a CDDI dataset from DrugBank for systematic evaluation.
- **Chirality-structural interaction learning.** We propose CSIL to integrate hierarchical motif representations with configuration-sensitive chirality representations and employ chirality-motif collaborative cross-attention to capture their cross-drug dependencies for CDDI prediction.
- **CDDI interpretability analysis.** We analyze the learned interaction attention to identify interaction-relevant motifs and characterize configuration-associated differences in motif prioritization under matched drug pairs.

2 Related Works

Drug-Drug Interaction Prediction. Existing structure-based DDI methods have progressively evolved from global molecular descriptors to finer-grained substructure, motif, multi-scale, and three-dimensional representations Ahmad et al. (2025); Liu et al. (2026). Early methods such as CASTER extracted frequent substructures through sequence pattern mining and embedded the resulting functional representation using an autoencoder Huang et al. (2020). Subsequent methods

learned interaction-relevant substructures directly from molecular graphs. SSI-DDI represented intermediate node features from successive GAT layers as substructures and matched them across drug pairs through co-attention Nyamabo et al. (2021), while GMPNN-CS further modeled substructures of different sizes and shapes using gated message passing Nyamabo et al. (2022). SA-DDI estimated the interaction relevance of individual substructures Yang et al. (2022), whereas MSAN and DGNN-DDI explicitly modeled pairwise dependencies between substructures of two drugs through similarity-based or attention-based interaction mechanisms Zhu et al. (2022); Ma & Lei (2023). FG-DDI incorporated functional-group knowledge into intra- and inter-molecular interaction modeling Zhou & Uddin (2026). More recent methods organize molecular representations around chemically meaningful motifs and richer structural information. MeTDDI modeled motif-based graphs with local-global self-attention Zhong et al. (2024), ImageDDI combined motif sequences with global molecular image features He et al. (2026), and MS-DDI integrated atom-level and motif-level representations through hierarchical cross-attention Zou et al. (2026). HDN-DDI further employed multi-level molecular graphs for hierarchical structural representation Sun & Zheng (2025). This progression has also extended to explicit three-dimensional structure, with LSA-DDI incorporating molecular coordinates, distances, and angles into DDI representation learning Wang et al. (2025b).

Despite these advances, many existing DDI methods still rely primarily on molecular connectivity and substructures without explicitly considering chirality. These methods may struggle to distinguish drugs with identical connectivity but different chiral configurations, limiting their ability to capture configuration-associated differences and accurately predict CDDIs.

Interpretability in DDI Prediction. Interpretability in DDI prediction is crucial for elucidating the pharmacological mechanisms underlying drug interactions Sun et al. (2025); Ru et al. (2025); Wu et al. (2026). Existing methods mainly interpret predictions by identifying relevant molecular structures or relational evidence. SA-DDI attributes predictions to interaction-relevant substructures Yang et al. (2022), while MeTDDI localizes chemically meaningful motifs associated with specific DDI events Zhong et al. (2024). EmerGNN provides relational explanations by identifying important paths and biomedical concepts between drug pairs in a biomedical network Zhang et al. (2023). More recently, MS-DDI employs atom- and motif-level attention to highlight structural regions and functional groups relevant to DDI predictions Zou et al. (2026).

Despite providing useful structural or relational insights, these methods lack explicit consideration of chirality, failing to explain how chiral configuration contributes to differences in CDDIs.

3 PRELIMINARIES

3.1 MOLECULAR STRUCTURE AND CHIRALITY

A molecular instance is represented as an atom-level graph $G = (\mathcal{V}^A, \mathcal{E}^A)$, where $\mathcal{V}^A$ and $\mathcal{E}^A$ denote atoms and chemical bonds, respectively. The graph is further decomposed into chemically meaningful BRICS motifs $\mathcal{M} = \{M_k\}_{k=1}^{K}$. Let $\mathcal{Z}$ denote the set of potential chiral centers, whose spatial arrangements determine the chiral configuration. Molecules with identical connectivity can differ in configuration, whereas conformers correspond to different spatial arrangements under the same configuration.

3.2 CHIRAL MOLECULAR ENSEMBLE

A drug structure may contain multiple chiral centers whose configurations are only partially assigned Southan et al. (2013); Orsi & Reymond (2026). For each drug d, let n_d^{pot} and n_d^{asg} denote the numbers of potential and assigned chiral centers, respectively. Based on assignment completeness and relationships among drugs sharing the same parent structure, we categorize each drug as *fully specified*, *inferred-racemate*, *partial*, or *unspecified*. A drug is considered fully specified when all potential chiral centers are assigned and partial when only a subset is assigned. A fully unassigned drug is labeled as inferred-racemate when a fully assigned form with the same molecular connectivity exists, and unspecified otherwise. The inferred-racemate label is based on molecular records rather than experimentally verified racemic composition.

Each drug d is associated with a candidate chiral configuration set $\mathcal{S}_d$. For fully specified drugs, $\mathcal{S}_d$ contains only the specified configuration; otherwise, candidate configurations are generated by

resolving unassigned chiral centers while preserving existing assignments. Each configuration $s \in \mathcal{S}_d$ may further admit multiple conformers $\mathcal{C}_{d,s}$. Accordingly, drug d is represented as

$$\mathcal{E}_d = \{\mathcal{G}_{d,s,c} \mid s \in \mathcal{S}_d,\ c \in \mathcal{C}_{d,s}\}, \tag{1}$$

where $\mathcal{G}_{d,s,c}$ denotes a molecular instance under configuration s and conformer c. This hierarchy aggregates conformers within a configuration before combining configurations.

3.3 Problem Formulation

Let $\mathcal{V}_{\text{tr}}$ denote the drugs observed in training and $\Pi_{\text{tr}} = \pi(\mathcal{V}_{\text{tr}})$ the set of parent structures observed in training. We define

$$\mathcal{V}_k = \mathcal{V}_{\text{tr}}, \quad \mathcal{V}_i = \{d \in \mathcal{V} \setminus \mathcal{V}_{\text{tr}} \mid \pi(d) \in \Pi_{\text{tr}}\}, \quad \mathcal{V}_p = \{d \in \mathcal{V} \mid \pi(d) \notin \Pi_{\text{tr}}\} \tag{2}$$

where $\mathcal{V}_k$, $\mathcal{V}_i$, and $\mathcal{V}_p$ denote known drugs, unseen chiral forms of known parents, and drugs with unseen parent structures, respectively.

We consider four prediction tasks:

$$\begin{aligned} \mathcal{T}_1 &= \mathcal{V}_k \times \mathcal{V}_k, \quad \mathcal{T}_2 = (\mathcal{V}_k \times \mathcal{V}_i) \cup (\mathcal{V}_i \times \mathcal{V}_k), \\ \mathcal{T}_3 &= (\mathcal{V}_k \times \mathcal{V}_p) \cup (\mathcal{V}_p \times \mathcal{V}_k), \quad \mathcal{T}_4 = \mathcal{V}_p \times \mathcal{V}_p. \end{aligned} \tag{3}$$

Here, Tasks 1-4 correspond to *Known-Known*, *Known-Unseen Chiral Form*, *Known-Unseen Parent*, and *Unseen Parent-Unseen Parent*, respectively. In particular, Task 2 evaluates generalization to unseen chiral forms while keeping the underlying parent structure observed during training. For each pair (d_i, d_j), the goal is to predict its interaction event $\hat{y}_{ij} \in \mathcal{Y}$.

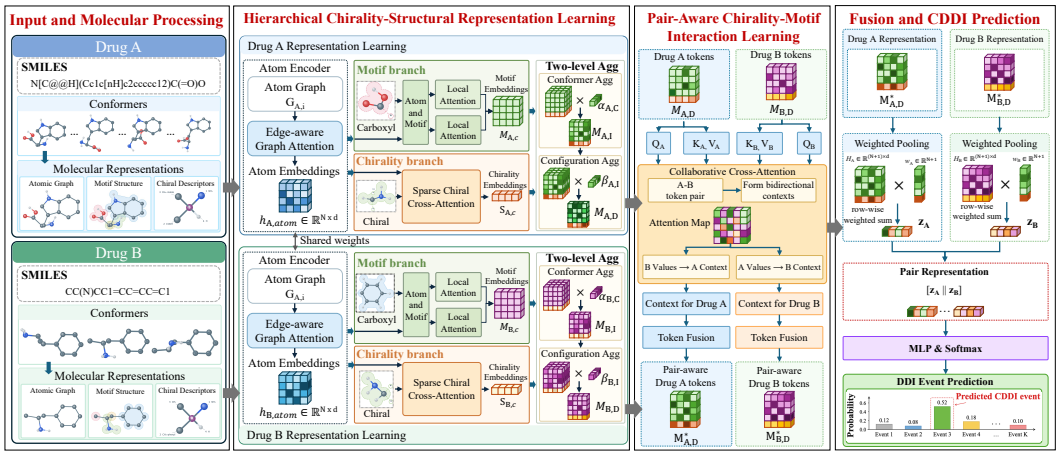


Figure 2: **Overview of the CSIL architecture**

4 Methodology

4.1 Overview

The overall framework of the proposed CSIL is illustrated in Figure 2. It jointly captures hierarchical molecular substructures and stereochemical information, and models pairwise interactions between two drugs to learn pair-aware representations for DDI event prediction. Algorithm 1 presents the pseudocode of the proposed CSIL model. The following subsections introduce the main components of the framework in detail.

4.2 Hierarchical Chirality-Structural Representation Learning

To jointly characterize molecular structure and chiral configuration, we learn drug representations hierarchically from configuration-conformer instances. For each conformer, a shared atom encoder first captures its local chemical and geometric environment, followed by two complementary

branches that extract motif-based structural features and chirality-related features. These conformer-level representations are subsequently integrated through conformer-level and configuration-level aggregation, yielding drug-level structural and chirality-related representations.

Conformer-level chirality-structural representation. For each atom i, its molecular features are combined with a role embedding that distinguishes a specified chiral center, an atom directly associated with a chiral center, and the remaining molecular environment:

$$\mathbf{h}_i^{(0)} = \mathrm{LN}\left(f_{\mathrm{atom}}(\mathbf{x}_i) + \mathbf{e}_{r_i}\right) \tag{4}$$

where $\mathbf{x}i$ denotes the atom features, r_i its atom role, and $f\mathrm{atom}(\cdot)$ a learnable projection. For each conformer, local geometry is captured by atom pairs within a spatial cutoff r_{cut}, with interatomic distances and chemical-bond information encoded as edge features. The atom representations are then processed by stacked edge-aware Graph Transformer blocks:

$$\mathbf{H}_{d,s,c}^{A,(l+1)} = f_{\mathrm{GT}}^{(l)}\left(\mathbf{H}_{d,s,c}^{A,(l)}, \mathcal{E}_{d,s,c}^{A}\right), \qquad l = 0, \ldots, L-1 \tag{5}$$

where $\mathcal{E}_{d,s,c}^{A}$ denotes the edge information of the atom-level graph, including chemical-bond attributes and local geometric information, $f_{\mathrm{GT}}^{(l)}$ denotes the l-th edge-aware Graph Transformer block. The resulting $\mathbf{H}_{d,s,c}^{A}$ is shared by the motif and chirality branches.

In the *motif branch*, the molecular graph is decomposed into chemically meaningful BRICS motifs, and the shared atom representations are aggregated to construct initial motif representations. A local-global motif attention block is then applied, where local attention captures dependencies between connected motifs, while global attention captures long-range dependencies among all motifs within the conformer. The resulting features are fused to obtain the conformer-specific motif representations $\mathbf{H}_{d,s,c}^{M}$.

In parallel, the *chirality branch* encodes the specified chiral centers. Each chiral center is initialized from its shared atom representation and enriched with explicit chirality-related descriptors, including the Cahn–Ingold–Prelog (CIP) assignment, signed chiral volume, and determinant-based geometric features, which provide explicit reflection-sensitive information, complementing the connectivity- and distance-based components. Sparse cross-attention is then applied between the chiral centers and their local atomic contexts:

$$\mathbf{Z}_{d,s,c}^{\chi} = f_{\mathrm{chiral}}\left(\mathbf{H}_{d,s,c}^{A}, \mathbf{X}_{d,s,c}^{\chi}, \mathcal{E}_{d,s,c}^{A}\right) \tag{6}$$

where $\mathbf{X}_{d,s,c}^{\chi}$ denotes the chirality-related descriptors and $\mathbf{Z}_{d,s,c}^{\chi}$ denotes the resulting chiral-center representations. These representations are further integrated through gated aggregation to obtain the conformer-level chirality representation $\mathbf{h}_{d,s,c}^{\chi}$. Accordingly, each conformer is represented by the motif token set $\mathbf{H}_{d,s,c}^{M}$ and the chirality-related representation $\mathbf{h}_{d,s,c}^{\chi}$.

Hierarchical conformer-configuration aggregation. Each configuration may admit multiple conformers, while a drug may contain one or multiple candidate configurations. Therefore, we organize the molecular representations hierarchically from conformers to configurations and finally to the drug level.

For each conformer, the motif representations are first summarized by attention pooling to obtain $\mathbf{u}_{d,s,c}$. A learnable scoring function then determines the relative contribution of each conformer within configuration s:

$$\alpha_{d,s,c} = \mathrm{softmax}_{c \in \mathcal{C}_{d,s}} f_{\mathrm{conf}}\left(\mathbf{u}_{d,s,c}\right), \qquad \mathbf{u}_{d,s} = \sum_{c \in \mathcal{C}_{d,s}} \alpha_{d,s,c} \mathbf{u}_{d,s,c}, \tag{7}$$

where $\mathcal{C}_{d,s}$ denotes conformer set of configuration s and $\mathbf{u}_{d,s}$ its configuration-level representation.

The configuration representations are further combined at the drug level. For inferred-racemate drugs with unknown configuration proportions, we assign equal weights to candidate configurations, while learning their relative contributions for the remaining stereochemical states:

$$\beta_{d,s} = \begin{cases} \dfrac{1}{|\mathcal{S}_d|}, & d \text{ is labeled as inferred-racemate,} \\ \mathrm{softmax}_{s \in \mathcal{S}_d} f_{\mathrm{cfg}}\left(\mathbf{u}_{d,s}\right), & \text{otherwise,} \end{cases} \tag{8}$$

where $\mathcal{S}_d$ denotes the candidate configuration set of drug d. The conformer-level and configuration-level weights are jointly applied to the motif and chirality-related representations:

$$\mathbf{M}_d = \sum_{s\in\mathcal{S}_d}\sum_{c\in\mathcal{C}_{d,s}} \alpha_{d,s,c}\beta_{d,s}\mathbf{H}^M_{d,s,c}, \qquad \mathbf{h}^\chi_d = \sum_{s\in\mathcal{S}_d}\sum_{c\in\mathcal{C}_{d,s}} \alpha_{d,s,c}\beta_{d,s}\mathbf{h}^\chi_{d,s,c}. \tag{9}$$

The resulting $\mathbf{M}_d$ and $\mathbf{h}^\chi_d$ provide the drug-level motif-based structural and chirality-related representations for subsequent pairwise interaction modeling.

4.3 Pair-Aware Chirality-Motif Interaction Learning

The interaction between two drugs depends not only on their individual molecular features but also on how these features respond to the paired drug. Therefore, we perform token-level interaction learning to adapt the motif-based structural and chirality-related features of each drug to its interaction partner.

For each drug d, the aggregated motif representations $\mathbf{M}_d$ and chirality-related representation $\mathbf{h}^\chi_d$ are organized into a token sequence $\mathbf{X}_d$, where the chirality-related representation is introduced as an explicit chirality token alongside the motif tokens. For a drug pair (d_a, d_b), collaborative cross-attention is performed in both directions to capture cross-drug dependencies:

$$\mathbf{A}_{ab} = \mathrm{softmax}\Big(\mathbf{Q}_a\mathbf{K}_b^\top/\sqrt{d_c}\Big),\ \mathbf{C}_a = \mathbf{A}_{ab}\mathbf{V}_b, \quad \mathbf{A}_{ba} = \mathrm{softmax}\Big(\mathbf{Q}_b\mathbf{K}_a^\top/\sqrt{d_c}\Big),\ \mathbf{C}_b = \mathbf{A}_{ba}\mathbf{V}_a \tag{10}$$

where $\mathbf{Q}_d$, $\mathbf{K}_d$, and $\mathbf{V}_d$ are projected from $\mathbf{X}_d$. The resulting contexts $\mathbf{C}_a$ and $\mathbf{C}_b$ selectively incorporate information from the paired drug, allowing both motif and chirality-related tokens to capture interaction-specific dependencies.

The partner context is then fused with the original token representation:

$$\mathbf{X}^*_d = f_{\mathrm{fus}}\left([\mathbf{X}_d \| \mathbf{C}_d \| \mathbf{X}_d \odot \mathbf{C}_d \| \,|\mathbf{X}_d - \mathbf{C}_d|]\right), \qquad d \in \{a, b\}. \tag{11}$$

This produces pair-aware token representations $\mathbf{X}^*_a$ and $\mathbf{X}^*_b$, in which the structural and chirality-related features of each drug are jointly adapted to its interaction partner.

4.4 Pairwise Fusion and CDDI Prediction

Given the pair-aware token representations $\mathbf{X}^*_a$ and $\mathbf{X}^*_b$, we first summarize the interaction-enhanced features of each drug through weighted pooling. For drug $d \in \{a, b\}$, the token weights are learned from its pair-aware representations and used to obtain a compact drug representation. The pooled representation is further fused with the retained chirality-related feature, yielding

$$\mathbf{z}_d = f_{\mathrm{drug}}\left([\mathrm{AttnPool}\left(\mathbf{X}^*_d\right) \| \mathbf{h}^\chi_d]\right), \qquad d \in \{a, b\}. \tag{12}$$

The two drug representations are then concatenated to construct the pair representation and fed into an MLP classifier:

$$\mathbf{z}_{ab} = [\mathbf{z}_a \| \mathbf{z}_b], \qquad \mathbf{o}_{ab} = f_{\mathrm{cls}}\left(\mathbf{z}_{ab}\right), \qquad P_\theta(y \mid d_a, d_b) = \mathrm{softmax}\left(\mathbf{o}_{ab}\right). \tag{13}$$

The predicted CDDI event is determined by the class with the highest probability.

To alleviate the class imbalance among DDI events, we optimize the model using a balanced-softmax cross-entropy loss. Let n_k denote the number of training samples belonging to event type k, and let π_k denote its empirical class prior. The classification logits are adjusted as follows:

$$\widetilde{o}_{ab,k} = o_{ab,k} + \tau \log \pi_k, \qquad \pi_k = \frac{n_k}{\sum_{j\in\mathcal{Y}} n_j} \tag{14}$$

where τ controls the strength of the prior adjustment. The training objective is

$$\mathcal{L} = -\frac{1}{|\mathcal{B}|}\sum_{(d_a,d_b,y)\in\mathcal{B}} \log \frac{\exp(\widetilde{o}_{ab,y})}{\sum_{k\in\mathcal{Y}}\exp(\widetilde{o}_{ab,k})} \tag{15}$$

where $\mathcal{B}$ denotes a training mini-batch.

5 EXPERIMENT

In the subsequent experiments, we evaluate CSIL on four CDDI tasks: Task 1 (*Known-Known*), Task 2 (*Known-Unseen Chiral Form*), Task 3 (*Known-Unseen Parent*), and Task 4 (*Unseen Parent-Unseen Parent*), and address the following research questions (RQs):

- RQ.1: How does the proposed model improve CDDI prediction compared with existing DDI prediction methods?
- RQ.2: What patterns are highlighted by the proposed method? (Interpretability Study)
- RQ.3: How do modality combinations and fusion schemes affect proposed method's performance? (Ablation Study)
- RQ.4: How do hyperparameters affect the proposed method's performance?

5.1 GENERAL SETTINGS

Dataset. We construct a dataset for chiral DDI prediction based on DrugBank 5.1.3. We identify drugs with explicit or potentially relevant chiral information and retain interactions in which both drugs belong to the resulting chiral drug set. Bidirectional duplicate records are consolidated irrespective of input order, and the retained drug pairs are subsequently ordered according to their semantic roles in the corresponding DDI descriptions. Each interaction is mapped to a standardized DDI event category, with empty categories removed and the remaining labels re-indexed consecutively. The resulting dataset contains 2,130 drugs, 394,475 unique DDI records, and 82 DDI event types, with each sample consisting of an ordered drug pair and its interaction category. Detailed dataset statistics, task-specific splits, and chirality-state definitions and distributions are provided in Appendix A.2.

Evaluation Metrics. For the purpose of assessing the proposed model, we employ a collection of evaluation metrics suitable for multi-class classification tasks. This set of metrics encompasses accuracy (ACC), the area under the precision-recall curve (AUPR), the area under the receiver operating characteristic curve (AUC), F1 score, precision (Pre), and recall (Rec).

Baselines. We compare the proposed model with seven representative baselines: SSI-DDI Nyamabo et al. (2021), HDN-DDI Sun & Zheng (2025), MOLGAECL Li et al. (2025), DMCL-DDI Gao et al. (2025), MSHGCL Ge et al. (2026), ImageDDI He et al. (2026), and MS-DDI Zou et al. (2026).

5.2 ANALYSIS OF MODEL PERFORMANCE (RQ.1)

Table 1 reports the performance comparison between CSIL and representative DDI methods across four CDDI tasks. Overall, CSIL outperforms the comparison methods in most cases and achieves the best Friedman mean rank Oladugba & Ikebuife (2026). Specifically, CSIL records higher mean scores in 165 of the 168 comparisons across the 168 pairwise comparisons, and achieves the lowest Friedman mean rank of 1.125. Furthermore, CSIL outperforms the best baseline in 22 out of 24 task-metric pairs, yielding an average improvement of 1.39%. These results indicate that CSIL ranks first overall, although the margins over the strongest baseline are small on several metrics.

5.3 INTERPRETABILITY STUDY (RQ.2)

Attention Overlaps with Known Warfarin Metabolism-Associated Regions. To examine whether the learned motif attention is consistent with established warfarin metabolism-associated regions, we analyze R-warfarin and S-warfarin paired with 152 common co-medications. Previous studies have reported that hydroxylation at the aromatic C7-H bond of R-warfarin and S-warfarin is primarily mediated by CYP1A2 and CYP2C9, respectively Yamazaki & Shimada (1997), while CYP3A4, CYP3A5, and CYP3A7 can catalyze hydroxylation at the aliphatic C10-H bond of both enantiomers R Jones et al. (2010). Accordingly, we define the coumarin-fused benzene-ring motif containing the C7-H site and the ketone-side-chain motif containing the C10-H site as the target regions. The target regions are first annotated by atom indices and then verified using the corresponding SMARTS patterns. For both CSIL and MeTDDI, the Top-1 attended motif/substructure is mapped to its constituent atom indices, and a hit is counted when it overlaps with either predefined target region.

Table 1: Comparison with representative DDI methods across four CDDI tasks

Setting	Metric	SSI-DDI	HDN-DDI	MOLGAECL	DMCL-DDI	MSHGCL	ImageDDI	MS-DDI	**Ours**	Gains vs. Best
Task 1	Accuracy	0.9162 ± 0.0025	0.9494 ± 0.0035	0.8960 ± 0.0008	0.8297 ± 0.0088	0.9694 ± 0.0024	0.9701 ± 0.0055	0.9674 ± 0.0017	**0.9758 ± 0.0031**	0.59%
	AUPR	0.9709 ± 0.0013	0.9857 ± 0.0015	0.9622 ± 0.0007	0.9110 ± 0.0080	0.9945 ± 0.0009	0.9944 ± 0.0007	0.9932 ± 0.0005	**0.9960 ± 0.0008**	0.15%
	AUC	0.9994 ± 0.0000	0.9996 ± 0.0000	0.9992 ± 0.0000	0.9981 ± 0.0002	0.9998 ± 0.0000	0.9998 ± 0.0000	0.9998 ± 0.0000	**0.9999 ± 0.0000**	0.01%
	F1	0.8926 ± 0.0098	0.9166 ± 0.0072	0.7281 ± 0.0105	0.7089 ± 0.0163	0.9402 ± 0.0100	0.9392 ± 0.0090	0.9409 ± 0.0058	**0.9516 ± 0.0092**	1.14%
	Precision	0.9064 ± 0.0095	0.9312 ± 0.0093	0.8074 ± 0.0113	0.8065 ± 0.0121	0.9546 ± 0.0120	0.9533 ± 0.0086	0.9523 ± 0.0117	**0.9618 ± 0.0089**	0.75%
	Recall	0.8852 ± 0.0104	0.9073 ± 0.0067	0.7078 ± 0.0061	0.6731 ± 0.0165	0.9326 ± 0.0091	0.9389 ± 0.0098	0.9341 ± 0.0025	**0.9469 ± 0.0126**	0.85%
Task 2	Accuracy	0.6067 ± 0.0255	0.7125 ± 0.0062	0.6652 ± 0.0227	0.6544 ± 0.0183	0.7033 ± 0.0140	0.7348 ± 0.0089	0.7428 ± 0.0101	**0.7456 ± 0.0214**	0.38%
	AUPR	0.6297 ± 0.0306	0.7553 ± 0.0122	0.6919 ± 0.0239	0.6910 ± 0.0250	0.7482 ± 0.0220	0.7600 ± 0.0135	0.7528 ± 0.0138	**0.7718 ± 0.0096**	1.55%
	AUC	0.9824 ± 0.0027	0.9801 ± 0.0023	0.9815 ± 0.0017	0.9840 ± 0.0028	0.9842 ± 0.0034	0.9832 ± 0.0018	0.9793 ± 0.0020	**0.9855 ± 0.0023**	0.13%
	F1	0.6039 ± 0.0295	0.6905 ± 0.0193	0.5755 ± 0.0111	0.5782 ± 0.0221	0.6768 ± 0.0291	0.6829 ± 0.0135	0.7121 ± 0.0122	**0.7199 ± 0.0242**	1.10%
	Precision	0.6392 ± 0.0209	0.7290 ± 0.0281	0.6113 ± 0.0188	0.6509 ± 0.0090	0.7193 ± 0.0228	0.7104 ± 0.0178	0.7351 ± 0.0185	**0.7582 ± 0.0348**	3.14%
	Recall	0.6353 ± 0.0359	0.7131 ± 0.0188	0.5888 ± 0.0192	0.5782 ± 0.0337	0.6955 ± 0.0435	0.7287 ± 0.0151	0.7408 ± 0.0096	**0.7495 ± 0.0245**	1.17%
Task 3	Accuracy	0.5710 ± 0.0232	0.6247 ± 0.0140	0.6273 ± 0.0186	0.5920 ± 0.0228	0.6432 ± 0.0295	0.5021 ± 0.0197	0.6450 ± 0.0222	**0.6569 ± 0.0252**	1.84%
	AUPR	0.5851 ± 0.0288	0.6580 ± 0.0178	0.6609 ± 0.0239	0.6163 ± 0.0313	0.6837 ± 0.0351	0.5139 ± 0.0193	0.6647 ± 0.0313	**0.6845 ± 0.0377**	0.12%
	AUC	0.9766 ± 0.0034	0.9712 ± 0.0033	0.9764 ± 0.0026	0.9761 ± 0.0036	0.9711 ± 0.0024	0.9742 ± 0.0063	0.9692 ± 0.0054	**0.9780 ± 0.0059**	0.14%
	F1	0.5154 ± 0.0322	0.5630 ± 0.0416	0.4838 ± 0.0355	0.4643 ± 0.0387	0.5562 ± 0.0487	0.3269 ± 0.0180	0.5596 ± 0.0355	**0.5659 ± 0.0484**	0.52%
	Precision	0.5701 ± 0.0251	0.6156 ± 0.0506	0.5722 ± 0.0303	0.5680 ± 0.0257	0.6173 ± 0.0324	0.3836 ± 0.0226	0.5975 ± 0.0348	**0.6263 ± 0.0308**	1.46%
	Recall	0.5081 ± 0.0433	0.5451 ± 0.0386	0.4632 ± 0.0365	0.4401 ± 0.0447	0.5638 ± 0.0546	0.3429 ± 0.0411	0.5739 ± 0.0428	**0.5861 ± 0.0628**	2.13%
Task 4	Accuracy	0.3714 ± 0.0295	0.3960 ± 0.0387	0.4258 ± 0.0385	0.4057 ± 0.0322	0.4003 ± 0.0339	0.1624 ± 0.0143	0.3907 ± 0.0307	**0.4290 ± 0.0359**	0.75%
	AUPR	0.3462 ± 0.0402	0.3788 ± 0.0463	0.3801 ± 0.0471	0.3762 ± 0.0442	0.3737 ± 0.0411	0.1011 ± 0.0131	0.3386 ± 0.0396	**0.3781 ± 0.0590**	−0.53%
	AUC	0.9449 ± 0.0088	0.9538 ± 0.0115	0.9381 ± 0.0105	0.9453 ± 0.0095	0.9452 ± 0.0081	0.8270 ± 0.0238	0.9100 ± 0.0113	**0.9544 ± 0.0134**	0.06%
	F1	0.2412 ± 0.0302	0.2738 ± 0.0327	0.2857 ± 0.0271	0.2461 ± 0.0340	0.2990 ± 0.0447	0.0482 ± 0.0110	0.2572 ± 0.0347	**0.3181 ± 0.0585**	6.38%
	Precision	0.2836 ± 0.0324	0.3173 ± 0.0332	0.3376 ± 0.0348	0.3020 ± 0.0329	0.3763 ± 0.0403	0.0973 ± 0.0249	0.2971 ± 0.0385	**0.3724 ± 0.0753**	−1.04%
	Recall	0.2345 ± 0.0289	0.2754 ± 0.0367	0.2821 ± 0.0232	0.2356 ± 0.0325	0.2843 ± 0.0477	0.0508 ± 0.0083	0.2613 ± 0.0379	**0.3143 ± 0.0676**	10.55%
Win/Loss		24/0	23/1	23/1	24/0	23/1	24/0	24/0	165/3	Avg Gains:
F-rank		6.125	4.167	5.375	6.250	3.417	5.458	4.083	**1.125**	**1.39%**

∗ Underlining marks the best baseline per metric. **Win/Loss**: Ours versus each baseline across 24 task-metric pairs. **F-rank**: Friedman mean rank. **Gains vs. Best**: percentage improvement over the best baseline per metric.

As shown in Figure 3a, CSIL achieves Top-1 target coverage in 145/152 R-warfarin pairs and 109/152 S-warfarin pairs. For reference, applying the same atom-overlap criterion to MeTDDI yields 3/152 Top-1 overlaps for both R-warfarin and S-warfarin. These results show that the highest-attention motifs frequently overlap with the predefined metabolism-associated regions of warfarin.

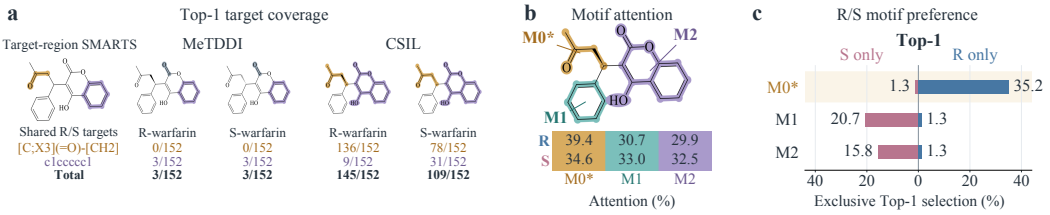


Figure 3: **CSIL prioritizes different motifs in R/S-warfarin**

Motif Attention Differs Between R-Warfarin and S-Warfarin. To examine configuration-associated structural attention, we pair R-warfarin and S-warfarin with the same 152 co-medications and partition warfarin into three chemically interpretable motifs according to its established structural moieties Park (1988); El-Haj et al. (2018). These motifs include the chiral-center-containing ketone side chain (M0), the phenyl group (M1), and the 4-hydroxycoumarin scaffold (M2). To focus on chiral configuration differences, we keep molecular connectivity and interatomic distances comparable between the R- and S-forms and normalize the attention assigned to each motif.

As shown in Figure 3b, M0 receives the highest average attention for both R- and S-warfarin (39.4% and 34.6%), compared with M1 (30.7% and 33.0%) and M2 (29.9% and 32.5%). Thus, R-warfarin shows stronger attention toward the chiral-center-containing side chain, whereas S-warfarin exhibits a more balanced distribution across the three motifs, showing that attention is distributed differently between the two configurations.

R/S-Warfarin Show Different Top-1 Motif Selection Patterns. To compare motif prioritization between R-warfarin and S-warfarin paired with the same co-medication, we examine their Top-1 motif selections across the 152 paired samples. As shown in Figure 3c, M0 is selected exclusively by R-warfarin in 35.2% of the pairs but exclusively by S-warfarin in only 1.3%. In contrast, M1 and M2 exhibit higher S-only rates (20.7% and 15.8%) than R-only rates (1.3% for both). These results show configuration-associated differences in motif prioritization under matched co-medications.

Table 2: Ablation results for Tasks 1-4

Task	Metric	Full model	w/o atom encoder	w/o motif branch	w/o chirality branch	w/o pair cross-attention	Single configuration	Single conformer
Task 1	ACC	0.9758 ± 0.0031	0.9720 ± 0.0016	0.9672 ± 0.0028	0.9759 ± 0.0011	0.9717 ± 0.0051	0.9758 ± 0.0020	0.9744 ± 0.0045
	AUPR	0.9960 ± 0.0008	0.9951 ± 0.0006	0.9945 ± 0.0007	0.9956 ± 0.0003	0.9958 ± 0.0012	0.9961 ± 0.0003	0.9958 ± 0.0012
	AUC	0.9999 ± 0.0000	0.9999 ± 0.0000	0.9998 ± 0.0000	0.9998 ± 0.0000	0.9999 ± 0.0000	0.9999 ± 0.0000	0.9999 ± 0.0000
	F1	0.9516 ± 0.0092	0.9496 ± 0.0110	0.9390 ± 0.0107	0.9500 ± 0.0066	0.9418 ± 0.0056	0.9562 ± 0.0074	0.9436 ± 0.0109
	Precision	0.9618 ± 0.0089	0.9606 ± 0.0076	0.9445 ± 0.0088	0.9616 ± 0.0057	0.9579 ± 0.0083	0.9653 ± 0.0136	0.9546 ± 0.0133
	Recall	0.9469 ± 0.0126	0.9459 ± 0.0131	0.9401 ± 0.0172	0.9451 ± 0.0092	0.9343 ± 0.0106	0.9526 ± 0.0030	0.9383 ± 0.0113
Task 2	ACC	0.7456 ± 0.0214	0.7143 ± 0.0250	0.7358 ± 0.0189	0.7331 ± 0.0079	0.7427 ± 0.0122	0.7359 ± 0.0126	0.7409 ± 0.0117
	AUPR	0.7718 ± 0.0096	0.7482 ± 0.0250	0.7719 ± 0.0271	0.7537 ± 0.0246	0.7701 ± 0.0208	0.7577 ± 0.0166	0.7982 ± 0.0170
	AUC	0.9855 ± 0.0023	0.9857 ± 0.0023	0.9870 ± 0.0020	0.9836 ± 0.0047	0.9812 ± 0.0039	0.9834 ± 0.0030	0.9848 ± 0.0040
	F1	0.7199 ± 0.0242	0.7034 ± 0.0445	0.7183 ± 0.0107	0.7115 ± 0.0103	0.7092 ± 0.0110	0.7001 ± 0.0265	0.7156 ± 0.0165
	Precision	0.7582 ± 0.0348	0.7585 ± 0.0559	0.7526 ± 0.0135	0.7505 ± 0.0265	0.7468 ± 0.0226	0.7494 ± 0.0283	0.7545 ± 0.0235
	Recall	0.7495 ± 0.0245	0.7260 ± 0.0311	0.7569 ± 0.0177	0.7481 ± 0.0210	0.7386 ± 0.0191	0.7356 ± 0.0208	0.7459 ± 0.0049
Task 3	ACC	0.6569 ± 0.0252	0.6397 ± 0.0192	0.6516 ± 0.0370	0.6489 ± 0.0176	0.6501 ± 0.0267	0.6592 ± 0.0318	0.6568 ± 0.0220
	AUPR	0.6845 ± 0.0377	0.6587 ± 0.0220	0.6763 ± 0.0457	0.6707 ± 0.0201	0.6780 ± 0.0337	0.6820 ± 0.0402	0.6838 ± 0.0360
	AUC	0.9780 ± 0.0059	0.9756 ± 0.0059	0.9743 ± 0.0045	0.9735 ± 0.0059	0.9721 ± 0.0037	0.9748 ± 0.0064	0.9733 ± 0.0088
	F1	0.5659 ± 0.0484	0.5377 ± 0.0292	0.5609 ± 0.0503	0.5576 ± 0.0292	0.5600 ± 0.0567	0.5650 ± 0.0602	0.5665 ± 0.0312
	Precision	0.6263 ± 0.0308	0.6029 ± 0.0300	0.6136 ± 0.0448	0.6277 ± 0.0322	0.6315 ± 0.0449	0.6466 ± 0.0690	0.6188 ± 0.0266
	Recall	0.5861 ± 0.0628	0.5558 ± 0.0306	0.5773 ± 0.0646	0.5706 ± 0.0309	0.5637 ± 0.0626	0.5765 ± 0.0503	0.5838 ± 0.0553
Task 4	ACC	0.4290 ± 0.0359	0.3936 ± 0.0511	0.4090 ± 0.0435	0.4158 ± 0.0441	0.4225 ± 0.0364	0.4062 ± 0.0281	0.4214 ± 0.0415
	AUPR	0.3781 ± 0.0590	0.3316 ± 0.0856	0.3111 ± 0.0650	0.3295 ± 0.0538	0.3314 ± 0.0562	0.3008 ± 0.0392	0.3321 ± 0.0407
	AUC	0.9544 ± 0.0134	0.9261 ± 0.0493	0.9242 ± 0.0283	0.9303 ± 0.0292	0.9344 ± 0.0299	0.9147 ± 0.0366	0.9173 ± 0.0420
	F1	0.3181 ± 0.0585	0.2871 ± 0.0540	0.3133 ± 0.0632	0.3010 ± 0.0603	0.3022 ± 0.0459	0.3142 ± 0.0523	0.3154 ± 0.0670
	Precision	0.3724 ± 0.0753	0.3340 ± 0.0357	0.3763 ± 0.0785	0.3617 ± 0.0518	0.3878 ± 0.0595	0.3890 ± 0.0449	0.3699 ± 0.0617
	Recall	0.3143 ± 0.0676	0.3050 ± 0.0687	0.3162 ± 0.0686	0.3167 ± 0.0550	0.2996 ± 0.0654	0.3153 ± 0.0623	0.3118 ± 0.0672
F-rank		**2.021**	5.250	4.396	4.479	4.646	3.563	3.646

5.4 Ablation Study (RQ.3)

To assess the contribution of each component, we conduct ablation studies across Tasks 1-4 by removing the atom encoder, motif branch, chirality branch, pair cross-attention, and further evaluate the use of a single configuration or conformer. As shown in Table 2, the full CSIL model achieves the lowest F-rank of 2.021, indicating the best overall performance. The chirality branch provides clearer benefits under the generalization settings than on Task 1. Although individual metrics may favor simplified variants, using the full set of components, configurations, and conformers yields more consistent performance across tasks.

5.5 Parameter Sensitivity Analysis (RQ.4)

We investigated four hyperparameters: the drug embedding dimension, number of molecular conformers, local spatial cutoff, and logit-adjustment strength. We varied each parameter while holding the others fixed across four evaluation tasks. A unified setting with five conformers is used for all main experiments to ensure consistency across tasks. Detailed results are provided in Appendix A.4.

6 Conclusion

In this work, we introduce CSIL, a chirality-structural interaction learning framework for Chiral Drug-Drug Interaction (CDDI) prediction. CSIL jointly models motif-based structural information and chirality-related representations, and captures their cross-drug dependencies through pair-aware interaction learning. We construct a CDDI dataset from DrugBank and evaluate CSIL across four settings, including generalization to unseen chiral forms and unseen parent structures. CSIL achieves the best overall Friedman mean rank among the compared methods, while ablation results show that the complete model provides more consistent performance across tasks. Interpretability analysis further shows that the highest-attention motifs frequently overlap with predefined warfarin metabolism-associated regions, and R-warfarin and S-warfarin exhibit different motif-prioritization patterns under matched co-medications. Overall, these findings underscore the importance of jointly modeling structure, chirality, and pair-specific interactions for CDDI prediction.

AI USE STATEMENT

(This section is **required** and does not count toward the page limit.)

In this work, we used generative AI tools to assist with translation. We have not used generative AI tools to generate synthetic data sets, help develop theoretical models or conceptual frameworks, formulate mathematical claims, provide critical ingredients for proving mathematical claims, assist in the writing of proofs, propose or refine hypotheses, design or provide feedback on research methodology or experiments, implement methods, clean and reformat datasets, support qualitative and thematic data analysis, or interpret results, and the rest of the required disclosure tasks are not applicable to this work. Additionally, we used generative AI tools to edit the research paper to improve readability. We have reviewed all AI-assisted work. Specifically, all AI-assisted translations and language edits were manually checked and revised by the authors to ensure that they accurately reflect the intended technical meaning and do not alter the scientific content. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

(This section is **recommended** and does not count toward the page limit.)

This work is a computational study and does not involve human subjects, clinical interventions, or personally identifiable information; therefore, IRB approval is not applicable. The drug and interaction data used in this study are derived from DrugBank and are used in accordance with the applicable licensing terms. The datasets included in the Supplementary Materials contain only processed or derived data that can be shared under the applicable terms and do not redistribute restricted DrugBank source content. The proposed method is intended for research on drug–drug interaction prediction and should not be used as a substitute for clinical judgment. Model predictions may contain errors and require appropriate experimental or clinical validation before being considered for real-world medical applications.

REPRODUCIBILITY STATEMENT

(This section is **recommended** and does not count toward the page limit.)

To support the reproducibility of our work, we have submitted the source code, datasets, implementation details, and necessary instructions as the Supplementary Materials. We believe that these resources provide sufficient information for reproducing the main experimental results presented in this paper.

REFERENCES

Keir Adams, Lagnajit Pattanaik, and Connor W. Coley. Learning 3d representations of molecular chirality with invariance to bond rotations. In *International Conference on Learning Representations*, 2022.

Mir Mansoor Ahmad, Zuraini Binti Ali Shah, and Hui Wen Nies. A systematic review of molecular structures, knowledge graphs, and cold-start scenario in drug-drug interaction prediction. *Computers in Biology and Medicine*, 190:110122, 2025. ISSN 0010-4825.

Shivani Chawla. Stereoisomerism in chemistry and drug development: Optical, geometrical, and conformational isomers. *Medicinal Chemistry*, 21:1–20, 2025. ISSN 1573-4064/1875-6638.

Fangni Chen, Tingting Jiang, Shuai Yang, Xiaohui Yuan, Chao Wang, and Lichuan Gu. Dsmv-ddi: A dual-level pharmacological semantic and stereochemical visual representation learning framework for drug-drug interaction prediction. *Journal of Molecular Graphics and Modelling*, pp. 109530, 2026.

Babiker M. El-Haj, Samrein B. M. Ahmed, Mousa A. Garawi, and Heyam S. Ali. Linking aromatic hydroxy metabolic functionalization of drug molecules to structure and pharmacologic activity. *Molecules*, 23(9), 2018. ISSN 1420-3049.

Piotr Gaiński, Michał Koziarski, Jacek Tabor, and Marek Śmieja. Chienn: embracing molecular chirality with graph neural networks. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 36–52. Springer, 2023.

Shiwei Gao, Tianli Wang, Ruiping Hu, Tianqi Chen, and Jing Liu. Dmcl-ddi: A dual-view molecular contrastive learning framework for drug-drug interaction prediction. *Expert Systems with Applications*, pp. 130815, 2025.

Daohui Ge, Xueyan Song, Guangshun Zhang, Jiaying Yan, Jiahao Li, and Jin-xing Liu. Mshgcl: Multi-scale hierarchical graph contrastive learning for drug-drug interactions. *IEEE Journal of Biomedical and Health Informatics*, 30(6):4492–4502, 2026.

Qadeer Hashir, Muhammad Asfand E Yar, Asad Ullah, Shahid Kamal, Fasee Ullah, and Zul Hilmi Abdullah. Computational approaches for drug–drug interaction prediction: a systematic review of data sources, modeling strategies, and evaluation frameworks. *Frontiers in Pharmacology*, Volume 17 - 2026, 2026. ISSN 1663-9812.

Yuqin He, Tengfei Ma, Chaoyi Li, Pengsen Ma, Hongxin Xiang, Jianmin Wang, Yiping Liu, Bosheng Song, and Xiangxiang Zeng. Imageddi: Image-enhanced molecular motif sequence representation for drug–drug interaction prediction. *Information Fusion*, 126:103574, 2026. ISSN 1566-2535.

Kexin Huang, Cao Xiao, Trong Hoang, Lucas Glass, and Jimeng Sun. Caster: Predicting drug interactions with chemical substructure representation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 702–709, 2020.

Burcu Karayavuz and Cigdem Karaaslan Kirmizioglu. Chiral (stereoselective) drugs, asymmetric synthesis, and racemic resolution methods. *Chirality*, 38(7):e70109, 2026.

Craig Knox, Mike Wilson, Christen M Klinger, Mark Franklin, Eponine Oler, Alex Wilson, Allison Pon, Jordan Cox, Na Eun Chin, Seth A Strawbridge, et al. Drugbank 6.0: the drugbank knowledgebase for 2024. *Nucleic acids research*, 52(D1):D1265–D1275, 2024.

Yu Li, Lin-Xuan Hou, Hai-Cheng Yi, Zhu-Hong You, Shi-Hong Chen, Jia Zheng, Yang Yuan, and Cheng-Gang Mi. Molgaecl: Molecular graph contrastive learning via graph auto-encoder pretraining and fine-tuning based on drug–drug interaction prediction. *Journal of Chemical Information and Modeling*, 65(6):3104–3116, 2025.

Xiaoqing Liu, Xue Yu, and Qi Dai. Graph-based deep learning for drug-drug interaction prediction: a systematic review. *Frontiers in Chemistry*, 14:1806549, 2026.

Yuqing Lu, Jing Chen, Nini Fan, Wenchao Song, Haiyang Sheng, Yinfeng Yang, and Jinghui Wang. Machine learning models for drug-drug interaction prediction from computational discovery to clinical application. *NPJ Digital Medicine*, 9(1):198, 2026.

Mei Ma and Xiujuan Lei. A dual graph neural network for drug–drug interactions prediction based on molecular structure and interactions. *PLOS Computational Biology*, 19(1):e1010812, 2023.

Arnold K Nyamabo, Hui Yu, and Jian-Yu Shi. Ssi–ddi: substructure–substructure interactions for drug–drug interaction prediction. *Briefings in Bioinformatics*, 22(6):bbab133, 2021.

Arnold K Nyamabo, Hui Yu, Zun Liu, and Jian-Yu Shi. Drug–drug interaction prediction with learnable size-adaptive molecular substructures. *Briefings in Bioinformatics*, 23(1):bbab441, 2022.

Grace A. Okunlola and Godwin A. Aleku. Stereoselective biotransformation: Transfer of learning to advance drug metabolism and biocatalysis. *Angewandte Chemie International Edition*, 65(25): e26152, 2026.

Abimibola Victoria Oladugba and Uchenna Valentine Ikebuife. A monte carlo simulation study of the type-i-error and power of anova f-test, friedman test and randomization test for randomized complete block design. *Quality & Quantity*, 60(1):2781–2803, 2026.

Markus Orsi and Jean-Louis Reymond. Assigning the stereochemistry of natural products by machine learning. *Journal of Cheminformatics*, 18(1):76, 2026.

Maxim Papusha and Kai Leonhard. Stereomolgraph: Stereochemistry-aware molecular and reaction graphs. *Journal of chemical information and modeling*, 66(7):3830–3839, 2026.

B.K. Park. Warfarin: Metabolism and mode of action. *Biochemical Pharmacology*, 37(1):19–27, 1988. ISSN 0006-2952.

Drew R Jones, So-Young Kim, Gunnar Boysen, Chul-Ho Yun, and Grover P Miller. Contribution of three cyp3a isoforms to metabolism of r-and s-warfarin. *Drug metabolism letters*, 4(4):213–219, 2010.

Xiaoqing Ru, Zhen Li, Leyi Wei, Yuanan Liu, and Quan Zou. Drug–drug interaction prediction: Paradigm shifts, key bottlenecks, and future directions. *WIREs Computational Molecular Science*, 15(6):e70056, 2025. e70056 CMS-1220.R1.

Jae Yong Ryu, Hyun Uk Kim, and Sang Yup Lee. Deep learning improves prediction of drug–drug and drug–food interactions. *Proceedings of the national academy of sciences*, 115(18):E4304–E4311, 2018.

Runhan Shi, Zhicheng Zhang, Letian Chen, Gufeng Yu, and Yang Yang. Learning molecular chirality via chiral determinant kernels. In *The Fourteenth International Conference on Learning Representations*, 2026.

Christopher Southan, Markus Sitzmann, and Sorel Muresan. Comparing the chemical structure and protein content of chembl, drugbank, human metabolome database and the therapeutic target database. *Molecular Informatics*, 32(11-12):881–897, 2013.

Jinchen Sun and Haoran Zheng. Hdn-ddi: a novel framework for predicting drug-drug interactions using hierarchical molecular graphs and enhanced dual-view representation learning. *BMC bioinformatics*, 26(1):28, 2025.

Zhaoyue Sun, Jiazheng Li, Gabriele Pergola, and Yulan He. Exddi: Explaining drug-drug interaction predictions with natural language. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 25228–25236, 2025.

Bowen Wang, Junyou Li, Donghao Zhou, Lanqing Li, Jinpeng Li, Ercheng Wang, Jianye Hao, Liang Shi, Chengqiang Lu, Jiezhong Qiu, et al. Unified and explainable molecular representation learning for imperfectly annotated data from the hypergraph view. *Nature Communications*, 16 (1):8717, 2025a.

Shanshan Wang, Chen Yang, and Lirong Chen. Lsa-ddi: Learning stereochemistry-aware drug interactions via 3d feature fusion and contrastive cross-attention. *International Journal of Molecular Sciences*, 26(14):6799, 2025b.

Di Wu, Hongyi Sun, Haichao Xu, Jia Chen, Zhong Chen, and Jie Yang. Cmf-eln: A cross-modal-fused end-to-end learning network for cold-start drug-drug interaction prediction. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pp. 5419–5430, 2026.

Hiroshi Yamazaki and Tsutomu Shimada. Human liver cytochrome p450 enzymes involved in the 7-hydroxylation of r- and s-warfarin enantiomers. *Biochemical Pharmacology*, 54(11):1195–1203, 1997. ISSN 0006-2952.

Jiaxin Yan, Haiyuan Wang, Wensheng Yang, Xiaonan Ma, Yajing Sun, and Wenping Hu. Chignn: Interpretable algorithm framework of molecular chiral knowledge-embedding and stereosensitive property prediction. *Journal of Chemical Information and Modeling*, 65(7):3239–3247, 2025.

Ziduo Yang, Weihe Zhong, Qiujie Lv, and Calvin Yu-Chian Chen. Learning size-adaptive molecular substructures for explainable drug–drug interaction prediction by substructure-aware graph neural network. *Chemical science*, 13(29):8693–8703, 2022.

Yongqi Zhang, Quanming Yao, Ling Yue, Xian Wu, Ziheng Zhang, Zhenxi Lin, and Yefeng Zheng. Emerging drug interaction prediction enabled by a flow-based graph neural network with biomedical network. *Nature Computational Science*, 3(12):1023–1033, 2023.

Yi Zhong, Gaozheng Li, Ji Yang, Houbing Zheng, Yongqiang Yu, Jiheng Zhang, Heng Luo, Biao Wang, and Zuquan Weng. Learning motif-based graphs for drug–drug interaction prediction via local–global self-attention. *Nature Machine Intelligence*, 6(9):1094–1105, 2024.

Fangyu Zhou and Shahadat Uddin. Fg-ddi: Functional group-aware graph neural networks for drug-drug interaction prediction. *Journal of Biomedical Informatics*, pp. 104981, 2026.

Xinyu Zhu, Yongliang Shen, and Weiming Lu. Molecular substructure-aware network for drug-drug interaction prediction. In *Proceedings of the 31st ACM international conference on information & knowledge management*, pp. 4757–4761, 2022.

Zhilin Zou, Peiying Li, Bin Sheng, and Liming Xin. Ms-ddi: A multi-scale representation learning approach for drug-drug interaction prediction. *Knowledge-Based Systems*, 344:116094, 2026. ISSN 0950-7051.

A APPENDIX

This appendix provides supplementary materials in the following order. First, we present the CSIL algorithm and analyze its computational complexity. Second, we provide detailed CDDI dataset statistics, task-specific splits, and chirality-state distributions. Third, we introduce the baseline models. Fourth, we report the parameter sensitivity analysis across the four CDDI tasks. Finally, we analyze the learned drug-pair representations under different model configurations.

A.1 ALGORITHM DESIGN AND TIME COMPLEXITY ANALYSIS

Algorithm 1 summarizes the procedure of CSIL. For each chiral drug, the model first determines its chirality state, constructs a bounded set of candidate configurations and conformers, and builds the corresponding atom graphs, chirality descriptors, and motif graphs (lines 1-8). The class priors used for logit adjustment are then computed from the training data (line 9).

During training (lines 10-40), the atom encoder first extracts atom-level representations, from which the chirality branch encodes chiral-center information and the motif branch learns motif-level structural representations (lines 12-19). CSIL then learns conformer weights and aggregates conformer-level representations within each configuration (lines 20-21). Configuration-level weights are subsequently determined, with uniform weighting for inferred-racemate drugs and learned weighting for the remaining chirality states (lines 22-25). The resulting representations are hierarchically aggregated into the drug-level motif representation M_d and chirality representation h_d^{χ}, which are organized into the token sequence X_d (lines 26-28).

For each drug pair, bidirectional cross-attention captures cross-drug dependencies between motif and chirality tokens (lines 29-32). The resulting partner contexts are fused with the original token representations and pooled to obtain the pair-aware drug representations z_a and z_b (lines 33-36). These representations are then used to obtain the classification logits and their logit-adjusted counterparts (lines 37-38). Finally, the classification loss is computed and the model parameters are updated (lines 39-40). The learned model outputs the CDDI event probabilities through softmax prediction (line 41).

Let N_a and E_a denote the average numbers of atoms and atom-level edges, N_m the average number of motifs, S the average number of candidate configurations per drug, C the average number of conformers per configuration, T the average number of motif and chirality tokens, d the hidden dimension, d_c the cross-attention dimension, U the number of drug-encoding operations per epoch, and Ω the number of DDI pairs. The per-epoch time complexity of CSIL is

$$O\left(USC\left(E_a d + N_a d^2 + N_m^2 d\right) + \Omega\left(T^2 d_c + T d^2\right)\right).$$

The first term corresponds to hierarchical drug representation learning over candidate configurations and conformers, while the second mainly arises from cross-drug attention and pairwise feature fusion.

A.2 DATASET STATISTICS AND CHIRALITY STATUS

Dataset overview. We construct the CDDI dataset from DrugBank records with explicit or potentially relevant chiral information. After restricting the interaction records to the selected chiral-drug set and removing duplicated undirected drug pairs, the resulting dataset contains 394,475 unique DDI records involving 2,130 drugs. Each record is represented as (d_A, d_B, y), where d_A and d_B denote the two interacting drugs and y denotes one of 82 DDI events. Duplicate interaction records are consolidated irrespective of the original drug order, while the retained samples are ordered according to the semantic roles specified in the corresponding DDI descriptions.

To evaluate different generalization scenarios, we construct four prediction tasks according to whether each test drug itself and/or its parent structure has been observed in the corresponding training split. Task 1 (*Known-Known*) considers interactions between two drugs observed during training. Task 2 (*Known-Unseen Chiral Form*) considers a known drug interacting with an unseen chiral form whose parent structure is represented in training. Task 3 (*Known-Unseen Parent*) considers a known drug interacting with a drug whose parent structure is unseen during training. Task 4 (*Unseen Parent-Unseen Parent*) considers interactions between two drugs whose parent structures

are both unseen during training. Five data splits are constructed for each task. Table 3 summarizes the record counts and label coverage across the five splits.

For every split, the training, test, and excluded subsets together contain all 394,475 DDI records. The ranges in Table 3 are independently calculated across five splits, so the minimum and maximum values in different columns do not necessarily come from the same split. In the cold-start settings, test records whose event classes are absent from the corresponding training set are excluded to ensure that each evaluated event category has supervision during training.

Chirality status. Following the definitions in Section 3.2, for each drug d, n_d^{pot} and n_d^{asg} denote the numbers of potential and assigned chiral centers, respectively. Based on assignment completeness and the relationship among drugs sharing the same parent structure, each drug is categorized as *fully specified*, *inferred-racemate*, *partial*, or *unspecified*. Table 4 summarizes their definitions and distributions among the 2,130 drugs in the CDDI dataset.

These chirality states are operational labels derived from molecular records. Following prior studies on fully, partially, and unspecified stereochemical assignments (Southan et al., 2013; Orsi & Reymond, 2026), *fully specified* denotes drugs whose potential chiral centers are fully assigned, while *partial* denotes those with only a subset assigned. Drugs with no assigned chiral centers are further categorized as *inferred-racemate* when a fully assigned form with the same parent structure exists, and as unspecified otherwise. The *inferred-racemate* label is an operational inference from molecular records rather than experimentally verified racemic composition.

A.3 Baseline Models

Due to the length of the text, we will introduce some details of the baseline model here. We will introduce seven state-of-the-art models: SSI-DDI, HDN-DDI, MOLGAECL, DMCL-DDI, MSHGCL, ImageDDI, and MS-DDI. Details of the comparison model are shown in Table 5.

A.4 Parameter Sensitivity Analysis

We varied the four hyperparameters individually across the four tasks (Figure 4). For consistency across all main experiments, we use $(d, N_c, r_{cut}, \tau) = (128, 5, 4.5\ \text{Å}, 0.5)$ for all four tasks. The sensitivity results show that these settings perform well overall, although Task 4 achieves its highest observed mean ACC with $N_c = 1$.

A.5 Drug-Pair Representation Analysis

To examine the learned drug-pair representations, we visualize the embeddings of different model configurations using t-SNE, including Atom, Atom + motif, Atom + chiral, Atom + motif + chiral, and the full CSIL model. For each task, we conduct the analysis on the first-fold test set. Silhouette scores are computed on the original high-dimensional pair representations before t-SNE projection, while t-SNE is used only for visualization. We first randomly sample 20 drug pairs to inspect the local organization of the representation space, as shown in Figures 5a 8a. The local results vary across tasks: CSIL achieves the highest Silhouette scores on Tasks 1 and 4, whereas component configurations obtain higher scores on Tasks 2 and 3. This variation indicates that the local structure observed from a small sampled subset may not fully reflect the overall representation space.

Therefore, we further examine the global drug-pair representations in Figures 5b-8b. Across all four tasks, CSIL consistently achieves the highest global Silhouette scores among the compared configurations. The resulting representations also show clearer grouping of drug pairs with the same DDI labels and better separation across different labels. These results suggest that integrating structural, chirality-related, and pair-aware interaction information enables CSIL to learn more separable drug-pair representations at the global level.

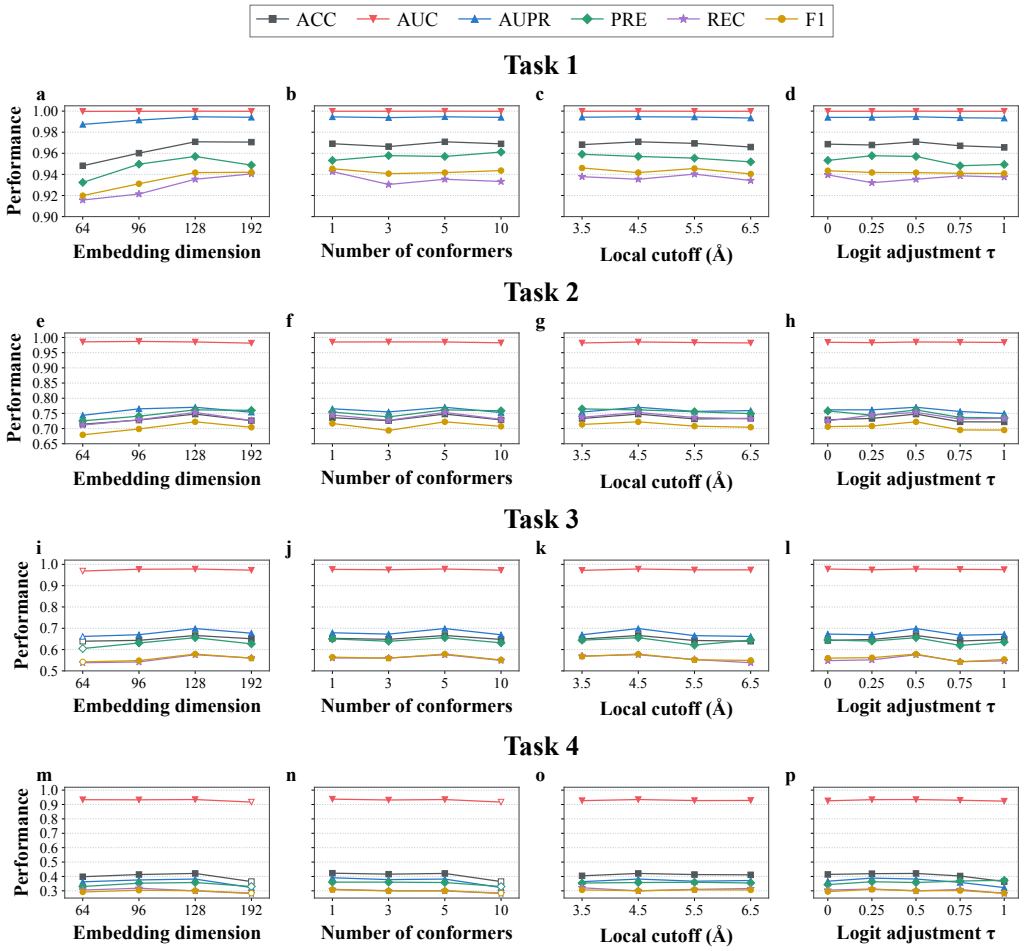


Figure 4: **Parameter Sensitivity Analysis**

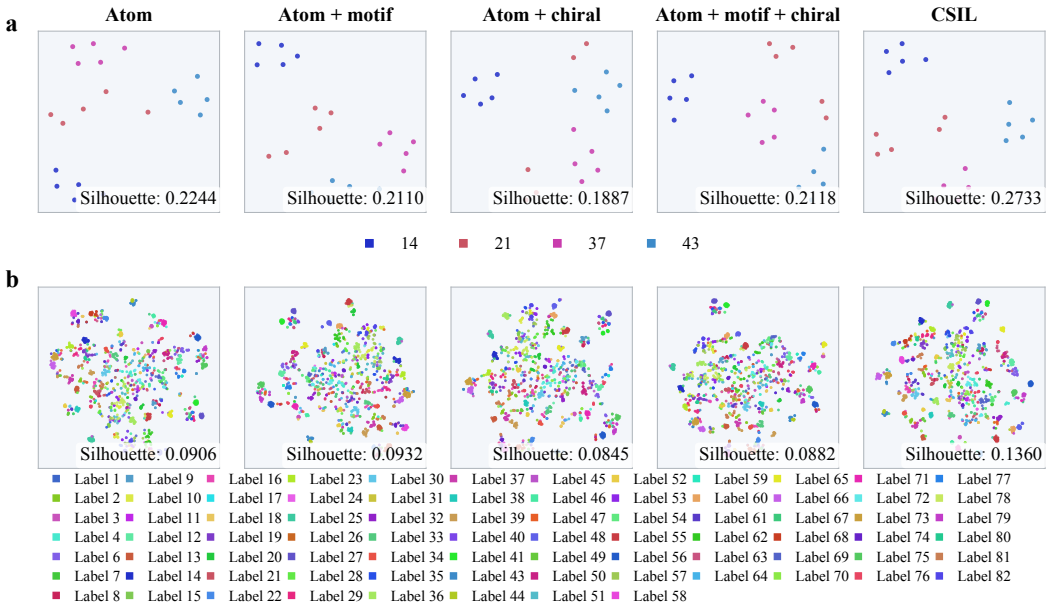


Figure 5: **t-SNE visualization of drug-pair representations on Task 1**

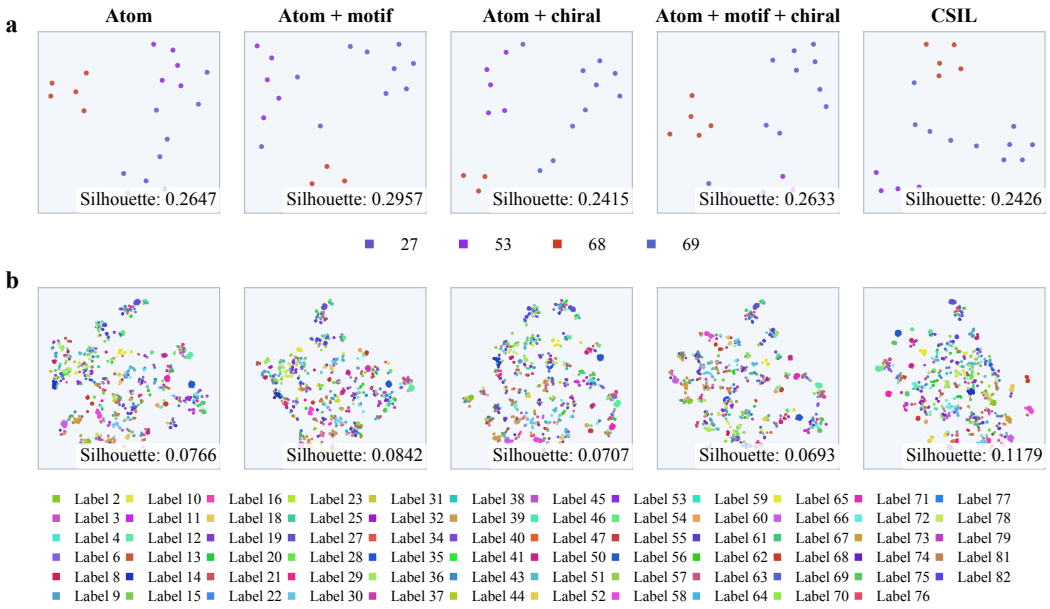


Figure 6: **t-SNE visualization of drug-pair representations on Task 2**

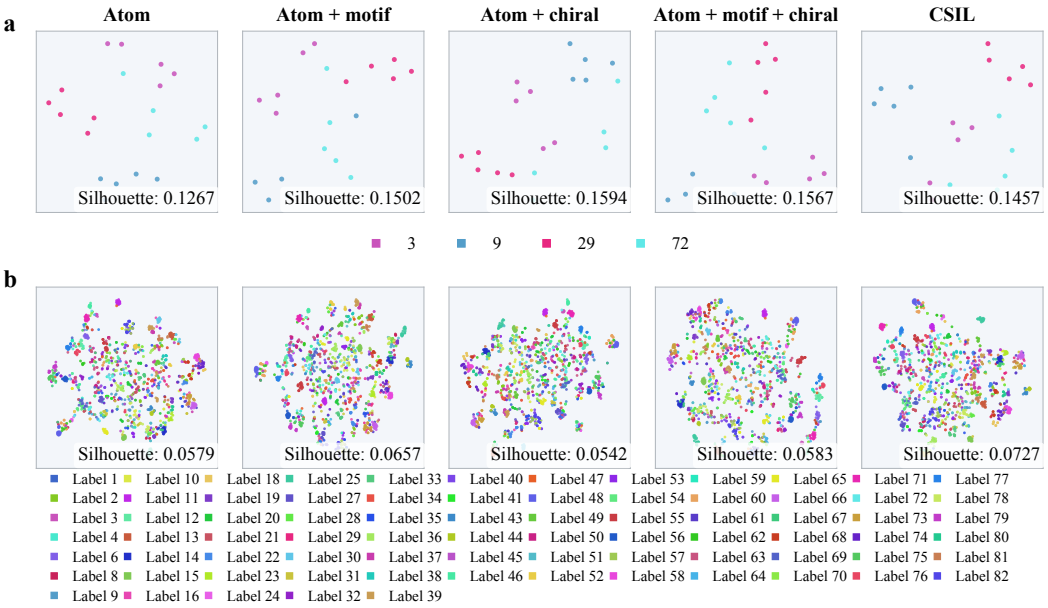


Figure 7: **t-SNE visualization of drug-pair representations on Task 3**

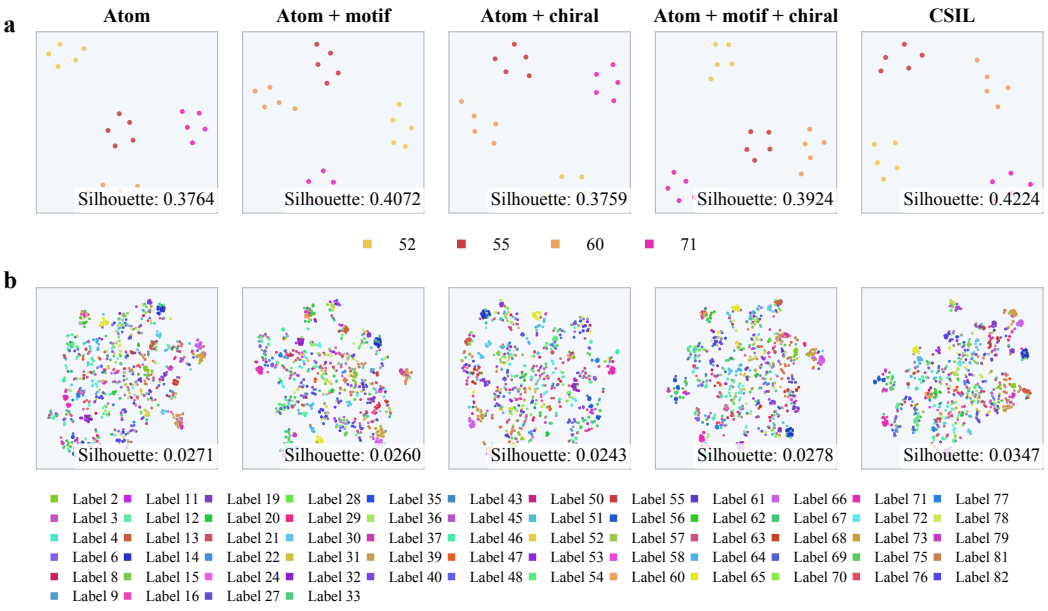


Figure 8: **t-SNE visualization of drug-pair representations on Task 4**

Algorithm 1: CSIL Algorithm for DDI Prediction between Chiral Drugs

Input: Training set $\mathcal{D}_{\mathrm{tr}} = \{(d_a, d_b, y_{ab})\}$; molecular records of drugs in $\mathcal{V}$; maximum numbers of configurations N_s and conformers N_c.

Output: Learned parameters Θ and $P_\Theta(y \mid d_a, d_b)$.

Initialize model parameters Θ;

foreach *drug* $d \in \mathcal{V}$ **do**
  determine its chirality state;
  construct candidate configuration set $\mathcal{S}_d$, with $|\mathcal{S}_d| \le N_s$;
  foreach *configuration* $s \in \mathcal{S}_d$ **do**
    generate conformer set $\mathcal{C}_{d,s}$, with $|\mathcal{C}_{d,s}| \le N_c$;
    foreach *conformer* $c \in \mathcal{C}_{d,s}$ **do**
      construct atom graph $G^A_{d,s,c}$, chirality descriptors $X^\chi_{d,s,c}$, and motif graph $G^M_{d,s,c}$;

Compute empirical class priors $\pi_k = \frac{n_k}{\sum_{j\in\mathcal{Y}} n_j}$ for $k \in \mathcal{Y}$;

while *CSIL not converged* **do**
  foreach *mini-batch* $\mathcal{B} \subset \mathcal{D}_{\mathrm{tr}}$ **do**
    foreach *unique drug* d *occurring in* $\mathcal{B}$ **do**
      foreach $s \in \mathcal{S}_d$ **do**
        foreach $c \in \mathcal{C}_{d,s}$ **do**
          $H^A_{d,s,c} \leftarrow \mathrm{AtomEncoder}(G^A_{d,s,c})$;
          $Z^\chi_{d,s,c} \leftarrow \mathrm{ChiralCrossAttn}(H^A_{d,s,c}, X^\chi_{d,s,c}, G^A_{d,s,c})$;
          $h^\chi_{d,s,c} \leftarrow \mathrm{GatedAggr}(Z^\chi_{d,s,c})$;
          $H^M_{d,s,c} \leftarrow \mathrm{MotifEncoder}(H^A_{d,s,c}, G^M_{d,s,c})$;
          $u_{d,s,c} \leftarrow \mathrm{AttnPool}(H^M_{d,s,c})$;
        $\alpha_{d,s,c} \leftarrow \mathrm{softmax}_{c\in\mathcal{C}_{d,s}} f_{\mathrm{conf}}(u_{d,s,c})$;
        $u_{d,s} \leftarrow \sum_{c\in\mathcal{C}_{d,s}} \alpha_{d,s,c} u_{d,s,c}$;
      if d *is labeled as inferred racemate* **then**
        $\beta_{d,s} \leftarrow \frac{1}{|\mathcal{S}_d|}$;
      else
        $\beta_{d,s} \leftarrow \mathrm{softmax}_{s\in\mathcal{S}_d} f_{\mathrm{cfg}}(u_{d,s})$;
      $M_d \leftarrow \sum_{s\in\mathcal{S}_d}\sum_{c\in\mathcal{C}_{d,s}} \beta_{d,s}\alpha_{d,s,c} H^M_{d,s,c}$;
      $h^\chi_d \leftarrow \sum_{s\in\mathcal{S}_d}\sum_{c\in\mathcal{C}_{d,s}} \beta_{d,s}\alpha_{d,s,c} h^\chi_{d,s,c}$;
      $X_d \leftarrow [M_d; h^\chi_d]$;
    foreach $((d_a, d_b), y_{ab}) \in \mathcal{B}$ **do**
      $A_{ab} \leftarrow \mathrm{softmax}\left(\frac{Q(X_a)K(X_b)^\top}{\sqrt{d_c}}\right)$;
      $A_{ba} \leftarrow \mathrm{softmax}\left(\frac{Q(X_b)K(X_a)^\top}{\sqrt{d_c}}\right)$;
      $C_a \leftarrow A_{ab}V(X_b), \quad C_b \leftarrow A_{ba}V(X_a)$;
      $X^*_a \leftarrow f_{\mathrm{fus}}\left([X_a \| C_a \| X_a \odot C_a \| |X_a - C_a|]\right)$;
      $X^*_b \leftarrow f_{\mathrm{fus}}\left([X_b \| C_b \| X_b \odot C_b \| |X_b - C_b|]\right)$;
      $z_a \leftarrow f_{\mathrm{drug}}\left([\mathrm{AttnPool}(X^*_a) \| h^\chi_a]\right)$;
      $z_b \leftarrow f_{\mathrm{drug}}\left([\mathrm{AttnPool}(X^*_b) \| h^\chi_b]\right)$;
      $o_{ab} \leftarrow f_{\mathrm{cls}}([z_a \| z_b])$;
      $\widetilde{o}_{ab,k} \leftarrow o_{ab,k} + \tau \log \pi_k, \quad k \in \mathcal{Y}$;
    $\mathcal{L} \leftarrow -\frac{1}{|\mathcal{B}|}\sum_{(d_a,d_b,y_{ab})\in\mathcal{B}} \log \frac{\exp(\widetilde{o}_{ab,y_{ab}})}{\sum_{k\in\mathcal{Y}}\exp(\widetilde{o}_{ab,k})}$;
    Update Θ using $\nabla_\Theta \mathcal{L}$;

$P_\Theta(y \mid d_a, d_b) \leftarrow \mathrm{softmax}(o_{ab})$;

Table 3: Statistics of the five data splits for the four CDDI prediction tasks

Task	Train	Test	Excluded	Train/Test Classes
Task 1	315,580-315,598	78,877-78,895	0	82 / 79-80
Task 2	363,197-364,374	29,519-30,655	582-623	82 / 71-72
Task 3	239,640-257,638	122,327-135,791	14,510-19,044	79-80 / 77-79
Task 4	239,640-257,638	14,397-18,993	122,440-135,842	79-80 / 59-69

Table 4: Chirality-state definitions and distributions in the CDDI dataset

Status	Definition	Drugs	Proportion
fully specified	$n_d^{\mathrm{pot}} > 0$ and $n_d^{\mathrm{asg}} = n_d^{\mathrm{pot}}$. All potential chiral centers are assigned.	1,479	69.44%
inferred-racemate	$n_d^{\mathrm{asg}} = 0$, and a fully assigned drug d' exists with the same parent structure, i.e., $\pi(d') = \pi(d)$.	36	1.69%
partial	$0 < n_d^{\mathrm{asg}} < n_d^{\mathrm{pot}}$. Only a subset of the potential chiral centers is assigned.	92	4.32%
unspecified	$n_d^{\mathrm{asg}} = 0$, and no *fully specified* drug with the same parent structure exists.	523	24.55%
Total		2,130	100.00%

Table 5: The Baseline Models.

Model	Description
SSI-DDI Nyamabo et al. (2021)	A substructure-level DDI prediction model employing graph attention networks and co-attention mechanisms to capture cross-drug substructure interactions for drug-drug interaction event prediction. *Briefings in Bioinformatics, 2021.*
HDN-DDI Sun & Zheng (2025)	A hierarchical molecular graph-based DDI model that integrates interpretable chemical substructures with enhanced dual-view representation learning to capture both molecular hierarchy and cross-drug interaction information. *BMC bioinformatics, 2025.*
MOLGAECL Li et al. (2025)	A molecular graph-based DDI model that learns drug representations through graph autoencoder pretraining and molecular graph contrastive learning, followed by full-parameter fine-tuning for DDI prediction. *Journal of Chemical Information and Modeling, 2025.*
DMCL-DDI Gao et al. (2025)	A dual-view DDI model that learns intra-molecular structural features and inter-molecular interaction representations through graph attention networks, and enhances their consistency via local and global contrastive learning with co-attention-based feature fusion. *Expert Systems with Applications, 2025.*
MSHGCL Ge et al. (2026)	A multi-scale hierarchical graph contrastive learning model that captures relationships among drug substructures through intra-layer and inter-layer contrastive learning for DDI prediction. *IEEE Journal of Biomedical and Health Informatics, 2026.*
ImageDDI He et al. (2026)	An image-enhanced DDI model that integrates motif sequence representations with global molecular image features through Transformer encoding and adaptive feature fusion. *Information Fusion, 2026.*
MS-DDI Zou et al. (2026)	A multi-scale DDI model that integrates atom-level and motif-level molecular representations through Transformer encoding and hierarchical cross-scale feature interaction. *Knowledge-Based Systems, 2026.*
CSIL	A chirality-structural interaction learning model that jointly captures motif-level structural features and chiral-center information, and employs cross-drug co-attention to model chirality-related interactions for DDI prediction.