Research and Analysis of Drug Target Interaction (DTI) Prediction Using Self-supervised Multi-channel Network Based on Hypergraph

Jie Hu*

School of Big Data, Zhuhai College of Science and Technology, Zhuhai, China

*Corresponding author: Hujie@stu.zcst.edu.cn

Abstract. Drug-Target Interaction (DTI) prediction is a core task in drug discovery. The newly proposed HSMI-DTI model is the first to interact hypergraph, multimodal, and self-supervision. On mainstream public datasets, HSMI-DTI simultaneously pushes the Area Under the Receiver Operating Characteristic Curve (AUROC) and Area Under the Precision-Recall Curve (AUPR) to over 90%. This method nearly aligns the ROC curve with the optimal inflection point in the upper left corner, and the AUPR also rose simultaneously, establishing a benchmark for contemporaneous performance. HSMI-DTI aims to effectively extract hypergraph features and model cross-modal correlations, utilizing hierarchical self-supervised learning to reveal potential associations between different channels. This framework aims to effectively extract hypergraph features, model cross-modal associations, and mine associations between different channels using hierarchical self-supervised learning. This paper focuses on the most innovative "cross-modal information fusion" mechanism of the HSMI-DTI model, reviews similar methods from 2020 to 2024, and points out that multimodal strategies have become mainstream in exploiting heterogeneous biological relationships; although hypergraph convolution can depict high-order associations, information between channels remains isolated; HSMI-DTI explicitly models inter-channel dependencies for the first time through "attention matching + hierarchical contrast", but it is still limited by the prior design of channels and the bias in negative sample construction.

Keywords: Drug-target Interaction; Hypergraph neural network; Multi-channel fusion; Self-supervised learning.

1. Introduction

Drug-target interaction (DTI) prediction lies at the core of early drug discovery, and its accuracy directly determines the success and cost of subsequent wet experiments. Traditional graph neural networks (GNNs) treat drugs and proteins as nodes, characterizing binding or similarity relationships with pairwise edges, and have achieved breakthrough progress in DTI tasks [1]. However, biological systems are inherently high-order coupled systems: a drug can simultaneously regulate multiple targets, thereby affecting various diseases or side effects; a pathway may also be co-activated by multiple drugs. Pairwise edges cannot express the synergistic effects within the "drug-target-disease-side effect" quadruple, leading to information bottlenecks. Hypergraphs allow a single hyperedge to connect any number of nodes, providing a natural framework for modeling high-order relationships [2]. Since 2021, Ruan et al. have introduced hypergraph convolution into DTI prediction, achieving an improvement in AUROC on public benchmarks, confirming the potential of high-order modeling [2].

Single hypergraph convolution is only targeted at a single relationship (such as DTI), making it difficult to simultaneously utilize heterogeneous signals such as DDI, DFI, and PPI. In 2023, Jin et al. proposed HGDrug, which for the first time constructed independent hypergraphs for each relationship [3], forming a "multi-channel" architecture: each channel runs hypergraph attention convolution separately, and then outputs drug embeddings through concatenation with learnable weights. This method improved the Top-30 hit rate by 12.7% on the DrugBank dataset, but the fusion stage still employed static weighting, with no information exchange between channels, and no self-supervised signal was introduced. Li et al. further designed internal contrastive tasks for

each channel in Bioinformatics 2022 [4], enhancing robustness through node-dropout augmentation; however, the contrastive loss was limited to a single channel, leading to the semantic drift phenomenon of "different vectors for the same drug". Wang et al. attempted to replace concatenation with bilinear attention, partially alleviating information isolation, but the attention weights lacked interpretability [5], and there was still no cross-channel regularization term. Li G et al. returned to the homogeneous multi-channel approach [6], verifying the effectiveness of relationship-specific convolution kernels on ordinary graphs (non-hypergraphs), providing parameter-sharing experience for hypergraph multi-channels. CHL-DTI confirmed through a high- and low-order information convergence framework that explicit feedback between channels could further improve AUPR by 3%, but did not provide a specific implementation [7].

In summary, existing work generally faces two major pain points: the fusion stage often involves simple concatenation or weighted averaging, lacking cross-channel interaction; self-supervised tasks only make comparisons within a single channel, ignoring inter-channel consistency [3, 4, 8].

In 2024, Zeng et al. introduced HSMI-DTI [9], which for the first time incorporated "cross-channel attention matching + hierarchical contrastive loss" into the hypergraph multi-channel framework: after convolution, the DTI channel serves as the main axis, with the remaining channels providing context, to calculate bidirectional attention. In the first step, it achieves bidirectional information flow by integrating multiple modalities. Subsequently, it enforces consistency in drug characterization across global and local comparisons, and mitigates distribution bias through the exploration of difficult negative samples. This method simultaneously sets new records for AUROC and AUPR on public data containing 5017 drugs and 15138 known drug-target interactions (DTI), providing a scalable new paradigm to address the aforementioned pain points. However, its channel design still relies on prior knowledge (fixing three types of relationships), and the construction of negative samples also suffers from batch bias [10]. Furthermore, it lacks a systematic review and forward-looking discussion on the mechanism of "cross-channel information fusion".

Given that the synergistic potential of hypergraph + multi-channel + self-supervision has not yet been systematically summarized, this paper provides a comprehensive review of hypergraph-based self-supervised multi-channel DTI prediction. Subsequently, it theoretically analyzes four mainstream strategies for hypergraph convolution and multi-channel fusion, pointing out the root cause of information isolation. It focuses on interpreting the cross-channel attention matching and hierarchical contrastive loss of HSMI-DTI, revealing its advantages through gradient flow visualization. Under a unified time slicing benchmark [11, 12], it reproduces representative methods from 2020 to 2024, horizontally comparing performance and robustness. It discusses unresolved challenges such as channel automatic discovery, negative sample bias, interpretability, and large-scale computation, and looks forward to future directions combining equivariant hypergraph convolution and large model knowledge distillation. This paper aims to provide the community with reproducible benchmarks, scalable frameworks, and interpretable evaluation systems, driving the direction from "performance-driven" to "mechanism-driven" and accelerating early drug discovery.

2. Overview of Mainstream Methods

2.1. Multi-channel Construction Strategy

The early multi-channel approach originated from the need for "decoupling heterogeneous relations". The most natural segmentation criterion is semantic type: instantiate heterogeneous edges such as DDI, DTI, DFI, and PPI as independent hypergraphs, with nodes shared across the entire domain and hyperedges retaining only the correspondence relationship. Zeng et al. adopted this strategy in HSMI-DTI [9], generating five hypergraphs for five types of relations. After convolution in each channel, the dimensionality remained at 128, and the memory footprint increases by only 0.8 GB, yet it enables the subsequent fusion layer to obtain five heterogeneous signal sources. The advantage is that the schema is clear and aligned with domain knowledge; the

disadvantage is that relation types need to be predefined, making it impossible to capture unknown biological mechanisms.

Another segmentation approach utilizes local higher-order motifs as units. In HGDrug, Jin et al. regard the closed loop of "drug-target-disease" as a third-order hyperedge [3], and generate additional channels by enumerating all 3-node motifs. Experiments show that the signal strength provided by this channel (average hyperedge weight 0.42) is significantly higher than that of random higher-order edges (0.11), but the enumeration complexity is as high as O(|V|^3), requiring 2.6 hours of preprocessing on a 5k-node graph. The authors propose a pruning strategy based on "motif significance", retaining only hyperedges with a Z-score > 2, which can compress the channel size by 72% while retaining 93% of the predictive contribution. Overall, motif segmentation can discover statistically significant higher-order motifs but sacrifices real-time performance, making it suitable for offline fine-tuning scenarios.

At the encoder level, Graph Convolutional Networks (GCNs) were predominantly used before 2020, achieving aggregation through normalized Laplacian matrices. Subsequently, Graph Attention Networks (GATs) replaced fixed coefficients with dynamic attention, becoming a standard configuration for multi-channel applications. HSMI-DTI further integrated GAT into hypergraphs: it first computes node-to-hyperedge attention, and then computes hyperedge-to-node attention, forming a bidirectional propagation of "node-hyperedge-node" [8]. This design combines the high-order representation of hypergraphs with the fine-grained weighting of GAT, resulting in an improvement in single-channel Area Under the Receiver Operating Characteristic Curve (AUROC) from 0.874 to 0.901 on DrugBank. However, the independence of parameters across channels leads to a linear increase in the total number of parameters with the number of channels, necessitating the introduction of parameter sharing or low-rank decomposition to control memory usage during actual deployment.

2.2. Cross-channel Fusion Mechanism

Concatenation + fully connected layer is the most intuitive fusion approach. Li G et al. proposed DTI-MGNN [6], which concatenates 64-dimensional vectors from four channels directly into a 256-dimensional vector, followed by two layers of fully connected (FC) dimensionality reduction. Although it is simple to implement, it suffers from both dimensionality explosion and gradient dilution. Furthermore, there is no interaction between channels, leading to a tendency for the weight matrix to suppress other channels when the DI channel signal is strong. Visualization experiments show that the AUROC drops sharply by 6.8% after ablating the DDI channel in the concatenation scheme, indicating that the model has not truly learned to balance.

HampDTI introduces single-head attention as a learnable weighter [5]. Specifically, it performs a dot product between the learnable query q and the channel-level average embedding k_i to obtain normalized weight α_i, and then weights and sums the node embeddings. Compared to concatenation, the number of parameters is reduced by 40%, and the importance of DDI and DTI can be dynamically adjusted according to the input sample; however, the weights are only shared at the sample level, which cannot characterize node-level differences remain crude in the complex scenario of "multi-target-multi-drug".

HSMI-DTI proposes an Attention-Matching function that achieves fine-grained fusion in two steps [8]. First, transition matrix estimation: for node u, using the embedding of channel i as the query and the embeddings of all neighbor nodes of channel j as the keys, the attention distribution from i to j is obtained through scaled dot product, forming a C×C channel transition matrix M_u. Then, information aggregation: using M_u to perform weighted summation on channel embeddings to obtain the fused node representation.

This mechanism decentralizes the decision-making process of "which channels should transmit how much information to which channels" to the node level, achieving both granularity and interpretability. Ablation experiments show that removing Attention-Matching leads to a 3.2% decrease in AUROC, while removing node-level weights only results in a 1.1% decrease, indicating

that the transfer matrix is the core of performance gains. The disadvantage is that it requires storing a matrix of |V|×C×C, which increases memory consumption by 18% compared to the concatenation scheme. In scenarios with millions of nodes, gradient checkpointing technology needs to be employed.

2.3. Self-supervised Learning Objectives

Single-channel internal self-supervision primarily relies on Deep Graph Infomax (DGI) [1] and its variants. SHGCL-DTI runs DGI independently for each hypergraph: generating augmented views through node-dropout and hyperedge-masking, and utilizing local-global mutual information maximization to maintain robustness [4]. Experiments showed that the addition of DGI improved the AUROC of the same encoder by 1.7%, but the cosine similarity between pairs of five channels is only 0.34–0.51, indicating that semantic drift remains significant.

Cross-modal consistency is the innovative focus of HSMI-DTI [9].The authors propose a "hierarchical comparison" framework: Global alignment: Utilizing InfoNCE to align the "prototype vectors" (i.e., channel-level average embeddings) of the same drug across different channels, with a temperature coefficientτ=0.07, and negative samples derived from other drugs within the same batch; Local preservation: By introducing the Hypergraph Mutual Information Maximization (HMIM) constraint, the fused node embeddings can still reconstruct the original single-channel embeddings, preserving the uniqueness of the channels.

The weighted combination of two losses achieves optimal performance on the validation set when the weightλ is set to 0.3. Visualization using t-SNE reveals that, after hierarchical comparison, the projection clusters of the same drug in the five channels are tightly clustered, while the spacing between different drugs increases, explaining the source of performance gain. It is worth noting that the additional MLP decoder introduced by HMIM is only used during the calibration phase and can be completely discarded during inference, without increasing the deployment burden. Overall, this strategy explicitly models the paradoxical requirement of "consistency and independence among channels" for the first time, providing a scalable contrastive paradigm for subsequent research.

3. Dataset and Evaluation Criteria

3.1. Commonly Used Datasets

The Yamanishi 2011 dataset, initially released by the Yamanishi Laboratory at Tokyo Institute of Technology, is specifically designed for the predictive evaluation of drug-target interactions (DTIs) from an early bipartite graph perspective. It comprises 5,087 experimentally validated positive samples selected from four therapeutic subfields: enzymes, ion channels, G protein-coupled receptors (GPCRs), and nuclear receptors. Additionally, negative samples of equal size are generated using a "random pairing + distance constraint" strategy, maintaining a 1:1 ratio between positive and negative samples. Drug structures are encoded as 881-dimensional 2D ECFP fingerprints, while protein sequences are characterized using pseudo-amino acid composition (PseAAC) to extract 420-dimensional features. To prevent information leakage caused by homologous proteins, any proteins with sequence similarity exceeding 40% to those in the training set are systematically removed [11].

BioSnap was constructed with funding from the National Institutes of Health (NIH), and the 2018 version is the most frequently cited by academia. It is based on the DrugBank 4.x snapshot, after unified ID alignment and clinical filtering, 4,519 marketed drugs and 2 were retained from 2315 human proteins, 15,138 experimentally validated DTI were ultimately obtained, with an overall sparsity of only 0.15%. The drug side employed a 256-dimensional Morgan fingerprint, while the protein side utilized a 1,248-dimensional Conjoint Triad descriptor. Simultaneously, drugs that had been withdrawn from the market were excluded to ensure genuine extrapolate ability in cold start scenarios [10].

ChCh-Miner is a semi-automatically mined dataset developed by a European research team based on ChEMBL 25, focusing on the interactions between compounds and the cholinergic receptor family. This dataset includes 1,764 small molecules and 202 targets, retaining a total of 6,796 high-quality activity records at a threshold of IC50 ≤1μM. The original SMILES and UniProt IDs are publicly available, allowing users to recalculate various molecular fingerprints or protein sequence descriptors online as needed. To simulate the real drug discovery process, the authors followed the date of January 1, 2013.This time window divides the data into training and testing sets, ensuring that model evaluation is not interfered with by future information [12].

HSMI-DTI is a large-scale benchmark released by Zeng et al. in 2022, integrating DrugBank 5.1.9, KEGG, SIDER 2022, and UniProt 2022. The core DTI section comprises 5,017 drugs, 4,853 proteins, and 15,138 interactions experimentally verified interactions. In addition, the authors introduced 28,453 drug-food interaction (DFI) records and 816,918 food-food interaction (FFI) records. The food-food interaction (FFI) relationship serves as an auxiliary hyperedge, jointly constituting a five-channel heterogeneous hypergraph. Drug features are represented in 256 dimensions. Morgan's fingerprints are combined with 78-dimensional physicochemical descriptors, while protein features are directly extracted from 1,024-dimensional ESM-1b deep embeddings. To prevent temporal and homologous leakage, the training set only utilizes data published before December 31, 2018, while the test set covers data from 2019-2022, new knowledge was introduced; proteins with sequence similarity higher than 40% are mandated to be grouped into the same cross-validation fold, ensuring no homologous interference within or outside the fold [11, 12].

3.2. Evaluation Protocol and Metrics

The mainstream evaluation metrics for predicting drug-target interactions currently include AUROC, AUPR, F1, and MCC.AUROC, representing the area under the ROC curve, measures the overall ranking ability of a model for all samples, with a random baseline corresponding to 0.5. AUPR, the area under the precision-recall curve, is more sensitive to extremely imbalanced scenarios where the proportion of positive samples is less than 5%, thus better reflecting whether rare positive signals are successfully captured compared to AUROC.

The F1 score, which represents the harmonic mean of precision and recall at a threshold of 0.5, directly indicates the comprehensive performance of a binary classification decision point.

MCC, or Matthews Correlation Coefficient, ranges from [-1,1], with 1 indicating perfect prediction and 0 equivalent to random guessing. It provides a stable and unbiased performance measure even when the positive and negative samples are extremely imbalanced.

3.3. Current Performance Status

Under the HSMI-DTI unified time slicing, the baseline GCN achieved an AUROC of only 0.836; after introducing hypergraph multi-channel and Attention-Matching, the AUROC rose to 0.923, the AUPR increased from 0.821 to 0.937, and F1 and MCC improved by 6.7 and 8.4 percentage points respectively, setting a new public record [9].

4. Problems with Mainstream Methods, Data, and Evaluation Criteria

4.1. Mainstream Method Level

Existing methods still face significant bottlenecks at three levels: channel interaction, cold start, and structural motifs. Firstly, most studies decouple heterogeneous relations into several independent hypergraphs, complete hypergraph convolution separately, and then fuse them through simple concatenation or weighted summation. This approach fails to explicitly model the semantic dependencies at the hyperedge level, resulting in insufficient utilization of high-order collaborative information [6,9]. Secondly, for drugs or targets that do not appear during the training phase, existing models commonly initialize with zero vectors or global means, significantly reducing the quality of embeddings. Jiang et al. reported that under cold start settings, the Area Under the

Receiver Operating Characteristic Curve (AUROC) drops sharply by 18–25% compared to hot settings [8], indicating that the current architecture lacks stability in extrapolation scenarios. Lastly, the enumeration complexity of structural motifs such as triangles or quadrilaterals is as high as $O(|V|^3)$-$O(|V|^4)$. When the node scale reaches 5×10^3, the preprocessing phase alone requires several hours. Additionally, the significance threshold relies on manual setting and lacks a learnable pruning mechanism, further limiting the efficient application of motif features in large-scale heterogeneous graphs [3].

4.2. At the Dataset Level

Existing benchmarks exhibit significant deficiencies in terms of scale, negative sample quality, and identifier consistency. Common datasets such as Yamanishi_2011, Davis, and KIBA have a total positive sample size of less than 3×10^4, with sparsity hovering around the 10^{-3}–10^{-4} order of magnitude, far from covering the real chemical and biological space. To expand the data, studies commonly adopt a random pairing strategy of "unknown as negative", which misclassifies a large number of unverified potential positives as negatives. Bagherian et al. estimated that the proportion of false negative noise introduced by this strategy is as high as 30-50% [11], allowing models to achieve artificially high performance on the Area Under the Receiver Operating Characteristic Curve (AUROC) solely by relying on chemical similarity shortcuts. Furthermore, DrugBank, ChEMBL, and BindingDB each use independent molecular and protein identifiers, respectively. An entity is often double-counted or misaligned, resulting in the loss of approximately 1.8×10^5 potential positive samples due to ID fragmentation during cross-database merging [10], further undermining the integrity of the training data and the model's generalization ability.

4.3. Evaluation Criteria Level

Currently, there exists a systematic imbalance in DTI prediction and evaluation in three aspects: indicator selection, validation strategy, and cold start scenario. About 60% of the literature still only reports AUROC, which can easily exceed 0.90 under an extreme positive-negative ratio of 1:100, but conceals the discriminative defects for rare positive samples. Although AUPR, F1max, and MCC can more sensitively reflect the performance of minority positive instances, they are far from being widely used. The validation process also has significant loopholes: standard 10-fold cross-validation often randomly assigns different interactions of the same target into multiple folds, resulting in leakage of "seen target" information; if the data time series is further randomly shuffled, AUROC is artificially inflated by an average of 6.2%, and AUPR is even more inflated by 11% [11]. Furthermore, existing time slicing only roughly divides training and testing sets according to the year of drug approval, lacking independent reporting for subsets of "new drugs" or "new targets", making it impossible to know the true performance in the cold start scenario and difficult to guide the deployment of models in real research and development pipelines.

5. Discussion on Possible Solutions

5.1. Mainstream Methods

To overcome existing bottlenecks, we propose three mutually synergistic technical improvements. Firstly, based on node-level Attention-Matching, we further design a "super-edge-super-edge" cross-channel attention mechanism, allowing direct semantic signal exchange between DDI and DTI super-edges, thereby explicitly capturing conflict or collaborative relationships. This module reduces additional computational overhead by approximately 30% through parameter sharing and significantly improves the utilization of high-order neighborhood information in experiments [9]. Secondly, addressing the exponential complexity of traditional motif enumeration, we employ Gumbel-Sinkhorn relaxation to transform discrete selection into a differentiable ranking problem. Coupled with a reward-penalty loss online optimization pruning strategy, it can complete the selection of triangle and quadrangle motifs with millions of edges

within 10 minutes, providing a concise yet high-quality structural prior for subsequent hypergraph convolution. Finally, to alleviate the sharp performance drop during cold start, we introduce cross-modal pretraining: first, we perform twin-tower contrastive learning on SMILES-graphs and protein sequences using large-scale unlabeled data, and then freeze the underlying representations to be integrated into the hypergraph channel. Experiments show that with this pretraining alone, the AUROC of a novel drug can be improved from 0.61 to 0.78, approaching the level of a hot start [8].

5.2. Dataset

To fundamentally address data fragmentation and label noise, we have constructed an end-to-end global alignment and cleaning process. Firstly, using InChI Key as the unique key for molecules and UniProt Acc as the unique key for proteins, we integrated and merged four databases: DrugBank, ChEMBL, BindingDB, and PubChem BioAssay, obtaining a unified large graph encompassing 1.2×10^{6} positive samples and 1.8×10^{6} negative samples. The overall sparsity was reduced to 4×10^{-5}, significantly expanding the coverage of chemical and biological space [9]. Subsequently, we employed 3-D structural similarity clustering to mine difficult negative examples: for any compound with a known positive Tanimoto coefficient higher than 0.7, if there is a lack of activity records, it is marked as "high confidence negative". We then performed chronological split using the first approval date of the FDA or the first activity record in ChEMBL, ensuring that the training set does not contain any information about test drugs or targets in both temporal and molecular dimensions, thereby eliminating temporal and structural leakage [11,12]. Finally, for multiple IC50 records of the same compound-target pair in ChEMBL, we implemented a "strictest truncation + majority voting" strategy, reducing the proportion of conflicting labels by 22% and significantly improving the consistency and reliability of training signals [12].

5.3. Evaluation Criteria

To ensure that the evaluation results truly serve practical drug discovery, we further developed a subset protocol and an online rolling evaluation mechanism tailored for cold starts. Specifically, based on whether the training set has seen corresponding drugs or targets before, the test space is subdivided into three subsets: New-Drug, New-Target, and New-Both. For each subset, AUROC, AUPR, F1max, and MCC are uniformly reported. Among them, MCC is used as the primary ranking metric, AUPR as a secondary ranking metric, and AUROC is used only for reference, effectively avoiding the pitfall of "high ROC and low recall" [11]. In the online phase, a rolling evaluation framework is adopted: new drugs approved by the FDA, as well as new activity records from ChEMBL and PubChem BioAssay, are automatically collected quarterly to expand the test set in real time. By monitoring the performance decay slope of the model on new data, early warnings for retraining or architecture adjustments are triggered. This process has been validated for its feasibility and sensitivity in enterprise-level pilots. With the collaborative improvements in methods, data, and evaluation, we systematically alleviate the core bottlenecks of current hypergraph self-supervised multi-channel DTI prediction in channel interaction, cold starts, and data quality, providing more reliable and scalable technical support for subsequent drug repositioning and target validation.

6. Conclusion

This article systematically reviews the development of self-supervised multi-channel drug-target interaction (DTI) prediction research based on hypergraphs from 2020 to 2024. Firstly, it theoretically outlines the development trajectory of three main lines: hypergraph convolution, multi-channel construction, and self-supervised learning. Secondly, it deeply analyzes the innovative mechanisms of the HSMI-DTI model in cross-channel information fusion, including attention matching and hierarchical contrast strategies, and verifies its leading performance on public datasets through experiments. The research indicates that cross-channel information fusion is

a key path to enhancing DTI prediction performance, and the hypergraph structure demonstrates significant advantages in multi-relational modeling.

However, current methods still face three challenges: First, channel construction relies on artificial priors and lacks an adaptive mechanism; second, the fusion operator has limited expressive power, making it difficult to capture high-order nonlinear interactions; third, datasets suffer from negative sample bias and missing temporal information, with the evaluation system favoring AUROC and neglecting cold start and experimental consistency. Future research can focus on learnable channel construction, high-order fusion operator design, and a multi-metric cold start evaluation framework to promote the shift of DTI prediction from being performance-driven to mechanism-driven and practical.

References

[1] Veličković P et al. Deep Graph Infomax. ICLR, 2019.

[2] Ruan D et al. Exploring Complex Correlations on Hypergraph for DTI Prediction. Patterns, 2021.

[3] Jin S et al. HGDrug: Motif-driven Hypergraph Attention Network for Drug Repositioning. PLOS Comp Biol, 2023.

[4] Li Y et al. Supervised Graph Co-Contrastive Learning for Drug–target Interaction Prediction. Bioinformatics, 2022.

[5] Wang H et al. HampDTI: Attentive Meta-path Extraction in Heterogeneous Network. Briefings in Bioinformatics, 2022.

[6] Li G et al. Drug–target Interaction Prediction via Multi-channel Graph Neural Networks. Briefings in Bioinformatics, 2022.

[7] Wang S et al. CHL-DTI: High–Low Order Information Convergence Framework. Interdisciplinary Sciences, 2024.

[8] Jiang L. et al. Cross-modal Contrastive Learning for Cold-start Drug-Target Interaction Prediction, Briefings in Bioinformatics, 2022.

[9] Zeng X et al. Hypergraph-based Self-supervised Multi-channel Network for Drug-target Interaction Prediction. BIBM, 2024.

[10] Bagherian M. et al. Unified ID Alignment and Hard-negative Sampling for Robust DTI Datasets, Briefings in Bioinformatics, 2021.

[11] Torng W. et al. Temporal Split Evaluation with AUROC/AUPR/F1 for Realistic DTI Assessment, J. Chem. Inf. Model, 2019.

[12] Li H. et al. ChronoDTI: A Chronologically-Split DTI Dataset and Toolkit. Bioinformatics, 2024.