

# ZeroDDI: A Zero-Shot Drug-Drug Interaction Event Prediction Method with Semantic Enhanced Learning and Dual-Modal Uniform Alignment

Ziyan Wang<sup>†</sup>, Zhankun Xiong<sup>†</sup>, Feng Huang, Xuan Liu and Wen Zhang<sup>\*</sup>

College of Informatics, Huazhong Agricultural University, Wuhan 430070, China

{wangziyan, xiongzk, fhuang233, lx666}@webmail.hzau.edu.cn, zhangwen@mail.hzau.edu.cn

## Abstract

Drug-drug interactions (DDIs) can result in various pharmacological changes, which can be categorized into different classes known as DDI events (DDIEs). In recent years, previously unobserved/unseen DDIEs have been emerging, posing a new classification task when unseen classes have no labelled instances in the training stage, which is formulated as a zero-shot DDIE prediction (ZS-DDIE) task. However, existing computational methods are not directly applicable to ZS-DDIE, which has two primary challenges: obtaining suitable DDIE representations and handling the class imbalance issue. To overcome these challenges, we propose a novel method named ZeroDDI for the ZS-DDIE task. Specifically, we design a biological semantic enhanced DDIE representation learning module, which emphasizes the key biological semantics and distills discriminative molecular substructure-related semantics for DDIE representation learning. Furthermore, we propose a dual-modal uniform alignment strategy to distribute drug pair representations and DDIE semantic representations uniformly in a unit sphere and align the matched ones, which can mitigate the issue of class imbalance. Extensive experiments showed that ZeroDDI surpasses the baselines and indicate that it is a promising tool for detecting unseen DDIEs. Our code has been released in <https://github.com/wzy-Sarah/ZeroDDI>.

## 1 Introduction

Taking multiple drugs concurrently can cause drug-drug interactions (DDIs), which may lead to various pharmacological changes [Vilar *et al.*, 2014]. According to the different specific consequences, the huge number of DDIs can be categorized into hundreds of classes [Ryu *et al.*, 2018], named DDI events (DDIEs), as shown in Figure 1(a). In recent years, researchers have paid much attention to DDIE prediction, i.e., classifying a DDI into a specific DDIE, which is helpful for researchers and clinicians to investigate the mechanism behind the consequences of polypharmacy. Note that to avoid

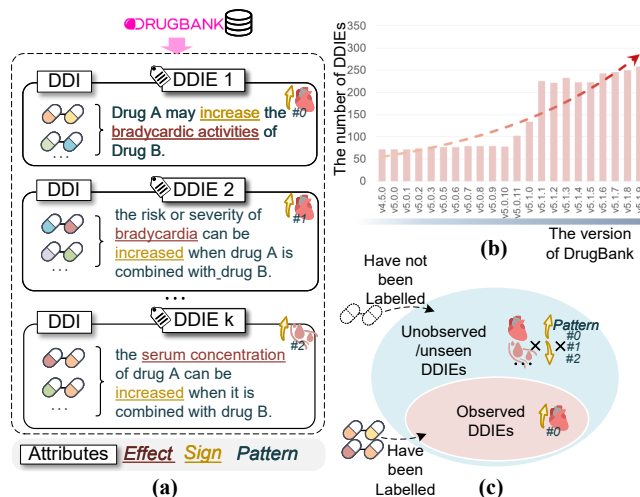


Figure 1: (a) DrugBank is a standard database containing the textual descriptions of DDIEs. Almost every description can be split into *Effect*, *Sign*, and *Pattern* attributes. (b) With the updates of the Drugbank version, the number of DDIEs is also increasing. (c) Unseen DDIEs may come from the composition of existing attributes.

confusion, we will use the term "drug pairs" instead of DDIs in this paper.

With further clinical observation and the development of therapeutic drugs, the number of DDIEs is increasing, as shown in Figure 1 (b). Moreover, inspired by [Deng *et al.*, 2020; Feng *et al.*, 2023], we notice that the DDIE contains three key concepts: the effect (such as "bradycardia"), the "increased" or "decreased" pharmacological changes, and the directional interaction between two drugs (which is reflected in the sentence pattern). To simplify, we term them as three kinds of attributes, i.e., *Effect*, *Sign*, and *Pattern*, respectively. We observed that annotated DDIEs (seen DDIEs) are less than all compositions of three kinds of attributes, as shown in Figure 1 (c), thus we can infer the existence of unseen DDIEs without labelled drug pairs. Predicting unseen DDIEs has a high practical value because novel drug pairs may come from unseen DDIEs. We formulate classifying drug pairs of unseen DDIEs as a zero-shot DDIE prediction (ZS-DDIE) task.

Although many DDIE prediction methods have been proposed, they all neglect the ZS-DDIE task. When no trainable

instances are used to transfer knowledge from seen to unseen classes, the semantic information (or auxiliary information) is needed to be attached to every class for constructing semantic relatedness between seen and unseen classes. However, most DDIE prediction methods view DDIEs as either labels [Zhu *et al.*, 2023], random initialized vectors [Lin *et al.*, 2022] or one-hot vectors [Zhu *et al.*, 2022a], which do not have realistic semantic significance. Therefore, they cannot be applied to the ZS-DDIE task. An intuitive solution for ZS-DDIE is to learn the experience from zero-shot learning (ZSL) classification methods, which are developed by computer vision. There is a commonly used compatibility framework in ZSL [Xian *et al.*, 2018], which transfers knowledge from seen to unseen classes by projecting all classes into a common semantic space, and realizes the unseen class prediction by aligning the representations of instances and matched classes. Nevertheless, applying this framework to the ZS-DDIE task needs to resolve two major challenges. **The first challenge** is to obtain suitable DDIE representations, which are required to construct reasonable relatedness between different classes. For example, based on DDIE textual descriptions, the DDIE representations with the same *Pattern* may be closer than that with the same *Effect* in semantic space (as shown in Appendix A), which is inconsistent with the biological significance and will hinder the prediction performance. **The second challenge** is the class imbalance. Researchers have revealed that the top three most frequent DDIEs contain more than half of all instances [Deng *et al.*, 2021], while a great number of DDIEs have less than 10 instances. The class imbalance problem may lead to unclear decision boundaries of classification, causing a decline of discriminative ability for the prediction of classes with few or even no instances [Ji *et al.*, 2022].

In this work, we propose a novel method, called ZeroDDI, for zero-shot DDIE prediction and resolve the challenges above. **For the first challenge**, we design a biological semantic enhanced DDIE representation learning module (BRL) to obtain suitable DDIE representation. Since *Effect* attributes are the key biological semantics, BRL first extracts attribute-level semantics of *Effect* and together with class-level semantics of DDIE textual descriptions. Then, considering that molecular substructure plays an important role in drug properties [Jin *et al.*, 2023], BRL establishes the fine-grained interaction between substructures and semantic tokens of text to guide the fusion of bi-level semantics, which can distill discriminative semantics for learning DDIE representations. **For the second challenge**, we design a dual-modal uniform alignment strategy (DUA) to enable the representations of both drug pairs (structure modal) and DDIEs (text modal) uniformly distributed on the unit sphere and then align the matched ones, which can mitigate the problem of unclear decision boundaries of classification that caused by class imbalance [Li *et al.*, 2022] and further improve the ZS-DDIE performance.

In summary, the main contributions of this paper are described as follows:

- We investigate a novel problem for DDIE prediction: zero-shot DDIE prediction (ZS-DDIE), and propose a novel method ZeroDDI for the ZS-DDIE task, along

with a ZS-DDIE dataset.

- We design a novel biological semantic enhanced DDIE representation learning module (BRL) with the extraction of class-level and attribute-level semantics, and substructure-guided bi-level semantic fusion to learn suitable DDIE representations with reasonable relatedness between different classes.
- We introduce a dual-modal uniform alignment strategy (DUA) to enable structure and text modal representations (i.e., drug pair and DDIE representations) to be uniformly distributed in unit sphere, aiming to handle the class imbalance issue and thus improve the discrimination of the model.
- The extensive experiments show that ZeroDDI can achieve superior performances in the ZS-DDIE task than baselines and reveal that ZeroDDI is a promising tool for detecting unseen DDIEs.

## 2 Related Work

### 2.1 Drug-Drug Interaction Event Prediction

Existing DDIE prediction methods can be classified into three categories: DNN-based methods [Deng *et al.*, 2021; Lin *et al.*, 2021], Tensor factorization-based methods [Jin *et al.*, 2017; Yu *et al.*, 2022], and GNN-based methods [Zhu *et al.*, 2022a; Lv *et al.*, 2023; Xiong *et al.*, 2023]. However, none of the DDIE prediction methods are specifically tailored to tackle the ZS-DDIE task. Most of them cannot be applied to the ZS-DDIE task in that they neglect to encode DDIE or lack semantic significance in their DDIE embeddings [Zhu *et al.*, 2023; Zhu *et al.*, 2022a; Lin *et al.*, 2022]. We find that 3DGT-DDI [Lv *et al.*, 2023], which utilized a pre-trained language model (SCIBERT) to learn DDIE representations from textual descriptions, shows potential adaptability for handling the ZS-DDIE task. Despite this, zero-shot DDIE prediction remains uncharted territory and needs to be explored. More discussions of related work are in Appendix B.

### 2.2 Zero-Shot Learning

Zero-Shot learning (ZSL) methods, which were introduced in computer vision, achieve transferability by projecting all classes into a common semantic space, and facilitate the classification of instances from unseen classes through a compatibility function [Nawaz *et al.*, 2022]. For learning semantic representation, current ZSL methods rely on three primary information sources: class names, class attributes, and class textual descriptions. Word-based methods [Frome *et al.*, 2013; Li *et al.*, 2015] represent the classes by simple words of the class name. Attribute-based methods represent the classes by more accurate and specific shared characteristics among classes [Lampert *et al.*, 2013; Al-Halah and Stiefelhagen, 2017]. Textual description-based methods represent the classes by textual descriptions of classes [Reed *et al.*, 2016; Zhu *et al.*, 2018], which can provide more context information about classes. For constructing compatibility functions, ZSL methods utilized/developed different metric functions to pull an instance close to its matched class and push it away

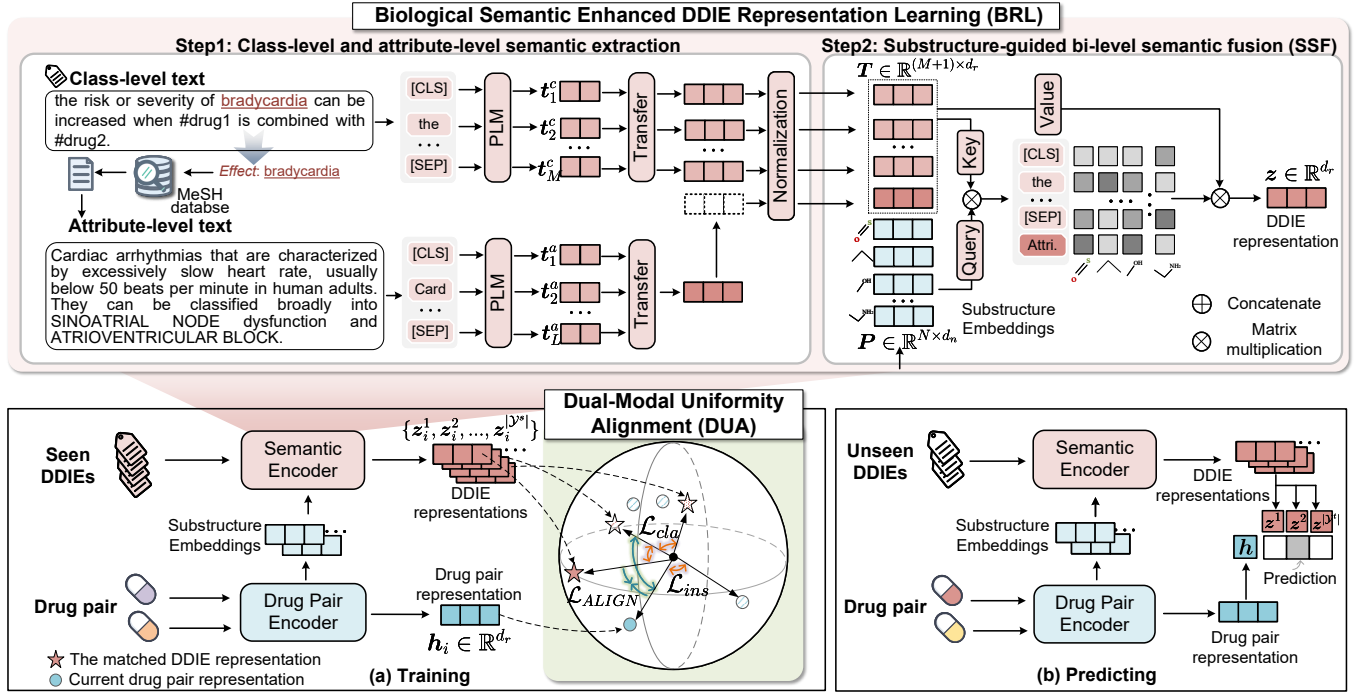


Figure 2: The overall framework of ZeroDDI.

from other classes. For example, DUET [Chen *et al.*, 2023], Sumbul *et al.* [Sumbul *et al.*, 2017] use cross-entropy loss to enforce the instance to have the highest compatibility score with its matched class; DeVise [Frome *et al.*, 2013], ALE [Akata *et al.*, 2016] and Cacheux *et al.* [Cacheux *et al.*, 2019] use rank loss to push the matched pairs and the mismatched pairs away by a margin.

In the ZS-DDIE task, although textual descriptions contain more context information, the key semantics are not emphasized, thus we add *Effect* attribute semantics to enhance the biological semantic. Moreover, inspired by the success of CLIP [Radford *et al.*, 2021], we utilize a contrastive loss to maximize the similarity of a drug pair and its matched DDIE for classification.

### 3 Methodology

#### 3.1 Preliminaries

Given the training data  $\mathcal{S} = \{((g_i, g'_i), y_i)\}_{i=1}^I$ , where each element includes a pair of drugs  $(g_i, g'_i)$  as an instance and a matched DDIE  $y_i \in \mathcal{Y}^s$  as the label, and  $\mathcal{Y}^s$  denotes the set of seen DDIE labels. Let  $\mathcal{Y}^u$  denote the set of unseen DDIE labels, which is disjoint from  $\mathcal{Y}^s$ , i.e.  $\mathcal{Y}^s \cap \mathcal{Y}^u = \emptyset$ . We denote  $\mathcal{T} = \{((g_j, g'_j), y_j)\}_{j=1}^J$  as the testing data and  $\mathcal{Y}^t$  as the set of test DDIE labels,  $y_j \in \mathcal{Y}^t$ . Herein we test the model in two scenarios: one is that  $\mathcal{Y}^t = \mathcal{Y}^u$ , i.e., all test classes come from the unseen class set, named conventional ZSL (CZSL); the other is that  $\mathcal{Y}^t = \mathcal{Y}^s \cup \mathcal{Y}^u$ ,  $\mathcal{Y}^s \subseteq \mathcal{Y}^t$ , i.e., the test classes come from both seen class set and unseen class set, named generalized ZSL (GZSL). GZSL is a more

realistic scenario where the model lacks information about whether a novel instance belongs to the seen or unseen class. The goal of the ZS-DDIE task is to leverage the training data  $\mathcal{S}$  from seen classes to learn a model that predicts the label of each instance in  $\mathcal{T}$ , which is a multi-class classification task.

#### 3.2 Overview

The architecture of ZeroDDI is shown in Figure 2. Based on the compatibility framework, we first obtain the drug pair representations and molecular substructure embeddings of drug pairs from a Drug Pair Encoder (based on [Zhu *et al.*, 2022b] and details can be seen in Appendix C). Then we get the DDIE representations by our designed BRL. Thereafter, DUA is designed to align the drug pair representations and their matched DDIE representations in unit sphere space and train the model. Finally, in the prediction phase, the unseen class whose DDIE representation has the maximum dot-product similarity score with the drug pair is the prediction.

#### 3.3 Biological Semantic Enhanced DDIE Representation Learning (BRL)

To obtain DDIE representations, we design a BRL module, which is composed of two steps: class-level and attribute-level semantic extraction, and substructure-guided bi-level semantic fusion (SSF). The former extracts the bi-level semantics of DDIEs from the attribute-level and class-level text. The latter adaptive fuses the bi-level semantic to obtain DDIE representations, preserving discriminative information as much as possible.

**Class-level and attribute-level semantic extraction.** For the class-level texts, we extract the DDIE textual descriptions from the DrugBank database [Wishart *et al.*, 2018]. For the attribute-level texts, we utilize a StanfordNLP tool [Qi *et al.*, 2019] to locate the *Effect* words in every DDIE textual description, and we extract the textual description of the *Effect* from the MeSH database<sup>1</sup>. After that, we employ a language model BioBERT [Lee *et al.*, 2020] pre-trained on large-scale biomedical corpora, to get the token features  $\{\mathbf{t}_m^c \in \mathbb{R}^{d_t}\}_{m=1}^M$  and  $\{\mathbf{t}_l^a \in \mathbb{R}^{d_t}\}_{l=1}^L$  from class-level text and attribute-level text, respectively, where  $M$  and  $L$  are the numbers of the tokens and  $d_t$  denote the token feature dimension. Since the attribute is part of a DDIE, we take the average of attribute-level token features as extra token information attached to class-level token features. Formally, we can obtain the bi-level token-wise feature  $\mathbf{T} \in \mathbb{R}^{(M+1) \times d_r}$  by:

$$\mathbf{T} = LN \left[ \phi_1([\mathbf{t}_1^c; \mathbf{t}_2^c; \dots; \mathbf{t}_M^c]); \frac{1}{L} \sum_l \phi_2(\mathbf{t}_l^a) \right] \quad (1)$$

where  $\phi_1$  and  $\phi_2$  are MLPs transferring  $\mathbb{R}^{d_t} \rightarrow \mathbb{R}^{d_r}$  to reduce the token feature dimension to  $d_r$ . The symbol  $;$  denotes concatenate.  $LN$  is Layer Normalization [Ba *et al.*, 2016], which is used to normalize the features that come from two-level semantic sources.

**Substructure-guided bi-level semantic fusion (SSF).** Different from averaging the tokens  $\mathbf{T}$  directly, we utilize a cross-modality attention mechanism to fuse the tokens discriminatively, i.e.,  $\mathbb{R}^{(M+1) \times d_r} \rightarrow \mathbb{R}^{1 \times d_r}$ . In particular, we take the molecular substructure embeddings of drug pair  $\mathbf{P} \in \mathbb{R}^{N \times d_n}$  (learned from Drug Pair Encoder) to establish the fine-grained interaction with  $\mathbf{T}$  and distill the substructure-related token semantics, where  $N$  is the number of substructures. More specifically,  $\mathbf{P}$  and  $\mathbf{T}$  are transferred to:

$$\mathbf{Q} = \mathbf{P}\mathbf{W}_Q, \mathbf{K} = \mathbf{T}\mathbf{W}_K, \mathbf{V} = \mathbf{T}\mathbf{W}_V \quad (2)$$

where  $\mathbf{Q} \in \mathbb{R}^{N \times d_r}$  denotes the queries,  $\mathbf{K} \in \mathbb{R}^{(M+1) \times d_r}$  and  $\mathbf{V} \in \mathbb{R}^{(M+1) \times d_r}$  denote keys and values.  $\mathbf{W}_Q$ ,  $\mathbf{W}_K$  and  $\mathbf{W}_V$  are learnable linear transformations. Then the cross-modality attention can be calculated by:

$$\mathbf{A} = \text{softmax} \left( \frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{r}} \right), \mathbf{O} = \mathbf{A}\mathbf{V}, \quad (3)$$

where  $\mathbf{A} \in \mathbb{R}^{N \times (M+1)}$ ,  $\mathbf{O} \in \mathbb{R}^{N \times d_r}$ . Finally, averaging over all the rows of  $\mathbf{O}$ , we can obtain the DDIE representation  $\mathbf{z} \in \mathbb{R}^{d_r}$ .

### 3.4 Dual-Modal Uniform Alignment (DUA)

We introduce a DUA as the compatibility function, which includes an alignment function to align the drug pair representations and its matched DDIE representations for classification, and a dual-modal uniformity loss to constraint both modal representations uniformly distributed in unit sphere.

<sup>1</sup> MeSH (Medical Subject Headings) is the NLM controlled vocabulary thesaurus used for indexing articles for PubMed

To realize the classification of drug pairs, we use the dot-product to measure the similarity of the drug pair representation with all DDIE semantic representations and then employ contrastive loss [Han *et al.*, 2021] to enforce the drug pair to have the highest similarity score with its matched DDIE:

$$\mathcal{L}_{ALIGN} = - \sum_{i=1}^I \log \frac{\exp(\mathbf{h}_i \cdot \mathbf{z}_i^{y_i} / \tau)}{\sum_{j \in \mathcal{Y}^s} \exp(\mathbf{h}_i \cdot \mathbf{z}_i^j / \tau)} \quad (4)$$

where  $\mathbf{h}_i$  denotes the current drug pair representation,  $\mathbf{z}_i^{y_i}$  denotes the matched DDIE representation of  $\mathbf{h}_i$ ,  $\mathbf{z}_i^j \in \{\mathbf{z}_i^1, \mathbf{z}_i^2, \dots, \mathbf{z}_i^{|\mathcal{Y}^s|}\}$  denotes the  $j$ -th DDIE representations of all DDIE representations corresponding to  $\mathbf{h}_i$ .  $\tau$  is a temperature coefficient.

To realize the uniform distribution, i.e., maximize the representation distance between classes in unit sphere space, inspired by [Lu *et al.*, 2022], we take the center of all the DDIE representations as the center of the sphere. Then we utilize the class uniformity loss  $\mathcal{L}_{cla}$  to pull every DDIE representation to unit sphere:

$$\mathcal{L}_{cla} = \frac{1}{I|\mathcal{Y}^s|} \sum_{i=1}^I \sum_{j=1}^{|\mathcal{Y}^s|} \left\{ 1 + \max_{k \in \mathcal{Y}^s \setminus \{j\}} \left[ \frac{(\mathbf{z}_i^j - \mathbf{c}_i) \cdot (\mathbf{z}_i^k - \mathbf{c}_i)}{\|\mathbf{z}_i^j - \mathbf{c}_i\| \cdot \|\mathbf{z}_i^k - \mathbf{c}_i\|} \right] \right\} \quad (5)$$

where  $\mathbf{c}_i = 1/|\mathcal{Y}^s| \cdot \sum_{j=1}^{|\mathcal{Y}^s|} \mathbf{z}_i^j$  and denotes the center of all DDIE representations of drug pair  $i$ .  $k \in \mathcal{Y}^s \setminus \{j\}$  means sampling all DDIEs except the  $j$ -th DDIE. This loss function uses the cosine similarity to pull every distance between the DDIE representations and the center representation approaching to same. Thus all DDIE representations are unified and distributed on a unit sphere.

To make the drug pair representation get consistency constraints with DDIE representations. We introduce an instance uniformity loss  $\mathcal{L}_{ins}$  to unify every drug pair  $i$  to the sphere with the center of  $\mathbf{c}_i$ :

$$\mathcal{L}_{ins} = \frac{1}{I} \sum_{i=1}^I \left\{ 1 + \max_{k \in \mathcal{S} \setminus \{i\}} \left[ \frac{(\mathbf{h}_i - \mathbf{c}_i) \cdot (\mathbf{h}_k - \mathbf{c}_i)}{\|\mathbf{h}_i - \mathbf{c}_i\| \cdot \|\mathbf{h}_k - \mathbf{c}_i\|} \right] \right\} \quad (6)$$

where  $k \in \mathcal{S} \setminus \{i\}$  is the sample randomly drawn from a batch set of training data excluding  $i$ . According to Eq. (5) and Eq. (6), the dual-modal uniformity loss  $\mathcal{L}_{UNI}$  can be expressed as:

$$\mathcal{L}_{UNI} = \mathcal{L}_{cla} + \mathcal{L}_{ins} \quad (7)$$

### 3.5 Training and Prediction

For training our model ZeroDDI, we optimize the total loss that combines Eq. (4) and Eq. (7):

$$\mathcal{L} = \mathcal{L}_{ALIGN} + \lambda \mathcal{L}_{UNI} \quad (8)$$

where  $\lambda$  is hyper-parameters.

For the zero-shot DDIE prediction, we use the trained Drug Pair Encoder to obtain drug pair representation  $\mathbf{h}$ . Its DDIE  $\hat{y}$

can be predicted by researching the highest dot-product similarity between  $\mathbf{h}$  and all DDIE representations in test set:

$$\hat{y} = \operatorname{argmax}_{j \in \mathcal{Y}^t} \mathbf{h} \cdot \mathbf{z}^j, \quad (9)$$

where  $\mathcal{Y}^t$  is the set of test DDIEs,  $\mathbf{z}^j$  denotes the  $j$ -th DDIE representation learned by the trained BRL.

## 4 Experiments

### 4.1 Experimental Setup

#### Datasets

DrugBank [Wishart *et al.*, 2018] is a widely used data source for DDIE prediction, which contains drugs, DDIs, textual descriptions of DDIEs, and the molecular structures of drugs. Based on raw data from DrugBank v5.1.9, we construct a ZS-DDIE dataset by annotating attributes for every DDI manually, which contains a total of 2,004 approved drugs, 394,118 DDIs, 175 DDIEs with unique textual descriptions, 2 *Signs*, 3 *Patterns* and 114 *Effects* with corresponding attribute-level text. We rank the classes of DDIEs in the order of the number of instances from more to less and split DDIEs into seen classes and unseen classes, in which we take the less number of DDIEs as unseen classes to simulate the real situation of DDIE that are newly observed. 107 DDIEs are viewed as seen classes  $\mathcal{Y}^s$ , which contains all attributes, and the rest 68 DDIEs are viewed as unseen classes  $\mathcal{Y}^u$ . Additionally, every DDI is associated with a DDIE in this dataset. Drug pairs in the dataset are asymmetric, in that the roles of drugs are different in one DDIE [Feng *et al.*, 2023]. More details of the dataset can be found in Appendix D.1.

#### Baselines

As mentioned in the related work, only one DDIE method (3DGT-DDI) has the potential to be applied to the ZS-DDIE task. In addition, under the compatibility framework, we also compare several popular compatibility functions in computer vision and equip them with two kinds of DDIE semantic representations. One learned from a binary vector of attributes (namely attribute-based representation) and the other learned from a PLM embedding of class textual description (namely class-based representation), respectively. These baselines have the same Drug Pair Encoder as our method.

- **3DGT-DDI** [Lv *et al.*, 2023]. The DDIE semantic encoder of 3DGT-DDI is composed of a PLM and a CNN. To enable this model to deal with the ZS-DDIE task, we replace its binary classifier with a compatibility function incorporating cross-entropy loss.
- **ZSLHinge**. ZSLHinge utilizes a hinge rank loss from the classic ZSL method DeViSE [Frome *et al.*, 2013], which can produce a higher dot-product similarity between an instance and the matched class than between the instance and other randomly chosen classes by a margin.
- **ZSLCE**. ZSLCE utilizes the cross-entropy loss from [Chen *et al.*, 2023] and [Sumbul *et al.*, 2017], which maximizes the similarity score of the instance and its matching class.

- **ZSLTriplet**. ZSLTriplet utilizes a ZSL triplet loss from [Cacheux *et al.*, 2019], which improves the standard triplet loss by adding a flexible semantic margin, partial normalization, and relevance weighting.

More details of baselines can be found in Appendix D.2.

#### Evaluation Protocols

The experiments are performed on two common zero-shot learning scenarios: CZSL and GZSL, which are described in Preliminaries. We train the model on data from seen classes. **In the CZSL scenario**, we validate and test the model in 3-fold of the unseen classes, i.e., in each fold, one-third of unseen classes are used for the validation, and the rest are used for the testing. **In the GZSL scenario**, to evaluate the performance of seen and unseen prediction simultaneously, a part of the seen classes and instances are added into the GZSL validation and testing set. Note that all instances from seen classes in the validation and testing set are not included in the training set.

#### Metrics

In the CZSL scenario, the performance is measured with top-k accuracy ( $k=1, 3, 5$  and denoted as  $\mathbf{acc}_{@1}^u$ ,  $\mathbf{acc}_{@3}^u$  and  $\mathbf{acc}_{@5}^u$ ). In addition, we calculate the average of top-1 accuracy across all unseen classes, denoted as  $\mathbf{acc}_{ave}^u$ . In the GZSL scenario, we calculate the  $\mathbf{acc}_{@1}$  or  $\mathbf{acc}_{ave}$  on seen classes (denoted as  $\mathbf{acc}_{@1}^s$  or  $\mathbf{acc}_{ave}^s$ ) and unseen classes, and calculate the harmonic mean of them:  $\mathbf{acc}^H = 2 * (\mathbf{acc}^s * \mathbf{acc}^u) / (\mathbf{acc}^s + \mathbf{acc}^u)$  [Xian *et al.*, 2018]. In addition,  $\mathbf{acc}_{b1}^s$  or  $\mathbf{acc}_{b1}^u$  is used to evaluate the seen-unseen binary classification results. We also utilize the proportion  $\mathbf{P}^s = \mathbf{acc}_{@1}^s / \mathbf{acc}_{b1}^s$  to evaluate the multi-class classification performance under the correct binary classification.

The implementation details, hyper-parameter sensitivity analysis, and model configuration can be found in Appendix D due to space constraints.

### 4.2 Comparison with Baselines

We compare ZeroDDI with baselines in both CZSL and GZSL scenarios, and the results are shown in Table 1. Overall, ZeroDDI achieves the best performances on all unseen metrics (i.e.,  $\mathbf{acc}^u$ ,  $\mathbf{P}^u$  and  $\mathbf{acc}^H$ ), which indicates that our method has a superior ability to deal with unseen DDIE prediction problem. Specifically, we have the following observations: (1) Compared with 3DGT-DDI, our method has salient advantages in the ZS-DDIE task, as evidenced by the 9.32% performance gain on the unseen metrics  $\mathbf{acc}_{ave}^u$  of CZSL and 7.92% performance gain on  $\mathbf{acc}_{ave}^u$  of GZSL. This not only shows the promising unseen prediction performance of our method but also underscores the necessity of specifically tailoring methods for ZS-DDIE tasks. (2) Comparing the performances of class-based and attribute-based models, we find that, in most cases (i.e., ZSLHinge, ZSLCE, and ZeroDDI), the class-based representations perform better. This demonstrates that, in the ZS-DDIE task, the transferability of class-based representations learned by a PLM may be greater than that of attribute-based representations encoded by the shared attributes among classes. Thus our BRL module is constructed based on the PLM. (3) ZeroDDI

Model	CZSL					GZSL						
	$acc_{ave}^u$	$acc_{@1}^u$	$acc_{@3}^u$	$acc_{@5}^u$	$acc_{ave}^u$	$acc_{@1}^u$	$acc_{ave}^s$	$acc_{@1}^s$	$acc_{ave}^H$	$acc_{@1}^H$	$P^u$	$P^s$
3DGT-DDI	8.25	13.50	27.93	39.27	8.21	12.55	48.30	45.29	14.03	19.65	14.89	90.15
ZSLHinge <sup>1</sup>	5.47	5.75	21.17	30.45	5.39	5.69	23.14	22.65	8.06	9.09	5.87	<b>98.05</b>
ZSLHinge <sup>2</sup>	12.80	10.69	28.74	42.46	10.01	10.1	54.93	53.02	16.87	15.86	11.08	95.48
ZSLCE <sup>1</sup>	11.35	11.32	30.21	44.72	7.55	8.90	54.57	61.75	12.97	15.14	13.65	75.19
ZSLCE <sup>2</sup>	15.08	16.86	39.04	53.34	<u>12.73</u>	14.84	76.20	71.03	<u>21.78</u>	24.27	17.76	97.01
ZSLTriplet <sup>1</sup>	<u>15.09</u>	14.77	40.51	<u>59.84</u>	<u>10.23</u>	10.92	<b>78.71</b>	<b>74.20</b>	18.05	18.61	15.09	94.15
ZSLTriplet <sup>2</sup>	13.51	18.80	28.36	51.27	11.75	<u>16.22</u>	66.48	62.37	19.89	20.19	19.13	96.25
ZeroDDI <sup>1</sup>	13.16	15.79	40.04	52.94	10.65	13.31	72.70	69.35	18.19	20.95	15.48	<u>97.87</u>
ZeroDDI <sup>2</sup>	14.66	<u>19.22</u>	<u>42.65</u>	56.28	11.72	15.10	74.09	<u>73.66</u>	20.21	<u>24.83</u>	<u>19.42</u>	96.93
ZeroDDI	<b>17.57</b>	<b>21.35</b>	<b>47.81</b>	<b>66.16</b>	<b>16.13</b>	<b>19.12</b>	<u>77.25</u>	70.67	<b>26.48</b>	<b>29.14</b>	<b>22.09</b>	97.02

Table 1: Performance (in %) comparisons of ZeroDDI with baselines in the CZSL and GZSL scenario. The best and suboptimal results are highlighted in **bold** and underline, respectively. Note that Model<sup>1</sup> (or Model<sup>2</sup>) denotes using attribute-based (or class-based) representations learned from a binary vector of attributes (or the PLM of class textual description where the semantic encoder is the same as that in ZeroDDI).




Figure 3: Performance (in %) comparisons of ZeroDDI with its variants in the CZSL and GZSL scenarios.

outperforms ZeroDDI<sup>1</sup> and ZeroDDI<sup>2</sup> demonstrates that the designed BRL, which emphasizes the *Effect* attribute with key biological semantics on the class-based representation, can combine the advantages of both attributes and DDIE textual description for unseen DDIE prediction. (4) Compared with ZSLTriplet<sup>1or2</sup>, ZeroDDI<sup>1or2</sup> also get competitive advantages, which reveals that our DUA strategy benefits unseen prediction.

### 4.3 Ablation Study

We conduct an ablation study to validate the contribution of important components in ZeroDDI, i.e., attribute-level semantics, the SSF component, and dual-modal uniformity loss. The results are shown in Figure 3. Specifically, (1) after removing the input of attribute-level semantics (w/o Attri.), the performance drops, which indicates that the attribute-level semantics extracted by our method can improve the performance of the model for the ZS-DDIE task. We will further discuss this in Section 4.4. (2) The model shows a significant decline after replacing the fusion component SSF with the

PLM	Attri. Level Semantic	CZSL	GZSL	
		$acc_{ave}^u$	$acc_{ave}^u$	$acc_{ave}^H$
SCIBERT	✓	13.79	12.47	21.31
	-	10.84	8.58	14.65
PubMedBERT	✓	15.56	11.97	20.44
	-	9.60	6.94	12.40
BioBERT	✓	<b>17.57</b>	<b>16.13</b>	<b>26.48</b>
	-	14.66	11.72	20.21

Table 2: Performance (in %) comparisons of ZeroDDI with its variants that replace the BioBERT by different PLM models in CZSL and GZSL scenarios.

average operation of bi-level token-wise features (w/o SSF), which proves the necessity of selective bi-level fusion and the effectiveness of molecular substructure for distilling discriminative semantic information. (3) Removing the dual-modal uniformity loss (w/o  $\mathcal{L}_{UNI}$ ) or any one part of this loss function (w/o  $\mathcal{L}_{cla}$  or w/o  $\mathcal{L}_{ins}$ ) undermines the accuracy, demonstrating that the uniformity loss assist in improve the discriminative ability, and  $\mathcal{L}_{cla}$  and  $\mathcal{L}_{ins}$  are all useful for ZS-DDI task. We will further discuss this in Section 4.5.

### 4.4 Effectiveness of Bio-Enhanced Representation Learning

In this section, we further discuss the effectiveness of two components (attribute-level semantics and SSF) of BRL in unseen DDIE prediction.

For attribute-level semantics, we compare the variants of ZeroDDI (by replacing the BioBERT with SCIBERT [Beltagy *et al.*, 2019] or PubMedBERT [Gu *et al.*, 2022] whose versions are shown in Appendix D.3) equipped with/without the attribute-level semantic information. Table 2 shows that our attribute-level semantics can provide additional stable gains of transferability, leading to better performance in unseen DDIE prediction no matter which PLM is based. For SSF, we visualize an example to verify that SSF can con-






Figure 4: The visualization of an example of DDIE textual description with its corresponding drug pair molecular structures. The attention scores between the word "hypertensive" and substructures are highlighted in red colour.

struct a reasonable relationship between molecular substructures and the bi-level text tokens. We randomly select two drug pairs from the hits of a DDIE in the test set, then visualize them and highlight one word "hypertensive" with its most relevant substructures in Figure 4. As the website<sup>2</sup> reports that "coadministration with *adrenergic agonists* with hypertensive potential may lead to increased cardiac output and blood pressure due to the vasoconstricting effects of *ergot derivatives*". In the example, Cabergoline and Dihydroergotamine are *ergot derivatives*, and the SSF can highlight the part of their parent structure of ergot derivatives, which proves the effectiveness of the constructed relationship between tokens and substructures and indicate the potential interpretability of SSF.

#### 4.5 Effectiveness of Dual-Modal Uniform Distribution

We further discuss the effectiveness of dual-modal uniformity loss for the class imbalance issue in seen and unseen DDIE prediction by comparing ZeroDDI with ZeroDDI (w/o  $\mathcal{L}_{UNI}$ ). To simulate the relatively balanced scenario, we reconstruct the training set with the imbalance ratio  $\rho$  is 1:100, where  $\rho$  is defined as the number of samples in the most frequent class divided by that of the least frequent class [Li *et al.*, 2022] (here we denote least frequent class has ten instances at least). We randomly select five unseen and seen DDIE classes from the results of the test set, respectively, and visualize them in Figure 5. From the results, we have the following observations: (1) For unseen DDIE prediction, the drop of accuracy from the scenario  $\rho = 1:100$  to 1:10000 shows that the class imbalance has an impact on unseen DDIE prediction to some extent, which demonstrates that the necessity of tackle class unbalance challenge in unseen DDIE prediction. (2) The class imbalance can cause a drop of performance in both seen and unseen classes and lead to unclear and inseparable decision boundaries as shown in the right column of Figure 5(a) and (b), while our method can decelerate this problem. This indicates that  $\mathcal{L}_{UNI}$  can improve the discriminability of the model and mitigate the adverse effects of class imbalance.

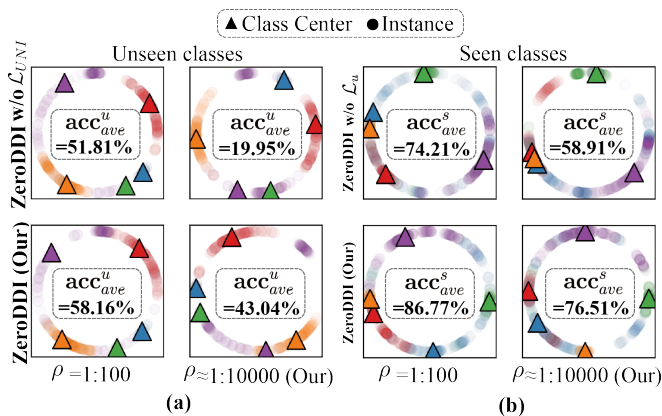


Figure 5: The visualization of drug pair representation distribution of test unseen and seen classes in GZSL scenario. Class Center denotes the center of all instances in a class.  $\rho$  denotes imbalance ratios of training data. The  $\text{acc}_{ave}$  here is the average accuracy of five DDIE classes.

#### 4.6 Zero-Shot DDIE Application Analysis

In this section, we conduct an application analysis to validate the factuality of the zero-shot DDIE setting in this work and evaluate the practical capabilities of ZeroDDI. Firstly, we choose the latest dataset (i.e., DrugBank v5.1.11) as the Novel dataset and take the dataset used in this work as the Existing Dataset. Through the data process for the Novel dataset, we find that, in the Novel Dataset, 6 DDIEs are composed of the attributes from the Existing Dataset but are not included in the Existing Dataset. It not only proves that novel DDIEs are increasing but also proves that novel DDIEs can be composed of existing attributes and further indicates the practical significance of our evaluation setting. Then, we use the Existing Dataset to train the ZeroDDI and use the trained model to predict 6 novel DDIEs in the Novel Dataset. When there are no labelled instances in the training set, ZeroDDI can also achieve 61.11 % in  $\text{acc}_{ave}^u$ , which shows the application ability of our method. More details and experiments are shown in Appendix E.

## 5 Conclusion

This is the first work that pays attention to the zero-shot DDIE prediction and proposes a novel method, called ZeroDDI, to predict zero-shot DDIEs. Specifically, we designed a biological semantic enhanced DDIE representation learning module to learn suitable DDIE representations containing enhanced key biological semantic and substructure-guided discriminative semantics, leading to better knowledge transferability from seen DDIEs to unseen DDIEs. Moreover, we design a dual-modal uniformity alignment strategy to uniform the distribution of drug pair representations and DDIE representations in unit sphere and thus mitigate the class imbalance issue. The extensive experiments show that ZeroDDI can produce superior performance in zero-shot DDIE prediction, our designed module can effectively promote the performances, and ZeroDDI is a promising tool for practice application.

<sup>2</sup> <https://go.drugbank.com/drug-interaction-checker>

## Acknowledgments

This work was supported by the National Natural Science Foundation of China (62372204, 62072206, 61772381, 62102158); Huazhong Agricultural University Scientific & Technological Self-innovation Foundation; Fundamental Research Funds for the Central Universities (2662021JC008, 2662022JC004). The funders have no role in study design, data collection, data analysis, data interpretation, or writing of the manuscript.

## Contribution Statement

Wen Zhang is the corresponding author of this study. Ziyang Wang is responsible for the paper writing, data annotating, experimental code, figure drawing, and experimental design. Zhankun Xiong helped with paper writing, figure drawing, and experimental design, which contributed equally to this paper. Feng Huang provided help in paper writing and Xuan Liu provided help in data annotating.

## References

- [Akata *et al.*, 2016] Zeynep Akata, Florent Perronnin, Zaid Harchaoui, and Cordelia Schmid. Label-Embedding for Image Classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(7):1425–1438, 2016.
- [Al-Halah and Stiefelwagen, 2017] Ziad Al-Halah and Rainer Stiefelwagen. Automatic discovery, association estimation and learning of semantic attributes for a thousand categories. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 614–623, 2017.
- [Ba *et al.*, 2016] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [Beltagy *et al.*, 2019] Iz Beltagy, Kyle Lo, and Arman Cohen. SciBERT: A Pretrained Language Model for Scientific Text. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3615–3620, Hong Kong, China, 2019. Association for Computational Linguistics.
- [Cacheux *et al.*, 2019] Yannick Le Cacheux, Herve Le Borgne, and Michel Crucianu. Modeling inter and intra-class relations in the triplet loss for zero-shot learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10333–10342, 2019.
- [Chen *et al.*, 2023] Zhuo Chen, Yufeng Huang, Jiaoyan Chen, Yuxia Geng, Wen Zhang, Yin Fang, Jeff Z Pan, and Huajun Chen. Duet: Cross-modal semantic grounding for contrastive zero-shot learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 405–413, 2023.
- [Deng *et al.*, 2020] Yifan Deng, Xinran Xu, Yang Qiu, Jingbo Xia, Wen Zhang, and Shichao Liu. A multimodal deep learning framework for predicting drug–drug interaction events. *Bioinformatics*, 36(15):4316–4322, 05 2020.
- [Deng *et al.*, 2021] Yifan Deng, Yang Qiu, Xinran Xu, Shichao Liu, Zhongfei Zhang, Shanfeng Zhu, and Wen Zhang. META-DDIE: predicting drug–drug interaction events with few-shot learning. *Briefings in Bioinformatics*, 23(1), December 2021.
- [Feng *et al.*, 2023] Yue-Hua Feng, Shao-Wu Zhang, Yi-Yang Feng, Qing-Qing Zhang, Ming-Hui Shi, and Jian-Yu Shi. A social theory-enhanced graph representation learning framework for multitask prediction of drug–drug interactions. *Briefings in Bioinformatics*, 24(1), January 2023.
- [Frome *et al.*, 2013] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. *Advances in neural information processing systems*, 26, 2013.
- [Gu *et al.*, 2022] Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. Domain-Specific Language Model Pretraining for Biomedical Natural Language Processing. *ACM Transactions on Computing for Healthcare*, 3(1):1–23, January 2022.
- [Han *et al.*, 2021] Zongyan Han, Zhenyong Fu, Shuo Chen, and Jian Yang. Contrastive embedding for generalized zero-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2371–2381, 2021.
- [Ji *et al.*, 2022] Zhong Ji, Xuejie Yu, Yunlong Yu, Yanwei Pang, and Zhongfei Zhang. Semantic-guided class-imbalance learning model for zero-shot image classification. *IEEE Transactions on Cybernetics*, 52(7):6543–6554, 2022.
- [Jin *et al.*, 2017] Bo Jin, Haoyu Yang, Cao Xiao, Ping Zhang, Xiaopeng Wei, and Fei Wang. Multitask dyadic prediction and its application in prediction of adverse drug–drug interaction. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31, 2017.
- [Jin *et al.*, 2023] Shuting Jin, Yue Hong, Li Zeng, Yinghui Jiang, Yuan Lin, Leyi Wei, Zhuohang Yu, Xiangxiang Zeng, and Xiangrong Liu. A general hypergraph learning algorithm for drug multi-task predictions in micro-to-macro biomedical networks. *PLOS Computational Biology*, 19(11):1–27, November 2023. Publisher: Public Library of Science.
- [Lampert *et al.*, 2013] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Attribute-based classification for zero-shot visual object categorization. *IEEE transactions on pattern analysis and machine intelligence*, 36(3):453–465, 2013. Publisher: IEEE.
- [Lee *et al.*, 2020] Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240, 2020. Publisher: Oxford University Press.
- [Li *et al.*, 2015] Xirong Li, Shuai Liao, Weiyu Lan, Xiaoyong Du, and Gang Yang. Zero-shot image tagging by



- hierarchical semantic embedding. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 879–882, 2015.
- [Li *et al.*, 2022] Tianhong Li, Peng Cao, Yuan Yuan, Lijie Fan, Yuzhe Yang, Rogerio S Feris, Piotr Indyk, and Dina Katabi. Targeted supervised contrastive learning for long-tailed recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6918–6928, 2022.
- [Lin *et al.*, 2021] Shenggeng Lin, Yanjing Wang, Lingfeng Zhang, Yanyi Chu, Yatong Liu, Yitian Fang, Mingming Jiang, Qiankun Wang, Bowen Zhao, Yi Xiong, et al. Mdfsa-ddi: predicting drug–drug interaction events based on multi-source drug fusion, multi-source feature fusion and transformer self-attention mechanism. *Briefings in Bioinformatics*, 2021.
- [Lin *et al.*, 2022] Jiacheng Lin, Lijun Wu, Jinhua Zhu, Xiaobo Liang, Yingce Xia, Shufang Xie, Tao Qin, and Tie-Yan Liu. R2-DDI: relation-aware feature refinement for drug–drug interaction prediction. *Briefings in Bioinformatics*, December 2022.
- [Lu *et al.*, 2022] Bin Lu, Xiaoying Gan, Lina Yang, Weinan Zhang, Luoyi Fu, and Xinbing Wang. Geometer: Graph Few-Shot Class-Incremental Learning via Prototype Representation. *KDD2022*, 2022.
- [Lv *et al.*, 2023] Qiuji Lv, Jun Zhou, Ziduo Yang, Haohuai He, and Calvin Yu-Chian Chen. 3d graph neural network with few-shot learning for predicting drug–drug interactions in scaffold-based cold start scenario. *Neural Networks*, 165:94–105, 2023.
- [Nawaz *et al.*, 2022] Shah Nawaz, Jacopo Cavazza, and Alessio Del Bue. Semantically grounded visual embeddings for zero-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4589–4599, 2022.
- [Qi *et al.*, 2019] Peng Qi, Timothy Dozat, Yuhao Zhang, and Christopher D Manning. Universal dependency parsing from scratch. *arXiv preprint arXiv:1901.10457*, 2019.
- [Radford *et al.*, 2021] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, and others. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [Reed *et al.*, 2016] Scott Reed, Zeynep Akata, Honglak Lee, and Bernt Schiele. Learning deep representations of fine-grained visual descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 49–58, 2016.
- [Ryu *et al.*, 2018] 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 of the United States of America*, 115(18):E4304–E4311, May 2018.
- [Sumbul *et al.*, 2017] Gencer Sumbul, Ramazan Gokberk Cinbis, and Selim Aksoy. Fine-grained object recognition and zero-shot learning in remote sensing imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 56(2):770–779, 2017. Publisher: IEEE.
- [Vilar *et al.*, 2014] Santiago Vilar, Eugenio Uriarte, Lourdes Santana, Tal Lorberbaum, George Hripcsak, Carol Friedman, and Nicholas P. Tatonetti. Similarity-based modeling in large-scale prediction of drug–drug interactions. *Nature protocols*, 9(9):2147–2163, September 2014.
- [Wishart *et al.*, 2018] David S Wishart, Yannick D Feunang, An C Guo, Elvis J Lo, Ana Marcu, Jason R Grant, Tanvir Sajed, Daniel Johnson, Carin Li, Zinat Sayeeda, et al. Drugbank 5.0: a major update to the drugbank database for 2018. *Nucleic acids research*, 46(D1):D1074–D1082, 2018.
- [Xian *et al.*, 2018] Yongqin Xian, Christoph H Lampert, Bernt Schiele, and Zeynep Akata. Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *IEEE transactions on pattern analysis and machine intelligence*, 41(9):2251–2265, 2018. Publisher: IEEE.
- [Xiong *et al.*, 2023] Zhankun Xiong, Shichao Liu, Feng Huang, Ziyang Wang, Xuan Liu, Zhongfei Zhang, and Wen Zhang. Multi-relational contrastive learning graph neural network for drug–drug interaction event prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 5339–5347, 2023.
- [Yu *et al.*, 2022] Hui Yu, ShiYu Zhao, and JianYu Shi. STNN-DDI: a Substructure-aware Tensor Neural Network to predict Drug–Drug Interactions. *Briefings in Bioinformatics*, 23(4):bbac209, 2022. Publisher: Oxford University Press.
- [Zhu *et al.*, 2018] Yizhe Zhu, Mohamed Elhoseiny, Bingchen Liu, Xi Peng, and Ahmed Elgammal. A generative adversarial approach for zero-shot learning from noisy texts. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1004–1013, 2018.
- [Zhu *et al.*, 2022a] 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*, pages 4757–4761, 2022.
- [Zhu *et al.*, 2022b] 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*, pages 4757–4761, 2022.
- [Zhu *et al.*, 2023] Fangqi Zhu, Yongqi Zhang, Lei Chen, Bing Qin, and Ruifeng Xu. Learning to Describe for Predicting Zero-shot Drug–Drug Interactions. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14855–14870, 2023.