

BERTKG-DDI: Towards Incorporating Entity-specific Knowledge Graph Information in Predicting Drug-Drug Interactions

Ishani Mondal

Microsoft Research Lab,
Lavelle Road, Bengaluru, India
ishani340@gmail.com

Abstract

Off-the-shelf biomedical embeddings obtained from the recently released various pre-trained language models (such as BERT, XLNET) have demonstrated state-of-the-art results (in terms of accuracy) for the various natural language understanding tasks (NLU) in the biomedical domain. Relation Classification (RC) falls into one of the most critical tasks. In this paper, we explore how to incorporate domain knowledge of the biomedical entities (such as drug, disease, genes), obtained from Knowledge Graph (KG) Embeddings, for predicting Drug-Drug Interaction from textual corpus. We propose a new method, **BERTKG-DDI**, to combine drug embeddings obtained from its interaction with other biomedical entities along with domain-specific BioBERT embedding-based RC architecture. Experiments conducted on the DDIExtraction 2013 corpus clearly indicate that this strategy improves other baselines architectures by 4.1% macro F1-score.

Introduction

During the concurrent administration of multiple drugs to a patient, there seems to be a possibility in which an ailment might get cured or it can lead to serious side-effects. These type of interactions are known as Drug-Drug Interactions (DDIs). Predicting drug-drug interactions (DDI) is a difficult task as it requires to understand the underlying action principle of the interacting drugs. Numerous efforts by the researchers have been observed recently in terms of automatic extraction of DDIs from the textual corpus (Sahu and Anand 2018), (Liu et al. 2016), (Sun et al. 2019), (Li and Ji 2019), (Mondal 2020) and predicting unknown DDI from KG (Purkayastha et al. 2019). Automatic extraction of DDI from texts helps to maintain large-scale databases and thereby facilitate the medical experts in their diagnosis.

In parallel to the progress of DDI extraction from the textual corpus, some efforts have been observed recently where the researchers came up with various strategies of augmenting chemical structure information of the drugs and textual description of the drugs (Zhu et al. 2020) to improve Drug-Drug Interaction prediction performance from corpus and Knowledge Graphs. The DDI Prediction from the textual corpus has been framed by the earlier researchers as relation

classification problem (Sahu and Anand 2018), (Liu et al. 2016), (Sun et al. 2019), (Li and Ji 2019) using CNN or RNN-based neural networks.

Recently, with the massive success of the pre-trained language models (Devlin et al. 2019), (Yang et al. 2019) in many NLP classifications, we formulate the problem of DDI classification as a relation classification task by leveraging both entities and contextual information. We propose a model that leverages both domain-specific contextual embeddings (Bio-BERT) (Lee et al. 2019) from the target entities (drugs) and also its external information. In the recent years, representation learning has played a pivotal role in solving various machine learning tasks.

In this work, we explore the direction of augmenting graph embeddings to predict relation between two drugs from the textual corpus. We have made use of an in-house Knowledge Graph (*Bio-KG*) after curating the interactions among drugs, diseases, genes from multiple ontologies. In order to understand the complex underlying mechanism of interactions among the biomedical entities, we employ translation-based and semantics preserving heterogeneous graph embeddings on *Bio-KG* and augment the entities representation jointly to train the relation classification model. Experiments conducted on the DDIExtraction 2013 corpus (Herrero-Zazo et al. 2013) reveals that this method outperforms the existing baseline models and is in line with the new direction of research of fusing various information to DDI prediction. In a nutshell, the major contributions of this work are summarized as follows:

1. We propose a novel method that jointly leverages textual and external Knowledge information to classify relation type between the drug pairs mentioned in the text showing the efficacy of external entity specific information.
2. Our method achieves new state-of-the-art performance on DDI Extraction 2013 corpus.

Problem Statement

Given an input instance or sentence s with two target drug entities d_1 and d_2 , the task is to classify the type of relation (y) the drugs hold between them, $y \in (y_1, \dots, y_N)$. Here N denotes the number of relation types.

Methodology

Text-based Relation Classification

Our model for extracting DDIs from texts is based on the pre-trained BERT-based relation classification model by (Wu and He 2019). Given a sentence s with drugs d_1 and d_2 , let the final hidden state output from BERT module is H . Let the vectors H_i to H_j are the final hidden state vectors from BERT for entity d_1 , and H_k to H_m are the final hidden state vectors from BERT for entity d_2 . An average operation is applied to obtain the vector representation for each of the drug entities. An activation operation \tanh is applied followed by a fully connected layer to each of the two vectors, and the output for d_1 and d_2 are H'_1 and H'_2 respectively.

$$H'_1 = W[\tanh(\frac{1}{(j-i+1)} \sum_{t=i}^j H_t)] + b \quad (1)$$

$$H'_2 = W[\tanh(\frac{1}{(m-k+1)} \sum_{t=k}^m H_t)] + b \quad (2)$$

The weights (W) and bias (b) parameters are shared. For the final hidden state vector of the first token ('[CLS]'), we also add an activation operation and a fully connected layer, which is formally expressed as:

$$H'_0 = W_0(\tanh(H_0)) + b_0 \quad (3)$$

Matrices W_0, W_1, W_2 have the same dimensions, i.e. $W_0 \in R^{d \times d}, W_1 \in R^{d \times d}, W_2 \in R^{d \times d}$, where d is the hidden state size from BERT. We concatenate H'_0, H'_1 and H'_2 and then add a fully connected layer and a softmax layer, which is expressed as :

$$h'' = W_3[\text{concat}(H'_0, H'_1, H'_2)] + b_3 \quad (4)$$

$$y'_t = \text{softmax}(h'') \quad (5)$$

$W_3 \in R^{N \times 3d}$, and y'_t is the softmax probability output over N . In Equations (1), (2), (3), (4) the bias vectors are b_0, b_1, b_2, b_3 . We use cross entropy as the loss function. We denote this text-based architecture as *BERT-Text-DDI*.

Entity Representation from KG

To infuse external information of the entities in relation classification task, we obtain the representation of two *Drug* entities mentioned in each input instance of the relation classification task. We use an in-house heterogeneous biomedical Knowledge Graph (*Bio-KG*) consisting of the interactions of target-target, drug-drug, drug-disease, drug-target, disease-disease, disease-target interactions from a large number of ontologies such as : DrugBank¹, BioSNAP², UniProt³ (The UniProt Consortium 2016). The overall statistics of *Bio-KG* has been enumerated in table 1. The real-world information/facts observed in the *Bio-KG* are stored as a collection of triples in the form (h, r, t) . Each triple is composed of a head entity $h \in E$, a tail entity $t \in E$, and a relation $r \in$

¹<https://go.drugbank.com/>

²<http://snap.stanford.edu/biodata/>

³<https://www.uniprot.org/>

Node Types	Count	Edge Types	Count
Drug	6512	Drug-Target	15245
Target	30098	Target-Target	77108
Disease	23458	Drug-Disease	84745
		Disease-Disease	35382
		Disease-Target	31161
Total Nodes	60068	Total Edges	243641

Table 1: Statistics of Bio-KG.

R between them, e.g., (*paracetamol, treats, fever*). The fact that *paracetamol* is effective in curing *fever* is being stored in *Bio-KG*. In this case, E denotes set of entities, and R denotes the set of relations. There are three different types of E in *Bio-KG* such as *drugs, diseases, targets* and five different types of R such as *target-target, drug-disease, drug-target, disease-disease, disease-target interactions*.

The aim of a Knowledge Graph embedding is to embed the entities and relations into a low-dimensional continuous vector space, so as to simplify the computations on the KG. They mostly use facts in the KG to perform the embedding task, enforcing embedding to be compatible with the facts. They provide a generalizable context about the overall Knowledge Graph (KG) that can be used to infer the relations. In this work, we employ some off-the-shelf KG embeddings to encode the representation of each of the drugs (in terms of their relationship with other entities). The knowledge graph embeddings are computed so that they satisfy certain properties; i.e., they follow a given KGE model. These KGE models define different score functions that measure the distance of two entities relative to its relation type in the low-dimensional embedding space. These score functions are used to train the KGE models so that the entities connected by relations are close to each other while the entities that are not connected are far away. Some of the KGEs used in our experiments as explained below:

- **TransE** (Bordes et al. 2013): Given a fact (h, r, t) , the relation in TransE is interpreted as a translation vector r so that the embedded entities h and t can be connected by r , i.e., $h + r \approx t$ when (h, r, t) holds. The scoring function is defined as (negative) distance between $h + r$ and t , i.e.,

$$f_r(h, t) = \| h + r - t \| \quad (6)$$

- **TransR** (Lin et al. 2015): Given a fact (h, r, t) , TransR first projects the entity representations h and t into the space specific to relation r , Here M_r is a projection matrix from the entity space to the relation space of r , the scoring function is:

$$h_t = M_r h, t_t = M_r t \quad (7)$$

- **RESCAL** (Nickel, Tresp, and Krieger 2011): Each relation in RESCAL is represented as a matrix which models pairwise interactions between latent factors. The score of a fact (h, r, t) is defined by a bi-linear function where h, t are vector representations of the entities, and M_r is a matrix associated with the relation. This score captures

pairwise interactions between all components of h and t :

$$f_r(h, t) = h^T M_r t = \sum_{i=0}^{d-1} \sum_{j=0}^{d-1} [M_r]_{ij} * [h]_i * [t]_j \quad (8)$$

- **DistMult** (Yang et al. 2015): DistMult simplifies RESCAL by restricting M_r to diagonal matrices. For each relation r , it introduces a vector embedding r and requires $M_r = \text{diag}(r)$. The scoring function is defined as:

$$f_r(h, t) = h^T \text{diag}(r) t = \sum_{i=0}^{d-1} r_i * [h]_i * [t]_j \quad (9)$$

This score captures pairwise interactions between only the components of h and t along the same dimension, and reduces the number of parameters to $O(d)$ per relation.

From *Bio-KG*, we train these KG Embeddings and obtain the representation of all the nodes. In our case, we are only interested in obtaining the representation of *drug* nodes. We denote the KG representation of drug d as kge .

BERTKG-DDI

From the input instance s with two tagged target drug entities d_1 and d_2 , we obtain the KG embedding representation of two drugs kge_1 and kge_2 respectively using *Bio-KG*. We concatenate these two embeddings kge_1 and kge_2 and pass those through a fully connected layer as represented below:

$$kge = W[\text{concat}(kge_1, kge_2)] + b \quad (10)$$

W and b are the parameters of the fully-connected layer of the KG representation of kge_1 and kge_2 . The final layer of *BERTKG-DDI model* contains concatenation of all the previous text-based outputs and drug representation from KG as expressed below:

$$o' = W_3[\text{concat}(H'_0, H'_1, H'_2, kge)] + b_3 \quad (11)$$

$$y'_t = \text{softmax}(o') \quad (12)$$

Finally the training optimization is achieved using the cross-entropy loss.

Experimental Setup

Dataset and Pre-processing

We have followed the task setting of Task 9.2 in the DDIExtraction 2013 shared task (Herrero-Zazo et al. 2013) for evaluation. It consists of MEDLINE documents annotated with the drug mentions and five types of interactions: *Mechanism*, *Effect*, *Advice*, *Interaction* and *Other*. The task is a multi-class classification to classify each of the drug pairs in the sentences into one of the types and we evaluate using three standard evaluation metrics such as: Precision (P), Recall (R) and F1-score (F1).

During pre-processing, we obtain the DRUG mentions in the corpus and map those into unique DrugBank⁴ identifiers. This is a step for converting the drug mentions into

⁴<https://go.drugbank.com/>

Embeddings on BERT-Text-DDI	Test set Macro F1
<i>bert-base-cased</i>	0.806
<i>scibert-scivocab-uncased</i>	0.812
<i>biobert v1.0 pubmed pmc</i>	0.818
<i>biobert v1.1 pubmed</i>	0.822

Table 2: Ablation of the contextual embeddings.

KG Embeddings on BERTKG-DDI	Test set Macro F1
<i>BERTKG-DDI w/ TransE</i>	0.826
<i>BERTKG-DDI w/ TransR</i>	0.829
<i>BERTKG-DDI w/ RESCAL</i>	0.834
<i>BERTKG-DDI w/ DistMult</i>	0.840

Table 3: Ablation of the influence of KG embeddings.

Models	Contextual Embeddings	Macro F1
<i>BERT-Text-DDI</i>	<i>biobert v1.0 pubmed pmc</i>	0.818
<i>BERTKG-DDI</i>	<i>biobert v1.0 pubmed pmc</i>	0.831
<i>BERT-Text-DDI</i>	<i>biobert v1.1 pubmed</i>	0.822
<i>BERTKG-DDI</i>	<i>biobert v1.1 pubmed</i>	0.840

Table 4: Probing deeper into influence of KG embeddings into BERT-based models for DDI Relation Classification.

their respective DrugBank ID, a step of entity linking (Mondal et al. 2019), (Leaman, Dogan, and lu 2013). This mention normalization has been performed based on the longest overlap of drug mentions in DrugBank and map the drugs to different Knowledge sources used to construct *Bio-KG*.

Training Details

For the purpose of experiments, we use the initialization of various pre-trained contextual embeddings. For instance, we use the embeddings such as *bert-base-cased*⁵, *scibert-scivocab-uncased* (Beltagy, Lo, and Cohan 2019)⁶ and domain-specific *biobert v1.0 pubmed pmc* and *biobert v1.0 pubmed*⁷ as the initialization of the transformer encoder in *BERTKG-DDI*. We uniformly keep the maximum sequence length as 300 for all the embedding ablations and trained for 5 epochs. For the KG embeddings, we use word embeddings dimensions to be 200. Stochastic Gradient Descent (SGD) was used for optimization with an initial learning rate of 0.0001 and the model is trained for 300 epochs. After training the embeddings, we obtain the final representation of each drug. For the drugs mentioned in the input instance, we make use of the obtained embeddings as shown in the equation 11. We initialize the non-normalized drugs using pre-trained word2vec (of dimension 200 same as the KG embedding) trained on PubMed⁸.

⁵<https://huggingface.co/bert-base-cased>

⁶<https://github.com/allenai/scibert>

⁷<https://github.com/dmis-lab/biobert>

⁸<http://evexdb.org/pmresources/ngrams/PubMed/>

Methods	Advice F1 Score	Effect F1 Score	Mechanism F1 Score	Interaction F1 Score	Total F1 Score
(Zhang et al. 2017)	0.80	0.71	0.74	0.54	0.72
(Vivian et al. 2017)	0.85	0.76	0.77	0.57	0.77
(Asada, Miwa, and Sasaki 2018)	0.81	0.71	0.73	0.45	0.72
(Sun et al. 2019)	0.80	0.73	0.78	0.58	0.75
(Zhu et al. 2020)	0.86	0.80	0.84	0.56	0.80
Our method (BERTKG-DDI)	0.88	0.81	0.87	0.59	0.84

Table 5: Comparison of F1 scores across relation types using baselines on the test set of the DDI Benchmark Corpus, containing five types of relations between drug pairs (*Advice, Mechanism, Effect, Interaction, Others*).

Results and Discussion

In this section, we provide a detailed analysis of the various results and findings that we have observed during experiments. We show empirical results based on *BERTKG-DDI* for both text and KG information.

Ablation of Embeddings on BERT-Text-DDI: During ablation analysis, we observe that the incorporation of domain-specific information in *biobert v.1 pubmed* boosts up the predictive performance in terms of macro-F1 score (across all relation types) by 2.3% compared to *bert-base-cased*. Moreover, the *scibert-vocab-cased* embeddings due to the scientific details obtained during fine-tuning achieves reasonable boost in performance. *biobert v.1 pubmed based BERT-Text-DDI* is the best-performing text-based relation classification model. The results are enumerated in Table 2.

Ablation analysis of KG Embeddings on BERTKG-DDI: We compare the different KG embeddings for drugs obtained from *Bio-KG* after augmenting with the *BERT-Text-DDI* model in Table 3. The semantic-matching models such as *RESCAL* and *DistMult* measure plausibility of facts by matching the latent semantics of both relations and entities in their vector space. In our experiments, they seem to outperform the translation-based KGE such as *TransE* and *TransR* by an average of 1% macro F1-score.

Advantage of KG information on BERTKG-DDI: During empirical analysis of the *BERTKG-DDI* model, we observe how much performance gain can be achieved by augmenting KG embeddings. From the results enumerated in terms of macro F1-score on all the relation types in Table 4, we observe that the best-performing *BERT-Text-DDI* model achieves a performance boost of 1.8% after augmenting KG information in *BERTKG-DDI*.

Comparison with the existing baselines: We compare our best-performing model with some of the best-performing existing baselines. (Asada, Miwa, and Sasaki 2018) proposed a novel neural method to extract drug-drug interactions (DDIs) from texts using external drug molecular structure information. They encode textual drug pairs with convolutional neural networks and their molecular pairs with graph convolutional networks (GCNs), and then concatenate the outputs of these two networks. (Vivian et al. 2017) proposed an effective model that classifies DDIs from the lit-

erature by combining an attention mechanism and a recurrent neural network with long short-term memory (LSTM) units. (Zhang et al. 2017) has presented a hierarchical recurrent neural networks (RNNs)-based method to integrate the SDP and sentence sequence for DDI extraction task. (Sun et al. 2019) has proposed a novel recurrent hybrid convolutional neural network (RHCNN) for DDI extraction from biomedical literature. In the embedding layer, the texts mentioning two entities are represented as a sequence of semantic embeddings and position embeddings. In particular, the complete semantic embedding is obtained by the information fusion between a word embedding and its contextual information which is learnt by recurrent structure. Recently, (Zhu et al. 2020) proposed multiple entity-aware attentions with various entity information to strengthen the representations of drug entities in sentences. They integrate drug descriptions from Wikipedia and DrugBank to our model to enhance the semantic information of drug entities. Also, they modified the output of the BioBERT model and the results show that it is better than using the BioBERT model directly. On the contrary, our method achieves the state-of-the-art performance based on the results on the DDI Extraction 2013 corpus (in terms of F1-scores of all the relation types) as shown in Table 5.

Conclusion

In this paper, we propose an approach, *BERTKG-DDI*, for DDI relation classification based on pre-trained language models and Knowledge Graph Embedding of the drug entities. Experiments conducted on a benchmark DDI dataset proves the effectiveness of our proposed method. Possible directions of further research might be to explore other external drug representation such as chemical structure, textual description in predicting DDI from textual corpus.

References

- Asada, M.; Miwa, M.; and Sasaki, Y. 2018. Enhancing Drug-Drug Interaction Extraction from Texts by Molecular Structure Information. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 680–685. Melbourne, Australia: Association for Computational Linguistics. doi:10.18653/v1/P18-2108. URL <https://www.aclweb.org/anthology/P18-2108>.
- Beltagy, I.; Lo, K.; and Cohan, A. 2019. SciBERT:

- A Pretrained Language Model for Scientific Text. In *EMNLP/IJCNLP*.
- Bordes, A.; Usunier, N.; García-Durán, A.; Weston, J.; and Yakhnenko, O. 2013. Translating Embeddings for Modeling Multi-relational Data. In *NIPS*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT*.
- Herrero-Zazo, M.; Segura-Bedmar, I.; Martínez, P.; and De-clerck, T. 2013. The DDI corpus: An annotated corpus with pharmacological substances and drug-drug interactions. *Journal of Biomedical Informatics* 46(5): 914 – 920. ISSN 1532-0464. doi:<https://doi.org/10.1016/j.jbi.2013.07.011>. URL <http://www.sciencedirect.com/science/article/pii/S1532046413001123>.
- Leaman, R.; Dogan, R.; and lu, Z. 2013. DNorm: Disease Name Normalization with Pairwise Learning to Rank. *Bioinformatics (Oxford, England)* 29. doi:10.1093/bioinformatics/btt474.
- Lee, J.; Yoon, W.; Kim, S.; Kim, D.; Kim, S.; So, C. H.; and Kang, J. 2019. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics* 36(4): 1234–1240. ISSN 1367-4803. doi:10.1093/bioinformatics/btz682. URL <https://doi.org/10.1093/bioinformatics/btz682>.
- Li, D.; and Ji, H. 2019. Syntax-aware Multi-task Graph Convolutional Networks for Biomedical Relation Extraction. In *Proceedings of the Tenth International Workshop on Health Text Mining and Information Analysis (LOUHI 2019)*, 28–33. Hong Kong: Association for Computational Linguistics. doi:10.18653/v1/D19-6204. URL <https://www.aclweb.org/anthology/D19-6204>.
- Lin, Y.; Liu, Z.; Sun, M.; Liu, Y.; and Zhu, X. 2015. Learning Entity and Relation Embeddings for Knowledge Graph Completion. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, AAAI'15*, 2181–2187. AAAI Press. ISBN 0262511290.
- Liu, S.; Tang, B.; Chen, Q.; and Wang, X. 2016. Drug-Drug Interaction Extraction via Convolutional Neural Networks. *Computational and Mathematical Methods in Medicine* 2016: 1–8. doi:10.1155/2016/6918381.
- Mondal, I. 2020. BERTChem-DDI : Improved Drug-Drug Interaction Prediction from text using Chemical Structure Information. In *Proceedings of Knowledgeable NLP: the First Workshop on Integrating Structured Knowledge and Neural Networks for NLP*, 27–32. Suzhou, China: Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/2020.knlp-1.4>.
- Mondal, I.; Purkayastha, S.; Sarkar, S.; Goyal, P.; Pillai, J.; Bhattacharyya, A.; and Gattu, M. 2019. Medical Entity Linking using Triplet Network. In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, 95–100. Minneapolis, Minnesota, USA: Association for Computational Linguistics. doi:10.18653/v1/W19-1912. URL <https://www.aclweb.org/anthology/W19-1912>.
- Nickel, M.; Tresp, V.; and Kriegel, H.-P. 2011. A Three-Way Model for Collective Learning on Multi-Relational Data. In *Proceedings of the 28th International Conference on International Conference on Machine Learning, ICML'11*, 809–816. Madison, WI, USA: Omnipress. ISBN 9781450306195.
- Purkayastha, S.; Mondal, I.; Sarkar, S.; Goyal, P.; and Pillai, J. K. 2019. Drug-Drug Interactions Prediction Based on Drug Embedding and Graph Auto-Encoder. In *2019 IEEE 19th International Conference on Bioinformatics and Bio-engineering (BIBE)*, 547–552.
- Sahu, S. K.; and Anand, A. 2018. Drug-drug interaction extraction from biomedical texts using long short-term memory network. *Journal of Biomedical Informatics* 86: 15 – 24. ISSN 1532-0464. doi:<https://doi.org/10.1016/j.jbi.2018.08.005>. URL <http://www.sciencedirect.com/science/article/pii/S1532046418301606>.
- Sun, X.; Dong, K.; Ma, L.; Sutcliffe, R.; He, F.; Chen, S.; and Feng, J. 2019. Drug-Drug Interaction Extraction via Recurrent Hybrid Convolutional Neural Networks with an Improved Focal Loss. *Entropy* 21(1): 37. ISSN 1099-4300. doi:10.3390/e21010037. URL <http://dx.doi.org/10.3390/e21010037>.
- The UniProt Consortium. 2016. UniProt: the universal protein knowledgebase. *Nucleic Acids Research* 45(D1): D158–D169. ISSN 0305-1048. doi:10.1093/nar/gkw1099. URL <https://doi.org/10.1093/nar/gkw1099>.
- Vivian, V.; Lin, H.; Luo, L.; Zhao, Z.; Zhengguang, I.; Yijia, Z.; Yang, Z.; and Wang, J. 2017. An attention-based effective neural model for drug-drug interactions extraction. *BMC Bioinformatics* 18. doi:10.1186/s12859-017-1855-x.
- Wu, S.; and He, Y. 2019. Enriching Pre-trained Language Model with Entity Information for Relation Classification. *CoRR* abs/1905.08284. URL <http://arxiv.org/abs/1905.08284>.
- Yang, B.; tau Yih, W.; He, X.; Gao, J.; and Deng, L. 2015. Embedding Entities and Relations for Learning and Inference in Knowledge Bases. *CoRR* abs/1412.6575.
- Yang, Z.; Dai, Z.; Yang, Y.; Carbonell, J.; Salakhutdinov, R.; and Le, Q. V. 2019. XLNet: Generalized Autoregressive Pretraining for Language Understanding. In *NeurIPS*.
- Zhang, Y.; Zheng, W.; Lin, H.; Wang, J.; Yang, Z.; and Dumontier, M. 2017. Drug-drug interaction extraction via hierarchical RNNs on sequence and shortest dependency paths. *Bioinformatics* 34(5): 828–835. ISSN 1367-4803. doi:10.1093/bioinformatics/btx659. URL <https://doi.org/10.1093/bioinformatics/btx659>.
- Zhu, Y.; Li, L.; Lu, H.; Zhou, A.; and Qin, X. 2020. Extracting drug-drug interactions from texts with BioBERT and multiple entity-aware attentions. *Journal of Biomedical Informatics* 106: 103451. ISSN 1532-0464. doi:<https://doi.org/10.1016/j.jbi.2020.103451>. URL <http://www.sciencedirect.com/science/article/pii/S1532046420300794>.