



BiOnt: Deep Learning Using Multiple Biomedical Ontologies for Relation Extraction

Diana Sousa^(✉)  and Francisco M. Couto 

LASIGE, Departamento de Informática, Faculdade de Ciências,
Universidade de Lisboa, 1749–016 Lisbon, Portugal
dfsousa@lasige.di.fc.ul.pt

Abstract. Successful biomedical relation extraction can provide evidence to researchers and clinicians about possible unknown associations between biomedical entities, advancing the current knowledge we have about those entities and their inherent mechanisms. Most biomedical relation extraction systems do not resort to external sources of knowledge, such as domain-specific ontologies. However, using deep learning methods, along with biomedical ontologies, has been recently shown to effectively advance the biomedical relation extraction field. To perform relation extraction, our deep learning system, BiOnt, employs four types of biomedical ontologies, namely, the Gene Ontology, the Human Phenotype Ontology, the Human Disease Ontology, and the Chemical Entities of Biological Interest, regarding gene-products, phenotypes, diseases, and chemical compounds, respectively. We tested our system with three data sets that represent three different types of relations of biomedical entities. BiOnt achieved, in F-score, an improvement of 4.93% points for drug-drug interactions (DDI corpus), 4.99% points for phenotype-gene relations (PGR corpus), and 2.21% points for chemical-induced disease relations (BC5CDR corpus), relatively to the state-of-the-art. The code supporting this system is available at <https://github.com/lasigeBioTM/BiONT>.

Keywords: Relation extraction · Biomedical ontologies · Deep learning · Text mining

1 Introduction

The description of the mechanisms that are responsible for the behavior of biological systems is non-trivial, and each step towards the understanding of those mechanisms often constitutes a scientific achievement [2, 26]. Typical examples describe diseases that are associated with mechanisms that originate phenotypic

This work was supported by FCT through project DeST: Deep Semantic Tagger, ref. PTDC/CCI-BIO/28685/2017, LASIGE Research Unit, ref. UIDP/00408/2020, and PhD Scholarship, ref. SFRH/BD/145221/2019.

abnormalities as a result of modified gene expression, as well as the action of drugs on those diseases [4], among others. One significant step to fully understand biological systems mechanisms is to extract and classify the relations that exist between the different biomedical entities, namely chemicals, diseases, genes, and phenotypes. In literature, authors classify this problem as a Relation Extraction (RE) task. Biomedical RE aims to extract and classify relations between biomedical entities in highly heterogeneous or unstructured scientific or clinical text.

Deep learning is widely used to solve problems such as speech recognition, visual object recognition, and object detection. Lately, deep learning based-systems have started to tackle RE problems. These systems are becoming increasingly more complex, namely the MIMLCNN [10], and PCNN + Att [17] systems, that mark recent turning points in the deep learning RE field. Both of these systems use Word2Vec [19] that aims to capture the syntactic and semantic information about the word [11]. However, deep learning methods that effectively extract and classify relations between biomedical entities in the text are still scarce [13, 15].

Ontologies play an important role in biomedical research through a variety of applications and are used primarily as a source of vocabulary for standardization and integration purposes [3]. Word embeddings can learn how to detect relations between entities but manifest difficulties in grasping the semantics of each entity and their specific domain. Domain-specific ontologies provide and formalize this knowledge. Thus, a structured representation of the semantics between entities and their relations, an ontology, allows us to use it as an added feature to a machine learning classifier. Some of the biomedical entities structured in publicly available ontologies are genes properties/attributes (Gene Ontology (**GO**)) (45003 terms) [1, 23], phenotypes (Human Phenotype Ontology (**HPO**)) (25810 terms) [12], diseases (Human Disease Ontology (**DO**)) (18114 terms) [21], and drugs/chemicals (Chemical Entities of Biological Interest (**ChEBI**)) (133104 terms) [8]¹.

This work presents the **BiOnt** system, a biomedical RE system built using bidirectional Long Short-Term Memory (LSTM) networks. The BiOnt system incorporates the state-of-the-art Word2Vec word embeddings [19] and makes use of different combinations of input channels to maximize performance. Our system is based on the work of Lamurias et al. [13] and Xu et al. [25]. Both of these models make use of biomedical resources as embedding layers for their respective systems. Lamurias et al. [13] uses the Xu et al. [25] model has a baseline with an added ontological embedding layer (BO-LSTM model). However, the BO-LSTM model is limited to two types of relations, namely, drug-drug, and gene-phenotype relations.

External sources of knowledge, such as biomedical ontologies, can provide highly valuable information for the detection of relations between entities in the text, as described previously by Lamurias et al. [13]. These knowledge-bases provide not only relevant characteristics about the respective entities but also

¹ Term counts at 09/09/2019.

about the underlying semantics of the relations they establish. This information is not expressed directly in the training data but usually reinforces a relation between two entities in the text. The novelty of our system is that expands the previous work done by Lamurias et al. [13] by using four types of domain-specific ontologies, and combine them to extract new types of relations, along with word embeddings [19] and WordNet hypernyms [5]. BiOnt successfully replicates the results of the BO-LSTM application, using different types of ontologies. Our system can extract new relations between four biomedical entities, namely, genes, phenotypes, diseases, and chemicals. Figure 1 shows how these four types of biomedical ontologies can be combined to aid the relation extraction of ten different combinations of biomedical entities. The BiOnt system also explores the use of entities that are not direct entries in an ontology (e.g., genes), linking each entity to their most informative annotation concept within a corresponding ontology (e.g., GO). Our method incorporates more ontologies than the previously mentioned systems and is evaluated using three state-of-the-art data sets. The BiOnt system can be used to effectively populate knowledge bases regarding gold standard relations. Ultimately, it can be used to explore new experimental hypotheses providing evidence to researchers and clinicians about possible unknown associations between biomedical entities.

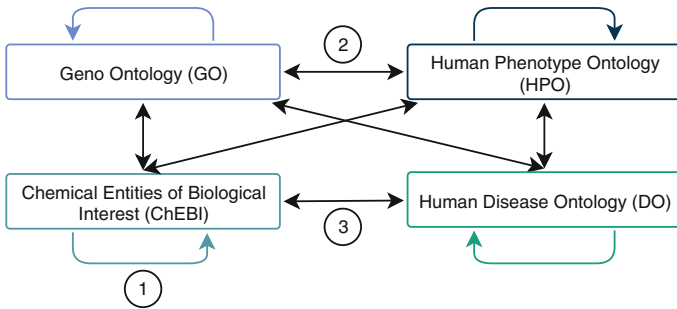


Fig. 1. The ten possible combinations between the four biomedical ontologies (the Gene Ontology (GO) [1], the Human Phenotype Ontology (HPO) [12], the Human Disease Ontology (DO) [21], and the Chemical Entities of Biological Interest (ChEBI) [8]). The **1** represents the DDI corpus, the **2** the PGR corpus, and the **3** the BC5CDR corpus (described in Sect. 3).

2 Methodology

This section describes the BiOnt model with an emphasis on the enhancements done to BO-LSTM [13] model to allow multi-ontology integration, expanding the number of different type candidate pairs from two to ten. The BiOnt model uses a combination of different language and entity related data representations, that feed individual channels creating a multichannel architecture. The input data is used to generate instances to be classified by the model. Each instance

corresponds to a candidate pair of entities in a sentence. To each instance, the model assigns a positive or negative class. A positive class corresponds to an identified relation between two biomedical concepts, where the nature of this relation depends of the data set being used to perform the evaluation, and a negative class implies no relation between the different entities.

An instance should condense all relevant information to classify a candidate pair. To create an instance the BiOnt model relies on three primary data information layers. After sentence tokenization, these layers are: Shortest Dependency Path (SDP) [18,20], WordNet Classes [5], and **Ontology Embeddings**. The latter represents the relations between the ancestors for each ontology concept corresponding to an entity (Fig. 2). The model assumes that the input data already has the offsets of the relevant entities identified and their respective concept ID, the Named-Entity Recognition and Linking tasks. However, while most entities already corresponded to an ontology concept ID, some entities, such as genes do not have a direct entry in an ontology. BiOnt matches these entities to their most representative concept in the Gene Ontology [1]. To match the gene to their most representative GO term the priority was given to concepts inferred from experiments, for having a more sustained background and usually be more descriptive. For tie-breaking, if we have several GO terms inferred from experiments, the choice is the term that is the most specific (i.e., with the longer ancestry line).




Fig. 2. BiOnt ontology embedding illustration based on the HPO and the GO ontologies, for the candidate relation between the human phenotype *blindness* and the gene *CRB1* (represented by the GO term *heterophilic cell-cell adhesion via plasma membrane cell adhesion molecules*).

As stated previously, our system expands the work done by Lamurias et al. [13] by using four types of domain-specific ontologies, and combine them to extract new types of relations. Therefore, to allow this diversity of relations, we adapted the BO-LSTM model common ancestors and the concatenation of ancestors channels. Since the common ancestors' channel could only be used for relations between the same type of biomedical entities, we only use the concatenation of ancestors channel for the relations between different biomedical entities.

3 Evaluation

To showcase our systems' performance, we used three different state-of-the-art data sets. These data sets represent three out of the ten possible combinations of the biomedical entities used in this work, drug-drug interactions, phenotype-gene relations, and chemical-induced disease relations. With these data sets, we intend to show the flexibility of our model to the different types of biomedical entities represented by biomedical ontologies. Figure 1 illustrates how the entities present in the three data sets (1, 2, and 3) are connected to the different biomedical entities.

Drug-Drug Interactions (1). The SemEval 2013: Task 9 DDI Extraction Corpus [9] is a corpus that describes drug-drug interactions (DDIs) focused on both pharmacokinetic (PK) and pharmacodynamic (PD) DDIs. The manually annotated corpus created by Herrero et al. [9] combines 5028 DDIs, from selected texts of the DrugBank database and Medline abstracts.

Phenotype-Gene Relations (2). The Phenotype-Gene Relations Corpus (PGR) [22] is a corpus that describes human phenotype-gene relations, created in a fully automated manner. Due to being a silver standard corpus is not expected to be as reliable as manually annotated corpora. Nonetheless, the authors show the system efficiency by training two state-of-the-art relation extraction deep learning systems. The PGR corpus combines 4283 human phenotype-gene relations.

Chemical-Induced Disease Relations (3). The BioCreative V CDR Corpus (BC5CDR) [16] is a corpus of chemical-induced disease (CID) relations. The BC5CDR corpus consists of 3116 chemical-disease interactions annotated from PubMed articles. To use the BC5CDR corpus, we had to preprocess the documents linking the annotations of the relations to their sentences. We assumed that if two entities share a relation in the document, they will continue to share that relation if present in the same sentence of that document.

4 Results and Discussion

Table 1 presents the relation extraction results of our system, BiOnt, for each data set. For all three data sets, our system performs better using the ontological embeddings layer (+ Ontologies), than just using the word embeddings and WordNet classes layers (State-of-the-art), by an average of 0.0404. The most relevant contribution for this metric was an increase in recall for the DDI and PGR corpus, and in precision for the BC5CDR corpus. The ontology embeddings contribute to the identification of more correct relations, with a small trade-off in precision, for the DDI corpus. For the other two data sets, the ontological embedding layer does not damage the precision, while more correct relations are identified.

Table 1. Relation extraction results with the BiOnt system, for each data set, expressing drug-drug interactions (DDI Corpus), phenotype-gene relations (PGR Corpus), and chemical-induced disease relations (BC5CDR Corpus).

Data Set	Configuration	Precision	Recall	F-score
DDI Corpus	State-of-the-art	0.7134	0.6410	0.6753
	+ Ontologies	0.6784	0.7775	0.7246
PGR Corpus	State-of-the-art	0.8421	0.6666	0.7442
	+ Ontologies	0.8438	0.7500	0.7941
BC5CDR Corpus	State-of-the-art	0.5371	0.7264	0.6175
	+ Ontologies	0.5770	0.7173	0.6396

For the DDI corpus, the BiOnt system, due to the inherent variability of the preprocessing phase (by randomizing the division between training and test sets), when comparing with the BO-LSTM system, performed slightly worse (0.7246 in F-score) than the previously reported results (0.7290 in F-score). The paper supporting the PGR corpus [22] reported some deep learning applications results, including with the BERT [7] based BioBERT [14] pre-trained biomedical language representation model (0.6716 in F-score). Our system outperformed those results with an F-score of 0.7941. Regarding the BC5CDR corpus, our system outperformed the best system (0.5703 in F-score) in the challenge task chemical-induced disease (CID) relation extraction of BioCreative V, by 0.0693 [24], with 0.6396 in F-score. The differences in F-score, for the distinct data sets, are mostly due to how they were built, and the completeness and complexity of the respective ontologies. For instance, the PGR corpus is a silver standard corpus, therefore, could have entities that were poorly identified, not identified at all, or not linked to the right identifier. The BC5CDR corpus was annotated for documents, not regarding the offsets of the entities that shared a relation in each document, which is also a possible limitation.

5 Conclusions and Future Work

This work showed that the knowledge encoded in biomedical ontologies plays a vital part in the development of learning systems, providing semantic and ancestry information for entities, such as genes, phenotypes, chemicals, and diseases. We evaluated BiOnt using three state-of-the-art data sets (DDI, PGR, and BC5CDR corpus), obtaining improvements in F-score (4.93, 4.99, and 2.21% points, respectively), by using an ontological information layer. Our system successfully enhances the results of Lamurias et al. [13] to other entities and ontologies. BiOnt shows that integrating biomedical ontologies instead of relying solely on the training data for creating classification models will allow us not only to find relevant information for a particular problem quicker but possibly also to find unknown associations between biomedical entities.

Regarding future work, it is possible to integrate more ontological information, and in different ways. For instance, one could consider only the relations between the ancestors with the highest information content (more relevant for the candidate pair they characterize). The information content could be inferred from the probability of each term in each ontology or resorting to an external data set. Also, a semantic similarity measurement could account for non-transitive relations (within the same ontology). Relatively to biomedical concepts that do not constitute ontology entries, we could explore quantitative evidence values, choose more than one representative term, and we could also employ semantic similarity measures [6].

References

1. Ashburner, M., et al.: Gene ontology: tool for the unification of biology. The gene ontology consortium. *Nature Genet.* **25**(1), 25–29 (2000)
2. Bechtel, W.: Biological mechanisms: organized to maintain autonomy. In: Boogerd, F.C., Bruggeman, F.J., Hofmeyr, J.H.S., Westerhoff, H.V. (eds.) *Systems Biology*, pp. 269–302. Elsevier, Amsterdam (2007)
3. Bodenreider, O.: Biomedical ontologies in action: role in knowledge management, data integration and decision support. In: *IMIA Yearbook Medical Informatics*, pp. 67–79 (2008)
4. Campaner, R.: Understanding mechanisms in the health sciences. *Theor. Med. Bioeth.* **32**(1), 5–17 (2011). <https://doi.org/10.1007/s11017-010-9166-5>
5. Ciaramita, M., Altun, Y.: Broad-coverage sense disambiguation and information extraction with a supersense sequence tagger. In: *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing, EMNLP 2006*, pp. 594–602. Association for Computational Linguistics, Stroudsburg (2006)
6. Couto, F.M., Lamurias, A.: Semantic similarity definition. In: Ranganathan, S., Gribskov, M., Nakai, K., Schönbach, C. (eds.) *Encyclopedia of Bioinformatics and Computational Biology*, pp. 870–876. Academic Press, Oxford (2019)
7. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186. Association for Computational Linguistics, Minneapolis (2019)
8. Hastings, J., et al.: ChEBI in 2016: improved services and an expanding collection of metabolites. *Nucleic Acids Res.* **44**(D1), D1214–D1219 (2015)
9. Herrero-Zazo, M., Segura-Bedmar, I., Martínez, P., Declerck, T.: The DDI corpus: an annotated corpus with pharmacological substances and drug-drug interactions. *J. Biomed. Inform.* **46**(5), 914–920 (2013)
10. Jiang, X., Wang, Q., Li, P., Wang, B.: Relation extraction with multi-instance multi-label convolutional neural networks. In: *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pp. 1471–1480. The COLING 2016 Organizing Committee, Osaka (2016)
11. Kumar, S.: A survey of deep learning methods for relation extraction. *CoRR* abs/1705.03645 (2017)
12. Köhler, S., Vasilevsky, N., Engelstad, M., Foster, E., et al.: The human phenotype ontology. *Nucleic Acids Res.* **45**, D865–D876 (2017)

13. Lamurias, A., Sousa, D., Clarke, L.A., Couto, F.M.: BO-LSTM: classifying relations via long short-term memory networks along biomedical ontologies. *BMC Bioinform.* **20**(1), 10 (2019)
14. Lee, J., et al.: BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *arXiv e-prints preprint arXiv:1901.08746* (2019)
15. Li, F., Zhang, M., Fu, G., Ji, D.H.: A neural joint model for entity and relation extraction from biomedical text. *BMC Bioinform.* **18**, 198 (2017)
16. Li, J., et al.: BioCreative V CDR task corpus: a resource for chemical disease relation extraction. *Database* **2016**, 1–10 (2016)
17. Lin, Y., Shen, S., Liu, Z., Luan, H., Sun, M.: Neural relation extraction with selective attention over instances. In: *ACL* (2016)
18. Mikolov, T., Sutskever, I., Chen, K., Corrado, G., Dean, J.: Distributed representations of words and phrases and their compositionality. In: *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2, NIPS 2013*, pp. 3111–3119. Curran Associates Inc., USA (2013)
19. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: Burges, C.J.C., Bottou, L., Welling, M., Ghahramani, Z., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems 26*, pp. 3111–3119. Curran Associates Inc., New York (2013)
20. Pyysalo, S., Ginter, F., Moen, H., Salakoski, T., Ananiadou, S.: Distributional semantics resources for biomedical text processing. In: *Proceedings of LBM 2013*, pp. 39–44 (2013)
21. Schriml, L.M., et al.: Human Disease Ontology 2018 update: classification, content and workflow expansion. *Nucleic Acids Res.* **47**(D1), D955–D962 (2018)
22. Sousa, D., Lamurias, A., Couto, F.M.: A silver standard corpus of human phenotype-gene relations. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 1487–1492. Association for Computational Linguistics, Minneapolis (2019)
23. The Gene Ontology Consortium: The Gene Ontology Resource: 20 years and still GOing strong. *Nucleic Acids Res.* **47**(D1), D330–D338 (2018)
24. Wei, C.H., et al.: Overview of the BioCreative V chemical disease relation (CDR) task. In: *Proceedings of the Fifth BioCreative Challenge Evaluation Workshop*, vol. 14 (2015)
25. Xu, B., Shi, X., Zhao, Z., Zheng, W.: Leveraging biomedical resources in Bi-LSTM for drug-drug interaction extraction. *IEEE Access* **6**, 33432–33439 (2018)
26. Yu, A.C.: Methods in biomedical ontology. *J. Biomed. Inform.* **39**(3), 252–266 (2006)