

A Data-driven Approach for Noise Reduction in Distantly Supervised Biomedical Relation Extraction

Saadullah Amin* Katherine Ann Dunfield* Anna Vechkaeva Günter Neumann

German Research Center for Artificial Intelligence (DFKI)
Multilinguality and Language Technology Lab

{saadullah.amin, katherine.dunfield, anna.vechkaeva, guenter.neumann}@dfki.de

Abstract

Fact triples are a common form of structured knowledge used within the biomedical domain. As the amount of unstructured scientific texts continues to grow, manual annotation of these texts for the task of relation extraction becomes increasingly expensive. Distant supervision offers a viable approach to combat this by quickly producing large amounts of labeled, but considerably noisy, data. We aim to reduce such noise by extending an entity-enriched relation classification BERT model to the problem of multiple instance learning, and defining a simple data encoding scheme that *significantly* reduces noise, reaching state-of-the-art performance for distantly-supervised biomedical relation extraction. Our approach further encodes knowledge about the direction of relation triples, allowing for increased focus on relation learning by reducing noise and alleviating the need for joint learning with knowledge graph completion.

1 Introduction

Relation extraction (RE) remains an important natural language processing task for understanding the interaction between entities that appear in texts. In supervised settings (GuoDong et al., 2005; Zeng et al., 2014; Wang et al., 2016), obtaining fine-grained relations for the biomedical domain is challenging due to not only the annotation costs, but the added requirement of domain expertise. Distant supervision (DS), however, provides a meaningful way to obtain large-scale data for RE (Mintz et al., 2009; Hoffmann et al., 2011), but this form of data collection also tends to result in an increased amount of noise, as the target relation may not always be expressed (Takamatsu et al., 2012; Ritter et al., 2013). Exemplified in Figure 1, the last two sentences can be seen as potentially *noisy* evidence, as they do not explicitly express the given relation.

* Equal contribution.

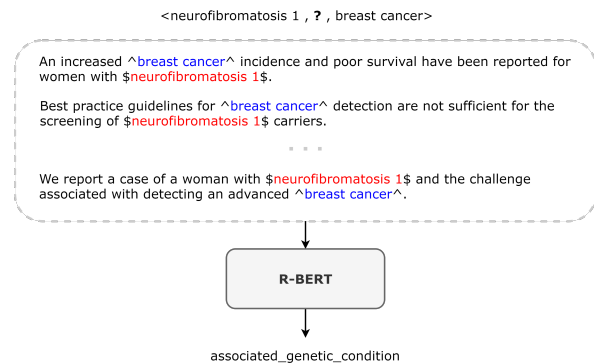


Figure 1: Example of a distantly supervised bag of sentences for a knowledge base tuple (*neurofibromatosis 1, breast cancer*) with special order sensitive entity markers to capture the position and the latent relation direction with BERT for predicting the missing relation.

Since individual instance labels may be unknown (Wang et al., 2018), we instead build on the recent findings of Wu and He (2019) and Soares et al. (2019) in using positional markings and latent relation direction (Figure 1), as a signal to mitigate noise in *bag-level* multiple instance learning (MIL) for distantly supervised biomedical RE. Our approach greatly simplifies previous work by Dai et al. (2019) without the need for additional tasks. Our contributions follow:

- We extend *sentence-level* relation enriched BERT (Wu and He, 2019) to *bag-level* MIL.
- We demonstrate that the simple applications of this model under-perform and require knowledge base order-sensitive markings, *k-tag*, to achieve state-of-the-art performance. This data encoding scheme captures the latent relation direction and provides a simple way to reduce noise in distant supervision.
- We make our code and data creation pipeline publicly available: <https://github.com/suamin/umls-medline-distant-re>

2 Related Work

In MIL-based distant supervision for *corpus-level* RE, earlier works rely on the assumption that at least one of the evidence samples represent the target relation in a triple (Riedel et al., 2010; Hoffmann et al., 2011; Surdeanu et al., 2012). Recently, piecewise convolutional neural networks (PCNN) (Zeng et al., 2014) have been applied to DS (Zeng et al., 2015), with notable extensions in selective attention (Lin et al., 2016) and the modelling of noise dynamics (Luo et al., 2017). Han et al. (2018a) proposed a joint learning framework for knowledge graph completion (KGC) and RE with mutual attention, showing that DS improves downstream KGC performance, while KGC acts as an indirect signal to filter textual noise. Dai et al. (2019) extended this framework to biomedical RE, using improved KGC models, ComplEx (Trouillon et al., 2017) and Simple (Kazemi and Poole, 2018), as well as additional auxiliary tasks of entity-type classification and named entity recognition to mitigate noise.

Pre-trained language models, such as BERT (Devlin et al., 2019), have been shown to improve the downstream performance of many NLP tasks. Relevant to distant RE, Alt et al. (2019) extended the OpenAI Generative Pre-trained Transformer (GPT) model (Radford et al., 2019) for *bag-level* MIL with selective attention (Lin et al., 2016). Sun et al. (2019) enriched pre-training stage with KB entity information, resulting in improved performance. For *sentence-level* RE, Wu and He (2019) proposed an entity marking strategy for BERT (referred to here as R-BERT) to perform relation classification. Specifically, they mark the entity boundaries with special tokens following the order they appear in the sentence. Likewise, Soares et al. (2019) studied several data encoding schemes and found marking entity boundaries important for *sentence-level* RE. With such encoding, they further proposed a novel pre-training scheme for distributed relational learning, suited to few-shot relation classification (Han et al., 2018b).

Our work builds on these findings, in particular, we extend the BERT model (Devlin et al., 2019) for *bag-level* MIL, similar to Alt et al. (2019). More importantly, noting the significance of *sentence-ordered* entity marking in *sentence-level* RE (Wu and He, 2019; Soares et al., 2019), we introduce the *knowledge-based* entity marking strategy suited to *bag-level* DS. This naturally encodes the information stored in KB, reducing the inherent noise.

3 Bag-level MIL for Distant RE

3.1 Problem Definition

Let \mathcal{E} and \mathcal{R} represent the set of entities and relations from a knowledge base \mathcal{KB} , respectively. For $h, t \in \mathcal{E}$ and $r \in \mathcal{R}$, let $(h, r, t) \in \mathcal{KB}$ be a fact triple for an ordered tuple (h, t) . We denote all such (h, t) tuples by a set \mathcal{G}^+ , i.e., there exists some $r \in \mathcal{R}$ for which the triple (h, r, t) belongs to the KB, called *positive groups*. Similarly, we denote by \mathcal{G}^- the set of *negative groups*, i.e., for all $r \in \mathcal{R}$, the triple (h, r, t) does not belong to KB. The union of these groups is represented by $\mathcal{G} = \mathcal{G}^+ \cup \mathcal{G}^-$. We denote by $\mathcal{B}_g = [s_g^{(1)}, \dots, s_g^{(m)}]$ an unordered sequence of sentences, called *bag*, for $g \in \mathcal{G}$ such that the sentences contain the group $g = (h, t)$, where the bag size m can vary. Let f be a function that maps each element in the bag to a low-dimensional relation representation $[\mathbf{r}_g^{(1)}, \dots, \mathbf{r}_g^{(m)}]$. With o , we represent the bag aggregation function, that maps instance level relation representation to a final bag representation $\mathbf{b}_g = o(f(\mathcal{B}_g))$. The goal of distantly supervised *bag-level* MIL for *corpus-level* RE is then to predict the missing relation r given the bag.

3.2 Entity Markers

Wu and He (2019) and Soares et al. (2019) showed that using special markers for entities with BERT in the order they appear in a sentence encodes the positional information that improves the performance of *sentence-level* RE. It allows the model to focus on target entities when, possibly, other entities are also present in the sentence, implicitly doing entity disambiguation and reducing noise. In contrast, for *bag-level* distant supervision, the noisy channel be attributed to several factors for a given triple (h, r, t) and bag \mathcal{B}_g :

1. Evidence sentences may not express the relation.
2. Multiple entities appearing in the sentence, requiring the model to disambiguate target entities among other.
3. The direction of missing relation.
4. Discrepancy between the order of the target entities in the sentence and knowledge base.

To address (1), common approaches are to learn a negative relation class NA and use better bag aggregation strategies (Lin et al., 2016; Luo et al.,

¹The sets are disjoint, $\mathcal{G}^+ \cap \mathcal{G}^- = \emptyset$

2017; Alt et al., 2019). For (2), encoding positional information is important, such as, in PCNN (Zeng et al., 2014), that takes into account the relative positions of *head* and *tail* entities (Zeng et al., 2015), and in (Wu and He, 2019; Soares et al., 2019) for *sentence-level* RE. To account for (3) and (4), multi-task learning with KGC and mutual attention has proved effective (Han et al., 2018a; Dai et al., 2019). Simply extending sentence sensitive marking to *bag-level* can be adverse, as it enhances (4) and even if the composition is uniform, it distributes the evidence sentence across several bags. On the other hand, expanding relations to multiple sub-classes based on direction (Wu and He, 2019), enhances class imbalance and also distributes supporting sentences. To jointly address (2), (3) and (4), we introduce KB sensitive encoding suitable for *bag-level* distant RE.

Formally, for a group $g = (h, t)$ and a matching sentence $s_g^{(i)}$ with tokens $(x_0, \dots, x_L)^2$, we add special tokens \$ and ^ to mark the entity spans as: **Sentence ordered:** Called *s-tag*, entities are marked in the order they appear in the sentence. Following Soares et al. (2019), let $s_1 = (i, j)$ and $s_2 = (k, l)$ be the index pairs with $0 < i < j - 1, j < k, k \leq l - 1$ and $l \leq L$ delimiting the entity mentions $e_1 = (x_i, \dots, x_j)$ and $e_2 = (x_k, \dots, x_l)$ respectively. We mark the boundary of s_1 with \$ and s_2 with ^. Note, e_1 and e_2 can be either *h* or *t*.

KB ordered: Called *k-tag*, entities are marked in the order they appear in the KB. Let $s_h = (i, j)$ and $s_t = (k, l)$ be the index pairs delimiting head (*h*) and tail (*t*) entities, irrespective of the order they appear in the sentence. We mark the boundary of s_h with \$ and s_t with ^.

The *s-tag* annotation scheme is followed by Soares et al. (2019) and Wu and He (2019) for span identification. In Wu and He (2019), each relation type $r \in \mathcal{R}$ is further expanded to two sub-classes as $r(e_1, e_2)$ and $r(e_2, e_1)$ to capture direction, while holding the *s-tag* annotation as fixed. For DS-based RE, since the ordered tuple (h, t) is given, the task is reduced to relation classification without direction. This side information is encoded in data with *k-tag*, covering (2) but also (3) and (4). To account for (1), we also experiment with selective attention (Lin et al., 2016) which has been widely used in other works (Luo et al., 2017; Han et al., 2018a; Alt et al., 2019).

² $x_0 = [\text{CLS}]$ and $x_L = [\text{SEP}]$

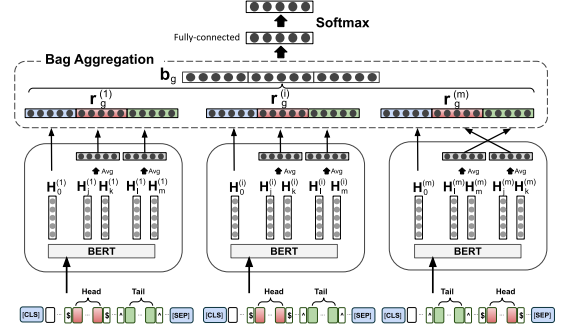


Figure 2: Multiple instance learning (MIL) based *bag-level* relation classification BERT with KB ordered entity marking (Section 3.2). Special markers \$ and ^ always delimit the span of *head* (h_s, h_e) and *tail* (t_s, t_e) entities regardless of their order in the sentence. The markers capture the *positions* of entities and latent relation *direction*.

3.3 Model Architecture

BERT (Devlin et al., 2019) is used as our base sentence encoder, specifically, BioBERT (Lee et al., 2020), and we extend R-BERT (Wu and He, 2019) to *bag-level* MIL. Figure 2 shows the model’s architecture with *k-tag*. Consider a bag \mathcal{B}_g of size m for a group $g \in \mathcal{G}$ representing the ordered tuple (h, t) , with corresponding spans $S_g = [(s_h^{(1)}, s_t^{(1)}), \dots, (s_h^{(m)}, s_t^{(m)})]$ obtained with *k-tag*, then for a pair of sentences in the bag and spans, $(s^{(i)}, (s_h^{(i)}, s_t^{(i)}))$, we can represent the model in three steps, such that the first two steps represent the map f and the final step o , as follows:

1. **SENTENCE ENCODING:** BERT is applied to the sentence and the final hidden state $\mathbf{H}_0^{(i)} \in \mathbb{R}^d$, corresponding to the [CLS] token, is passed through a linear layer³ $\mathbf{W}^{(1)} \in \mathbb{R}^{d \times d}$ with $\tanh(\cdot)$ activation to obtain the global sentence information in $\mathbf{h}_0^{(i)}$.
2. **RELATION REPRESENTATION:** For the head entity, represented by the span $s_h^{(i)} = (j, k)$ for $k > j$, we apply average pooling $\frac{1}{k-j+1} \sum_{n=j}^k \mathbf{H}_n^{(i)}$, and similarly for the tail entity with span $s_t^{(i)} = (l, m)$ for $m > l$, we get $\frac{1}{m-l+1} \sum_{n=l}^m \mathbf{H}_n^{(i)}$. The pooled representations are then passed through a shared linear layer $\mathbf{W}^{(2)} \in \mathbb{R}^{d \times d}$ with $\tanh(\cdot)$ activation to get $\mathbf{h}_h^{(i)}$ and $\mathbf{h}_t^{(i)}$. To get the final latent relation representation, we concatenate the pooled entities representation with [CLS] as $\mathbf{r}_g^{(i)} = [\mathbf{h}_0^{(i)}; \mathbf{h}_h^{(i)}; \mathbf{h}_t^{(i)}] \in \mathbb{R}^{3d}$.

³Each linear layer is implicitly assumed with a bias vector

3. **BAG AGGREGATION:** After applying the first two steps to each sentence in the bag, we obtain $[\mathbf{r}_g^{(1)}, \dots, \mathbf{r}_g^{(m)}]$. With a final linear layer consisting of a relation matrix $\mathbf{M}_r \in \mathbb{R}^{|\mathcal{R}| \times 3d}$ and a bias vector $\mathbf{b}_r \in \mathbb{R}^{|\mathcal{R}|}$, we aggregate the bag information with o in two ways:

Average: The bag elements are averaged as:

$$\mathbf{b}_g = \frac{1}{m} \sum_{i=1}^m \mathbf{r}_g^{(i)}$$

Selective attention (Lin et al., 2016): For a row \mathbf{r} in \mathbf{M}_r representing the relation $r \in \mathcal{R}$, we get the attention weights as:

$$\alpha_i = \frac{\exp(\mathbf{r}^T \mathbf{r}_g^{(i)})}{\sum_{j=1}^m \exp(\mathbf{r}^T \mathbf{r}_g^{(j)})}$$

$$\mathbf{b}_g = \sum_{i=1}^m \alpha_i \mathbf{r}_g^{(i)}$$

Following \mathbf{b}_g , a *softmax* classifier is applied to predict the probability $p(r|\mathbf{b}_g; \theta)$ of relation r being a true relation with θ representing the model parameters, where we minimize the cross-entropy loss during training.

4 Experiments

4.1 Data

Similar to (Dai et al., 2019), UMLS⁴ (Bodenreider, 2004) is used as our KB and MEDLINE abstracts⁵ as our text source. A data summary is shown in Table 1 (see Appendix A for details on the data creation pipeline). We approximate the same statistics as reported in Dai et al. (2019) for relations and entities, but it is important to note that the data does not contain the same samples. We divided triples into train, validation and test sets, and following (Weston et al., 2013; Dai et al., 2019), we make sure that there is no overlapping facts across the splits. Additionally, we add another constraint, i.e., there is no sentence-level overlap between the training and held-out sets. To perform groups negative sampling, for the collection of evidence sentences supporting NA relation type bags, we extend KGC open-world assumption to *bag-level* MIL (see A.3). 20% of the data is reserved for testing, and of the remaining 80%, we use 10% for validation and the rest for training.

⁴We use 2019 release: umls-2019AB-full

⁵<https://www.nlm.nih.gov/bsd/medline.html>

Table 1: Overall statistics of the data.

Triples	Entities	Relations	Pos. Groups	Neg. Groups
169,438	27,403	355	92,070	64,448

4.2 Models and Evaluation

We compare each tagging scheme, *s-tag* and *k-tag*, with average (*avg*) and selective attention (*attn*) bag aggregation functions. To test the setup of Wu and He (2019), which follows *s-tag*, we expand each relation type (*exprels*) $r \in \mathcal{R}$ to two sub-classes $r(e_1, e_2)$ and $r(e_2, e_1)$ indicating relation direction from first entity to second and vice versa. For all experiments, we used batch size 2, bag size 16 with sampling (see A.4 for details on bag composition), learning rate $2e^{-5}$ with linear decay, and 3 epochs. As the standard practice (Weston et al., 2013), evaluation is performed through constructing candidate triples by combining the entity pairs in the test set with all relations (except NA) and ranking the resulting triples. The extracted triples are matched against the test triples and the precision-recall (PR) curve, area under the PR curve (AUC), F1 measure, and Precision@ k , for k in $\{100, 200, 300, 2000, 4000, 6000\}$ are reported.

4.3 Results

Performance metrics are shown in Table 2 and plots of the resulting PR curves in Figure 3. Since our data differs from Dai et al. (2019), the AUC cannot be directly compared. However, Precision@ k indicates the general performance of extracting the true triples, and can therefore be compared. Generally, models annotated with *k-tag* perform significantly better than other models, with *k-tag+avg* achieving state-of-the-art Precision@ $\{2k, 4k, 6k\}$ compared to the previous best (Dai et al., 2019). The best model of Dai et al. (2019) uses PCNN sentence encoder, with additional tasks of Simple (Kazemi and Poole, 2018) based KGC and KG-attention, entity-type classification and named entity recognition. In contrast our data-driven method, *k-tag*, greatly simplifies this by directly encoding the KB information, i.e., order of the *head* and *tail* entities and therefore, the latent relation direction. Consider again the example in Figure 1 where our source triple (h, r, t) is (*neurofibromatosis 1, associated_genetic_condition, breast cancer*), and only last sentence has the same order of entities as KB. This discrepancy is conveniently resolved (note in Figure 2, for last sentence the extracted entities

Table 2: Relation extraction results for different model configurations and data splits.

Model	Bag Agg.	AUC	F1	P@100	P@200	P@300	P@2k	P@4k	P@6k
Dai et al. (2019)	-	-	-	-	-	-	.913	.829	.753
<i>s-tag</i>	<i>avg</i>	.359	.468	.791	.704	.649	.504	.487	.481
	<i>attn</i>	.122	.225	.587	.563	.547	.476	.441	.418
<i>s-tag+exprels</i>	<i>avg</i>	.383	.494	.508	.519	.521	.507	.508	.511
	<i>attn</i>	.114	.216	.459	.476	.482	.504	.496	.484
<i>k-tag</i>	<i>avg</i>	.684	.649	.974	.983	.986	.983	.977	.969
	<i>attn</i>	.314	.376	.967	.941	.925	.857	.814	.772

sentence order is flipped to KG order when concatenating, unlike *s-tag*) with *k-tag*. We remark that such knowledge can be seen as learned, when jointly modeling with KGC, however, considering the task of *bag-level* distant RE only, the KG triples are *known* information and we utilize this information explicitly with *k-tag* encoding.

As PCNN (Zeng et al., 2015) can account for the relative positions of head and tail entities, it also performs better than the models tagged with *s-tag* using sentence order. Similar to Alt et al. (2019)⁶, we also note that the pre-trained contextualized models result in sustained long tail performance. *s-tag+exprels* reflects the direct application of Wu and He (2019) to *bag-level* MIL for distant RE. In this case, the relations are explicitly extended to model entity direction appearing first to second in the sentence, and vice versa. This implicitly introduces independence between the two sub-classes of the same relation, limiting the gain from shared knowledge. Likewise, with such expanded relations, class imbalance is further enhanced to more fine-grained classes.

Though selective attention (Lin et al., 2016) has been shown to improve the performance of distant RE (Luo et al., 2017; Han et al., 2018a; Alt et al., 2019), models in our experiments with such an attention mechanism significantly underperformed, in each case bumping the area under the PR curve and making it flatter. We note that more than 50% of bags are under-sized, in many cases, with only 1-2 sentences, requiring repeated over-sampling to match fixed bag size, therefore, making it difficult for attention to learn a distribution over the bag with repetitions, and further adding noise. For such cases, the distribution should ideally be close to uniform, as is the case with averaging, resulting in better performance.

⁶Their model does not use any entity marking strategy.




Figure 3: Precision-Recall (PR) curve for different models. We see that the models with *k-tag* perform better than the *s-tag* with average aggregation showing consistent performance for long-tail relations.

5 Conclusion

This work extends BERT to *bag-level* MIL and introduces a simple data-driven strategy to reduce the noise in distantly supervised biomedical RE. We note that the *position* of entities in sentence and the *order* in KB encodes the latent *direction* of relation, which plays an important role for learning under such noise. With a relatively simple methodology, we show that this can sufficiently be achieved by reducing the need for additional tasks and highlighting the importance of data quality.

Acknowledgements

The authors would like to thank the anonymous reviewers for helpful feedback. The work was partially funded by the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 777107 through the project Precise4Q and by the German Federal Ministry of Education and Research (BMBF) through the project DEEPLLEE (01IW17001).

References

- Christoph Alt, Marc Hübner, and Leonhard Hennig. 2019. Fine-tuning pre-trained transformer language models to distantly supervised relation extraction. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1388–1398.
- Olivier Bodenreider. 2004. The unified medical language system (UMLS): integrating biomedical terminology. *Nucleic acids research*, 32(suppl_1):D267–D270.
- Qin Dai, Naoya Inoue, Paul Reisert, Ryo Takahashi, and Kentaro Inui. 2019. Distantly supervised biomedical knowledge acquisition via knowledge graph based attention. In *Proceedings of the Workshop on Extracting Structured Knowledge from Scientific Publications*, pages 1–10.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. 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)*, pages 4171–4186.
- Zhou GuoDong, Su Jian, Zhang Jie, and Zhang Min. 2005. Exploring various knowledge in relation extraction. In *Proceedings of the 43rd annual meeting on association for computational linguistics*, pages 427–434. Association for Computational Linguistics.
- Xu Han, Tianyu Gao, Yuan Yao, Deming Ye, Zhiyuan Liu, and Maosong Sun. 2019. OpenNRE: An open and extensible toolkit for neural relation extraction. In *Proceedings of EMNLP-IJCNLP: System Demonstrations*, pages 169–174.
- Xu Han, Zhiyuan Liu, and Maosong Sun. 2018a. Neural knowledge acquisition via mutual attention between knowledge graph and text. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Xu Han, Hao Zhu, Pengfei Yu, Ziyun Wang, Yuan Yao, Zhiyuan Liu, and Maosong Sun. 2018b. FewRel: A large-scale supervised few-shot relation classification dataset with state-of-the-art evaluation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4803–4809.
- Raphael Hoffmann, Congle Zhang, Xiao Ling, Luke Zettlemoyer, and Daniel S Weld. 2011. Knowledge-based weak supervision for information extraction of overlapping relations. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 541–550. Association for Computational Linguistics.
- Seyed Mehran Kazemi and David Poole. 2018. Simple embedding for link prediction in knowledge graphs. In *Advances in neural information processing systems*, pages 4284–4295.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240.
- Adam Lerer, Ledell Wu, Jiajun Shen, Timothee Lacroix, Luca Wehrstedt, Abhijit Bose, and Alex Peysakhovich. 2019. PyTorch-BigGraph: A large-scale graph embedding system. In *Proceedings of the 2nd SysML Conference*, Palo Alto, CA, USA.
- Yankai Lin, Shiqi Shen, Zhiyuan Liu, Huanbo Luan, and Maosong Sun. 2016. Neural relation extraction with selective attention over instances. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2124–2133.
- Bingfeng Luo, Yansong Feng, Zheng Wang, Zhanxing Zhu, Songfang Huang, Rui Yan, and Dongyan Zhao. 2017. Learning with noise: Enhance distantly supervised relation extraction with dynamic transition matrix. *arXiv preprint arXiv:1705.03995*.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2-Volume 2*, pages 1003–1011. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9.
- Sebastian Riedel, Limin Yao, and Andrew McCallum. 2010. Modeling relations and their mentions without labeled text. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 148–163. Springer.
- Alan Ritter, Luke Zettlemoyer, Mausam, and Oren Etzioni. 2013. Modeling missing data in distant supervision for information extraction. *Transactions of the Association for Computational Linguistics*, 1:367–378.
- Livio Baldini Soares, Nicholas FitzGerald, Jeffrey Ling, and Tom Kwiatkowski. 2019. Matching the blanks: Distributional similarity for relation learning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2895–2905.

- Yu Sun, Shuohuan Wang, Yukun Li, Shikun Feng, Xuyi Chen, Han Zhang, Xin Tian, Danxiang Zhu, Hao Tian, and Hua Wu. 2019. ERNIE: Enhanced representation through knowledge integration. *arXiv preprint arXiv:1904.09223*.
- Mihai Surdeanu, Julie Tibshirani, Ramesh Nallapati, and Christopher D Manning. 2012. Multi-instance multi-label learning for relation extraction. In *Proceedings of the 2012 joint conference on empirical methods in natural language processing and computational natural language learning*, pages 455–465. Association for Computational Linguistics.
- Shingo Takamatsu, Issei Sato, and Hiroshi Nakagawa. 2012. Reducing wrong labels in distant supervision for relation extraction. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*, pages 721–729. Association for Computational Linguistics.
- Théo Trouillon, Christopher R Dance, Éric Gaussier, Johannes Welbl, Sebastian Riedel, and Guillaume Bouchard. 2017. Knowledge graph completion via complex tensor factorization. *The Journal of Machine Learning Research*, 18(1):4735–4772.
- Linlin Wang, Zhu Cao, Gerard De Melo, and Zhiyuan Liu. 2016. Relation classification via multi-level attention cnns. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1298–1307.
- Xinggang Wang, Yongluan Yan, Peng Tang, Xiang Bai, and Wenyu Liu. 2018. Revisiting multiple instance neural networks. *Pattern Recognition*, 74:15–24.
- Jason Weston, Antoine Bordes, Oksana Yakhnenko, and Nicolas Usunier. 2013. Connecting language and knowledge bases with embedding models for relation extraction. In *Conference on Empirical Methods in Natural Language Processing*, pages 1366–1371.
- Shanchan Wu and Yifan He. 2019. Enriching pre-trained language model with entity information for relation classification. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pages 2361–2364.
- Daojian Zeng, Kang Liu, Yubo Chen, and Jun Zhao. 2015. Distant supervision for relation extraction via piecewise convolutional neural networks. In *Proceedings of the 2015 conference on empirical methods in natural language processing*, pages 1753–1762.
- Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, and Jun Zhao. 2014. Relation classification via convolutional deep neural network. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2335–2344.

A Data Pipeline

In this section, we explain the steps taken to create the data for distantly-supervised (DS) biomedical relation extraction (RE). We highlight the importance of a data creation pipeline as the quality of data plays a key role in the downstream performance of our model. We note that a pipeline is likewise important for generating reproducible results, and contributes toward the possibility of having either a benchmark dataset or a repeatable set of rules.

A.1 UMLS processing

The fact triples were obtained for English concepts, filtering for RO relation types only (Dai et al., 2019). We collected 9.9M (CUI.head, relation.text, CUI.tail) triples, where CUI represents the concept unique identifier in UMLS.

A.2 MEDLINE processing

From 34.4M abstracts, we extracted 160.4M unique sentences. To perform fast and scalable search, we use the Trie data structure⁷ to index all the textual descriptions of UMLS entities. In obtaining a clean set of sentences, we set the minimum and maximum sentence character length to 32 and 256 respectively, and further considered only those sentences where matching entities are mentioned only once. The latter decision is to lower the noise that may come when only one instance of multiple occurrences is marked for a matched entity. With these constraints, the data was reduced to 118.7M matching sentences.

A.3 Groups linking and negative sampling

Recall the entity groups $\mathcal{G} = \mathcal{G}^+ \cup \mathcal{G}^-$ (Section 3.1). For training with NA relation class, we generate hard negative samples with an open-world assumption (Soares et al., 2019; Lerer et al., 2019) suited to bag-level multiple instance learning (MIL). From 9.9M triples, we removed the relation type and collected 9M CUI groups in the form of (h, t) . Since each CUI is linked to more than one textual form, all of the text combinations for two entities must be considered for a given pair, resulting in 531M textual groups \mathcal{T} for the 586 relation types.

Next, for each matched sentence, let \mathcal{P}_s^2 denote the size 2 permutations of entities present in the sentence, then $\mathcal{T} \cap \mathcal{P}_s^2$ return groups which are present in KB and have matching evidence (positive

⁷<https://github.com/vi3k6i5/flashtext>

groups, \mathcal{G}^+). Simultaneously, with a probability of $\frac{1}{2}$, we remove the h or t entity from this group and replace it with a novel entity e in the sentence, such that the resulting group (e, t) or (h, e) belongs to \mathcal{G}^- . This method results in sentences that are seen both for the true triple, as well as for the invalid ones. Further using the constraints that the relation group sizes must be between 10 to 1500, we find 354⁸ relation types (approximately the same as Dai et al. (2019)) with 92K positive groups and 2.1M negative groups, which were reduced to 64K by considering a random subset of 70% of the positive groups. Table 1 provides these summary statistics.

A.4 Bag composition and data splits

For bag composition, we created bags of constant size by randomly under- or over-sampling the sentences in the bag (Han et al., 2019) to avoid larger bias towards common entities (Soares et al., 2019). The true distribution had a long tail, with more than 50% of the bags having 1 or 2 sentences. We defined a bag to be *uniform*, if the special markers represent the same entity in each sentence, either h or t . If the special markers can take on both h or t , we consider that bag to have a *mix* composition. The k -tag scheme, on the other hand, naturally generates uniform bags. Further, to support the setting of Wu and He (2019), we followed the s -tag scheme and expanded the relations by adding a suffix to denote the directions as $r(e_1, e_2)$ or $r(e_2, e_1)$, with the exception of the NA class, resulting in 709 classes. For fair comparisons with k -tag, we generated uniform bags with s -tag as well, by keeping e_1 and e_2 the same per bag. Due to these bag composition and class expansion (in one setting, *exprels*) differences, we generated three different splits, supporting each scheme, with the same test sets in cases where the classes are not expanded and a different test set when the classes are expanded. Table A.1 shows the statistics for these splits.

Table A.1: Different data splits.

Model	Set Type	Triples	Triples (w/o NA)	Groups	Sentences (Sampled)
k -tag	train	92,972	48,563	92,972	1,487,552
	valid	13,555	8,399	15,963	255,408
	test	33,888	20,988	38,860	621,760
s -tag	train	91,555	47,588	125,852	2,013,632
	valid	13,555	8,399	22,497	359,952
	test	33,888	20,988	55,080	881,280
s -tag+ <i>exprels</i>	train	125,155	71,402	125,439	2,007,024
	valid	22,604	16,298	22,607	361,712
	test	55,083	39,282	55,094	881,504

⁸355 including NA relation