

# FewRel: A Large-Scale Supervised Few-Shot Relation Classification Dataset with State-of-the-Art Evaluation

Xu Han<sup>1,\*</sup> Hao Zhu<sup>1,\*</sup> Pengfei Yu<sup>2,\*</sup> Ziyun Wang<sup>1,†,\*</sup>  
Yuan Yao<sup>1</sup> Zhiyuan Liu<sup>1,‡</sup> Maosong Sun<sup>1</sup>

<http://zhuhao.me/fewrel>

Institute for Artificial Intelligence

State Key Laboratory of Intelligent Technology and Systems

<sup>1</sup>Department of CST, <sup>2</sup>Department of EE, Tsinghua University, Beijing, China

{hanxu17, zhuhao15, yupf15, yy18}@mails.tsinghua.edu.cn,  
ziyunw@nyu.edu, {lzy, sms}@tsinghua.edu.cn

## Abstract

We present a Few-Shot Relation Classification Dataset (FewRel), consisting of 70,000 sentences on 100 relations derived from Wikipedia and annotated by crowdworkers. The relation of each sentence is first recognized by distant supervision methods, and then filtered by crowdworkers. We adapt the most recent state-of-the-art few-shot learning methods for relation classification and conduct thorough evaluation of these methods. Empirical results show that even the most competitive few-shot learning models struggle on this task, especially as compared with humans. We also show that a range of different reasoning skills are needed to solve our task. These results indicate that few-shot relation classification remains an open problem and still requires further research. Our detailed analysis points multiple directions for future research. All details and resources about the dataset and baselines are released on <http://zhuhao.me/fewrel>.

## 1 Introduction

Relation classification (RC) is an important task in NLP, aiming to determine the correct relation between two entities in a given sentence. Many works have been proposed for this task, including kernel methods (Zelenko et al., 2002; Mooney and Bunescu, 2006), embedding methods (Gormley et al., 2015), and neural methods (Zeng et al., 2014). The performance of these conventional models heavily depends on time-consuming and labor-intensive annotated data, which make themselves hard to generalize well. Adopting distant supervision is a primary approach to alleviate this problem for RC (Mintz et al.; Riedel et al.; Hoffmann et al., 2011; Surdeanu et al., 2012; Zeng

Supporting Set	
(A) capital_of	(1) <i>London</i> is the capital of <i>the U.K.</i> (2) <i>Washington</i> is the capital of <i>the U.S.A.</i>
(B) member_of	(1) <i>Newton</i> served as the president of <i>the Royal Society</i> . (2) <i>Leibniz</i> was a member of <i>the Prussian Academy of Sciences</i> .
(C) birth_name	(1) <i>Samuel Langhorne Clemens</i> , better known by his pen name <i>Mark Twain</i> , was an American writer. (2) <i>Alexei Maximovich Peshkov</i> , primarily known as <i>Maxim Gorky</i> , was a Russian and Soviet writer.
Test Instance	
(A) or (B) or (C)	<i>Euler</i> was elected a foreign member of <i>the Royal Swedish Academy of Sciences</i> .

Table 1: An example for a 3 way 2 shot scenario. Different colors indicate different entities, blue for head entity, and red for tail entity.

et al., 2015; Lin et al., 2016), which heuristically aligns knowledge bases (KBs) and text to automatically annotate adequate amounts of training instances. We evaluate the model proposed by Lin et al. (2016), which is followed by the recent state-of-the-art methods (Zeng et al., 2017; Ji et al., 2017; Huang and Wang, 2017; Wu et al., 2017; Liu et al., 2017; Feng et al., 2018; Zeng et al., 2018), on the benchmark dataset NYT-10 (Riedel et al.). Though it achieves promising results on common relations, the performance of a relation drops dramatically when its number of training instances decrease. About 58% of the relations in NYT-10 are long-tail with fewer than 100 instances. Furthermore, distant supervision suffers from the wrong labeling problem, which makes it harder to classify long-tail relations. Hence, it is necessary to study training RC models with insufficient training instances.

We formulate RC as a few-shot learning task in this paper, which requires models capable of handling classification task with a handful of training instances, as shown in Table 1. Many efforts have devoted to few-shot learning. The early works (Caruana, 1995; Bengio, 2012; Donahue et al.,

\* The first four authors contribute equally. The order is determined by dice rolling.

† Z. Wang is now at New York University.

‡ Correspondence author.

2014) apply transfer learning methods to fine-tune pre-trained models from the common classes containing adequate instances to the uncommon classes with only few instances. Then metric learning methods (Koch et al., 2015; Vinyals et al., 2016; Snell et al., 2017) have been proposed to learn the distance distributions among classes. Similar classes are adjacent in the distance space. The metric methods also take advantage of non-parametric estimation to make models efficient and general. Recently, the idea of meta-learning is proposed, which encourages the models to learn fast-learning abilities from previous experience and rapidly generalize to new concepts. Many meta-learning models (Ravi and Larochelle, 2017; Santoro et al., 2016; Finn et al., 2017; Munkhdalai and Yu, 2017) achieve the state-of-the-art results on several few-shot benchmarks.

Though meta-learning methods develop fast, most of these works evaluate on two popular datasets, Omniglot (Lake et al., 2015) and mini-ImageNet (Vinyals et al., 2016). Both the datasets concentrate on image classification. Many works in NLP mainly focus on the zero-shot/semi-supervised scenario (Xie et al., 2016; Ma et al., 2016; Carlson et al., 2009), which incorporate extra information to classify objects never appearing in the training sets. However, the few-shot scenario needs models to classify objects with few instances without any extra information. Recently, Yu et al. (2018) propose a multi-metric method for few-shot text classification. However, there lack systematic researches about adopting few-shot learning for NLP tasks. We propose FewRel: a new large-scale supervised **Few-shot Relation Classification** dataset. To address the wrong labeling problem in most distantly supervised RC datasets, we apply crowd-sourcing to manually remove the noise.<sup>i</sup>

Besides constructing the dataset, we systematically implement the most recent state-of-the-art few-shot learning methods and adapt them for RC. We conduct a detailed evaluation for all these models on our dataset. Though the state-of-the-art few-shot learning methods have much lower results than humans on our challenging dataset, they significantly outperform the vanilla RC mod-

els, indicating that incorporating few-shot learning is promising and needs further research. In summary, our contribution is three-fold:

(1) We formulate RC as a few-shot learning task, and propose a new large supervised few-shot RC dataset.

(2) We systematically adapt the most recent state-of-the-art few-shot learning methods for RC, which may further benefit other NLP tasks.

(3) We conduct a comprehensive evaluation of few-shot learning methods on our dataset, which indicates some promising research directions for RC.

## 2 FewRel Dataset

In this section, we describe the process of creating FewRel in detail. The whole procedure can be divided into two steps: (1) We create a large candidate set of sentences aligned to relations via distant supervision. (2) We ask human annotators to filter out the wrong labeled sentences for each relation to finally achieve a clean RC dataset.

### 2.1 Distant Supervision

For the first step, We use Wikipedia as the corpus<sup>ii</sup> and Wikidata as the KB. Wikidata is a large-scale KB where many entities are already linked to Wikipedia articles. The articles in Wikipedia also contain anchors linking to each other. Thus it is convenient to align sentences in Wikipedia articles to KB facts in Wikidata. We also employ entity linking technique to extract more unanchored entities in articles. We first adopt named entity recognition via spaCy<sup>iii</sup> to find possible entity mentions, then match each mention with the name of an entity in KBs, and link the mention to the entity if successfully matched.

For each sentence  $s$  in Wikipedia articles containing head and tail entities  $e_1$  and  $e_2$ , if there exists a Wikidata statement  $(e_1, e_2, r)$  meaning  $e_1$  and  $e_2$  have the relation  $r$ , we denote the  $(s, e_1, e_2, r)$  tuple as an instance and add it to the candidate set. Empirically, many instances of a given relation contain the same entity pair. For such relation, classifiers may prefer memorizing the entity pairs in the training instances rather than grasping the sentence semantics. Therefore, in the candidate set of each relation, we only keep 1 instance for each unique entity pair. Finally, we

<sup>i</sup> Many previous works, such as (Roth et al., 2013; Luo et al., 2017; Xin et al., 2018) have worked on automatically removing noise from distantly supervision. Instead, we use crowd-sourcing methods to achieve a high accuracy.

<sup>ii</sup> We use whole Wikipedia articles as corpus, not just the first sentence.

<sup>iii</sup> <https://spacy.io/>

remove relations with fewer than 1000 instances, and randomly keep 1000 instances for the rest of the relations. As a result, we get a candidate set of 122 relations and 122,000 instances.

## 2.2 Human Annotation

Next, we invite some well-educated annotators to filter the raw data on a platform similar to Amazon MTurk developed by ourselves. The platform presents each annotator with one instance each time, by showing the sentence, two entities in the sentence, and the corresponding relation labeled by distant supervision. The platform also provides the name of the entities and relation in Wikidata accompanied with the detailed description of that relation. Then the annotator is asked to judge whether the relation could be deduced only from the sentence semantics. We also ask the annotator to mark an instance as negative if the sentence is not complete, or the mention is falsely linked with the entity.

Relations are randomly assigned to annotators from the candidate set, and each annotator will consecutively annotate 20 instances of the same relation before switching to next relation. To ensure the labeling quality, each instance is labeled by at least two annotators. If the two annotators have disagreements on this instance, it will be assigned to a third annotator. As a result, each instance has at least two same annotations, which will be the final decision. After the annotation, we remove relations with fewer than 700 positive instances. For the remaining 105 relations, we calculate the inter-annotator agreement for each relation using the free-marginal multirater kappa (Randolph, 2005), and keep the top 100 relations.

## 2.3 Dataset Statistics

The final FewRel dataset consists of 100 relations, each has 700 instances. A full list of relations, including their names and descriptions, is provided in Appendix A.2. The average number of tokens in each sentence is 24.99, and there are 124,577 unique tokens in total. Following recent meta-learning tasks (Vinyals et al., 2016), which use separate sets of classes for training and testing, we use 64, 16, and 20 relations for training, validation, and testing respectively. Table 2 provides a comparison of our FewRel dataset to two other popular few-shot classification datasets, Omniglot and mini-ImageNet. Table 3 provides a comparison of FewRel to the previous RC datasets, in-

Dataset	#cls.	#inst./cls	#insts.
Omniglot	1,623	20	32,460
mini-ImageNet	100	600	60,000
FewRel	100	700	70,000

Table 2: Comparison of FewRel with Omniglot and mini-ImageNet.

Dataset	#cls.	#insts.
SemEval-2010 Task 8	9	6,674
ACE 2003-2004	24	16,771
TACRED	42	21,784
NYT-10	57	143,391
FewRel	100	70,000

Table 3: Comparison of FewRel with existing RC datasets. Note that negative (no relation) instances in some datasets are ignored.

cluding SemEval-2010 Task 8 dataset (Hendrickx et al., 2009), ACE 2003-2004 dataset (Strassel et al., 2008), TACRED dataset (Zhang et al., 2017), and NYT-10 dataset (Riedel et al., 2010). While some RC datasets contain instances with no relations (negative), we ignore such instances for comparison.

## 3 Experiments

We conduct comprehensive evaluations of vanilla RC models with simple strategies such as finetune or kNN on our new dataset. We also evaluate the recent state-of-the-art few-shot learning methods.

### 3.1 Task Formulation

In few-shot relation classification, we intend to obtain a function  $F : (\mathcal{R}, \mathcal{S}, x) \mapsto y$ . Here  $\mathcal{R} = \{r_1, \dots, r_m\}$  defines the relations that the instances are classified into.  $\mathcal{S}$  is a support set

$$\mathcal{S} = \{(x_1^1, r_1), (x_1^2, r_1), \dots, (x_1^{n_1}, r_1), \dots, (x_m^1, r_m), (x_m^2, r_m), \dots, (x_m^{n_m}, r_m)\} \quad (1)$$

including  $n_i$  instances for each relation  $r_i \in \mathcal{R}$ . For relation classification, a data instance  $x_i^j$  is a sentence accompanied with a pair of entities. The query data  $x$  is an unlabeled instance to classify, and  $y \in \mathcal{R}$  is the prediction of  $x$  given by  $F$ .

In recent research on few-shot learning,  $N$  way  $K$  shot setting is widely adopted. We follow this setting for the few-shot relation classification problem. To be exact, for  $N$  way  $K$  shot learning

$$N = m = |\mathcal{R}|, K = n_1 = \dots = n_m \quad (2)$$

Model	5 Way 1 Shot	5 Way 5 Shot	10 Way 1 Shot	10 Way 5 Shot
Finetune (CNN)	44.21 $\pm$ 0.44	68.66 $\pm$ 0.41	27.30 $\pm$ 0.28	55.04 $\pm$ 0.31
Finetune (PCNN)	45.64 $\pm$ 0.62	57.86 $\pm$ 0.61	29.65 $\pm$ 0.40	37.43 $\pm$ 0.42
kNN (CNN)	54.67 $\pm$ 0.44	68.77 $\pm$ 0.41	41.24 $\pm$ 0.31	55.87 $\pm$ 0.31
kNN (PCNN)	60.28 $\pm$ 0.43	72.41 $\pm$ 0.39	46.15 $\pm$ 0.31	59.11 $\pm$ 0.30
Meta Network (CNN)	64.46 $\pm$ 0.54	80.57 $\pm$ 0.48	53.96 $\pm$ 0.56	69.23 $\pm$ 0.52
GNN (CNN)	66.23 $\pm$ 0.75	81.28 $\pm$ 0.62	46.27 $\pm$ 0.80	64.02 $\pm$ 0.77
SNAIL (CNN)	67.29 $\pm$ 0.26	79.40 $\pm$ 0.22	53.28 $\pm$ 0.27	68.33 $\pm$ 0.25
Prototypical Network (CNN)	69.20 $\pm$ 0.20	84.79 $\pm$ 0.16	56.44 $\pm$ 0.22	75.55 $\pm$ 0.19
Human performance	92.22 $\pm$ 5.53	-	85.88 $\pm$ 7.40	-

Table 4: Accuracies (%) of all models on FewRel under four different settings.

### 3.2 Experiment Settings

We consider four types of few-shot tasks in our experiments: 5 way 1 shot, 5 way 5 shot, 10 way 1 shot, 10 way 5 shot. Under this setting, we evaluate different few-shot training strategies and state-of-the-art few-shot learning methods built upon two widely used instance encoders, CNN (Zeng et al., 2014) and PCNN (Zeng et al., 2015).

For both CNN and PCNN, the sentence is first represented to the input vectors by transforming each word into concatenation of word embeddings and position embeddings. In CNN, the input vectors pass a convolution layer, a max-pooling layer, and a non-linear activation layer to get the final output sentence embedding. PCNN is a variant of CNN, which replaces the max-pooling operation with a piecewise max-pooling operation.

To evaluate this two vanilla models in few-shot RC task, we first consider two training strategies, namely Finetune and kNN. For the Finetune baseline, it learns to classify all relations on the training set with CNN/PCNN, and tune parameters on the support set. We only tune the parameters of output layer, and keep other parameters unchanged. For the kNN baseline, it also jointly classifies all relations during training, while at the test time, it uses the neural networks to embed all the instances and then adopts k-nearest-neighbor (kNN) to classify the test instances.

By adapting them to relation classification, we also evaluate four recently proposed few-shot learning methods, including Meta Network (Munkhdalai and Yu, 2017), GNN (Satorras and Estrach, 2018), SNAIL (Mishra et al., 2018), and Prototypical Network (Snell et al., 2017). We describe briefly about these baselines in Sec. 3.3. If you are familiar with these methods, you can safely skip that subsection. The hyperparameters of each model are selected via grid search against

the validation set.

Human performance is also evaluated under 5 way 1 shot setting and 10 way 1 shot setting. A human labeler is given 5/10 instances from different relations and one extra test instance. Human labelers are asked to decide which relation the test instance belongs to. Note that these labelers are not provided the name of the relations and any extra information. Since 5 way 5 shot and 10 way 5 shot settings are easier, we only evaluate performance of 5 way 1 shot and 10 way 1 shot.

### 3.3 Baselines of Few-shot Learning Models

**Meta Network** Meta Network (Munkhdalai and Yu, 2017) is a meta learning algorithm utilizing a high level *meta learner* on top of the traditional classification model, or *base learner*, to supervise the training process. The weights of base learner are divided into two groups, fast weights and slow weights. Fast weights are generated by the meta learner, whereas slow weights are simply updated by minimizing classification loss. The fast weights are expected to help the model generalize to new tasks with very few training instances.

**GNN** GNN (Satorras and Estrach, 2018) tackles the few-shot learning problem by considering each supporting instance or query instance as a node in the graph. For those instances in the support sets, label information is also embedded into the corresponding node representations. Graph neural networks are then employed to propagate the information between nodes. A query instance is expected to receive information from support sets in order to make the classification. In our adaption, while the instances are encoded by CNNs, labels are represented by one-hot encoding.

**SNAIL** SNAIL (Mishra et al., 2018) is a meta learning model that utilizes temporal convolu-



tional neural networks and attention modules for fast learning from past experience. SNAIL arranges all the supporting instance-label pairs into a sequence and appends the query instance behind them. Such an order agrees with the temporal order of learning process where we learn information by reading supporting instances before making predictions for unlabeled instances. Temporal convolution (a 1-D convolution) is then performed along the sequence to aggregate information across different time steps and a causally masked attention model is used over the sequence to aggregate useful information from former instances to latter ones.

**Prototypical Networks** Prototypical Network (Snell et al., 2017) is a few-shot classification model based on the assumption that for each class there exists a prototype. The model tries to find the prototypes for classes from supporting instances, and compares the distance between the query instance and each prototype under certain distance metric. Prototypical network learns an embedding function  $u$  to embed each class’s instances, and computes each prototype by averaging over all the output embeddings of instances in the support set  $S$  that are labeled with the corresponding class.

#### 4 Result Analysis and Future Work

We report evaluation results in Table 4. From our preliminary experiments, PCNN with few-shot learning methods perform 3-10 percentages worse than CNN, therefore only CNN results are shown in our experimental results. From the results, we observe that integrating few-shot learning methods into CNN significantly outperforms CNN/PCNN with finetune or kNN, which means adapting few-shot learning methods for RC is promising. However, there are still huge gaps between their performance and humans’, which means our dataset is a challenging testbed for both relation classification and few-shot learning.

In this paper, we propose a new large and high quality dataset, FewRel, for few-shot relation classification task. This dataset provides a new point of view for RC, and also a new benchmark for few-shot learning. Through the evaluation of different few-shot learning methods, we find even the best model performs much worse than humans, which suggests there is still large space for few-shot learning methods to improve.

The most challenging characteristic of our

Sentence	Reasoning
<i>Chris Bohjalian</i> graduated from <i>Amherst College</i> Summa Cum Laude, where he was a member of the Phi Beta Kappa Society.	Simple Pattern
<i>James Alry</i> obtained a 1st class honours (Physics) at <i>Liverpool University</i> .	Common-sense Reasoning
He was a professor at <i>Reed College</i> , where he taught <i>Steve Jobs</i> , and replaced Lloyd J. Reynolds as the head of the calligraphy program.	Logical Reasoning
He and <i>Cesare Borgia</i> were thought to be close friends since childhood, going on to accompany one another during their studies at the <i>University of Pisa</i> .	Co-reference Reasoning

Table 5: Examples from relation “educated\_at”. Different colors indicate different entities, blue for head entity, and red for tail entity.

dataset is the diversity in expressing the same relation. We provide some examples from FewRel in Table 5, showing different reasoning modes needed for classifying some instances. Future researches may consider incorporating common-sense knowledge or improved causal modules.

#### Acknowledgement

This work is supported by the National Natural Science Foundation of China (NSFC No. 61572273, 61532010). This work is also funded by the Natural Science Foundation of China (NSFC) and the German Research Foundation (DFG) in Project Crossmodal Learning, NSFC 61621136008 / DFC TRR-169. Hao Zhu is supported by Tsinghua University Initiative Scientific Research Program. We thank all annotators for their hard work. We also thank all members from Tsinghua NLP Lab for their strong support for annotator recruitment.

#### References

- Yoshua Bengio. 2012. Deep learning of representations for unsupervised and transfer learning. In *Proceedings of ICML*.
- Andrew Carlson, Justin Betteridge, Estevam R Hruschka Jr, and Tom M Mitchell. 2009. Coupling semi-supervised learning of categories and relations. In *Proceedings of the NAACL HLT 2009 Workshop on Semi-supervised Learning for Natural Language Processing*, pages 1–9. Association for Computational Linguistics.
- Rich Caruana. 1995. Learning many related tasks at the same time with backpropagation. In *Proceedings of NIPS*.
- Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell.

2014. Decaf: A deep convolutional activation feature for generic visual recognition. In *Proceedings of ICML*.
- Jun Feng, Minlie Huang, Li Zhao, Yang Yang, and Xiaoyan Zhu. 2018. Reinforcement learning for relation classification from noisy data. In *Proceedings of AAAI*.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of ICML*.
- Matthew R. Gormley, Mo Yu, and Mark Dredze. 2015. Improved relation extraction with feature-rich compositional embedding models. In *Proceedings of EMNLP*.
- Iris Hendrickx, Su Nam Kim, Zornitsa Kozareva, Preslav Nakov, Diarmuid Ó Séaghdha, Sebastian Padó, Marco Pennacchiotti, Lorenza Romano, and Stan Szpakowicz. 2009. Semeval-2010 task 8: Multi-way classification of semantic relations between pairs of nominals. In *Proceedings of SemEval@ACL*.
- 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 ACL*.
- Yi Yao Huang and William Yang Wang. 2017. Deep residual learning for weakly-supervised relation extraction. In *Proceedings of EMNLP*.
- Guoliang Ji, Kang Liu, Shizhu He, Jun Zhao, et al. 2017. Distant supervision for relation extraction with sentence-level attention and entity descriptions. In *Proceedings of AAAI*.
- Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. 2015. Siamese neural networks for one-shot image recognition. In *Proceedings of ICML*.
- Brenden M. Lake, Ruslan Salakhutdinov, and Joshua B. Tenenbaum. 2015. Human-level concept learning through probabilistic program induction. *Science*.
- Yankai Lin, Shiqi Shen, Zhiyuan Liu, Huanbo Luan, and Maosong Sun. 2016. Neural relation extraction with selective attention over instances. In *Proceedings of ACL*.
- Tianyu Liu, Kexiang Wang, Baobao Chang, and Zhi-fang Sui. 2017. A soft-label method for noise-tolerant distantly supervised relation extraction. In *Proceedings of EMNLP*.
- 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. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 430–439.
- Yukun Ma, Erik Cambria, and Sa Gao. 2016. Label embedding for zero-shot fine-grained named entity typing. In *COLING*.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. Distant supervision for relation extraction without labeled data. In *Proceedings of ACL-IJCNLP*.
- Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. 2018. A simple neural attentive meta-learner. In *Proceedings of ICLR*.
- Raymond J Mooney and Razvan C Bunescu. 2006. Subsequence kernels for relation extraction. In *Proceedings of NIPS*.
- Tsendsuren Munkhdalai and Hong Yu. 2017. Meta networks. In *Proceedings of ICML*.
- Justus J Randolph. 2005. Free-marginal multirater kappa (multirater  $\kappa_{\text{free}}$ ): an alternative to fleiss’ fixed-marginal multirater kappa. In *Proceedings of JLIS*.
- Sachin Ravi and Hugo Larochelle. 2017. Optimization as a model for few-shot learning. In *Proceedings of ICLR*.
- Sebastian Riedel, Limin Yao, and Andrew McCallum. Modeling relations and their mentions without labeled text. In *Proceedings of ECML-PKDD*.
- Sebastian Riedel, Limin Yao, and Andrew D McCallum. 2010. Modeling relations and their mentions without labeled text. In *Proceedings of ECML-PKDD*.
- Benjamin Roth, Tassilo Barth, Michael Wiegand, and Dietrich Klakow. 2013. A survey of noise reduction methods for distant supervision. In *Proceedings of the 2013 workshop on Automated knowledge base construction*, pages 73–78. ACM.
- Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. 2016. Meta-learning with memory-augmented neural networks. In *Proceedings of ICML*.
- Victor Garcia Satorras and Joan Bruna Estrach. 2018. Few-shot learning with graph neural networks. In *Proceedings of ICLR*.
- Jake Snell, Kevin Swersky, and Richard S. Zemel. 2017. Prototypical networks for few-shot learning. In *Proceedings of NIPS*.
- Stephanie Strassel, Mark A. Przybocki, Kay Peterson, Zhiyi Song, and Kazuaki Maeda. 2008. Linguistic resources and evaluation techniques for evaluation of cross-document automatic content extraction. In *Proceedings of LREC*.
- Mihai Surdeanu, Julie Tibshirani, Ramesh Nallapati, and Christopher D Manning. 2012. Multi-instance multi-label learning for relation extraction. In *Proceedings of EMNLP*.

- Oriol Vinyals, Charles Blundell, Tim Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. 2016. Matching networks for one shot learning. In *Proceedings of NIPS*.
- Yi Wu, David Bamman, and Stuart Russell. 2017. Adversarial training for relation extraction. In *Proceedings of EMNLP*.
- Ruobing Xie, Zhiyuan Liu, Jia Jia, Huanbo Luan, and Maosong Sun. 2016. Representation learning of knowledge graphs with entity descriptions. In *AAAI*.
- Ji Xin, Hao Zhu, Xu Han, Zhiyuan Liu, and Maosong Sun. 2018. Put it back: Entity typing with language model enhancement. In *Proceedings of EMNLP*.
- Mo Yu, Xiaoxiao Guo, Jinfeng Yi, Shiyu Chang, Saloni Potdar, Yu Cheng, Gerald Tesauro, Haoyu Wang, and Bowen Zhou. 2018. Diverse few-shot text classification with multiple metrics. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, volume 1, pages 1206–1215.
- Dmitry Zelenko, Chinatsu Aone, and Anthony Richardella. 2002. Kernel methods for relation extraction. *JMLR*.
- Daojian Zeng, Kang Liu, Yubo Chen, and Jun Zhao. 2015. Distant supervision for relation extraction via piecewise convolutional neural networks. In *Proceedings of EMNLP*.
- Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, and Jun Zhao. 2014. Relation classification via convolutional deep neural network. In *Proceedings of COLING*.
- Wenyuan Zeng, Yankai Lin, Zhiyuan Liu, and Maosong Sun. 2017. Incorporating relation paths in neural relation extraction. In *Proceedings of EMNLP*.
- Xiangrong Zeng, Shizhu He, Kang Liu, and Jun Zhao. 2018. Large scaled relation extraction with reinforcement learning. In *Proceedings of AAAI*.
- Yuhao Zhang, Victor Zhong, Danqi Chen, Gabor Angeli, and Christopher D. Manning. 2017. Position-aware attention and supervised data improve slot filling. In *Proceedings of EMNLP*.