

FewRel 2.0: Towards More Challenging Few-Shot Relation Classification

Tianyu Gao¹, Xu Han¹, Hao Zhu^{1,2}, Zhiyuan Liu^{1*}, Peng Li³,
Maosong Sun¹, Jie Zhou³

¹Department of Computer Science and Technology, Tsinghua University, Beijing, China

²Carnegie Mellon University, Pittsburgh, PA, USA

³Pattern Recognition Center, WeChat AI, Tencent Inc, China

{gty16, hanxu17}@mails.tsinghua.edu.cn, zhuhao@cmu.edu

liuzy@tsinghua.edu.cn, patrickpli@tencent.com

sms@tsinghua.edu.cn, withtomzhou@tencent.com

Abstract

We present FewRel 2.0, a more challenging task to investigate two aspects of few-shot relation classification models: (1) Can they adapt to a new domain with only a handful of instances? (2) Can they detect none-of-the-above (NOTA) relations? To construct FewRel 2.0, we build upon the FewRel dataset (Han et al., 2018) by adding a new test set in a quite different domain, and a NOTA relation choice. With the new dataset and extensive experimental analysis, we found (1) that the state-of-the-art few-shot relation classification models struggle on these two aspects, and (2) that the commonly-used techniques for domain adaptation and NOTA detection still cannot handle the two challenges well. Our research calls for more attention and further efforts to these two real-world issues. All details and resources about the dataset and baselines are released at <https://github.com/thunlp/fewrel>.

1 Introduction

Few-shot learning, which requires models to handle new classification tasks with only a handful of training instances, has drawn much attention in recent years (Ravi and Larochelle, 2017; Vinyals et al., 2016; Munkhdalai and Yu, 2017; Snell et al., 2017). To advance this field in NLP, Han et al. (2018) propose FewRel, a large-scale dataset to explore few-shot learning in relation classification. Many efforts (Gao et al., 2019; Soares et al., 2019) have been devoted to the new task and some of the methods even exceed human performance¹ on FewRel. Based on the dataset FewRel, we propose FewRel 2.0, a new task containing two real-world issues that FewRel ignores: (1) few-shot domain adaptation, and (2) few-shot none-of-the-above detection.

* Corresponding author

¹<https://thunlp.github.io/fewrel.html>




Figure 1: The comparison between the best results of the current models on FewRel, few-shot DA and few-shot NOTA. From the figures we can see that even the state-of-the-art models struggle on the new tasks.

Few-shot domain adaptation (few-shot DA) aims to evaluate the abilities of few-shot models to transfer across domains, which is crucial for real-world applications, since the test domains usually lack of annotations and could differ vastly from the training domains. To this end, we construct a new test set sharing great disparities with the original FewRel dataset, and carry out extensive experiments on the state-of-the-art few-shot models and commonly-used domain adaptation methods. Some prior experimental results in Figure 1 show that even the performance of the most effective methods on FewRel drops drastically on the new test set, proving that few-shot DA is challenging and requires further investigations.

Few-shot none-of-the-above detection (few-shot NOTA) is an advanced version of the existing N -way K -shot setting in few-shot learning. The original N -way K -shot setting samples N classes, as well as K supporting instances and several queries from each class for each test batch, assuming that all queries belong to the sampled N classes. However, in few-shot NOTA, queries could also be none-of-the-above (NOTA), which brings one more option in classification and challenges existing few-shot methods. Consider-

Training Phase (Famous Person from Wikipedia)		
Supp. Set	(A) date_of_birth	<i>Mark Twain</i> was born in <i>1835</i> .
	(B) place_of_birth	<i>Elvis Presley</i> was born in <i>Memphis, Tennessee</i> .
Query	(A) or (B) or NOTA	<i>William Shakespeare</i> passed away at age 52 (around <i>1616</i>).
Test Phase (Biomedicine)		
Supp. Set	(A) may_treat	<i>Ribavirin</i> remains essential to <i>Chronic Hepatitis C</i> treatment.
	(B) manifestation_of	Boys with <i>Prader-Willi syndrome</i> often have <i>undescended testicle</i> .
Query	<u>(A)</u> or (B) or NOTA	<i>Thiabendazole</i> was effective in eradicating the <i>strongyloides infection</i> .

Table 1: An example for a 2-way 1-shot scenario, including both few-shot DA and few-shot NOTA. Different colors indicate different entities, blue for head entities, and red for tail entities. For few-shot DA, instances in the training phase and test phase come from different domains. For few-shot NOTA, it requires models to detect the none-of-the-above (NOTA) relation.

ing few-shot NOTA has not yet been widely explored, we propose several solutions based on the state-of-the-art few-shot models and evaluate them with few-shot NOTA setting. Figure 1 shows that though achieving promising results, there is still a room of improvements for few-shot NOTA.

In the following sections, we first describe the two newly-added challenges in FewRel 2.0, then introduce possible directions for addressing these two issues, and finally present results and observations from our experiments.

2 FewRel 2.0

Formulation for N -Way K -Shot Setting

The original FewRel task adopts the N -way K -shot setting. The whole dataset is divided into training, validation and test subsets, which have no intersection in relation types. Models are evaluated with batches sampled from the test set, each of which consists of $(\mathcal{R}, \mathcal{S}, x, r)$, where $\mathcal{R} = \{r_1, r_2, \dots, r_N\}$ is the sampled relation set, $r \in \mathcal{R}$ is the correct relation label for the query x , and \mathcal{S} is the supporting set containing K instances for each relation,

$$\mathcal{S} = \{(x_{r_i}^j, r_i)\}, 1 \leq i \leq N, 1 \leq j \leq K. \quad (1)$$

Models should predict the relation label $y \in \mathcal{R}$ for the query instance x based on the given \mathcal{S} and \mathcal{R} . Both of the following two challenges are based on this N -way K -shot setting.

Few-Shot Domain Adaptation

Both the training and test sets of the original FewRel dataset are constructed by manually annotating the distantly supervised (Bunescu and Mooney, 2007; Mintz et al., 2009) results on Wikipedia corpus and Wikidata (Vrandečić and Krötzsch, 2014) knowledge bases. In other words, they are from the same domain, yet in a real-world scenario, we might train models on one domain and perform few-shot learning on a different one. For example, we may train models on Wikipedia, which has large amounts of data and adequate annotations, and then perform few-shot learning on some domains suffering data sparsity, like literature, finance and medicine. Note that, not only do these corpora differ vastly from each other in morphology and syntax, but there are wide disparities between the relation sets defined on these domains as well, which makes transferring knowledge across different domains more challenging.

To explore few-shot DA, we construct a new test set by aligning PubMed², a database containing large amounts of biomedical literature, with UMLS³, a large-scale knowledge base in the biomedical sciences. Then we let the annotators classify whether each instance we get from the distant supervision is correct. Every sentence is assigned to at least two annotators, and if their annotation results do not agree with each other, the third annotator is assigned. In the end, we gather a valid dataset with 25 relations and 100 instances for each relation.

For few-shot DA, we adopt the original FewRel training set for training, and the newly-annotated dataset for test, as shown in Table 1. Besides, we use SemEval-2010 task 8 dataset (Hendrickx et al., 2009) as the validation set, since both the corpora and the schema of SemEval-2010 task 8 are in different domains from the original FewRel dataset and the newly-annotated test set.

Few-Shot None-of-the-Above Detection

In a N -way K -shot, all queries are assumed to be in the given relation set, yet sentences expressing no specific relations or relations not in the given set should also be taken into consideration, for they make up the vast majority of text. This calls for the none-of-the-above (NOTA) relation, which

²<https://www.ncbi.nlm.nih.gov/pubmed/>

³UMLS represents the Unified Medical Language System®, which is the trademark of U.S. National Library of Medicine.

indicates that the query instance does not express any of the given relations. Though it is common in some conventional classification tasks, where NOTA is usually regarded as an extra class, detecting NOTA could be hard in few-shot learning, because the given relation sets are not fixed so that the NOTA relation requires to cover a different semantic space each time. An example of NOTA is given in Table 1.

We formalize few-shot NOTA based on the N -way K -shot setting. For the query instance x , the correct relation label becomes $r \in \{r_1, r_2, \dots, r_N, \text{NOTA}\}$ rather than $r \in \{r_1, r_2, \dots, r_N\}$. We use the parameter NOTA rate to describe the proportion of NOTA queries during the whole test phase. For example, 0% NOTA rate means no queries are NOTA and 50% NOTA rate means half of the queries have the label NOTA.

The NOTA queries are sampled from those relations outside the given N relations. To be more specific, denoting the whole test set as $\mathcal{D}_{\text{test}}$, the set containing all instances in the relation set \mathcal{R} as $\mathcal{D}_{\mathcal{R}}$ and the NOTA rate as α , α of the query instances (NOTA queries) are from $\mathcal{D}_{\text{test}} \setminus \mathcal{D}_{\mathcal{R}}$ and $1 - \alpha$ of the instances are from $\mathcal{D}_{\mathcal{R}}$.

Note that during the test phase, all the queries are from the test set, though models can sample instances from the training set as supporting instances for NOTA relation (this method is described explicitly in Section 4). Also note that to better demonstrate the effects of the NOTA relation, we use the original FewRel dataset for few-shot NOTA, instead of the new test set, which can get rid of the influence of domain adaptation.

3 Approaches for Few-Shot DA

Many efforts have been devoted for domain adaptation, like subspace mapping (Pan et al., 2010; Fernando et al., 2013), finding domain-invariant spaces (Baktashmotlagh et al., 2013; Ganin et al., 2016), feature augmentation (Blitzer et al., 2006) and minimax estimators (Provost and Fawcett, 2001). Among them, adversarial training (Goodfellow et al., 2015; Ganin et al., 2016; Wang et al., 2018) has been proved to be efficient in finding domain-invariant features. It is a game process between an encoder and a discriminator, where the encoder tries to generate domain-invariant features while the discriminator tries to tell which domain the features are from.

Here we follow the adversarial training setting

in Wang et al. (2018), where a two-layer perceptron network is used as the discriminator. While training the few-shot learning task, we feed the sentence encoder E and the discriminator D with the corpora from the training domain and the test domain, and optimize the min-max game,

$$\min_{\theta_E} \max_{\theta_D} \sum_{x \in \mathcal{C}_0} \log[D(E(x))]_0 + \sum_{x \in \mathcal{C}_1} \log[D(E(x))]_1, \quad (2)$$

where $[\cdot]_i$ is the i -th element of the vector, \mathcal{C}_0 is the training corpus and \mathcal{C}_1 is the test corpus.

4 Approaches for Few-Shot NOTA

A simple way to handle NOTA is to regard it as an extra class in the N -way K -shot setting. To be more specific, we can sample instances outside the N relations as the supporting data of NOTA, and perform the $(N + 1)$ -way K -shot learning. As compared to the current methods ignoring NOTA, this approach does not bring much improvements, since the supporting data for NOTA actually belong to several different relations and are scattered in the feature space, making it hard to perform classification.

To better address few-shot NOTA, we propose a model named **BERT-PAIR** based on the sequence classification model in BERT (Devlin et al., 2019). We pair each query instance with all the supporting instances, concatenate each pair as one sequence, and send the concatenated sequence to the BERT sequence classification model to get the score of the two instances expressing the same relation. Denote the BERT model as B , the query instance as x and the paired supporting instance as x_r^j (the j -th supporting instance for the relation r), $B(x, x_r^j)$ outputs a two-element vector corresponding to scores of the pair sharing the same relation and not sharing the same relation. The probability over each relation in the few-shot scenario, including NOTA, is addressed as follows,

$$p(y = r|x) = \frac{\exp(o_r)}{\sum_{r' \in \mathcal{R}} \exp(o_{r'})}, r \in \mathcal{R} \quad (3)$$

where y is the predicted label and $\mathcal{R} = \{r_1, \dots, r_N, \text{NOTA}\}$ is the relation set including NOTA. For $r \in \{r_1, \dots, r_N\}$, o_r is calculated by averaging,

$$o_r = \frac{1}{K} \sum_{j=1}^K [B(x, x_r^j)]_1. \quad (4)$$

Model	5-Way 1-Shot		5-Way 5-Shot	
	On 1.0	On 2.0	On 1.0	On 2.0
GNN (CNN)	66.23 ± 0.75	27.94 ± 0.03	81.28 ± 0.62	29.33 ± 0.11
Proto (CNN)	74.52 ± 0.07	35.09 ± 0.10	88.40 ± 0.06	49.37 ± 0.10
Proto-ADV (CNN)	70.28 ± 0.15	42.21 ± 0.09	84.63 ± 0.07	58.71 ± 0.06
Proto (BERT)	80.68 ± 0.28	40.12 ± 0.19	89.60 ± 0.09	51.50 ± 0.29
Proto-ADV (BERT)	73.35 ± 0.95	41.90 ± 0.44	82.30 ± 0.53	54.74 ± 0.22
BERT-PAIR	88.32 ± 0.64	56.25 ± 0.40	93.22 ± 0.13	67.44 ± 0.54

Model	10-Way 1-Shot		10-Way 5-Shot	
	On 1.0	On 2.0	On 1.0	On 2.0
GNN (CNN)	46.27 ± 0.80	16.44 ± 0.04	64.02 ± 0.77	18.26 ± 0.03
Proto (CNN)	62.38 ± 0.06	22.98 ± 0.05	80.45 ± 0.08	35.22 ± 0.06
Proto-ADV (CNN)	56.34 ± 0.08	28.91 ± 0.10	74.67 ± 0.12	44.35 ± 0.09
Proto (BERT)	71.48 ± 0.15	26.45 ± 0.10	82.89 ± 0.11	36.93 ± 0.01
Proto-ADV (BERT)	61.49 ± 0.69	27.36 ± 0.50	72.60 ± 0.38	37.40 ± 0.36
BERT-PAIR	80.63 ± 0.17	43.64 ± 0.46	87.02 ± 0.12	53.17 ± 0.09

Table 2: Accuracies (%) on few-shot DA. “On 1.0” represents the results on the original FewRel dataset and “On 2.0” represents the results on the new test set. The models with “-ADV” use adversarial training described in Section 3.

The score for NOTA o_{NOTA} is calculated by the equation,

$$o_{\text{NOTA}} = \min_{r \in \{r_1, \dots, r_N\}} \frac{1}{K} \sum_{j=1}^K [B(x, x_r^j)]_0. \quad (5)$$

Then we can treat NOTA the same as other relations and optimize the model with the cross entropy loss, which is commonly-used in few-shot learning and other classification tasks.

5 Experiments

5.1 Baseline Models for Few-Shot Learning

We pick the two best models from the results in Han et al. (2018), GNN (Satorras and Estrach, 2018) and **Prototypical Networks** (Snell et al., 2017), as our baseline models. As for the encoders, besides the CNN encoder used in Han et al. (2018), we also adopt BERT since it achieves the state-of-the-arts in multiple tasks (Devlin et al., 2019). For all models and encoders, we follow the parameter settings from Han et al. (2018) and Devlin et al. (2019).

5.2 Evaluation Results on Few-Shot DA

Table 2 demonstrates the evaluation results of few-shot DA on the existing FewRel test set and the new test set. Besides the baselines, we also evaluate Prototypical Networks with adversarial training described in Section 3 and our proposed




Figure 2: 5-way K -shot results under different NOTA rates. Models with * simply ignore the NOTA setting and assume all queries can be classified as one of the N relations.

BERT-PAIR model in Section 4. We get three observations from the results:

- (1) All few-shot models suffer dramatic performance falls when tested on a different domain.
- (2) Adversarial training does improve the results on the new test domain, yet still has large space for growth.
- (3) BERT-PAIR outperforms all other few-shot models on both 1.0 and 2.0 test set.

Besides, to see where the growth boundary is, we split 10 relations, 1,000 instances out of the 2.0 test set and add them to the training set, then train and evaluate BERT-PAIR on the new data. We get 72.30% for 5-way 1-shot and 80.50% for 5-way 5-shot, 16 and 13 points higher than the current best results. Note that only 1,000 training instances can lead to such an enormous gap, indicating that

Model	5-Way-1-Shot			
	0% NOTA	15% NOTA	30% NOTA	50% NOTA
Proto (CNN)*	74.52 ± 0.07	62.18 ± 0.22	53.38 ± 0.14	37.26 ± 0.04
Proto (CNN)	69.17 ± 0.07	60.59 ± 0.05	53.18 ± 0.12	40.00 ± 0.10
Proto (BERT)*	80.68 ± 0.28	67.92 ± 0.31	58.22 ± 0.20	40.64 ± 0.14
Proto (BERT)	81.65 ± 0.97	70.02 ± 0.23	61.08 ± 0.28	45.94 ± 0.50
BERT-PAIR*	88.32 ± 0.64	73.60 ± 0.51	63.00 ± 0.47	43.99 ± 0.09
BERT-PAIR	76.73 ± 0.55	77.67 ± 0.14	78.49 ± 0.21	80.31 ± 0.12

Model	5-Way-5-Shot			
	0% NOTA	15% NOTA	30% NOTA	50% NOTA
Proto (CNN)*	88.40 ± 0.06	73.64 ± 0.11	62.95 ± 0.12	44.20 ± 0.05
Proto (CNN)	85.23 ± 0.07	77.79 ± 0.03	71.96 ± 0.14	61.66 ± 0.08
Proto (BERT)*	89.60 ± 0.09	75.03 ± 0.17	64.44 ± 0.18	45.22 ± 0.03
Proto (BERT)	88.74 ± 0.83	83.79 ± 0.44	81.17 ± 0.48	75.21 ± 0.52
BERT-PAIR*	93.22 ± 0.13	77.58 ± 0.42	66.41 ± 0.24	46.58 ± 0.09
BERT-PAIR	83.32 ± 0.38	84.19 ± 0.46	84.64 ± 0.13	86.06 ± 0.43

Table 3: Accuracies (%) on few-shot NOTA. Models with * simply ignore the NOTA setting and assume all queries can be classified as one of the given relations.

there is still a huge room for improvements.

5.3 Evaluation Results on Few-Shot NOTA

We evaluate Prototypical Networks with the naive NOTA solution described in Section 4 and BERT-PAIR under the NOTA setting. All models are trained given 50% NOTA queries and tested under four different NOTA rates: 0%, 15%, 30%, 50%. To show how accuracy falls if ignoring the NOTA relation, we also demonstrate the results of models without considering NOTA (marked with * in Figure 2). We demonstrate the evaluation results in Figure 2. For detailed numbers of results on few-shot NOTA, please refer to Table 3. From Figure 2 we can conclude that:

(1) Treating NOTA as the $N + 1$ relation is beneficial for handling Few-Shot NOTA, though the results still fall fast when the NOTA rate increases.

(2) BERT-PAIR works better under the NOTA setting for its binary-classification style model, and stays stable with rising NOTA rate.

(3) Though BERT-PAIR achieves promising results, huge gaps still exist between the conventional (0% NOTA rate) and NOTA settings (gaps of 8 points for 5-way 1-shot and 7 points for 5-way 5-shot with 50% NOTA rate), which calls for further research to address the challenge.

6 Conclusion

In this paper, we propose FewRel 2.0, a more challenging few-shot relation classification task with

a new test set from the biomedical domain and the none-of-the-above setting. The purpose of the new task is to explore two aspects which are ignored in the previous work: few-shot domain adaptation (few-shot DA) and few-shot none-of-the-above detection (few-shot NOTA). Extensive experiments demonstrate that the existing state-of-the-art few-shot models struggle on the new task. We also point out some possible directions to handle these two issues, implement several new models and evaluate them with the new task. Though achieving promising improvements, these commonly-used techniques are still not the satisfactory solutions for few-shot DA and few-shot NOTA, which requires further explorations in these two real-world challenges.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (NSFC No. 61572273, 61661146007) and Tsinghua University Initiative Scientific Research Program (20151080406). This work is also supported by the Pattern Recognition Center, WeChat AI, Tencent Inc. Han and Gao are supported by 2018 and 2019 Tencent Rhino-Bird Elite Training Program respectively. Gao is also supported by Tsinghua University Initiative Scientific Research Program. We also thank Xiaozhi Wang for his insightful ideas and suggestions.

References

- Mahsa Baktashmotlagh, Mehrtash T Harandi, Brian C Lovell, and Mathieu Salzmann. 2013. Unsupervised domain adaptation by domain invariant projection. In *Proceedings of ICCV*, pages 769–776.
- John Blitzer, Ryan McDonald, and Fernando Pereira. 2006. Domain adaptation with structural correspondence learning. In *Proceedings of EMNLP*, pages 120–128.
- Razvan Bunescu and Raymond Mooney. 2007. Learning to extract relations from the web using minimal supervision. In *Proceedings of ACL*, pages 576–583.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.
- Basura Fernando, Amaury Habrard, Marc Sebban, and Tinne Tuytelaars. 2013. Unsupervised visual domain adaptation using subspace alignment. In *Proceedings of ICCV*, pages 2960–2967.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *JMLR*, 17(1):2096–2030.
- Tianyu Gao, Xu Han, Zhiyuan Liu, and Maosong Sun. 2019. Hybrid attention-based prototypical networks for noisy few-shot relation classification. In *Proceedings of AAAI*, pages 6407–6414.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and harnessing adversarial examples. In *ICLR*.
- Xu Han, Hao Zhu, Pengfei Yu, Ziyun Wang, Yuan Yao, Zhiyuan Liu, and Maosong Sun. 2018. Fewrel: A large-scale supervised few-shot relation classification dataset with state-of-the-art evaluation. In *Proceedings of EMNLP*, pages 4803–4809.
- 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 SEW-2009*, pages 94–99.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of ACL-IJCNLP*, pages 1003–1011.
- Tsendsuren Munkhdalai and Hong Yu. 2017. Meta networks. In *Proceedings of ICML*, pages 2554–2563.
- Sinno Jialin Pan, Ivor W Tsang, James T Kwok, and Qiang Yang. 2010. Domain adaptation via transfer component analysis. *TNNLS*, 22(2):199–210.
- Foster Provost and Tom Fawcett. 2001. Robust classification for imprecise environments. *ML*, 42(3):203–231.
- Sachin Ravi and Hugo Larochelle. 2017. Optimization as a model for few-shot learning. In *Proceedings of ICLR*.
- 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 Zemel. 2017. Prototypical networks for few-shot learning. In *Proceedings of NIPS*, pages 4077–4087.
- Livio Baldini Soares, Nicholas FitzGerald, Jeffrey Ling, and Tom Kwiatkowski. 2019. Matching the blanks: Distributional similarity for relation learning. In *Proceedings of ACL*, pages 2895–2905.
- Oriol Vinyals, Charles Blundell, Tim Lillicrap, Daan Wierstra, et al. 2016. Matching networks for one shot learning. In *Proceedings of NIPS*, pages 3630–3638.
- Denny Vrandečić and Markus Krötzsch. 2014. Wiki-data: a free collaborative knowledgebase. *Communications of the ACM*, 57(10):78–85.
- Xiaozhi Wang, Xu Han, Yankai Lin, Zhiyuan Liu, and Maosong Sun. 2018. Adversarial multi-lingual neural relation extraction. In *Proceedings of COLING*, pages 1156–1166.