

Pre-training to Match for Unified Low-shot Relation Extraction

Fangchao Liu^{1,3}, Hongyu Lin^{1,*}, Xianpei Han^{1,2,4}, Boxi Cao^{1,3}, Le Sun^{1,2}

¹Chinese Information Processing Laboratory ²State Key Laboratory of Computer Science
Institute of Software, Chinese Academy of Sciences, Beijing, China

³University of Chinese Academy of Sciences, Beijing, China

⁴Beijing Academy of Artificial Intelligence, Beijing, China

{fangchao2017, hongyu, xianpei, boxi2020, sunle}@iscas.ac.cn

Abstract

Low-shot relation extraction (RE) aims to recognize novel relations with very few or even no samples, which is critical in real scenario application. Few-shot and zero-shot RE are two representative low-shot RE tasks, which seem to be with similar target but require totally different underlying abilities. In this paper, we propose Multi-Choice Matching Networks to unify low-shot relation extraction. To fill in the gap between zero-shot and few-shot RE, we propose the triplet-paraphrase meta-training, which leverages triplet paraphrase to pre-train zero-shot label matching ability and uses meta-learning paradigm to learn few-shot instance summarizing ability. Experimental results on three different low-shot RE tasks show that the proposed method outperforms strong baselines by a large margin, and achieve the best performance on few-shot RE leaderboard¹.

1 Introduction

Relation extraction (RE) aims to extract the relation between two given entities in the context. The most popular approaches to build RE models are based on supervised learning (Zeng et al., 2014; Baldini Soares et al., 2019). Despite the superior performance, supervised relation extraction approaches severely suffer from the data bottleneck, which restricts their application to more relation types in real scenarios.

Consequently, low-shot relation extraction has become a recent research hotspot in RE area. There are two mainstream learning paradigms widely explored in low-shot relation extraction, namely zero-shot RE (Levy et al., 2017) and few-shot RE (Han et al., 2018). Few-shot relation extraction aims to identify instances of novel relation type with only a few illustrative instances, while zero-shot RE is more progressive, which only uses external

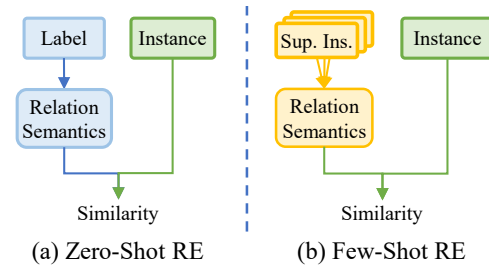


Figure 1: Difference between zero-shot RE and few-shot RE. (a) Zero-shot requires for the ability of label semantic matching, while (b) Few-shot requires for the ability of support instance (Sup. Ins.) summarizing.

knowledge and the name or definition of the novel relations to recognize them. Because low-shot RE only requires very limited manually annotated data, it can effectively alleviate data bottlenecks in conventional RE and therefore attached great attention.

However, even with similar goals, zero-shot RE and few-shot RE actually require different fundamental abilities. Specifically, zero-shot RE is built on *label semantic matching* ability, which requires models to sufficiently exploit the label semantic of given novel relations, and matches relations and queries based on their underlying semantics. While few-shot RE is built on *instance semantic summarizing* ability, which requires a model to quickly generalize to novel relations by summarizing critical information from few-shot instances. Due to this fundamental difference, current state-of-the-art architectures are separately learned to deal with these two low-shot RE tasks. For zero-shot RE, the most popular solution is to transform it into a textual entailment (Obamuyide and Vlachos, 2018; Sainz et al., 2021), word prediction (Brown et al., 2020) or MRC problem (Levy et al., 2017; Bragg et al., 2021) and use external resources from these tasks to pre-training the label semantic matching ability. However, the divergence between relation extraction and these tasks will inevitably undermine the performance. Besides, MRC and tex-

*Corresponding authors.

¹https://thunlp.github.io/2/fewrel2_nota.html

tual entailment architecture can only deal with one novel relation each time, which significantly increases the computational and storage cost of deploying such models in real-world scenarios. For few-shot RE, current methods mainly focus on summarizing better prototypes from a few illustrative instances (Snell et al., 2017), or learning a model that can generalize to novel types within a few steps (Finn et al., 2017). These approaches require few-shot examples to fine-tune or summarize prototypes, and therefore can not be directly applied to zero-shot RE. As a result, current relation extraction models can not be effectively and efficiently to apply to all low-shot RE settings.

In this paper, we propose to unify low-shot relation extraction by returning to the essence of relation extraction. Fundamentally, relation extraction can be viewed as a multiple choice task. Given two entities in context, a RE system needs to match the most appropriate relation – or *others* for none-of-the-above – from a set of pre-defined relation types. The information required to accomplish the multiple-choice matching can be summarized from either the surface form of relation name or from few-shot instances. Motivated by this, we propose Multi-Choice Matching Network (MCMN) for unified low-shot RE, which is shown in Figure 2. Specifically, MCMN converts all candidate relation descriptions into a multi-choice prompt. Then the input instance is concatenated with the multi-choice prompt and passes through a pre-trained encoder to obtain the semantic representations of the input instance and candidate relations. Finally, MCMN conduct relation extraction by directly matching the relation representations and the instance representation.

To equip MCMN with both label semantic matching ability and instance semantic summarizing ability, we propose to pre-train MCMN via triplet-paraphrase meta pre-training, which contains the following two critical components: 1) a text-triple-text paraphrase module, which can generate large-scale pseudo relation extraction data to pre-train the label semantic matching ability of MCMN; 2) a meta-learning style training algorithm, which enriches MCMN with instance semantic summarizing ability to quickly generalize across different relation extraction tasks. Specifically, given large-scale raw texts, triplet-paraphrase first extracts (subject, predicate, object) triplets via OpenIE (Cui et al., 2018) toolkit. Then based on

the extracted triplets, paraphrases of the original texts is generated using an RDF-to-Text generation model. In this way, we can obtain large-scale pseudo annotations by collecting the generated sentences and the predicate in the triples. Such corpus can be used to effectively pre-train the label semantic matching ability of MCMN by matching the paraphrases to the corresponding predicate. Furthermore, to enrich MCMN with the instance semantic summarizing ability, such pre-training is conducted in a meta-learning paradigm. That is, MCMN is asked to learn different relation extraction tasks at each iteration, so that the MCMN can not over-fit the pre-training corpus by directly memorizing specific target relations.

To evaluate our methods, we conduct experiments on three fairly different RE tasks: zero-shot RE, few-shot RE, and few-shot RE with none-of-the-above relation. Experiments show that the proposed method outperform previous methods on all these three tasks. Our source code is available at <https://github.com/fc-liu/MCMN>.

The main contributions of this work are:

- We propose MCMN, a unified architecture for low-shot relation extraction by fundamentally formulating relation extract using a multiple-choice matching paradigm.
- We propose to pre-train MCMN with triplet-paraphrase meta training, which enriches MCMN with label semantic matching ability and instance semantic summarizing ability for both zero-shot RE and few-shot RE.
- We comprehensively study the performance of MCNN on three different relation extraction tasks, including zero-shot, few-shot, and few-shot with none-of-the-above relation extraction, where MCMN outperforms strong baseline models.

2 Background

In this section, we formulate relation extraction task and the low-shot RE settings including zero-shot RE and few-shot learning RE.

Relation Extraction. Suppose the input text $T = [t_1, t_2, \dots, t_n]$ contains n tokens, $e_1 = [i, j]$ and $e_2 = [k, l]$ indicate the entity pair spans, where $1 \leq i \leq j, j < k \leq l$ and $l \leq n$. A relation instance is defined as $x = (T, e_1, e_2)$. For example, the




Figure 2: Illustration of our multi-choice matching networks (MCMN).

tuple (“*Tim Cook is the CEO of Apple Inc.*”, “*Tim Cook*”, “*Apple Inc.*”) is a relation instance. The aim of relation extraction is to learn a mapping function: $f : x \rightarrow y$, where y is the relation class. For example, we want mapping (“*Tim Cook is the CEO of Apple Inc.*”, “*Tim Cook*”, “*Apple Inc.*”) to its relation class “*CEO_of*”. Traditional RE tasks typically pre-define the class space Y and annotate a large set of instances to train the model. However, in real scenarios, the targeting relation types vary in different tasks, and most of the novel relations lack annotations, rendering the supervised paradigms inapplicable. In that regard, how to transfer models to novel tasks becomes critical.

Low-shot Relation Extraction. Low-shot relation extraction requires models to recognize novel relations with very few samples. There are two mainstream low-shot RE tasks, including:

Zero-shot RE. This task aims to conduct relation extraction without any annotated instance other than some external knowledge z (or side information), such as relation descriptions. Models are supposed to transfer the knowledge and extract the targeting relation y_t for input instance x through only the external knowledge.

Few-shot RE. This task aims to conduct relation extraction with only a few annotated instances per novel relations. Each few-shot RE task contains a support set $S = S_1, \dots, S_N$ for N novel relations. And for relation i , $S_i = S_i^0, \dots, S_i^K$ contains K annotated instances. Models are supposed to learn to transfer the knowledge and extract the targeting relation y_t for instance x through the N -way K -shot support set.

3 Multi-Choice Matching Networks

In this section, we introduce our multi-choice matching networks (MCMN). Different from previous unifying models, MCMN adopts a much more efficient and lightweight decoding module. Following are the detailed descriptions.

3.1 Multi-choice Prompt

Fundamentally, relation extraction can be viewed as a multiple choice task. Inspired by recent advances of prompt learning (Brown et al., 2020; Schick and Schütze, 2021), we construct a multi-choice prompt for each relation extraction task by directly concatenate all relation names or descriptions. Formally, the multi-choice prompts are in the following form:

$$[C] \text{ rel 1 } [C] \text{ rel 2 } \dots [C] \text{ rel N}$$

where $[C]$ is the placeholder separator for the following relation. For example in Figure 2, the target RE task contains three novel relations: *employee_of*, *ceo_of*, and *others*, of which the relation descriptions are then concatenated altogether to form the multi-choice prompt “[C] employee of [C] ceo of [C] others”. After obtaining the multi-choice prompt, we then feed it accompanied with the input sentence into the instance encoder, and the representations at separator $[C]$ is regarded as the representation of its following relation.

3.2 Instance Encoder

Before instance encoding, we concatenate the multi-choice prompt with each input instance into a single sentence, and separate them with a $[SEP]$ token. Besides, we follow (Baldini Soares et al., 2019) and wrap the given entity pair with $[e1]$, $[/e1]$, $[e2]$ and $[/e2]$ respectively. For the example in Figure 2, the entire input to encoder is: “[CLS] [C] employee of [C] ceo of [C] others [SEP] [e1] Tim Cook [/e1] is the CEO of [e2] Apple [/e2] . [SEP]”. Then we encode the entire sentence x through a Transformer (Vaswani et al., 2017) encoder:

$$\mathbf{h}_{[CLS]}, \mathbf{h}_{[C]}, \dots, \mathbf{h}_{[SEP]} = \mathcal{H}(x), \quad (1)$$

where \mathbf{h} is the output embedding for each token in x , d is the dimension of hidden states. These token embeddings are then used for multi-choice matching and model prediction.

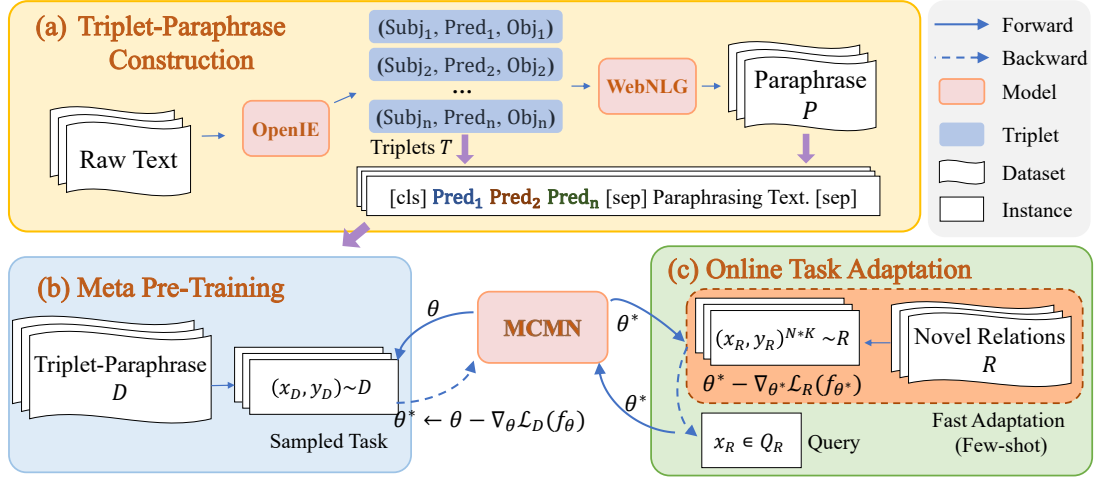


Figure 3: The framework of multi-choice matching training strategy. (a) **Triplet-Paraphrase Construction** conducts triplet-paraphrase pairs for meta training; (b) **Meta Training** on the triplet-paraphrase triplets; (c) **Online Task Meta-Training** performs an online meta-training for each test tasks.

3.3 Multi-choice Matching and Prediction

The multi-choice matching module matches the input instance to the corresponding relation. For each relation type, we use hidden states of $[C]$ marker to represent each following relation:

$$\mathbf{h}_{rel_i} = \mathbf{h}_{[C]_i}, \quad (2)$$

where \mathbf{h}_{rel_i} is the representation for relation i and $\mathbf{h}_{[C]_i}$ is the hidden state for the i^{th} $[C]$ token. For the input text, we simply average hidden states of $[e1]$ and $[e2]$ to obtain the instance representation \mathbf{X} :

$$\mathbf{X} = avg(\mathbf{h}_{[e1]}, \mathbf{h}_{[e2]}). \quad (3)$$

Then we perform matching operation between the instance and each relation:

$$\mathcal{D}(x, y_i) = \|\mathbf{X} - \mathbf{h}_{rel_i}\|_2. \quad (4)$$

In this equation, we adopt the Euclidean distance to measure the similarity, and the corresponding probability for each relation is:

$$P(y_i|x; \theta) = \frac{\exp(-\mathcal{D}(x, y_i))}{\sum_{j=1}^N \exp(-\mathcal{D}(x, y_j))}. \quad (5)$$

Finally, we choose the relation \hat{y} with the maximal probability as the prediction:

$$\hat{y} = \arg \max_i P(y_i|x; \theta). \quad (6)$$

3.4 Training Loss

We adopt an end-to-end training manner by minimizing the following loss function:

$$\mathcal{L}_{(x,y)}(\theta) = - \sum_{i=1}^N \mathbf{I}(y_i) \log P(y_i|x_i; \theta), \quad (7)$$

where $\mathbf{I}(\cdot)$ equals 1 if y_i is the golden class, otherwise $\mathbf{I}(\cdot) = 0$. The three-period training process will be detailed described in the following section.

4 Training Strategies for Multi-Choice Matching Networks

As mentioned above, the required abilities for zero-shot and few-shot are different. In this paper, we propose triplet-paraphrase meta pre-training, which jointly learn the label semantic matching ability required by zero-shot RE and instance summarizing ability required by few-shot RE. Following is the detailed description of the pre-training framework.

4.1 Triplet-Paraphrase Construction

To endow the label semantic matching ability to MCMN, it is required to incorporate large-scale data of both relational sentences and relation types to pre-train the model. Unfortunately, the highly limited relation types in existing RE datasets may lead to overfitting on specific relations and impair the generalization of MCMN. In this paper, we propose the triplet-paraphrase to generate large-scale pre-training data for MCMN from raw texts. The overall procedure of triplet-paraphrase module is demonstrated in Figure 3(a), which extracts predicates from large-scaled raw texts as the relation descriptions. Then we utilize the extracted relational triplets to generate paraphrase sentences for further multi-choice matching pre-training. The elaboration is presented below.

Relational Triplet Extraction. Most complete sentences contain at least one relational triplet,

which includes the subject, predicate, and object. The predicate in a sentence corresponds to property or relation between the subject and object, which can be regarded as a concrete expression of one relationship. Therefore, To extract large-scaled triplets from open domain texts, we use OpenIE model² to extract on article collections of Wikipedia. Considering the example sentence: *The service traces its history to an online service known as PlayNET*. OpenIE model extracts all the possible triplets: *(an online service, known as, PlayNET)* and *(The service, traces, its history)*. We collect all extracted predicates from raw texts to represent the corresponding relations, preventing the models from overfitting specific relation types. These triplets are further used for paraphrase generation and pre-training.

Paraphrase Generation. One drawback of matching predicate as the relation is that the predicate extracted by OpenIE is commonly a span from current sentence, which may lead models to take the shortcut by directly matching through words co-occurrence. To eliminate this shortcut, we follow several recent works (Agarwal et al., 2021; Liu et al., 2021) to generate paraphrase texts to match the predicate. Specifically, for extracted triplets, we first wrap them with special markers “[H], [R], [T]” correspond to subject, predicate and object. Then we input the wrapped triplet texts to generate the paraphrase texts. In our implementation, we adopt T5³ (Raffel et al., 2020) as the generator, and pre-train it on WebNLG dataset (Gardent et al., 2017). For example, we wrap *(an online service, known as, PlayNET)* to “[H] an online service [R] known as [T] PlayNET” then generate the paraphrase text *playnet is an online service*. After generating the paraphrase, we then match it to the corresponding predicate for pre-training.

4.2 Triplet-Paraphrase Meta Pre-training

Each instance in the pre-training batch contains the paraphrase text and the corresponding predicate span. In addition, as shown in Figure 3(a), we concatenate all predicates in the current mini-batch as the multi-choice prompt and follow the training loss in Equation 7 to pre-train MCMN, where $I(y_i)$ equals to 1 when y_i is the corresponding predicate, otherwise, $I(y_i) = 0$.

²<https://github.com/dair-iitd/OpenIE-standalone>

³<https://github.com/UKPLab/plms-graph2text>

Algorithm 1 MCMN for Few-shot Prediction

Require: n : fine-tuning epochs in online period

Require: θ^* : meta-learned model parameters

Require: \mathcal{S} : support set, x_q : query instance

Require: α : learning rate

```

1:  $\theta' = \theta^*$  # save original model
2: for epoch in range( $n$ ) do
3:   # compute loss of the support set:
4:    $\mathcal{L}_{\mathcal{S}} = \mathbb{E}_{(x,y) \in \mathcal{S}} \mathcal{L}_{(x,y)}(\theta^*)$ 
5:   # update model parameters:
6:    $\theta^* \leftarrow \theta^* - \alpha \nabla_{\theta^*} \mathcal{L}_{\mathcal{S}}$ 
7: end for
8:  $y = f_{\theta^*}(x_q)$  # predict the query instance
9:  $\theta^* = \theta'$  # restore the original model
10: return  $y$ 

```

4.3 Online Task Adaptation

In online learning or testing period, we adopt different adaptation strategies for different low-shot tasks. For zero-shot RE, we directly use the trained MCMN to conduct the task. For few-shot RE, we perform an online task meta-training on the support set, as shown in Algorithm 1. For each few-shot task with support set \mathcal{S} and query instance x_q , we first update the model with all support instances:

$$\theta^* \leftarrow \theta^* - \alpha \nabla_{\theta^*} \mathbb{E}_{(x,y) \in \mathcal{S}} \mathcal{L}_{(x,y)}(\theta^*), \quad (8)$$

where α is the learning rate, $\mathcal{L}_{(x,y)}(\theta^*)$ is the loss defined in Equation 7. To avoid overfitting, we use an early-stop criterion controlled by an adaptation epoch threshold that once the adaptation epoch is over the threshold, we exit the online fine-tuning and give the prediction for current query instance x_q :

$$y = f_{\theta^*}(x_q). \quad (9)$$

Finally, we restore the model parameter $\theta^* = \theta'$ and repeat the procedure to the next task.

5 Experiments

5.1 Datasets and Task Settings

We conduct experiments on three low-shot relation extraction tasks: zero-shot RE (Bragg et al., 2021), few-shot RE (Bragg et al., 2021) and the more challenging few-shot RE with none-of-the-above (NOTA) (Gao et al., 2019b). These tasks are all conducted based on FewRel dataset (Han et al., 2018), which is constructed through distantly aligning WikiData triplets to Wikipedia articles. In total, FewRel dataset consists of 100 relation types and

Model	Zero-shot		Few-shot		Avg.
	Acc.±ci.	Std.	Acc.±ci.	Std.	
UniFew (Bragg et al., 2021)	52.5± 2.0	9.7	79.2± 1.5	7.5	65.9
UniFew-meta (Bragg et al., 2021)	79.4 ± 1.9	9.2	87.2± 1.2	5.7	83.3
Our	66.6±1.7	8.7	74.4±1.5	7.6	70.5
Our-meta	82.9±1.3	6.6	87.4±1.2	5.6	85.1
Few-shot with NOTA					
Model	5-way 1-shot 0.15	5-way 5-shot 0.15	5-way 1-shot 0.5	5-way 5-shot 0.5	Avg.
Proto (CNN) (Gao et al., 2019b)	60.59	77.79	40.00	61.66	60.01
Proto (BERT) (Gao et al., 2019b)	70.02	83.79	45.94	75.21	68.74
Bert-Pair (Gao et al., 2019b)	77.67	84.19	80.31	86.06	82.06
2 nd on Leaderboard (anonymous)	79.53	86.31	79.99	81.92	81.94
3 rd on Leaderboard (anonymous)	67.97	81.94	74.85	78.12	75.72
Our-meta (1 st on Leaderboard)	88.40	89.91	84.56	85.32	87.05

Table 1: Results (%) on zero-shot, few-shot, and few-shot with NOTA RE tasks. We report accuracy (Acc.), confidence interval (ci.), and standard deviation (Std., lower is better) for zero- and few-shot RE, and only accuracy for few-shot with NOTA task.

700 instances per type. Standard FewRel settings adopt a split of 64/16/20 fraction corresponding to train/validation/test set, where the train and validation sets are publicly accessible while the test set is not. Following are the detailed settings for each evaluation task.

Zero- and Few-shot Relation Extraction Settings. We follow the standard Flex benchmark settings, which separate the train and validation sets from FewRel into a train set of 65 relations, a validation set of 5 relations and a test set of 10 relations. The test tasks are sampled and processed through the FLEX official toolkit ⁴.

Few-shot RE with NOTA Relation Settings. A drawback of conventional few-shot RE tasks is that they neglect the existence of other relations, that is all query instances are assumed to express one of the given relations in the support set. Gao et al. (2019b) point out this problem and add the “*none-of-the-above*” (NOTA) relation to consider the situation where query instance does not express any of the given relations. In our experiment, we follow the default settings of FewRel benchmark and evaluate our methods on 5-way 1/5-shot tasks with a 15% or 50% NOTA rate.

5.2 Baseline and Evaluation Metrics

Baseline Methods. For zero-shot and few-shot RE tasks, we compare our model with UniFew (Bragg et al., 2021), a unified few-shot learning model based on T5 (Raffel et al., 2020). This model converts each few-shot classification

⁴<https://github.com/allenai/flex>

task into a machine reading comprehension format and predicts the results through generation. With a pre-training period on large-scaled MRC data, this model reaches strong performance on both zero- and few-shot tasks. For few-shot RE with NOTA relation task, we compare our model with Bert-Pair (Gao et al., 2019b), an instance-pair matching framework for few-shot RE tasks. This model computes a similarity and a dissimilarity score simultaneously between query instance and each support instance, then aggregates the similarity score for each relation and dissimilarity score for NOTA relation. And the results of CNN and BERT based prototypical networks from Gao et al. (2019b) are also reported.

Evaluation Metrics. For zero-shot and few-shot RE tasks, we follow FLEX benchmark and report the accuracy, confidence interval and standard deviation correspondingly. All these results reported are from the official Flex toolkits. For few-shot RE with NOTA relation task, we follow FewRel 2.0 benchmark and report the corresponding accuracy for four different settings.

5.3 Hyperparameters and Implementation Details

In the triplet-paraphrase construction period, we extract relation triplets from articles in Wikipedia and generate the counterpart paraphrase texts. Overall, we generate about one million triplet and paraphrase text pairs. In triplet-paraphrase meta-training periods, we use a learning rate of 5e-6, weight decay of 1e-6, dropout rate of 0.5, and a linear learning schedule with a 0.95 weight decay.

In the online task meta-training period, we use learning rate of $5e-6$, and the adaptation epoch of 1 or 2 for FewRel NOTA tasks, epochs of 45 for FLEX tasks, while keep other hyperparameters the same. We use RoBERTa-large (Liu et al., 2019) to initialize our model. Furthermore, to better endow the low-shot capability to our model, we adopt annotated FewRel data (Han et al., 2018) as an additionally supervised meta-training procedure.

5.4 Overall Results

Table 1 shows the overall results on three different RE tasks. From this table, we can see that:

- **MCMN with triplet-paraphrase pre-training outperforms previous methods in all three RE tasks and achieve state-of-the-art performance.** Compared with the strong baseline methods, MCMN achieves remarkable performance improvements. In zero-shot and few-shot RE tasks, MCMN with triplet paraphrase pre-training outperforms the baseline methods by at least 1.8% in average. In few-shot RE with NOTA task, our method outperforms previous best method by at least 4.99% in average and achieve the best performance in the leaderboard.
- **Our triplet-paraphrase pre-training achieves promising results on low-shot RE tasks.** Comparing with other pre-training strategies such as UniFew model pre-trained with large annotated MRC datasets, triplet-paraphrase pre-training achieves much better performance on zero-shot RE tasks. Besides, triplet-paraphrase can further enhance MCMN to achieves the new state-of-the-art results on all three low-shot RE tasks with supervised meta-training procedure, which are detailed analyzed in the next section.
- **MCMN performs more robust than previous methods.** In zero-shot and few-shot tasks, our methods perform a lower standard deviation and more shallow confidence interval than baseline methods, which means the prediction of our methods is more stable across different tasks.

5.5 Detailed Analysis

In this section, we conduct several experiments for in-depth analysis of our methods.

Ablation Studies on Zero- and Few-shot RE Tasks. To evaluate the effect of each part of our methods on zero- and few-shot RE tasks, we conduct separate experiments on triplet-paraphrase

Model	Zero-shot		Few-shot		Avg.
	acc	std.	acc.	std.	
RoBERTa	15.6	5.1	21.4	7.3	18.5
Triplet-Para Pre-train	66.6	8.7	74.4	7.6	70.5
MCMN w/o Pre-train	81.0	6.7	85.3	5.7	83.2
MCMN	82.9	6.6	87.4	5.6	85.1

Table 2: Ablation study results (%) of our methods on FLEX RE benchmark, Sup. Meta stands for supervised meta training.

pre-training, MCMN and MCMN without triplet-paraphrase pre-training on Flex test set. As shown in Table 2, we can see that the pure triplet-paraphrase pre-training model outperforms RoBERTa-large model with a remarkable margin as well as leverages the MCMN model with an improvement of at least 1.9% compared with MCMN without triplet-paraphrase pre-training on both zero-shot and few-shot settings. These results demonstrate that triplet-paraphrase pre-training method can significantly improve the generalization and performance of our model, and the framework of multi-choice matching network is quite applicable in low-shot RE tasks. Besides, we notice the performance of pure triplet-paraphrase pre-training model is lower than MCMN without triplet-paraphrase pre-training. To study this issue, we analyze the triplet-paraphrase data, and find that many of the generated texts still consist of words in predicates, though the expression is quite different from the original sentences. This may still lead to the shortcut learning problem. On top of that, the expression of predicates is much different from the relation name, and the negative predicates are much easier to distinguish than the real test cases. These issues altogether result in poor performance. Fortunately, the triplet-paraphrase pre-training period can properly initialize MCMN and leverage the final performance.

Ablation Studies on Few-shot NOTA RE tasks.

We also conduct detailed analyses of our methods in few-shot NOTA RE tasks. As shown in Table 3, the pure triplet-paraphrase pre-trained model can also boost the performance of roberta-large initialized model and leverage the supervised meta-trained MCMN by at least 0.9% in average. Although we do not consider the NOTA relation in the triplet-paraphrase pre-training period, this period can also contribute to the further supervised meta-training period, which indicates that the matching-pattern learned in triplet-paraphrase pre-training

Model	Few-shot with NOTA				Avg.
	5-way 1-shot 0.15	5-way 5-shot 0.15	5-way 1-shot 0.5	5-way 5-shot 0.5	
RoBERTa	27.37	27.88	16.38	16.50	22.03
Triplet-Para Pre-train	69.00	70.59	43.99	43.66	56.81
MCMN w/o. Pre-train	87.89	90.36	83.22	83.10	86.14
MCMN	88.40	89.91	84.56	85.32	87.05
MCMN w/o. Pre-train (0-shot)	83.08	84.10	83.61	83.45	83.56
MCMN (0-shot)	85.11	85.45	82.72	82.16	83.86

Table 3: Ablation study results (%) of MCMN on FewRel NOTA benchmark. Triplet-Para corresponds to triplet-paraphrase.

period is generalized and robust to down-stream tasks. Besides, we notice that in NA rate of 0.5 tasks, the pure triplet-paraphrase pre-trained model suffers from serious performance drops. This may be caused by the large proportion of negative instances in test tasks. Fortunately, this issue can be alleviated by the online adaptation period.

Zero-shot NOTA RE tasks. This experiment studies the zero-shot performance of our methods on FewRel NOTA tasks. From Table 3, we surprisingly found that our methods also outperform previous state-of-the-art few-shot NOTA models even in zero-shot conditions. This also indicates that our methods are effective in low-shot RE tasks and are robust enough across different settings.

Computing Efficiency of Multi-Choice Matching Networks. This experiment compares the computing efficiency of our method with MRC-based method. Each model is tested on the Flex test set, including both zero-shot and few-shot RE tasks. Models in zero-shot setting only need inference while both models in few-shot setting require fine-tuning on the support set which involves time-consuming back-propagation operations. For fair comparison, we use a single TITAN RTX GPU for each model and keep other computing environments the same. As a result, UniFew takes **647** minutes (more than 10 hours) to finish the test prediction, while our method takes about **80** minutes to obtain the results in Table 1, which improves the speed of roughly an order of magnitude. The main reason for such an efficiency discrepancy is that UniFew, as a generative model, involves an autoregressive decoder to generate the results, whereas our method directly matches the relation and instance representations to give the results.

6 Related Works

Relation Extraction. Recent success of supervised relation extraction methods (Zeng et al., 2014;

Zhou et al., 2016) heavily depends on large amount of annotated data. However, the bottleneck on data annotation severely limits the adaptation of these supervised methods to real scenarios. Recent works reply to this dilemma from the perspective of low-shot learning, which mainly focuses on zero- and few-shot RE tasks. In this work, we shed light on three representative sub-fields tasks, including zero-shot RE, few-shot RE and few-shot RE with NOTA relation to evaluate our methods

Zero-shot Relation Extraction. Levy et al. (2017) firstly introduce the zero-shot relation extraction task and adjust the machine reading comprehension (MRC)-based paradigm for it. Following this line, other MRC-based methods have been proposed (Cetoli, 2020; Bragg et al., 2021). Another paradigm for zero-shot RE is matching-based (Socher et al., 2013), which falls into the text-entailment-based methods (Obamuyide and Vlachos, 2018; Sainz et al., 2021), and the representation matching-based methods (Chen and Li, 2021; Dong et al., 2021). Text-entailment-based methods concatenate the relation description with the input sentence to assess whether they entail the same semantic relationship; Representation matching-based methods separately encode the relation and instance into the same semantic space but are not capable of handling the NOTA relation.

Few-shot Relation Extraction. Han et al. (2018) firstly propose the few-shot relation extraction task and adopt several meta-learning methods (Munkhdalai and Yu, 2017; Snell et al., 2017; Satorras and Estrach, 2018; Mishra et al., 2018) for it. Recent works on few-shot RE mostly centers around the metric-based methods (Vinyals et al., 2016), such as prototype-based methods (Baldini Soares et al., 2019; Ye and Ling, 2019; Gao et al., 2019a) and meta learning-based methods (Finn et al., 2017). Besides, Gao et al. (2019b) extend the FewRel challenge with few-shot domain

adaptation (DA) and none-of-the-above (NOTA) tasks, which are more challenging and closer to real-world application.

Few-shot RE with NOTA. Although NOTA relation is common in conventional supervised RE tasks (Zhang et al., 2017), it is quite different in few-shot scenarios due to the label inconsistency problem. As an example, consider an instance that expresses relation r . In task A, relation r is not included in the support set, and thus model learns the semantic mapping between this instance and the NOTA relation. But in another task B where relation r is included in the support set, the model learned from task A may continue to match this instance to NOTA relation. Because of the difficulty, attempts to resolve this problem are scarce. To the best of our knowledge, Bert-Pair (Gao et al., 2019b) is the only public method for this task, and our work is the first method to unify the zero-shot, few-shot and few-shot with NOTA tasks.

7 Conclusions

In this paper, we propose Multi-Choice Matching Networks to unify low-shot relation extraction. MCMN introduces a multi-choice prompt to formulate relation extraction as in a multi-choice paradigm. To equip MCMN with different zero-shot and few-shot abilities, we propose the triplet-paraphrase meta pre-training, which leverages triplet paraphrase to pre-train zero-shot label matching ability and uses meta-learning paradigm to learn few-shot instance summarizing ability. Experimental results on three different RE tasks show MCMN outperforms strong baseline models by large margins.

Acknowledgements

We thank all reviewers for their insightful suggestions. Moreover, this research work is supported by the Strategic Priority Research Program of Chinese Academy of Sciences under Grant No. XDA27020200, and the National Natural Science Foundation of China under Grants no. 62106251 and 62076233.

Ethics Consideration

This paper has no particular ethic consideration.

References

- Oshin Agarwal, Heming Ge, Siamak Shakeri, and Rami Al-Rfou. 2021. [Knowledge graph based synthetic corpus generation for knowledge-enhanced language model pre-training](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3554–3565, Online. Association for Computational Linguistics.
- 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, Florence, Italy. Association for Computational Linguistics.
- Jonathan Bragg, Arman Cohan, Kyle Lo, and Iz Beltagy. 2021. [FLEX: Unifying evaluation for few-shot NLP](#). In *Advances in Neural Information Processing Systems*.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). *CoRR*, abs/2005.14165.
- Alberto Cetoli. 2020. [Exploring the zero-shot limit of FewRel](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1447–1451, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Chih-Yao Chen and Cheng-Te Li. 2021. [ZS-BERT: Towards zero-shot relation extraction with attribute representation learning](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3470–3479, Online. Association for Computational Linguistics.
- Lei Cui, Furu Wei, and Ming Zhou. 2018. [Neural open information extraction](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 407–413, Melbourne, Australia. Association for Computational Linguistics.
- Manqing Dong, Chunguang Pan, and Zhipeng Luo. 2021. [Mapre: An effective semantic mapping approach for low-resource relation extraction](#).
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. [Model-agnostic meta-learning for fast adaptation of deep networks](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70

- of *Proceedings of Machine Learning Research*, pages 1126–1135. PMLR.
- Tianyu Gao, Xu Han, Zhiyuan Liu, and Maosong Sun. 2019a. [Hybrid attention-based prototypical networks for noisy few-shot relation classification](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Tianyu Gao, Xu Han, Hao Zhu, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. 2019b. [FewRel 2.0: Towards more challenging few-shot relation classification](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6250–6255, Hong Kong, China. Association for Computational Linguistics.
- Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017. [The WebNLG challenge: Generating text from RDF data](#). In *Proceedings of the 10th International Conference on Natural Language Generation*, pages 124–133, Santiago de Compostela, Spain. Association for Computational Linguistics.
- 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 the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4803–4809, Brussels, Belgium. Association for Computational Linguistics.
- Omer Levy, Minjoon Seo, Eunsol Choi, and Luke Zettlemoyer. 2017. [Zero-shot relation extraction via reading comprehension](#). In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 333–342, Vancouver, Canada. Association for Computational Linguistics.
- Fangchao Liu, Lingyong Yan, Hongyu Lin, Xianpei Han, and Le Sun. 2021. [Element intervention for open relation extraction](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4683–4693, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).
- Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. 2018. [A simple neural attentive meta-learner](#). In *International Conference on Learning Representations*.
- Tsendsuren Munkhdalai and Hong Yu. 2017. [Meta networks](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2554–2563. PMLR.
- Abiola Obamuyide and Andreas Vlachos. 2018. [Zero-shot relation classification as textual entailment](#). In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 72–78, Brussels, Belgium. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Oscar Sainz, Oier Lopez de Lacalle, Gorra Labaka, Ander Barrena, and Eneko Agirre. 2021. [Label verbalization and entailment for effective zero- and few-shot relation extraction](#).
- Victor Garcia Satorras and Joan Bruna Estrach. 2018. [Few-shot learning with graph neural networks](#). In *International Conference on Learning Representations*.
- Timo Schick and Hinrich Schütze. 2021. [Exploiting cloze-questions for few-shot text classification and natural language inference](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 255–269, Online. Association for Computational Linguistics.
- Jake Snell, Kevin Swersky, and Richard Zemel. 2017. [Prototypical networks for few-shot learning](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Richard Socher, Milind Ganjoo, Christopher D Manning, and Andrew Ng. 2013. [Zero-shot learning through cross-modal transfer](#). In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Oriol Vinyals, Charles Blundell, Timothy Lillicrap, koray kavukcuoglu, and Daan Wierstra. 2016. [Matching networks for one shot learning](#). In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 3630–3638. Curran Associates, Inc.
- Zhi-Xiu Ye and Zhen-Hua Ling. 2019. [Multi-level matching and aggregation network for few-shot relation classification](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2872–2881, Florence, Italy. Association for Computational Linguistics.

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, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.

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 the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP 2017)*, pages 35–45.

Peng Zhou, Wei Shi, Jun Tian, Zhenyu Qi, Bingchen Li, Hongwei Hao, and Bo Xu. 2016. [Attention-based bidirectional long short-term memory networks for relation classification](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 207–212, Berlin, Germany. Association for Computational Linguistics.