

Meta-Learning Adversarial Domain Adaptation Network for Few-Shot Text Classification

ChengCheng Han¹ Zeqiu Fan¹ Dongxiang Zhang² Minghui Qiu³
Ming Gao^{1*} Aoying Zhou¹

¹School of Data Science and Engineering, East China Normal University

²College of Computer Science and Technology, Zhejiang University

³Alibaba Group

{51195100009, 51195100007}@stu.ecnu.edu.cn

zhangdongxiang@zju.edu.cn

minghui.qmh@alibaba-inc.com

mgao@dase.ecnu.edu.cn

ayzhou@sei.ecnu.edu.cn

Abstract

Meta-learning has emerged as a trending technique to tackle few-shot text classification and achieved state-of-the-art performance. However, existing solutions heavily rely on the exploitation of lexical features and their distributional signatures on training data, while neglecting to strengthen the model's ability to adapt to new tasks. In this paper, we propose a novel meta-learning framework integrated with an adversarial domain adaptation network, aiming to improve the adaptive ability of the model and generate high-quality text embedding for new classes. Extensive experiments are conducted on four benchmark datasets and our method demonstrates clear superiority over the state-of-the-art models in all the datasets. In particular, the accuracy of 1-shot and 5-shot classification on the dataset of 20 Newsgroups is boosted from 52.1% to 59.6%, and from 68.3% to 77.8%, respectively¹.

1 Introduction

Few-shot text classification (Yu et al., 2018; Geng et al., 2019) is a task in which a model will be adapted to predict new classes not seen in training. For each of these new classes, we only have a few labeled examples. To be specific, we are given lots of training data with a set of classes \mathcal{Y}_{train} . After training, our goal is to get accurate classification results on the testing data with a set of new classes \mathcal{Y}_{test} , which is disjoint to \mathcal{Y}_{train} . Only a small labeled support set will be available in the testing

stage. If the support set contains K labeled examples for each of the N unique classes, we refer to the task as a N -way K -shot classification.

Existing approaches for few-shot text classification mainly fall into two categories: (1) transfer-learning based methods (Howard and Ruder, 2018; Pan et al., 2019; Gupta et al., 2020), which aim to transfer knowledge learned from a task to a new task or leverage general-domain pretraining and fine-tuning techniques for few-shot classification. (2) meta-learning based methods (Jamal et al., 2018; Yu et al., 2018; Geng et al., 2019, 2020; Bao et al., 2020), which aim to learn generic information (meta-knowledge) by recreating training episodes, so that it can classify new classes through only a few labeled examples. Among these methods, Bao et al. (2020) leveraged distributional signatures (e.g. word frequency and information entropy) to train a model within a meta-learning framework, and achieved state-of-the-art performance. However, the method pays more attention to statistical information and ignores other implicit information such as correlation between words. Furthermore, existing meta-learning methods heavily rely on the exploitation of lexical features and their distributional signatures on training data, while neglecting to strengthen the model's ability to adapt to new tasks.

In this paper, we propose an adversarial domain adaptation network to enhance meta-learning framework, with the objective of improving the model's adaptive ability for new tasks in new domains. We first utilize two neural networks competing against each other, separately playing the roles of a domain discriminator and a meta-knowledge

*Corresponding author

¹The source code of the paper is available at <https://github.com/hccngu/MLADA>.

generator. The adversarial network is able to strengthen the adaptability of the meta-learning architecture. Moreover, we aggregate transferable features generated by the meta-knowledge generator with sentence-specific features to produce high-quality sentence embeddings. Finally, we utilize a ridge regression classifier to obtain final classification results. To the best of our knowledge, we are the first to combine adversarial domain adaptation with meta-learning for few-shot text classification.

We evaluate our model on four popular datasets for few-shot text classification. Experimental results demonstrate that our method outperforms state-of-the-art models in all datasets, for both in 1-shot and 5-shot classification tasks. Especially on the 20 Newsgroups dataset, our model outperforms DS-FSL (Bao et al., 2020) by 7.5% in 1-shot classification and 9.5% in 5-shot classification. In addition, we conduct visualization analysis to verify the adaptability of our model and capability to recognize important lexical features for unseen classes.

2 Related Work

The mainstream approaches for few-shot text classification are based on meta-learning or transfer learning. In this section, we first briefly introduce the preliminary background of these two technologies, and then review how they are applied to support few-shot text classification.

Meta-learning Meta-learning, also known as “learning to learn”, refers to improving the learning ability of a model through multiple training episodes so that it can learn new tasks or adapt to new environments quickly with a few training examples. Existing approaches mainly fall into two categories: (1) Optimization-based methods, including developing a meta-learner as optimizer to output search steps for each learner directly (Andrychowicz et al., 2016; Ravi and Larochelle, 2017; Mishra et al., 2018; Gordon et al., 2019) and learning an optimized initialization of model parameters, which can be later adapted to new tasks by a few steps of gradient descent (Finn et al., 2017; Yoon et al., 2018; Grant et al., 2018; Bao et al., 2020). (2) Metric-based methods, including Matching Network (Vinyals et al., 2016), PROTO (Snell et al., 2017), Relation Network (Sung et al., 2018), TapNet (Yoon et al., 2019) and Induction Network (Geng et al., 2019), which aim to learn an appropriate distance metric to compare validation points

with training points and make prediction through matching training points.

Transfer learning Few-shot text classification relates closely to transfer learning (Zhuang et al., 2021) that aims to leverage knowledge from a related domain (a.k.a. source domain) to improve the learning performance and reduce the reliance on the number of labeled examples required in a target domain. Compared to meta-learning designed to aggregate the knowledge learned from many tasks, transfer learning typically involves a few tasks. In addition, we aim to directly reuse or fine-tune some existing representation in transfer learning, while a meta-learner is typically optimized at adapting to new tasks. Domain adaptation (Ganin et al., 2016; Tzeng et al., 2017; Khaddaj and Hajj, 2020) is a type of transfer learning, which aims to bridge the gap between the source and target domains by learning domain-invariant feature representations. Pre-trained model (Devlin et al., 2019; Yang et al., 2019; Brown et al., 2020) can also be viewed as a type of transfer learning. The parameters pre-trained in the source domain are fine-tuned in the target domain, with faster training convergence.

Few-shot text classification To tackle few-shot text classification, a straightforward idea is to apply BERT (Devlin et al., 2019) or XLNet (Yang et al., 2019), which have achieved strong performance in text classification by fine-tuning with a small number of training examples. Their performances can be less dependent on the number of training samples for the new classes. Some other approaches are based on transfer learning. Pan et al. (2019) proposed a modified hierarchical pooling strategy over pre-trained word embeddings to transfer knowledge obtained from some source domains to the target domain. Gupta et al. (2020) developed a binary classifier on the source domain to classify new classes by prefixing class identifiers to input texts.

Meta-learning (Jamal et al., 2018; Yu et al., 2018; Geng et al., 2019, 2020; Bao et al., 2020) can also be utilized to solve few-shot text classification, and has achieved state-of-the-art performance. Yu et al. (2018) proposed an adaptive metric learning approach that automatically determines the best weighted combination from meta-training tasks for few-shot tasks. Geng et al. (2019, 2020) leveraged the dynamic routing algorithm in meta-learning for few-shot text classification. (Bao et al., 2020) lever-

aged distributional signatures (e.g. word frequency and information entropy) to train a model within a meta-learning framework.

3 Method

In this section, we first present the preliminary background on episode-based meta-learning framework (Vinyals et al., 2016). After that, we explicitly describe the proposed MLADA (Meta-Learning Adversarial Domain Adaptation) Network.

3.1 Episode-based meta-learning

The goal of meta-training is to train a classifier that can learn meta-knowledge from training data. In this way, the classifier can quickly learn from a few annotations when classifying unseen classes. The “episode” training strategy that Vinyals et al. (2016) proposed has proved to be effective. The episode-based meta-learning consists of two main stages:

Meta-training Firstly, N classes are sampled from training data \mathcal{Y}_{train} . For each of these N classes, two subsets of examples are sampled separately as the support set S and the query set Q . Next, input the support set S and the query set Q to the model and update the parameters by minimizing the loss in the query set Q . The procedure above is called a training episode, which will be repeated multiple times.

Meta-testing After meta-training is finished, the performance of the model will be evaluated by the same episode-based mechanism. In a testing episode, N new classes will be sampled from \mathcal{Y}_{test} , which is disjoint to \mathcal{Y}_{train} . Then the support set and the query set will be sampled from the N classes. The model parameters can be fine-tuned through the small support set. The performance of the model will be evaluated through the average classification accuracy on the query set across all testing episodes.

We found that only a small subset of training data are accessible per training episode in the standard episode-based meta-training (Vinyals et al., 2016). To solve this problem, we build domain adversarial tasks to utilize more training data per training episode. Details of our model are described in the next section.

3.2 Meta-Learning Adversarial Domain Adaptation Network (MLADA)

Overview Our goal is to improve the performance of few-shot classification by combining adversarial domain adaptation and episode-based meta-learning. Figure 1 gives an overview of our model. In the rest of this section, we will introduce the main components of the model.

Word Representation Layer The goal of this layer is to represent each word with a d -dimensional vector. Following Bao et al. (2020), we construct the d -dimensional vector with the word embeddings, which is pre-trained with fast-Text (Joulin et al., 2016).

Domain Discriminator We refer to the support set and the query set as the target domain and the rest of the training data as the source domain. We sample a subset of examples from the source domain as the source set. The goal of this module is to distinguish whether the sample is from the source domain or the target domain. The discriminator is a three layer feed-forward neural network. We apply the *softmax* function in the output layer to evaluate the probability distribution $Pr(y|\lambda)$. $y = 0$ or 1 represents that the sample is from the query set or the source set.

Meta-knowledge Generator This module is mainly composed of a bi-directional LSTM (BiLSTM) and a fully connected layer. We utilize a BiLSTM to encode contextual embeddings for each time-step. The input of the module is a sequence of word vectors $P : [p_1, \dots, p_m]$, where m represents the number of words in a sentence. The output is a matrix $h_{d \times m}^p$, which is composed of contextual embeddings.

$$\vec{h}_i^p = LSTM(\vec{h}_{i-1}^p, p_i) \quad i = 1, \dots, m \quad (1)$$

$$\overleftarrow{h}_i^p = LSTM(\overleftarrow{h}_{i+1}^p, p_i) \quad i = m, \dots, 1 \quad (2)$$

$$h_i^p = Concat(\vec{h}_i^p, \overleftarrow{h}_i^p) \quad i = 1, \dots, m \quad (3)$$

$$h^p = [h_1^p, h_2^p, \dots, h_m^p] \quad (4)$$

Next, we employ a single layer feed-forward neural network and apply the *softmax* function to get the output k^p .

$$k^p = Softmax(\omega \cdot h^p + b) \quad (5)$$

k^p is an n -dimensional vector, which represents the meta-knowledge included in the sentence. n denotes the length of the sentence.

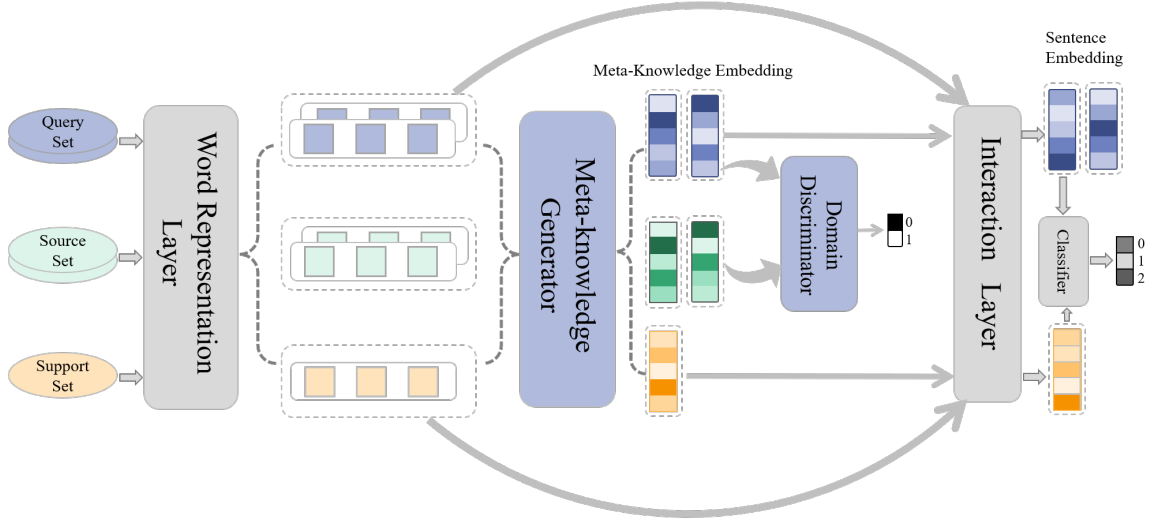


Figure 1: MLADA Network architecture for a N -way K -shot ($N = 3, K = 2$) problem

The goal of the meta-knowledge generator is not only to make the final classification results better, but also to confuse the domain discriminator as much as possible, so that the discriminator can not distinguish between samples from query set or source set. The theory on domain adaptation suggests that, for effective domain transfer to be achieved, predictions must be made based on features that cannot discriminate between the source domain and target domain, which is the motivation for us to build the meta-knowledge generator.

Interaction Layer We consider that the vector generated by the meta-knowledge generator is the transferable features, and word embeddings is the specific features of sentences. The role of the interaction layer is to fuse transferable features and sentence-specific features to produce the output as sentence embeddings, which will be used as the input of the classifier to obtain the final classification results. Suppose that the length of the sentence p is m , the word vectors is $w_i^p (i \in [1, m])$, the dimension of the word vector is d and the meta-knowledge of the sentence is k^p , then the final sentence vector is s^p :

$$s^p = W_{d \times m}^p \cdot k^p \quad (6)$$

where $W^p = [w_1^p, w_2^p, \dots, w_m^p]$.

Classifier The classifier is trained by the support set from scratch for each episode. We choose the *ridge regression* as the classifier. The reason

why we adopt the ridge regression to fit the support set are as follows: 1) If we choose neural networks as the classifier, it will be trained inadequately because the number of samples in the support set is too small. 2) The ridge regression admits a closed-form solution and it reduces over-fitting on the small support set through proper regularization. Specifically, we minimize regularized squared loss:

$$\mathcal{L}^{RR}(\theta) = \frac{1}{2m} \sum_{i=1}^m [((f_{\theta}(x^{(i)})) - y^{(i)})^2 + \lambda \sum_{j=1}^n \theta_j^2] \quad (7)$$

where m represents the number of samples in the support set, $f_{\theta}(x^{(i)})$ represents the prediction of the ridge regressor, $y^{(i)}$ represents the label of the sample, $\sum_{j=1}^n \theta_j^2$ denotes the squared Frobenius norm and $\lambda > 0$ controls the extent of the regularization.

Loss Function In each training episode, we first fix the parameters of the generator and the discriminator to update the classifier's parameters by the support set. The classifier's loss function is shown in Eq.7.

Next, we fix the parameters of the generator and the classifier to update the discriminator's parameters by the query set and the source set. We use the cross-entropy loss as the discriminator's loss function, which is shown in Eq.8.

$$\begin{aligned} \mathcal{L}^D(\mu) = & -\frac{1}{2m} \sum_{i=1}^{2m} [y_d^{(i)} \log D_{\mu}(k^{(i)}) \\ & + (1 - y_d^{(i)}) \log(1 - D_{\mu}(k^{(i)}))] \quad (8) \end{aligned}$$

Algorithm 1 MLADA Training Procedure

Input: Training data $\{\mathcal{X}_{train}, \mathcal{Y}_{train}\}$; T episodes and ep epochs; N classes in support set or query set; K samples in each class in the support set and L samples in each class in the query set; The generator’s parameters β , the discriminator’s parameters μ and the classifier’s parameters θ .

Output: Parameters β and μ after training;

```
1: Randomly initialize the model parameters  $\beta$ ,  
    $\mu$  and  $\theta$ ;  
2: for each  $i \in [1, ep]$  do  
3:    $\mathcal{Y} \leftarrow \Lambda(\mathcal{Y}_{train}, N)$ ;1  
4:   for each  $j \in [1, T]$  do  
5:      $S, Q, \Phi \leftarrow \emptyset, \emptyset, \emptyset$ ;  
6:     for  $y \in \mathcal{Y}$  do  
7:        $S \leftarrow S \cup \Lambda(\mathcal{X}_{train}\{y\}, K)$ ;2  
8:        $Q \leftarrow Q \cup \Lambda(\mathcal{X}_{train}\{y\} \setminus S, L)$ ;  
9:        $\Phi \leftarrow \Phi \cup \Lambda(\mathcal{X}_{train} \setminus \mathcal{X}_{train}\{y\}, L)$ ;  
10:    end for  
11:    Input  $S$  to the model;  
12:    Fix  $\mu, \beta$ . Update  $\theta$  by minimizing the  
    Eq.7;  
13:    Input  $Q, \Phi$  to the model;  
14:    Fix  $\beta, \theta$ . Update  $\mu$  by minimizing the  
    loss of the discriminator (Eq.8);  
15:    Fix  $\mu, \theta$ . Update  $\beta$  by minimizing the  
    loss of the generator (Eq.9);  
16:  end for  
17: end for
```

where μ denotes the parameters of the discriminator, m represents the number of samples of the query set or the source set. $y_d = 0$ or 1 denotes whether the sample is from the source set or the query set. k represents the meta-knowledge vector.

Finally, we fix the parameters of the discriminator and the classifier to update the generator’s parameters by the query set and the source set. The loss function of the generator is composed of two components. The first one is a cross-entropy loss for the final classification results, and the second one is the opposite of the discriminator’s loss, which is to confuse the discriminator.

$$\mathcal{L}^G(\beta) = CE_{Loss}(f(W \cdot G_\beta(W)), y) - \mathcal{L}^D \quad (9)$$

¹ $\Lambda(\mathcal{Y}, N)$ denotes selecting N elements from \mathcal{Y} randomly.

² $\mathcal{X}_{train}\{y\}$ denotes samples labeled y in \mathcal{X}_{train} .

where β represents the generator’s parameters. f denotes the ridge regressor. W represents the matrix of word vectors in a sentence. y denotes the real labels of samples. \mathcal{L}^D is shown in Eq.8.

Training Procedure It is remarkable that the meta-knowledge generator is optimized over all training episodes, while the classifier is trained from scratch for each episode. In each training episode, we first utilize the support set to update the parameters in the classifier. Next, we use the query set and source set to update the parameters of the meta-knowledge generator and the domain discriminator. The details of training procedure of our model are shown in Algorithm 1.

4 Experiments

In this section, we perform comprehensive experiments to compare our proposed model with five competitive baselines, and evaluate the performance on four text classification datasets.

4.1 Datasets

We use four benchmark datasets for text classification, whose statistics are summarized in Table 1.

HuffPost headlines contains 41 classes of news headlines from the year 2012 to 2018 obtained from HuffPost (Misra, 2018). Its text is less abundant (i.e., with smaller text length) than the other datasets and considered to be more challenging for text classification.

Amazon product data contains product reviews from 24 product categories, including 142.8 million reviews spanning 1996-2014 (He and McAuley, 2016). Our task is to identify the product categories of the reviews. Since the original dataset is proverbially large, we sample a subset of 1,000 reviews from each category.

Reuters-21578 is collected from Reuters newswire in 1987. We use the standard ApteMode version of the dataset. Following Bao et al. (2020), we consider 31 classes and remove multi-labeled articles. Each class contains at least 20 articles.

20 Newsgroups is a collection of approximately 20,000 newsgroup documents (Lang, 1995), partitioned (nearly) evenly across 20 different newsgroups.

4.2 Experiment Setup

Baselines We compare our MLADA with multiple competitive baselines, which are briefly summarized in the following:

Dataset	Avg. text length	vocab size	# samples	# train / val / test classes
HuffPost	11	8218	36900	20 / 5 / 16
Amazon	140	17062	24000	10 / 5 / 9
Reuters	168	2234	620	15 / 5 / 11
20 Newsgroups	340	32137	18820	8 / 5 / 7

Table 1: Statistics of the four benchmark datasets.

Method	HuffPost		Amazon		Reuters		20 News		Average	
	1 shot	5 shot	1 shot	5 shot	1 shot	5 shot	1 shot	5 shot	1 shot	5 shot
MAML(2017)	35.9	49.3	39.6	47.1	54.6	62.9	33.8	43.7	40.9	50.8
PROTO(2017)	35.7	41.3	37.6	52.1	59.6	66.9	37.8	45.3	42.7	51.4
Induct(2019)	38.7	49.1	34.9	41.3	59.4	67.9	28.7	33.3	40.4	47.9
HATT(2019)	41.1	56.3	49.1	66.0	43.2	56.2	44.2	55.0	44.4	58.4
DS-FSL(2020)	43.0	63.5	62.6	81.1	81.8	96.0	52.1	68.3	59.9	77.2
MLADA(ours)	45.0	64.9	68.4	86.0	82.3	96.7	59.6	77.8	63.9	81.4

Table 2: Mean accuracy (%) of 5-way 1-shot and 5-way 5-shot classification over four datasets.

- **MAML** (Finn et al., 2017) is trained by maximizing the sensitivity of the loss functions of new tasks, so that it can rapidly adapt to new tasks after the parameters have been up-dated through few gradient steps.
- **Prototypical Networks** (Snell et al., 2017), abbreviated as **PROTO**, is a metric-based method for few-shot classification by using sample averages as class prototypes.
- **Induction Networks** (Geng et al., 2019) learns a class-wise representation by leveraging the dynamic routing algorithm in meta-learning.
- **HATT** (Gao et al., 2019) extends **PROTO** by adding a hybrid attention mechanism to the prototypical network.
- **DS-FSL** (Bao et al., 2020) is trained within a meta-learning framework to map the distribution signatures into attention scores so as to extract more transferable features.

Implementation Details Following Bao et al. (2020), we use pre-trained fastText (Joulin et al., 2016) for word embedding. In the meta-knowledge generator, we use a BiLSTM with 128 hidden units. In the domain discriminator, the numbers of hidden units for the two feed-forward layers are set to 256 and 128, respectively. All parameters are

optimized using Adam with a learning rate of 0.001 (Kingma and Ba, 2015).

During meta-training, we perform 100 training episodes ($T = 100$) per epoch. Meanwhile, we apply early stopping when the accuracy on the validation set fails to improve for 20 epochs. We evaluate the model performance based on 1,000 testing episodes and report the average accuracy over 5 different random seeds. All the experiments are conducted on a NVIDIA v100 GPU.

4.3 Experimental Results

The experimental results are reported in Table 2. Our model achieves the best performance across all datasets, with an average accuracy of 63.9% in 1-shot classification and 81.4% in 5-shot classification, outperforming the state-of-the-art model DS-FSL (Bao et al., 2020) by a notable 4% improvement. For DS-FSL, it extracts transferable features via certain distribution signatures (e.g., word frequency or information entropy), but ignores other information of sentences, including implicit interaction between words. In contrast, we does not limit the transferable knowledge to statistical information. Our strategy is to combine the proposed domain adversarial network with meta-learning, generating more comprehensive transferable features.

Furthermore, our model improves dramatically 7.5% and 9.5% on 20 Newsgroups in 1-shot and

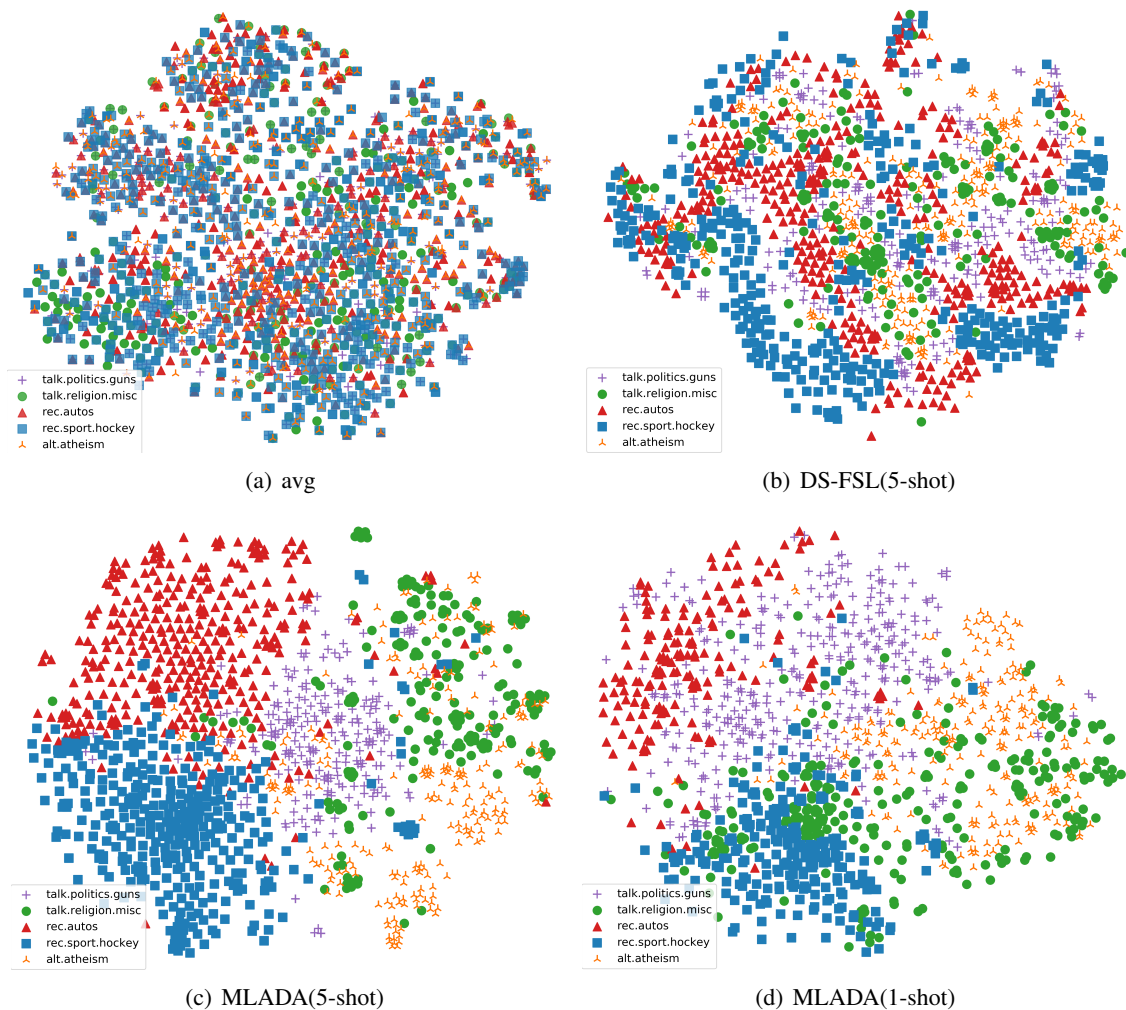


Figure 2: t-SNE visualization of the input representation of the classifier for a testing episode ($N = 5$, $K = 5$, $L = 500$) sampled from 20 Newsgroups. Note that the 5 classes is not seen in training set. The input representation of the classifier given by (a) the average of word embeddings (b) DS-FSL and (c) MLADA(ours). (d) is the t-SNE visualization of MLADA on 5-way 1-shot classification.

5-shot classification. The average length of texts in the 20 Newsgroups is longer than the other datasets. The empirical results clearly demonstrate that our model is more suitable for longer texts, which contain more abundant text information.

4.4 Ablation Study

We conduct an ablation study to examine the effectiveness of the proposed domain adversarial network as well as the interaction layer and the source set. The results of Amazon dataset are reported in Table 3.

Firstly, we use a bi-directional LSTM instead of the proposed domain adversarial network (including the meta-knowledge generator and the domain discriminator) for sentence encoding. The performances in the tasks of 1-shot classification and

5-shot classification decrease by 6.5% and 5.3%, respectively. This verifies the effectiveness of the proposed domain adversarial network.

Secondly, we study how the interaction layer contributes to the performance of our model. We concatenate the vector generated by the meta-knowledge generator directly with the average sentence embedding instead of the interaction layer. From the result in Table 3, we can see that our proposed interaction layer to combine the transferable features with the sentence-specific information are indeed more effective.

Finally, we remove the source set and utilize the discriminator to distinguish the true classes of samples. We observe that the source set is also important to performance. Due to the removal of the source set, the model has only access to the sup-

Seen classes	<i>Politics, Entertainment, Food&Drink, College, Arts</i>							Prediction
DS-FSL	Senate	committee	advances	bill to protect	Robert	Mueller	.	<i>politics</i> ✓
MLADA(ours)	Senate	committee	advances	bill to protect	Robert	Mueller	.	<i>politics</i> ✓
Unseen classes	<i>Sports, Education, Media, Tech, Environment</i>							Prediction
DS-FSL	Olympic	committee	CEO	resigns	cites	health	issues.	<i>environment</i> ✗
MLADA(ours)	Olympic	committee	CEO	resigns	cites	health	issues.	<i>sports</i> ✓

Figure 3: The visualization of attention weights generated by DS-FSL and the meta-knowledge generator of our model.

port set and the query set in each training episode. Therefore, it cannot learn cross-domain transferable features.

Models	Accuracy(%)	
	1 shot	5 shot
– Domain Adversarial Network	61.9	80.7
– Interaction Layer	66.6	83.0
– Source Set	67.1	84.2
MLADA	68.4	86.0

Table 3: Ablation study results of 5-way 1-shot and 5-way 5-shot classification on the Amazon dataset.

4.5 Visualization

We utilize visualization experiments to demonstrate that our model can generate high-quality sentence embeddings and identify important lexical features for unseen classes.

We first use t-SNE (Van der Maaten and Hinton, 2008) visualization of sentence embeddings generated by different methods on the query set, as shown in Figure 2. Compared to 2(a) average word embeddings and 2(b) DS-FSL, our method produces better separation both in 1-shot and 5-shot classification, demonstrating the effectiveness of MLADA in leveraging the supervised learning experience to generate high-quality sentence embeddings for few-shot text classification.

Moreover, we visualize the weight vectors generated by the meta-knowledge generator and compare it with DS-FSL, as shown in Figure 3. Our model reduces the weight of “committee” while increasing the weight of “Olympic”, which demonstrates that our model can recognize important lexical features in the new task, rather than simply transferring features obtained from experience.

5 Conclusion

In this paper, we propose a novel meta-learning approach called Meta-Learning Adversarial Domain Adaptation Network (MLADA), which can recognize important lexical features and generate high-quality sentence embeddings in new classes (not seen in training data). Specifically, we design an adversarial domain adaptation network in meta-training episodes, which aims to extract domain-invariant features and improve the adaptability of the meta-learner in new classes. We demonstrate that our method outperforms the existing state-of-the-art approaches on four standard text classification datasets. Future work includes applying MLADA to other fields including computer vision and speech recognition, and exploring the combination between adversarial domain adaptation network and other FSL algorithms.

6 Acknowledgments

This work has been supported by the National Key Research and Development Program of China under Grant 2016YFB1000905, the National Natural Science Foundation of China under Grant No. U1811264, U1911203, 61877018, 61672234, 61672384 and Alibaba Group through Alibaba Innovative Research Program.

References

- Marcin Andrychowicz, Misha Denil, Sergio Gomez Colmenarejo, Matthew W. Hoffman, David Pfau, Tom Schaul, and Nando de Freitas. 2016. [Learning to learn by gradient descent by gradient descent](#). In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 3981–3989.
- Yujia Bao, Menghua Wu, Shiyu Chang, and Regina Barzilay. 2020. [Few-shot text classification with distributional signatures](#). In *8th International Confer-*

- ence on Learning Representations, *ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- 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](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- 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, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 1126–1135. PMLR.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor S. Lempitsky. 2016. [Domain-adversarial training of neural networks](#). *J. Mach. Learn. Res.*, 17:59:1–59:35.
- Tianyu Gao, Xu Han, Zhiyuan Liu, and Maosong Sun. 2019. [Hybrid attention-based prototypical networks for noisy few-shot relation classification](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 6407–6414. AAAI Press.
- Ruiying Geng, Binhua Li, Yongbin Li, Jian Sun, and Xiaodan Zhu. 2020. [Dynamic memory induction networks for few-shot text classification](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 1087–1094. Association for Computational Linguistics.
- Ruiying Geng, Binhua Li, Yongbin Li, Xiaodan Zhu, Ping Jian, and Jian Sun. 2019. [Induction networks for few-shot text 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 2019, Hong Kong, China, November 3-7, 2019*, pages 3902–3911. Association for Computational Linguistics.
- Jonathan Gordon, John Bronskill, Matthias Bauer, Sebastian Nowozin, and Richard E. Turner. 2019. [Meta-learning probabilistic inference for prediction](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Erin Grant, Chelsea Finn, Sergey Levine, Trevor Darrell, and Thomas L. Griffiths. 2018. [Recasting gradient-based meta-learning as hierarchical bayes](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Aakriti Gupta, Kapil Thadani, and Neil O’Hare. 2020. [Effective few-shot classification with transfer learning](#). In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pages 1061–1066. International Committee on Computational Linguistics.
- Ruining He and Julian J. McAuley. 2016. [Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering](#). In *Proceedings of the 25th International Conference on World Wide Web, WWW 2016, Montreal, Canada, April 11 - 15, 2016*, pages 507–517. ACM.
- Jeremy Howard and Sebastian Ruder. 2018. [Universal language model fine-tuning for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 328–339. Association for Computational Linguistics.
- Muhammad Abdullah Jamal, Guo-Jun Qi, and Mubarak Shah. 2018. [Task-agnostic meta-learning for few-shot learning](#). *CoRR*, abs/1805.07722.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, Matthijs Douze, Hervé Jégou, and Tomáš Mikolov. 2016. [Fasttext.zip: Compressing text classification models](#). *CoRR*, abs/1612.03651.
- Alaa Khaddaj and Hazem M. Hajj. 2020. [Representation learning for improved generalization of adversarial domain adaptation with text classification](#). In *IEEE International Conference on Informatics, IoT, and Enabling Technologies, ICIoT 2020, Doha, Qatar, February 2-5, 2020*, pages 525–531. IEEE.

- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Ken Lang. 1995. [Newsweeder: Learning to filter net-news](#). In *Machine Learning, Proceedings of the Twelfth International Conference on Machine Learning, Tahoe City, California, USA, July 9-12, 1995*, pages 331–339. Morgan Kaufmann.
- Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(11).
- Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. 2018. [A simple neural attentive meta-learner](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Rishabh Misra. 2018. [News category dataset](#).
- Chongyu Pan, Jian Huang, Jianxing Gong, and Xingsheng Yuan. 2019. [Few-shot transfer learning for text classification with lightweight word embedding based models](#). *IEEE Access*, 7:53296–53304.
- Sachin Ravi and Hugo Larochelle. 2017. [Optimization as a model for few-shot learning](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Jake Snell, Kevin Swersky, and Richard S. Zemel. 2017. [Prototypical networks for few-shot learning](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 4077–4087.
- Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip H. S. Torr, and Timothy M. Hospedales. 2018. [Learning to compare: Relation network for few-shot learning](#). In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 1199–1208. IEEE Computer Society.
- Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. 2017. [Adversarial discriminative domain adaptation](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 2962–2971. IEEE Computer Society.
- Oriol Vinyals, Charles Blundell, Timothy Lillicrap, koray kavukcuoglu, and Daan Wierstra. 2016. [Matching networks for one shot learning](#). In *Advances in Neural Information Processing Systems*, volume 29, pages 3630–3638. Curran Associates, Inc.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. [Xlnet: Generalized autoregressive pretraining for language understanding](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 5754–5764.
- Jaesik Yoon, Taesup Kim, Ousmane Dia, Sungwoong Kim, Yoshua Bengio, and Sungjin Ahn. 2018. [Bayesian model-agnostic meta-learning](#). In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 7343–7353.
- Sung Whan Yoon, Jun Seo, and Jaekyun Moon. 2019. [Tapnet: Neural network augmented with task-adaptive projection for few-shot learning](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 7115–7123. PMLR.
- Mo Yu, Xiaoxiao Guo, Jinfeng Yi, Shiyu Chang, Saloni Potdar, Yu Cheng, Gerald Tesaro, 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, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 1206–1215. Association for Computational Linguistics.
- Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. 2021. [A comprehensive survey on transfer learning](#). *Proc. IEEE*, 109(1):43–76.