

# MCML: A Novel Memory-based Contrastive Meta-Learning Method for Few Shot Slot Tagging

Hongru Wang, Zezhong Wang, Wai Chung Kwan, Kam-Fai Wong  
MoE Key Laboratory of High Confidence Software Technologies  
The Chinese University of Hong Kong  
{hrwang, kfwong}@se.cuhk.edu.hk

## Abstract

Meta-learning is widely used for few-shot slot tagging in task of few-shot learning. The performance of existing methods is, however, seriously affected by *sample forgetting issue*, where the model forgets the historically learned meta-training tasks while solely relying on support sets when adapting to new tasks. To overcome this predicament, we propose the **Memory-based Contrastive Meta-Learning** (aka, MCML) method, including *learn-from-the-memory* and *adaption-from-the-memory* modules, which bridge the distribution gap between training episodes and between training and testing respectively. Specifically, the former uses an explicit memory bank to keep track of the label representations of previously trained episodes, with a contrastive constraint between the label representations in the current episode with the historical ones stored in the memory. In addition, the *adaption-from-memory* mechanism is introduced to learn more accurate and robust representations based on the shift between the same labels embedded in the testing episodes and memory. Experimental results show that the MCML outperforms several state-of-the-art methods on both SNIPS and NER datasets and demonstrates strong scalability with consistent improvement when the number of shots gets more.

## 1 Introduction

Slot tagging (Tur and De Mori, 2011), is a key part of natural language understanding, which is usually modeled as a sequence labeling problem with BIO format as shown in Figure 1 (Chen et al., 2019). However, rapid domain transfer and scarce labeled data in the target domain introduce new challenges (Bapna et al., 2017a; Zhang et al., 2020). To this end, significant efforts have been made to develop few-shot techniques (Li Fei-Fei et al., 2006; Snell et al., 2017; Vinyals et al., 2016), which aim to recognize a set of novel classes with only a few labeled samples (i.e, less than 50-shot) by knowledge

transfer from a set of base classes with abundant annotated samples.

Among several few-shot learning approaches (Hospedales et al., 2022), metric-based meta-learning has been widely used in slot tagging because they are model-agnostic, effective, and easily applicable (Snell et al., 2017; Vinyals et al., 2016; Zhu et al., 2020; Hou et al., 2020). To cope with the data scarcity of novel classes, metric-based methods split data into different episodes deliberately while each episode consists of one support set and one query set. The model classifies a (query) item according to its similarity with the representation of each label learned from the support set in this episode.

However, this kind of setting has shown several limitations, where the similarity calculation conducted only at the episode level hinders the learning of the original representations, resulting in the *sample forgetting problem* (Toneva et al., 2018). On the one hand, this cripples the model’s ability to learn consistent representations for the same labels across different episodes. Here, the same labels may occur during different episodes at the meta-training stage and also possibly span the meta-training and meta-testing stages. For example, *B-Object\_name* occurs in both the meta-training stage and meta-testing stage as shown in Figure 1. On the other hand, the similarity calculation is conducted between the query and support set only in one episode under the few-shot setting, resulting in the representation shift while ignoring the same label representation in previous episodes. First of all, the representation of each label is not accurate due to the limited labeled samples. Besides that, the locally closest label in one episode is not necessarily the globally closest. Furthermore, with the number of shots increasing, the sample forgetting problem becomes worse and the model performance saturates quickly (Cao et al., 2019).

To overcome the above limitations, we propose

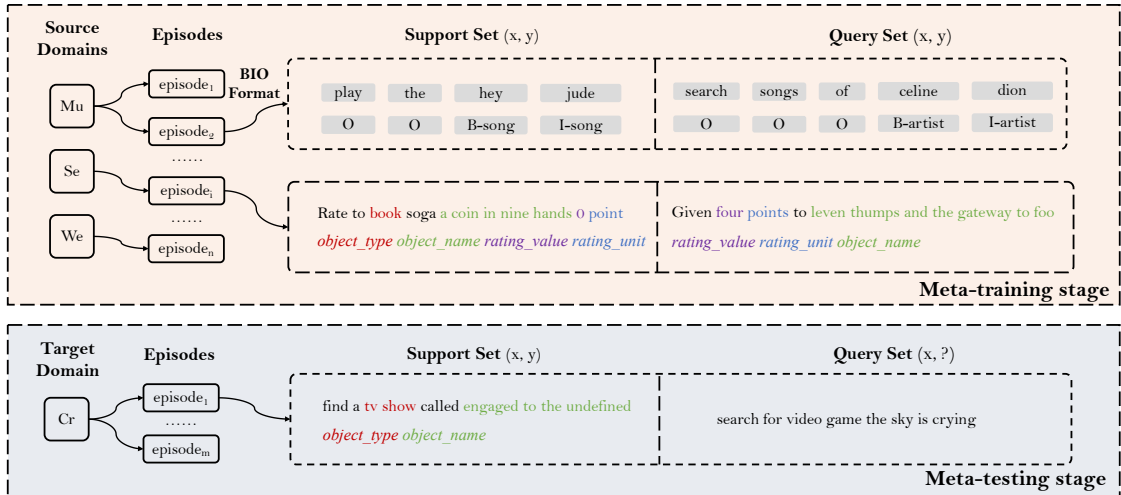


Figure 1: Episode-setting of Metric-based meta-learning, where each domain contains multiple episodes and each episode consists of a support-query-set pair. Different colors indicate different labels.

the Memory-based Contrastive Meta-learning (aka, MCML) method, marrying the benefits of *learn-from-the-memory* and *adaption-from-the-memory* to capture more transferable and informative label representations. Specifically, during meta-training, we use an explicit memory bank to keep track of the label representations from the historical episodes. Then a contrastive constraint is added to pull together semantically similar (i.e, positive) samples in the embedding space while pushing apart dissimilar (i.e, negative) samples. This is what we call the *learn-from-the-memory* technique. Secondly, during meta-testing, we use the *adaption-from-the-memory* technique to bridge the shift between the input labels embedded in the test episodes and the label anchors in the memory. In addition, an indicator is used to control how much information we want to acquire from the memory. The combination of *learn-from-the-memory* and *adaption-from-the-memory* helps the model to learn consistent representations for the same labels and distinguished representations for different labels concurrently across different episodes. To summarize, our contributions are three-fold:

- This is the first work to tackle the *sample forgetting problem* of metric-based methods. We propose a novel Memory-based Contrastive Meta-learning (MCML) method to bridge the gap between different episodes.
- We propose two model-agnostic methods including *learn-from-the-memory* and *adaption-*

from-the-memory, which can be applied in different stages separately. The combination of them achieves the best performance even with the number of shots increasing.

- The experimental results confirm the effectiveness of our model with very favorable performance over several state-of-the-art methods on both *SNIPS* and *NER* datasets.

## 2 Related Work

Few-shot learning was first proposed as a transfer method using a Bayesian approach on low-level visual features (Li Fei-Fei et al., 2006). Over the past few years, researchers have developed alternative techniques to build domain-specific modules for low-resource cross-domain natural language understanding (Bapna et al., 2017b; Lee and Jha, 2019; Fritzier et al., 2019; Shah et al., 2019). Most recent works have tried to model the transition possibility or similarity function between different labels with the metric-based meta-learning framework as backbone (Hou et al., 2020; Zhu et al., 2020; Wang et al., 2022). Nevertheless, episode-level relationships are still under-explored in previous works, except for a number of methods on image classification (Li et al., 2019; Sun et al., 2019; Ouali et al., 2020; Fei et al., 2021). Fei et al. (2021) proposed a novel method to learn more robust representations by sampling two episodes containing the same set of classes for meta-training while Ouali et al. (2020) used intra-episode spatial contrastive

learning (SCL) as an auxiliary pre-training objective to learn general-purpose visual embeddings for image classification.

Distinguishing from prior work, we first exploit the inter-episode relationship for natural language understanding by using an explicit memory bank. Most researchers choose to store the encoded contextual information in each meta episode under the few-shot setting (Kaiser et al., 2017; Cai et al., 2018). Another alternative method adopts parameterized memory network to implicitly save historical information (Geng et al., 2020). Our work also keeps in line with Momentum Contrast (MoCo) which utilizes an external memory module to store positive or negative samples for contrastive learning (He et al., 2020). Similarly, with a relatively large size of samples, unilateral representations from one episode in the metric-based methods can be alleviated.

### 3 Preliminaries

Before introducing our proposed framework, we provide the problem definition and an illustration of the basic framework of metric-based meta-learning to solve few-shot slot tagging in this section.

#### 3.1 Problem Definition

We denote each sentence  $x = (x_1, x_2, x_3, \dots, x_p)$  and the corresponding label  $y = (y_1, y_2, y_3, \dots, y_p)$ . Usually, we are provided with lots of labeled data (i.e.  $(x, y)$  pairs) of source domains  $D_s$ , and few-shot (less than 50) labeled data as well as plenty of unlabeled data in the target domain  $D_t$  under the few shot setting. We split the data as episodes  $e = (S, Q)$  in which  $S = \{x_i^j, y_i^j\}_{i=1 \dots N}^{j=1 \dots K}$  and  $Q = \{x^j, y^j\}_{j=1}^{|Q|}$  respectively.  $S$ , as the support set, contains  $K$  examples (K-shot) for each of  $N$  labels (N-way) while  $Q$  contains several unlabeled samples<sup>1</sup>. Thus, the few-shot model is trained based on many episodes  $E_{tr} = (e_1, e_2, e_3, \dots, e_n)$  initially. The trained model is then directly evaluated on the target domain  $E_{te} = (e_1, e_2, \dots, e_m)$ . The objective is formulated as follows:

$$y^* = \operatorname{argmax}_y p_\theta(y|x, S) \quad (1)$$

where  $\theta$  refers the parameters of the slot tagging mode, the  $(x, y)$  pair and the support set from the target domain,  $E_{tr}$  and  $E_{te}$  represent different episodes during meta-training and meta-testing respectively.

<sup>1</sup>We have labels during meta-training.

#### 3.2 Metric-based Meta Learning

Given an episode consisting of a support-query set pair, the basic idea of metric-based meta-learning (Zhu et al., 2020; Hou et al., 2020) is to classify an item (a sentence or token) in the query set based on its similarity with the representation of each label, which is learned from the few labeled data of the support set. Some representative works are matching network (Vinyals et al., 2016) and prototypical network (Snell et al., 2017). More specifically, given an input episode  $(S, Q)$  pair, the model encodes these two parts to get the sample vector and query vector respectively:

$$S, Q = \text{Encoder}(S, Q) \quad (2)$$

After that, various models can be used to extract label representations  $c_{y_i}$ . Take the prototypical network as an example, each prototype (label representation) is defined as the average vector of the embedded samples which have the same labels:

$$c_{y_i} = \frac{1}{N_{y_i}} \sum_{i=1}^N \sum_{j=1}^K I\{y_i^j = y_i\} s_i^j \quad (3)$$

while  $I$  is an indicator function which equals to True when  $y_i^j == y_i$  else False;  $s_i^j$  is the corresponding sample vector from  $S$ .

Lastly, we calculate the distance between the label representation and the sample vector from the query set. The most popular distance function is the dot product function which is defined as follows:

$$\text{SIM}(x_i, c_k) = x_i^T c_k \quad (4)$$

$$y = \text{Softmax}(\text{SIM}(x_i, c_k)) \quad (5)$$

The label of instance (i.e.  $x_i$ ) from the query set is the label whose embedding is closest with the instance vector (i.e.  $c_i$ ). This can be calculated through a softmax layer. However, in this way, the learned prototype of labels may lack general discriminative semantic features since it ( $c_i$ ) only needs to be closer to instance ( $x_i$ ) compared with other prototypes (i.e.  $c_j$  where  $j \neq i$ ) in the same episode without considering the global prototypes. Since the support set may only contain a few instances with the same label, the representation becomes imprecise and fragile (Ouali et al., 2020).

### 4 Model

In this section, we first illustrate the overview of our proposed framework (Section 4.1), and then we discuss how to learn and adaption from memory (Section 4.2 and 4.3) respectively.

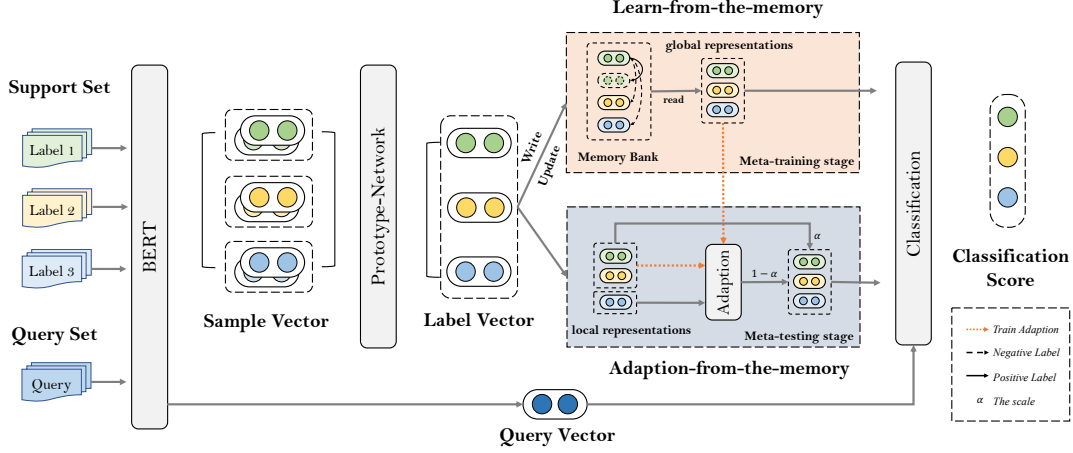


Figure 2: The overview of Our Framework, including three modules 1) Memory Bank, 2) Learn-from-the-memory, and 3) Adaption-from-the-memory, while the memory bank stores global representations across different episodes, learn-from-the-memory is used to learn more accurate and robust representations during the meta-training stage, and adaption-from-the-memory to bridge the gap between training and testing for local representations in test episode.

#### 4.1 Framework

Due to data scarcity and domain transfer, *sample forgetting problem* seriously hinders the model to learn robust representation, resulting in worse adaptability. To overcome this problem during meta-training and meta-testing stages, we propose *learn-from-the-memory* and *adaption-from-the-memory* techniques respectively as shown in Figure 2 to reuse the learned representations (Raghu et al., 2019).

**Learn-from-the-memory:** During the meta-training stage, the model will continuously train on different episodes. We utilize an external memory bank to store all learned label representations from the support set. These representations form different clusters naturally according to their original labels. When a newly seen label appears, a contrastive loss is computed on these dimensional representations by attracting positive samples, which have the same label, and by repelling the negative samples which have different labels. If this label has not been encountered before, we just write it into our memory.

**Adaption-from-the-memory:** During the meta-testing stage, we first learn an adaption layer by using these overlapped labels during meta-training and meta-testing, and then we use the learned adaption layer to project these unseen (i.e. not overlap) labels from testing space to training space in order to bridge the shift between testing space and training space. In addition, we use the skip connection to control how much information we want to

acquire from the memory.

#### 4.2 Learn from Memory

To consider all prototypes and learn consistent representations during meta-training, we first use a memory bank to store all prototypes of different labels from the support set. We design three basic operations in the memory bank: *write*, *update*, and *read*.

**Write.** Specifically, starting from the first episode  $e_1$  to the last episode  $e_n$  in  $E_{tr}$ , we store the label representations from the corresponding support set  $C_i = (c_1, c_2, \dots, c_k)$  into external memory bank  $M$  with the label name as key, where  $k$  is the number of labels for the current episode.  $M$  increases as the episode continue on.

$$\bar{k} \leq M \leq m * \bar{k} \quad (6)$$

while  $\bar{k}$  represents average number of labels for all episodes, and  $m$  is the number of episodes. For the  $i$ th episode, we first calculate the prototypical embedding of seen-label clusters from memory. Theoretically, this step is unnecessary but we choose to do so to save computational resources<sup>2</sup>.

$$c_k^* = \frac{1}{N_k} \sum_{i=1}^{N_k} I\{c_i = c_k\} c_i \quad (7)$$

<sup>2</sup>We call the prototypical embedding of label clusters by centroid (i.e. prototype in historical episodes), and prototype as average label embedding in one episode. Here if we skip the calculation of centroid, then we can directly use these prototypes.




Figure 3: The high-level overview of our *learn-from-the-memory*. There are multiple label representations in the memory bank. We will attract representations with the same label, repel otherwise, and also update accordingly.

We use  $c_k^*$  to represent the centroid of the  $k$ th cluster and we also store it in the memory bank, and then we define a distance function following (Ding et al., 2021) as follows:

$$d(c_i, c_j) = 1 / (1 + \exp(\frac{c_i}{\|c_i\|} \cdot \frac{c_j}{\|c_j\|})) \quad (8)$$

**Read.** For the coming labels and corresponding representations, there are two situations: (1) the label is new which means it never appears in the previous episodes, and (2) the label has already been stored in the memory bank. We extract all centroid representations from the memory bank and impose a contrastive learning constraint accordingly.

$$L_{memory} = -\frac{1}{K} \sum_{c_i \in S^c, c_j \in S^{-c}} [\log d(c_i, c_k^*) + \log(1 - d(c_j, c_k^*))] \quad (9)$$

For the new label, there are no positive pairs, and we increase the distances between its representation and all extracted representations. For the same label, we draw the same centroid representation but repel different centroids.

This objective effectively serves as regularization to learn more consistent and transferable label representation as they evolve during meta-training (Ding et al., 2021; He et al., 2020). We emphasize that the parameters of models do not change at this stage, and we do not need to modify the architecture of traditional metric-based meta-learning models. As such, the model can be easily trained together with other components in an end-to-end fashion.

**Update.** At the last, we need to re-calculate the prototypical embeddings of seen-label clusters in the memory following Equation 7. In this way, the




Figure 4: The high-level overview of our *adaption-from-memory*. We extract global representations from the memory bank which stores label representations during the meta-training stage and use them as grounded truth to learn the adaption layer (aka, MLP) with label representations from the current episode (aka, local representations) as input. We pass all labels, including overlap labels and new labels, from the current episode to the adaption layer to get the final representations.

distribution shift across different episodes during meta-training will be alleviated and thus more general discriminative representations can be learned. Figure 3 demonstrates the whole processing of *learn-from-the-memory*.

### 4.3 Adaption from Memory

To address the forgetting problem during the meta-testing stage, we take advantage of stored representations in the memory bank to build a bridge connecting the testing space and training space. With overlapped labels between meta-training and meta-testing, two types of representations can be observed: 1) one from memory during meta-training; 2) the other from the current episode during meta-testing. It is noted that labels overlap frequently in practice, e.g. *B-person* and *B-city* almost appear in every slot tagging dataset.

As shown in Figure 4, we decompose the whole process into two steps. First of all, we use the overlapped labels during meta-training and meta-testing to learn the adaption function  $f$  which minimizes the representation gap between meta-training and meta-testing<sup>3</sup>.

$$y_i = f(R_{test\_overlap}^i) \quad (10)$$

Where  $f$  can be implemented by Multilayer Perceptron (MLP) or one linear layer with the following loss function. Here  $R_{test\_overlap}^i$  means  $i$ th

<sup>3</sup>We emphasize this operation is conducted at episode-level to comply with the few-shot setting.

label appears in both the memory bank and current test episode. The learning objective of adaption layers  $L_{ada}$  can be defined as follows:

$$L_{ada} = \sum_{i=1}^{|\text{Overlap}|} \|R_{train\_overlap}^i - y_i\|^2 \quad (11)$$

Where  $R_{train\_overlap}^i$  here can be directly extracted from the memory bank as the ground truth representation. We then use the learned adaption function to project the new labels (i.e, not overlap) to the training space based on the assumption that the training space should be more accurate than the testing space which consists of more labeled data.

$$R_{train\_new} = f(R_{test\_new}) \quad (12)$$

In this case, we can get original representations in the testing space and representations in the training space after adaption for both overlap labels and new labels. Our final representation for each label can be the combination of these two kinds of representation from different spaces.

$$R_{fin} = \alpha * R_{ori} + (1 - \alpha) * R_{ada} \quad (13)$$

where  $\alpha \in (0, 1)$  is a hyper-parameter that controls the percentage of information from the original testing space and from adaption. By adaption from the memory, the distribution shift in the testing episodes rooting in domain transfer and few-shot setting will be de-biased by the global representations in the memory.

#### 4.4 Training Objective

The learning objective of our methods is the sum of three parts. It is noted that these losses are not optimized simultaneously.

$$L_{ner} = \sum_{c=1}^K y_c \log(P_c) \quad (14)$$

$$L = L_{ner} + L_{memory} + L_{ada} \quad (15)$$

while  $L_{ner}$  represents the traditional cross-entropy loss of sequence labeling (see Eq. 14) and is optimized with  $L_{memory}$  during training (see Eq. 9).  $L_{ada}$  is optimized during testing (see Eq. 11).

## 5 Experiments

### 5.1 Datasets

We evaluate the proposed methods following the data split setting provided by (Hou et al., 2020) on NER and SNIPS datasets (Coucke et al., 2018). It is in the episode data setting (Vinyals et al., 2016),

where each episode contains a support set (1-shot or 5-shot) and a batch of labeled samples. For NER, we followed the setting as same as (Zhu et al., 2020), which contains 4 different datasets: CoNLL-2003 (i.e. News) (Tjong Kim Sang and De Meulder, 2003), GUM (i.e. Wiki) (Zeldes, 2017), WNUT-2017 (i.e. Social) (Derczynski et al., 2017) and OntoNotes (i.e. Mixed) (Pradhan et al., 2013). For SNIPS, it consists of 7 domains with different label sets: Weather (We), Music (Mu), PlayList (Pl), Book (Bo), Search Screen (Se), Restaurant (Re) and Creative Work (Cr). And also, we extend our method to more shots (10-shot and 20-shot) to further demonstrate the effectiveness and robust generalization capability of our approach.

### 5.2 Baselines

**SimBERT.** It assigns labels to words according to the cosine similarity of word embedding of a fixed BERT. For each word  $x_i$ , SimBERT finds the most similar word  $x_k$  in the support set and assigns  $x_k$ 's label to  $x_i$ .

**TransferBERT.** It directly transfers the knowledge from the source domain to the target domain by parameter sharing. We train it on the source domain and select the best model on the same validation set of our model. Before evaluation, we fine-tune it on the target domain support set.

**L-TapNet+CDT+PWE** (Hou et al., 2020) one of the strong baselines for few-shot slot tagging, which enhances the WarmProtoZero(WPZ) (Fritzier et al., 2019) model with label name representation and incorporate it into the proposed CRF framework.

**L-ProtoNet+CDT+VPB** (Zhu et al., 2020) current state-of-the-art metric-based meta-learning, which investigates the different distance functions and utilizes the distance function VPB to boost the performance of the model.

**Coach** (Liu et al., 2020) Coarse-to-fine approach (Coach) for cross-domain slot filling, which is a current state-of-the-art few-shot fine-tuning method incorporating template regular loss and slot description information.

### 5.3 Implementation Details

We take the pre-trained uncased BERT-Base (Devlin et al., 2019) as an encoder to embed words into contextually related vectors in all experiments. Following the setting in (Zhu et al., 2020), we use ADAM (Kingma and Ba, 2015) to train the model with a learning rate of 1e-5, a weight decay of

N-shot	Model	Domain							Avg.
		We	Mu	Pl	Bo	Se	Re	Cr	
1-shot	SimBERT	36.10	37.08	35.11	68.09	41.61	42.82	23.91	40.67
	TransferBERT	55.82	38.01	45.65	31.63	21.96	41.79	38.53	39.06
	L-TapNet+CDT+PWE	71.53	<b>60.56</b>	66.27	84.54	76.27	70.79	62.89	70.41
	L-ProtoNet+CDT+VPB	<b>73.08</b>	58.50	68.81	82.41	75.88	<b>73.17</b>	70.27	71.73
	Coach (Liu et al., 2020)	55.81	38.72	41.60	41.44	35.25	54.38	47.74	44.99
	Ours	72.30	58.33	<b>69.64</b>	82.90	<b>77.23</b>	72.79	<b>79.57</b>	<b>73.25</b>
5-shot	SimBERT	53.46	54.13	42.81	75.54	57.10	55.30	32.38	52.96
	TransferBERT	59.41	42.00	46.07	20.74	28.20	67.75	58.61	46.11
	L-TapNet+CDT+PWE	71.64	67.16	75.88	84.38	82.58	70.05	73.41	75.01
	L-ProtoNet+CDT+VPB	<b>82.54</b>	69.52	80.45	91.03	86.14	80.75	75.95	80.91
	Coach (Liu et al., 2020)	73.56	45.85	47.23	61.61	65.82	69.99	57.28	60.19
	Ours	81.79	<b>69.70</b>	<b>80.78</b>	<b>91.53</b>	<b>87.09</b>	<b>82.49</b>	<b>81.07</b>	<b>82.06</b>

Table 1:  $F_1$  scores on few-shot slot tagging of the SNIPS dataset

Model	1-shot					5-shot				
	News.	Wiki	Social	Mixed	Avg.	News	Wiki	Social	Mixed	Avg.
SimBERT	19.22	6.91	5.18	13.99	11.32	32.01	10.63	8.20	21.14	18.00
TransferBERT	4.75	0.57	2.71	3.46	2.87	15.36	3.62	11.08	35.49	16.39
L-TapNet+CDT+PWE	44.30	12.04	20.80	15.17	23.08	45.35	11.65	23.30	20.95	25.31
L-ProtoNet+CDT+VPB	42.23	11.36	<b>27.72</b>	<b>31.17</b>	28.10	56.30	19.17	34.95	<b>43.30</b>	38.43
Ours	<b>42.70</b>	<b>13.20</b>	26.75	29.86	<b>28.13</b>	<b>56.89</b>	<b>22.09</b>	<b>35.27</b>	42.08	<b>39.08</b>

Table 2:  $F_1$  scores on few-shot slot tagging of the NER dataset

5e-5. And we set the distance function as  $\text{VPB}$  (Zhu et al., 2020). To prevent the impact of randomness, we test each experiment 10 times with different random seeds following (Hou et al., 2020). For adaption from memory, we set the iteration as 1000, and  $\alpha$  from [0.1, 0.3, 0.5, 0.7, 0.9] and report the best result.

## 5.4 Main Result

Table 1 and Table 2 show the results of both 1-shot and 5-shot slot tagging of SNIPS and NER datasets respectively. Our method reaches comparable results with the state-of-the-art and outperforms in 3 out of 7 domains under 1-shot setting, and 6 under 5-shot setting at SNIPS dataset. Specifically, our method achieves about 13% ( $70.27 \rightarrow 79.57$ ) and 7% ( $75.95 \rightarrow 81.07$ ) improvements in the Cr domain under 1-shot and 5-shot respectively. Besides that, the improvement keeps consistent on the NER dataset while adding additional shots leads to greater improvement. It’s obvious that our method demonstrates strong scalability and flexibility with the number of shots increasing.

When comparing Coach (Liu et al., 2020) with L-TapNet+CDT+PWE (Hou et al., 2020) and L-TapNet+CDT+VPB (Zhu et al., 2020), it is also interesting to see that fine-tuning is not as competitive as metric-based approaches when the shot is smaller.

## 6 Ablation Study and Analysis

### 6.1 Ablation Result

We borrow the result from Zhu et al. (2020) as baseline (i.e. L-ProtoNet+CDT+VPB) here since it reaches the best performance out of all baselines. Table 3 shows the ablation study of learning and adaption from memory. Comparing the result between 1-shot, 5-shot, 10-shot, and 20-shot, we find that the *learn-from-the-memory* (i.e. M) module gets more important as the number of shots increases. We attribute this phenomenon to the more transferable representations due to more labeled data brought by more shots. However, the *adaption-from-the-memory* cannot keep consistent improvement, we think this is caused by noise introduced by the adaption layer. After combining

	1-shot			5-shot			10-shot			20-shot		
	B	A	M	B	A	M	B	A	M	B	A	M
<b>We</b>	73.08	72.30	71.83	82.54	81.13	81.79	79.09	78.75	<b>79.12</b>	82.06	80.92	<b>82.79</b>
<b>Mu</b>	58.50	56.58	58.33	69.52	67.95	<b>69.70</b>	65.71	64.75	<b>66.65</b>	68.94	67.47	<b>70.03</b>
<b>Pl</b>	68.81	<b>69.64</b>	68.16	80.45	<b>80.78</b>	79.62	74.43	<b>75.08</b>	<b>77.89</b>	75.90	<b>76.16</b>	<b>77.09</b>
<b>Bo</b>	82.41	81.95	<b>82.90</b>	91.03	89.99	<b>91.53</b>	88.38	87.31	<b>89.65</b>	89.10	88.16	<b>90.94</b>
<b>Se</b>	75.88	<b>77.23</b>	74.45	86.41	86.35	86.95	86.96	<b>87.37</b>	<b>87.70</b>	88.33	88.08	<b>88.48</b>
<b>Re</b>	73.17	71.64	72.79	80.75	78.21	<b>82.49</b>	77.06	74.95	<b>78.00</b>	79.32	76.90	<b>80.31</b>
<b>Cr</b>	70.27	<b>79.57</b>	70.77	75.95	<b>81.07</b>	76.61	80.82	<b>84.91</b>	77.31	77.37	<b>82.02</b>	75.88
<b>Avg.</b>	71.73	<b>72.70</b>	71.32	80.91	80.78	<b>81.24</b>	78.92	79.02	<b>79.96</b>	80.15	79.96	<b>80.79</b>

Table 3: Ablation Study of *adaption-from-the-memory* and *learn-from-the-memory* on 1-shot, 5-shot, 10-shot and 20-shot respectively on SNIPS dataset. B, A, and M stand for the strongest baseline *L-ProtoNet+CDT+VPB*, only adaption-from-memory, and only learn-from-memory respectively.

these two modules, the model can reach the best performance as reported in Table 1 and Table 2. Compared with the strongest baseline, the averaged F1 score further improved (More analysis can be found in Appendix A).

## 6.2 t-SNE Visualization Analysis

We present a *t-SNE* visualization of label representations of trained metric-based meta-learning methods as shown in Figure 5 and we additionally draw the *t-SNE* visualization of label representations after adding contrastive learning constraint in Figure 7. On the one hand, it is observed from Figure 5 that: 1) the representations of *B-object\_type* and *I-object\_type* at the meta-training stage are separated into distant groups; and 2) the representations at the meta-testing stage are shifted compared with those at the meta-training stage. For the first observation, we can conclude that the model can not remember what it already learned, failing to capture a consistent representation of the same label. A similar problem still happens at the meta-testing stage due to the presence of poorly sampled shots (Fei et al., 2021). On the other hand, in Figure 7, it is found that the distance between the representations of *B-object\_type* (also *I-object\_type*) during the meta-training stage is much closer, which proves the effectiveness of learn-from-the-memory to alleviate the sample forgetting problem.

## 6.3 The Impact of different value of scale

To investigate the effects of  $\alpha$  during adaption-from-the-memory, we report the performance of different assignments of this scale. The result can be found in Figure 6. Since the larger the value, the more information from the meta-testing space,




Figure 5: The t-SNE visualization of label representations during meta-training and meta-testing learned by Metric-based Meta-learning Model (Target Domain: Cr 5-shot).

the less information from adaption. Thus, as long as the graph is monotonically increasing, the less useful the adaption is. However, as we can observe, it is obvious that not all domains show this trend. Specifically, under the 1-shot setting, "Pl", "Se" and "Cr" gets higher performance because of the adaption, and "Pl" and "Cr" continue this trend in the 5-shot. This shows the adaption layer is highly domain-sensitive and prefers the domain which has more overlapped labels. More analysis can be found in Appendix A.2.

## 7 Conclusion

In this paper, we address the *sample forgetting problem* during meta-training and meta-testing stages in






Figure 6: The impact of different scales on the performance, we suggest zooming in to check more details.




Figure 7: The t-SNE visualization of label representations with our proposed learn-from-the-memory (Target Domain: Cr 5-shot).

the metric-based meta-learning framework by capturing more transferable and informative label representations. To this end, we propose the Memory-based Contrastive Meta-learning (MCML) method, which consists of two modules: *learn-from-the-memory* and *adaption-from-the-memory* to function at different stages. Experimental results on both NER and SNIPS datasets demonstrate the advantages of our MCML framework in terms of scalability and robustness.

## Limitations

This paper tackles the issues of the sample forgetting problem in the metric-based meta-learning framework. We mainly focus on the few-shot slot tagging tasks but our proposed method is motivated by the unique setting of metric-based meta-learning

which can be applied to other text classification tasks such as intent detection or news classification. We left this in our future work.

## Acknowledgement

We thank all reviewers for their insightful comments and suggestions. This research work is partially supported by ITF Project No. PRP/054/21FX and CUHK under Project No. 3230366.

## References

- Ankur Bapna, Gokhan Tur, Dilek Hakkani-Tur, and Larry Heck. 2017a. Towards zero-shot frame semantic parsing for domain scaling. *arXiv preprint arXiv:1707.02363*.
- Ankur Bapna, Gokhan Tur, Dilek Hakkani-Tur, and Larry Heck. 2017b. Towards zero-shot frame semantic parsing for domain scaling. *arXiv preprint arXiv:1707.02363*.
- Qi Cai, Yingwei Pan, Ting Yao, Chenggang Yan, and Tao Mei. 2018. [Memory matching networks for one-shot image recognition](#). In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 4080–4088. IEEE Computer Society.
- Tianshi Cao, Marc Law, and Sanja Fidler. 2019. A theoretical analysis of the number of shots in few-shot learning. *arXiv preprint arXiv:1909.11722*.
- Qian Chen, Zhu Zhuo, and Wen Wang. 2019. Bert for joint intent classification and slot filling. *arXiv preprint arXiv:1902.10909*.
- Alice Coucke, Alaa Saade, Adrien Ball, Théodore Bluche, Alexandre Caulier, David Leroy, Clément Doumouro, Thibault Gisselbrecht, Francesco Calta-girone, Thibaut Lavril, et al. 2018. Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces. *arXiv preprint arXiv:1805.10190*.

- Leon Derczynski, Eric Nichols, Marieke van Erp, and Nut Limsopatham. 2017. [Results of the WNUT2017 shared task on novel and emerging entity recognition](#). In *Proceedings of the 3rd Workshop on Noisy User-generated Text*, pages 140–147, Copenhagen, Denmark. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ning Ding, Xiaobin Wang, Yao Fu, Guangwei Xu, Rui Wang, Pengjun Xie, Ying Shen, Fei Huang, Hai-Tao Zheng, and Rui Zhang. 2021. Prototypical representation learning for relation extraction. *arXiv preprint arXiv:2103.11647*.
- Nanyi Fei, Zhiwu Lu, Tao Xiang, and Songfang Huang. 2021. Melr: Meta-learning via modeling episode-level relationships for few-shot learning. In *Proc. Int. Conf. Learn. Represent.*, pages 1–20.
- Alexander Fritzier, Varvara Logacheva, and Maksim Kretov. 2019. [Few-shot classification in named entity recognition task](#). *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing*.
- 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*, pages 1087–1094, Online. Association for Computational Linguistics.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738.
- Timothy Hospedales, Antreas Antoniou, Paul Micaelli, and Amos Storkey. 2022. [Meta-learning in neural networks: A survey](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9):5149–5169.
- Yutai Hou, Wanxiang Che, Yongkui Lai, Zhihan Zhou, Yijia Liu, Han Liu, and Ting Liu. 2020. [Few-shot slot tagging with collapsed dependency transfer and label-enhanced task-adaptive projection network](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1381–1393, Online. Association for Computational Linguistics.
- Lukasz Kaiser, Ofir Nachum, Aurko Roy, and Samy Bengio. 2017. [Learning to remember rare events](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- 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*.
- Sungjin Lee and Rahul Jha. 2019. [Zero-shot adaptive transfer for conversational language understanding](#). 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 6642–6649. AAAI Press.
- Huaiyu Li, Weiming Dong, Xing Mei, Chongyang Ma, Feiyue Huang, and Bao-Gang Hu. 2019. Lgm-net: Learning to generate matching networks for few-shot learning. In *International conference on machine learning*, pages 3825–3834. PMLR.
- Li Fei-Fei, R. Fergus, and P. Perona. 2006. [One-shot learning of object categories](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(4):594–611.
- Zihan Liu, Genta Indra Winata, Peng Xu, and Pascale Fung. 2020. [Coach: A coarse-to-fine approach for cross-domain slot filling](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 19–25, Online. Association for Computational Linguistics.
- Yassine Ouali, Céline Hudelot, and Myriam Tami. 2020. Spatial contrastive learning for few-shot classification. *arXiv preprint arXiv:2012.13831*.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Hwee Tou Ng, Anders Björkelund, Olga Uryupina, Yuchen Zhang, and Zhi Zhong. 2013. [Towards robust linguistic analysis using OntoNotes](#). In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, pages 143–152, Sofia, Bulgaria. Association for Computational Linguistics.
- Aniruddh Raghu, Maithra Raghu, Samy Bengio, and Oriol Vinyals. 2019. Rapid learning or feature reuse? towards understanding the effectiveness of maml. *arXiv preprint arXiv:1909.09157*.
- Darsh Shah, Raghav Gupta, Amir Fayazi, and Dilek Hakkani-Tur. 2019. [Robust zero-shot cross-domain slot filling with example values](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5484–5490, Florence, Italy. Association for Computational Linguistics.
- 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.

- Qianru Sun, Yaoyao Liu, Tat-Seng Chua, and Bernt Schiele. 2019. Meta-transfer learning for few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 403–412.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition](#). In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Mariya Toneva, Alessandro Sordani, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J Gordon. 2018. An empirical study of example forgetting during deep neural network learning. *arXiv preprint arXiv:1812.05159*.
- Gokhan Tur and Renato De Mori. 2011. *Spoken language understanding: Systems for extracting semantic information from speech*. John Wiley & Sons.
- Oriol Vinyals, Charles Blundell, Tim Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. 2016. [Matching networks for one shot learning](#). In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 3630–3638.
- Zezhong Wang, Hongru Wang, Wai Chung Kwan, and Kam-Fai Wong. 2022. [Prior omission of dissimilar source domain\(s\) for cost-effective few-shot learning](#). In *Proceedings of the 5th International Conference on Natural Language and Speech Processing (IC-NLSP 2022)*, pages 30–39, Trento, Italy. Association for Computational Linguistics.
- Amir Zeldes. 2017. [The gum corpus: Creating multi-layer resources in the classroom](#). *Lang. Resour. Eval.*, 51(3):581–612.
- Tao Zhang, Congying Xia, Chun-Ta Lu, and Philip Yu. 2020. [MZET: Memory augmented zero-shot fine-grained named entity typing](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 77–87, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Su Zhu, Ruisheng Cao, Lu Chen, and Kai Yu. 2020. Vector projection network for few-shot slot tagging in natural language understanding. *arXiv preprint arXiv:2009.09568*.

## A Analysis

### A.1 Less-shot or More-shot ?

Table 3 shows the result of 10-shot and 20-shot on the SNIPS dataset which is generated following the method proposed by Hou et al. (2020).

**More Shots.** Compare 10-shot with 20-shot, we can find that all domains are improved with the help of *learn-from-the-memory* when the number of shots increases except “SearchCreativeWork”. Since this is the only domain which has 100% overlap labels during meta-training and meta-testing, we attribute this phenomenon caused by poor representations from meta-testing without *adaption-from-memory*.

**Fewer Shot v.s More Shot.** Compare 1-shot and 5-shot (less-shot) with 10-shot and 20-shot (more-shot), there are some interesting findings: 1) *learn-from-the-memory* can boost 6 out of 7 domains in more-shot instead of 3 in less-shot. This demonstrates the importance and effectiveness of this module when the number of shots gets more; 2) *adaption-from-memory* shows exactly the same gains whether or not there are more shots. This is reasonable since the number of shots does not affect the number of labels, and also the accuracy of adaption. We conclude that *learn-from-the-memory* is always worth trying, and *adaption-from-the-memory* highly depends on a specific domain.

### A.2 The Impact of overlap

The performance of the adaption function highly depends on the number of overlap labels. Since the more overlap labels between training and testing, we will get a more accurate adaption function. Figure 8 shows the percentage of overlap labels between training data and validation or test data.

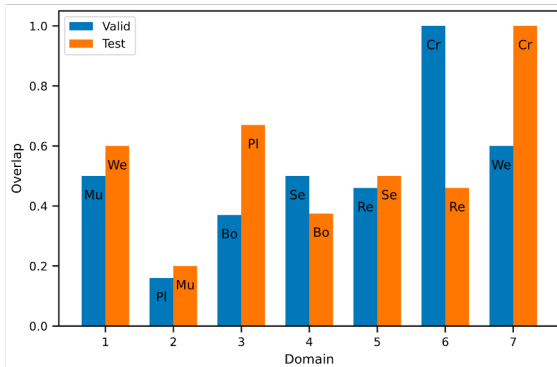


Figure 8: The Percentage of Overlap Labels between train and valid or test

To further investigate to what extent the influence of overlap labels on adaption performance, we utilize the Pearson correlation coefficient to analyze the relationship between these two variables. The calculated result is 0.83 which shows these two variables are highly related.




Figure 9: Improvement of different shots with different percentage overlapped labels.

The performance improved by the adaption of different domains can be found in Figure 9. It is noted that although "GetWeather" domain has 60% overlap labels with training data, the performance declines surprisingly. We further investigate the specific overlap labels of this domain, and we find most of them are "state", "country" and "city", common regular entity types which appear in almost every corpus. When the number of shots is less, these common entities cannot be represented accurately during meta-training, much less during adaption. This explains the poor performance of 1-shot and 5-shot and the higher performance of 10-shot and 20-shot. For the above reasons, we argue it is worth trying adaption once the overlap exceeds 50% as long as the overlaps labels have some domain-specific features.