

A Two-Stage Framework with Self-Supervised Distillation for Cross-Domain Text Classification

Yunlong Feng, Bohan Li, Libo Qin, Xiao Xu, Wanxiang Che*

Research Center for Social Computing and Information Retrieval
Harbin Institute of Technology, China
{ylfeng,bhli,lbqin,xxu,car}@ir.hit.edu.cn

Abstract

Cross-domain text classification is a crucial task as it enables models to adapt to a target domain that lacks labeled data. It leverages or reuses rich labeled data from the different but related source domain(s) and unlabeled data from the target domain. To this end, previous work focuses on either extracting domain-invariant features or task-agnostic features, ignoring domain-aware features that may be present in the target domain and could be useful for the downstream task. In this paper, we propose a two-stage framework for cross-domain text classification. In the first stage, we finetune the model with mask language modeling (MLM) to learn from the source domain. In the second stage, we further fine-tune the model with *self-supervised distillation* (SSD) and unlabeled data to adapt to the target domain. We evaluate its performance on a public cross-domain text classification benchmark and the experiment results show that our method achieves new state-of-the-art results for both single-source domain adaptations (94.17% \uparrow 1.03%) and multi-source domain adaptations (95.09% \uparrow 1.34%).

Keywords: Cross-domain text classification, Unsupervised domain adaptation

1. Introduction

In the era of large models, neural network models have achieved remarkable results in a myriad of tasks. However, a prevalent challenge arises when these models, often trained on source domains, are deployed in different, target domains, leading to a domain shift (Gretton et al., 2006). Unsupervised Domain Adaptation (UDA) emerges as a vital solution by aiming to adapt the models trained on source domains with labeled data to a target domain laden with unlabeled data. The significance of UDA is pronounced in the age of large models, which, despite their prowess, frequently require abundant labeled data for fine-tuning to attain optimal performance. By leveraging labeled data from source domains, UDA substantially mitigates this dependency, thereby eliminating the need for expensive and time-consuming annotation processes in the target domain. In this light, our paper delves into the subdomain of UDA, specifically focusing on cross-domain text classification.

Cross-domain text classification is encumbered by domain discrepancy emanating from variations in expressions across different domains. Addressing this conundrum, a substantial body of work (Clinchant et al., 2016; Ben-David et al., 2020; Zhou et al., 2020; Du et al., 2020; Wu and Shi, 2022) has been dedicated to extracting domain-invariant features between domains to bolster classification models' performance across multiple domains. Concurrently, another strand of work (Du et al., 2020; Karouzos et al., 2021) explores the utilization of language modeling to aid models in har-

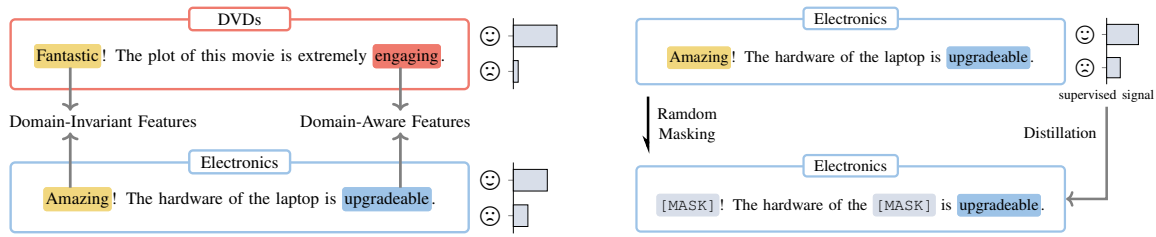
nessing task-agnostic features in the target domain, thus enhancing their performance in cross-domain text classification tasks.

In spite of these advancements, not all features conducive to a given task exhibit domain-invariance, as illustrated in Figure 1a. For instance, while expressions like “fantastic” and “amazing” are domain-invariant and can convey positive sentiments universally, terms such as “upgradeable” pertain to specific contexts like electronic products but not to DVDs, representing what we term as domain-aware features. The exploration of such features and their relation to the task at hand is often overlooked by existing methods, leaving a gap in addressing domain-aware features in the target domain.

To bridge this gap, Figure 1b outlines our proposed approach which ingeniously constructs a self-supervised signal, enabling models to attend to the latent domain-aware features of the target domain. This is pivotal for large models, which often grapple with new data from domains different from their training corpus. By masking domain-invariant features, our approach forces the model to establish a correlation between the predictions and the remaining domain-aware features. Subsequently, the model reinforces this relationship when domain-aware features are masked, allowing it to focus on latent domain-aware features in the target domain through a process that we denote as *self-supervised distillation*.

In this paper, we propose a novel cross-domain text classification model comprising a two-stage learning procedure: (1) *learning from the source domain* and (2) *adapting to the target domain*. This two-stage learning procedure significantly aug-

* Corresponding author.



(a) The model's predictions of DVDs (source domain) and Electronics (target domain) without adaptation. (b) An overview of the self-supervised signal we constructed to guide model.

Figure 1: The colors mean domain-invariant (●), source domain aware (●) and target domain aware (●). (a) The model can exploit the domain-invariant features but lacks the use of domain-aware features when predicting the target domain. (b) The supervised signal we construct, is designed to force the model to make a connection between predictions and latent domain-aware features of the target domain.

ments the model's performance and stability, rendering it a promising approach for cross-domain text classification tasks. Our experiment results on the Amazon reviews benchmark [Blitzer et al. \(2007\)](#) substantiate that our proposed method sets new state-of-the-art results for both single-source domain adaptations (94.17% \uparrow 1.03%) and multi-source domain adaptations (95.09% \uparrow 1.34%). Furthermore, a detailed analysis accentuates the generalization and effectiveness of our method, heralding a significant stride in the realm of UDA and cross-domain text classification.

To summarize, our contributions are as follows:

- We introduce *self-supervised distillation*, a simple yet effective method that helps models better capture domain-aware features from unlabeled data in the target domain.
- We propose a two-stage learning procedure that enables existing classification models to adapt to the target domain effectively.
- The experiments on the Amazon reviews benchmark for cross-domain classification show that our proposed model achieves new state-of-the-art results.

2. Background

2.1. Problem Formulation

To establish basic notations for our study, we define a domain $\mathcal{D} = \{\mathcal{X}, P(\mathcal{X})\}$, where \mathcal{X} represents the input feature space (e.g., the text representations), and $P(\mathcal{X})$ denotes the marginal probability distribution over that feature space. Let \mathcal{T} define a task (e.g., sentiment classification) as $\mathcal{T} = \{\mathcal{Y}, P(Y|X)\}$, where \mathcal{Y} is the label space. Moreover, a dataset is denoted by $\mathcal{D}^T = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$, where $\mathbf{x}_i \in \mathcal{D}$ and $y_i \in \mathcal{Y}$.

In this paper, we focus on cross-domain text classification, which is a subdomain of UDA ([Ramponi and Plank, 2020](#)). Specifically, we aim to learn a function \mathcal{F} trained with labeled dataset \mathcal{D}_S^T and un-

labeled dataset \mathcal{D}_T , which can effectively perform the task \mathcal{T} in the domain \mathcal{D}_T . Here we respectively denote S and T as the source domain and the target domain, and $P_S(\mathcal{X}) \neq P_T(\mathcal{X})$.

2.2. Prompt Tuning

We use prompt tuning as the formula for the text classification task, which reformulates the downstream task into cloze questions through a textual prompt x_p ([Petroni et al., 2019](#); [Brown et al., 2020](#)). Specifically, a textual prompt consists of an input sentence, a template containing [MASK], and two special tokens ([CLS] and [SEP]):

$$x_p = "[\text{CLS}] \mathbf{x}. \text{ It is } [\text{MASK}]. [\text{SEP}]", \quad (1)$$

where \mathbf{x} is the input sentence.

The Pretrained Language Model (PLM) takes the textual prompt x_p as input and utilizes contextual information to fill in the [MASK] token with a word from the vocabulary as the output. The output word is subsequently mapped to a label \mathcal{Y} . Following [Wu and Shi \(2022\)](#), we use "{good,bad}" as the label words. Finally, given an labeled dataset $\mathcal{D}^T = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$, the PLM is finetuned by minimizing the cross-entropy loss. The objective of prompt tuning can be defined as the following formula:

$$\mathcal{L}_{pmt}(\mathcal{D}^T; \theta_{\mathcal{M}}) = - \sum_{\mathbf{x}, y \in \mathcal{D}} y \log p_{\theta_{\mathcal{M}}}(\hat{y}|\mathbf{x}_p), \quad (2)$$

where y denotes the gold label, and $\theta_{\mathcal{M}}$ denotes the overall trainable parameters of the PLM.

2.3. Mask Language Modeling

We use mask language modeling to avoid shortcut learning¹ ([Geirhos et al., 2020](#)) and adapt to the

¹Shortcuts learning means that model learned decision rules (or shortcut features) that perform well on standard benchmarks but fail to transfer to more challenging testing conditions, such as real-world scenarios.

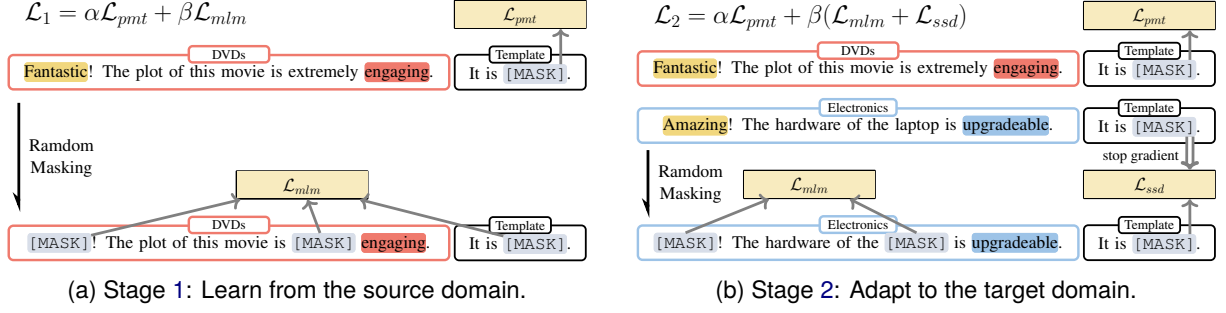


Figure 2: An overview of the proposed method. The \rightarrow means the model’s output. The \Rightarrow means the model’s output which has no gradient. The \mathcal{L}_{pmt} , \mathcal{L}_{mlm} , \mathcal{L}_{ssd} mean the losses of prompt tuning for classification, mask language modeling, and self-supervised distillation. The training process of the method consists of two stages. (a) In Stage 1, we apply mask language modeling and classification task on the source domain, to prevent the model from over-focusing on overfitting features or shortcut features. (b) In Stage 2, we continually train the classification task in the source domain, while we do mask language modeling and self-supervised distillation in the target domain.

distribution of the target domain. Specifically, we construct a *masked* textual prompt x_{pm} which is similar to the textual prompt x_p in Eq 1. The masked textual prompt consists of a *masked* input sentence, a template containing [MASK], and two special tokens ([CLS] and [SEP]):

$$x_{pm} = \text{“[CLS] } x_m \text{. It is [MASK]. [SEP]”}, \quad (3)$$

where x_m is the masked version of x in Eq 1.

Here we fine-tune the PLM through the masked language modeling task on x_m . Given an labeled dataset $\mathcal{D}^T = \{x_i, y_i\}_{i=1}^n$, the loss of each sentence in \mathcal{D}^T is the mean masked LM likelihood of all [MASK] in the sentence. Furthermore, the overall loss of \mathcal{D}^T is the summation of the individual sentence loss in the dataset:

$$\mathcal{L}_{mlm}(\mathcal{D}; \theta_{\mathcal{M}}) = - \sum_{x \in \mathcal{D}} \sum_{\hat{x} \in m(x_m)} \frac{\log p_{\theta_{\mathcal{M}}}(\hat{x}|x_{pm})}{len_m(x_m)}, \quad (4)$$

where $m(y_m)$ and $len_m(x_m)$ denote the masked words and counts in x_m , respectively, $\theta_{\mathcal{M}}$ denotes the overall trainable parameters of the PLM.

3. Method

In this section, we introduce the self-supervised distillation method and the two-stage learning procedure for cross-domain text classification. Figure 2 illustrates the overall learning procedure.

3.1. Self-Supervised Distillation (SSD)

The method needs a trained model which can perform the task. In addition, the model has the ability to generalize to the target domain through domain-invariant features, but not all useful features for the task are domain-invariant. We use the model itself

to construct a soft self-supervised signal. This signal enables the model to establish a connection between predictions and latent domain-aware features of the target domain. This process, which we refer to as **Self-Supervised Distillation (SSD)**, constitutes one of our core contributions.

During prediction, the model can only utilize the unmasked features of a masked sentence (x_{pm}), but it can use all features of the original version of the sentence (x_p). The model will be forced to make the connection between predictions of x_p in Eq 1 and the unmasked words of x_{pm} , which can include domain-invariant, domain-aware features, or both in the target domain. Recall that they contain the original (x_p) and masked versions (x_{pm}) of the same sentence (x), respectively. We perform knowledge distillation between the model predictions of $p_{\theta}(y|x_{pm})$ and $p_{\theta}(y|x_p)$. The objective of SSD can be defined as the following formula:

$$\mathcal{L}_{ssd}(\mathcal{D}; \theta_{\mathcal{M}}) = \sum_{x \in \mathcal{D}} \text{KL}(p_{\theta_{\mathcal{M}}}(y|x_{pm}) || p_{\theta_{\mathcal{M}}}(y|x_p)), \quad (5)$$

where x_m and x_{pm} is processed from the same input sentence x .

3.2. Learning Procedure

Our learning procedure comprises two stages, which are summarized in Algorithm 1 and Algorithm 2. We use the vanilla prompt tuning method without masking during inference.

3.2.1. Stage 1: Learn from the source domain

In Stage 1, our objective is to obtain a fine-tuned model with the ability to perform the downstream task effectively. We use a variant of the mask language modeling as an auxiliary task for prompt tuning to prevent the model from over-focusing overfitting features or shortcut features (Geirhos et al.,

Algorithm 1 Stage 1: Learn from the source domain

Input: Training samples of source domain labeled dataset \mathcal{D}_S^T

Output: Configurations of finetuned model $\theta_{\mathcal{M}}$

Initialize: PLM $\theta_{\mathcal{M}}$; learning rate η ; trade-off parameter α, β

```
1: while Training epoch not end do
2:   for  $x$  in  $\mathcal{D}_S^T$  do
3:      $\triangleright$  Minimizing the classification loss in
       source domain
4:      $\mathcal{L}'_1 \leftarrow \alpha \mathcal{L}_{pmt}(x; \theta_{\mathcal{M}})$ 
5:      $\theta_{\mathcal{M}} = \theta_{\mathcal{M}} - \eta \nabla_{\theta_{\mathcal{M}}} \mathcal{L}'_1$ 
6:      $\triangleright$  Minimizing the MLM loss in source do-
       main
7:      $\mathcal{L}''_1 \leftarrow \beta \mathcal{L}_{mlm}(x; \theta_{\mathcal{M}})$ 
8:      $\theta_{\mathcal{M}} = \theta_{\mathcal{M}} - \eta \nabla_{\theta_{\mathcal{M}}} \mathcal{L}''_1$ 
9:   end for
10: end while
```

Algorithm 2 Stage 2: Adapt to the target domain

Input: Training samples of source domain labeled dataset \mathcal{D}_S^T and target domain dataset \mathcal{D}_T

Output: Configurations of Final Model $\theta_{\mathcal{M}}$

Initialize: Model $\theta_{\mathcal{M}}$ already tuned in Stage 1; learning rate η ; trade-off parameter α, β

```
1: while Training epoch not end do
2:   for  $x^s, x^t$  in  $\mathcal{D}_S^T, \mathcal{D}_T$  do
3:      $\triangleright$  Minimizing the classification loss in
       source domain
4:      $\mathcal{L}'_2 \leftarrow \alpha \mathcal{L}_{pmt}(x^s; \theta_{\mathcal{M}})$ 
5:      $\theta_{\mathcal{M}} = \theta_{\mathcal{M}} - \eta \nabla_{\theta_{\mathcal{M}}} \mathcal{L}'_2$ 
6:      $\triangleright$  Minimizing the SSD loss and MLM in
       target domain
7:      $\mathcal{L}''_2 \leftarrow \beta (\mathcal{L}_{mlm}(x^t; \theta_{\mathcal{M}}) + \mathcal{L}_{ssd}(x^t; \theta_{\mathcal{M}}))$ 
8:      $\theta_{\mathcal{M}} = \theta_{\mathcal{M}} - \eta \nabla_{\theta_{\mathcal{M}}} \mathcal{L}''_2$ 
9:   end for
10: end while
```

2020). We refer to this approach as **MLM Enhanced Prompt Tuning (MEPT)**.

We initialize the parameters of the model \mathcal{M} using a Pretrained Language Model (PLM) such as BERT or RoBERTa. During each iteration, we sample new examples from the source domain \mathcal{D}_S^T to train the model \mathcal{M} . Firstly, we calculate the classification loss of those sentences and update the parameters with the loss, as shown in line 5 of Algorithm 1. Then we mask the same sentence and calculate mask language modeling loss to update the parameters, as depicted in line 8 of Algorithm 1. The parameters of the model will be updated together by these two losses.

In summary, the objective of Stage 1, given a labeled dataset \mathcal{D}^T , is obtained using the weighted cross-entropy loss for prompt tuning classification (\mathcal{L}_{pmt}) and mask language modeling loss (\mathcal{L}_{mlm}):

$$\begin{aligned} \mathcal{L}'_1(\mathcal{D}^T; \theta_{\mathcal{M}}) &= \alpha \mathcal{L}_{pmt}(\mathcal{D}^T; \theta_{\mathcal{M}}), \\ \mathcal{L}''_1(\mathcal{D}^T; \theta_{\mathcal{M}}) &= \beta \mathcal{L}_{mlm}(\mathcal{D}; \theta_{\mathcal{M}}), \end{aligned} \quad (6)$$

where α, β is the loss weight. As the Algorithm 1 shows, we alternate between \mathcal{L}'_1 and \mathcal{L}''_1 optimizations during training.

3.2.2. Stage 2: Adapt to the target domain

In Stage 2, we adapt the model trained in Stage 1 to the target domain. We refer to the resulting model **Two-stage Adapted MLM Enhanced Prompt Tuning (TAMEPT)**, which is our proposed model.

We initialize the parameters of our TAMEPT model using the model already tuned in Stage 1. Firstly, we sample labeled data from the source domain \mathcal{D}_S^T and calculate sentiment classification loss. The model parameters are updated using this loss in line 5 of Algorithm 2. Next, we sample unlabeled

data from the target domain \mathcal{D}_T and mask the unlabeled data to do a masking language model and self-supervised distillation with the previous prediction. It should be noted that the self-supervised distillation requires obtaining the prediction of the original sentence before masking. Finally, the model parameters are updated using the mask language modeling loss and self-supervised distillation loss of target domain examples, as shown in line 8 of Algorithm 2. The model parameters are updated together using the three aforementioned losses.

In conclusion, the training objective for Stage 2 is obtained using the weighted cross-entropy loss for classification (\mathcal{L}_{pmt}), mask language modeling loss (\mathcal{L}_{mlm}) and self-supervised distillation loss (\mathcal{L}_{ssd}). Given an labeled dataset \mathcal{D}_S^T and unlabeled datasets \mathcal{D}_T . the loss can be defined as:

$$\begin{aligned} \mathcal{L}'_2(\mathcal{D}_S^T, \mathcal{D}_T; \theta_{\mathcal{M}}) &= \alpha \mathcal{L}_{pmt}(\mathcal{D}_S^T; \theta_{\mathcal{M}}), \\ \mathcal{L}''_2(\mathcal{D}_S^T, \mathcal{D}_T; \theta_{\mathcal{M}}) &= \beta (\mathcal{L}_{mlm}(\mathcal{D}_T; \theta_{\mathcal{M}}) \\ &\quad + \mathcal{L}_{ssd}(\mathcal{D}_T; \theta_{\mathcal{M}})), \end{aligned} \quad (7)$$

where α, β is the loss weight. As the Algorithm 2 shows, we alternate between \mathcal{L}'_2 and \mathcal{L}''_2 optimizations during training.

3.3. Summary

The proposed method consists of two stages, as shown in Figure 2. **MLM Enhanced Prompt Tuning (MEPT)** means that we use mask language modeling as an auxiliary task for prompt tuning. **Two-stage Adapted MLM Enhanced Prompt Tuning (TAMEPT)** means that we use the model tuned in the source domain to adapt to the target domain with mask language modeling and self-supervised distillation.

Domain	# Positive	# Negative	# Unlabeled
Books (B)	1,000	1,000	6,000
DVDs (D)	1,000	1,000	34,741
Electronics (E)	1,000	1,000	13,153
Kitchen (K)	1,000	1,000	16,785

Table 1: Statistics for the Amazon reviews multi-domain classification dataset.

4. Experiments

In this section, we begin by introducing the standard benchmark used for cross-domain text classification on which we conduct experiments. We then present a series of baseline models that we use for comparison purposes. Subsequently, we provide a detailed description of our method’s training procedure. Finally, we present the results of the experiment and the analysis of our method.

4.1. Dataset

We evaluate the effectiveness of our proposed method on the Amazon reviews dataset [Blitzer et al. \(2007\)](#), which is a widely-used benchmark dataset for cross-domain text classification. The dataset contains reviews in four different domains: Books (B), DVDs (D), Electronics (E), and Kitchen appliances (K) with 2,000 manually labeled reviews in each domain, equally balanced between positive and negative sentiments. Additionally, the dataset also provides a certain amount of unlabeled data for each domain, as shown in Table 1.

(1) In the single-source domain adaptation experiments, we adopt the setting used in [\(Karouzos et al., 2021\)](#) and [\(Wu and Shi, 2022\)](#) to construct 12 cross-domain text classification tasks, each corresponding to a distinct ordered domain pair. In each of these 12 adaptation scenarios, we apply 20% of both labeled source and unlabeled target data for validation, while the labeled target data are used exclusively for testing and are not seen during training or validation. (2) In the multi-source domain adaptation experiments, we follow [\(Wu and Shi, 2022\)](#) to construct 4 cross-domain text classification tasks. Specifically, we choose one as the target domain and the remaining three domains as multiple source domains, resulting in tasks such as “BDE \rightarrow K”, “BDK \rightarrow E”.

4.2. Baselines

We present several strong baselines in our experiments and demonstrate the effectiveness of our proposed methods.

1. **R-PERL** ([Ben-David et al., 2020](#)): Utilize BERT for cross-domain text classification with pivot-based fine-tuning.

2. **DAAT** ([Du et al., 2020](#)): Employ BERT post-training for cross-domain text classification through adversarial training.
3. **p+CFd** ([Ye et al., 2020](#)): Leverage XLM-R for cross-domain text classification employing class-aware feature self-distillation (CFd).
4. **SENTIX_{Fix}** ([Zhou et al., 2020](#)): Pre-train a sentiment-aware language model via multiple pre-training tasks.
5. **UDALM** ([Karouzos et al., 2021](#)): Conduct fine-tuning with a mixed classification and MLM loss on domain-adapted PLMs.
6. **AdSPT** ([Wu and Shi, 2022](#)): Execute soft prompt tuning with an adversarial training object on vanilla PLMs.

We present our proposed method, denoted as **TAMEPT**, which is comprehensively introduced in Section 3. To thoroughly evaluate the performance, in alignment with [\(Karouzos et al., 2021\)](#) and [\(Wu and Shi, 2022\)](#), we employ *accuracy* as the selected evaluation metric.

4.3. Implementation Details

We adopt a 12-layer Transformer ([Vaswani et al., 2017](#); [Devlin et al., 2019](#)) initialized with RoBERTa_{base} ([Liu et al., 2019](#)) as the PLM.

(1) During Stage 1, we conduct training over 10 epochs with a batch size of 4, employing early stopping (patience = 3) based on the accuracy metric. The chosen optimizer for this phase is AdamW ([Loshchilov and Hutter, 2017](#)), with a learning rate of 1×10^{-5} . Additionally, we implement a strategy to halve the learning rate every 3 epochs. For this stage, we set $\alpha = 1.0, \beta = 0.6$ for Eq. 6. (2) During Stage 2, we conduct training over 10 epochs with a batch size of 4, with early stopping (patience = 3) on the mixing loss encompassing both classification loss and mask language modeling loss. The optimization utilizes AdamW with a learning rate of 1×10^{-6} without learning rate decay. The parameter settings for this stage are adjusted to $\alpha = 0.5, \beta = 0.5$ for Eq. 7.

Furthermore, for both the mask language modeling objective and the self-supervised distillation objective, we adopt a strategy wherein 30% of tokens are randomly replaced with [MASK] or random tokens. To manage the input size, the maximum sequence length is set to 512 through truncation of inputs. In a notable measure, during Stage 2, we randomly select an equal number of unlabeled data from the target domain of every epoch.

All the models and the accompanying analysis are meticulously implemented utilizing the PyTorch framework ([Paszke et al., 2019](#)), along with Hydra framework ([Yadan, 2019](#)), PyTorch Lightning ([Falcon and The PyTorch Lightning team, 2019](#)), HuggingFace transformers ([Wolf et al., 2020](#)) and

S → T	R-PERL	DAAT	p+CFd	UDALM	SENTIX _{Fix}	AdSPT	TAMEPT
B → D	87.80	89.70	87.65±0.10	90.97±0.22	91.30	92.00	93.27±0.49
B → E	87.20	89.57	91.30±0.20	91.69±0.31	93.25	93.75	94.82±0.23
B → K	90.20	90.75	92.45±0.60	93.21±0.22	96.20	93.10	95.75±0.40
D → B	85.60	90.86	91.50±0.40	91.00±0.42	91.15	92.15	94.83±0.31
D → E	89.30	89.30	91.55±0.30	92.30±0.47	93.55	94.00	94.57±0.18
D → K	90.40	87.53	92.45±0.20	93.66±0.37	96.00	93.25	95.84±0.24
E → B	90.20	88.91	88.65±0.40	90.61±0.30	90.40	92.70	93.20±0.63
E → D	84.80	90.13	88.20±0.40	88.83±0.61	91.20	93.15	92.63±0.34
E → K	91.20	93.18	93.60±0.50	94.43±0.24	96.20	94.75	96.16±0.07
K → B	83.00	87.98	89.75±0.80	90.29±0.51	89.55	92.35	92.18±0.84
K → D	85.60	88.81	87.80±0.40	89.54±0.59	89.85	92.55	91.77±0.68
K → E	91.20	91.72	92.60±0.50	94.34±0.26	93.55	93.95	95.06±0.43
AVG	87.50	90.12	90.63±0.40	91.74±0.38	92.68	93.14	94.17±0.40

Table 2: Results of single-source domain adaptation on Amazon reviews. There are four domains, B: Books, D: DVDs, E: Electronics, K: Kitchen appliances. In the table header, S: Source domain; T: Target domain. The TAMEPT is our proposed method, which is described in Section 3. We report mean performances and standard errors over 5 seeds.

S → T	AdSPT	TAMEPT
BDE → K	93.75	96.13±0.12
BDK → E	94.25	95.68±0.11
BEK → D	93.50	93.98±0.16
DEK → B	93.50	94.57±0.37
AVG	93.75	95.09±0.19

Table 3: Results of multi-source domain adaptation on Amazon reviews. The AdSPT is the only one of the baselines that do experiments in the multi-source domain adaptation settings. We report mean performances and standard errors over 5 seeds.

datasets (Lhoest et al., 2021) for a streamlined workflow.

4.4. Experiment Results

In this section, we concentrate on the results of single-source domain adaptation (Table 2) and multi-source domain adaptation (Table 3). Our proposed method, TAMEPT, achieves new state-of-the-art performance for both tasks, with an accuracy of 94.17% (+1.03%) for single-source domain adaptation and 95.09% (+1.34%) for multi-source domain adaptation. These results demonstrate the effectiveness of our approach for adapting PLMs to cross-domain text classification tasks.

4.4.1. Single-Source Domain Adaptation

The main experiment results in Table 2 demonstrate that our proposed method, TAMEPT, outper-

forms other state-of-the-art methods in most single-source domain adaptation settings. Specifically, compared to previous state-of-the-art methods, TAMEPT achieves significantly higher average accuracy (1.03% absolute improvement over AdSPT, 1.49% absolute improvement over SENTIX_{Fix}, 2.43% absolute improvement over UDALM, and 4.05% absolute improvement over DAAT). However, the AdSPT achieves better performance in experiments “E → D” and “K → D”. Furthermore, SENTIX_{Fix} achieves the best performance when the target domain is “K”, but our method still achieves comparable performance. It is mainly because the extra training data of SENTIX_{Fix} is closer to the domain “K”.

4.4.2. Multi-Source Domain Adaptation

The results presented in Table 3 demonstrate the superior performance of our proposed TAMEPT method in all multi-source domain adaptations. Compared to the previous state-of-the-art model AdSPT, TAMEPT achieves significantly higher average accuracy (1.34% absolute improvement). Notably, our method also achieves better performance than the single-domain adaptation method in most cases, with a lower standard error, indicating its ability to maintain stability as the amount of data and domain increases. However, when the target domain is “K”, the result of “E → K” (in Table 2) is superior to that of “BDE → K” (96.16% v.s. 96.13%). A similar situation occurs in AdSPT (94.75% v.s. 93.75%). It is mainly because the feature distribution of “E” and “K” is closer.

S → T	T _{AMEPT}	Stage 1		Stage 2	
		w/o Stage 2	w/o mlm	w/o mlm	w/o ssd
B → D	93.27 ±0.49	92.88±0.37	92.77±0.19	93.21±0.36	93.14±0.11
B → E	94.82 ±0.23	94.30±0.18	94.22±0.44	94.60±0.19	94.29±0.17
B → K	95.75 ±0.40	95.19±0.48	94.86±0.51	95.70±0.38	95.39±0.24
D → B	94.83 ±0.31	94.37±0.29	93.76±0.42	94.38±0.36	94.25±0.31
D → E	94.57 ±0.18	94.11±0.34	93.82±0.38	94.47±0.51	94.33±0.24
D → K	95.84 ±0.24	94.99±0.25	94.87±0.15	95.77±0.17	95.10±0.31
E → B	93.20±0.63	92.19±0.66	92.54±0.37	93.24 ±0.23	92.70±0.45
E → D	92.63 ±0.34	90.71±0.69	91.29±0.46	92.20±0.43	91.19±0.25
E → K	96.16 ±0.07	95.70±0.58	95.40±0.46	95.84±0.19	95.77±0.23
K → B	92.18±0.84	92.09±0.53	91.89±0.37	91.94±1.99	92.55 ±0.26
K → D	91.77 ±0.68	90.49±0.70	89.85±0.44	89.91±4.18	91.38±0.36
K → E	95.06±0.43	94.55±0.26	94.58±0.32	95.15 ±0.15	95.02±0.04
AVG	94.17 ±0.40	93.46±0.44	93.32±0.38	93.87±0.76	93.76±0.25

Table 4: The ablation experiments of our method. There are four domains, B: Books, D: DVDs, E: Electronics, K: Kitchen appliances. In the table header, S: Source domain; T: Target domain. The “mlm” and “ssd” mean self-supervised distillation and mask language modeling. The “w/o Stage 2” is also called MEPT, which is described in Section 3.2.1. We report mean performances and standard errors over 5 seeds.

4.5. Analysis

Table 4 indicates the results of our ablation experiments, which are conducted to evaluate the contributions of different components in our proposed method. Figure 3 shows the case study of our method. Additionally, Table 5 presents the results of experiments that validate the generality of our two-stage adaptation approach on different methods and pre-trained models. These experiments demonstrate the effectiveness and flexibility of our proposed method.

4.5.1. Ablation Study

Our proposed method comprises two stages, namely Stage 1 and Stage 2. To evaluate the effectiveness of each stage, we conduct ablation experiments by removing the corresponding components and comparing the performance of the resulting model with the original model. The results of the ablation experiments are presented in Table 4.

In Stage 1, we observe that the use of **mask language modeling** is crucial for achieving high accuracy. Specifically, when we remove the mask language modeling, the performance of the Stage 1 model dropped by an average of 0.14% (from 93.46% to 93.32%), as shown in Table 4. With the exception of the “E → D” experiment, the model including mask language modeling consistently outperforms the one without mask language modeling.

In Stage 2, we find that using both **self-supervised distillation** and **mask language modeling** is critical for achieving high accuracy and stability. When we remove the self-supervised dis-

tillation, the performance of the Stage 2 model decreases by an average of 0.41% (from 94.17% to 93.76%). Similarly, when we remove the mask language modeling, the performance of the Stage 2 model decreases by an average of 0.30% (from 94.17% to 93.87%), while the standard error increases by 0.36 (from 0.40 to 0.76). Notably, the standard error for experiments “K → B” and “K → D” increases significantly by 1.15 (from 0.84 to 1.99) and 3.50 (from 0.68 to 4.18), respectively. These results indicate that mask language modeling is a crucial factor for achieving high accuracy and, in particular, stability in the proposed method.

In summary, our experiments confirm the effectiveness of both self-supervised distillation and mask language modeling in achieving high accuracy and stability.

4.5.2. Case Study

As shown in Figure 3, we conduct a case study to demonstrate the effectiveness of our method in capturing domain-aware features from unlabeled data in the target domain. Specifically, we used the sentence “*Plain Vanilla Wireless adapter for you. Slow but steady and inexpensive.*” as an example to analyze the gradients from MEPT² and T_{AMEPT} in the “D → E” setting. The gradients from MEPT and T_{AMEPT} are depicted in Figures 3(a-b), respectively. Notably, compared to the gradient of MEPT, the gradient from T_{AMEPT} places more emphasis on the words “slow” and “steady”, which are domain-aware features. This result demonstrates that our

²MEPT: The proposed model T_{AMEPT} without Stage 2.

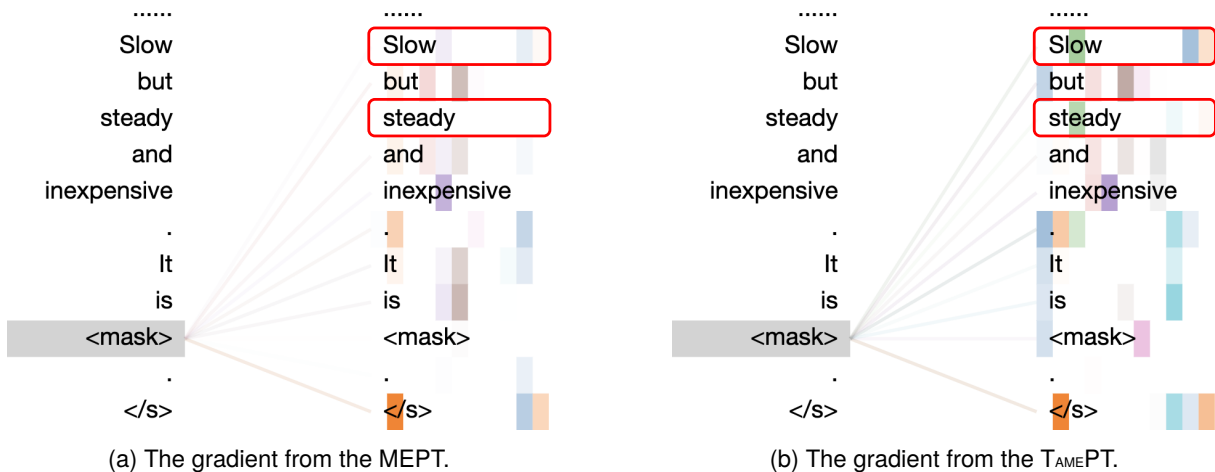


Figure 3: Visualization for the sentence “Plain Vanilla Wireless adapter for you. Slow but steady and inexpensive.” in the “D → E” setting. The different colors mean the gradient of different heads. Compared with the gradient from MEPT, the gradient from TAMEPT pays more attention to the domain-aware features (“slow” and “steady”).

Method	BERT		RoBERTa	
	Stage 1	+Stage 2	Stage 1	+Stage 2
FT	90.00±0.51	91.10±1.77	93.31±0.47	94.09±0.38
PT	90.78±0.78	91.62±1.59	93.63±0.36	94.38±0.36
MEPT	91.27±0.58	92.54±0.26	93.74±0.43	94.40±0.35
AVG	91.02±0.62	91.75±1.21	93.56±0.42	94.29±0.36

Table 5: The experiments for validating the generality of our method. We report mean performances and standard error of 12 single-source domain adaptations and 4 multi-domain adaptations. The MEPT initialized with RoBERTa and adapted by Stage 2 is denoted as “TAMEPT”, which is our proposed model.

method is capable of assisting models in capturing domain-aware features from unlabeled data in the target domain.

4.5.3. Generality Study

To validate the generality of our method, we conduct experiments on different pre-trained models and methods, as summarized in Table 5. Specifically, we apply Algorithm 2 to the fine-tuning (FT), prompt tuning (PT), and MEPT methods. The results demonstrate the effectiveness of our method on different pre-trained models and methods, with an average improvement of 0.73% (from 91.02% to 91.75% with BERT and from 93.56% to 94.29% with RoBERTa). These findings support the generality of our proposed method.

4.5.4. Sensitive Analysis

To evaluate the sensitivity of the hyperparameters, we conduct a sensitive analysis of the mask ratio, which is the ratio of the number of masked tokens to




Figure 4: The sensitive analysis of the hyperparameters mask ratio in the $B \rightarrow D$ setting.

the total number of tokens. The results are shown in Figure 4. The results in the $B \rightarrow D$ setting demonstrate that the performance of our method is relatively stable across a wide range of mask ratios, with the maximum difference in average accuracy being less than 0.5%. This finding indicates that our method is robust to changes in the mask ratio.

5. Related Work

Unsupervised Domain Adaptation is a technique that addresses domain shift issues by learning labeled data of the source domain(s) and unlabeled data of the target domain, which is typically available for both source and the target domains (Ramponi and Plank, 2020).

As mentioned in Section 1, current works can be roughly categorized into two groups. One group aims to capture domain-invariant features, including pivot-based methods (Ben-David et al., 2020), domain adversarial training (Wu and Shi, 2022),

class-aware feature self-distillation (Ye et al., 2020), and sentiment-aware language model (Zhou et al., 2020). Another group employs pre-trained models to exploit the task-agnostic features of the target domain during domain adaptation by language modeling (Du et al., 2020; Karouzos et al., 2021). Some works use both domain-invariant and task-agnostic features (Du et al., 2020). However, previous work only focuses on extracting domain-invariant features or task-agnostic features. In contrast, we are the first to consider the domain-aware features of the target domain.

6. Conclusion

In this paper, we propose a two-stage learning procedure for cross-domain text classification that leverages self-supervised distillation to capture domain-aware features in the target domain. We demonstrate that this procedure outperforms previous state-of-the-art models in most cases, achieving a significant improvement in average accuracy. Our experiments also highlight the significance of self-supervised distillation and mask language modeling in achieving high performance and stability. Moreover, the two-stage learning procedure can be easily applied to existing trained models for cross-domain text classification.

Acknowledgments

We gratefully acknowledge the support of the National Natural Science Foundation of China (NSFC) via grant 62236004 and 62206078, and the support of Du Xiaoman (Beijing) Science Technology Co., Ltd.

7. Bibliographical References

Eyal Ben-David, Carmel Rabinovitz, and Roi Reichart. 2020. [PERL: Pivot-based Domain Adaptation for Pre-trained Deep Contextualized Embedding Models](#). *Transactions of the Association for Computational Linguistics*, 8:504–521.

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.

Stéphane Clinchant, Gabriela Csurka, and Boris Chidlovskii. 2016. [A domain adaptation regularization for denoising autoencoders](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 26–31, Berlin, Germany. 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.

Chunning Du, Haifeng Sun, Jingyu Wang, Qi Qi, and Jianxin Liao. 2020. [Adversarial and domain-aware BERT for cross-domain sentiment analysis](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4019–4028, Online. Association for Computational Linguistics.

William Falcon and The PyTorch Lightning team. 2019. [PyTorch Lightning](#).

Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. 2020. [Shortcut Learning in Deep Neural Networks](#). *arXiv e-prints*, page arXiv:2004.07780.

Arthur Gretton, Karsten Borgwardt, Malte Rasch, Bernhard Schölkopf, and Alex Smola. 2006. A kernel method for the two-sample-problem. *Advances in neural information processing systems*, 19.

Constantinos Karouzos, Georgios Paraskevopoulos, and Alexandros Potamianos. 2021. [UDALM: Unsupervised domain adaptation through language modeling](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2579–2590, Online. Association for Computational Linguistics.

Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario

- Šaško, Gunjan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clément Delangue, Théo Matussière, Lysandre Debut, Stas Bekman, Pierric Cistac, Thibault Goehringer, Victor Mustar, François Lagunas, Alexander Rush, and Thomas Wolf. 2021. [Datasets: A community library for natural language processing](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–184, Online and Punta Cana, Dominican Republic. 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 pre-training approach](#). *CoRR*, abs/1907.11692.
- Ilya Loshchilov and Frank Hutter. 2017. [Fixing weight decay regularization in adam](#). *CoRR*, abs/1711.05101.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. [Language models as knowledge bases?](#) 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 2463–2473, Hong Kong, China. Association for Computational Linguistics.
- Alan Ramponi and Barbara Plank. 2020. [Neural unsupervised domain adaptation in NLP—A survey](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6838–6855, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- 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.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clément Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Hui Wu and Xiaodong Shi. 2022. [Adversarial soft prompt tuning for cross-domain sentiment analysis](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2438–2447, Dublin, Ireland. Association for Computational Linguistics.
- Omry Yadan. 2019. [Hydra - a framework for elegantly configuring complex applications](#). Github.
- Hai Ye, Qingyu Tan, Ruidan He, Juntao Li, Hwee Tou Ng, and Lidong Bing. 2020. [Feature adaptation of pre-trained language models across languages and domains with robust self-training](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7386–7399, Online. Association for Computational Linguistics.
- Jie Zhou, Junfeng Tian, Rui Wang, Yuanbin Wu, Wenming Xiao, and Liang He. 2020. [SentiX: A sentiment-aware pre-trained model for cross-domain sentiment analysis](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 568–579, Barcelona, Spain (Online). International Committee on Computational Linguistics.

8. Language Resource References

- Blitzer, John and Dredze, Mark and Pereira, Fernando. 2007. *Biographies, Bollywood, Boomboxes and Blenders: Domain Adaptation for Sentiment Classification*. Association for Computational Linguistics.