

Towards Unified Prompt Tuning for Few-shot Text Classification

Jianing Wang^{1,*}, Chengyu Wang^{2,*}, Fuli Luo², Chuanqi Tan², Minghui Qiu², Fei Yang³,
Qihui Shi⁴, Songfang Huang², Ming Gao^{1,5†}

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

² Alibaba Group, Hangzhou, China ³ Zhejiang Lab, Hangzhou, China

⁴ Ant Group, Hangzhou, China

⁵ Shanghai Key Laboratory of Mental Health and Psychological Crisis Intervention,
School of Psychology and Cognitive Science, East China Normal University, Shanghai, China

lygwjn@gmail.com,

{chengyu.wcy, lf1259702, chuanqi.tcq}@alibaba-inc.com,

minghui.qmh@alibaba-inc.com, yangf@zhejianglab.com,

qihui.sqh@antgroup.com, songfang.hsf@alibaba-inc.com,

mgao@dase.ecnu.edu.cn

Abstract

Prompt-based fine-tuning has boosted the performance of Pre-trained Language Models (PLMs) on few-shot text classification by employing task-specific prompts. Yet, PLMs are unfamiliar with prompt-style expressions during pre-training, which limits the few-shot learning performance on downstream tasks. It would be desirable if the models can acquire some prompting knowledge before adapting to specific NLP tasks. We present the *Unified Prompt Tuning (UPT)* framework, leading to better few-shot text classification for BERT-style models by explicitly capturing prompting semantics from non-target NLP datasets. In *UPT*, a novel paradigm *Prompt-Options-Verbalizer* is proposed for joint prompt learning across different NLP tasks, forcing PLMs to capture task-invariant prompting knowledge. We further design a self-supervised task named *Knowledge-enhanced Selective Masked Language Modeling* to improve the PLM’s generalization abilities for accurate adaptation to previously unseen tasks. After multi-task learning across multiple tasks, the PLM can be better prompt-tuned towards any dissimilar target tasks in low-resourced settings. Experiments over a variety of NLP tasks show that *UPT* consistently outperforms state-of-the-arts for prompt-based fine-tuning.¹

1 Introduction

The emergence of Pre-trained Language Models (PLMs) has boosted the performance of a variety

of NLP tasks (Qiu et al., 2020; Han et al., 2021a). However, during fine-tuning, PLMs can perform poorly with few training samples due to model over-fitting (Gao et al., 2021).

To alleviate this problem for low-resourced scenarios, natural language prompts have been applied to enable few-shot or zero-shot learning with PLMs (Liu et al., 2021a). To make prompts more flexible and task-adaptive, *prompt tuning* freezes the PLM backbone and adjusts the representations of prompts (Lester et al., 2021). This type of method is especially suitable for ultra-large PLMs that are difficult to tune. For BERT-style PLMs, *prompt-based fine-tuning* has been proposed, transforming text classification tasks into cloze-style problems (Schick and Schütze, 2021a,b; Gao et al., 2021). To specify, task-specific discrete templates with masked language tokens are added to input texts. The result tokens of the masked positions predicted by the Masked Language Modeling (MLM) head are used for class label prediction². Therefore, the pre-trained knowledge acquired by PLMs can be better utilized by “re-using” the MLM training objective. Witnessing the successful usage of prompts for few-shot learning, various following-up works have been conducted, such as continuous prompt encoding (Liu et al., 2021c), knowledgeable prompt learning (Hu et al., 2021), and prompt generation (Shin et al., 2020).

Recently, a few works (Wei et al., 2021; Zhong et al., 2021a; Mishra et al., 2021) focus on multi-

* J. Wang and C. Wang contributed equally to this work.

† Corresponding author.

¹All datasets are publicly available. Source codes will be released in EasyNLP (Wang et al., 2022). URL: <https://github.com/alibaba/EasyNLP>

²For example, in the review analysis task, given an input “It is a wonderful movie.”, one can add the prompt template “Based on the review, it is [MASK].” to the input. The output of the masked token “great” and “terrible” can be mapped to the positive and the negative class, respectively.

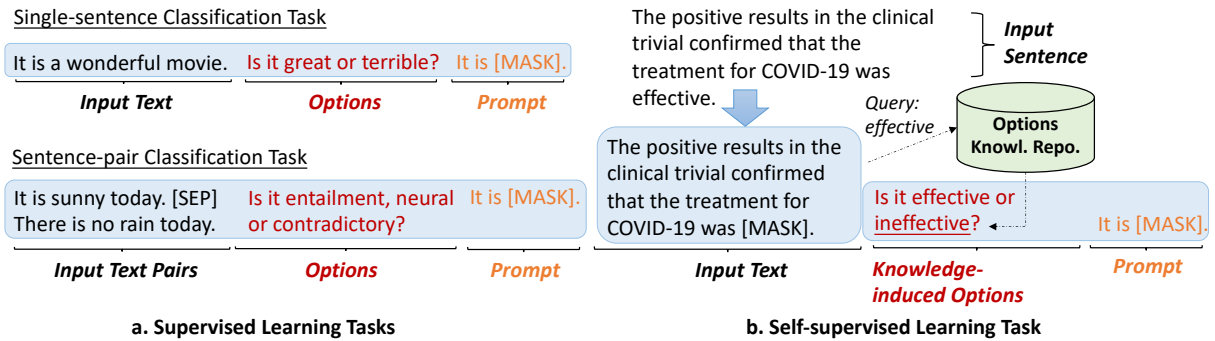


Figure 1: *UPT* is a unified framework that learns prompting knowledge from non-target NLP datasets to improve the performance on target tasks, in the format of *Prompt-Options-Verbalizer* (Sect. 2.2). Figures a) and b) show examples of supervised and self-supervised learning tasks (i.e., *Knowledge-enhanced Selective MLM*, Sect. 2.3).

task prompt tuning on ultra-large PLMs. Specifically, they tune PLMs on full training samples from different tasks to force PLMs to learn more prompting knowledge, and directly make predictions over the target task by zero-shot learning. Yet, we observe that for BERT-style PLMs, the performance is not satisfactory for two reasons. 1) These PLMs are sensitive to different designs of prompt templates and verbalizers (Liu et al., 2021c), which fail to adapt to target tasks with new prompts and verbalizers. 2) There are word distribution differences between prompt-style texts and sentences in pre-training corpora. It would be better if BERT-style PLMs can acquire some prompting knowledge before they are adapted to downstream tasks. Therefore, a natural question arises: *how can we make BERT-style PLMs adapt to target NLP tasks accurately with more prompting knowledge?*

To address these issues, we introduce a novel framework named *Unified Prompt Tuning (UPT)*, facilitating better few-shot text classification performance for BERT-style models by explicitly capturing general prompting semantics from non-target datasets. Specially, we propose a unified paradigm named *Prompt-Options-Verbalizer (POV)*, which enables mixture prompt-tuning over a series of *non-target NLP tasks* of varied types. To further improve the model’s generalization abilities on previously unseen tasks, we propose a novel auxiliary task named *Knowledge-enhanced Selective MLM (KSMLM)*, which mimics the behavior of MLM with explicit usage of prompts following the *POV* paradigm. After multi-task training is completed, the underlying PLM can be fine-tuned to fit any few-shot tasks using the same prompting paradigm.

In the experiments, we verify the effectiveness of *UPT* over public NLP datasets of various tasks.

Experimental results show that *UPT* consistently outperforms state-of-the-art approaches for prompt-based few-shot fine-tuning. In summary, we make the following major contributions:

- We introduce the novel *UPT* framework to improve prompt-based fine-tuning for BERT-style models, which captures unified prompting semantics from multiple source tasks of various types for few-shot text classification on new target tasks.
- In *UPT*, a new paradigm *POV* is proposed for joint prompt tuning across different NLP tasks. We further design the self-supervised *KSMLM* task to improve the PLM’s generalization abilities for accurate task adaptation.
- Extensive experiments over various NLP datasets show that *UPT* consistently outperforms state-of-the-arts for prompt-based few-shot fine-tuning by a relatively large margin.

2 *UPT*: The Proposed Framework

We start with a brief overview of the *UPT* framework, followed by its detailed techniques.

2.1 A Brief Overview of *UPT*

For clarity, we introduce some basic notations. Let \mathcal{D}^* be the N -way- K -shot training set of a target NLP task \mathcal{T}^* . The underlying PLM is parameterized by Θ . The basic goal of few-shot learning is to obtain a high-performing model for \mathcal{T}^* based on \mathcal{D}^* , with parameters initialized from Θ . As the size of \mathcal{D}^* is only $N \times K$, the model performance would be highly limited. Here, we assume that there are M other NLP tasks that are *dissimilar* to \mathcal{T}^* , i.e., $\mathcal{T}^{(1)}, \dots, \mathcal{T}^{(M)}$, with their (usually non few-shot) training sets denoted as $\mathcal{D}^{(1)}, \dots, \mathcal{D}^{(M)}$,

respectively³. The *UPT* framework seeks to explore how to employ $\mathcal{D}^{(1)}, \dots, \mathcal{D}^{(M)}$ to enhance the performance of the PLM on a new task (such as \mathcal{T}^*) based on its own few-shot training set \mathcal{D}^* .

In *UPT*, the model is firstly trained over all the source tasks $\mathcal{T}^{(1)}, \dots, \mathcal{T}^{(M)}$, aiming to learn the semantics of prompts and the general methodology of solving downstream tasks by prompting. After that, it is prompt-tuned over a specific target task \mathcal{T}^* in the low-resourced scenario. To unify the learning process, each training sample i in all different tasks (either $\mathcal{T}^{(1)}, \dots, \mathcal{T}^{(M)}$ or \mathcal{T}^*) is augmented in the same format, by means of the *Prompt-Options-Verbalizer (POV)* triple (P_i, O_i, V_i) . Here, P_i is the prompt. O_i is the expression containing all possible options of the masked language token appearing in the prompt P_i (i.e., the collection of label words). V_i is the verbalizer that maps the target token predicted by the MLM head of the PLM to the class label. Readers can also refer to the examples of supervised learning tasks in Figure 1.

In addition, we observe that the diversity of label words in original labeled tasks $\mathcal{T}^{(1)}, \dots, \mathcal{T}^{(M)}$ is limited. For previously unseen tasks, the optimization of these tasks alone often leads to a poorly generalized model that is biased towards these tasks. Therefore, we further introduce the self-supervised *Knowledge-enhanced Selective MLM (KSMLM)* $\tilde{\mathcal{T}}$ as an auxiliary task. Specifically, take the sentences from source tasks training data $\tilde{\mathcal{D}} = \mathcal{D}^{(1)} \cup \mathcal{D}^{(2)} \cup \dots \cup \mathcal{D}^{(M)}$ as inputs. These sentences are selectively masked, with options generated by rich knowledge mined from a massive corpus. An example is also in Figure 1. Hence, the model has better generalization abilities and avoids catastrophic forgetting of pre-training knowledge.

2.2 The Unified Prompting Paradigm

A fundamental challenge for prompt-based training across $\mathcal{D}^{(1)}, \dots, \mathcal{D}^{(M)}$ for BERT-style models is that different NLP tasks have diverse sets of label words w.r.t. masked language tokens. When dealing with a mixture of training samples, a naive solution is to build a unified output prediction space, consisting of candidate label words from all tasks. However, the enlarged output space makes it chal-

³Note that we constrain that $\mathcal{T}^{(1)}, \dots, \mathcal{T}^{(M)}$ are *dissimilar* to \mathcal{T}^* to deal with true low-resourced scenarios where no training sets of similar tasks are available. If $\mathcal{T}^{(1)}, \dots, \mathcal{T}^{(M)}$ are similar to \mathcal{T}^* , one can directly apply transfer learning techniques to train the model, which is considered a relatively trivial problem and not the major focus of this work.

lenging for the PLM to optimize. Additionally, the output prediction space may not cover the label words of all possible unseen NLP tasks.

Here, we propose a unified prompting paradigm that augments each sample i by a *Prompt-Options-Verbalizer (POV)* triple (P_i, O_i, V_i) . P_i is the prompt that provides task guidance (in line with PET (Schick and Schütze, 2021a,b)). O_i is a fixed expression that explicitly provides selection for the model over all its candidate label words⁴. To facilitate the fast adaptation of arbitrary tasks, the verbalizer V_i maps the output of the masked language token to the entire vocabulary \mathcal{V} . We can see that the options are crucial as they give strong indications on the possible outputs of the PLM (i.e., the candidates). Overall, the output probability $q(v|i, P_i, O_i, \Theta)$ of the token $v \in \mathcal{V}$ w.r.t. the training sample i is computed as follows:

$$q(v|i, P_i, O_i, \Theta) = \frac{\exp(s(v|i, P_i, O_i, \Theta))}{\sum_{v' \in \mathcal{V}} \exp(s(v'|i, P_i, O_i, \Theta))}$$

where $s(v|i, P_i, O_i, \Theta)$ is the un-normalized score of the MLM head (before the softmax function) for generating token v at the position of the masked language token with i, P_i and O_i as inputs. Denote the entire prediction vector (of the length $|\mathcal{V}|$) as $Q(\mathcal{V}|i, P_i, O_i, \Theta)$. The *multi-task prompting loss* (denoted as \mathcal{L}_{MP}) can be written as follows:

$$\mathcal{L}_{MP} = - \sum_{i \in \mathcal{D}} P(\mathcal{V}|i, P_i, O_i, \Theta) \cdot \log Q(\mathcal{V}|i, P_i, O_i, \Theta)$$

where $\mathcal{D} = \bigcup_{k=1}^M \mathcal{D}^{(k)}$, and $P(\mathcal{V}|i, P_i, O_i, \Theta)$ is the one-hot ground-truth prediction vector.

In addition, we notice that $\mathcal{D}^{(1)}, \dots, \mathcal{D}^{(M)}$ can be arbitrary labeled datasets with varied sizes. Optimizing \mathcal{L}_{MP} directly on their original datasets would make the few-shot learner more likely to be biased towards larger datasets. In our work, we do stratified sampling to form a batch where a training sample i from $\mathcal{D}^{(1)}, \dots, \mathcal{D}^{(M)}$ is picked with the probability proportional to its own dataset size (denoted as w_i), i.e., $w_i = \frac{\log |\mathcal{D}^{(k)}| + \gamma}{M \cdot \gamma + \sum_{k'=1}^M \log |\mathcal{D}^{(k')}|}$ where $\gamma > 0$ is a smoothing factor and $i \in \mathcal{D}^{(k)}$. Hence, we re-formulate \mathcal{L}_{PT} as the *weighted multi-task*

⁴Note that our framework is not restricted to binary classification. For NLP tasks with many labels, we can also directly list all the labels in options. More details and the external experiments can be found in Appendix.

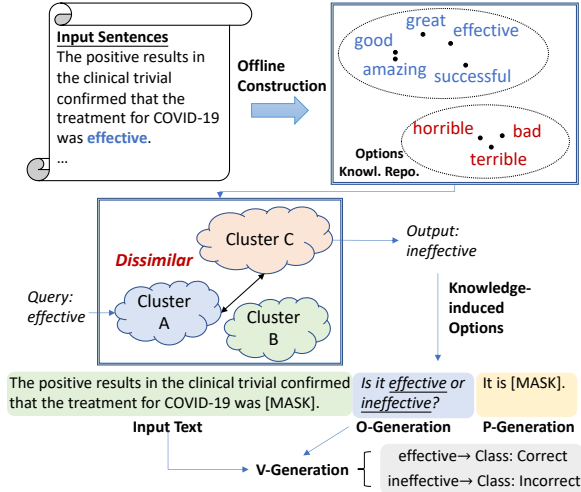


Figure 2: An illustrated example of the *POV* generation process for the *KSMLM* task.

prompting (WMP) loss \mathcal{L}_{WMP} :

$$\mathcal{L}_{WMP} = - \sum_{i \in \mathcal{D}} w_i \cdot P(\mathcal{V}|i, P_i, O_i, \Theta) \cdot \log Q(\mathcal{V}|i, P_i, O_i, \Theta)$$

2.3 Extending Unified Prompting to Self-supervised Learning

One drawback of the above approach is that the diversity of label words in these supervised learning tasks is usually limited, covering a narrow spectrum of the vocabulary \mathcal{V} . The model would not be well generalized for tasks with new label words. Hence, we leverage the idea of MLM pre-training, formulated by the *POV* paradigm.

As a naive approach, given a sentence, we can randomly mask a word and generate the options of the correct and a randomly selected word, and then ask the model to make the prediction. Unfortunately, the seemingly feasible approach may ruin the training process, because not all words are suitable label words. For example, stop words and a large number of verbs and adverbs have not been used in any verbalizers in downstream tasks. The alternatives used in options should be reasonable, in order to make the model learn truly useful knowledge. To address the issue, we present the self-supervised *KSMLM* task, with an example shown in Figure 2. In the following, we describe the *POV* construction process for *KSMLM*. After that, the loss function of the task is given.

P-Generation. This process aims to generate a template with a [MASK] token for each sentence, which is fixed to be “It is [MASK].” during the

multi-task training stage. In the task-specific fine-tuning stage, we follow LM-BFF (Gao et al., 2021) to automatically generate templates for each task. During training, the PLM is asked to predict the actual word of the masked position.

O-Generation. From Gao et al. (2021), we can see that most label words for language understanding tasks are adjectives⁵ (such as “great” and “terrible” for sentiment analysis). Thus in our work, we detect all adjectives in the corpus by part-of-speech tagging models⁶ and filter out low-frequency adjectives. The adjectives are then clustered by K-Means, with their token representations generated from the underlying PLM as features. Formally, We construct a knowledge repository named *Options Knowledge Repository (OKR)*, in the form of triples $\mathcal{R} = \{(v, \vec{v}, c_v)\}$, where v is a candidate label word. \vec{v} and c_v denote the representation vector and the cluster membership of v , respectively. The cluster centroids are also stored. We do not use existing lexicons such as WordNet (Miller, 1995) because they may have limited coverage of label words. Additionally, the automatic process enables the extension of our algorithm to arbitrary languages and domains.

With the availability of \mathcal{R} , we can generate knowledge-induced options. Given a sentence with the masked word as v , we query v against \mathcal{R} for the most *dissimilar* cluster w.r.t. v , denoted as \tilde{c}_v , where the cosine similarity of the vector representation \vec{v} and the cluster centroid is employed as the similarity measure. Finally, we randomly select one adjective from \tilde{c}_v as the alternative label word to generate the *knowledge-induced options*. The text expressions of options is fixed, i.e., “Is it [x1] or [x2]?”. Readers can further refer to the example in Figure 2.

V-Generation. For verbalizers, we map the true and the generated label words in the options to two classes, namely *Class: Correct* and *Class: Incorrect*. For instance, the verbalizers of the sample sentence in Figure 2 are:

It is “effective”. → “Class: Correct”

It is “ineffective”. → “Class: Incorrect”

Loss Function. The *KSMLM* loss is significantly different from the auxiliary MLM loss used

⁵In fact, we can also take into account the nouns if the label word space of target tasks is related to topics. Without loss of generality, we only consider adjectives in the experiments.

⁶We use the *spacy* toolkit in our work. URL: <https://spacy.io/>.

in Schick and Schütze (2021a,b). In $\tilde{\mathcal{D}}$, each training sample i can be directly extended to the training example for *KSMLM* by *POV* construction process with exactly one masked token, the *knowledge-induced options* O_i and the prompt P_i . The PLM is trained to predict the correct masked word in the sentence, with the loss function: $\mathcal{L}_{KSMLM} = -\sum_{i \in \tilde{\mathcal{D}}} P(\mathcal{V}|i, P_i, O_i, \Theta) \log Q(\mathcal{V}|i, P_i, O_i, \Theta)$. Overall, the loss function of *UPT* \mathcal{L} is defined as the summation of the WMP and KSMLM loss:

$$\mathcal{L} = \mathcal{L}_{WMP} + \lambda \cdot \mathcal{L}_{KSMLM}$$

where $\lambda \geq 0$ is the balancing hyper-parameter.

Discussion. To our knowledge, external knowledge has also been applied to other prompt-based methods, such as KPT (Hu et al., 2021). The major difference between KPT and ours is that *UPT* uses the knowledge for options creation of the self-supervised task *KSMLM* that we proposed, in order to improve the model generalization abilities for accurate adaptation on new tasks. In contrast, previous works consider the expansion of verbalizers for specific downstream NLP tasks.

2.4 Few-shot Fine-tuning

For a specific downstream task \mathcal{T}^* , the samples in the target few-shot training set \mathcal{D}^* can be processed and computed in the same way as those supervised tasks used during *UPT*. The learning consistency in the two stages ensures that the underlying PLM has already acquired prompting knowledge for \mathcal{T}^* . In addition, one can prompt-tune a single PLM over various tasks and uses it to fine-tune over any target tasks, making it computationally efficient to produce models for these applications.

3 Experiments

3.1 Experimental Settings

In the experiments, we employ 9 public text classification datasets to evaluate the proposed *UPT* framework, which are divided into three groups: sentiment analysis (Sentiment) (SST-2 (Socher et al., 2013), MR (Hu and Liu, 2004), CR (Pang and Lee, 2005)), Natural Language Inference (NLI) (MNLI (Williams et al., 2018), SNLI (Bowman et al., 2015), QNLI (Wang et al., 2019b), RTE (Dagan et al., 2005)) and Paraphrase (MRPC (Dolan and Brockett, 2005), QQP⁷). The data statistics are shown in the Appendix. In default, $K = 16$ (training instances per class).

⁷<https://www.quora.com/q/quoradata/>.

As mentioned above, during *UPT*, we only leverage full training data from all *dissimilar* task groups, and then prompt-tune the model on the target task in the low-resourced setting. For example, when the target task is SST-2, the training data during *UPT* is from NLI and Paraphrase. The underlying PLM is the RoBERTa-large model (with 335M parameters) (Liu et al., 2019), unless otherwise specified. The baselines include standard *fine-tuning*, and four recently proposed few-shot learning algorithms: PET (Schick and Schütze, 2021a)⁸, LM-BFF (Gao et al., 2021)⁹, P-tuning (Liu et al., 2021c)¹⁰ and PPT (Gu et al., 2021). To make a fair comparison with these single-task baselines, a variant of our approach (denoted as *UPT*-Single) is also implemented by only fine-tuning over the few-shot target task based on *POV* without the usage of *dissimilar* supervised source datasets.

As we use other *dissimilar* datasets to train our model, we also include two multi-task methods that are *meta-tuned* using the same *dissimilar* datasets as strong baselines, namely MT (Zero-shot) and MT (Few-shot) (Zhong et al., 2021a)¹¹. We also implement the zero-shot version of *UPT*, denote as *UPT* (Zero-shot). In addition, given a supervised NLP task, multiple prompts can be manually crafted. By augmenting one training sample with these prompts, we can automatically realize *self-ensemble learning*. For the self-ensemble version of *UPT*, we employ five different prompts. For each input sample, we randomly select one expression of options and one set of verbalizers. We denote this method as *UPT*-SE. The designed prompts, options, and verbalizers are listed in the Appendix. All the results of these models are evaluated in terms of averaged accuracy and its standard deviation, over 5 random seeds.

Our *UPT* framework is implemented in PyTorch

⁸<https://github.com/timoschick/pet>

⁹For a fair comparison with other approaches, we train the underlying models by LM-BFF with manual-compiled prompts without demonstration learning. URL: <https://github.com/princeton-nlp/LM-BFF>

¹⁰<https://github.com/THUDM/P-tuning>

¹¹In Zhong et al. (2021a), the authors only conduct zero-shot learning using larger PLMs. To make their work comparable to ours, we re-implement their algorithm over the Roberta model on our datasets under two settings. MT (Zero-shot) refers to the model tuned only using *dissimilar* full datasets. MT (Few-shot) further tunes the entire model over the target few-shot training set based on the prompts. Note that a few contemporaneous works (such as Wei et al. (2021)) also consider multi-task zero-shot learning. Because the settings and model scales are significantly different from ours, they are not directly comparable.

Paradigm	Method	Group 1: Sentiment.			MNLI	Group 2: NLI.			RTE	Group 3: Paraphrase.		Avg.
		SST-2	MR	CR		SNLI	QNLI	MRPC		QQP		
<i>Single-task methods w/o. the usage of dissimilar datasets (K = 16)</i>												
FT	Fine-tuning	81.1 ±4.1	78.2±5.4	75.4±3.3	45.8±6.0	48.4±4.8	60.9±5.8	54.0±6.1	74.4±2.5	61.0±4.1	64.4±4.7	
PT	PET	91.8±1.3	86.4±2.9	90.5±1.9	58.4±2.2	59.4±2.9	61.3±1.8	65.7±2.0	74.5±1.6	67.6±3.1	72.8±2.2	
	LM-BFF	92.0±1.7	87.4±0.7	90.8±1.0	65.2±2.6	71.7 ±4.9	69.1±2.8	69.5±2.0	74.2±2.3	63.5±1.2	75.9±2.4	
	P-Tuning	92.6±1.6	87.0±1.2	91.7±1.4	62.4±2.3	70.2±2.1	68.8±3.5	70.8 ±2.5	73.4±1.9	67.6±0.8	76.0±1.6	
	PPT	92.3±0.5	87.1±1.6	90.9±1.3	64.9±2.0	71.4±1.5	68.8±2.9	67.9±2.6	74.8±2.1	67.2±1.2	76.1±1.8	
	UPT-Single	92.9 ±1.0	87.7 ±1.5	91.8 ±0.7	65.6 ±1.4	71.2±2.3	70.1 ±1.6	68.9±1.7	75.1 ±0.9	72.1 ±2.0	77.2 ±1.5	
<i>Multi-task methods w. the usage of dissimilar datasets (K = 16)</i>												
PT	MT(Zero-shot)	58.7±1.6	59.0±3.6	58.9±2.8	36.3±3.3	39.2±3.2	40.9±2.5	54.9±1.4	70.6±2.6	42.8±2.5	51.3±2.2	
	MT(Few-shot)	92.1±1.4	86.5±1.3	91.0±2.2	69.6±1.1	67.1±2.7	68.9±2.3	68.6±1.2	71.0±1.4	74.8±2.1	76.7±1.7	
	UPT(Zero-shot)	74.5±1.2	73.9±1.3	72.4±1.4	43.7±2.0	46.0±2.1	53.9±1.9	57.1±1.0	70.7±0.9	56.5±1.3	61.0±1.5	
	UPT	93.5 ±0.6	88.1±0.9	91.4±1.2	70.1±1.4	68.2±1.2	69.9±1.5	73.5±1.5	77.0 ±1.1	78.8±1.7	78.9±1.4	
	UPT-SE	93.1±0.4	88.4 ±0.9	92.1 ±1.0	71.4 ±1.1	73.6 ±0.6	70.5 ±1.6	75.8 ±0.8	76.2±0.4	79.6 ±1.3	80.1 ±1.1	

Table 1: Comparison between *UPT* and baselines over all testing sets in terms of accuracy (%) and standard deviation. “FT” and “PT” refer to the *fine-tuning* and *prompt-based fine-tuning* paradigm, respectively. The methods in bold refer to our approach and its variants. The scores of baselines are re-produced using their open-source codes.

BERT Scale	SST-2	MR	CR	Avg.
Base	82.6±3.8	71.1±9.3	78.1±8.9	77.2±7.3
Medium	68.0±3.0	63.4±4.2	70.2±6.1	67.2±4.4
Small	66.3±3.7	58.1±4.6	68.2±5.5	64.2±4.6
Mini	58.8±3.1	59.4±7.6	65.8±7.5	61.3±6.1
Tiny	54.2±3.8	54.0±1.3	54.4±5.2	54.2±3.4

Table 2: Results of model scale analysis. We report the accuracy (%) of *UPT* based on BERT with other scales, and relative improvements, compared to the models w/o. prompt learning over *dissimilar* datasets.

and run with NVIDIA V100 GPUs. Specifically, we train our model with the Adam optimizer. The learning rate for all training stages is fixed to be $1e-5$. We set the default hyper-parameters as $\gamma = 0.001$ and $\lambda = 0.1$, which are also tuned over the development sets. The parameter regularizers are the same as in Gao et al. (2021).

3.2 Main Results

In Table 1, we report the general experimental results of *UPT* and all the baselines. The results show that: 1) Prompt-learning based methods (i.e., PET (Schick and Schütze, 2021a), LM-BFF (Gao et al., 2021), P-tuning (Liu et al., 2021c) and PPT (Gu et al., 2021)) have large improvements over standard *fine-tuning*. 2) *UPT*-Single outperforms previous few-shot learning models in average, which indicates that the utilization of *POV* is better than vanilla prompts (Schick and Schütze, 2021a). 3) *UPT* (both the vanilla and the ensemble version) consistently outperforms all baselines on all tasks, which demonstrates that our framework possesses better generalization by learning

from *dissimilar* groups of tasks¹². 4) MT (Zero-shot) (Zhong et al., 2021a) and *UPT* (Zero-shot) do not yields satisfactory results on BERT-style models. Different from ultra-large models, we suggest that few-shot prompt-tuning is necessary for BERT-style models to produce good results over these tasks. 5) By comparing *UPT* against MT (Few-shot), we can see that the proposed *POV* paradigm and the self-supervised *KSMLM* task are more effective for few-shot learning. 6) Generally, *UPT*-SE improves the averaged accuracy on all tasks by 1.2% than *UPT*. It means that self-ensemble learning can enhance model generalization, but the improvement is not consistent across all tasks. A possible cause is that some prompts and options are not optimal for the target task.

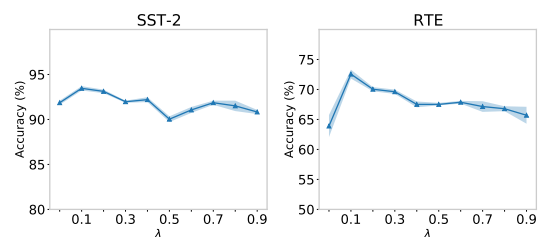


Figure 3: Parameter analysis w.r.t. hyper-parameter λ .

3.3 Model Analysis

Parameter Analysis. We conduct parameter analysis to investigate the best choice of the balance coefficient λ . Results over SST-2 and RTE are shown in Figure 3. We have the best performance when $\lambda = 0.1$, which indicates that our proposed *UPT*

¹²We also conduct the single-tail paired t-test to compare our approach against few-shot baselines across tasks. The result is $p < 0.05$, indicating the statistical significance.

Method/Group	Group 1	Group 2	Group 3
MT (Few-shot)	89.9	68.6	72.9
<i>UPT</i>	91.0	70.2	77.9
w/o. <i>POV</i>	90.2	68.9	74.2
w/o. <i>KSMLM</i>	90.9	69.1	73.7
w/o. <i>POV&KSMLM</i>	89.6	68.7	73.5
w/o. <i>OKR</i>	90.7	69.9	76.8

Table 3: Ablation study in terms of accuracy (%). Standard deviations are omitted here to save space.

possess generalization when it is jointly trained over the self-supervised *KSMLM* task. We also observe that the performance decreases when λ becomes larger. This means *KSMLM* is a suitable regularization task, but also may introduce a lot of prompts and options that are irrelevant to downstream tasks. This opens up new opportunities for model improvement.

Ablation Study. To clearly verify the contributions of each component in *UPT*, we conduct an ablation study over all groups and report the mean accuracy. As shown in Table 3, w/o. *POV* denotes the method with manually designed prompts without the usage of any options. w/o. *KSMLM* equals the setting with $\lambda = 0$, which is the same as *UPT*-Single. w/o. *OKR* means that we randomly choose the alternative label words in the options without knowledge guidance when we optimize the *KSMLM* task. w/o. *POV & KSMLM* denotes the method without any options and the auxiliary *KSMLM* task. The results show that no matter which module is removed, the model performance is affected. Particularly, when we remove both *POV* and *KSMLM*, the performance is decreased by 1.4%, 1.5%, 4.4%, respectively. The accuracy values of this setting are lower than w/o. *POV* and w/o. *KSMLM*, which suggests that both two components highly contribute to the high performance of our framework. We also find that w/o. *POV* or w/o. *KSMLM* both outperform MT (Few-shot) over all groups. Additionally, we find that if we use *KSMLM* but remove *OKR*, the results decrease over all these tasks, but are still higher than w/o. *KSMLM*. It means that the options knowledge that we mine from the corpus is suitable for the self-supervised learning task.

Sample Efficiency. We further explore the model effects with different numbers of training samples per class (K) from 16 to 512. We also use standard *fine-tuning* as the reference. As shown in Figure 4, each point refers to the averaged score across 5 randomly sampled datasets. We observe that our *UPT* consistently achieves higher scores regardless of

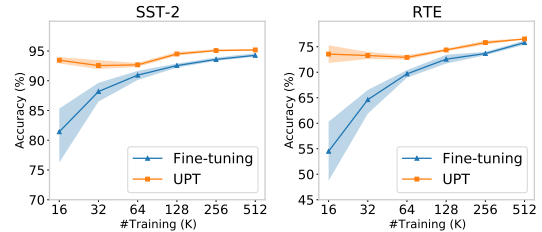


Figure 4: Results of sample efficiency analysis. We compare *UPT* with standard *fine-tuning* with different numbers of training samples K over two tasks.

the number of training samples. In addition, the variance of *UPT* is lower than *fine-tuning*, meaning that the stability of our method is better. This is different from other prompt-based methods (Schick and Schütze, 2021a,b; Gao et al., 2021).

Model Scale Analysis. To further show that *UPT* can improve the model performance regardless of the scales, we regard multiple small-scale BERT as model backbones¹³. Due to space limitations, we only illustrate the results in Table 2 over SST-2, MR, and CR. To make a fair comparison, we also test the performance without the usage of *dissimilar* NLP datasets and show the relative improvements. The results demonstrate that the model scale plays an important role in the ability of model generalization. We also find that *UPT* that uses *dissimilar* datasets can highly improve the effectiveness, especially on small-scale PLMs. Therefore, our method is better suitable for producing high-performing small PLMs for online applications.

Adaptation Efficiency of Task Groups. Because we focus on multi-task training before prompt-tuning over the target task in low-resourced settings. Therefore, it is worth exploring which/how many groups of tasks have a better effect on the adaptation improvement. Specifically, when given a target task (e.g., MNLI), we only choose one group of tasks (e.g., MRPC and QQP of Group 3 (Paraphrase)) for multi-task prompt-tuning, and then fine-tune the model on the target task. As shown in Figure 5, the cell in the i -th row and j -th column denotes the relative improvement from single-task learning over the j -th task to the setting where the i -th group is added for multi-task prompt learning. For visualization, we normalize the values of each column to show the percentage of influence of each group. The results show that the performance of a target task improves the most when we add data

¹³<https://github.com/google-research/bert>

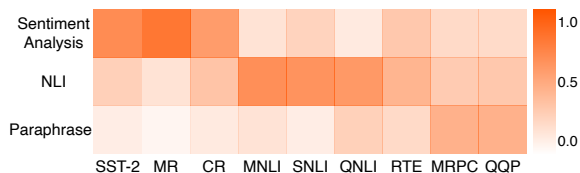


Figure 5: Adaptation efficiency between task groups. The shade of color indicates the degree of adaptation.

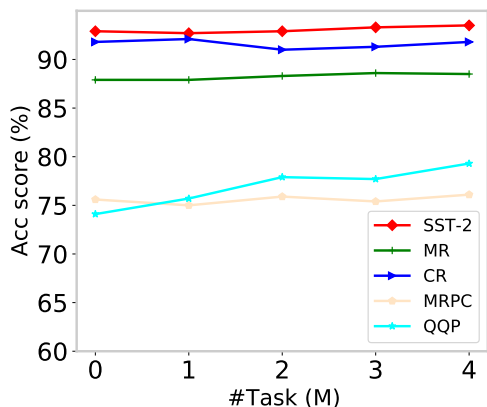


Figure 6: Adaptation efficiency between the different numbers of NLI tasks (M) and each target task from Sentiment and Paraphrase.

samples from other datasets within the same task group. However, in low-resourced scenarios, similar datasets are not available. By using *UPT*, we can even transfer the knowledge from datasets from *dissimilar* tasks to the target task.

Specifically, taking NLI as the source group, we randomly choose M dataset(s) from the group as our source tasks and then prompt-tune the model on each target task. The results from Figure 6 demonstrate that the accuracy is further improved when we increase the value M . We also find that the improvements over MRPC and QQP are more obvious. We suggest that NLI is easier to be adapted to paraphrase tasks because they both model the relations between sentence pairs.

4 Related Work

Pre-trained Language Models. Recently, benefited from the powerful modeling abilities of PLMs and computational resources, we have witnessed the qualitative improvement of multiple NLP tasks (Qiu et al., 2020; Han et al., 2021a). For examples, the large GPT model series (Radford et al., 2019; Brown et al., 2020) utilizes multi-layer transformer decoders to capture left-to-right semantics of natural languages. BERT (Devlin

et al., 2019) focuses on the learning of bidirectional contextual representations. Other notable PLMs include Transformer-XL (Dai et al., 2019), ELMo (Peters et al., 2018), RoBERTa (Liu et al., 2019), ALBERT (Lan et al., 2020), XLNet (Yang et al., 2019), StructBERT (Wang et al., 2019d), T5 (Raffel et al., 2020), etc. As model architecture is not the focus of our work, we do not elaborate.

Prompt-based Learning. Fine-tuning PLMs directly by learning the [CLS] head may perform poorly with few training samples (Liu et al., 2021a). Recently, the huge GPT-3 model (Brown et al., 2020) has been proposed to enable in-context learning, which introduces handcrafted prompts and demonstrations. Schick and Schütze (2021a) apply handcrafted prompts to prompt-based fine-tuning for BERT-style models. To facilitate the automatic prompt generation, Gao et al. (2021) present LM-BFF to generate discrete templates (Raffel et al., 2020). Other works (Shin et al., 2020; Han et al., 2021b; Scao and Rush, 2021; Utama et al., 2021) mine prompts from the training corpus based on heuristic rules/semantic relations. However, these methods are time-consuming for mining optimized prompts for target tasks. A series of methods are proposed to learn continuous/soft prompt embeddings, such as P-tuning (Liu et al., 2021c), P-tuning-V2 (Liu et al., 2021b), OptiPrompt (Zhong et al., 2021b), Prefix-tuning (Li and Liang, 2021). Zhao and Schütze (2021); Gu et al. (2021) focus on the hybrid training with both discrete and continuous prompts. Hu et al. (2021) consider the automatic expansion of label words and presents Knowledgeable Prompt-tuning (KPT) to utilize knowledge for the construction of verbalizers. Sun et al. (2021) and Wang et al. (2021b) prompt the PLMs to make language inference in zero-shot learning. In addition, Wang et al. (2021a); Vu et al. (2021) consider transfer learning on continuous prompt-tuning. Li et al. (2021); Chen et al. (2021); Ma et al. (2021) focus on prompts for specific NLP tasks, such as sentiment analysis and information extraction.

Recently, Wei et al. (2021); Zhong et al. (2021a); Min et al. (2021); Mishra et al. (2021) tune PLMs on mixed data samples drawn from different NLP tasks with manually designed task-specific prompts. The resulting PLMs are then utilized to solve unseen tasks by zero-shot learning. These methods successfully work for large PLMs such as GPT-3 (Brown et al., 2020) and T5 (Raffel et al., 2020), but consume a large amount of computation re-

sources. We further leverage data from non-target NLP tasks to make prompt-tuned PLMs have better capacities of adapting to unseen NLP tasks.

5 Conclusion and Future Work

In this paper, we present the *Unified Prompt Tuning* framework (*UPT*) that enables better few-shot text classification for BERT-style models by explicitly capturing prompting semantics from non-target datasets. We introduce a novel *POV* paradigm to unify the task format, and then extend the unified prompting to the self-supervised learning with the knowledge-enhanced selective MLM task. Experiments show that *UPT* consistently outperforms state-of-the-arts for prompt-based fine-tuning over multiple text classification scenarios. As for future work, we seek to extend *UPT* to other tasks such as named entity recognition, text generation, and machine translation. In addition, we will explore continuous prompt-tuning for *UPT*.

Limitations

Our work focuses on the prompt-based fine-tuning for text classification. It is also possible to extend our work to other NLP task (such as question answering, relation extraction, text generation, etc.) which will be addressed in the future work.

Ethical Considerations

Our contribution in this work is fully methodological, namely a unified prompt tuning (*UPT*) to boost the prompt-tuned PLMs. Hence, there are no direct negative social impacts of this contribution. However, as PLMs may have some negative impacts, such as the existence of social and gender biases, the tuned models produced by our *UPT* would unavoidably suffer from these issues. We suggest that users should carefully deal with the potential risks before deploying the models online.

Acknowledgments

This work has been supported by the National Natural Science Foundation of China under Grant No. U1911203, Alibaba Group through the Alibaba Innovation Research Program, the National Natural Science Foundation of China under Grant No. 61877018, the Research Project of Shanghai Science and Technology Commission (20dz2260300) and the Fundamental Research Funds for the Central Universities.

References

- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *EMNLP*, pages 632–642.
- 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 *NeurIPS*.
- Xiang Chen, Ningyu Zhang, Xin Xie, Shumin Deng, Yunzhi Yao, Chuanqi Tan, Fei Huang, Luo Si, and Huajun Chen. 2021. [Knowprompt: Knowledge-aware prompt-tuning with synergistic optimization for relation extraction](#). *CoRR*, abs/2104.07650.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2005. [The PASCAL recognising textual entailment challenge](#). In *MLCW*, volume 3944 of *Lecture Notes in Computer Science*, pages 177–190.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G. Carbonell, Quoc Viet Le, and Ruslan Salakhutdinov. 2019. [Transformer-xl: Attentive language models beyond a fixed-length context](#). In *ACL*, pages 2978–2988.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: pre-training of deep bidirectional transformers for language understanding](#). In *NAACL*, pages 4171–4186.
- William B. Dolan and Chris Brockett. 2005. [Automatically constructing a corpus of sentential paraphrases](#). In *IWP@IJCNLP 2005*.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. [Making pre-trained language models better few-shot learners](#). In *ACL*, pages 3816–3830.
- Yuxian Gu, Xu Han, Zhiyuan Liu, and Minlie Huang. 2021. [PPT: pre-trained prompt tuning for few-shot learning](#). *CoRR*, abs/2109.04332.
- Xu Han, Zhengyan Zhang, Ning Ding, Yuxian Gu, Xiao Liu, Yuqi Huo, Jiezhong Qiu, Liang Zhang, Wentao Han, Minlie Huang, Qin Jin, Yanyan Lan, Yang Liu, Zhiyuan Liu, Zhiwu Lu, Xipeng Qiu, Ruihua Song, Jie Tang, Ji-Rong Wen, Jinhui Yuan, Wayne Xin Zhao, and Jun Zhu. 2021a. [Pre-trained models: Past, present and future](#). *CoRR*, abs/2106.07139.
- Xu Han, Weilin Zhao, Ning Ding, Zhiyuan Liu, and Maosong Sun. 2021b. [PTR: prompt tuning with rules for text classification](#). *CoRR*, abs/2105.11259.

- Minqing Hu and Bing Liu. 2004. [Mining and summarizing customer reviews](#). In *KDD 2004*, pages 168–177.
- Shengding Hu, Ning Ding, Huadong Wang, Zhiyuan Liu, Juanzi Li, and Maosong Sun. 2021. [Knowledgeable prompt-tuning: Incorporating knowledge into prompt verbalizer for text classification](#). *CoRR*, abs/2108.02035.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. [ALBERT: A lite BERT for self-supervised learning of language representations](#). In *ICLR*.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#). *CoRR*, abs/2104.08691.
- Chengxi Li, Feiyu Gao, Jiajun Bu, Lu Xu, Xiang Chen, Yu Gu, Zirui Shao, Qi Zheng, Ningyu Zhang, Yongpan Wang, and Zhi Yu. 2021. [Sentiprompt: Sentiment knowledge enhanced prompt-tuning for aspect-based sentiment analysis](#). *CoRR*, abs/2109.08306.
- Xiang Lisa Li and Percy Liang. 2021. [Prefix-tuning: Optimizing continuous prompts for generation](#). In *ACL*, pages 4582–4597.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021a. [Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing](#). *CoRR*, abs/2107.13586.
- Xiao Liu, Kaixuan Ji, Yicheng Fu, Zhengxiao Du, Zhilin Yang, and Jie Tang. 2021b. [P-tuning v2: Prompt tuning can be comparable to fine-tuning universally across scales and tasks](#). *CoRR*, abs/2110.07602.
- Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2021c. [GPT understands, too](#). *CoRR*, abs/2103.10385.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.
- Ruotian Ma, Xin Zhou, Tao Gui, Yiding Tan, Qi Zhang, and Xuanjing Huang. 2021. [Template-free prompt tuning for few-shot NER](#). *CoRR*, abs/2109.13532.
- George A. Miller. 1995. Wordnet: A lexical database for english. *Commun. ACM*, 38(11):39–41.
- Sewon Min, Mike Lewis, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2021. [Metaicl: Learning to learn in context](#). *CoRR*, abs/2110.15943.
- Swaroop Mishra, Daniel Khashabi, Chitta Baral, Yejin Choi, and Hannaneh Hajishirzi. 2021. [Reframing instructional prompts to gptk’s language](#). *CoRR*, abs/2109.07830.
- Bo Pang and Lillian Lee. 2005. [Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales](#). In *ACL*, pages 115–124.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *NAACL*, pages 2227–2237.
- Xipeng Qiu, Tianxiang Sun, Yige Xu, Yunfan Shao, Ning Dai, and Xuanjing Huang. 2020. [Pre-trained models for natural language processing: A survey](#). *CoRR*, abs/2003.08271.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. [Language models are unsupervised multitask learners](#). *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Teven Le Scao and Alexander M. Rush. 2021. [How many data points is a prompt worth?](#) In *NAACL*, pages 2627–2636.
- Timo Schick and Hinrich Schütze. 2021a. [Exploiting cloze-questions for few-shot text classification and natural language inference](#). In *EACL*, pages 255–269.
- Timo Schick and Hinrich Schütze. 2021b. [It’s not just size that matters: Small language models are also few-shot learners](#). In *NAACL*, pages 2339–2352.
- Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh. 2020. [Autoprompt: Eliciting knowledge from language models with automatically generated prompts](#). In *EMNLP*, pages 4222–4235.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *EMNLP*, pages 1631–1642.
- Yi Sun, Yu Zheng, Chao Hao, and Hangping Qiu. 2021. [NSP-BERT: A prompt-based zero-shot learner through an original pre-training task-next sentence prediction](#). *CoRR*, abs/2109.03564.
- Prasetya Ajie Utama, Nafise Sadat Moosavi, Victor Sanh, and Iryna Gurevych. 2021. [Avoiding inference heuristics in few-shot prompt-based finetuning](#). *CoRR*, abs/2109.04144.
- Tu Vu, Brian Lester, Noah Constant, Rami Al-Rfou, and Daniel Cer. 2021. [Spot: Better frozen model adaptation through soft prompt transfer](#). *CoRR*, abs/2110.07904.

- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019a. Superglue: A sticker benchmark for general-purpose language understanding systems. In *NeurIPS*, pages 3261–3275.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019b. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *ICLR*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019c. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *ICLR*.
- Chengyu Wang, Minghui Qiu, Taolin Zhang, Tingting Liu, Lei Li, Jianing Wang, Ming Wang, Jun Huang, and Wei Lin. 2022. EasyNLP: A comprehensive and easy-to-use toolkit for natural language processing. *CoRR*, abs/2205.00258.
- Chengyu Wang, Jianing Wang, Minghui Qiu, Jun Huang, and Ming Gao. 2021a. [Transprompt: Towards an automatic transferable prompting framework for few-shot text classification](#). In *EMNLP*, pages 2792–2802.
- Sinong Wang, Han Fang, Madian Khabsa, Hanzi Mao, and Hao Ma. 2021b. [Entailment as few-shot learner](#). *CoRR*, abs/2104.14690.
- Wei Wang, Bin Bi, Ming Yan, Chen Wu, Zuyi Bao, Jiangnan Xia, Liwei Peng, and Luo Si. 2019d. [Structbert: incorporating language structures into pre-training for deep language understanding](#). *arXiv preprint arXiv:1908.04577*.
- Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2021. [Finetuned language models are zero-shot learners](#). *CoRR*, abs/2109.01652.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *NAACL*, pages 1112–1122.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. [Xlnet: Generalized autoregressive pretraining for language understanding](#). In *NeurIPS*, pages 5754–5764.
- Mengjie Zhao and Hinrich Schütze. 2021. [Discrete and soft prompting for multilingual models](#). In *EMNLP*, pages 8547–8555.
- Ruiqi Zhong, Kristy Lee, Zheng Zhang, and Dan Klein. 2021a. [Adapting language models for zero-shot learning by meta-tuning on dataset and prompt collections](#). In *EMNLP*, pages 2856–2878.
- Zexuan Zhong, Dan Friedman, and Danqi Chen. 2021b. [Factual probing is \[MASK\]: learning vs. learning to recall](#). In *NAACL*, pages 5017–5033.

A Dataset Statistics

In the main experiments, we employ 9 different NLP datasets for evaluation. As shown in Table 4, we divided all datasets into three groups, i.e., Sentiment, NLI, and Paraphrase. During multi-task training, we select two groups of tasks with full training data for *POV* prompt-tuning with the auxiliary *KSMLM* objective. After that, we prompt-tune the model over the target task in the few-shot learning setting. The corresponding group of the target task is unseen during multi-task training.

B POV Examples

As shown in Table 5, we list the designed *POVs* for all the tasks. Note that for each task group, the options are the same, but verbalizers of these tasks may be different. For example, SST-2, MR, and CR have the same schema of options, but with different verbalizers.

C Detailed Experiments for the *KSMLM* Task

We further conduct experiments over each group to evaluate the effectiveness of different settings in *KSMLM*. The baselines for comparison include:

- ***UPT w/o. KSMLM***: It means the training process on source tasks without the *KSMLM* learning objective before prompt-tuning over the target task.
- ***MLM***: It means that we directly train the vanilla MLM based on the full training data from source tasks.
- ***KSMLM (w/o. OKR)***: It means that we randomly select options without the K-Means algorithm and the knowledge-guided options construction process.
- ***KSMLM (w/o. Options)***: It means that we directly remove the options in *POV*.
- ***KSMLM (w/o. Verbalizer)***: It means that the prediction search space at each masked position is the whole BERT vocabulary rather than the designed limited collection of label words (expressed by options).

As shown in Table 7, we follow the same settings with the ablation study in Table 3 to report the mean accuracy values of each group. We can draw the following conclusions: 1) Compared to

Group	Category	Task	#Training	#Testing	N	Class Labels
Group 1: Sentiment	Single Sentence	SST-2	6,920	872	2	positive, negative
		MR	8,662	2,000	2	positive, negative
		CR	1,775	2,000	2	positive, negative
Group 2: NLI	Sentence Pair	MNLI	392,702	9,815	3	entailment, neutral, contradiction
		SNLI	549,367	9,842	3	entailment, neutral, contradiction
		QNLI	104,743	5,463	2	entailment, not entailment
		RTE	2,490	277	2	entailment, not entailment
Group 3: Paraphrase	Sentence Pair	MRPC	3,668	408	2	equivalent, not equivalent
		QQP	363,846	40,431	2	equivalent, not equivalent

Table 4: Dataset statistics. We only sample $N \times K$ instances from the original training sets to form the few-shot training and development sets. The testing sets used in the experiments are full datasets.

Task	Prompt	Option	Label word
SST-2	Template 1: [$\langle s1 \rangle$]. It was [MASK].	Option 1: Is $\langle x1 \rangle$ or $\langle x2 \rangle$? Option 2: Does $\langle x1 \rangle$ or $\langle x2 \rangle$? Option 3: $\langle x1 \rangle$ or $\langle x2 \rangle$?	Verbalizer 1: Negative (Bad), Positive (Wonderful) Verbalizer 2: Negative (Silly), Positive (Solid) Verbalizer 3: Negative (Pathetic), Positive (Irresistible)
	Template 2: [$\langle s1 \rangle$]. I thought it was [MASK].		
	Template 3: [$\langle s1 \rangle$]. It is [MASK].		
	Template 4: [$\langle s1 \rangle$]. The review is [MASK].		
	Template 5: [$\langle s1 \rangle$]. A [MASK] one.		
MR	Template 1: [$\langle s1 \rangle$]. It was [MASK].	Option 1: Is $\langle x1 \rangle$ or $\langle x2 \rangle$? Option 2: Does $\langle x1 \rangle$ or $\langle x2 \rangle$? Option 3: $\langle x1 \rangle$ or $\langle x2 \rangle$?	Verbalizer 1: Negative (Horrible), Positive (Exquisite) Verbalizer 2: Negative (Silly), Positive (Solid) Verbalizer 3: Negative (Bad), Positive (Wonderful)
	Template 2: [$\langle s1 \rangle$]. A [MASK] piece of work.		
	Template 3: [$\langle s1 \rangle$]. It is [MASK].		
	Template 4: [$\langle s1 \rangle$]. The film is [MASK].		
	Template 5: [$\langle s1 \rangle$]. A really [MASK] movie.		
CR	Template 1: [$\langle s1 \rangle$]. It was [MASK].	Option 1: Is $\langle x1 \rangle$ or $\langle x2 \rangle$? Option 2: Does $\langle x1 \rangle$ or $\langle x2 \rangle$? Option 3: $\langle x1 \rangle$ or $\langle x2 \rangle$?	Verbalizer 1: Negative (Horrible), Positive (Fantastic) Verbalizer 2: Negative (Silly), Positive (Solid) Verbalizer 3: Negative (Bad), Positive (Wonderful) Verbalizer 4: Negative (Pointless), Positive (Neat)
	Template 2: [$\langle s1 \rangle$]. It looks [MASK].		
	Template 3: [$\langle s1 \rangle$]. It is [MASK].		
	Template 4: [$\langle s1 \rangle$]. The quality is [MASK].		
	Template 5: [$\langle s1 \rangle$]. I thought it was [MASK].		
MNLI	Template 1: [$\langle s1 \rangle$]. You are right, [MASK], [$\langle s2 \rangle$].	Option 1: Is $\langle x1 \rangle$ or $\langle x2 \rangle$ or $\langle x3 \rangle$? Option 2: Based on the paragraph above, is the following $\langle x1 \rangle$ or $\langle x2 \rangle$ or $\langle x3 \rangle$?	Verbalizer 1: Contradiction (Next), Entailment (Exactly), Neutral (Indeed) Verbalizer 2: Contradiction (Wrong), Entailment (True), Neutral (Uncertain) Verbalizer 3: Contradiction (Otherwise), Entailment (Fine), Neutral (Plus) Verbalizer 4: Contradiction (Otherwise), Entailment (Exactly), Neutral (Naturally)
	Template 2: [$\langle s1 \rangle$]. It was [MASK], [$\langle s2 \rangle$].		
	Template 3: [$\langle s1 \rangle$], [$\langle s2 \rangle$]. It is [MASK].		
	Template 4: [$\langle s1 \rangle$]. It is true that [MASK], [$\langle s2 \rangle$].		
	Template 5: [$\langle s1 \rangle$]. [MASK]. Then, [$\langle s2 \rangle$].		
SNLI	Template 1: [$\langle s1 \rangle$]. [MASK], no, [$\langle s2 \rangle$].	Option 1: Is $\langle x1 \rangle$ or $\langle x2 \rangle$ or $\langle x3 \rangle$? Option 2: Based on the paragraph above, is the following $\langle x1 \rangle$ or $\langle x2 \rangle$ or $\langle x3 \rangle$?	Verbalizer 1: Contradiction (Next), Entailment (Exactly), Neutral (Indeed) Verbalizer 2: Contradiction (Wrong), Entailment (True), Neutral (Uncertain) Verbalizer 3: Contradiction (Instead), Entailment (Indeed), Neutral (Basically) Verbalizer 4: Contradiction (Except), Entailment (Alright), Neutral (Watch)
	Template 2: [$\langle s1 \rangle$]. [MASK], in this case, [$\langle s2 \rangle$].		
	Template 3: [$\langle s1 \rangle$]. [MASK], I think, [$\langle s2 \rangle$].		
	Template 4: [$\langle s1 \rangle$], [$\langle s2 \rangle$]. It was [MASK].		
	Template 5: [$\langle s1 \rangle$]. [MASK], [$\langle s2 \rangle$].		
QNLI	Template 1: Question: [$\langle s1 \rangle$]? [$\langle s2 \rangle$]. The answer: [MASK].	Option 1: Is $\langle x1 \rangle$ or $\langle x2 \rangle$? Option 2: Based on the question, is the following $\langle x1 \rangle$ or $\langle x2 \rangle$? Option 3: Is the answer $\langle x1 \rangle$ or $\langle x2 \rangle$?	Verbalizer 1: Entailment (Yes), Not Entailment (No) Verbalizer 2: Entailment (Okay), Not Entailment (Nonetheless) Verbalizer 3: Entailment (Notably), Not Entailment (Yet)
	Template 2: Question: [$\langle s1 \rangle$]? [$\langle s2 \rangle$]. [MASK].		
	Template 3: Question: [$\langle s1 \rangle$]? [MASK], Yes, [$\langle s2 \rangle$].		
	Template 4: [$\langle s1 \rangle$]? [MASK], it is known that [$\langle s2 \rangle$].		
	Template 5: [$\langle s1 \rangle$]? [MASK]. Then, [$\langle s2 \rangle$].		
RTE	Template 1: [$\langle s1 \rangle$], [$\langle s2 \rangle$]. The answer: [MASK].	Option 1: Is $\langle x1 \rangle$ or $\langle x2 \rangle$? Option 2: Based on the question, the answer is $\langle x1 \rangle$ or $\langle x2 \rangle$? Option 3: Is the answer $\langle x1 \rangle$ or $\langle x2 \rangle$?	Verbalizer 1: Entailment (So), Not Entailment (Meanwhile) Verbalizer 2: Entailment (Yes), Not Entailment (No) Verbalizer 3: Entailment (Notably), Not Entailment (Yet)
	Template 2: [$\langle s1 \rangle$], [$\langle s2 \rangle$]. [MASK].		
	Template 3: [$\langle s1 \rangle$]. [MASK], I think, [$\langle s2 \rangle$].		
	Template 4: [$\langle s1 \rangle$]. The question: [$\langle s2 \rangle$]? It is [MASK].		
	Template 5: [$\langle s1 \rangle$]. [MASK]. I believe, [$\langle s2 \rangle$].		
MRPC	Template 1: [$\langle s1 \rangle$], [$\langle s2 \rangle$]. The answer: [MASK].	Option 1: Is $\langle x1 \rangle$ or $\langle x2 \rangle$? Option 2: Are two question $\langle x1 \rangle$ or $\langle x2 \rangle$? Option 3: $\langle x1 \rangle$ or $\langle x2 \rangle$?	Verbalizer 1: 0 (Alas), 1 (Rather) Verbalizer 2: 0 (Different), 1 (Same) Verbalizer 3: 0 (Wrong), 1 (Right)
	Template 2: [$\langle s1 \rangle$], [$\langle s2 \rangle$]. [MASK].		
	Template 3: [$\langle s1 \rangle$]. [MASK], however, [$\langle s2 \rangle$].		
	Template 4: [$\langle s1 \rangle$], [$\langle s2 \rangle$]. In fact [MASK].		
	Template 5: [$\langle s1 \rangle$]. [MASK]. that's right, [$\langle s2 \rangle$].		
QQP	Template 1: [$\langle s1 \rangle$], [$\langle s2 \rangle$]. The answer: [MASK].	Option 1: Is $\langle x1 \rangle$ or $\langle x2 \rangle$? Option 2: Are two question $\langle x1 \rangle$ or $\langle x2 \rangle$? Option 3: $\langle x1 \rangle$ or $\langle x2 \rangle$?	Verbalizer 1: 0 (Alas), 1 (Rather) Verbalizer 2: 0 (Different), 1 (Same) Verbalizer 3: 0 (Wrong), 1 (Right)
	Template 2: [$\langle s1 \rangle$], [$\langle s2 \rangle$]. [MASK].		
	Template 3: [$\langle s1 \rangle$]. [MASK], however, [$\langle s2 \rangle$].		
	Template 4: [$\langle s1 \rangle$], [$\langle s2 \rangle$]. In fact [MASK].		
	Template 5: [$\langle s1 \rangle$]. [MASK]. that's right, [$\langle s2 \rangle$].		

Table 5: The Prompts, Options and Verbalizers (POV) for each task. $\langle s1 \rangle$ and $\langle s2 \rangle$ denote the input sentences. $\langle x1 \rangle$, $\langle x2 \rangle$ and $\langle x3 \rangle$ denote the label words.

Paradigm	Method	AX-b	AX-g	BoolQ	CB	SST-5	TREC	Subj	Avg.
FT	Fine-tuning	47.51±1.8	60.83±2.0	65.96±1.3	73.21±1.3	40.10±3.4	59.30±1.8	73.00±2.0	59.99
PT	PET	60.28±1.2	64.08±0.8	70.54±1.6	82.09±2.0	44.10±1.7	84.90±1.9	89.30±1.0	70.76
	LM-BFF	61.53±1.4	63.89±1.9	71.30±2.1	82.14±2.6	46.10±1.3	84.80±1.5	89.25±1.0	71.29
	P-Tuning	62.23±0.8	63.19±1.2	72.88±0.9	83.08±1.8	48.20±1.5	85.10±1.9	89.35±1.1	72.00
	<i>UPT</i>	64.25±1.2	69.44±1.4	74.06±1.6	83.92±0.9	48.35±1.0	85.90±0.8	90.15±1.2	73.72

Table 6: Additional experiments for comparison between *UPT* and baselines over all testing sets in terms of accuracy (%) and standard deviation.

Method/Group	Group 1	Group 2	Group 3
<i>UPT</i>	91.0	70.2	77.9
<i>UPT</i> w/o. <i>KSMLM</i>	90.9	69.1	73.7
MLM	87.1	67.4	72.0
<i>KSMLM</i> (w/o. <i>OKR</i>)	90.7	69.9	76.8
<i>KSMLM</i> (w/o. Options)	90.1	68.2	76.3
<i>KSMLM</i> (w/o. Verbalizer)	85.0	62.4	66.7

Table 7: The ablation analysis of the *KSMLM* task in terms of accuracy (%).

Paradigm/Task	SST-2	MR	CR	Avg.
<i>POV</i>	92.9	87.7	91.8	90.8
Multiple-choice	82.7	73.9	80.9	79.2
Yes/No	92.6	87.9	91.6	90.7

Table 8: The comparison between different paradigms in terms of accuracy (%).

vanilla MLM, the results indicate that *KSMLM* is an irreplaceable task for the improvement of the model generalization power. 2) We also find that if we ignore the verbalizer construction, the results decrease to a large degree, and lower than *UPT* w/o. *KSMLM*. It means that verbalizers are crucial for template-based prompt-tuning. 3) When *OKR* or options are removed, the results also decline, indicating the effectiveness of these techniques.

D Comparing *POV* with Other Paradigms

To compare the proposed *POV* paradigm with other paradigms, we perform experiments over SST-2, MR, and CR tasks. The alternative paradigms are as follows:

- **Multiple-choice.** It is a unified template to list all the candidate results. For example, an input sentence can be “The Disney cartoons are very interesting for children to enrich their extracurricular life. A. great; B. terrible. It is [MASK].”. This paradigm is closely in line with PPT (Gu et al., 2021).

- **Yes/No.** We can reformulate the multi-class classification tasks into a series of binary classification. Take NLI for example. We can design three templates for each class, i.e “Are these descriptions are entailment?”, “Are these descriptions are neutral?”, and “Are these descriptions are contradiction?”. We follow Zhong et al. (2021a) to add an MLP layer on the top of the PLM to obtain the output of the [MASK] token to classify the answer to be “Yes” or “No”.

Experimental results in Table 8 show that in average, *POV* outperforms all baselines. For Multiple-choice, we find the results decline a lot. We guess that the PLM is hard to understand and generate the items number, such as “A, B, C, D”. In addition, we find the paradigm “Yes/No” has a similar performance to *POV*. Overall, the experiments prove the effectiveness of *POV*, which is easy to implement and avoids the transformation to multiple binary classification tasks for tasks with multiple classes.

E Additional Evaluation Results over Other Tasks

In this part, we further present additional experiments over other tasks from GLUE (Wang et al., 2019c) and SuperGLUE (Wang et al., 2019a), including AX-b, AX-g, BoolQ, CB, SST-5, TREC and Subj. The data statistics can be found in the original papers. We choose standard fine-tuning, PET (Schick and Schütze, 2021a), LM-BFF (Liu et al., 2021c) as our baselines to make comparison. In this experiment, we only conduct task-specific single-task learning to evaluate the efficiency of the *POV* paradigm. We also set $K = 16$. As shown in Table 6, we can draw the following conclusions. 1) Our *UPT* framework outperforms strong baselines over these tasks. 2) SST-5 and TREC are challenging tasks with many labels, which consist of 5 and 6 classes, respectively. Experiments show that our *POV* paradigm can achieve the best performances.