

# AdaPrompt: Adaptive Model Training for Prompt-based NLP

Yulong Chen<sup>♣\*</sup>, Yang Liu<sup>♣</sup>, Li Dong<sup>♣</sup>, Shuohang Wang<sup>♣</sup>,  
Chenguang Zhu<sup>♣</sup>, Michael Zeng<sup>♣</sup>, Yue Zhang<sup>♡◇</sup>

<sup>♣</sup> Zhejiang University    <sup>♡</sup> Westlake University

<sup>♣</sup> Microsoft Research    <sup>◇</sup> Westlake Institute for Advanced Study

[yulongchen1010@gmail.com](mailto:yulongchen1010@gmail.com)    [yaliu10@microsoft.com](mailto:yaliu10@microsoft.com)    [yue.zhang@wias.org.cn](mailto:yue.zhang@wias.org.cn)

## Abstract

Prompt-based learning, with its capability to tackle zero-shot and few-shot NLP tasks, has gained much attention in community. The main idea is to bridge the gap between NLP downstream tasks and language modeling (LM), by mapping these tasks into natural language prompts, which are then filled by pretrained language models (PLMs). However, for prompt learning, there are still two salient gaps between NLP tasks and pretraining. First, prompt information is not necessarily sufficiently present during LM pretraining. Second, task-specific data are not necessarily well represented during pretraining. We address these two issues by proposing AdaPrompt, adaptively retrieving external data for continual pretraining of PLMs by making use of both task and prompt characteristics. In addition, we make use of knowledge in Natural Language Inference models for deriving adaptive verbalizers. Experimental results on five NLP benchmarks show that AdaPrompt can improve over standard PLMs in few-shot settings. In addition, in zero-shot settings, our method outperforms standard prompt-based methods by up to 26.35% relative error reduction.

## 1 Introduction

Prompt-based methods (Brown et al., 2020; Liu et al., 2021; Schick and Schütze, 2021a; Li and Liang, 2021) have received increasing attention in Natural Language Processing (NLP) recently. The main idea is to make the most use of pre-trained language models (PLMs) by adapting an NLP task into a natural language prompt, which can then be filled by PLMs. Take sentiment classification (Socher et al., 2013; Bai et al., 2021) for example. Given the sentence “*I love the movie.*”, the standard task is to make a binary classification on its sentiment polarity (i.e., positive or negative).

\*Yulong Chen completed this work during his internship at Microsoft.

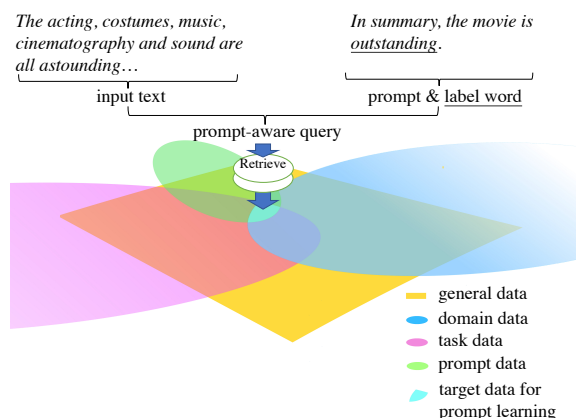


Figure 1: The distributions of data in prompt-based models. Task data, domain data, prompt data, and general data (for LM pretraining) are usually sampled from different distributions while remaining certain overlap (target data for prompt training). We aim to explore data from the overlapping area to bridge the gap between PLM and downstream tasks in prompt-based systems.

Prompt-based methods first transform the sentence into “*I love the movie. The movie is <mask>.*” (the underlined text is called prompt), and then identify its polarity by checking whether PLMs tends to predict “good” or “bad” for the <mask> token (where the predicted words are then verbalized into class labels). The prompt-based task formulation is close to masked language modeling (Schick and Schütze, 2021a,b), which is the mainstream pretraining strategy, allowing PLMs to provide rich language knowledge seamlessly. Prompt-based methods have been shown particularly useful in zero-shot and few-shot settings (Petroni et al., 2019; Yin et al., 2019; Min et al., 2022), where with limited direct task data, prompt-based inference benefits more from large-scale pretraining than task-oriented fine-tuning.

Existing methods, however, still suffer from several potential limitations. First, large raw text data used for pretraining do not necessarily contain sufficient patterns that are directly related to

task specific prompts (Illustrated in Figure 1). For instance, the prompt for a question classification task is “*Can you tell me the  $\langle \text{mask} \rangle$ : What are the twin cities?*”, where  $\langle \text{mask} \rangle$  should be a class label word, e.g., *location, person* and etc (the correct label for this sample is *definition*). However, LM pretraining data are typically BOOKCORPUS (Zhu et al., 2015) plus WIKIPEDIA corpus, where such prompts can occur scarcely in the literal or paraphrased form. As a result, directly using PLMs to fill such handcrafted prompts across domains can lead to poor performance. Second, to project label words to task labels, most existing work (Schick and Schütze, 2021a,b; Cui et al., 2021) uses a pre-defined verbalizer. However, it often requires expert knowledge to build a verbalizer that can thoroughly cover candidate words and a poorly-designed verbalizer limits the accuracy of predictions. These problems become even more serious under zero-shot or very-few-shot settings, where prompt-based models highly rely on the generalization ability of PLMs to new tasks and domains.

We propose AdaPrompt, a framework that adapts PLMs for end tasks considering both the prompts and the verbalizer. We are interested in addressing the above issues under a zero-shot setting, where little or no labeled training data are available for a particular task. The main idea is to adapt a PLM to a strong prompt-based model for an end task by exploring knowledge from its raw input data. In particular, as shown in Figure 2, given a raw test set without labels, we first ask a PLM to fill a prompt template for each input (e.g., “*In summary, the movie is great.*”, where “*great*” is filled by PLMs). Then, we use the resulting text (input text + prompt + PLM output) as a prompt-aware query to retrieve relevant data from a large unlabeled corpus. In this manner, we can obtain a large dataset that contain both task and prompt characteristics, and we adaptively continual pretrain (Gururangan et al., 2020) the PLM on the retrieved data, which can substantially benefit prompt-based methods on downstream NLP tasks.

Meanwhile, we found current way of building verbalizers is also not optimal. Given a specific task, different words can be verbalized into the same class labels. For example, a large number of adjectives can express the positive sentiment, and the best-performing candidates depend on the domain, PLM and context. In AdaPrompt, we propose to adaptively augment verbalizers by making

use of knowledge from PLMs and Natural Language Inference (NLI) models. Take sentiment analysis for example, given “*good*” and “*bad*” as seed verbalizers, we first let PLMs to predict more candidate words, such as “*amazing*” and “*great*”. Then, to identify if these candidates are suitable to verbalizer, we refer to an NLI model to predict whether “*This movie is amazing.*” entails the meaning of “*This movie is good.*”. In this way, we can automatically expand the verbalizers.

Experiments on five text classification tasks show that AdaPrompt outperforms baseline prompt-based methods by 2.29%-5.79% in very-few-shot setting and 2.46%-15.00% in zero-shot setting on accuracy. To our knowledge, we are the first to consider how to bridge the gap between LM pretraining and NLP downstream tasks for prompt-based NLP. We release our code and data at <https://github.com/cylnlp/AdaPrompt>.

## 2 Related work

### 2.1 Zero/Few-shot Prompt-based NLP

Although prompt-based methods have been used for multiple NLP tasks (Brown et al., 2020; Raffel et al., 2020; Brown et al., 2020; Cui et al., 2021), most of existing work focus on text classification (Shin et al., 2020; Gao et al., 2021; Min et al., 2022; Hu et al., 2022). A typical related work is PET (Schick and Schütze, 2021a), where Schick and Schütze (2021a) formally define *pattern-verbalizer pairs* that have been widely adopted by successive works. By using such pairs, Schick and Schütze (2021a,b) develop a series of work to explore the potential of PLMs, including annotating soft labels for raw training data, and data augmentation iteratively. However, different from PET that assumes the availability of large silver training set for downstream tasks, we focus on zero and very-few-shot settings, where even unannotated task-relevant dataset is also limited (Perez et al., 2021). Therefore, following Hu et al. (2022), we simply focus on standard pattern-verbalizer pairs for text classification.

Prompt engineering (Jiang et al., 2020; Gao et al., 2021) focuses on how to create prompts that can better induce PLMs to make correct predictions. Discrete prompt engineering works by replacing, deleting, inserting or paraphrasing parts of the prompt (Wallace et al., 2019; Yuan et al., 2021). Those methods can efficiently adapt PLMs to end tasks, but they highly rely on annotated data for

tuning parameters. Different from the above studies, we are interested in narrowing the gap between LM pretraining and NLP tasks for prompting learning in zero or very-few-shot settings.

It has been shown that using different verbalizers can also be a key factor for prompt learning (Hu et al., 2022; Cui et al., 2021). However, manually exploring label words is time-consuming and may neglect potential candidates. Recently, Hu et al. (2022) uses multiple external knowledge bases, such as related words and sentiment dictionaries, to augment verbalizers for corresponding tasks. Different from them, we focus on exploring knowledge in PLMs themselves. By making use of external NLI models AdaPrompt can select verbalizers automatically without the need of labeled task data, which is useful in zero-shot settings.

## 2.2 Continual Pretraining for Domain Adaptation

Continual pretraining (Gururangan et al., 2020) has shown benefit of optimizing a PLM to a target domain before further fine-tuning. It can be categorised into domain adaptive continual pretraining and task adaptive continual pretraining. The difference is that, domain adaptive pretraining (DAPT) uses domain relevant data while task adaptive pretraining (TAPT) uses task-specific data.

Similar to continual pretraining, many recent methods highlight the merits of relying on language modeling objectives for domain adaptation. Chronopoulou et al. (2019) and Radford et al. (2018) propose to train task-specific parameters for PLMs by using an auxiliary LM loss on target domains. Models like SciBERT (Beltagy et al., 2019), DialogLM (Zhong et al., 2021), AMRBART (Bai et al., 2022a), SARA-BERT (Bai et al., 2022b) and Dict-BERT (Yu et al., 2022) are PLMs that are continually pretrained on large amounts of domain/task-specific corpora.

Data selection is a common practice in domain adaption for NLP models (Moore and Lewis, 2010; Ruder and Plank, 2017; van der Wees et al., 2017). It has been used in machine translation (van der Wees et al., 2017; Wang et al., 2018), parsing (Plank and van Noord, 2011; Ruder and Plank, 2017) and sentiment analysis (Ruder et al., 2017). The main idea is to have a selection model that can distinguish in-domain and out-of-domain data. The selection model can be a supervised classifier (Aharoni and Goldberg, 2020), similarity-

based metric (Plank and van Noord, 2011) or language model perplexity (Moore and Lewis, 2010). Very recently, Yao et al. (2021) propose to retrieve a small set of training data from general corpora with labeled task data as queries, finding that using LM objective on this data as an auxiliary loss can help train task-specific NLP models without pretraining.

## 3 Method

Our method is based on prompt-based text classification methods (Section 3.1). The overall procedure of AdaPrompt is shown in Figure 2, which can be divided into two parts: PLM adaptation (Section 3.2) and verbalizer adaptation (Section 3.4). In Section 3.3, we introduce a method that adapts both PLMs and verbalizers in an iterative way for continual improvements.

### 3.1 Prompt-based Text Classification

Given an input text,  $\mathbf{x} = (x_0, x_1, \dots, x_n)$ , we consider various tasks to classify the sentence into a class label  $l \in \mathcal{L}$ . As mentioned in Section 1, the standard prompt-based method reformulates the input into a cloze-style question and identifies its label by checking PLMs’ predictions. Table 1 shows the prompt templates and verbalizer patterns for the SST-2 (Socher et al., 2013), Yelp (Zhang et al., 2015), AGNews (Zhang et al., 2015), TREC (Voorhees and Tice, 2000) and DBPedia (Lehmann et al., 2015) datasets, which cover sentiment classification, topic classification and question classification tasks. Formally, let  $\mathcal{M}$  be a language model pretrained on large-scale general data, and  $\langle \text{mask} \rangle$  be the mask token. The prompt-based method first defines a *pattern* function, *Prompt*, that converts  $\mathbf{x}$  into a cloze-style question containing  $\langle \text{mask} \rangle$ . Then, it defines a *verbalizer* function  $v$ , which maps a small set of pre-defined verbalizer words ( $\mathcal{Y}$ ) predicted at the position of  $\langle \text{mask} \rangle$  into class labels, i.e.,  $v : \mathcal{Y} \mapsto \mathcal{L}$ .

Take sentiment classification for movie review for instance. The task is to classify the sentiment polarity, where  $\mathcal{L} = \{\text{positive}, \text{negative}\}$ . For an input  $\mathbf{x}$ , we choose the pattern:

*Prompt* = “ $\mathbf{x}$ . In summary, the movie is  $\langle \text{mask} \rangle$ .”

Then we define a verbalizer that maps  $\mathcal{Y} = \{\text{“good”}, \text{“bad”}\}$  into  $\mathcal{L}$ :

$v(\text{“good”}) = \text{positive}$ ;  
 $v(\text{“bad”}) = \text{negative}$

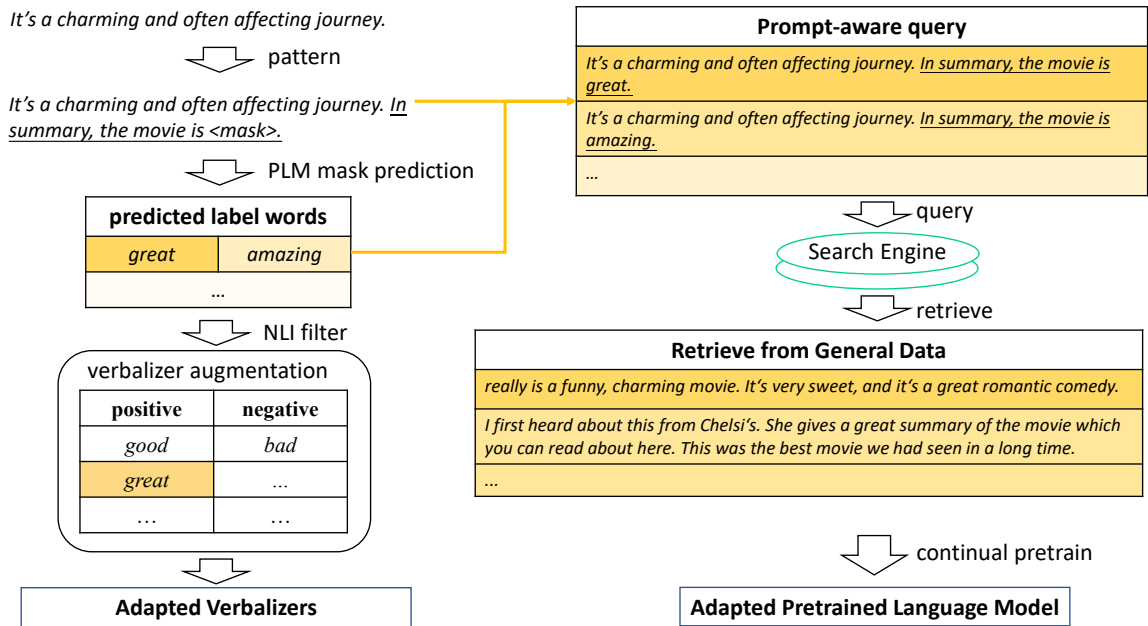


Figure 2: Overall framework of AdaPrompt.

Given an example:

$\mathbf{x} = \text{“}It’s a charming journey.\text{”}$ ,

we can convert the input into a cloze-style question using *Prompt*:

$Prompt(\mathbf{x}) = \text{“}It’s a charming journey. In summary, the movie is \langle mask \rangle.\text{”}$

Using such *pattern-verbalizer* pairs, we ask  $\mathcal{M}$  to directly give scores  $s$  for each label  $l \in \mathcal{L}$  as:

$$s(l|\mathbf{x}) = Pr[\langle mask \rangle = y | Prompt(\mathbf{x}), \mathcal{M}] \quad (1)$$

where  $l = v(y)$ . The predicted label is:

$$\hat{l} = \arg \max_{l \in \mathcal{L}} s(l|\mathbf{x}) \quad (2)$$

### 3.2 Adaptively Retrieve Data for Continual Pretraining

As discussed in the Section 1, the lack of domain adaptation can be a potential challenge for prompt-based NLP models, especially under zero-shot and very-few-shot settings. To tackle this problem, we propose to build a continual pretraining dataset by retrieving from general corpora, with unannotated test texts, designed prompts and label words as queries. In this way, we can obtain task-relevant data for any tasks or domains, using only test input. Meanwhile, prompt and verbalizer information is also considered during the retrieval process, leading to a more comprehensive dataset for prompt-aware continual pretraining.

Formally, given a retrieval query  $q$ , a retrieval engine  $\mathcal{E}_D$  indexed on a large general dataset  $\mathcal{D}$  can

return a set of similar text  $d_q = \mathcal{E}_D(q)$ . To obtain prompt-aware data that can not only adapt PLMs to target domains but also make PLMs more sensitive to prompts, we include both task and prompt characteristics when building queries. As shown in Figure 2, for a raw input text  $\mathbf{x}$  in text data, we first convert it into  $Prompt(\mathbf{x})$ , and obtain a set of predicted label words using a PLM  $\mathcal{M}$ :

$$\mathcal{O} = \mathcal{M}(Prompt(\mathbf{x})) \quad (3)$$

where  $\mathcal{O} = \{o_1, o_2, \dots, o_{|\mathcal{O}|}\}$  are the top- $|\mathcal{O}|$  predictions. We replace the mask token in  $P(\mathbf{x})$  with  $o_i$ , to form a list  $Q$  of queries. For example:

$$Q = \{q_1, \dots, q_{|\mathcal{O}|}\}, \quad (4)$$

where  $q_i = \text{“}\mathbf{x}. In summary, the movie is o_i.\text{”}$

With this set of prompt-based queries, we retrieve prompt-aware data  $\mathcal{D}_p$ , which is a small subset of the general data. In this work, we use ElasticSearch<sup>1</sup> indexed on a large general corpus as the search engine and we ask it to return a list of top- $k$  texts that match the query. As shown in Figure 2, one test input can lead to multiple prompt-aware queries because the masked token in the prompt can be replaced by the  $|\mathcal{O}|$  predictions. In addition, given one query, ElasticSearch can also give multiple returns with demanded  $k$ .

We continue to pretrain the PLM  $\mathcal{M}$  on  $\mathcal{D}_p$  with masked language modeling loss and obtain

<sup>1</sup><https://www.elastic.co>

---

**Algorithm 1** Verbalizer Adaptation

---

**Input:** prompt  $P$ , seed verbalizer words  $y \in \mathcal{Y}_l$ , candidate words  $c \in \mathcal{C}$  and an NLI system  $\mathcal{N}$   
**for**  $c$  **in**  $\mathcal{C}$  **do**  
  **if**  $\mathcal{N}(f(P, y), fill(P, c)) = Entail$   
  **or**  $\mathcal{N}(fill(P, c), f(P, y)) = Entail$  **then**  
    add  $c$  in  $\mathcal{Y}_l$   
  **end if**  
**end for**  
Return  $\mathcal{Y}_l$

---

an adapted PLM  $\mathcal{M}_{\mathcal{D}_p}$ .  $\mathcal{M}_{\mathcal{D}_p}$  now contains richer knowledge of both the target domain and the prompts. It can be used to replace  $\mathcal{M}$  in Eq. 1 for zero-shot text classification.

### 3.3 Iterative Adaptation

After obtaining  $\mathcal{M}_{\mathcal{D}_p}$ , we can iterate the process by replacing  $\mathcal{M}$  with  $\mathcal{M}_{\mathcal{D}_p}$  in Eq. 3, and obtain an iterative set of predicted words and a list of queries marked as  $\mathcal{O}'$  and  $Q'$ . Given that  $\mathcal{O}'$  contains more in-domain knowledge, we can retrieve higher quality pretraining data with more task relevant information, using  $Q'$  to query the  $\mathcal{E}_D$ . In this way, we obtain a new version of  $\mathcal{D}'_p$ , and a new continual pretrained PLM  $\mathcal{M}'_{\mathcal{D}_p}$ , which can also be used for zero-shot predictions using Eq. 1. In this work, we conduct this procedure twice.

### 3.4 Adaptive Verbalizer Augmentation

As described in Section 3.1, the regular prompt-based method defines the *verbalizer* that maps predicted label word into task classes, such as “good” for positive and “bad” for negative. However, predefined verbalizer can be limited. To expand this verbalizer, we first infer top- $|\mathcal{O}|$  label words at mask token position over all inputs in test set. We filter the predicted words and obtain a set of high frequent words  $\mathcal{C}$  as candidates for verbalizer augmentation. Then, we propose a new method for exploring useful verbalizer words by using knowledge from a Natural Language Entailment model.

Specifically, given a seed verbalizer word  $y_l \in \mathcal{Y}_l$  for label  $l$ , and a candidate word  $c \in \mathcal{C}$ , we compare whether a prompt filled by  $y_l$  is entailed with the prompt filled by  $c$ . The pseudo code is shown in Algorithm 1. If entailment relation holds for this pair, we add  $c$  to  $\mathcal{Y}_l$ . And the new  $\mathcal{Y}$  which can be considered as an augmented verbalizer.

After obtaining the augmented set of verbalizer

words, Eq. 1 can be rewritten as:

$$s(l|\mathbf{x}) = \frac{1}{|\mathcal{Y}_l|} \sum_{y \in \mathcal{Y}_l} Pr[\langle \text{mask} \rangle = y | Prompt(\mathbf{x}), \mathcal{M}] \quad (5)$$

and we can still use Eq. 2 for prediction.

## 4 Experiments

### 4.1 Datasets and Prompts

To evaluate our methods, we conduct experiments on five benchmarks: SST-2 (Socher et al., 2013), Yelp (Zhang et al., 2015), AGNews (Zhang et al., 2015), TREC (Voorhees and Tice, 2000) and DBPedia (Lehmann et al., 2015) datasets. Table 1 shows prompt templates and seed verbalizer words that we use for each dataset. For AGNews and YELP, we adapt patterns and verbalizers from PET (Schick and Schütze, 2021a) since it is the basic prompt-based method that has been mostly widely used.

**AGNews** is a text classification dataset in the domain of News. Given a headline and a main text body, the model is required to classify the news into one of the classes: (1) World, (2) Sports, (3) Business or (4) Science/Tech.

**YELP** is a sentiment analysis dataset. Given a restaurant review, the task is to predict whether the review is positive or negative.

**SST-2** is a sentiment analysis dataset similar to YELP but its domain is movie reviews. Thus, we use the same seed prompt and verbalizer words as for YELP, but change “restaurant” in prompt template to “movie”.

**DBPedia 2014** is an ontology classification dataset, extracted from DBPedia 2014 with 14 non-overlap classes, such as Educational Institution and Office Holder. We define two patterns for this task:

$$P1(\mathbf{x}) = \text{“Description to the } \langle \text{mask} \rangle \mathbf{x} \text{”}$$

$$P2(\mathbf{x}) = \text{“Introduction to the } \langle \text{mask} \rangle \mathbf{x} \text{”}$$

and we use  $P2$  as the seed pattern.

**TREC-10** is a question classification dataset. Given a question, the task is to identify the objective that the question asks, and classify it into one of six classes, such as a definition question or a numeric question. We define two patterns for this task:

$$P1(\mathbf{x}) = \text{“Tell me the } \langle \text{mask} \rangle \mathbf{x} \text{”}$$

$$P2(\mathbf{x}) = \text{“Can you tell me the } \langle \text{mask} \rangle \mathbf{x} \text{”}$$

and  $P2$  as the seed prompt.

### 4.2 Settings

In this work, we take ROBERTA-large (Liu et al., 2019) as our foundation PLM and adopt pattern-verbalizer pairs from (Schick and Schütze, 2021a)

Dataset	Class	Objective	Prompt Template	Verbalizer
SST-2	2	sentiment	Text <i>In summary, this movie is</i> $\langle mask \rangle$ .	“good”, “bad”
Yelp	2	sentiment	Text <i>In summary, this restaurant is</i> $\langle mask \rangle$ .	“good”, “bad”
AGNews	4	news topic	[Category: $\langle mask \rangle$ ] Title, Body	“Sport”, “Tech”, “Business”, “World”
TREC	6	question	<i>Can you tell me the</i> $\langle mask \rangle$ Text	“explanation”, “description”, “person”, “location”, “number”, “entity”
DBPedia	14	ontology	<i>Introduction to the</i> $\langle mask \rangle$ Text	“company”, “school”, “artist”, “film”, “book”, “plan”, “building”, “village”, “animal”, “sport”, “album”, “officer”, “scenery”, “transportation”

Table 1: Datasets used in this paper with seed prompts and verbalizer words. Each seed verbalizer word corresponds to a class label.

Dataset	test set	Top- $ \mathcal{O} $	$E_{space}$	Resulting Data
TREC	500	20	100	60k
SST-2	872	20	100	205k
AGNews	7,600	10	50	414k
YELP	38,000	1	50	267k
DBPedia	70,000	1	50	1,301k

Table 2: Data statistics for datasets.  $E_{space}$  corresponds to the ElasticSearch space. Note that the resulting data size is calculated after data de-duplication.

(Section 3.1) as the baseline setting which is widely used and can be easily extended to other methods (Shin et al., 2020).

We conduct experiments in zero-shot and few-shot settings. In the zero-shot setting, we directly use PLMs to infer label words at masked positions. Under the few-shot setting, we follow Schick and Schütze (2021a) and Hu et al. (2022) and use prompt-tuning, which directly fine-tunes a LM given a small set of annotated data and prompts.

For zero-shot settings, the choice of hyperparameters is based on previous work (Gao et al., 2021; Schick and Schütze, 2021a,b). For all continual pretraining, we use a learning rate of  $1e^{-5}$ , batch size of 96. We train each model for 3 epochs and use the checkpoint at 500 steps for evaluation.

For few-shot settings, we evaluate our models with 10, 50, 100 training samples. We follow previous work (Hu et al., 2022; Schick and Schütze, 2021a; Gao et al., 2021) and repeat the training and evaluation for 5 times using different seed, and report the averaged scores for each datasets.

**Prompt-Aware Data Retrieval** We take pre-train data of the ROBERTA model (BOOKCORPUS (Zhu et al., 2015), WIKIPEDIA, CCNEWS (Nagel, 2016), STORIES (Trinh and Le, 2018), and OPENWEBTEXT (Gokaslan and Cohen, 2019)) as the general dataset to query from. We index them on sentence level with ElasticSearch and consider TF-IDF as the similarity metric.

Table 2 presents the statistics of evaluation datasets used in this paper. TREC and SST contain smaller test sets, while YELP and DBPedia contain much larger test sets. To balance the retrieved data size, we set different top- $|\mathcal{O}|$  for predicted words and ElasticSearch space ( $k$ ) for different datasets based on our practical experience. In other words, given one test input, we have  $|\mathcal{O}| \times k$  data. After de-duplication, the resulting retrieved data sizes are shown in Table 2.

**Verbalizer Augmentation** To obtain possible verbalizers that can better represent classes, we first obtain top- $N$  predicted words given a test sample ( $N = 20$  for SST-2 and TREC,  $N = 10$  for AGNews and  $N = 5$  for YELP and DBPedia, considering their test set sizes). We set the number of candidate words  $|\mathcal{C}| = 20 \times |\mathcal{L}|$ , where  $|\mathcal{L}|$  is number of classes. We use a ROBERTA-large model fine-tuned on MNLI (Williams et al., 2018), as the entailment model for identifying potential verbalizer words for augmentation. Candidate with probability higher than a threshold  $t$  is then added to the augmented verbalizer. We set  $t = 0.4$  by experiments.

For comparison, we also use Word2Vec (Mikolov et al., 2013) to obtain word vectors and explore potential verbalizer words by their similarity with the seed verbalizer words.

## 4.3 Results

### 4.3.1 Main Results

**Zero-shot Performance** In zero-shot setting, we compare AdaPrompt with prompt-based methods using ROBERTA (Schick and Schütze, 2021a), GPT-2 (Gao et al., 2021) and GPT-3 (Zhao et al., 2021), respectively. The Channel refers to noisy channel model (Min et al., 2022) based on GPT-2. Table 3 presents the results under zero-shot setting. Following previous work (Schick and Schütze,

Models	SST-2	Yelp	AGNEWS	DBPedia	TREC	Avg.
GPT-2	63.00/ NA(NA)	--	59.80/ NA(NA)	32.30/ NA(NA)	38.70/ NA(NA)	--
Channel	77.10/ NA(NA)	--	61.80/ NA(NA)	51.40/ NA(NA)	30.50/ NA(NA)	--
GPT-3	75.80/ 0.00(75.80)	--	73.90/0.00(73.90)	59.70/0.00(59.70)	57.40/0.00(57.40)	--
R.	64.56/16.77(88.99)	72.63/ 6.34(87.97)	69.52/6.96(78.76)	56.32/0.49(56.67)	45.50/0.14(45.60)	61.71
Ada	75.92/17.36(91.28)	75.09/17.57(89.25)	76.55/7.28(84.95)	70.95/8.80(77.17)	60.50/3.54(63.00)	71.80
iAda	77.18/17.96(91.74)	75.81/18.05(90.41)	74.28/9.00(83.37)	73.01/6.70(77.92)	61.10/1.27(62.00)	72.28

Table 3: Zero-shot results. We report average accuracy and standard deviation of different patterns here. Results of the best patterns are shown in brackets. The Avg. reports the overall averaged results. R. stands for ROBERTA-large. Ada and iAda denote to AdaPrompt and iterative AdaPrompt based on ROBERTA-large, respectively. The results of GPT-2 large and Channel are from (Min et al., 2022), and Channel is based on GPT-2 large. GPT-3 results are reported by Zhao et al. (2021), using GPT-3 (175B). NA denotes to that results are not reported. For GPT-3 (Zhao et al., 2021), they only use a fixed prompt format.

T	Models	SST-2	Yelp	AGNEWS	DBPedia	TREC	Avg.
10	ROBERTA	84.97 ± 9.88	86.84 ± 16.08	78.42 ± 6.23	86.78 ± 1.10	45.56 ± 9.55	76.51
	AdaPrompt	90.42 ± 1.63	89.13 ± 13.30	84.21 ± 2.00	91.68 ± 1.84	57.56 ± 7.85	82.60
50	ROBERTA	92.56 ± 1.31	95.87 ± 0.57	85.50 ± 1.36	94.72 ± 0.49	73.88 ± 3.13	88.51
	AdaPrompt	92.75 ± 1.03	95.74 ± 0.89	86.29 ± 0.80	94.59 ± 0.71	78.42 ± 6.17	89.56
100	ROBERTA	92.40 ± 1.04	95.89 ± 0.68	87.29 ± 1.31	95.59 ± 0.52	86.30 ± 2.14	91.49
	AdaPrompt	92.75 ± 0.68	95.93 ± 0.95	87.98 ± 0.65	95.60 ± 0.51	87.58 ± 1.38	91.97

Table 4: Average accuracy and standard deviation on SST-2, YELP, AGNews, DBPedia and TREC under few-shot settings. |T| is the training set size. Each experiment is repeated 5 times using different seeds.

2021a,b), we report average accuracy, standard deviation and accuracy of the best pattern over different patterns.

First, compared with our foundation model, ROBERTA-large, we see that AdaPrompt consistently outperforms regular prompt-based methods on all datasets with better average performance and best pattern performance, bringing a 2.46 ~ 14.63 improvement. It is noticeable that AdaPrompt outperforms GPT-3 in zero-shot setting, which is a huge model with 175B parameters pretrained on a gigantic corpus. This confirms the effectiveness of AdaPrompt in domain adaptation. We observe that iterative AdaPrompt can further bring improvements on most datasets (SST-2, YELP and DBPedia). This directly demonstrates that PLMs continual pretrained on the retrieved data can be more adaptive to downstream tasks, and thus generate more task relevant label words, which can serve as a source to find better texts. Performance of iterative AdaPrompt (iAda) decreases on AGNEWS, we believe this is because this news dataset is similar with general data used for pretraining ROBERTA, and thus continual pretraining on such retrieved data can be less useful. Finally, we see that AdaPrompt improves over 10.09 accuracy of the overall performance.

**Few-shot Performance** Table 4 reports the experimental results in few shot setting. Each experiment

is repeated 5 times using different seeds and we report the average accuracy and standard deviation. To explore whether AdaPrompt can consistently bring improvement to ROBERTA, we conduct experiments using 10, 50, 100 samples, respectively.

Compared with ROBERTA-large baseline, under few-shot setting, AdaPrompt can still improve model performance. Although the relative improvement decreases as the size of training set improves, we can see that AdaPrompt outperforms ROBERTA over all tasks in all few-shot settings. In particular, AdaPrompt outperforms standard ROBERTA models by 2.29 ~ 5.79% in 10-shot setting, showing that it is useful in the very-few-shot setting.

### 4.3.2 Ablation Study

To study the effectiveness of continual pretraining on prompt-aware data and verbalizer augmentation, we conduct ablation experiments by removing continual pretraining (CP) or verbalizer augmentation (va). As shown in Table 5, We can see that compared with foundation model (-CP-va, 61.71 acc. on average), continual pretraining and verbalizer augmentation can both bring improvement to model performance (5.31 and 5.89 acc. on average, respectively), and the model has the best results when two methods are combined together (AdaPrompt), suggesting these two methods can benefit each other.

Models	SST-2	Yelp	AGNEWS	DBPedia	TREC	Avg.
AdaPrompt	75.92 ± 17.36	75.09 ± 17.57	76.55 ± 07.28	70.95 ± 08.80	60.50 ± 03.54	71.80
-va	71.07 ± 13.58	71.04 ± 15.57	72.16 ± 05.78	65.90 ± 02.71	45.40 ± 01.13	65.11
-CP	72.16 ± 16.35	75.72 ± 17.79	75.70 ± 07.88	50.95 ± 00.09	58.70 ± 03.25	66.65
-PR	71.22 ± 15.55	74.85 ± 17.51	75.12 ± 05.71	70.40 ± 07.48	58.60 ± 00.57	70.04
-CP-va	64.56 ± 16.77	72.63 ± 16.34	69.52 ± 06.96	56.32 ± 00.49	45.50 ± 00.14	61.71

Table 5: Experimental results of ablation study. “-” means “without” here. va: verbalizer augmentation, CP: Continual Pretraining, PR: Prompt-aware Retrieval. Note that -PR means we do not use prompt-aware retrieval, but simply use raw test input data for retrieval and continual pretraining, referred as *in-domain adaptation*.

Model	SST-2	DBPedia
RoBERTA	64.82 ± 11.62	56.49 ± 00.41
AdaPrompt	73.05 ± 13.08	70.97 ± 08.87

Table 6: Model performance tested on *unseen* test set. We report averaged accuracy and standard deviation.

SST-2				
$E_{space}$	1	10	50	100
Size	3k	23k	98k	205k
Accuracy	73.54 ± 16.77	75.06 ± 17.34	75.95 ± 17.73	75.92 ± 17.36
DBPedia				
$E_{space}$	1	5	25	50
Size	58k	235k	708k	1,301k
Accuracy	70.64 ± 9.66	71.39 ± 10.78	74.13 ± 7.51	70.95 ± 8.80

Table 7: Analysis on retrieved data size. Data sizes are calculated after de-duplication.

In addition, we investigate the influence on model performance by removing prompt-aware retrieval and only retrieving with raw texts. From the table we can see that on all datasets, using prompt-augmented queries (AdaPrompt) give substantially stronger results. Take SST-2 for example, the accuracy is 71.22 (SST-2 -PR) given only raw input queries, but 75.92 with prompt-augmented queries, with a 4.7 absolute improvement. This shows that continual pretraining using prompt-aware data is highly beneficial to zero-shot prompt-based NLP.

#### 4.4 Analysis

**Generalization Capability** For experiments in section 4.3.1, we use task test set as the sources to build queries for retrieving pretraining data. However, in a more general setting, we want to learn when the query data and test set are different, whether AdaPrompt can still generalize to this test set. To this end, we build an *unseen* test set by using the original training set of SST-2 and DBPedia. We then evaluate models (trained using queries from the origin test set) on this unseen test set. As shown in Table 6, AdaPrompt achieves 73.05 and 70.97 accuracy on SST-2 and DBPedia,

respectively. Compared with performance on original test set (Table 3), although the performance of AdaPrompt slightly decreases when evaluated on SST-2 unseen test set, it can still outperform RoBERTA by a large margin (+8.23). It demonstrates that AdaPrompt has a strong generalization ability when query data and test set are different.

**Size of Retrieved Data** As stated, Elasticsearch returns top- $k$  texts in the order of matching scores. Using a smaller  $k$ , the retrieved data are more textual related to the query, while using a larger  $k$ , the retrieved data can contain certain noise. To compare the effects of different sizes of retrieved data for continual pretraining, We set  $k$  to 1, 10, 50 100 for the SST-2 and set  $k$  to 1, 5, 25, 50 for DBPedia, respectively. As shown in Table 7, we see that accuracy rises in the beginning when retrieval size increases. But as the retrieval size grows bigger, the accuracy starts to decrease slightly. This can be explained by that the lower-ranked retrieved data have a lower relevance to the target task, which introduces more noise in continual pretraining. We use fixed  $k$  for our experiments in zero-shot settings (Section 4.2), due to lack of a validation set. In few-shot settings, in practice,  $k$  can be considered as a hyperparameter and tuned over validation data.

**The Effect of Verbalizer Strategies** Table 8 compares the model performance when using different verbalizer augmentation strategies, namely using NLI model and word similarity (Section 4.2). Additional, we compare AdaPrompt with a verbalizer augmentation method using knowledge base (KB) (Hu et al., 2022)<sup>2</sup>. To set a fair comparison, we limit the verbalizer word set for each label within 5. We report average accuracy and standard deviation here.

<sup>2</sup>For sentiment analysis tasks, we take sentiment words shown in (Hu et al., 2022), which are adopted from <https://www.enchantedlearning.com/wordlist/>; for other tasks, we use most related words: <https://relatedwords.org/>.



Dataset	SST-2	YELP	AGNEWS	DBPedia	TREC	Avg.
$va_w$	74.91 $\pm$ 11.71	75.39 $\pm$ 17.47	69.07 $\pm$ 06.70	55.32 $\pm$ 11.33	60.60 $\pm$ 03.39	67.06
$va_m$	75.92 $\pm$ 17.36	75.09 $\pm$ 17.57	76.55 $\pm$ 07.28	70.95 $\pm$ 08.80	60.50 $\pm$ 03.54	71.80
$va_k$	69.07 $\pm$ 15.80	74.64 $\pm$ 17.55	60.15 $\pm$ 07.79	74.85 $\pm$ 17.50	24.00 $\pm$ 00.57	60.54

Table 8: Model performance of AdaPrompt using different verbalizer augmentation strategies.  $va_w$ : using word2vec similarity.  $va_m$ : using ROBERTA trained on MNLI.  $va_k$ : using most related words/sentiment dictionary. Avg. refers to overall averaged results.

Model	Size	SST-2
Albert	17M	54.67 $\pm$ 3.30(58.94)
Albert+AdaPrompt	17M	58.51 $\pm$ 5.79(63.99)
Bert	340M	58.03 $\pm$ 6.18(63.53)
Bert+AdaPrompt	340M	68.89 $\pm$ 16.11(85.67)
ROBERTA	355M	64.56 $\pm$ 16.77(88.99)
ROBERTA+AdaPrompt	355M	77.18 $\pm$ 17.96(91.74)

Table 9: We report average accuracy and standard deviation here. Results of best patterns are shown in the bracket.

Results show that, compared with using word similarity to select candidate words and directly using KBs to augment verbalizer words, using NLI to augment verbalizer words gives better performance on most tasks, and is also more stable. We also find that using KBs to augment verbalizer words gives better performance on the DBPedia tasks, but much worse performance on the TREC task. This can be because TREC is less close to topic classification (Min et al., 2022), and directly using the most related words can be noisy. This also suggests that more sophisticated strategy that cares of tasks and prompt information can be useful, which we leave for future work.

**AdaPrompt with different PLMs** We apply AdaPrompt with different PLMs (Bert-large, Albert-large and ROBERTA-large). We report experimental results on the SST-2 dataset in Table 9. Although the performance of different models varies, we observe that AdaPrompt can consistently bring huge improvement over all models. We also find that model performance increases with model size. AdaPrompt using ROBERTA-large outperforms other models overall performance by a large margin (8.29  $\sim$  18.67) and achieves 91.74 accuracy with the best pattern.

## 5 Conclusion

We investigated AdaPrompt, a zero-shot prompt-based method for NLP that makes use of test input data and prompts for adaptive continual pretraining and verbalizer selection. Results on five classification datasets show that AdaPrompt improves

over a standard prompt method by large margins. In particular, retrieving relevant data for continual pretraining of a language model can serve to warm-up the model for both domain adaptation and prompt-filling tasks. In addition, an NLI model allows effective selection of filled tokens to achieve improved performance.

## Limitation

We acknowledge two major limitations of this work:

1. We only tested AdaPrompt on text classification tasks. The intention is to use this clear setting to compare with other prompt-based models. However, it is possible to extend AdaPrompt to other natural language understanding tasks or languages, which we leave for future exploration.
2. We only tested with ElasticSearch as the search method. However, there are signals showing the quality of retrieved text is constrained to the search engines. A better configuration or model of the search method might further improve AdaPrompt.

## Acknowledgements

Yue Zhang is the corresponding author. We appreciate all reviewers for their comments. This paper receives a support from the Pioneer and ‘‘Leading Goose’’ RD Program of Zhejiang under Grant Number 2022SDXHDX0003.

## References

- Roei Aharoni and Yoav Goldberg. 2020. [Unsupervised domain clusters in pretrained language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7747–7763, Online. Association for Computational Linguistics.
- Xuefeng Bai, Yulong Chen, and Yue Zhang. 2022a. [Graph pre-training for AMR parsing and generation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6001–6015, Dublin, Ireland. Association for Computational Linguistics.
- Xuefeng Bai, Pengbo Liu, and Yue Zhang. 2021. [Investigating typed syntactic dependencies for targeted sentiment classification using graph attention neural network](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:503–514.
- Xuefeng Bai, Linfeng Song, and Yue Zhang. 2022b. [Semantic-based pre-training for dialogue understanding](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 592–607, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [SciBERT: A pretrained language model for scientific text](#). 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 3615–3620, Hong Kong, China. Association for Computational Linguistics.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Alexandra Chronopoulou, Christos Baziotis, and Alexandros Potamianos. 2019. [An embarrassingly simple approach for transfer learning from pretrained language models](#). 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 2089–2095, Minneapolis, Minnesota. Association for Computational Linguistics.
- Leyang Cui, Yu Wu, Jian Liu, Sen Yang, and Yue Zhang. 2021. [Template-based named entity recognition using BART](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1835–1845, Online. Association for Computational Linguistics.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. [Making pre-trained language models better few-shot learners](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3816–3830, Online. Association for Computational Linguistics.
- Aaron Gokaslan and Vanya Cohen. 2019. [Openweb-text corpus](#).
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don’t stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.
- Shengding Hu, Ning Ding, Huadong Wang, Zhiyuan Liu, Jingang Wang, Juanzi Li, Wei Wu, and Maosong Sun. 2022. [Knowledgeable prompt-tuning: Incorporating knowledge into prompt verbalizer for text classification](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2225–2240, Dublin, Ireland. Association for Computational Linguistics.
- Zhengbao Jiang, Frank F. Xu, Jun Araki, and Graham Neubig. 2020. [How can we know what language models know?](#) *Transactions of the Association for Computational Linguistics*, 8:423–438.
- Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick Van Kleef, Sören Auer, et al. 2015. Dbpedia—a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic web*, 6(2):167–195.
- Xiang Lisa Li and Percy Liang. 2021. [Prefix-tuning: Optimizing continuous prompts for generation](#). *arXiv preprint arXiv:2101.00190*.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021. [Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing](#). *arXiv preprint arXiv:2107.13586*.
- 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](#). *ArXiv preprint*, abs/1907.11692.

- Tomas Mikolov, Kai Chen, Greg S Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space.
- Sewon Min, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. [Noisy channel language model prompting for few-shot text classification](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5316–5330, Dublin, Ireland. Association for Computational Linguistics.
- Robert C. Moore and William Lewis. 2010. [Intelligent selection of language model training data](#). In *Proceedings of the ACL 2010 Conference Short Papers*, pages 220–224, Uppsala, Sweden. Association for Computational Linguistics.
- Sebastian Nagel. 2016. [Cc-news](#).
- Ethan Perez, Douwe Kiela, and Kyunghyun Cho. 2021. [True few-shot learning with language models](#). *ArXiv preprint*, abs/2105.11447.
- 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.
- Barbara Plank and Gertjan van Noord. 2011. [Effective measures of domain similarity for parsing](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1566–1576, Portland, Oregon, USA. Association for Computational Linguistics.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21:1–67.
- Sebastian Ruder, Parsa Ghaffari, and John G Breslin. 2017. [Data selection strategies for multi-domain sentiment analysis](#). *ArXiv preprint*, abs/1702.02426.
- Sebastian Ruder and Barbara Plank. 2017. [Learning to select data for transfer learning with Bayesian optimization](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 372–382, Copenhagen, Denmark. Association for Computational Linguistics.
- Timo Schick and Hinrich Schütze. 2021a. [Exploiting cloze-questions for few-shot text classification and natural language inference](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 255–269, Online. Association for Computational Linguistics.
- Timo Schick and Hinrich Schütze. 2021b. It’s not just size that matters: Small language models are also few-shot learners. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2339–2352.
- Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. 2020. Eliciting knowledge from language models using automatically generated prompts. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4222–4235.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Trieu H Trinh and Quoc V Le. 2018. [A simple method for commonsense reasoning](#). *ArXiv preprint*, abs/1806.02847.
- Marlies van der Wees, Arianna Bisazza, and Christof Monz. 2017. [Dynamic data selection for neural machine translation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1400–1410, Copenhagen, Denmark. Association for Computational Linguistics.
- Ellen M Voorhees and Dawn M Tice. 2000. Building a question answering test collection. In *Proceedings of the 23rd annual international ACM SIGIR conference on Research and development in information retrieval*, pages 200–207.
- Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. 2019. [Universal adversarial triggers for attacking and analyzing NLP](#). 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 2153–2162, Hong Kong, China. Association for Computational Linguistics.
- Wei Wang, Taro Watanabe, Macduff Hughes, Tetsuji Nakagawa, and Ciprian Chelba. 2018. [Denosing neural machine translation training with trusted data and online data selection](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 133–143, Brussels, Belgium. Association for Computational Linguistics.

- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Xingcheng Yao, Yanan Zheng, Xiaocong Yang, and Zhilin Yang. 2021. [Nlp from scratch without large-scale pretraining: A simple and efficient framework](#). *ArXiv preprint*, abs/2111.04130.
- Wenpeng Yin, Jamaal Hay, and Dan Roth. 2019. [Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach](#). 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 3914–3923, Hong Kong, China. Association for Computational Linguistics.
- Wenhao Yu, Chenguang Zhu, Yuwei Fang, Donghan Yu, Shuohang Wang, Yichong Xu, Michael Zeng, and Meng Jiang. 2022. [Dict-BERT: Enhancing language model pre-training with dictionary](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1907–1918, Dublin, Ireland. Association for Computational Linguistics.
- Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. [Bartscore: Evaluating generated text as text generation](#). *ArXiv preprint*, abs/2106.11520.
- Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 649–657.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. [Calibrate before use: Improving few-shot performance of language models](#). In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 12697–12706. PMLR.
- Ming Zhong, Yang Liu, Yichong Xu, Chenguang Zhu, and Michael Zeng. 2021. [Dialoglm: Pre-trained model for long dialogue understanding and summarization](#). *ArXiv preprint*, abs/2109.02492.
- Yukun Zhu, Ryan Kiros, Richard S. Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. [Aligning books and movies: Towards story-like visual explanations by watching movies and reading books](#). In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 19–27. IEEE Computer Society.