

# Prompt Combines Paraphrase: Teaching Pre-trained Models to Understand Rare Biomedical Words

Haochun Wang<sup>1</sup>, Chi Liu<sup>1</sup>, Nuwa Xi<sup>1</sup>, Sendong Zhao<sup>1\*</sup>, Meizhi Ju<sup>2</sup>, Shiwei Zhang<sup>2</sup>  
Ziheng Zhang<sup>2</sup>, Yefeng Zheng<sup>2</sup>, Bing Qin<sup>1</sup> and Ting Liu<sup>1</sup>

<sup>1</sup>Research Center for Social Computing and Information Retrieval,  
Harbin Institute of Technology, China

<sup>2</sup>Tencent Jarvis Lab, Shenzhen, China

{hcwang, cliu, nwxi, sdzhao, bqin, tliu}@ir.hit.edu.cn

meizhi.ju@outlook.com, {zswvizhang, zihengzhang, yefengzheng}@tencent.com

## Abstract

Prompt-based fine-tuning for pre-trained models has proven effective for many natural language processing tasks under few-shot settings in general domain. However, tuning with prompt in biomedical domain has not been investigated thoroughly. Biomedical words are often rare in general domain, but quite ubiquitous in biomedical contexts, which dramatically deteriorates the performance of pre-trained models on downstream biomedical applications even after fine-tuning, especially in low-resource scenarios. We propose a simple yet effective approach to helping models learn rare biomedical words during tuning with prompt. Experimental results show that our method can achieve up to 6% improvement in biomedical natural language inference task without any extra parameters or training steps using few-shot vanilla prompt settings.

## 1 Introduction

Pre-trained models have achieved a great success in natural language processing (NLP) and become a new paradigm for various tasks (Peters et al., 2018; Devlin et al., 2019; Liu et al., 2019; Qiu et al., 2020). Many studies have paid attention to pre-trained models in biomedical NLP tasks (Lee et al., 2020; Lewis et al., 2020; Zhao et al., 2021). However, plain pre-trained models sometimes cannot do very well in biomedical NLP tasks. In general, there are two challenges to fully exploit the potential of the pre-trained models for biomedical NLP tasks, i.e., (1) **limited data** and (2) **rare biomedical words**. Firstly, it is common that the amount of biomedical labeled data is limited due to strict privacy policy constraints (Šuster et al., 2017), high cost and professional requirement for data annotation. Pre-trained models perform poorly with few samples since abundant training samples are essential to optimize task-related parameters (Liu et al.,

2021a). Secondly, biomedical words are usually low-frequency words but critical to understanding biomedical texts. As an example of natural language inference (NLI) task in Figure 1, the model goes wrong during tuning when faced with a rare word “*afebrile*”<sup>1</sup> in the premise, whose meaning is “*having no fever*”. It can be no easy for the pre-trained models to predict the correct label if the models haven’t seen the rare biomedical words for enough times during pre-training or tuning stage. Thus, pre-trained models cannot capture the precise semantics of biomedical texts in the scenario of low-resource tasks.

With very few annotated samples available for a new task, it is hard to effectively fine-tune pre-trained models with the additional task-specific parameters, which is even more of a challenge to biomedical domain as mentioned above. Prompt technique has been introduced to smooth the fine-tuning process in the few-shot settings by narrowing down the gap between pre-training stage and the downstream task in general domain (Liu et al., 2021a), as demonstrated in Figure 2. Therefore, it is beneficial to adapt prompt-based tuning to biomedical NLP tasks.

Although the challenge of rare words is a critical problem for the biomedical pre-trained models, only a handful of works have studied the issue and most of them focus on enriching the representation of rare words through pre-training stage (Schick and Schütze, 2020; Yu et al., 2021; Wu et al., 2020). Thus, it naturally requires them to involve a second-round pre-training or further training steps with biomedical knowledge to achieve the above goal, which is highly time-consuming and inefficient. Alternatively, we emphasize on tuning stage instead of pre-training to resolve these issues. When coming across an unknown word, human may seek the

<sup>1</sup>There are around 4 billion words in the selected biomedical pre-training texts while “*afebrile*” appears only about 100,000 times, accounting for 0.0025%.

\*Corresponding author

<b>Task:</b>	Medical Natural Language Inference	
<b>Premise:</b>	Lactate only 1.3 and pt <span style="border: 1px solid black; border-radius: 5px; padding: 2px;">afebrile</span> .	<b>Hypothesis:</b> Temperature was within normal range.
<b>Gold label:</b>	Entailment <span style="color: green;">✔</span>	<b>Model Prediction:</b> Neutral <span style="color: red;">✘</span>
<b>Rare word:</b>	<span style="border: 1px solid black; border-radius: 5px; padding: 2px;">afebrile</span>	<b>Paraphrase:</b> <span style="border: 1px solid black; border-radius: 5px; padding: 2px;">afebrile - having no fever</span>
<b>Premise:</b>	<span style="border: 1px solid black; border-radius: 5px; padding: 2px;">HCV</span> diagnosed when he was in his 20's with a history of recurrent <span style="border: 1px solid black; border-radius: 5px; padding: 2px;">HCV</span> 2.	
<b>Hypothesis:</b>	He has no history of liver disease.	
<b>Gold label:</b>	Contradiction <span style="color: green;">✔</span>	<b>Model Prediction:</b> Entailment <span style="color: red;">✘</span>
<b>Rare word:</b>	<span style="border: 1px solid black; border-radius: 5px; padding: 2px;">HCV</span>	<b>Paraphrase:</b> <span style="border: 1px solid black; border-radius: 5px; padding: 2px;">HCV - hepatitis c virus</span>

Figure 1: Failures of a biomedical pre-trained model (Lewis et al., 2020) on the biomedical NLI task.

dictionary for its paraphrase. Enlightened by this phenomenon, we propose to explain rare biomedical words with the paraphrases on the basis of prompt-based tuning. The new approach could enhance tuning capability in understanding biomedical words. Furthermore, as a generic plug-in module for non-specific datasets, our approach is model-agnostic and can be easily transferred to other domains.<sup>2</sup>

In summary, our contributions are as follows:

- We investigate a valuable problem of the adaptation of pre-trained models to few-shot scenarios to enhance biomedical text understanding with a focus on rare biomedical words.
- We propose a novel method to combine the prompt paradigm and paraphrases of rare biomedical words in the tuning stage of pre-trained models to address the limitation caused by “rare but key words” in biomedical texts.
- We evaluate on six pre-trained models over two biomedical natural language understanding datasets—MedNLI and MedSTS. Our approach can improve the performance by up to 6% in the few-shot settings. Moreover, we discuss how the paraphrases improve the pre-trained models and provide a perspective about task-related rare words.

## 2 Related Work

**Representation learning of rare words in pre-trained models.** Words in the vocabulary list fol-

<sup>2</sup>We release our code at [https://github.com/s65b40/prompt\\_n\\_paraphrase](https://github.com/s65b40/prompt_n_paraphrase)

low a Zipf distribution (Zipf, 2016) by and large. Previous works have shown that the word representation space of the pre-trained models is anisotropic and high-frequency words dominate the representation of a sentence which can induce semantic bias (Gao et al., 2019; Li et al., 2020; Yan et al., 2021). Meanwhile, it has also been proven that rare words limit the performance of pre-trained models as the rare words can play a decisive role in the sentence understanding (Schick and Schütze, 2020; Wu et al., 2020; Yu et al., 2021). Schick and Schütze (2020) introduced one-token approximation to infer the embedding of arbitrary rare word by a single token. Wu et al. (2020) proposed to take notes on the fly to maintain a note dictionary for rare words to save the contextual information that helps enhance the representation during pre-training.

**Biomedical pre-trained models.** With the booming trend of pre-trained models in NLP tasks (Peters et al., 2018; Devlin et al., 2019; Liu et al., 2019), various trials have been made to investigate the pre-trained models in biomedical domain (Peng et al., 2019; Lee et al., 2020; Huang et al., 2019). Lewis et al. (2020) and Gu et al. (2021a) further built the domain-specific vocabulary to amend the representation of biomedical words. More recent works guided biomedical pre-trained models with domain knowledge. For example, Zhang et al. (2021) amplified the biomedical entities with type information from neighbor entities. Michalopoulos et al. (2021) learned clinical word embeddings with the association of synonyms in the Unified Medical Language System (UMLS) Metathesaurus.

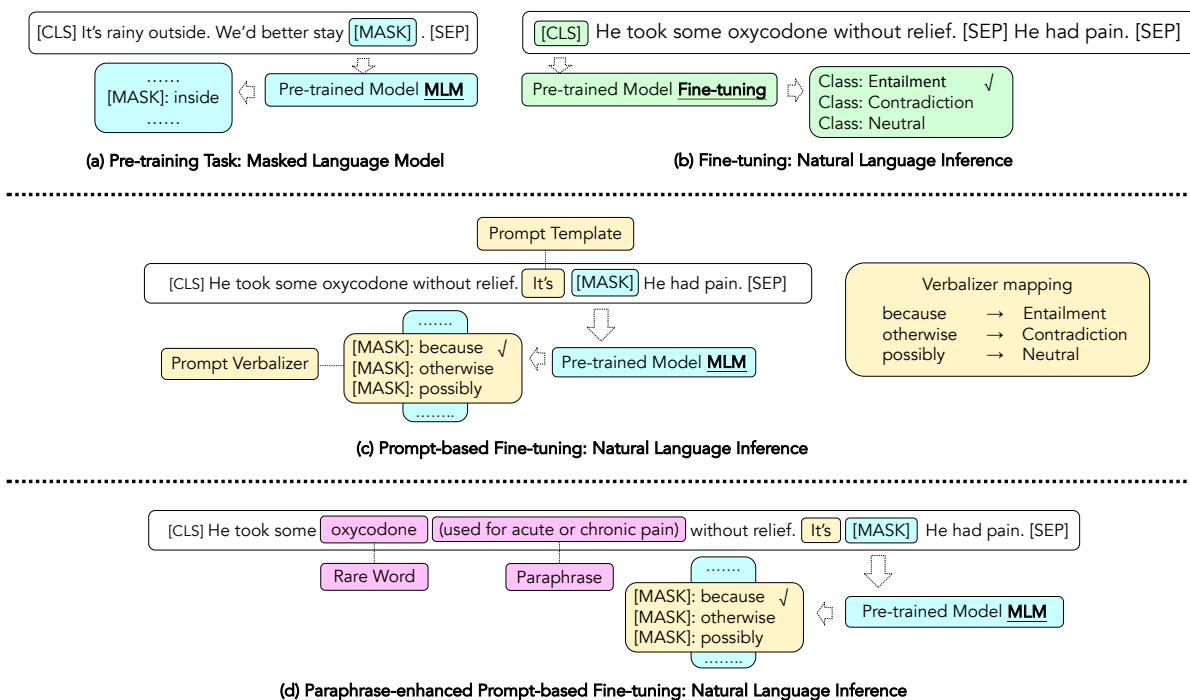


Figure 2: Examples for paradigms of (a) Masked Language Model (MLM) pre-training; (b) Task-specific fine-tuning; (c) Prompt-based fine-tuning, with same task as pre-training process; (d) Paraphrase-enhanced prompt-based fine-tuning. Best viewed in color.

**Tuning pre-trained models with prompt.** Many works were dedicated to applying prompt in fine-tuning by adapting the downstream tasks to the paradigm of pre-training tasks. Prompts that have been employed by now can be categorized into two groups: (1) discrete prompt in natural language (Schick and Schütze, 2021; Gao et al., 2021) and (2) continuous prompt in representation based on trainable vectors (Li and Liang, 2021; Shin et al., 2020b). Discrete prompt follows the settings of pre-training tasks and converts the downstream tasks into a cloze question format without requiring additional parameters. Continuous prompt inserts prompt embeddings into the models, which could perform better than discrete prompt but at the expense of explainability and extra training cost on additional data for the prompt embeddings (Wei et al., 2021; Gu et al., 2021b). Recently, Liu et al. (2021b) considered training continuous prompts as a parameter-efficient method instead of tuning the parameters of the entire pre-trained model. In this study, we follow the paradigm of discrete prompt to avoid introducing more ambiguity from the prompt embeddings or training costs on additional training data.

### 3 Method

In this section, we introduce how we find the rare biomedical words and append paraphrases to the rare biomedical words with the prompt-based tuning of pre-trained models in a model-agnostic plug-in manner.

#### 3.1 Rare Words

The rarity of a word mostly depends on its frequency in a certain corpus, which can vary from context to context. A rare word in the pre-training corpora is possibly not *that* rare in the down-stream tasks. In this work, we define the “rare words” as the words whose frequency is under a specific threshold in the pre-training corpora as aforementioned.

Meanwhile, although the pre-trained models tokenize the input words into tokens, tokenizers based on byte-pair encoding (Sennrich et al., 2016) or WordPiece (Schuster and Nakajima, 2012) split words into sub-words by frequency or likelihood, which is both dominated by the common words. Thus, although the rare words can be split into possible non-rare tokens, there is not much semantics from the original rare words retained after being tokenized into common tokens for the pre-trained

models. Also, tokenizers of different pre-trained models can tokenize the same rare word into different tokens and consequentially make rare tokens model-related. For example, BERT-Large (Devlin et al., 2019) model tokenizes “afebrile” into “af-eb-ril-e” while Biomedical-Clinical-RoBERTa-Large (Lewis et al., 2020) model tokenizes it into “a-fe-brile”.

### 3.2 Selection of Rare Biomedical Words

To obtain the frequency of words, we adopt the biomedical corpora including PubMed abstract,<sup>3</sup> PubMed Central<sup>4</sup> (PMC) full-text and MIMIC-III dataset,<sup>5</sup> which are widely used for pre-training biomedical language models, such as BC-RoBERTa (Lewis et al., 2020), BioBERT (Lee et al., 2020), and PubMedBERT (Gu et al., 2020). We loop the above corpora to obtain the frequency of each word in the pre-training phase. The rare words found in biomedical corpora are likely to contain words not only in the biomedical domain but also in the general domain. Instead of including all rare words, we consider rare words from biomedical domain with the following two reasons: (1) **Domain-specific distribution**: unlike the general domain, distribution of words in the biomedical domain is shaped with domain-specific terms, such as disease, medicine, diagnosis and treatment (Lee et al., 2020). (2) **Task-specific words**: rare words from the biomedical domain can contribute more to biomedical tasks than that from the general domain. Therefore, we introduce a threshold, an empirical hyper-parameter, to assist the selection of rare words following Yu et al. (2021), and we also experiment with different thresholds for the rare words in Section 4.3. We retrieve the paraphrases of the rare biomedical words from an online dictionary “Wiktionary”.<sup>6</sup> To optimize the selection, we only keep the rare words that are tagged with medical-related categories from the Wiktionary, i.e. *medical*, *medicine*, *disease*, *symptom* and *pharmacology*.

### 3.3 Selection of Paraphrases

There can be more than just one paraphrase for a rare biomedical word and it is tricky to choose the most appropriate paraphrases. Therefore, to

avoid introducing noise from the inappropriate paraphrases, we exclude rare biomedical words with more than one corresponding paraphrase. In addition, we ignore paraphrases that contain additional rare words whose frequencies are below the set threshold since it only replaces one rare word with another. Meanwhile, considering that biomedical abbreviations are likely to be tokenized into separate letters with no meaningful semantic information, we retrieve and append the paraphrases to all the biomedical abbreviations.

### 3.4 Prompt-based Fine-Tuning with Paraphrases

When coming across new words during reading, humans habitually seek dictionaries for the corresponding paraphrases to help us understand. Following the same idea, we guide the pre-trained models with paraphrases, where rare words are followed by the parenthesis punctuation, as shown in Figure 2(d). In this way, given a pre-trained model, paraphrases of biomedical rare words can be considered as a portable plug-in module and generated for any dataset instantly before prompt-based fine-tuning.

## 4 Experiments

### 4.1 Setup

**Models.** To demonstrate that our approach is model-agnostic, we adopt six pre-trained models in both general and biomedical domains, namely (1) BERT-Large (Devlin et al., 2019), (2) RoBERTa-Large (Liu et al., 2019), (3) BioBERT-Base (Lee et al., 2020), (4) PubMedBERT-Base (Gu et al., 2020), (5) SciBERT-Base (Beltagy et al., 2019) and (6) BC-RoBERTa-Large (Biomedical-Clinical RoBERTa-Large) (Lewis et al., 2020). During prompt-based tuning, we use the same set of hyper-parameters for all the six pre-trained models, including learning rate of  $1 \times 10^{-5}$ , batch size of 2 and max epoch of 10.

**Datasets.** Previous work demonstrates that rare words have more impact on Natural Language Understanding (NLU) tasks than Information Extraction (IE) (Schick and Schütze, 2020), while most biomedical NLP tasks fall into the category of IE (Shin et al., 2020a; Gu et al., 2021a). To better demonstrate the method effectiveness, we perform evaluation over the two biomedical NLU datasets, namely MedNLI (Romanov and Shivade, 2018)

<sup>3</sup><https://pubmed.ncbi.nlm.nih.gov>

<sup>4</sup><https://www.ncbi.nlm.nih.gov/pmc>

<sup>5</sup><https://physionet.org/content/mimiciii/1.4/>

<sup>6</sup>Wiktionary - <https://en.wiktionary.org/>

and MedSTS<sup>7</sup> (Wang et al., 2020).

MedNLI is a natural language inference dataset where premises are selected from real clinical notes in MIMIC-III (Johnson et al., 2016). And MedSTS is a semantic textual similarity dataset gathered from a clinical corpus at Mayo Clinic and the ground-truth label of the similarity is the mean of the subjectively annotated scores from multiple annotators. As MedSTS is actually a regression task, we adapt the task following Gao et al. (2021) and convert it into a classification task. We use the same data splitting of training, development and test sets as the original two datasets. Statistics of datasets can be found in Table 1.

Dataset	Train	Dev	Test
MedNLI	11,232	1,395	1,422
MedSTS	750	/	318

Table 1: Statistics of the MedNLI and MedSTS datasets. We use the ClinicalSTS-2018 subset of MedSTS.

**Few-shot datasets.** Initialized with 10 different random seeds, we randomly sample instances within the range of 16 to 256 from corresponding training and development sets as the few-shot training and development sets. The original test set is directly used as the few-shot test set. Note that there is no development set in MedSTS, so we sample the few-shot development set from the original training set with the same quantity of samples as the few-shot training set with no overlapping instances. Accuracy and Pearson correlation coefficients are used as the evaluation metrics for MedNLI and MedSTS, respectively.

**Prompt settings.** We prepare discrete prompts using the same prompt settings from Schick and Schütze (2021) and Gao et al. (2021), which correspond to the NLI and STS tasks respectively in Table 2.

Task	Template	Verbalizer
MedNLI	<Sent1>. [MASK]. <Sent2>	Yes/No/maybe
MedSTS	<Sent1>. [MASK]. <Sent2>	Yes/No

Table 2: Prompt settings for MedNLI and MedSTS.

**Rare biomedical words and paraphrases.** We find that a threshold of rare biomedical words at 200,000, corresponding to a frequency less than

<sup>7</sup>We use ClinicalSTS-2018 which is a sub-dataset of MedSTS provided by the maintainers of the MedSTS project.

0.005% in the pre-training corpora, can yield better results in most scenarios (details in Section 4.3), so we consider the biomedical words that appear less than 200,000 times as “rare biomedical words” and prepare the same rare word and paraphrase sets for all the models to validate the generalization ability of our approach.

## 4.2 Few-shot learning results

We report the mean accuracy for the MedNLI task and Pearson correlation coefficient for the MedSTS task over 10 sampled few-shot datasets based on different random seeds, along with standard deviation and  $p$ -value of paired  $t$ -test. Table 3 and Table 4 show the results for the two tasks on six pre-trained models.

**Results on MedNLI.** The pre-trained models with the paraphrases for rare biomedical words can outperform the baselines in all cases and can bring about 6% improvement on average for few-shot learning with 16 training samples and 2% with 256 training samples. All performance improvements are statistically significant with  $p$ -value less than 0.05 except only two out of ten cases from the RoBERTa-based model in general domain. Besides, small pre-trained biomedical models are comparable with large pre-trained models in the general domain under the few-shot settings. Furthermore, with more training samples up to 256, our approach is consistently effective.

**Results on MedSTS.** Similar to MedNLI, the incorporation of paraphrases improves the performances compared with baselines in general. For some cases, statistical significance is not as stable as that on MedNLI for two reasons: (1) Some ground-truth labels in the MedSTS task can be biased due to the subjectivity of annotation (Yang et al., 2020); (2) Rare biomedical words shared in the sentence pair of the same sample can be a shortcut for the pre-trained models, as the paraphrases increase the overlap between the two sentences and mislead the models to overlook the rest of the sentences (McCoy et al., 2019).

## 4.3 Thresholds for rare biomedical words

The number of paraphrases of rare biomedical words involved in the samples can directly affect the model performance, which is controlled by the pre-set threshold. To measure the influence of different thresholds, we conduct experiments over

		MedNLI				
#Samples	Model	16	32	64	128	256
	BERT-Large	38.9 (3.7)	44.5 (5.2)	50.1 (5.2)	54.8 (2.5)	59.9 (1.2)
	+ paraphrase	<b>40.8</b> (4.1)	<b>46.0</b> (5.5)	<b>53.3</b> (4.9)	<b>58.1</b> (1.4)	<b>61.9</b> (1.4)
	<i>p</i> -value	< 0.01	0.02	< 0.01	< 0.01	< 0.01
	RoBERTa-Large	43.2 (6.7)	52.1 (8.2)	63.6 (4.6)	69.2 (1.8)	72.7 (1.4)
	+ paraphrase	<b>49.5</b> (8.1)	<b>56.1</b> (7.6)	<b>65.6</b> (2.9)	<b>70.8</b> (0.7)	<b>74.0</b> (1.3)
	<i>p</i> -value	< 0.01	0.03	0.08	0.06	0.02
	BioBERT-Base	34.1 (1.5)	38.5 (3.3)	42.5 (4.6)	52.1 (2.5)	59.4 (1.6)
	+ paraphrase	<b>36.3</b> (1.8)	<b>40.9</b> (3.5)	<b>45.2</b> (4.3)	<b>54.0</b> (2.3)	<b>60.4</b> (4.7)
	<i>p</i> -value	< 0.01	< 0.01	< 0.01	0.02	0.03
	PubMedBERT-Base	40.5 (3.4)	46.8 (4.9)	53.9 (4.0)	62.9 (1.6)	69.2 (1.1)
	+ paraphrase	<b>45.0</b> (4.0)	<b>49.7</b> (5.0)	<b>56.4</b> (3.4)	<b>65.8</b> (1.6)	<b>71.0</b> (1.6)
	<i>p</i> -value	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01
	SciBERT-Base	36.8 (1.2)	41.4 (3.3)	49.0 (3.7)	54.7 (2.0)	60.9 (1.0)
	+ paraphrase	<b>38.2</b> (2.1)	<b>45.4</b> (5.2)	<b>50.1</b> (3.5)	<b>56.4</b> (2.2)	<b>61.9</b> (1.6)
	<i>p</i> -value	< 0.01	0.03	0.01	0.02	< 0.01
	BC-RoBERTa-Large	51.3 (5.9)	60.6 (6.7)	71.0 (3.7)	80.6 (1.3)	83.1 (1.3)
	+ paraphrase	<b>56.6</b> (5.0)	<b>62.3</b> (6.0)	<b>74.5</b> (3.0)	<b>81.1</b> (1.5)	<b>83.6</b> (1.0)
	<i>p</i> -value	< 0.01	0.05	< 0.01	0.02	0.01

Table 3: Few-shot results on the MedNLI task using various pre-trained models with training and development sets of different sizes. We report mean (standard deviation) performance of accuracy over 10 different random seeds, along with the *p*-value of the paired *t*-test. + paraphrase: with paraphrases of selected rare biomedical words.

		MedSTS				
#Samples	Model	16	32	64	128	256
	BERT-Large	14.1 (7.4)	24.8 (10.1)	43.6 (5.7)	60.2 (4.7)	72.2 (4.5)
	+ paraphrase	<b>18.5</b> (9.0)	<b>28.2</b> (12.1)	<b>48.7</b> (7.1)	<b>64.1</b> (5.1)	<b>72.7</b> (4.6)
	<i>p</i> -value	< 0.01	0.04	< 0.01	0.01	0.08
	RoBERTa-Large	29.5 (9.3)	41.7 (18.5)	55.1 (12.1)	67.6 (6.8)	76.0 (3.5)
	+ paraphrase	<b>34.6</b> (13.5)	<b>46.1</b> (13.0)	<b>57.7</b> (12.7)	<b>69.9</b> (6.7)	<b>77.2</b> (3.7)
	<i>p</i> -value	0.04	0.01	0.02	0.03	0.02
	BioBERT-Base	17.3 (14.4)	26.3 (13.7)	41.4 (9.0)	52.2 (10.6)	63.0 (7.3)
	+ paraphrase	<b>20.0</b> (12.9)	<b>28.2</b> (12.8)	<b>43.1</b> (8.6)	<b>53.7</b> (9.2)	<b>64.2</b> (7.1)
	<i>p</i> -value	0.03	0.02	0.01	0.02	0.02
	PubMedBERT-Base	10.3 (9.8)	22.8 (10.7)	36.9 (10.6)	48.1 (11.9)	<b>65.5</b> (9.5)
	+ paraphrase	<b>18.8</b> (13.7)	<b>27.3</b> (11.6)	<b>39.9</b> (10.4)	<b>49.4</b> (12.4)	65.1 (9.4)
	<i>p</i> -value	0.04	0.01	< 0.01	0.07	0.1
	SciBERT-Base	15.2 (16.9)	29.1 (18.2)	45.5 (14.0)	60.8 (9.0)	72.2 (7.3)
	+ paraphrase	<b>19.2</b> (15.5)	<b>32.6</b> (18.4)	<b>48.9</b> (14.4)	<b>62.5</b> (8.3)	<b>73.4</b> (6.8)
	<i>p</i> -value	0.01	< 0.01	0.02	< 0.01	< 0.01
	BC-RoBERTa-Large	<b>54.2</b> (8.1)	63.9 (9.2)	73.3 (3.8)	77.4 (2.7)	81.5 (1.5)
	+ paraphrase	53.0 (7.4)	<b>67.2</b> (6.6)	<b>74.5</b> (2.7)	<b>79.1</b> (1.6)	<b>81.8</b> (1.2)
	<i>p</i> -value	0.08	0.01	0.03	< 0.01	0.02

Table 4: Few-shot results on the MedSTS task using various pre-trained models with training and development sets of different sizes. We report mean (standard deviation) performance of Pearson correlation coefficient over 10 different random seeds, along with the *p*-value of the paired *t*-test. + paraphrase: with paraphrases of selected rare biomedical words.

threshold \ Dataset		16		32		64		128		256		Test
		Train	Dev	Train	Dev	Train	Dev	Train	Dev	Train	Dev	
$t = 20k$	*	1.1	1.1	1.9	1.7	3.4	3.9	6.8	7.6	13.2	14.3	29
	**	1.1	1.1	1.9	1.8	3.7	4.4	7.3	9.0	16.3	21.1	129
$t = 50k$	*	1.8	1.5	3.5	2.2	6.8	4.8	12.4	9.2	21.8	17.0	40
	**	1.8	1.5	3.5	2.3	7.1	5.4	13.1	11.2	26.6	24.9	196
$t = 100k$	*	1.8	1.6	3.8	2.6	7.2	5.9	14.0	11.1	25.3	20.2	41
	**	1.8	1.6	3.8	2.8	7.6	6.5	14.8	13.0	30.8	25.4	179
$t = 200k$	*	1.9	2.1	4.3	3.5	7.6	6.8	13.9	12.3	23.3	21.7	41
	**	1.9	2.1	4.3	3.7	8.2	7.8	15.3	15.5	33.4	32.6	185

Table 5: The number of selected rare biomedical words in the training, development and test sets for different few-shot datasets on MedNLI. “ $t = 20k$ ” means the threshold for rare biomedical words is 20k. “\*” is the number of different rare biomedical words within the threshold. “\*\*” denotes the total occurrences of rare biomedical words in the dataset (a rare biomedical word can appear more than once).

MedNLI task since its labelled data is more objective compared to MedSTS.

**Statistics of rare biomedical words.** Table 5 shows the mean of rare biomedical words in the training, development and test sets under different thresholds for the MedNLI task. Among the above thresholds, the number of rare biomedical words varies from 29 to 41 in the test set, while the total number of all rare biomedical words varies from 129 to 196. Each rare biomedical word appears around 5 times on average in the test set. Note that as the threshold increases, the number of paraphrases does not necessarily increase as fast as the number of rare biomedical words which is because that although higher threshold does include more rare biomedical words, not all are appended with paraphrases as some of them are excluded as mentioned in Section 3.3.

In addition to the fixed threshold of 200,000, we experiment with the thresholds ranging from 20,000 to 200,000. Table 6 further demonstrates the effectiveness of paraphrases of rare biomedical words in most cases, where 200,000 of the threshold performs the best.<sup>8</sup>

## 5 Discussion

**Train with more samples.** Apart from applying paraphrases in few-shot scenarios, we also attempt with more training samples for the MedNLI task, even with full-size training dataset. We attempt to sample the same number of development samples from the official MedNLI development set

<sup>8</sup>We also attempt with thresholds higher than 200,000 but it will not bring improvement as much as 200,000.

#Samples	16	32	64	128	256
<i>w/o</i> paraphrase	51.3	60.6	71.0	80.6	83.1
$t = 20k$	55.0	61.9	72.4	80.6	83.4
$t = 50k$	54.2	62.1	73.8	81.0	<b>83.7</b>
$t = 100k$	54.5	<b>62.4</b>	73.4	80.4	83.2
$t = 200k$	<b>56.6</b>	62.3	<b>74.5</b>	<b>81.1</b>	83.6

Table 6: Results for few-shot learning with different thresholds of rare biomedical words on MedNLI with the BC-RoBERTa-Large model. *w/o* paraphrase: without paraphrases for rare biomedical words.  $t$ : the threshold for rare biomedical words.

(with 1,395 samples) to match the number of training samples. If the training set is larger than the full-size development set, we just use the whole development set.

Table 7 demonstrates that our method outperforms the baseline on four out of six cases and performs comparably on the remaining two cases. When the whole training set is used, our model with paraphrases achieves 0.8% improvement. The minor improvement might be attributed to the “not so rare” biomedical words as the rarity of words decreases during the expansion of training set, which helps the pre-trained models learn the semantics of the rare biomedical words better even without paraphrases.

**Which to look up?** Paraphrases can be helpful in general. However, for individual samples, that is not always the case. In Table 8, we further scrutinize the cases where the model yields different predictions after adding the paraphrases. It is observed that the paraphrases can be less effective if the paraphrases of rare words are not task-related

		MedNLI					
Model \ #Training		512	1024	2048	4096	8192	full-size
BC-RoBERTa-Large		84.8 (0.7)	85.5 (0.8)	86.4 (0.5)	86.3 (0.7)	86.2 (0.6)	85.9 (0.6)
	+ paraphrases	<b>85.2</b> (0.8)	<b>86.3</b> (1.0)	86.4 (0.7)	86.3 (0.6)	<b>86.7</b> (0.5)	<b>86.7</b> (0.7)
	<i>p</i> -value	0.05	< 0.01	0.43	0.36	< 0.01	0.02

Table 7: Test results on the MedNLI dataset with larger size of training sets. We report mean (and standard deviation) accuracy. + paraphrase: with paraphrases of rare biomedical words.

Sentence Pairs	w/o paraphrases	w/ paraphrases
P: She was found to have <b>BRBPR</b> ( <i>bright red blood per rectum</i> ) on rectal exam. H: the patient had bright red blood per rectum	Neutral	Entailment ( <b>right answer</b> )
P: Antenatal history - pregnancy complicated by chronic hypertension with increased gestational hypertension leading to admission 3 days prior to delivery followed by cesarean section. H: The patient had <b>proteinuria</b> ( <i>The presence of protein in the urine</i> ) during pregnancy	Entailment	Neutral ( <b>right answer</b> )
P: Following this rehab admission she was sent to a different OSH on [**2725-10-26**], for acute <b>CHF</b> ( <i>congestive heart failure</i> ) and at least one PEA arrest. H: The patient has a poorly functioning heart.	Contradiction	Entailment ( <b>right answer</b> )
P: The patient was sent to the HD unit prior to coming to the floor for <b>workup</b> ( <i>A general medical examination to assess a persons health and fitness</i> ) of fever. H: The patient has an infection	Neutral ( <b>right answer</b> )	Contradiction
P: - <b>COPD</b> ( <i>chronic obstructive pulmonary disease</i> ) - obesity - unspecified hypoxemia - CNS lymphoma c/b CVAs x3 (posterior circulation) and seizure d/o - history of SAH while on coumadin - diastolic heart failure - coronary artery disease - atrial fibrillation - hypertension - hyperlipidemia - severe OSA (did not tolerate CPAP in the past) - primary hyperparathyroidism/25-vit D deficiency c/b nephrolithiasis - toxic multinodular goiter with <b>subclinical</b> ( <i>Less than is needed for clinical reasons</i> ) hyperthyroidism - neovascular glaucoma c/b right eye blindness H: Patient has a history of malignancy	Neutral ( <b>right answer</b> )	Entailment

Table 8: Cases that model predicts differently after the supplement of paraphrases for rare biomedical words in MedNLI. “P” for Premise and “H” for Hypothesis. Words in **bold** are rare biomedical words and expressions in *italic* inside the parentheses are the paraphrases of rare biomedical words.

and critical for sentence understanding. For instance, in the third case in Table 8, “*CHF*” is a rare biomedical word which means “*congestive heart failure*”. Pre-trained models can easily match the hypothesis of “*poorly functioning*” with the given premise. Otherwise, paraphrases can possibly introduce confusion. For example, in the wrong case in Table 8, although “*workup*” is a rare biomedical word in the pre-training corpora, it does not severely affect the semantics of the sentence and the pre-trained model turns to make wrong prediction with the misdirection from supplementary paraphrase.

Actually, humans tend to continue reading unless the unknown words hinder the understanding. With this motivation, we believe that it is worthwhile to explore how to attach *informative* paraphrases for the rare words, which will be investigated in future.

## 6 Conclusion

Rare biomedical words are pervasive in biomedical texts, and understanding domain-specific rare words remains a tough challenge for pre-trained models. In this work, we presented a simple yet effective method to help the pre-trained models grasp the semantics of rare biomedical words.

Enlightened by human reading behavior, we taught the pre-trained models to understand rare biomedical words by incorporating paraphrases of rare biomedical words. Our method can be regarded as a generic plug-in approach for prompt-based tuning without additional parameters. Experiments showed that our method could substantially improve the pre-trained models under few-shot settings.



## Acknowledgements

We thank the anonymous reviewers for their helpful and constructive comments and gratefully acknowledge the support of the National Key R&D Program of China (2021ZD0113302) and the National Natural Science Foundation of China Youth Fund (62206079).

## References

- 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.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT (1)*.
- Jun Gao, Di He, Xu Tan, Tao Qin, Liwei Wang, and Tiejun Liu. 2019. Representation degeneration problem in training natural language generation models. In *International Conference on Learning Representations*.
- 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.
- Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2020. Domain-specific language model pretraining for biomedical natural language processing.
- Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2021a. Domain-specific language model pretraining for biomedical natural language processing. *ACM Transactions on Computing for Healthcare (HEALTH)*, 3(1):1–23.
- Yuxian Gu, Xu Han, Zhiyuan Liu, and Minlie Huang. 2021b. Ppt: Pre-trained prompt tuning for few-shot learning. *arXiv preprint arXiv:2109.04332*.
- Kexin Huang, Jaan Altosaar, and Rajesh Ranganath. 2019. Clinicalbert: Modeling clinical notes and predicting hospital readmission. *arXiv preprint arXiv:1904.05342*.
- Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. 2016. Mimic-iii, a freely accessible critical care database. *Scientific data*, 3(1):1–9.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240.
- Patrick Lewis, Myle Ott, Jingfei Du, and Veselin Stoyanov. 2020. Pretrained language models for biomedical and clinical tasks: Understanding and extending the state-of-the-art. In *Proceedings of the 3rd Clinical Natural Language Processing Workshop*, pages 146–157.
- Bohan Li, Hao Zhou, Junxian He, Mingxuan Wang, Yiming Yang, and Lei Li. 2020. On the sentence embeddings from pre-trained language models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9119–9130, Online. Association for Computational Linguistics.
- Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. 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 4582–4597, Online. Association for Computational Linguistics.
- 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. *arXiv preprint arXiv: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. *arXiv preprint arXiv:2110.07602*.
- 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 arXiv:1907.11692*.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448.
- George Michalopoulos, Yuanxin Wang, Hussam Kaka, Helen Chen, and Alexander Wong. 2021. Umlsbert: Clinical domain knowledge augmentation of contextual embeddings using the unified medical language system metathesaurus. In *Proceedings of the 2021*

- Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1744–1753.
- Yifan Peng, Shankai Yan, and Zhiyong Lu. 2019. Transfer learning in biomedical natural language processing: An evaluation of bert and elmo on ten benchmarking datasets. In *Proceedings of the 18th BioNLP Workshop and Shared Task*, pages 58–65.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of NAACL-HLT*, 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. *Science China Technological Sciences*, pages 1–26.
- Alexey Romanov and Chaitanya Shivade. 2018. Lessons from natural language inference in the clinical domain. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1586–1596.
- Timo Schick and Hinrich Schütze. 2020. Rare words: A major problem for contextualized embeddings and how to fix it by attentive mimicking. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8766–8774.
- Timo Schick and Hinrich Schütze. 2021. 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.
- Mike Schuster and Kaisuke Nakajima. 2012. Japanese and korean voice search. In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5149–5152.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725.
- Hoo-Chang Shin, Yang Zhang, Evelina Bakhturina, Raul Puri, Mostofa Patwary, Mohammad Shoeybi, and Raghav Mani. 2020a. Bio-megatron: Larger biomedical domain language model. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4700–4706.
- Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. 2020b. 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.
- Simon Šuster, Stéphan Tulkens, and Walter Daelemans. 2017. A short review of ethical challenges in clinical natural language processing. *arXiv preprint arXiv:1703.10090*.
- Yanshan Wang, Naveed Afzal, Sunyang Fu, Liwei Wang, Feichen Shen, Majid Rastegar-Mojarad, and Hongfang Liu. 2020. Medsts: a resource for clinical semantic textual similarity. *Language Resources and Evaluation*, 54(1):57–72.
- 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. *arXiv preprint arXiv:2109.01652*.
- Qiyu Wu, Chen Xing, Yatao Li, Guolin Ke, Di He, and Tie-Yan Liu. 2020. Taking notes on the fly helps bert pre-training. *arXiv preprint arXiv:2008.01466*.
- Yuanmeng Yan, Rumei Li, Sirui Wang, Fuzheng Zhang, Wei Wu, and Weiran Xu. 2021. ConSERT: A contrastive framework for self-supervised sentence representation transfer. 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 5065–5075, Online. Association for Computational Linguistics.
- Xi Yang, Xing He, Hansi Zhang, Yinghan Ma, Jiang Bian, Yonghui Wu, et al. 2020. Measurement of semantic textual similarity in clinical texts: comparison of transformer-based models. *JMIR medical informatics*, 8(11):e19735.
- Wenhao Yu, Chenguang Zhu, Yuwei Fang, Donghan Yu, Shuohang Wang, Yichong Xu, Michael Zeng, and Meng Jiang. 2021. Dict-bert: Enhancing language model pre-training with dictionary. *arXiv preprint arXiv:2110.06490*.
- Taolin Zhang, Zerui Cai, Chengyu Wang, Minghui Qiu, Bite Yang, and Xiaofeng He. 2021. Smedbert: A knowledge-enhanced pre-trained language model with structured semantics for medical text mining. 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 5882–5893.
- Sendong Zhao, Chang Su, Zhiyong Lu, and Fei Wang. 2021. Recent advances in biomedical literature mining. *Briefings in Bioinformatics*, 22(3):bbaa057.
- George Kingsley Zipf. 2016. *Human behavior and the principle of least effort: An introduction to human ecology*. Ravenio Books.