

Making Pretrained Language Models Good Long-tailed Learners

Chen Zhang¹, Lei Ren², Jingang Wang¹, Wei Wu², Dawei Song¹

¹Beijing Institute of Technology
{czhang, dwsong}@bit.edu.cn

²Meituan NLP
{wangjingang02, wuwei30}@meituan.com
renlei_work@163.com

Abstract

Prompt-tuning has shown appealing performance in few-shot classification by virtue of its capability in effectively exploiting pre-trained knowledge. This motivates us to check the hypothesis that prompt-tuning is also a promising choice for long-tailed classification, since the tail classes are intuitively few-shot ones. To achieve this aim, we conduct empirical studies to examine the hypothesis. The results demonstrate that prompt-tuning makes pretrained language models at least good long-tailed learners. For intuitions on why prompt-tuning can achieve good performance in long-tailed classification, we carry out in-depth analyses by progressively bridging the gap between prompt-tuning and commonly used finetuning. The summary is that the classifier structure and parameterization form the key to making good long-tailed learners, in comparison with the less important input structure. Finally, we verify the applicability of our finding to few-shot classification.¹

1 Introduction

Pretrained language models (PLMs) with CLS-tuning (i.e., finetuning a PLM by applying a classifier head over the [CLS] representation) have achieved strong performance in a wide range of downstream classification tasks (Devlin et al., 2019; Liu et al., 2019a; Raffel et al., 2020; Wang et al., 2019b). However, they have been less promising in the long-tailed scenario (Li et al., 2020). The long-tailed scenario is different from common scenarios due to the long tail phenomenon exhibited in the class distribution, as illustrated in Figure 1. The long-tailed class distribution prevents PLMs from achieving good performance, especially in




Figure 1: An example long-tailed class distribution from IFLYTEK dataset (Xu et al., 2020), where we can distinguish tail classes from head classes.

tail classes that only allow learning with very few examples (dubbed tail bottleneck).

Recent advances on Prompt-tuning² have witnessed a surge of making PLMs better few-shot learners (Schick and Schütze, 2021; Gao et al., 2021; Scao and Rush, 2021). Prompt-tuning requires PLMs to perform classification in a cloze style, thus is superior to CLS-tuning in two aspects: 1) it aligns the input structure with that of masked language modeling (MLM); and 2) it reuses the classifier structure and parameterization from the pretrained MLM head. These two merits of Prompt-tuning equip PLMs to better exploit pre-trained knowledge and hence gain a better few-shot performance than that with CLS-tuning.

The success of Prompt-tuning in few-shot scenarios motivates us to hypothesize that Prompt-tuning can relieve the tail bottleneck and thus make PLMs good long-tailed learners. The reason why we make such a hypothesis is that the tail classes are intuitively few-shot ones. However, long-tailed

¹Jingang Wang and Dawei Song are the corresponding authors.

²Good long-tailed learners are abbreviated as GLEE. Code and data are available at <https://github.com/GenZC/Glee>.

²Prompt-tuning can be an ambiguous term regarding the parameter-efficient prompt tuning (Lester et al., 2021). However, we insist on the use of it to accord with CLS-tuning.

classification is different from few-shot classification to a certain extent, as it allows the possibility to transfer knowledge from head classes to tail ones.

We empirically examine the hypothesis by conducting empirical evaluations on three long-tailed classification datasets. The comparison results show that PLMs can be good long-tailed learners with Prompt-tuning, which outperforms PLMs with CLS-tuning by large margins. Besides, Prompt-tuning even exhibits a better performance than that of CLS-tuning with appropriate calibrations (e.g., focal loss [Lin et al. 2017](#)). The widely accepted decoupling property ([Kang et al., 2020](#)) claims that a good long-tailed learner should enjoy a nearly uniform distribution in terms of weight norms, otherwise the norms of head classes can be way much larger than the tail ones'. It is therefore expected that the weights tuned with Prompt-tuning own a flat distribution. We validate the property of Prompt-tuning by visualizing the norms of trained classification weights across classes. With the compelling results, we put that our hypothesis is valid.

We also provide further intuitions by asking why Prompt-tuning could be so promising, as shown in the above empirical investigations. Through in-depth analyses, we uncover that re-using the classifier structure and parameterization from the MLM head is a key for attaining the good long-tailed performance, largely outweighing the importance of aligning the input structure with that of MLM. CLS-tuning, with classifier structure derived and parameters partly initialized from the MLM head, approximates the performance of Prompt-tuning. We believe that this observation would as well shed light on related work that aims to improve Prompt-tuning itself. As such, we finally present the comparison results of the improved CLS-tuning and Prompt-tuning in the few-shot scenario, suggesting the applicability of the improved CLS-tuning to few-shot classification.

2 Background

2.1 Long-tailed Classification

Long-tailed classification basically follows a classification setting. Given a dataset $\mathcal{D} = \{(x_i, y_i)\}_i$ in which $(x, y) \sim P(\mathcal{X}, \mathcal{Y})$, a model \mathcal{M} is required to learn to approximate $P(\mathcal{Y} | \mathcal{X})$ as accurate as possible so that it can correctly predict the label from the input. However, the long-tailed classification differs from the common classification setting in that $P(\mathcal{Y})$ is a long-tailed one, prohibiting \mathcal{M}

from achieving a good optimization, especially on tail classes.

2.2 Finetuning

CLS-tuning Pretrained with the special token [CLS] for overall semantics, PLMs can be finetuned with classifiers over the [CLS] representations for classification (Figure 2(a) left).

The optimization objective can be depicted as:

$$\mathbb{L}_{\text{CLS}} = -\log P(y_i | x_i; \mathcal{M}), \quad (1)$$

which is exactly a cross entropy loss. Here, \mathcal{M} can be disassembled to a backbone \mathcal{E} and a classifier \mathcal{C} . While \mathcal{E} is a PLM producing [CLS] representation, \mathcal{C} is a Tanh-activated MLP. Here, the MLP typically consists of two feed-forward linear layers. To be more specific, we name the classifier for CLS-tuning as CLS head (Figure 2(a) right), and generally the last layer of the classifier as predictor. We also name the input as CLS input for brevity.

Prompt-tuning For PLMs that are pretrained with an MLM objective, it is natural to finetune the PLMs in an MLM-like cloze style for better exploitation of the pretrained knowledge.

To reach the goal, a template \mathcal{T} and a verbalizer \mathcal{V} are introduced ([Schick and Schütze, 2021](#)). The template converts the original input to an input with one [MASK] token that should be recovered, in other words MLM input. The verbalizer maps all labels to their corresponding tokens, and the model should predict the token corresponding to the correct label. In particular, for a label that is mapped to multiple tokens, one [MASK] should be faced with the issue of inability of multi-token completion. We are inspired by the average strategy in [Chen et al. \(2021\)](#); [Hu et al. \(2021\)](#), and treat the average of logit values for multiple tokens as the logit value for the label.

For example (Figure 2(b) left), the template for sentiment classification can be:

$$\mathcal{T}(x) = x. \text{ It was [MASK].}, \quad (2)$$

Accordingly, the verbalizer can be:

$$\mathcal{V}(y) = \begin{cases} \text{great} & y \text{ is label:positive} \\ \text{terrible} & y \text{ is label:negative} \end{cases}, \quad (3)$$

Thereby, the optimization objective is described as:

$$\mathbb{L}_{\text{Prompt}} = -\log P(\mathcal{V}(y_i) | \mathcal{T}(x_i); \mathcal{M}), \quad (4)$$




Figure 2: Illustration of two finetuning schemes.

where \mathcal{E} in \mathcal{M} generates the [MASK] representation, and \mathcal{C} is the pretrained MLM head. The MLM head (Figure 2(b) right) is activated with a GELU (Hendrycks and Gimpel, 2016) and normalized with a layer normalization (Ba et al., 2016; Vaswani et al., 2017; Devlin et al., 2019). Note that \mathcal{E} and \mathcal{C} share a part of parameters (i.e., the word embeddings in \mathcal{E} and the predictor over the vocabulary in \mathcal{C}).

2.3 Research Hypothesis

As discussed in the leading Section 1, we observe that, in a long-tailed class distribution, each of the tail classes is provided with very few examples, typically fewer than one tenth of the number of a common class. This brings challenges in long-tailed classification. Meanwhile, Prompt-tuning has been demonstrated to make PLMs better few-shot learners by exploiting pretrained knowledge.

Therefore, we are inspired to hypothesize that Prompt-tuning can make PLMs good long-tailed learners, as pretrained knowledge is intuitively learned from a long-tailed vocabulary. In the following, we present a series of empirical examinations to test whether our hypothesis is valid or not in Section 3, and why it is so in Section 4.

3 Empirical Examination

3.1 Setup

Datasets We conduct examinations on five long-tailed classification datasets, ranging from Chinese to English ones. The first one is a medical intent question detection dataset (CMID) (Chen et al., 2020). The second one is an application category classification dataset (IFLYTEK) maintained by CLUE (Xu et al., 2020). The third one is a clinical trial criterion categorization dataset (CTC) (Zong et al., 2021). The fourth one is an entity typing dataset (MSRA) originally released as a named entity recognition dataset (Levow, 2006). The last one is a document topic classification dataset (R52)

essentially derived from Reuters 21578 dataset (Debole and Sebastiani, 2004).

For datasets that originally do not include a test set (e.g., IFLYTEK), we use the development set as test set and randomly take 10% of the training set as development set. The statistics of these datasets are listed in Table 1.

Templates and Verbalizers For Prompt-tuning, example templates for the three datasets separately are shown as below:

- CMID: $x?$ The intent of the question is [MASK].
- IFLYTEK: x . The mentioned application belongs to [MASK].
- CTC: x . The category of the criterion is [MASK].
- MSRA: x . The e in the sentence is [MASK].
- R52: x . This is [MASK].

where x denotes the input, and e denotes the mentioned entity in the sentence offered in MSRA. Here, necessary English translations of Chinese templates are used, and the according Chinese templates are listed in Appendix A.

Since there are many classes for each dataset, we leave the details on verbalizers to Appendix A. Basically, the verbalizers are deduced from class descriptions after removal of some less meaningful tokens (e.g., punctuations).

Implementation Experiments are carried out on an Nvidia Tesla V100. All models are implemented with PyTorch³ and Transformers⁴ libraries. We initialize models with the Google-released bert-base-chinese and bert-base-uncased checkpoints⁵. For parameter settings, the batch size is 32, the learning rate is 1e-5, the weight decay is 0, and lastly the gradient norm is constrained

³<https://github.com/pytorch/pytorch>

⁴<https://github.com/huggingface/transformers>

⁵<https://github.com/google-research/bert>

Table 1: Statistics of the long-tailed datasets.

Dataset	#Train exam.	#Dev exam.	#Test exam.	#Avg. tokens	#Classes
CMID	8,678	1,226	2,450	32.0	36
IFLYTEK	10,920	1,213	2,599	289.2	119
CTC	20,666	2,296	7,682	27.2	44
MSRA	106,301	11,811	8,419	82.9	26
R52	5,943	617	2,570	114.8	52

to 1. We train the models for 10 epochs with patience of 2 epochs. In order to stabilize the training procedure, we add a linear warm-up for 1 epoch. The maximum sequence length is set according to the dataset, specifically, 64 for CMID, 512 for IFLYTEK, 64 for CTC, 128 for MSRA, and 256 for R52.

Metrics Since we are more concerned with model performance across different classes, we use the macro F1 scores as main performance metric. We also offer the macro F1 scores of head (Head scores) and tail classes (Tail scores) separately to gain a fine-grained understanding of the model performance. To separate head classes from the tail classes, we sort all classes in descending order according to the number of examples within each class. According to the power law, we should get head classes that take up 80% of all examples. However, we find some tail classes with very limited examples can be included in this manner. So we manually determine the percentage for each dataset, specifically, 55% for MSRA and R52, 65% for CMID and IFLYTEK, and 80% for CTC. In addition, we also gather the accuracy scores (Acc scores) for reference. We take average scores over 5 runs as the results, attached with variances.

3.2 Comparison Results

In order to examine whether our hypothesis that Prompt-tuning can make PLMs good long-tailed learners is valid or not, we first conduct an evaluation on the long-tailed datasets. The results are given in Table 2.

A key finding from the comparison results is that Prompt-tuning outperforms CLS-tuning by large margins across datasets in terms of F1 scores. And Prompt-tuning in fact owns performance with lower variances compared with CLS-tuning. Prompt-tuning even exhibits better F1 scores than calibrated CLS-tuning (e.g., focal loss in our case). Besides, focal loss does not bring further improvement over Prompt-tuning, implying Prompt-tuning is a suffi-

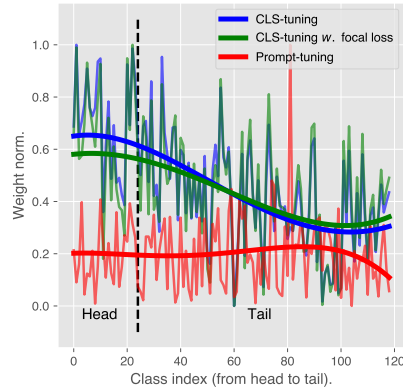


Figure 3: Weight norm visualization. Classification weight norms of the model trained on IFLYTEK.

ciently good long-tailed learner. Contrarily, calibration methods can unexpectedly degrade performance of CLS-tuning (on MSRA) due to inadequate hyperparameter search. Further, the comparison of Head and Tail scores hints that Prompt-tuning improves long-tailed performance by keeping a better trade-off between the head and tail, where Prompt-tuning achieves much better results on the tail. Yet, Prompt-tuning could slightly give a negative impact to Head scores.

Overall speaking, we can put that our hypothesis is valid, indicating a positive effect of Prompt-tuning.

3.3 Weight Norm Visualization

For sake of a deeper investigation of the hypothesis, we visualize the weight norms customized to different classes. The weights are essentially derived from the predictor of the classifier \mathcal{C} . The motivation behind the visualization is rooted on the widely accepted decoupling property (Kang et al., 2020), which claims that the learning of the backbone and classifier is in fact decoupled from each other. In other words, the long-tailed class distribution affects the classifier a lot, but might have little impact on the backbone. To this end, it is largely

Table 2: Comparison results. AVG denotes average results over all datasets. The best AVG scores are boldfaced. The variances are attached as subscripts.

Dataset	CMID		IFLYTEK		CTC		MSRA		R52		AVG	
Metric	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1
CLS-tuning	51.1 _{0.4}	37.3 _{2.3}	58.7 _{0.4}	33.7 _{1.6}	84.6 _{0.3}	77.2 _{2.9}	99.0 _{0.1}	97.5 _{1.0}	95.3 _{0.2}	67.3 _{1.3}	77.7	62.6
w/ η -norm	51.1 _{0.5}	37.4 _{2.0}	59.1 _{0.3}	35.7 _{1.6}	84.7 _{0.2}	77.3 _{3.1}	99.0 _{0.1}	97.4 _{0.9}	95.4 _{0.3}	68.9 _{1.9}	77.9	63.3
w/ focal loss	51.0 _{0.7}	42.1 _{1.3}	58.8 _{0.3}	36.0 _{1.6}	84.3 _{0.4}	78.5 _{2.4}	99.0 _{0.1}	96.8 _{1.2}	95.7 _{0.2}	72.8 _{2.3}	77.8	65.2
Prompt-tuning	49.3 _{0.7}	43.4 _{0.7}	61.2 _{0.6}	44.4 _{1.0}	84.2 _{0.1}	80.9 _{0.1}	99.1 _{0.0}	97.8 _{0.3}	95.7 _{0.1}	85.3 _{0.6}	77.9	70.4
w/ focal loss	48.6 _{0.6}	42.5 _{0.6}	59.7 _{0.6}	43.9 _{0.7}	83.5 _{0.6}	80.2 _{0.7}	99.0 _{0.1}	97.2 _{0.7}	95.5 _{0.3}	82.6 _{2.4}	77.3	69.3
Metric	Head	Tail	Head	Tail	Head	Tail	Head	Tail	Head	Tail	Head	Tail
CLS-tuning	50.3 _{1.0}	34.1 _{3.0}	61.8 _{0.6}	27.4 _{1.9}	87.7 _{0.2}	74.1 _{3.7}	99.2 _{0.1}	97.4 _{1.1}	99.0 _{0.1}	66.6 _{1.3}	79.6	59.9
w/ η -norm	50.3 _{0.9}	34.3 _{2.7}	62.1 _{0.4}	29.7 _{2.0}	87.8 _{0.2}	74.3 _{4.0}	99.2 _{0.1}	97.3 _{1.0}	99.0 _{0.1}	68.3 _{1.9}	79.7	60.8
w/ focal loss	49.8 _{0.8}	40.2 _{1.5}	62.0 _{0.4}	30.2 _{1.9}	87.5 _{0.3}	75.9 _{3.1}	99.3 _{0.0}	96.7 _{1.3}	99.0 _{0.0}	72.3 _{2.4}	79.5	63.1
Prompt-tuning	48.4 _{1.0}	42.2 _{0.7}	63.6 _{0.4}	40.1 _{1.2}	87.4 _{0.2}	79.0 _{0.2}	99.2 _{0.1}	97.7 _{0.3}	98.6 _{0.1}	85.0 _{0.6}	79.4	68.8
w/ focal loss	47.1 _{0.8}	41.4 _{0.8}	62.3 _{0.6}	39.8 _{0.8}	86.7 _{0.5}	78.3 _{0.7}	99.3 _{0.1}	97.1 _{0.7}	98.8 _{0.1}	82.3 _{2.4}	78.8	67.8

acknowledged in previous studies (Lin et al., 2017; Kang et al., 2020) that a good long-tailed learner should make weight norms roughly in the same scale, while a bad one sees a weight norm decay from the head to tail. The visualization is shown in Figure 3.

From the plot, we discover that CLS-tuning makes PLMs bad long-tailed learners, as it possesses large weight norms for the head while small ones for the tail. Applied to CLS-tuning, the focal loss slightly flattens the weight norm slope. As expected, Prompt-tuning makes PLMs own a way more flat distribution. This may be the reason why focal loss does not bring improvement over Prompt-tuning. While one may wonder whether weight norm regularization could boost long-tailed performance, η -norm is actually a calibration method to adjust weight norms so that they can be in similar scales, showcasing minor improvement of weight norm regularization.

The contrast indicates that Prompt-tuning makes PLMs good long-tailed learners.

4 Analyses

Although we have verified that our hypothesis is valid, we are more of thirst to see why Prompt-tuning can be so promising and provide further intuitions. To this demand, we consider three research questions here and carry out in-depth analyses to answer them. We also hope the analyses can shed light on the design of Prompt-tuning itself in related areas.

- **RQ1:** Does the shared embedding contribute to Prompt-tuning?

- **RQ2:** Does the input structure (i.e., MLM input) contribute to Prompt-tuning?
- **RQ3:** Does the classifier structure and parameterization (e.g., layer normalization used in MLM head) contribute to Prompt-tuning?

4.1 Impact of Shared Embedding

The first question comes into our mind is that whether it is the parameter sharing between the classifier and backbone that helps Prompt-tuning survive from collapsing. Hence, we decouple the parameters shared by the classifier and backbone (i.e., without shared embedding during optimization) and compare the results in Table 3 before and after the parameter decoupling.

We observe that the decoupling somehow has little impact on the performance (Prompt-tuning v.s. Prompt-tuning w/ ed.). Reversely, ed. even degrades the performance of Prompt-tuning. The phenomenon gives a possibly negative response to **RQ1**.

4.2 Impact of Input Structure

We explore whether the input structure is a significant factor regarding long-tailed performance. To this end, we arm CLS-tuning with MLM input so that CLS-tuning may share the input structure and representation for classification as Prompt-tuning.

The results in Table 3 demonstrate that the MLM input somewhat affects the performance of CLS-tuning, in terms of both Acc and F1 scores (CLS-tuning v.s. CLS-tuning w/ prompt). We attribute the performance detriment to mismatch of the CLS head and MLM input. That is, PLMs are not pre-

Table 3: Analysis results. AVG denotes average results over all datasets. The best AVG scores are boldfaced. The content after CLS-tuning indicates the activation that is being used, where T is Tanh and R is ReLU. LN stands for layer normalization. pt. is short for pretrained and ed. is short for embedding decoupling. The variances are attached as subscripts.

Dataset	CMID		IFLYTEK		CTC		MSRA		R52		AVG	
Metric	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1
CLS-tuning ◦ T	51.1 _{0.4}	37.3 _{2.3}	58.7 _{0.4}	33.7 _{1.6}	84.6 _{0.3}	77.2 _{2.9}	99.0 _{0.1}	97.5 _{1.0}	95.3 _{0.2}	67.3 _{1.3}	77.7	62.6
w/ η -norm	51.1 _{0.5}	37.4 _{2.0}	59.1 _{0.3}	35.7 _{1.6}	84.7 _{0.2}	77.3 _{3.1}	99.0 _{0.1}	97.4 _{0.9}	95.4 _{0.3}	68.9 _{1.9}	77.9	63.3
w/ focal loss	51.0 _{0.7}	42.1 _{1.3}	58.8 _{0.3}	36.0 _{1.6}	84.3 _{0.4}	78.5 _{2.4}	99.0 _{0.1}	96.8 _{1.2}	95.7 _{0.2}	72.8 _{2.3}	77.8	65.2
CLS-tuning ◦ R	50.9 _{0.4}	34.5 _{1.4}	58.7 _{0.3}	33.3 _{1.1}	84.4 _{0.4}	77.1 _{1.0}	99.0 _{0.1}	97.7 _{0.5}	94.2 _{0.4}	56.2 _{2.5}	77.4	59.8
w/ η -norm	50.9 _{0.5}	34.8 _{1.8}	58.4 _{0.3}	33.3 _{1.0}	84.6 _{0.4}	78.0 _{1.5}	99.1 _{0.0}	97.8 _{0.5}	94.3 _{0.3}	56.3 _{1.9}	77.5	60.0
w/ focal loss	51.0 _{0.5}	40.1 _{1.5}	58.8 _{0.4}	34.6 _{0.1}	84.6 _{0.3}	76.9 _{0.6}	99.0 _{0.1}	97.0 _{1.5}	95.1 _{0.3}	66.0 _{2.7}	77.7	62.9
CLS-tuning ◦ R												
w/ prompt	49.7 _{0.5}	33.1 _{0.4}	58.4 _{0.3}	32.8 _{1.0}	84.6 _{0.1}	77.2 _{3.0}	99.0 _{0.1}	96.9 _{0.3}	94.1 _{0.3}	54.5 _{3.1}	77.2	58.9
w/ LN	51.3 _{0.6}	42.0 _{1.4}	59.7 _{0.6}	39.1 _{0.8}	84.6 _{0.5}	79.4 _{2.2}	99.1 _{0.1}	97.1 _{0.8}	96.1 _{0.2}	77.7 _{3.5}	78.2	67.1
w/ pt. LN	50.8 _{0.6}	42.5 _{1.2}	59.4 _{0.4}	41.4 _{0.9}	84.4 _{0.5}	79.7 _{1.5}	99.1 _{0.1}	97.7 _{0.5}	96.2 _{0.2}	82.0 _{1.8}	78.0	68.7
Prompt-tuning	49.3 _{0.7}	43.4 _{0.7}	61.2 _{0.6}	44.4 _{1.0}	84.2 _{0.1}	80.9 _{0.1}	99.1 _{0.0}	97.8 _{0.3}	95.7 _{0.1}	85.3 _{0.6}	77.9	70.4
w/ ed.	49.4 _{0.7}	43.6 _{0.7}	61.0 _{0.7}	44.4 _{1.0}	84.2 _{0.4}	80.5 _{0.9}	99.0 _{0.2}	96.9 _{1.4}	95.7 _{0.2}	84.9 _{1.0}	77.9	70.1
Metric	Head	Tail	Head	Tail	Head	Tail	Head	Tail	Head	Tail	Head	Tail
CLS-tuning ◦ T	50.3 _{1.0}	34.1 _{3.0}	61.8 _{0.6}	27.4 _{1.9}	87.7 _{0.2}	74.1 _{3.7}	99.2 _{0.1}	97.4 _{1.1}	99.0 _{0.1}	66.6 _{1.3}	79.6	59.9
w/ η -norm	50.3 _{0.9}	34.3 _{2.7}	62.1 _{0.4}	29.7 _{2.0}	87.8 _{0.2}	74.3 _{4.0}	99.2 _{0.1}	97.3 _{1.0}	99.0 _{0.1}	68.3 _{1.9}	79.7	60.8
w/ focal loss	49.8 _{0.8}	40.2 _{1.5}	62.0 _{0.4}	30.2 _{1.9}	87.5 _{0.3}	75.9 _{3.1}	99.3 _{0.0}	96.7 _{1.3}	99.0 _{0.0}	72.3 _{2.4}	79.5	63.1
CLS-tuning ◦ R	50.7 _{0.4}	30.6 _{1.7}	61.7 _{0.6}	26.8 _{1.3}	87.5 _{0.2}	74.1 _{1.3}	99.2 _{0.1}	97.6 _{0.5}	99.0 _{0.1}	55.4 _{2.6}	79.6	56.9
w/ η -norm	50.4 _{0.7}	31.0 _{2.2}	61.7 _{0.7}	26.8 _{1.4}	87.7 _{0.3}	75.1 _{1.9}	99.2 _{0.1}	97.7 _{0.5}	99.0 _{0.2}	55.5 _{2.0}	79.6	57.2
w/ focal loss	50.3 _{0.5}	37.6 _{1.7}	61.9 _{0.3}	28.4 _{1.6}	87.6 _{0.3}	73.7 _{0.8}	99.2 _{0.1}	96.9 _{1.6}	99.0 _{0.2}	65.3 _{2.8}	79.6	60.4
CLS-tuning ◦ R												
w/ prompt	49.7 _{0.9}	29.1 _{0.5}	61.5 _{0.9}	26.3 _{1.3}	87.5 _{0.1}	74.2 _{3.9}	99.2 _{0.1}	96.8 _{0.3}	99.0 _{0.1}	53.6 _{3.2}	79.4	56.0
w/ LN	50.0 _{0.7}	40.0 _{1.6}	63.1 _{0.6}	33.7 _{0.9}	87.6 _{0.4}	77.0 _{2.7}	99.3 _{0.2}	97.0 _{0.8}	99.0 _{0.1}	77.3 _{3.6}	79.8	65.0
w/ pt. LN	49.5 _{0.5}	40.8 _{1.4}	62.7 _{0.3}	36.6 _{1.1}	87.5 _{0.4}	77.4 _{1.9}	99.2 _{0.1}	97.6 _{0.6}	99.0 _{0.1}	81.6 _{1.9}	79.6	66.8
Prompt-tuning	48.4 _{1.0}	42.2 _{0.7}	63.6 _{0.4}	40.1 _{1.2}	87.4 _{0.2}	79.0 _{0.2}	99.2 _{0.1}	97.7 _{0.3}	98.6 _{0.1}	85.0 _{0.6}	79.4	68.8
w/ ed.	48.2 _{1.0}	42.5 _{0.7}	63.6 _{0.4}	40.1 _{1.2}	87.4 _{0.3}	78.5 _{1.1}	99.3 _{0.1}	96.8 _{1.5}	98.7 _{0.2}	84.6 _{1.0}	79.4	68.5

trained in the way that CLS head should decode the MLM input. And the results naturally suggest *a possibly negative response to RQ2*.

4.3 Impact of Classifier Structure and Parameterization

We also investigate the impact of the classifier structure and parameterization, given structural and parameter differences between classifier heads used by CLS-tuning and Prompt-tuning. We aim to check whether it is the discrepancy between classifiers that biases the learning. To study the impact, we replace the Tanh in CLS-tuning with ReLU. We use ReLU here as an alternative of GELU owing to the fact that ReLU is more prevalent in the finetuning stage. Then, by adding a layer normalization after the ReLU activation, we fill the structural gap between two classifiers. We also perform a natural follow-on action, re-using the statistics of the layer normalization from the MLM head to further

enhance the classifier. The results are presented in Table 3.

It is revealed that the ReLU variant sometimes yields surprisingly deteriorated results when compared to CLS-tuning with Tanh. However, when the ReLU variant is additionally armed with a succeeding layer normalization (CLS-tuning w/ LN), it can surpass the original CLS-tuning by certain margins. Notably, CLS-tuning w/ LN has better Acc scores than Prompt-tuning does, potentially suggesting the balanced use of CLS-tuning w/ LN in real-world applications. Besides, by re-using the MLM layer normalization (CLS-tuning w/ pt. LN), CLS-tuning approximates Prompt-tuning at once. The results imply *an absolutely positive response to RQ3*. We conjecture the observation is underpinned by an information perspective towards regulated features.

For the ReLU variant, negative features will be zero out, leading to a cut-down of information. The

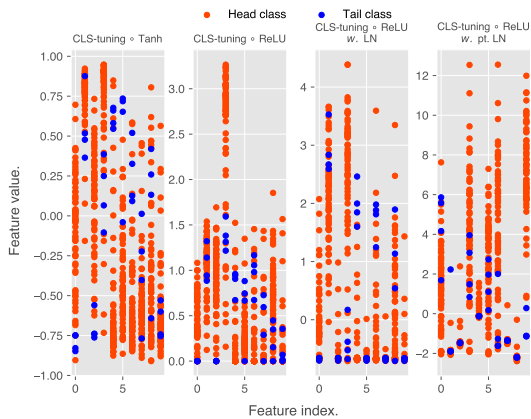


Figure 4: An illustration of sampled feature distributions, for both the head and tail.

information cut-down can be referred to as “dying ReLU” problem (Lu et al., 2019; Kayser, 2017) when negative features take a large portion. The information loss may be problematic to the head, and even more unfriendly to the tail that is under-represented (i.e., with much fewer examples). As a consequence, under-represented classes can be represented with high bias (and potentially high variance). In contrast, Tanh manipulates features without any tailoring, but restricts values to a constant range (i.e., from -1 to 1). Despite the reduced learning burden, Tanh suffers from large saturation area (Xu et al., 2016). Thereby, the information of some tail classes can be compressed with detriment. Existing literature (Girshick et al., 2014; He et al., 2015; Xu et al., 2015) calibrates the situation by relaxing negative features. On the other hand, the layer normalization can compensate the information loss caused by ReLU. With learned affinities (a.k.a., element-wise weights and biases), the layer normalization shall properly re-locate and re-scale the ReLU-activated features so that knowledge attained from the head can be transferred to the tail and the debuff of ReLU can be alleviated.

We add an illustration via Figure 4 for a more intuitive understanding of the above explanation. Taking examples of a sampled head class (*commerce*) and a sampled tail class (*delivery*) from IFLYTEK, we plot the distributions of the first 10 features in feature vectors (i.e., final hidden state vectors correspond to [CLS] tokens) derived from these examples. From the plot, we can see that ReLU certainly drops information of quite a few features of the tail (6 out of 10) and Tanh compulsively has features descend into the pre-defined range to involve themselves in optimization, both

Table 4: Statistics of few-shot datasets.

Dataset	#Test exam.	#Avg. tokens
ECOM	610	47.7
RTE	277	52.3
BOOLQ	3,270	105.3

limiting the expressiveness of features. The randomly initialized layer normalization re-activates the dead features from the ReLU by re-locating them, but largely leaving them in fixed scales (6 out of 10). Furthermore, the pretrained layer normalization from the MLM head transfers knowledge from the head to tail and makes the features of the tail diversely distributed by re-scaling them.

5 Applicability to Few-shot Classification

From the analyses above, we come up with an alternative to CLS-tuning that owns a comparable effectiveness with Prompt-tuning for long-tailed classification, by equipping CLS-tuning with the structure and the layer normalization affinities of the pretrained MLM head. With this much simpler surrogate classifier, we prefer to retrospectively explore how the classifier would perform in the few-shot scenario, especially when compared with Prompt-tuning. The intuition is that, while the layer normalization can not transfer knowledge from the head to tail any longer in few-shot classification, its innate friendliness to the tail can probably be applied to few-shot classes for generalization.

Therefore, we conduct experiments on three few-shot datasets, ECOM from FewCLUE (Xu et al., 2021), and RTE, BOOLQ from SuperGLUE (Wang et al., 2019a; Schick and Schütze, 2021). ECOM is a review sentiment classification dataset, RTE (Wang et al., 2019b) is a two-way entailment dataset, and BOOLQ (Clark et al., 2019) is a yes-or-no question answering dataset.

While ECOM has already been designed for few-shot learning, the others are not originally. For the latter two, we treat the original development set as the test set following Gao et al. (2021), and randomly sample 32 examples uniformly from the original training set as the training set following Schick and Schütze (2021). To strictly following a *true* few-shot setting (Perez et al., 2021), 32 examples that do not overlap with those in our training set are used to form the development set. The implementation is akin to the one we used for the long-tailed experiments. The batch size is re-

Table 5: Applicability to few-shot classification results. AVG denotes average results over all datasets. The best AVG scores are boldfaced. The variances are attached as subscripts.

Dataset	ECOM	RTE	BOOLQ	AVG
Metric	Acc	Acc	Acc	Acc
CLS-tuning ◦ T	63.6 _{10.5}	48.6 _{3.9}	60.1 _{1.3}	57.4
CLS-tuning ◦ R	64.7 _{2.7}	48.7 _{3.6}	55.2 _{7.5}	56.2
w/ prompt	60.0 _{5.5}	49.7 _{2.3}	59.9 _{4.3}	56.5
w/ pt. LN	68.7 _{5.0}	55.0 _{1.8}	59.1 _{1.7}	60.9
Prompt-tuning	77.3 _{9.3}	51.8 _{1.8}	60.7 _{1.0}	63.3

duced to 2 due to the small scale of training data. bert-base-uncased is leveraged as the backbone for English datasets. The maximum length is set to 64, 128, and 256 respectively for ECOM, RTE, and BOOLQ. The statistics of these datasets are listed in Table 4.

The example templates and verbalizers are listed in Appendix B.

Moreover, since few-shot experiments are sensitive to the choice of hyperparameters, we again take average accuracy scores over 5 runs as the results, attached with variances.

We can observe from Table 5 that CLS-tuning with ReLU and pretrained layer normalization can perform better than other CLS-tuning baselines, but may not necessarily hold across datasets. The filled gap between CLS-tuning and Prompt-tuning shows the scalability of our finding to few-shot classification. However, we encourage future work to explore the regime for a more comprehensive understanding.

6 Related Work

PLMs have brought classification tasks to a brand-new stage where the solutions to these tasks are way much simpler than ever (Devlin et al., 2019; Wang et al., 2019b). However, PLMs are still sub-optimal for some corner cases, such as few-shot classification (Zhang et al., 2021) and long-tailed classification (Li et al., 2020).

For few-shot classification, Prompt-tuning, which finetunes models in a language modeling fashion, is increasingly taking a central role in the mainstream methods (Liu et al., 2021). Since PLMs are mostly trained with language modeling objectives, Prompt-tuning becomes, as water is to fish, the key to unearthing the few-shot or even zero-shot learning capabilities of PLMs (Brown et al., 2020; Schick

and Schütze, 2021; Scao and Rush, 2021). Instantly after the success, the engineering of prompts drives related research on prompt search/generation (Jiang et al., 2020; Shin et al., 2020; Gao et al., 2021) and various downstream applications such as text generation (Li and Liang, 2021), relation extraction (Chen et al., 2020), and entity typing (Ding et al., 2021). Recently, Prompt-tuning also serves as an alternative way for parameter-efficient finetuning by only finetuning parameters of the inserted continuous prompts (Lester et al., 2021; Ma et al., 2022), in place of previously adopted adapter-tuning (Houlsby et al., 2019). The parameter efficiency brought by Prompt-tuning has blazed a trail for increasingly large language models.

In contrast, little work has been investigated to make PLMs good long-tailed learners. Intuitively, the tail classes are essentially few-shot ones. Thus, we speculate that Prompt-tuning is also a promising choice to make PLMs good long-tailed learners via transferring knowledge of head classes. As long-tailed classification is a long-standing problem in the general area of machine learning (Lin et al., 2017; Liu et al., 2019b; Zhou et al., 2020; Kang et al., 2020; Tang et al., 2020) and the long-tailed phenomenon also exists in the domain of natural language processing, we believe that a systematic exploration on whether and why Prompt-tuning can make PLMs good long-tailed learners will facilitate further advances in the related areas.

7 Conclusions

Inspired by the success of Prompt-tuning in few-shot learning, we empirically examine whether Prompt-tuning can make PLMs good long-tailed learners in this work. The results validate the hypothesis. We also conduct in-depth analyses on why Prompt-tuning benefits PLMs for long-tailed classification, from the perspectives of coupling, classifier, and input respectively, to offer further intuitions. Based on the analyses, we summarize that the classifier structure and parameterization are crucial for enhancing long-tailed performance of PLMs, in contrast to other factors. Extended empirical evaluation results on few-shot classification show that our finding would shed light on related work that seeks to boost Prompt-tuning.

Limitations

Since prompt-tuning is shown sensitive to small variations of templates, it should be performed with

reasonable templates. However, we do not study the impact of different templates since our work is not concerned with finding a good template for long-tailed classification.

Acknowledgements

We would like to thank Yujia Qin and Zhiyuan Liu from Tsinghua University for their helpful suggestions. This research was supported in part by Natural Science Foundation of Beijing (grant number: 4222036) and Huawei Technologies (grant number: TC20201228005).

References

- Lei Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. [Layer normalization](#). *arXiv*, 1607.06450.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *NeurIPS*.
- Nan Chen, Xiangdong Su, Tongyang Liu, Qizhi Hao, and Ming Wei. 2020. [A benchmark dataset and case study for chinese medical question intent classification](#). *BMC Medical Informatics Decision Making*, 20-S(3):125.
- Xiang Chen, Ningyu Zhang, Xin Xie, Shumin Deng, Yunzhi Yao, Chuanqi Tan, Fei Huang, Luo Si, and Huajun Chen. 2021. [Knowprompt: Knowledge-aware prompt-tuning with synergistic optimization for relation extraction](#). *arXiv*, 2104.07650.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. [Boolq: Exploring the surprising difficulty of natural yes/no questions](#). In *NAACL-HLT*, pages 2924–2936.
- Franca Debole and Fabrizio Sebastiani. 2004. [An analysis of the relative difficulty of reuters-21578 subsets](#). In *LREC*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: pre-training of deep bidirectional transformers for language understanding](#). In *NAACL-HLT*, pages 4171–4186.
- Ning Ding, Yulin Chen, Xu Han, Guangwei Xu, Pengjun Xie, Hai-Tao Zheng, Zhiyuan Liu, Juanzi Li, and Hong-Gee Kim. 2021. [Prompt-learning for fine-grained entity typing](#). *arXiv*, 2108.10604.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. [Making pre-trained language models better few-shot learners](#). In *ACL-IJCNLP*, pages 3816–3830.
- Ross B. Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. 2014. [Rich feature hierarchies for accurate object detection and semantic segmentation](#). In *CVPR*, pages 580–587.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. [Delving deep into rectifiers: Surpassing human-level performance on imagenet classification](#). In *ICCV*, pages 1026–1034.
- Dan Hendrycks and Kevin Gimpel. 2016. [Bridging non-linearities and stochastic regularizers with gaussian error linear units](#). *arXiv*, 1606.08415.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morroni, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for NLP](#). In *ICML*, volume 97 of *PMLR*, pages 2790–2799.
- Shengding Hu, Ning Ding, Huadong Wang, Zhiyuan Liu, Juanzi Li, and Maosong Sun. 2021. [Knowledgeable prompt-tuning: Incorporating knowledge into prompt verbalizer for text classification](#). *arXiv*, 2108.02035.
- Zhengbao Jiang, Frank F. Xu, Jun Araki, and Graham Neubig. 2020. [How can we know what language models know](#). *TACL*, 8:423–438.
- Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. 2020. [Decoupling representation and classifier for long-tailed recognition](#). In *ICLR*.
- Mike Kayser. 2017. What is the “dying relu” problem in neural networks? <https://qr.ae/pv2H1S>. Quora.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#). In *EMNLP*, pages 3045–3059.
- Gina-Anne Levow. 2006. [The third international chinese language processing bakeoff: Word segmentation and named entity recognition](#). In *SIGHAN@COLING/ACL*, pages 108–117.
- Xiang Lisa Li and Percy Liang. 2021. [Prefix-tuning: Optimizing continuous prompts for generation](#). In *ACL/IJCNLP*, pages 4582–4597.
- Xiaoya Li, Xiaofei Sun, Yuxian Meng, Junjun Liang, Fei Wu, and Jiwei Li. 2020. [Dice loss for data-imbalanced NLP tasks](#). In *ACL*, pages 465–476.
- Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. 2017. [Focal loss for dense object detection](#). In *ICCV*, pages 2999–3007.
- 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*, 2107.13586.

- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019a. [Roberta: A robustly optimized BERT pretraining approach](#). *arXiv*, 1907.11692.
- Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X. Yu. 2019b. [Large-scale long-tailed recognition in an open world](#). In *CVPR*, pages 2537–2546.
- Lu Lu, Yeonjong Shin, Yanhui Su, and George E. Karniadakis. 2019. [Dying relu and initialization: Theory and numerical examples](#). *arXiv*, abs/1903.06733.
- Fang Ma, Chen Zhang, Lei Ren, Jingang Wang, Qifan Wang, Wei Wu, Xiaojun Quan, and Dawei Song. 2022. [Xprompt: Exploring the extreme of prompt tuning](#). *arXiv*, abs/2210.04457.
- Ethan Perez, Douwe Kiela, and Kyunghyun Cho. 2021. [True few-shot learning with language models](#). *arXiv*, 2105.11447.
- 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](#). *JMLR*, 21:140:1–140:67.
- Teven Le Scao and Alexander M. Rush. 2021. [How many data points is a prompt worth?](#) In *NAACL-HLT*, pages 2627–2636.
- Timo Schick and Hinrich Schütze. 2021. [It’s not just size that matters: Small language models are also few-shot learners](#). In *NAACL-HLT*, pages 2339–2352.
- Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh. 2020. [Autoprompt: Eliciting knowledge from language models with automatically generated prompts](#). In *EMNLP*, pages 4222–4235.
- Kaihua Tang, Jianqiang Huang, and Hanwang Zhang. 2020. [Long-tailed classification by keeping the good and removing the bad momentum causal effect](#). In *NeurIPS*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *NeurIPS*, pages 5998–6008.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019a. [Superglue: A stickier benchmark for general-purpose language understanding systems](#). In *NeurIPS*, pages 3261–3275.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019b. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *ICLR*.
- Bing Xu, Ruitong Huang, and Mu Li. 2016. [Revise saturated activation functions](#). *arXiv*, abs/1602.05980.
- Bing Xu, Naiyan Wang, Tianqi Chen, and Mu Li. 2015. [Empirical evaluation of rectified activations in convolutional network](#). *arXiv*, abs/1505.00853.
- Liang Xu, Hai Hu, Xuanwei Zhang, Lu Li, Chenjie Cao, Yudong Li, Yechen Xu, Kai Sun, Dian Yu, Cong Yu, Yin Tian, Qianqian Dong, Weitang Liu, Bo Shi, Yiming Cui, Junyi Li, Jun Zeng, Rongzhao Wang, Weijian Xie, Yanting Li, Yina Patterson, Zuoyu Tian, Yiwen Zhang, He Zhou, Shaowei Hua Liu, Zhe Zhao, Qipeng Zhao, Cong Yue, Xinrui Zhang, Zhengliang Yang, Kyle Richardson, and Zhenzhong Lan. 2020. [CLUE: A chinese language understanding evaluation benchmark](#). In *COLING*, pages 4762–4772.
- Liang Xu, Xiaojing Lu, Chenyang Yuan, Xuanwei Zhang, Hu Yuan, Huilin Xu, Guoao Wei, Xiang Pan, and Hai Hu. 2021. [Fewclue: A chinese few-shot learning evaluation benchmark](#). *arXiv*, 2107.07498.
- Tianyi Zhang, Felix Wu, Arzoo Katiyar, Kilian Q. Weinberger, and Yoav Artzi. 2021. [Revisiting few-sample BERT fine-tuning](#). In *ICLR*.
- Boyan Zhou, Quan Cui, Xiu-Shen Wei, and Zhao-Min Chen. 2020. [BBN: bilateral-branch network with cumulative learning for long-tailed visual recognition](#). In *CVPR*, pages 9716–9725.
- Hui Zong, Jinxuan Yang, Zeyu Zhang, Zuofeng Li, and Xiaoyan Zhang. 2021. [Semantic categorization of chinese eligibility criteria in clinical trials using machine learning methods](#). *BMC Medical Informatics Decision-Making*, 21(1):128.

A Templates and Verbalizers for Long-tailed Datasets

The templates and verbalizers for three long-tailed datasets are separately listed in Table 6, Table 7, Table 8, Table 9, and Table 10.

B Templates and Verbalizers for Few-shot Datasets

The templates are:

- ECOM: x . 这个评论中的物品[MASK]好.
(x . The product mentioned in the review is [MASK] good.)
- RTE & BOOLQ: p . Question: q ? The answer: [MASK].

where p denotes the passage or premise, and q denotes the question or hypothesis. Accordingly, the verbalizers are:

- ECOM: {label:positive→很(very), label:negative→不(not)}
- RTE: {label:entailment→yes, label:not_entailment→no}
- BOOLQ: {label:true→yes,label:false→no}

Here, English translations are additionally used for ECOM.

Dataset	CMID
Template	x ?这个问题的意图是[MASK].
Verbalizer	<p>label:病症治疗方法→病症治疗方法(disease treatment), label:病症定义→病症定义(disease definition), label:病症临床表现(病症表现)→病症临床表现病症表现(disease symptom), label:药物适用症→药物适用症(medicine applicability), label:其他无法确定→其他无法确定(others undefined), label:病症禁忌→病症禁忌(disease contradiction), label:病症相关病症→病症相关病症(disease related diseases), label:其他对比→其他对比(others contrast), label:药物副作用→药物副作用(medicine side-effect), label:药物禁忌→药物禁忌(medicine contradiction), label:其他多问→其他多问(others multiple questions), label:病症病因→病症病因(disease cause), label:治疗方案化验/体检方案→治疗方案化验体检方案(treatment exmination), label:治疗方案恢复→治疗方案恢复(treatment recovery), label:病症严重性→病症严重性(disease severity), label:病症治愈率→病症治愈率(disease cure rate), label:药物用法→药物用法(medicine usage), label:药物作用→药物作用(medicine effect), label:其他两性→其他两性(others sex), label:治疗方案正常指标→治疗方案正常指标(treatment normal indication), label:其他养生→其他养生(others health), label:治疗方案方法→治疗方案方法(treatment method), label:病症传染性→病症传染性(disease infectivity), label:药物成分→药物成分(medicine component), label:病症预防→病症预防(disease prevention), label:治疗方案恢复时间→治疗方案恢复时间(treatment recovery duration), label:病症推荐医院→病症推荐医院(disease recommende hospital), label:治疗方案费用→治疗方案费用(treatment cost), label:治疗方案临床意义→治疗方案临床意义(treatment significance), label:其他设备用法→其他设备用法(others device usage), label:治疗方案疗效→治疗方案疗效(treatment efficacy), label:药物价钱→药物价钱(medicine price), label:治疗方案有效时间→治疗方案有效时间(treatment effective duration), label:其他整容→其他整容(others cosmetic), label:病症所属科室→病症所属科室(disease department), label:治疗方案治疗时间→治疗方案治疗时间(treatment duration)</p>

Table 6: Template and verbalizer for CMID (with necessary translations).

Dataset	IFLYTEK
Template	<i>x</i> .这句话描述的物品属于[MASK].
Verbalizer	<p>label:打车→打车(taxi),label:导航→导航(navigation),label:WIFI→wifi, label:租车→租车(car rent),label:同城→同城(urban service),label:快递→快递(express), label:婚庆→婚庆(wedding),label:家政→家政(house service), label:公共交通→公共交通(public transport),label:政务→政务(government affair), label:社区服务→社区服务(community service),label:薅羊毛→薅羊毛(deal hunter), label:魔幻→魔幻(magic game),label:仙侠→仙侠(warrior game), label:卡牌→卡牌(card game),label:空战→空战(flight game), label:射击→射击(shooting game),label:休闲→休闲(leisure game), label:动作→动作(action game),label:体育→体育(sports game), label:棋牌→棋牌(board game),label:养成→养成(simulation game), label:策略→策略(strategy game),label:MOBA→MOBA,label:辅助工具→辅助工具(aid tool), label:约会→约会(date),label:通讯→通讯(communication),label:工作→工作(work), label:论坛→论坛(BBS),label:婚恋→婚恋(marriage),label:情侣→情侣(lover), label:社交→社交(social),label:生活→生活(life),label:博客→博客(blog), label:新闻→新闻(news),label:漫画→漫画(cartoon),label:小说→小说(novel), label:技术→技术(technology),label:教辅→教辅(teaching),label:问答→问答(QA), label:搞笑→搞笑(fun),label:杂志→杂志(magazine),label:百科→百科(wikipedia), label:影视→影视(TV),label:求职→求职(job),label:兼职→兼职(part-time), label:视频→视频(video),label:短视频→短视频(clips),label:音乐→音乐(music), label:直播→直播(live),label:电台→电台(radio),label:K歌→K歌(KTV), label:成人→成人(adult),label:中小学→中小学(school),label:职考→职考(exam), label:公务员→公务员(civil servant),label:视频教育→视频教育(video edu.), label:高等教育→高等教育(advanced edu.),label:成人教育→成人教育(adult edu.), label:艺术→艺术(art),label:语言(非英语)→语言非英语(non-english), label:英语→英语(english),label:旅游→旅游(travel),label:预定→预定(preservation), label:民航→民航(flight),label:铁路→铁路(railway),label:酒店→酒店(hotel), label:行程→行程(route),label:民宿→民宿(BnB),label:出国→出国(abroad), label:工具→工具(general tool),label:亲子→亲子(kid),label:母婴→母婴(infant), label:驾校→驾校(driver),label:违章→违章(violation),label:汽车→汽车(vehicle), label:汽车交易→汽车交易(vehicle trade),label:养车→养车(vehicle maintainance), label:行车→行车(driving aid),label:租房→租房(house rent), label:买房→买房(house purchase),label:装修→装修(house decoration), label:问诊→问诊(inquiry),label:养生→养生(health preservation), label:医疗→医疗(medical),label:减肥→减肥(slim),label:美妆→美妆(makeup), label:菜谱→菜谱(menu),label:餐饮→餐饮(restaurant), label:健身→健身(fit),label:支付→支付(payment),label:保险→保险(insurance), label:股票→股票(stock),label:借贷→借贷(loan),label:理财→理财(finance management), label:彩票→彩票(lottery),label:记账→记账(account keeping),label:银行→银行(bank), label:剪辑→剪辑(film editing),label:修图→修图(photo editing), label:相机→相机(camera),label:绘画→绘画(painting),label:二手→二手(second-hand), label:电商→电商(e-commerce),label:团购→团购(group purchase), label:外卖→外卖(take-out),label:美颜→美颜(beauty),label:电子→电子(electronics), label:电影→电影(movies),label:超市→超市(market),label:购物→购物(shopping), label:笔记→笔记(notes),label:办公→办公(office),label:日程→日程(schedule), label:女性→女性(female),label:经营→经营(management),label:收款→收款(check), label:体育资讯→体育资讯(sports info.),label:其他→其他(others)</p>

Table 7: Template and verbalizer for IFLYTEK (with necessary translations).

Dataset	CTC
Template	<i>x</i> .这个标准的类别是[MASK].
Verbalizer	<p>label:治疗或手术→治疗或手术(therapy or surgery), label:体征→体征(body sign),label:成瘾行为→成瘾行为(addictive behavior), label:年龄→年龄(age),label:疾病→疾病(disease), label:器官组织状态→器官组织状态(organ or tissue status), label:过敏耐受→过敏耐受(allergy tolerance), label:依存性→依存性(compliance with protocol), label:风险评估→风险评估(risk assessment), label:怀孕相关→怀孕相关(pregnancy-related activity), label:诊断→诊断(diagnostic),label:复合→复合(multiple criteria), label:实验室检查→实验室检查(laboratory examination), label:知情同意→知情同意(consent),label:献血→献血(blood donation), label:参与其它试验→参与其它试验(enrollment in other studies), label:药物→药物(pharmaceutical substance or drug), label:能力→能力(capacity),label:饮食→饮食(diet), label:特殊病人体征→特殊病人体征(special patient characteristic), label:疾病分期→疾病分期(non-Neoplasm disease stage), label:研究者决定→研究者决定(researcher decision), label:数据可及性→数据可及性(data accessibility), label:预期寿命→预期寿命(life expectancy), label:肿瘤进展→肿瘤进展(neoplasm status), label:读写能力→读写能力(literacy), label:病例来源→病例来源(patience source), label:锻炼→锻炼(exercise),label:症状→症状(symptom), label:受体状态→受体状态(receptor status), label:口腔相关→口腔相关(oral related), label:种族→种族(ethnicity), label:健康群体→健康群体(the healthy), label:残疾群体→残疾群体(disabilities), label:设备→设备(device),label:性别→性别(gender), label:吸烟状况→吸烟状况(smoking status), label:性取向→性取向(sex related), label:护理→护理(nursing),label:睡眠→睡眠(bedtime), label:酒精使用→酒精使用(alcohol consumption), label:居住情况→居住情况(living condition), label:教育状况→教育状况(education), label:伦理审查→伦理审查(ethical audit)</p>

Table 8: Template and verbalizer for CTC (with necessary translations).

Dataset	MSRA
Template	x .这句话中的 e 是[MASK].
Verbalizer	<p>label:integer→整数(integer),label:ordinal→序数(ordinal), label:location→地点(location),label:date→日期(date), label:organization→机构(organization),label:person→人物(person), label:money→钱款(money),label:duration→时段(duration), label:time→时间(time),label:length→长度(length), label:age→年龄(age),label:frequency→频率(frequency), label:angle→角度(angle),label:phone→电话(phone), label:percent→百分数(percent),label:fraction→分数(fraction), label:weight→重量(weight),label:area→面积(area), label:capacity→容积(capacity),label:decimal→小数(decimal), label:measure→其他度量(other measure),label:speed→速度(speed), label:temperature→温度(temperature),label:postal code→邮政编码(postal code), label:rate→比率(rate),label:www→网址(website)</p>

Table 9: Template and verbalizer for MSRA.

Dataset	R52
Template	x . This is [MASK].
Verbalizer	<p>label:copper→copper,label:livestock→livestock, label:gold→gold,label:money-fx→money fx, label:tea→tea,label:ipi→ipi, label:trade→trade,label:cocoa→cocoa, label:iron-steel→iron steel,label:reserves→reserves, label:zinc→zinc,label:nickel→nickel, label:ship→ship,label:cotton→cotton, label:platinum→platinum,label:alum→alum, label:strategic-metal→strategic metal,label:instal-debt→instal debt, label:lead→lead,label:housing→housing, label:gnp→gnp,label:sugar→sugar, label:rubber→rubber,label:dlr→dlr, label:tin→tin,label:interest→interest, label:income→income,label:crude→crude, label:coffee→coffee,label:jobs→jobs, label:meal-feed→meal feed,label:lei→lei, label:lumber→lumber,label:gas→gas, label:nat-gas→nat gas,label:veg-oil→veg oil, label:orange→orange,label:heat→heat, label:wpi→wpi,label:cpi→cpi, label:earn→earn,label:jet→jet, label:potato→potato,label:bop→bop, label:money-supply→money supply,label:carcass→carcass, label:acq→acq,label:pet-chem→pet chem, label:grain→grain,label:fuel→fuel, label:retail→retail,label:cpu→cpu,</p>

Table 10: Template and verbalizer for R52.