

Task-guided Disentangled Tuning for Pretrained Language Models

Jiali Zeng^{1*}, Yufan Jiang¹, Shuangzhi Wu¹, Yongjing Yin², Mu Li¹

¹Tencent Cloud Xiaowei, Beijing, China

²Zhejiang University, Westlake University, Zhejiang, China

{lemonzeng, garyyfjiang, frostwu, ethanlli}@tencent.com

yinyongjing@westlake.edu.cn

Abstract

Pretrained language models (PLMs) trained on large-scale unlabeled corpus are typically fine-tuned on task-specific downstream datasets, which have produced state-of-the-art results on various NLP tasks. However, the data discrepancy issue in domain and scale makes fine-tuning fail to efficiently capture task-specific patterns, especially in the low data regime. To address this issue, we propose **Task-guided Disentangled Tuning (TDT)** for PLMs, which enhances the generalization of representations by disentangling task-relevant signals from the entangled representations. For a given task, we introduce a learnable confidence model to detect indicative guidance from context, and further propose a disentangled regularization to mitigate the over-reliance problem. Experimental results on GLUE and CLUE benchmarks show that TDT gives consistently better results than fine-tuning with different PLMs, and extensive analysis demonstrates the effectiveness and robustness of our method. Code is available at <https://github.com/lemon0830/TDT>.

1 Introduction

In recent years, pretrained language models (PLMs) trained in a self-supervised manner like mask language modeling have achieved promising results on various natural language processing (NLP) tasks (Devlin et al., 2019; Yang et al., 2019; Liu et al., 2019b), which learn general linguistic and semantic knowledge from massive general corpus. To adapt PLMs to specific NLP tasks, a commonly-used approach is fine-tuning, where the whole or part of model parameters are tuned by task-specific objectives. Despite its success, the fine-tuned models have been proven ineffective to capture task-specific patterns due to the gap between task-agnostic pre-training and the weak fine-tuning with limited task-specific data (Gu et al., 2020; Gururangan et al., 2020; Kang et al., 2020).

*Corresponding author.

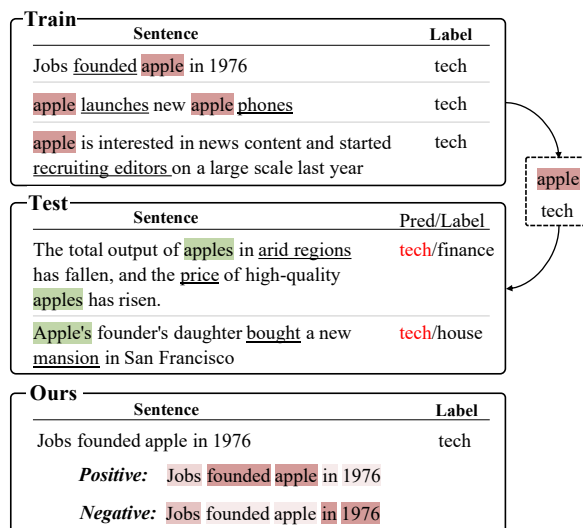


Figure 1: **An over-reliance example of news classification task.** The fine-tuned models tend to learn a simple rule that “Apple” (red) indicates “tech” class while ignore the real meaning of “apples” (green) A reliable model is expected to find out truly task-specific patterns (underlined words) instead of some high frequency but insignificant words (“apple”).

To address this problem, most existing methods focus on adapting PLMs to downstream tasks by continual pre-training on in-domain unsupervised data (Gururangan et al., 2020; Gu et al., 2020; Wu et al., 2021; Kang et al., 2020; Ye et al., 2021). For example, Gu et al. (2020) propose intermediate continual pre-training with a selective masking strategy, and Gururangan et al. (2020) adapt PLMs to in-domain tasks by domain-adaptive pretraining. Although straightforward, these kinds of methods heavily rely on the selection of large-scale additional domain corpora and the design of appropriate intermediate training tasks (Wang et al., 2019; Aghajanyan et al., 2021a).

In this paper, we propose a **Task-guided Disentangled Tuning (TDT)** for PLMs by automatically detecting task-specific informative inputs without the need of additional corpora and inter-

mediate training. The core component of TDT is a confidence model which assigns each token a confidence score, and we construct distilled samples by retaining informative tokens with high confidence scores while perturbing the rest. The confidence model performs a “deletion game” strategy, which encourages the model to perturb inputs as much as possible and to maintain the performance of downstream tasks to the greatest extent with the distilled samples. Although the informative tokens are important for downstream predictions, existing work shows that over-relying on part of these words may result in pool generalization, i.e., over-reliance problem (Moon et al., 2020; Geirhos et al., 2020; Sun et al., 2019). Take the sentences in Figure 1 as an example, when the context word “Apple” frequently co-occurs with the label “tech”, fine-tuned models may learn a spurious association by binding “Apple” and “tech”, leading to incorrect predictions of sentences which contain “apple” but belong to other categories.

Based on the observation, we further enhance our method with a disentangled regularization, aiming to distinguish task-relevant and task-irrelevant features. First, we construct two variants of the original input in a complementary view: (1) *positive* variant, which maintains the high-confidence keywords, and (2) *negative* variant, derived by a “cut-out-keyword” operation on the original input. Next, we propose a “triplet-style loss”, which makes predictions between the *original* input and the *positive* variant similar while the predictions between the *negative* variant and the other two different. To illustrate the mechanism of our disentangled regularization, we go back to Figure 1 and take the sentence “Jobs founded apple in 1976” as an example. Under the influence of the disentangled regularization, the *positive* variant tends to retain clue words for predictions (i.e., “founded apple”), while the *negative* variant, as the complement (i.e., “Jobs in 1976”), tends to be task-irrelevant.

We evaluate our TDT on a wide range of neural language understanding benchmark datasets in English and Chinese, i.e., GLUE and CLUE, and our TDT affords strong predictive performance compared with standard fine-tuning. Moreover, we conduct extensive analysis with respect to robustness to perturbation, domain generalization, and low-resource settings, from which we conclude:

- TDT learns reasonable confidence scores for input tokens.

- TDT is robust to input perturbation and domain shift by encouraging the model to learn more generalized features.
- TDT effectively captures the high-confidence decisive cues for downstream tasks, thus alleviating over-fitting in low-resource scenarios.

2 Method

In this section, we begin with a brief introduction of the vanilla Fine-tuning, and then introduce Task-guided Disentangled Tuning (TDT) in detail. Figure 2 shows the overall framework. TDT is composed of two parts: (1) token-level confidence model, which discovers the essential parts of inputs for the model prediction; (2) task-guided regularization, which promotes the model to decouple task-relevant keywords from non-keyword context.

2.1 Vanilla Fine-tuning

Given an example of training data $\langle X, y \rangle$, where $X = \{x_1, \dots, x_i, \dots, x_n\}$ is the input sequence and y is its corresponding label. We first map each token x_i to a real-valued vector e_i by an embedding layer. Then, the packed embedding output $E = \{e_i\}$ is fed into the PLM to get the contextualized sentence representations $H = \{h_{cls}, h_1, \dots, h_n\}$, and the hidden state h_{cls} is used to conduct classification with a MLP head. We fine-tune the parameters of the PLM with the cross entropy loss:

$$\mathcal{L}_{cla} = -\log P(y|H). \quad (1)$$

2.2 Token-level Confidence Model

For each token x_i , we generate a scalar $c_i \in [0, 1]$, coined confidence score, by stacking a single-layer feed-forward network with *sigmoid* activation on the top of the embedding layer:

$$c_i = \sigma(W e_i + b), \quad (2)$$

where W and b are trainable parameters. Based on the confidence score, we obtain a distilled sample $\{e_i^+\}$ defined as

$$e_i^+ = c_i \odot e_i + (1 - c_i) \odot \mu_0, \quad (3)$$

where μ_0 is a perturbation term and \odot denotes element-wise multiplication. Specifically, the perturbation term μ_0 can be a zero vector, a random Gaussian noise vector, or the average of the token

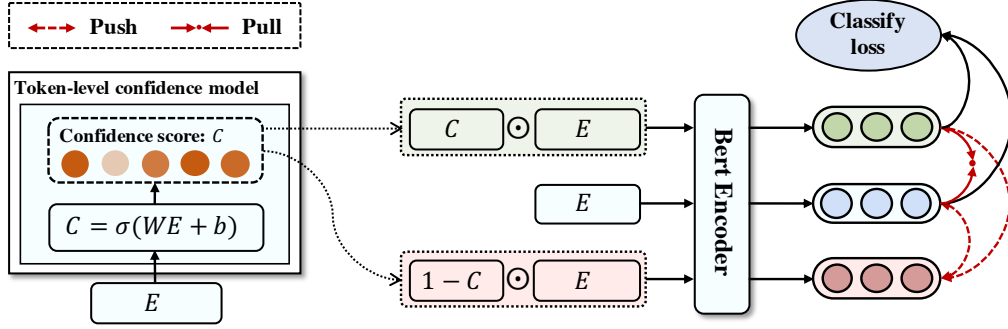


Figure 2: The overall framework of our proposed Task-guided Disentangled Tuning method.

embedding, and we choose the last one in this paper. In this manner, for the distilled sample of each training instance, the higher the c_i is, the more semantic information of the i -th token retains, while the tokens with lower scores are perturbed.

Then, the distilled sample $\{e_i^+\}$ is fed into the PLM to generate the sentence representations $H^+ = \{h_i^+\}$. Inspired by “deletion game” (Fong and Vedaldi, 2017; Voita et al., 2019), the objective function of the confidence model is

$$\mathcal{L}_C = -\log P(y|H^+) + \gamma \|C\|_2, \quad (4)$$

where $C = \{c_i\}$ is the set of confidence scores of X . The first term is the cross entropy loss of classification on the distilled sample to encourage the confidence model to assign higher scores to the more decisive part of the input, and the second term serves as a penalty to prevent the model from mode collapsing (i.e., always choosing $c_i=1$).

2.3 Task-Guided Regularization

It has been widely observed that the pretrained models tend to learn an easy-to-learn but not generalizable solution by vanilla fine-tuning on various NLP tasks (Sun et al., 2019; McCoy et al., 2019; Min et al., 2019; Niven and Kao, 2019). To alleviate this issue, we further propose a triplet-style loss on the model predictions.

Specifically, for each input sequence, we derive two different variants: a *positive* variant and a *negative* variant. The *positive* variant is expected to maintain the most informative tokens to task prediction and vice versa. As aforementioned, our confidence model removes the meaningless tokens by setting the corresponding confidence scores to zero. Based on the confidence scores, we directly treat the distilled sample generated by Eq. 3 as the *positive* variant and generate the *negative* variant as

$$e_i^- = (1 - c_i) \odot e_i. \quad (5)$$

Given the original input and the two derived variants, we feed them into the PLM with the classifier, and obtain three prediction distributions $P(y|H)$, $P(y|H^+)$, and $P(y|H^-)$. Finally, we regularize these distributions by a triplet ranking loss

$$\mathcal{L}_R = \max(m + d(P(y|H^+), P(y|H)) - d(P(y|H^-), P(y|H)) - d(P(y|H^-), P(y|H^+)), 0) \quad (6)$$

where m is a hyperparameter indicating a margin for the loss and $d(\cdot)$ denotes the Kullback-Leibler (KL) divergence. By minimizing \mathcal{L}_R , the positive variant will be closer to the original input while the negative variant will be farther from the other two. Thus, the model is encouraged to disentangle task-relevant signals from task-irrelevant factors, and generate more general representations.

2.4 Overall Training Objective

The final training objective is

$$\mathcal{L} = \mathcal{L}_{cla} + \alpha \mathcal{L}_C + \beta \mathcal{L}_R, \quad (7)$$

where α and β are non-negative hyper-parameters to balance the effect of each loss term.

3 Experiments

3.1 Datasets

We evaluate our proposed method by fine-tuning the pretrained models on the General Language Understanding Evaluation (GLUE) (Wang et al., 2018) and the Chinese Language Understanding Evaluation (CLUE) (Xu et al., 2020). Concretely, the GLUE benchmark has 8 different text classification or regression tasks including MNLI, MRPC, QNLI, QQP, RTE, SST-2, SST-B, and CoLA. The CLUE benchmark includes 9 tasks spanning several single-sentence/sentence-pair classification tasks, and we choose 5 tasks, OCNLI, IFLYTEK, CSL, TNEWS,

Model	MNLI	QQP	QNLI	SST-2	CoLA	STS-B	MRPC	RTE	Avg
<i>BERT-base</i>									
FineTuning	84.5	90.9	91.3	92.8	60.5	88.7	85.1	67.5	82.66
TDT	85.3	91.2	91.9	93.7	62.4	89.3	87.5	71.8	84.14
<i>BERT-large</i>									
FineTuning †	86.6	91.3	92.3	93.2	60.6	90.0	88.0	70.4	84.05
FineTuning	85.9	90.9	92.3	93.9	61.5	90.0	86.0	75.1	84.45
TDT	86.4	91.4	92.6	94.3	66.2	89.9	88.5	75.8	85.64
<i>RoBERTa-large</i>									
FineTuning †	90.2	92.2	94.7	96.4	68.0	92.4	90.9	86.6	88.92
FineTuning	90.5	92.3	94.4	96.6	67.4	92.2	91.9	87.7	89.13
TDT	90.6	91.9	94.7	97.0	69.3	92.5	93.1	91.0	90.01
XLNet †	90.8	92.3	94.9	97.0	69.0	92.5	90.8	85.9	89.15
ELECTRA †	90.9	92.4	95.0	96.9	69.1	92.6	90.8	88.0	89.46
DeBERTa †	91.1	92.4	95.3	96.8	70.5	92.6	91.9	88.3	89.86
ALBERT †	90.8	92.2	95.3	96.9	71.4	93.0	90.9	89.2	89.96

Table 1: **Experimental results on GLUE language understanding benchmark.** When take RoBERTA-large as the PLM, for RTE and STS, we follow Liu et al. (2019b) to finetune starting from the MNLI model instead of the baseline pretrained model. Methods with † denote that we directly report the scores from corresponding paper, and others are from our implementation.

Task	<i>BERT-wwm-base</i>		<i>MacBERT-large</i>		<i>RoBERTa-wwm-large</i>	
	FineTuning	TDT	FineTuning	TDT	FineTuning	TDT
OCNLI	74.6	75.3	78.3	79.8	78.1	79.5
IFLYTEK	60.8	62.2	61.5	61.8	61.8	62.9
CSL	84.7	85.5	86.8	87.0	86.1	87.2
TNEWS	56.9	57.3	58.5	58.7	59.0	59.2
AFQMC	74.0	75.0	76.2	76.8	76.0	76.2
Avg	70.20	71.06	72.26	72.82	72.20	73.00

Table 2: **Experimental results on CLUE language understanding benchmark.** For TNEWS, we only use the raw “sentence” for classification without the “keywords” information. For CSL, we only mask the “abst” sequence and keep the “keywords” sequence unchanged in our proposed method.

and AFQMC. The detailed data statistics and metrics are provided in Appendix A.

3.2 Model & Training

We use the pretrained models and codes provided by HuggingFace¹. We take BERT-base (Devlin et al., 2019), BERT-large (Devlin et al., 2019) and RoBERTa-large (Liu et al., 2019b) as our backbones on GLUE, while BERT-wwm-base (Cui et al., 2019), MacBERT-large (Cui et al., 2020), and RoBERTa-wwm-large (Cui et al., 2019) on CLUE. We tune the task specific hyper-parameters $m \in \{0, 2\}$, $\alpha \in \{0.5, 2, 4\}$ and $\beta \in \{0.5, 1\}$. Detailed experimental setups are shown in Appendix B. Following previous work (Lee et al., 2020; Agha-

janyan et al., 2020), we report results of the development sets, since the performance on the test sets is only accessible on the leaderboard with a limitation of the number of submissions.

3.3 Main Results

Results on GLUE. We illustrate the experimental results on the GLUE benchmark in Table 1. We can observe that the PLMs enhanced by *TDT* outperforms *FineTuning* by a large margin across all the tasks. Specifically, *TDT*s achieve 1.48 points, 1.19 points and 0.88 points (on average) improvement over BERT-base, BERT-large, and RoBERTa-large, respectively. In particular, *BERT-base+TDT* achieves competitive performance compared with *BERT-large+FineTuning*, showing that our method

¹<https://github.com/huggingface/transformers>

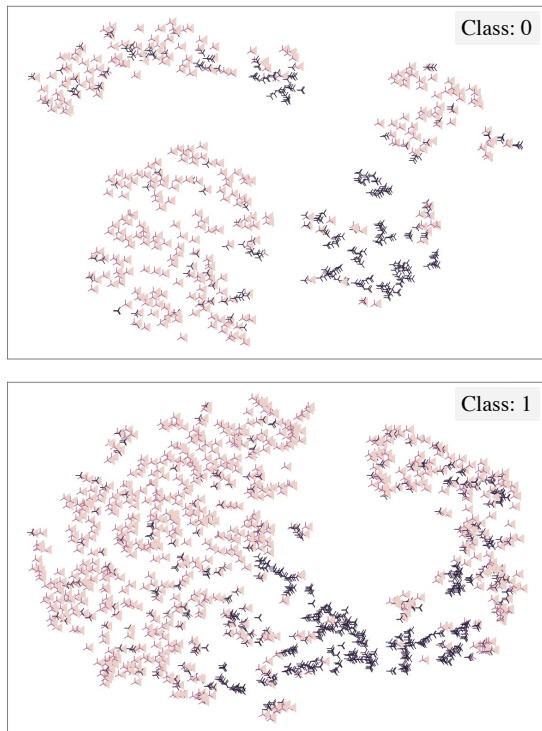


Figure 3: **Visualization of representations** of original input and two derived variants, where the triangle-shaped (pink), tri-up-shaped (purple), and tri-left-shaped (black) points denote the representations of *original input*, *positive variants*, and *negative variants*, respectively.

is more efficient to find task-specific information for downstream tasks. This may be because our training strategy prompts the models to predict with as little information as possible, isolating the task-related signals from the whole representations.

RoBERTa-large trained with *TDT* surpasses *XLNet-large* (Yang et al., 2019) *ALBERT-xxlarge* (Lan et al., 2019), *DeBERTa-large* (He et al., 2020), and *ELECTRA-large* (Clark et al., 2020), which are specially designed with different architectures and pre-training objectives.

Results on CLUE. Table 2 shows the overall results on the 5 tasks of CLUE benchmark. Concretely, *TDT* significantly outperforms *FineTuning* on CSL, IFLYTEK, AFQMC, and OCNLI, and shows competitive results on the short text classification task TNEWS, indicating the advantage of extracting important parts from long text or multiple input sequences. Note that TNEWS generally requires additional knowledge (e.g., keywords) as a supplement due to the short input, and thus cannot show the superiority of *TDT*.

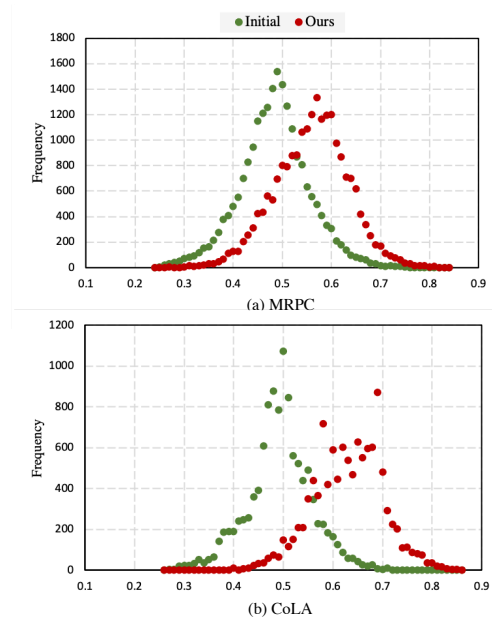


Figure 4: **Distribution of confidence scores** on MRPC and CoLA dev sets.

4 Analysis & Discussion

4.1 Visualization of Representations

In Figure 3, we plot t-SNE visualizations (van der Maaten and Hinton, 2008) of three kinds of representations generated by *BERT-large* trained with *TDT* on CoLA dev set. We can see that the representations of the original input are close to those of the positive variant in the same class. Although the negative variant representations are really similar to the original ones which derive the former, they are clearly separated from the other representations. The learned disentangled representations reveal that **the model trained with *TDT* is able to distinguish task-specific keywords and non-keyword context, which plays an important role in increasing models’ robustness.**

4.2 Distribution of Confidence Scores

We investigate the learned confidence score distributions in Figure 4. It shows that although the initial distribution is consistent, the model learns different task-specific patterns (confidence distributions) on different tasks.

4.3 Does our Confidence Model make a meaningful estimation for input tokens?

In section 2.2, we mention that *TDT* uses a scalar for evaluating the contribution of each input token. To analyze whether the strategy can successfully learn a meaningful importance estimation, we con-

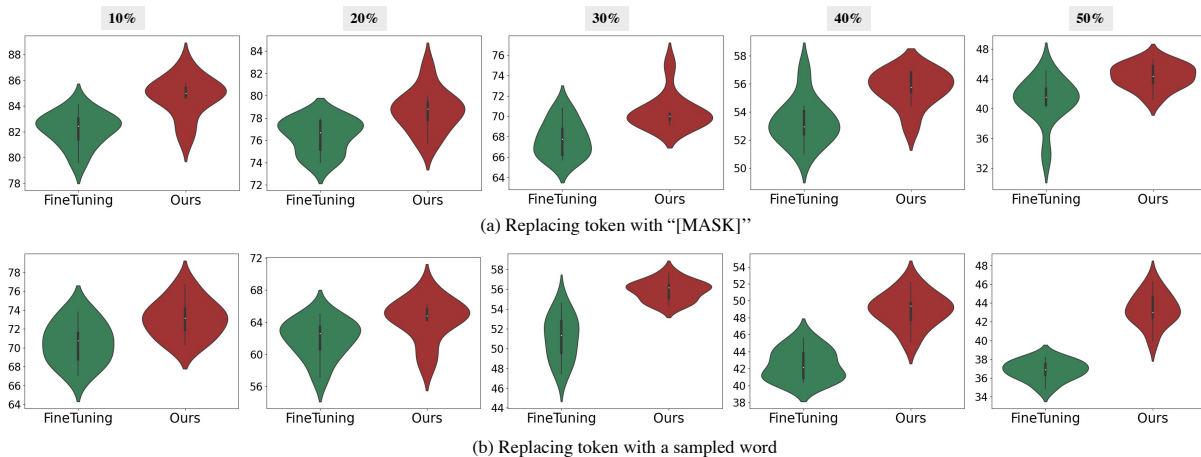


Figure 5: **Robustness to Input Perturbation.** The Y-axis is the accuracy on the development set.




Figure 6: **Accuracy of BERT-large trained with different methods and evaluated on MPRC dev set with different drop rates.** We denote vanilla fine-tuning as FineT. The solid lines indicate results on the datasets constructed by dropping tokens in descending order of confidence scores. The dotted lines denotes results on the datasets constructed by dropping tokens in increasing order of confidence scores.

struct two sets of datasets based on MRPC dev set and then evaluate the performance of BERT-large with TDT and standard fine-tuning. Specifically, we convert the confidence scores to probability distributions. We generate the first set of datasets by dropping input tokens in descending order of the distributions and generate the second set in ascending order. In order to ensure language fluency, we replace each dropped token with a “[MASK]” token. The results are shown in Figure 6 and we observe that:

- **TDT is more robust to incomplete input compared with Fine-tuning.** Specifically, although the performance of both *FineTuning* and *TDT* drops with the increase of dropout rate, our *TDT* achieves significantly better per-

Task	FineTuning	TDT	Δ
MNLI (<i>BERT-large</i>)			
MNLI-m	85.8	86.4	+0.6
QQP	73.1	74.2	+1.1
OCNLI (<i>MacBERT-large</i>)			
CMNLI	70.6	71.8	+1.2
BUSTM	64.8	66.4	+1.6

Table 3: **Performance of Domain Generalization.** The models are trained on MNLI/OCNLI but tested on out-of-domain data.

formance than *FineTuning* over all datasets.

- **Our learned confidence scores make reasonable assessments for each input token.** Particularly, regardless of the dropout rates and the training methods, dropping input tokens by the descending order of the masking scores always leads to worse performance.

4.4 Robustness to Input Perturbation

Based on the observation in Section 4.3, we further investigate the robustness of TDT on perturbed data. To construct perturbed data, we use the dev set of MRPC and possibly replace the input at each position with a “[MASK]” token or a token sampled from the input sequence. For each dropout rate, we construct 10 datasets with different random seeds and draw violin plots of the performance of BERT-large trained with TDT and fine-tuning (Figure 5). We can see that *Ours* is consistently better than *FineTuning* in all groups, indicating the superior robustness to noisy data.

Task	FineTuning	TDT	Δ
CLUE (MacBERT-large)			
OCNLI	60.85 (± 2.66)	63.38 (± 0.90)	+2.53
IFLYTEK	54.12 (± 0.75)	54.78 (± 0.94)	+0.66
CSL	80.25 (± 1.36)	81.45 (± 0.62)	+1.20
TNEWS	53.50 (± 0.58)	53.33 (± 0.25)	-0.17
AFQMC	64.77 (± 3.87)	66.45 (± 0.93)	+1.68
Avg	62.70	63.88	+1.18

Table 4: **Experimental results in low-resource scenarios.** We run 4 times for each task with different random seeds and report the average accuracy and the standard deviation.

4.5 Domain Generalization

We evaluate how well the trained models generalize to out-of-domain data on MNLI and OCNLI, Natural Language Inference (NLI) tasks of GLUE and CLUE respectively. In detail, we fine-tune BERT-large on MNLI, and test the accuracy of the fine-tuned models on other NLI datasets in different domains including MNLI-mismatch² and QQP. Besides, we fine-tune MacBERT-large on OCNLI and conduct an evaluation on CMNLI³ and BUSTM⁴. Detailed of Label Mapping is provided in Appendix C. As Table 3 illustrates, TDT outperforms vanilla fine-tuning across different out-of-domain datasets. The results suggest that **TDT encourages the model to learn more generalized features rather than some superficial contextual cues unique to training data.**

4.6 Results in Low-resource Scenarios

Fine-tuning PLMs on very small amount of training data can be challenging and result in unstable performance due to the serious over-fitting issue. In this section, we explore the effectiveness of TDT in such scenarios. For each dataset in CLUE, we use MacBERT-large and sample 1k training examples as its training data. As Table 4 demonstrates, *TDT* improves the accuracy by 1.18 on average and reduces the standard deviation by up to 2.94. It suggests that **our TDT is more stable and efficient than vanilla fine-tuning when training PLMs on limited data.**

²MNLI-mismatch has different domains from MNLI training data

³An NLI task of CLUE.

⁴A short text matching task of FewCLUE (Xu et al., 2021a)

4.7 Compared with Variants

Ablation Studies. We first conduct ablation studies to explore the effectiveness of two additional loss functions introduced in this paper and show the results in Table 5. We find that removing any of them leads to a performance drop, which indicates their effectiveness on regularization for training.

Soft Perturbation vs. Hard Perturbation. The confidence score in this paper is continuous value ranging from 0 to 1, and we perturb the input in a soft way. It is straightforward to investigate the discrete counterpart. To this end, we model the discrete confidence score with the Gumbel-Softmax trick (Jang et al., 2017). More detailed is introduced in Appendix D. We denote the model trained with the hard strategy as *TDT-hard* and show the comparison in Table 5. From the table, both *TDT-hard* and TDT yield better performance than vanilla fine-tuning. This observation supports our claim that different tokens or phrases contribute differently to the final results, which can be detected by task-guided signal and then used to model more reliable encoders by our proposed regularization. Moreover, the inferior performance of *TDT-hard* shows that naively removing tokens has an adverse effect on context modeling and thus it is better to regularize the over-reliance in a soft manner.

4.8 Compared with Previous Methods

TDT vs. Token Cutoff. Our method can also be viewed as a soft variant of token cutoff (Shen et al., 2020), which is a data augmentation strategy. Table 5 shows the results where we find that *TDT* performs better than *TokenCutoff*, which demonstrates that the improvement of our method is not entirely due to the effect of data augmentation but stems from the design of the training objectives.

TDT vs. R-drop & R3F. Recently, Liang et al. (2021) proposed R-drop to regularize the consistency of sub-models obtained through dropout. Aghajanyan et al. (2021b) introduced R3F rooted in trust region theory, which adds noise into the input embedding and minimize the KL divergence between prediction distributions given original input and noisy input. Both of them are task-agnostic, while our proposed method constructs two derived variants with task signal, and concentrates on how to disentangle the task-relevant and task-irrelevant factors. The better performance of TDT compared with the strong R-drop and R3F baselines (Table 5)

Model	GLUE (<i>RoBERTa-large</i>)					CLUE (<i>RoBERTa-www-large</i>)				
	SST-2	CoLA	MRPC	RTE	Avg	OCNLI	IFLYTEK	CSL	TNEWS	Avg
FineTuning	96.6	67.4	91.9	87.7	85.90	78.1	61.8	86.1	59.0	71.25
TokenCutoff †	96.9	70.0	90.9	90.6	87.10	78.2	61.8	86.1	59.2	71.33
R-drop †	96.9	70.0	91.4	88.4	86.67	78.9	61.6	86.6	58.9	71.50
R3F †	97.0	71.2	91.6	88.5	87.07	-	-	-	-	-
PostTraining	95.0	64.7	91.2	84.1	83.75	76.5	62.1	87.0	58.9	71.13
TDT w/o \mathcal{L}_C	96.4	69.3	91.9	89.5	86.77	78.6	61.9	86.9	59.0	71.60
TDT w/o \mathcal{L}_R	96.4	66.7	91.4	90.6	86.28	79.2	62.1	86.9	58.9	71.77
TDT-hard	96.7	67.6	92.2	90.3	86.70	79.1	62.5	87.0	59.1	71.93
TDT	97.0	69.3	93.1	91.0	87.60	79.5	62.9	87.2	59.2	72.20

Table 5: Results of RoBERTa-large trained with TDT, variants or previous methods on 4 GLUE tasks and 4 CLUE tasks. For GLUE, results with † are taken from the corresponding paper.

verify the advantage of task-driven regularization.

TDT vs. Post-Training. Post-training is an effective approach to reduce the objective gap between pretrained model and downstream tasks (Gu et al., 2020), which continues to train PLMs on task (or in-domain) training data with mask language model (MLM) loss. The difference lies in that we focus on the fine-tuning stage. Here, we compare TDT with the model first post-trained via MLM on training set of each task and then fine-tuned. It is surprising that post-training does not always have a positive effect on downstream fine-tuning, while TDT shows effective performance without additional post-training time consumption.

5 Related Work

Fine-tuning large-scale PLMs tends to be a popular paradigm of various NLP tasks (Devlin et al., 2019; Liu et al., 2019a; Yang et al., 2019). However, the fine-tuned models fail to capture task-specific patterns due to the imbalanced nature between the large number of parameters and limited training data (Aghajanyan et al., 2020). To address this issue, two main research lines are proposed: (1) continual pretraining after general pre-training, (2) regularization techniques in fine-tuning.

Continual pretraining of PLMs on unlabeled data of a given downstream domain or task has been proved effective for the end-task performance (Gururangan et al., 2020), and various continual pre-training objectives designed for different downstream tasks have been proposed (Tian et al., 2020; Wu et al., 2021). For example, Gu et al. (2020) propose a selective masking strategy to learn task-

specific patterns based on mid-scale in-domain data. However, such methods usually rely on extra in-domain data and manually designed training objectives.

Due to the overfitting problems of fine-tuning, lots of regularization techniques have been proposed. Lee et al. (2019) and Chen et al. (2020) regularize fine-tuned weights with original pretrained weights while others design adversarial training objectives or introduce noise into the input (Zhu et al., 2020; Jiang et al., 2020; Aghajanyan et al., 2020; Shen et al., 2020; Yu et al., 2021; Hua et al., 2021; Qu et al., 2020). Liang et al. (2021) regularize the training by minimizing the KL-divergence between the output distributions of two sub-models sampled by dropout and Xu et al. (2021b) only updates a sub-set of the whole network during fine-tuning by selectively masking out the gradients in both task-free and task-driven ways. Moon et al. (2020) handle the over-reliance problem by reconstructing keywords based on other words and making low-confidence predictions without enough context.

6 Conclusion

In this paper, we propose task-guided disentangled tuning for enhancing the efficiency and robustness of PLMs in downstream NLP tasks. Our method is able to efficiently distinguish task-specific features and task-agnostic ones, and bridges the gap between pretraining and adaptation without the need of immediate continual training. Experiments on GLUE and CLUE benchmarks demonstrate the effectiveness of our method, and extensive analysis shows the advantage in domain generalization and low-resource setting over fine-tuning.

References

- Armen Aghajanyan, Anchit Gupta, Akshat Shrivastava, Xilun Chen, Luke Zettlemoyer, and Sonal Gupta. 2021a. Muppet: Massive multi-task representations with pre-finetuning. *arXiv preprint arXiv:2101.11038*.
- Armen Aghajanyan, Akshat Shrivastava, Anchit Gupta, Naman Goyal, Luke Zettlemoyer, and Sonal Gupta. 2020. Better fine-tuning by reducing representational collapse. *arXiv preprint arXiv:2008.03156*.
- Armen Aghajanyan, Akshat Shrivastava, Anchit Gupta, Naman Goyal, Luke Zettlemoyer, and Sonal Gupta. 2021b. [Better fine-tuning by reducing representational collapse](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Sanyuan Chen, Yutai Hou, Yiming Cui, Wanxiang Che, Ting Liu, and Xiangzhan Yu. 2020. [Recall and learn: Fine-tuning deep pretrained language models with less forgetting](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7870–7881, Online. Association for Computational Linguistics.
- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2020. Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555*.
- Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, Shijin Wang, and Guoping Hu. 2020. [Revisiting pre-trained models for Chinese natural language processing](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 657–668, Online. Association for Computational Linguistics.
- Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, Ziqing Yang, Shijin Wang, and Guoping Hu. 2019. Pre-training with whole word masking for chinese bert. *arXiv preprint arXiv:1906.08101*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ruth C. Fong and Andrea Vedaldi. 2017. [Interpretable explanations of black boxes by meaningful perturbation](#). In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 3449–3457. IEEE Computer Society.
- Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard S. Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. 2020. [Shortcut learning in deep neural networks](#). *CoRR*, abs/2004.07780.
- Yuxian Gu, Zhengyan Zhang, Xiaozhi Wang, Zhiyuan Liu, and Maosong Sun. 2020. [Train no evil: Selective masking for task-guided pre-training](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6966–6974, Online. Association for Computational Linguistics.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don’t stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. Deberta: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*.
- Hang Hua, Xingjian Li, Dejing Dou, Chengzhong Xu, and Jiebo Luo. 2021. [Noise stability regularization for improving BERT fine-tuning](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3229–3241, Online. Association for Computational Linguistics.
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. [Categorical reparameterization with gumbel-softmax](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Tuo Zhao. 2020. [SMART: Robust and efficient fine-tuning for pretrained natural language models through principled regularized optimization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2177–2190, Online. Association for Computational Linguistics.
- Minki Kang, Moonsu Han, and Sung Ju Hwang. 2020. [Neural mask generator: Learning to generate adaptive word maskings for language model adaptation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6102–6120, Online. Association for Computational Linguistics.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.
- Cheolhyoung Lee, Kyunghyun Cho, and Wanmo Kang. 2019. Mixout: Effective regularization to fine-tune large-scale pretrained language models. *arXiv preprint arXiv:1909.11299*.

- Cheolhyoung Lee, Kyunghyun Cho, and Wanmo Kang. 2020. [Mixout: Effective regularization to finetune large-scale pretrained language models](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Xiaobo Liang, Lijun Wu, Juntao Li, Yue Wang, Qi Meng, Tao Qin, Wei Chen, Min Zhang, and Tie-Yan Liu. 2021. [R-drop: Regularized dropout for neural networks](#). *CoRR*, abs/2106.14448.
- Hairong Liu, Mingbo Ma, Liang Huang, Hao Xiong, and Zhongjun He. 2019a. [Robust neural machine translation with joint textual and phonetic embedding](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3044–3049, Florence, Italy. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. [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, Florence, Italy. Association for Computational Linguistics.
- Sewon Min, Eric Wallace, Sameer Singh, Matt Gardner, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2019. [Compositional questions do not necessitate multi-hop reasoning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4249–4257, Florence, Italy. Association for Computational Linguistics.
- Seung Jun Moon, Sangwoo Mo, Kimin Lee, Jaeho Lee, and Jinwoo Shin. 2020. [Masker: Masked keyword regularization for reliable text classification](#). *arXiv preprint arXiv:2012.09392*.
- Timothy Niven and Hung-Yu Kao. 2019. [Probing neural network comprehension of natural language arguments](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4658–4664, Florence, Italy. Association for Computational Linguistics.
- Yanru Qu, Dinghan Shen, Yelong Shen, Sandra Sajeew, Jiawei Han, and Weizhu Chen. 2020. [Coda: Contrast-enhanced and diversity-promoting data augmentation for natural language understanding](#). *arXiv preprint arXiv:2010.08670*.
- Dinghan Shen, Mingzhi Zheng, Yelong Shen, Yanru Qu, and Weizhu Chen. 2020. [A simple but tough-to-beat data augmentation approach for natural language understanding and generation](#). *CoRR*, abs/2009.13818.
- Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2019. [How to fine-tune BERT for text classification?](#) In *Chinese Computational Linguistics - 18th China National Conference, CCL 2019, Kunming, China, October 18-20, 2019, Proceedings*, volume 11856 of *Lecture Notes in Computer Science*, pages 194–206. Springer.
- Hao Tian, Can Gao, Xinyan Xiao, Hao Liu, Bolei He, Hua Wu, Haifeng Wang, and Feng Wu. 2020. [SKEP: Sentiment knowledge enhanced pre-training for sentiment analysis](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4067–4076, Online. Association for Computational Linguistics.
- Laurens van der Maaten and Geoffrey Hinton. 2008. [Visualizing data using t-sne](#). *Journal of Machine Learning Research*, 9(86):2579–2605.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Senrich, and Ivan Titov. 2019. [Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5797–5808, Florence, Italy. Association for Computational Linguistics.
- Alex Wang, Jan Hula, Patrick Xia, Raghavendra Pappagari, R. Thomas McCoy, Roma Patel, Najoung Kim, Ian Tenney, Yinghui Huang, Katherin Yu, Shuning Jin, Berlin Chen, Benjamin Van Durme, Edouard Grave, Ellie Pavlick, and Samuel R. Bowman. 2019. [Can you tell me how to get past sesame street? sentence-level pretraining beyond language modeling](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4465–4476, Florence, Italy. Association for Computational Linguistics.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Han Wu, Kun Xu, Linfeng Song, Lifeng Jin, Haisong Zhang, and Linqi Song. 2021. [Domain-adaptive pre-training methods for dialogue understanding](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 665–669, Online. Association for Computational Linguistics.
- 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 *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pages 4762–4772. International Committee on Computational Linguistics.

Liang Xu, Xiaojing Lu, Chenyang Yuan, Xuanwei Zhang, Hu Yuan, Huilin Xu, Guoao Wei, Xiang Pan, and Hai Hu. 2021a. Fewclue: A chinese few-shot learning evaluation benchmark. *CoRR*, abs/2107.07498.

Runxin Xu, Fuli Luo, Zhiyuan Zhang, Chuanqi Tan, Baobao Chang, Songfang Huang, and Fei Huang. 2021b. Raise a child in large language model: Towards effective and generalizable fine-tuning. *arXiv preprint arXiv:2109.05687*.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32.

Qinyuan Ye, Belinda Z Li, Sinong Wang, Benjamin Bolte, Hao Ma, Wen-tau Yih, Xiang Ren, and Madihan Khabisa. 2021. On the influence of masking policies in intermediate pre-training. *arXiv preprint arXiv:2104.08840*.

Yue Yu, Simiao Zuo, Haoming Jiang, Wendi Ren, Tuo Zhao, and Chao Zhang. 2021. Fine-tuning pre-trained language model with weak supervision: A contrastive-regularized self-training approach. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1063–1077, Online. Association for Computational Linguistics.

Chen Zhu, Yu Cheng, Zhe Gan, Siqu Sun, Tom Goldstein, and Jingjing Liu. 2020. FreeLB: Enhanced adversarial training for natural language understanding. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

A GLUE and CLUE Benchmark

In this paper, we conduct experiments on 8 datasets in GLUE benchmark (Wang et al., 2018), and 5 datasets in CLUE (Xu et al., 2020), including the short text classification task TNEWS, the long text classification tasks IFLYTEK and CSL, and sentence-pair classification tasks AFQMC and OCNLI. The data statistics and evaluate metrics are illustrated in Table 6.

Dataset	# Train	# Dev	Metrics
GLUE			
MNLI	393k	9.8k	Accuracy
QQP	364k	40k	Accuracy
QNLI	105k	5.5k	Accuracy
SST-2	67k	872	Accuracy
CoLA	8.5k	1.0k	Matthews Corr
STS-B	5.7k	1.5k	Spearman Corr
MRPC	3.7k	408	Accuracy
RTE	2.5k	277	Accuracy
CLUE			
OCNLI	50k	3k	Accuracy
IFLYTEK	12.1k	2.6k	Accuracy
CSL	20k	3k	Accuracy
TNEWS	53.3k	10k	Accuracy
AFQMC	34.3k	4.3k	Accuracy
CMNLI	391k	12k	Accuracy
CLUEWSC	1.2k	304	Accuracy

Table 6: Data Statistics and Evaluate Metrics.

Task	Batch Size	Steps	Warmup	lr
GLUE				
<i>BERT-base</i>				
MNLI	128	10000	1000	4e-5
QQP	128	10000	1000	4e-5
QNLI	64	3000	300	4e-5
SST-2	64	3000	300	4e-5
CoLA	64	2000	200	2e-5
STS-B	64	3000	300	4e-5
MRPC	64	2000	200	1e-5
RTE	64	2000	200	2e-5
<i>BERT-large & RoBERT-large</i>				
MNLI	64	10000	1000	2e-5
QQP	64	10000	1000	2e-5
QNLI	64	3000	300	2e-5
SST-2	64	3000	300	2e-5
CoLA	32	3000	300	2e-5
STS-B	64	3000	300	2e-5
MRPC	64	2000	200	2e-5
RTE	64	2000	100	2e-5
CLUE				
<i>BERT-wwm-base</i>				
OCNLI	64	3000	300	4e-5
IFLYTEK	16	5000	300	3e-5
CSL	32	3000	300	3e-5
TNEWS	64	5000	300	3e-5
AFQMC	32	3000	300	3e-5
<i>MacBERT-large & RoBERT-wwm-large</i>				
OCNLI	32	3000	300	1e-5
IFLYTEK	16	5000	300	1e-5
CSL	32	3000	300	1e-5
TNEWS	64	5000	300	1e-5
AFQMC	32	3000	300	1e-5

Table 7: Hyperparameters settings for different pre-trained models on variant tasks.

B Settings for Different Pretrained Models

In this paper, we fine-tuned different pretrained models with TDT, including BERT-base, BERT-

large, RoBERTa-large for GLUE and BERT-wwm-base, MacBERT-large, RoBERTa-wwm-large for CLUE. The batch size, training steps, warmup steps, and learning rate are listed in Table 7.

C Label Mapping in Domain Generalization

QQP has two labels, *duplicate* and *not duplicate*. We map *entailment* to *duplicate* and map both *neutral* and *contradiction* to *not duplicate*. BUSTM⁵ is a short text matching task of FewCLUE (Xu et al., 2021a). We use the public test set. BUSTM has two labels, *0* and *1*. We map *entailment* to label *1*, and map both *neutral* and *contradiction* to label *0*.

D Detailed of TDT-hard

Gumbel-Softmax trick (Jang et al., 2017) is an approximation to sampling from the *argmax*. Formally, we replace Eq. 2 by:

$$c_i = \operatorname{argmax}(\sigma_{\text{Gumbel}}(z(e_i))), \quad (8)$$

$$\sigma_{\text{Gumbel}}(z_i) = \frac{\exp((\log(z_i) + g_i)/\tau)}{\sum_{j=1}^K \exp((\log(z_j) + g_j)/\tau)}, \quad (9)$$

where $g_i \sim \text{Gumbel}(0,1)$, $z(\cdot)$ returns the logits produced for a given input, and τ is the temperature. By this way, if c_i is 0, the embedding of the i -th token is set to the embedding of the “[MASK]” token, otherwise the embedding remains unchanged.

⁵<https://github.com/xiaobu-coai/BUSTM>