

# PTP: Boosting Stability and Performance of Prompt Tuning with Perturbation-Based Regularizer

Anonymous ACL submission

## Abstract

Recent studies show that prompt tuning can better leverage the power of large language models than fine-tuning on downstream natural language understanding tasks. However, the existing prompt tuning methods have training instability issues, as the variance of scores under different random seeds is quite large. To address this critical problem, we first investigate and find that the loss landscape of vanilla prompt tuning is precipitous when it is visualized, where a slight change of input data can cause a big fluctuation in the loss landscape. This is an essential factor that leads to the instability of prompt tuning. Based on this observation, we introduce perturbation-based regularizers, which can smooth the loss landscape, into prompt tuning. We propose a new algorithm, called Prompt Tuning with Perturbation-based regularizer (PTP), which can not only alleviate training instability dramatically but also boost the performance of prompt tuning. We design two kinds of perturbation-based regularizers, including random-noise-based and adversarial-based. In particular, our proposed perturbations are flexible on both text space and embedding space. Extensive experiments show the effectiveness of our proposed methods in stabilizing the training. Our new algorithms improve the state-of-the-art prompt tuning methods by 1.94% and 2.34% on SuperGLUE and FewGLUE benchmarks, respectively.

## 1 Introduction

Releasing the burden of training models from scratch while keeping the outstanding performance on downstream tasks, pretrained Language Models (LMs) brought NLP to a new era (Raffel et al., 2020; He et al., 2021; Shoeybi et al., 2019). Since BERT (Vaswani et al., 2017), fine-tuning all the parameters of pretrained LMs becomes a common practice. However, it is memory-consuming to store a copy of the entire LM for each downstream task due to the number of parameters in LM can




Figure 1: A simple pipeline of our Prompt Tuning with Perturbation-based regularizer (PTP) algorithm.

be 10B or even 100B (Shoeybi et al., 2019; Brown et al., 2020).

Recently, inspired by the success of GPT-3 (Brown et al., 2020) on few-shot and zero-shot learning with manually created prompts, there has been a surging interest in prompting that freezes pre-trained LM and wraps up the input sequence with natural language templates. However, natural language prompts are handcrafted by experts and the performance is not comparable with fine-tuning methods. To tackle it, Lester et al. (2021); Li and Liang (2021) proposed prompt tuning, which prepends the input sequence with continuous embeddings and only tunes these embeddings during training. Liu et al. (2021b, 2022) verified the effectiveness of prompt tuning on natural language understanding (NLU) tasks under both few-shot learning and supervised learning settings, which is comparable to the fine-tuning methods but with much fewer ( $1000\times$  less) task-specific tunable parameters. However, under different random seeds, we observe that the current prompt tuning methods suffer from a high variance of scores, which indicates they suffer from training instability issues.

To investigate the factor that causes the instability of prompt tuning, we visualize the loss land-

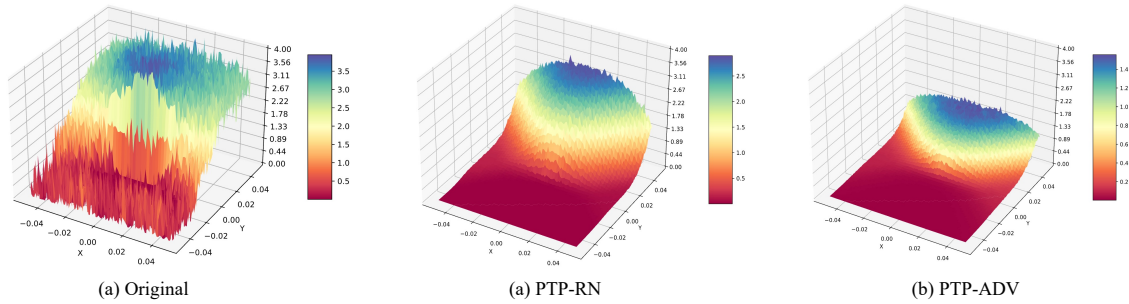


Figure 2: The loss landscapes on different continuous prompt tuning procedures. The X-axis and Y-axis denote the magnitude of perturbations (gradient direction) and another random orthogonal direction. Z-axis represents the cross-entropy loss of the different training methods.

scape of the vanilla prompt tuning and observe that there exist many sharp crests, as shown in Figure 2(a), which harms the training stability. Motivated by the recent study (Chen and Hsieh, 2020) which shows that perturbation-based regularizers are powerful tools to smooth the loss landscape and stabilize the training of machine learning systems, we introduce them into prompt tuning to address the lack of stability and generalization issues. To be specific, we propose Prompt Tuning with Perturbation-based regularizer (PTP) algorithm to make the training stable and boost the performance of prompt tuning.

Specifically, we consider two kinds of perturbations in PTP, Random-Noise-based (PTP-RN) perturbation and ADversarial-based (PTP-ADV) perturbation. PTP-RN is motivated by randomized smoothing (Cohen et al., 2019), which applies the neighborhood averaging and makes the neural network smoother while PTP-ADV is motivated by adversarial training, a method proposed to make the predictor capable of resisting the adversarial examples (Goodfellow et al., 2015) as well as boosting the clean accuracy of the predictors (Xie et al., 2020; Zhu et al., 2020). Moreover, in order to bring more flexibility and make our exploration more comprehensive, we apply perturbations to both text (discrete) and embedding (continuous) space, as depicted in Figure 1.

In the experiments, we conduct extensive experiments to evaluate our proposed algorithms, PTP-RN and PTP-ADV, on SuperGLUE (Wang et al., 2019) and FewGLUE (Schick and Schütze, 2021b) benchmark. By applying PTP algorithm on text or embedding space to the existing prompt tuning methods, we can boost the performance of prompt tuning on SuperGLUE and FewGLUE benchmarks

by 1.94% and 2.34% as well as make prompt tuning more stable. We also present a comprehensive ablation study and analysis of our algorithms with different perturbations.

Our contributions can be summarized as:

- We propose new PTP algorithm to tackle the training instability problem in prompt tuning, which can also boost the performance. Together with PTP algorithm, we design two types of perturbations as our implicit regularizers, which are Random-Noise-based perturbation (PTP-RN) and ADversarial-based perturbation (PTP-ADV).
- Moreover, as depicted in Figure 1, our proposed PTP-ADV and PTP-RN can be applied to both text space and embedding space, which makes our perturbations more flexible.
- We conduct extensive experiments to evaluate the effectiveness of our algorithms on SuperGLUE and FewGLUE benchmarks. The experimental results demonstrate our proposed PTP algorithm can boost the standard performance of prompt tuning on FewGLUE and SuperGLUE by 2.34% and 1.94%, respectively. It also shows the great power of our algorithm in improving training stability.

## 2 Preliminaries and Related Work

### 2.1 Prompt Tuning

**Discrete Prompt.** Discrete prompt, also known as hard prompt (Liu et al., 2021a; Davison et al., 2019; Jiang et al., 2020; Haviv et al., 2021), is typically a template composed of task descriptions and original input texts. Brown et al. (2020) created templates for GPT-3 based on their introspection and make it suitable for various downstream tasks,

such as machine translation, QA, *etc.* Utilizing discrete prompts, they can achieve stunning results on NLU tasks under few-shot learning settings. By employing the predefined templates and converting the tasks into cloze questions, Schick and Schütze (2021a,b) showed that even with ‘greener’ backbone (Lan et al., 2020), which has 100x fewer parameters than GPT-3, they can also reach prevailing results on the few-shot version of SuperGLUE benchmark (Wang et al., 2019) (also known as FewGLUE). Liu et al. (2021b) utilized the existing discrete templates and tuned embeddings of the selected tokens, which achieves SOTA results on FewGLUE. As the few-shot scenario is common and useful, in this paper, we also test the few-shot learning ability of our PTP in our experiments.

**Continuous Prompt Tuning.** As prompts aim to boost LM’s performance, it is not necessary to make tokens of prompts interpretable. Without the limit of tokens being natural words, Li and Liang (2021) proposed to prepend a series of tunable embeddings  $E_T$  to the input embedding sequence as prompt and optimize them with training data, which provides in-context information for LMs to condition on. Lester et al. (2021) prepended a sequence of special tokens  $T$  to the input sequence, and similarly, they tune the embeddings  $E_T$  (embeddings of special tokens) on the downstream tasks. Moreover, to further leverage the power of prompt embeddings, Liu et al. (2022) presented PT2, a method that adds the trainable prompt embeddings to every layer of pretrained LMs as prefix embeddings. To keep the consistency of the notation, we also apply the same prompt embedding representation  $E_T$  to represent trainable prompts in every layer of LMs.

## 2.2 Adversarial Training

**Continuous Space AT.** Over the past few years, Adversarial Training (AT) has demonstrated impressive results in improving model robustness (Goodfellow et al., 2015; Tramèr et al., 2018; Athalye et al., 2018). AT can be formulated as a min-max optimization problem:

$$\min_{\theta} \mathbb{E}_{(x_i, y_i) \sim \mathbb{D}} \left[ \max_{\|\delta\|_p \leq \epsilon} \mathcal{L}(\theta, x + \delta, y) \right], \quad (1)$$

where  $\mathcal{L}$  is the loss function,  $\|\cdot\|_p$  represents  $\ell_p$ -norm distance and  $\epsilon$  denotes the perturbation budget. Madry et al. (2018) proposed PGD algorithm

to compute an adversarial example (inner maximization problem) iteratively as:

$$\delta^{t+1} = \Pi_{\|\delta\|_{\infty} \leq \epsilon} \left( \delta^t + \alpha \text{sign}(\nabla_{\delta} L(\theta, \delta^t, y)) \right), \quad (2)$$

where  $t$  is the iteration step;  $\Pi_{\|\delta\|_{\infty} \leq \epsilon}$  projects perturbation  $\delta$  into  $\epsilon$ -ball.

Besides enhancing the robustness against adversarial examples, adversarial training has been shown great power in boosting the standard performance of image classifier (Xie et al., 2020), visual language representation learning (Gan et al., 2020), and GNN (Kong et al., 2022). In this paper, we also apply a similar idea to prompt tuning and focus on boosting its performance rather than its adversarial robustness.

**Discrete Space AT.** Different from adversarial attacks on images, in NLP attack, due to the discrete nature of text space, it is typically formulated as a combinatorial optimization problem to create adversarial input sequence  $s'$ , which is classically solved by heuristic search while maintaining the semantic similarity of the original input sequence  $s$  (Li et al., 2020; Ren et al., 2019; Morris et al., 2020). However, the searching algorithms of adversarial attacks, such as beam search (Ebrahimi et al., 2018), greedy search based on word importance (Ren et al., 2019), and deletion-based searching (Jin et al., 2019), are usually slow because of the high computation cost on sentence encoding and large search space (Yoo et al., 2020). Yoo and Qi (2021) proposed A2T algorithm to accelerate the heuristic search process, which replaces the slow USE encoder (Li et al., 2020) with DistilBERT (Sanh et al., 2019) to calculate the cosine similarity between original input text and perturbed text, and they obtain significant speedup comparing to the textfooler (Jin et al., 2019). Thus, in this paper, we adapt the attacking algorithm in (Yoo and Qi, 2021) to generate noises in text space.

## 3 Proposed Framework

### 3.1 Preliminary

Before introducing the proposed algorithms, we first briefly describe the word to embedding process of LMs as well as the final embedding input of the continuous prompt tuning. Given the  $n$  word input sequence  $s = [s_1, \dots, s_n]$  and word to embedding function  $f_V$ , where  $V$  denotes embedding matrix, the input embedding  $E_s = [e(s_1), \dots, e(s_n)] \in \mathbb{R}^{n \times d}$  can be obtained by  $E_s = f_V(s)$ , where

|      | RTE         | BoolQ       | WiC         |
|------|-------------|-------------|-------------|
| Acc. | 87.7 ± 1.81 | 83.9 ± 0.92 | 72.0 ± 1.38 |
| Var. | 1.45        | 0.74        | 1.16        |

Table 1: Re-implementation results of PT2 (Liu et al., 2022) with RoBERTa-large backbone on SuperGLUE benchmark. Acc.: mean accuracy. Var.: variance score computed by 5 runs with different random seeds.

---

### Algorithm 1: PGD on prompt tuning

---

1 **Require:** Perturbation iteration  $n$  and size  $\alpha$ . The bound of perturbation  $\epsilon$  ;  
2 **for**  $epoch = 1, \dots, n$  **do**  
3      $E_s.requires\_grad \leftarrow \text{True};$   
4      $\hat{y}_i \leftarrow \arg \max_y (\Pr [y | \mathcal{M} (I')]);$   
5      $\mathcal{L}(\hat{y}, y).backward();$   
6      $E'_s \leftarrow E_s + \alpha * E_s.grad.sign();$   
7      $\delta = \Pi_{\|\delta\|_\infty \leq \epsilon} (E'_s - E_s);$   
8      $E'_s \leftarrow E_s + \delta;$   
9      $\mathcal{M}.zero\_grad();$   
10 **return**  $E'_s$

---

$d$  denotes the dimension of the word embedding. Let continuous prompts be  $E_T = [e_t^1, \dots, e_t^m] \in \mathbb{R}^{m \times d}$ , the update of  $E_T$  can be either directly, or through a reparameterization encoder  $P$  like MLP/LSTM,  $P(E_T) = [h_0, \dots, h_m] \in \mathbb{R}^{m \times d}$ . For simplicity, we still use  $E_T$  to represent the output of  $P(E_T)$ , and the final input embedding sequence is written as  $[E_T; E_s] \in \mathbb{R}^{(m+n) \times d}$ .

The training objective of continuous prompt tuning can be formulated as:

$$\min_{\theta} \mathbb{E}_{(s,y) \sim \mathbb{D}} [\mathcal{L}(\mathcal{M}(\theta, s, y))], \quad (3)$$

where  $M$  denotes LM;  $\theta$  represents the trainable parameters of  $E_T$  and prompt encoder  $P$  while  $\mathbb{D}$  is the underlying data distribution.

### 3.2 Proposed Formulation

Although continuous prompt tuning (Liu et al., 2022) could achieve comparable performance with the fine-tuning method by only using 0.1% to 3% trainable parameters, it suffers from unstable training issues: as shown in Table 1, even only changing the random seed in different runs, the final performance is very unstable. To investigate the issue, we plot the loss landscape of the vanilla prompt tuning as Figure 2 and observe sharp crests in a small local region with a small noise, which means

---

### Algorithm 2: PTP

---

1 **Require:** Prompt embeddings  $E_T$ ; input embeddings  $E_s$ ; trainable parameter  $\theta$  for prompt encoder  $P$  and  $E_T$ ; Training data  $D$ ; Pre-trained LM  $\mathcal{M}$ ; Loss function  $\mathcal{L}$  ;  
2 Initialize parameters  $\theta$ ;  
3 **for**  $epoch = 1, \dots, K$  **do**  
4     */\* standard prompt tuning \*/*  
5     Sample a minibatch data  $(s, y)$  from  $D$ ;  
6      $\Theta.requires\_grad \leftarrow \text{True};$   
7      $I \leftarrow [E_T; E_s];$   
8      $\hat{y} \leftarrow \arg \max_y (\Pr [y | \mathcal{M} (I)]);$   
9      $\mathcal{L}(\hat{y}, y).backward()$  and update  $E_T$  ;  
10    */\* training with perturbed data \*/*  
11     $\Theta.requires\_grad \leftarrow \text{False};$   
12    Apply PTP-RN or PTP-ADV to  $s$  or  $E_s$ ;  
13     $I' \leftarrow [E_T; E'_s];$   
14     $\Theta.requires\_grad \leftarrow \text{True};$   
15     $\hat{y}_i \leftarrow \arg \max_y (\Pr [y | \mathcal{M} (I')]);$   
16     $\mathcal{L}(\hat{y}, y).backward()$  and update  $E_T$ ;

---

a small perturbation on embedding space would cause a significant reduction in the final accuracy. It is also known as the training instability problem (Chen and Hsieh, 2020). To address this challenge in prompt tuning, we propose perturbation-based regularizers to force the loss landscape to be smooth. Specifically, we introduce two versions of perturbation-based regularizers that can be formulated as follows:

$$\begin{aligned} & \min_{\theta} \mathbb{E}_{(s,y) \sim \mathbb{D}} [\mathcal{L}(\mathcal{M}(\theta, s + \delta, y))], \text{ s.t.} \\ & \text{PTP-RN: } \delta \sim \mathcal{N} \\ & \text{PTP-ADV: } \delta = \max_{\|\delta\| \leq \epsilon} \mathcal{L}(\theta, s + \delta, y), \end{aligned} \quad (4)$$

where  $\mathcal{N}$  denotes Gaussian distribution. For PTP-RN, we minimize  $\theta$  under small random perturbation, aiming to force the model to focus on perturbed pair  $(s + \delta, y)$  and have a robust prediction within the neighborhood of  $s$ . It is related to the idea of randomized smoothing (Cohen et al., 2019), which obtains a smoother predictor via randomly averaging the neighborhood of the given function. For PTP-ADV, the perturbation  $\delta$  is generated by adversarial attack algorithms such as PGD (Madry et al., 2018), A2T (Yoo and Qi, 2021), and the worst-case training loss is minimized under small perturbation bounded by  $\epsilon$ . The idea is motivated by adversarial training, which is usually applied as a form of adversarial defense. (Goodfellow et al., 2015; Cheng et al., 2021)



Since we are the first to investigate the training stability issue of the prompt tuning and it is still unknown which space to inject perturbation  $\delta$  is better, we apply it on both text and embedding space to bring more flexibility and make our exploration more comprehensive. The perturbed sequence  $s'$  or the perturbed embedding  $E'_s$  can be obtained as

$$\begin{aligned} s' &= s + \delta, \\ E'_s &= E_s + \delta, \end{aligned} \quad (5)$$

where  $s, E_s$  denote the input sequence and the input embedding, respectively. It is worth noticing that if the perturbation is on  $s$  (text space), through  $f_V$ , the perturbed  $s'$  will be converted into input embedding, which is also denoted as  $E'_s$ .

The main idea of our proposed formulation is that we force our algorithm to not only learn from the clean data pair  $(s, y)$  but also perturbed data pair  $(s + \delta, y)$  to make the training more smooth.

### 3.3 PTP-RN

#### 3.3.1 Embedding Space (RG Perturbation)

In embedding space, how to create perturbed examples is still an unsolved problem. But since the ultimate effects of PTP-RN are the only thing we care about, not the interpretability, it is feasible for us to add random-noise-based perturbation on word embeddings. Given the embeddings of the input sequence  $E_s = [e(s_0), e(s_1), \dots, e(s_n)]$ , where  $e(s_i) \in \mathbb{R}^d$ , PTP-RN in embedding space samples  $\delta$  from Gaussian distribution and randomly selects some embeddings to perturb, which make sure. The perturbation  $\delta$  can be formulated as:

$$\begin{aligned} E'_s &= E_s + \delta, \text{ s.t.} \\ \delta &= \{\delta_1, 0, \dots, \delta_i, 0\}, \end{aligned} \quad (6)$$

where  $\delta \in \mathbb{R}^{n \times d}$  has the same length as the input embeddings;  $i$  denotes the number of embeddings being perturbed and  $\delta_{n=1, \dots, i} \sim \mathcal{N}(0, \sigma \mathbb{I}_d)$ , with  $d$  denoting the dimension of the word embedding and  $\sigma$  controlling the magnitude of perturbation. We represent PTP-RN on embedding space as PTP+RG.

#### 3.3.2 Text Space (RM Perturbation)

In text space, similarly, our goal is to create label-preserving and perturbed input data to augment the training data and make the training stable. Given an input sequence  $s$ , PTP-RN randomly selects some tokens and converts them into [MASK]. The perturbed sequence  $s'$  can be formulated as:

$$s' = \{s_0, [\text{MASK}], \dots, [\text{MASK}], s_n\}, \quad (7)$$

where we perturb  $i$  tokens. It should be noticed that unlike BERT pretraining process (Vaswani et al., 2017), where the model predicts the label of [MASK], our model will not predict anything on the tokens we mask and we just use [MASK] token as a perturbation on discrete space. PTP+RM is used to denote PTP-RN on text embedding space.

### 3.4 PTP-ADV

#### 3.4.1 Embedding Space (PGD Perturbation)

Different from previous PGD training methods which focus on improving the models' robustness, we aim to smooth the loss landscapes and boost the performance of prompt tuning by adding adversarial-based regularization. Given the embedding sequence  $E_s$ , PTP-ADV adopts multi-step PGD to generate perturbations on embedding space. The perturbation  $\delta$  is computed iteratively as:

$$\begin{aligned} E'_s &= E_s + \delta^t, \text{ s.t.} \\ \delta^t &= \Pi_{\|\delta\|_\infty \leq \epsilon} (\delta^{t-1} + \alpha \nabla_\delta L) \end{aligned} \quad (8)$$

where  $\delta^t$  denotes the  $t$ -th iterations of PGD perturbation and it will be added to the input embedding sequence  $E_s$  after all the iterations are finished. Algorithm 1 shows the implementation details of our PGD attack on prompt tuning. PTP+PGD is applied to denote our PTP-ADV algorithm with perturbation on embedding space.

#### 3.4.2 Text Space (A2T Perturbation)

Furthermore, to enhance the flexibility of PTP-ADV and boost model generalization ability, we apply it to the text space: PTP-ADV adopts the attack algorithm in A2T (Yoo and Qi, 2021) to generate its perturbation  $\delta$ , which is an algorithm composed of NLP attack and adversarial training. Given the input sequence  $s$ , the perturbed sequence  $s'$ , with A2T perturbation, can be represented as:

$$s' = \{s_0, s'_1, \dots, s'_{n-1}, s_n\}, \quad (9)$$

where  $s'_i$  denotes the perturbed word. For simplicity, we also call it PTP+A2T.

Algorithm 2 provides the details about the whole training process of PTP algorithm. In the standard prompt tuning, the input of LM is composed of prompt embedding  $E_T$  and input embedding  $E_s$ , which is denoted as  $I = [E_T; E_s]$ . After LM gives a prediction of  $I$ , we backpropagate the loss to update  $E_T$ . In training with perturbed data part (Line 10-16, Algorithm 2), the discrete or continuous space perturbations of input data are generated by

| Method  | BoolQ                       |                             | CB                          |                          | WiC                         |                             | RTE                         |                             |
|---------|-----------------------------|-----------------------------|-----------------------------|--------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|
|         | BERT                        | RoBERTa                     | BERT                        | RoBERTa                  | BERT                        | RoBERTa                     | BERT                        | RoBERTa                     |
| FT      | 77.7                        | 86.9                        | 94.6                        | 98.2                     | 74.9                        | 75.6                        | 70.4                        | 86.6                        |
| PT2     | 75.8                        | 84.8                        | 94.6                        | 100                      | 75.1                        | 73.4                        | 78.3                        | 89.5                        |
| PTP+A2T | 76.4                        | 85.7                        | 94.5                        | 99.6                     | 75.8                        | 72.9                        | 78.6                        | 89.7                        |
| PTP+RG  | 77.3                        | 86.2                        | 95.8                        | 99.8                     | <b>76.7</b> <sub>↑1.6</sub> | 75.5                        | 79.9                        | 90.6                        |
| PTP+RM  | 77.4                        | 85.9                        | 95.7                        | 100                      | 76.4                        | 75.2                        | 79.6                        | 91.6                        |
| PTP+PGD | <b>78.3</b> <sub>↑2.5</sub> | <b>86.7</b> <sub>↑1.9</sub> | <b>96.1</b> <sub>↑1.5</sub> | <b>100</b> <sub>↑0</sub> | 76.6                        | <b>75.7</b> <sub>↑2.3</sub> | <b>80.3</b> <sub>↑2.0</sub> | <b>92.0</b> <sub>↑2.5</sub> |

| Method  | COPA                        |                             | MultiRC(F1a)                |                             | ReCoRD                      |                             | WSC                         |                             |
|---------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|
|         | BERT                        | RoBERTa                     | BERT                        | RoBERTa                     | BERT                        | RoBERTa                     | BERT                        | RoBERTa                     |
| FT      | 69.0                        | 94.0                        | 70.5                        | 85.7                        | 70.6                        | 89.0                        | 68.3                        | 63.5                        |
| PT2     | 73.0                        | 93.0                        | 70.6                        | 82.5                        | 72.8                        | 89.3                        | 68.3                        | 63.5                        |
| PTP+A2T | 73.3                        | 93.2                        | 71.4                        | 82.6                        | 73.6                        | 89.7                        | 68.5                        | 63.8                        |
| PTP+RG  | <b>75.1</b> <sub>↑2.1</sub> | 93.9                        | 72.6                        | <b>84.9</b> <sub>↑2.4</sub> | 74.9                        | 90.5                        | 69.4                        | 65.0                        |
| PTP+RM  | 74.6                        | 93.8                        | 72.9                        | 84.4                        | 74.8                        | 90.6                        | 69.2                        | 64.8                        |
| PTP+PGD | 74.7                        | <b>94.1</b> <sub>↑1.1</sub> | <b>73.4</b> <sub>↑2.8</sub> | 84.6                        | <b>75.1</b> <sub>↑2.1</sub> | <b>91.9</b> <sub>↑1.6</sub> | <b>69.7</b> <sub>↑1.4</sub> | <b>65.0</b> <sub>↑1.5</sub> |

Table 2: Results of our proposed PTP algorithm in fully-supervised learning settings. We employ the large-size version of BERT and RoBERTa models (BERT-Large size: 335M and RoBERTa-large size: 355M, respectively). We use bold font to mark the best and red subscript to mark the improvement compared to the PT2.

PTP-RN or PTP-ADV firstly. Then the perturbed input is employed to conduct training with original label  $y$ , which also plays a data-augmentation role to boost the performance of the prompt tuning.

## 4 Experiments

We conducted empirical studies on two popular natural language understanding (NLU) benchmarks: SuperGLUE benchmark (Wang et al., 2019) and FewGLUE benchmark (Schick and Schütze, 2021b). We tested the proposed framework in both fully-supervised and few-shot settings to verify the effectiveness of our proposed PTP-RN and PTP-ADV algorithm with perturbations on both text and embedding space.

### 4.1 Experimental Setup

**NLU Dataset.** SuperGLUE benchmark (Wang et al., 2019) contains 8 challenging natural language understanding (NLU) tasks. We also include the few-shot version of SuperGLUE, FewGLUE benchmark (Schick and Schütze, 2021b) to test the ability of our algorithm, which consists of 32 training samples in each dataset on SuperGLUE. Following (Schick and Schütze, 2021b; Liu et al., 2021b), we report results on 7 of 8 NLU tasks in few-shot settings.

**Fully & Few-shot Setting.** In fully-supervised setting, the full training set of each task in Super-

GLUE (Wang et al., 2019) is employed during the prompt tuning process. Besides, in the model selection part, we adopt the whole validation set. As few-shot learning ability of prompt tuning can reduce the cost of annotations in real-world applications, following (Schick and Schütze, 2021b; Liu et al., 2021b), we also test our algorithm under few-shot settings. To be specific, we use the training set provided by FewGLUE (Schick and Schütze, 2021b), the few-shot version of SuperGLUE, containing 32 training pairs in each task. Besides, we use the same version of the development set as (Liu et al., 2021b) to select models, which are created by randomly choosing 32 unused training pairs.

**Baseline Methods.** We include 2 prompt tuning methods P-tuning (Liu et al., 2021b) (PT) and P-tuning-v2 (Liu et al., 2022) (PT2) as baselines. PT is the state-of-the-art method in FewGLUE benchmark while PT2 also achieves excellent performance in SuperGLUE benchmark. We defer the hyperparameters such as learning rate and prompt length in Appendix A. We also leave some figures and tables in Appendix.

**Pretrained LMs.** Following the settings in (Liu et al., 2022, 2021b), we include BERT-large (Devlin et al., 2019) and RoBERTa-large (Liu et al., 2019) for fully-supervised settings and ALBERTA-xxlarge-v2 (Lan et al., 2020) for few-shot settings.

| (Dev 32)<br>Method | BoolQ<br>(Acc.)             | CB<br>(F1)                  | WiC<br>(Acc.)               | RTE<br>(Acc.)               | MultiRC<br>(EM) (F1a)       |                             | WSC<br>(Acc.)               | COPA<br>(Acc.)              |
|--------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|
| PET Best           | 75.1                        | 83.5                        | 52.6                        | 65.7                        | 35.2                        | 75.0                        | 80.4                        | 83.3                        |
| PT                 | 77.8                        | 92.3                        | 56.3                        | 76.5                        | 36.1                        | 75.0                        | 84.6                        | 87.0                        |
| PTP+RM             | 79.9                        | 93.2                        | 58.1                        | <b>78.6</b> <sup>↑2.1</sup> | 36.2                        | 78.3                        | 85.9                        | 88.6                        |
| PTP+RG             | 79.5                        | <b>93.7</b> <sup>↑1.4</sup> | 58.0                        | 77.7                        | 36.6                        | 78.1                        | 85.4                        | 88.3                        |
| PTP+A2T            | 78.6                        | 92.6                        | 56.6                        | 77.4                        | 36.5                        | 76.0                        | 84.7                        | 87.7                        |
| PTP+PGD            | <b>80.2</b> <sup>↑2.4</sup> | 93.5                        | <b>58.5</b> <sup>↑2.2</sup> | 78.5                        | <b>37.4</b> <sup>↑1.3</sup> | <b>78.9</b> <sup>↑3.9</sup> | <b>86.0</b> <sup>↑1.4</sup> | <b>88.9</b> <sup>↑1.9</sup> |
| PET(Full Dev)      | 79.4                        | 59.4                        | 52.4                        | 69.8                        | 37.9                        | 77.3                        | 80.1                        | 95.0                        |
| iPET(Full Dev)     | 80.6                        | 92.4                        | 52.2                        | 74.0                        | 33.0                        | 74.0                        | -                           | -                           |

Table 3: Results of our PTP algorithm in Few-shot learning(32 training examples) settings. PT: P-tuning (Liu et al., 2021b) and the backbone LM is alberta-xxl-v2. We use **bold** font to mark the best. The red subscript denotes the increase of our method compared with the baseline method PT. Dev 32: development set contains 32 unused examples from the training set, same as (Liu et al., 2021b). Full Dev: original development set.

To have a fair comparison with the baseline methods, for fully-supervised settings, all backbone LMs are frozen, except in fine-tuning, same as (Liu et al., 2022). For few-shot learning settings, backbone LMs are tuned with trainable prompt embeddings, same as (Liu et al., 2021b).

## 4.2 Results on Fully-supervised Setting

In fully-supervised settings, Table 2 demonstrates the results of our proposed PTP algorithm with 4 different perturbations on all 8 tasks of SuperGLUE benchmark. It is worth noticing that PTP+PGD achieves the best performance in almost all datasets except WiC (BERT), COPA(BERT), and MultiRC (RoBERTa). Overall, the best method PTP+PGD outperforms the baseline method PT2 by 1.94% (with BERT-large backbone) and 1.63% (with RoBERTa-large backbone) on average.

**Text vs. Embedding Perturbation.** PTP with PGD and RG perturbation on continuous space (embedding space) are perform better than PTP with RM and A2T, which indicates perturbing on continuous space is more effective than perturbing on discrete space in fully-supervised settings. As for pretrained LMs (BERT-large and RoBERTa-large), results show the superb learning ability of our PTP algorithm regardless of which backbone.

## 4.3 Results on Few-shot Setting

In few-shot learning settings, we employ FewGLUE, also known as few-shot version of SuperGLUE. PET (Schick and Schütze, 2021b) and iPET (Schick and Schütze, 2021a) are the methods using discrete prompts. We test the

previous SOTA method on FewGLUE, PT (Liu et al., 2021b), as our baseline method and validate on the same development set (Dev 32). As illustrated in (Liu et al., 2021b), for a fair comparison, the results of PET Besst (Dev 32) are reported as removing all the additional tricks like ensemble, distillation, etc. PET (Full Dev) and iPET (Full Dev) denote the methods with the original validation set.

Our main results are shown in Table 3. PTP achieves better results than the previous state-of-the-art method PT in all 7 tasks, which verifies the effectiveness of our algorithms in few-shot NLU tasks. Especially, PTP+PGD outperforms the previous PT by 2.34% on average. Comparing the PTP+PGD (Dev 32) to the methods with the original dev set, it still has an advantage on most of the tasks (5 of 7) while the results are similar in BoolQ (better than PET with full dev set but worse than iPET). The PTP with RM and RG perturbation method also achieve remarkable improvement when compared to the baseline method PT. Moreover, PTP with A2T perturbation can also boost the performance of the baseline by a small margin.

## 4.4 Results on Improving Training Stability

In addition to improved generalization performance, the proposed method could stable the training process. Figure 2 provides strong evidence that our proposed PTP-RN and PTP-ADV training methods have a much smoother loss landscape compared to the vanilla prompt tuning. For the few-shot learning setting, we demonstrate the results in Table 6. It could be seen that all our PTP methods have smaller variances than the baseline




Figure 3: The variance of the scores on the dev sets from SuperGLUE Benchmark. We compute it on 5 runs with different random seeds to report. The reported experiments are all using BERT-large models as backbones.

method. Specifically, PTP+PGD has the smallest variance in 5 runs, which indicates its training and generalization stability. Compared with the PTP-RN methods (RG, RM), PTP-ADV methods (A2T, PGD) achieve smaller variance. We also conduct the experiments under fully-supervised learning settings in Figure 3. It shows that in all 8 tasks from SuperGLUE benchmark, our proposed method can still reduce the training variance of different runs with the same hyperparameter except the seeds.

## 5 Ablation Study

In this section, we conduct comprehensive ablation studies of our perturbation methods under both fully-supervised and few-shot learning settings.

**RG Perturbation.** We investigate the strength of gaussian noise  $\delta$  and number of perturbed embeddings  $i$  affects the performance (see Eq. (6)). Specifically, we run experiments with  $\sigma$ , which is the variance of the added Gaussian noise, from  $\{10^{-4}, 10^{-3}, 10^{-2}\}$  and the number of word embeddings perturbed, which is denoted as  $i$  in Eq. (6), from  $\{1, 5, 10, 20\}$ . Under the fully-supervised learning setting, we report the results on COPA, RTE, WiC tasks in Figure 4. The results show the appropriate  $\sigma$  is supposed to be  $10^{-2}$  and the number of perturbed embeddings to be 5. With large  $\sigma$  and large  $i$ , PTP+RG is more likely to fail in comparison to the baseline method. Under few-shot learning settings, we select results in MultiRC task to report, as shown in Table 7. The encouraging result also demonstrates that the best choice of  $\sigma$  and number of embeddings perturbed is  $10^{-3}$  and 5, respectively.

**PGD Perturbation.** We investigate how different  $\alpha$  (See Eq. (8)) and PGD iterations affect the performance. We run experiments with  $\alpha$  from  $\{10^{-4}, 10^{-3}, 10^{-2}\}$  and PGD iterations from 1 to 5. Under fully-supervised learning settings, we present the the results of COPA dataset in Table 8. It shows that large  $\alpha$  in PGD will be detrimental to the performance. Under few-shot learning settings, Figure 5 demonstrates the results of different  $\alpha$  and different iterations of PTP+PGD on few-shot settings. In all 3 datasets, when  $\alpha$  is  $10^{-3}$ , not too small nor too large, and PGD iters is 4, PTP+PGD can achieve outstanding performance.

**RM Perturbation.** We investigate how different numbers of random [MASK] inserted affects the PTP. Formally, the number of [MASK] inserted is defined as  $i$  in Eq. (7). We conduct experiments with  $i$  from 1 to 10. Under few-shot learning settings, Figure 6 presents the results of PTP with RM perturbation on FewGLUE (MultiRC and RTE dataset). We observe that RM perturbation can boost the performance substantially in few-shot settings and the best choice of  $i$  is 8. Under fully-supervised settings, Table 9 presents the ablation of RM perturbation. It also shows that a large number of [MASK] inserted harms the performance, especially when the backbone is RoBERTa.

**A2T Perturbation.** We investigate how the minimum cosine similarity between normal input  $s$  and perturbed input  $s'$  of A2T perturbation affects the results. We run experiments with minimum cosine similarity from  $\{0.2, 0.4, 0.6, 0.8\}$  and show results in Table 10. It indicates small similarity may cause damage to the standard performance because the perturbation can be too large in this case.

## 6 Conclusion

In this paper, we first investigated the training instability issues on prompt tuning, which has a precipitous loss landscape in its visualization. To tackle the problem, we proposed PTP-RN and PTP-ADV algorithms, which include four different perturbations (RG, RM, ADV, A2T) on both discrete and continuous spaces, to smooth the loss landscape and make the training stable. Furthermore, our algorithms are also capable of boosting the performance of prompt tuning. The extensive experiments validate the effectiveness of our proposed algorithms on NLU tasks under both fully-supervised and few-shot settings.



## References

- 583
- 584 Anish Athalye, Nicholas Carlini, and David A. Wagner. 2018. [Obfuscated gradients give a false sense](#)
- 585 [of security: Circumventing defenses to adversarial](#)
- 586 [examples](#). In *Proceedings of the 35th International*
- 587 *Conference on Machine Learning, ICML 2018, Stock-*
- 588 *holmsmässan, Stockholm, Sweden, July 10-15, 2018*,
- 589 *volume 80 of Proceedings of Machine Learning Re-*
- 590 *search*, pages 274–283. PMLR.
- 591
- 592 Tom Brown, Benjamin Mann, Nick Ryder, Melanie
- 593 Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind
- 594 Neelakantan, Pranav Shyam, Girish Sastry, Amanda
- 595 Askell, Sandhini Agarwal, Ariel Herbert-Voss,
- 596 Gretchen Krueger, Tom Henighan, Rewon Child,
- 597 Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens
- 598 Winter, Chris Hesse, Mark Chen, Eric Sigler, Ma-
- 599 teusz Litwin, Scott Gray, Benjamin Chess, Jack
- 600 Clark, Christopher Berner, Sam McCandlish, Alec
- 601 Radford, Ilya Sutskever, and Dario Amodei. 2020.
- 602 [Language models are few-shot learners](#). In *Ad-*
- 603 *vances in Neural Information Processing Systems*,
- 604 *volume 33*, pages 1877–1901.
- 605
- 606 Xiangning Chen and Cho-Jui Hsieh. 2020. [Stabilizing](#)
- 607 [differentiable architecture search via perturbation-](#)
- 608 [based regularization](#). In *Proceedings of the 37th In-*
- 609 *ternational Conference on Machine Learning, ICML*
- 610 *2020, 13-18 July 2020, Virtual Event*, volume 119 of
- 611 *Proceedings of Machine Learning Research*, pages
- 1554–1565. PMLR.
- 612
- 613 Minhao Cheng, Pin-Yu Chen, Sijia Liu, Shiyu Chang,
- 614 Cho-Jui Hsieh, and Payel Das. 2021. [Self-](#)
- 615 [progressing robust training](#). In *Thirty-Fifth AAAI*
- 616 *Conference on Artificial Intelligence, AAAI 2021,*
- 617 *Thirty-Third Conference on Innovative Applications*
- 618 *of Artificial Intelligence, IAAI 2021, The Eleventh*
- 619 *Symposium on Educational Advances in Artificial In-*
- 620 *telligence, EAAI 2021, Virtual Event, February 2-9,*
- 621 *2021*, pages 7107–7115. AAAI Press.
- 622
- 623 Christopher Clark, Kenton Lee, Ming-Wei Chang,
- 624 Tom Kwiatkowski, Michael Collins, and Kristina
- 625 Toutanova. 2019. [Boolq: Exploring the surprising](#)
- 626 [difficulty of natural yes/no questions](#). *arXiv preprint*
- 627 *arXiv:1905.10044*.
- 628
- 629 Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. 2019.
- 630 [Certified adversarial robustness via randomized](#)
- 631 [smoothing](#). In *Proceedings of the 36th International*
- 632 *Conference on Machine Learning*, volume 97 of *Pro-*
- 633 *ceedings of Machine Learning Research*, pages 1310–
- 634 1320. PMLR.
- 635
- 636 Ido Dagan, Oren Glickman, and Bernardo Magnini.
- 637 2005. The pascal recognising textual entailment chal-
- 638 lenge. In *Machine learning challenges workshop*,
- 639 pages 177–190. Springer.
- 640
- 641 Joe Davison, Joshua Feldman, and Alexander M. Rush.
- 642 2019. [Commonsense knowledge mining from pre-](#)
- 643 [trained models](#). In *Proceedings of the 2019 Confer-*
- 644 *ence on Empirical Methods in Natural Language Pro-*
- 645 *cessing and the 9th International Joint Conference*
- 646 *on Natural Language Processing, EMNLP-IJCNLP*
- 647 *2019, Hong Kong, China, November 3-7, 2019*, pages
- 648 1173–1178. Association for Computational Linguistics.
- 649
- 650 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and
- 651 Kristina Toutanova. 2019. [BERT: Pre-training of](#)
- 652 [deep bidirectional transformers for language under-](#)
- 653 [standing](#). In *Proceedings of the 2019 Conference of*
- 654 *the North American Chapter of the Association for*
- 655 *Computational Linguistics: Human Language Tech-*
- 656 *nologies, Volume 1 (Long and Short Papers)*, pages
- 657 4171–4186, Minneapolis, Minnesota. Association for
- 658 Computational Linguistics.
- 659
- 660 Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing
- 661 Dou. 2018. [HotFlip: White-box adversarial exam-](#)
- 662 [ples for text classification](#). In *Proceedings of the 56th*
- 663 *Annual Meeting of the Association for Computational*
- 664 *Linguistics (Volume 2: Short Papers)*, pages 31–36,
- 665 Melbourne, Australia. Association for Computational
- 666 Linguistics.
- 667
- 668 Zhe Gan, Yen-Chun Chen, Linjie Li, Chen Zhu,
- 669 Yu Cheng, and Jingjing Liu. 2020. [Large-scale adver-](#)
- 670 [sarial training for vision-and-language representation](#)
- 671 [learning](#). In *Advances in Neural Information Process-*
- 672 *ing Systems*, volume 33, pages 6616–6628. Curran
- 673 Associates, Inc.
- 674
- 675 Ian J. Goodfellow, Jonathon Shlens, and Christian
- 676 Szegedy. 2015. [Explaining and harnessing adver-](#)
- 677 [sarial examples](#). In *3rd International Conference on*
- 678 *Learning Representations, ICLR 2015*.
- 679
- 680 Andrew Gordon, Zornitsa Kozareva, and Melissa Roeme-
- 681 mele. 2012. Semeval-2012 task 7: Choice of plau-
- 682 sible alternatives: An evaluation of commonsense
- 683 causal reasoning. In *\* SEM 2012: The First Joint*
- 684 *Conference on Lexical and Computational Semantics-*
- 685 *Volume 1: Proceedings of the main conference and*
- 686 *the shared task, and Volume 2: Proceedings of the*
- 687 *Sixth International Workshop on Semantic Evalua-*
- 688 *tion (SemEval 2012)*, pages 394–398.
- 689
- 690 Adi Haviv, Jonathan Berant, and Amir Globerson. 2021.
- 691 [Bertese: Learning to speak to BERT](#). In *Proceedings*
- 692 *of the 16th Conference of the European Chapter of*
- 693 *the Association for Computational Linguistics: Main*
- 694 *Volume, EACL 2021*, pages 3618–3623. Association
- 695 for Computational Linguistics.
- 696
- 697 Pengcheng He, Xiaodong Liu, Jianfeng Gao, and
- 698 Weizhu Chen. 2021. [Deberta: decoding-enhanced](#)
- 699 [bert with disentangled attention](#). In *9th International*
- 700 *Conference on Learning Representations, ICLR 2021,*
- 701 *Virtual Event, Austria, May 3-7, 2021*. OpenRe-
- 702 view.net.
- 703
- 704 Zhengbao Jiang, Frank F. Xu, Jun Araki, and Graham
- 705 Neubig. 2020. [How can we know what language](#)
- 706 [models know](#). *Trans. Assoc. Comput. Linguistics*,
- 707 8:423–438.

|     |  |     |
|-----|--|-----|
| 696 | Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2019. <a href="#">Is BERT really robust? natural language attack on text classification and entailment</a> . <i>CoRR</i> , abs/1907.11932.   |     |
| 697 |  |     |
| 698 |  |     |
| 699 |  |     |
| 700 | Daniel Khashabi, Snigdha Chaturvedi, Michael Roth, Shyam Upadhyay, and Dan Roth. 2018. Looking beyond the surface: A challenge set for reading comprehension over multiple sentences. In <i>Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)</i> , pages 252–262.  |     |
| 701 |  |     |
| 702 |  |     |
| 703 |  |     |
| 704 |  |     |
| 705 |  |     |
| 706 |  |     |
| 707 |  |     |
| 708 | Kezhi Kong, Guohao Li, Mucong Ding, Zuxuan Wu, Chen Zhu, Bernard Ghanem, Gavin Taylor, and Tom Goldstein. 2022. <a href="#">Robust optimization as data augmentation for large-scale graphs</a> . In <i>IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022</i> , pages 60–69. IEEE.   |     |
| 709 |  |     |
| 710 |  |     |
| 711 |  |     |
| 712 |  |     |
| 713 |  |     |
| 714 |  |     |
| 715 | Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. <a href="#">ALBERT: A lite BERT for self-supervised learning of language representations</a> . In <i>8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020</i> . OpenReview.net.   |     |
| 716 |  |     |
| 717 |  |     |
| 718 |  |     |
| 719 |  |     |
| 720 |  |     |
| 721 |  |     |
| 722 | Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. <a href="#">The power of scale for parameter-efficient prompt tuning</a> . In <i>Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing</i> , pages 3045–3059, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.  |     |
| 723 |  |     |
| 724 |  |     |
| 725 |  |     |
| 726 |  |     |
| 727 |  |     |
| 728 |  |     |
| 729 | Hector Levesque, Ernest Davis, and Leora Morgenstern. 2012. The winograd schema challenge. In <i>Thirteenth international conference on the principles of knowledge representation and reasoning</i> .   |     |
| 730 |  |     |
| 731 |  |     |
| 732 |  |     |
| 733 | Linyang Li, Ruotian Ma, Qipeng Guo, Xiangyang Xue, and Xipeng Qiu. 2020. <a href="#">BERT-ATTACK: Adversarial attack against BERT using BERT</a> . In <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)</i> , pages 6193–6202, Online. Association for Computational Linguistics.  |     |
| 734 |  |     |
| 735 |  |     |
| 736 |  |     |
| 737 |  |     |
| 738 |  |     |
| 739 |  |     |
| 740 | Xiang Lisa Li and Percy Liang. 2021. <a href="#">Prefix-tuning: Optimizing continuous prompts for generation</a> . In <i>Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021</i> , pages 4582–4597. Association for Computational Linguistics. |     |
| 741 |  |     |
| 742 |  |     |
| 743 |  |     |
| 744 |  |     |
| 745 |  |     |
| 746 |  |     |
| 747 |  |     |
| 748 | Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021a. <a href="#">Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing</a> . <i>CoRR</i> , abs/2107.13586.  |     |
| 749 |  |     |
| 750 |  |     |
| 751 |  |     |
| 752 |  |     |
|     | Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Tam, Zhengxiao Du, Zhilin Yang, and Jie Tang. 2022. <a href="#">P-tuning: Prompt tuning can be comparable to fine-tuning across scales and tasks</a> . In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)</i> , pages 61–68, Dublin, Ireland. Association for Computational Linguistics.                                      | 753 |
|     |  | 754 |
|     |  | 755 |
|     |  | 756 |
|     |  | 757 |
|     |  | 758 |
|     |  | 759 |
|     |  | 760 |
|     | Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2021b. <a href="#">GPT understands, too</a> . <i>CoRR</i> , abs/2103.10385.   | 761 |
|     |  | 762 |
|     |  | 763 |
|     | Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. <a href="#">Roberta: A robustly optimized BERT pretraining approach</a> . <i>CoRR</i> , abs/1907.11692.  | 764 |
|     |  | 765 |
|     |  | 766 |
|     |  | 767 |
|     |  | 768 |
|     | Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2018. <a href="#">Towards deep learning models resistant to adversarial attacks</a> . In <i>6th International Conference on Learning Representations, ICLR 2018</i> . OpenReview.net.  | 769 |
|     |  | 770 |
|     |  | 771 |
|     |  | 772 |
|     |  | 773 |
|     | John Morris, Jin Yong Yoo, and Yanjun Qi. 2020. <a href="#">TextAttack: Lessons learned in designing python frameworks for NLP</a> . In <i>Proceedings of Second Workshop for NLP Open Source Software (NLP-OSS)</i> , pages 126–131. Association for Computational Linguistics.   | 774 |
|     |  | 775 |
|     |  | 776 |
|     |  | 777 |
|     |  | 778 |
|     | Mohammad Taher Pilehvar and Jose Camacho-Collados. 2018. <a href="#">Wic: the word-in-context dataset for evaluating context-sensitive meaning representations</a> . <i>arXiv preprint arXiv:1808.09121</i> .  | 779 |
|     |  | 780 |
|     |  | 781 |
|     |  | 782 |
|     | Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. <a href="#">Exploring the limits of transfer learning with a unified text-to-text transformer</a> . <i>Journal of Machine Learning Research (JMLR)</i> .   | 783 |
|     |  | 784 |
|     |  | 785 |
|     |  | 786 |
|     |  | 787 |
|     |  | 788 |
|     | Shuhuai Ren, Yihe Deng, Kun He, and Wanxiang Che. 2019. <a href="#">Generating natural language adversarial examples through probability weighted word saliency</a> . In <i>Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics</i> , pages 1085–1097, Florence, Italy. Association for Computational Linguistics.   | 789 |
|     |  | 790 |
|     |  | 791 |
|     |  | 792 |
|     |  | 793 |
|     |  | 794 |
|     |  | 795 |
|     | Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. <a href="#">Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter</a> . <i>CoRR</i> , abs/1910.01108.   | 796 |
|     |  | 797 |
|     |  | 798 |
|     |  | 799 |
|     | Timo Schick and Hinrich Schütze. 2021a. <a href="#">Exploiting cloze-questions for few-shot text classification and natural language inference</a> . In <i>Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EAACL 2021</i> , pages 255–269. Association for Computational Linguistics.  | 800 |
|     |  | 801 |
|     |  | 802 |
|     |  | 803 |
|     |  | 804 |
|     |  | 805 |
|     |  | 806 |

807 Timo Schick and Hinrich Schütze. 2021b. **It’s not just**  
 808 **size that matters: Small language models are also few-**  
 809 **shot learners.** In *Proceedings of the 2021 Conference*  
 810 *of the North American Chapter of the Association*  
 811 *for Computational Linguistics: Human Language*  
 812 *Technologies, NAACL-HLT 2021, Online, June 6-11,*  
 813 *2021*, pages 2339–2352. Association for Computa-  
 814 *tional Linguistics.*

815 Mohammad Shoeybi, Mostofa Patwary, Raul Puri,  
 816 Patrick LeGresley, Jared Casper, and Bryan Catan-  
 817 zaro. 2019. **Megatron-lm: Training multi-billion**  
 818 **parameter language models using model parallelism.**  
 819 *CoRR*, abs/1909.08053.

820 Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian J.  
 821 Goodfellow, Dan Boneh, and Patrick D. McDaniel.  
 822 2018. **Ensemble adversarial training: Attacks and de-**  
 823 **fenses.** In *6th International Conference on Learning*  
 824 *Representations, ICLR 2018, Vancouver, BC, Canada,*  
 825 *April 30 - May 3, 2018, Conference Track Proceed-*  
 826 *ings.* OpenReview.net.

827 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob  
 828 Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz  
 829 Kaiser, and Illia Polosukhin. 2017. **Attention is all**  
 830 **you need.** *CoRR*, abs/1706.03762.

831 Alex Wang, Yada Pruksachatkun, Nikita Nangia, Aman-  
 832 preet Singh, Julian Michael, Felix Hill, Omer Levy,  
 833 and Samuel Bowman. 2019. **Superglue: A stickier**  
 834 **benchmark for general-purpose language understand-**  
 835 **ing systems.** In *NeurIPS 2019*.

836 Cihang Xie, Mingxing Tan, Boqing Gong, Jiang Wang,  
 837 Alan L. Yuille, and Quoc V. Le. 2020. **Adversarial**  
 838 **examples improve image recognition.** In *2020*  
 839 *IEEE/CVF Conference on Computer Vision and Pat-*  
 840 *tern Recognition, CVPR 2020, Seattle, WA, USA,*  
 841 *June 13-19, 2020*, pages 816–825. Computer Vision  
 842 *Foundation / IEEE.*

843 Jin Yong Yoo, John Morris, Eli Lifland, and Yanjun Qi.  
 844 2020. **Searching for a search method: Benchmarking**  
 845 **search algorithms for generating NLP adversarial ex-**  
 846 **amples.** In *Proceedings of the Third BlackboxNLP*  
 847 *Workshop on Analyzing and Interpreting Neural Net-*  
 848 *works for NLP*, pages 323–332, Online. Association  
 849 *for Computational Linguistics.*

850 Jin Yong Yoo and Yanjun Qi. 2021. **Towards improving**  
 851 **adversarial training of NLP models.** In *Findings*  
 852 *of the Association for Computational Linguistics:*  
 853 *EMNLP 2021, Virtual Event / Punta Cana, Domini-*  
 854 *can Republic, 16-20 November, 2021*, pages 945–956.  
 855 Association for Computational Linguistics.

856 Chen Zhu, Yu Cheng, Zhe Gan, Siqi Sun, Tom Gold-  
 857 stein, and Jingjing Liu. 2020. **FreeLb: Enhanced**  
 858 **adversarial training for natural language understand-**  
 859 **ing.** In *8th International Conference on Learning*  
 860 *Representations, ICLR 2020, Addis Ababa, Ethiopia,*  
 861 *April 26-30, 2020.*

## APPENDIX

### A Details of learning rate and prompt length

| Tasks   | LR1  | LR2  | PL1 | PL2 |
|---------|------|------|-----|-----|
| BoolQ   | 1e-3 | 5e-3 | 40  | 16  |
| COPA    | 5e-3 | 7e-3 | 16  | 16  |
| RTE     | 1e-2 | 5e-3 | 20  | 128 |
| WSC     | 3e-3 | 7e-3 | 16  | 8   |
| CB      | 7e-3 | 9e-3 | 16  | 16  |
| MultiRC | 1e-4 | 3e-3 | 40  | 20  |
| ReCoRD  | 3e-4 | 5e-3 | 16  | 40  |
| WiC     | 1e-4 | 5e-3 | 20  | 16  |

Table 4: Prompt Length and Learning Rate details for 8 tasks on SuperGLUE. LR1, PL1: learning rate and prompt length for continuous prompts with BERT-large backbone. LR2 and PL2: learning rate and prompt length for prompts with RoBERTa-large backbone.

| Tasks   | Learning Rate |
|---------|---------------|
| BoolQ   | 5e-5          |
| RTE     | 5e-5          |
| WiC     | 1e-5          |
| WSC     | 5e-5          |
| COPA    | 1e-5          |
| MultiRC | 1e-4          |
| CB      | 1e-5          |

Table 5: Prompts’ Learning Rate details for 7 tasks in FewGLUE.

865 Under fully-supervised settings, the prompt  
 866 length and learning rate details are presented as  
 867 Table 4. Under few-shot learning settings, we re-  
 868 port it as Table 5. The prompt length is exactly the  
 869 same as the PT (Liu et al., 2021b), thus we ignore  
 870 it here.

### B Supplement for Experiment

871 This section includes the tables and figures as a  
 872 supplement to our experiment and ablations. Ta-  
 873 ble 6 demonstrates the comparison of the variance  
 874 of different training methods on FewGLUE bench-  
 875 mark. Table 7 presents the ablation of our algo-  
 876 rithm with RG perturbation on MultiRC (Khashabi  
 877 et al., 2018) task. We show the ablation of PGD  
 878 perturbation in Table 8. The ablation of RM per-  
 879 turbation is presented as Table 9. Table 10 shows  
 880




Figure 4: Performance of PTP+RG on SuperGLUE (WiC, RTE, COPA datasets) with  $\sigma$  from  $\{10^{-4}, 10^{-3}, 10^{-2}\}$  and perturbed embeddings from  $\{1, 5, 10, 20\}$ . The dashed red line represents the performance of baseline method PT2 with BERT-large as backbone LM.




Figure 5: Performance of PTP+PGD on FewGLUE (WiC, BoolQ, WSC datasets) with different  $\alpha$  and PGD iterations. The dashed red line represents the performance of the baseline method PT (Liu et al., 2021b). It shows the best  $\alpha$  and PGD iterations are  $10^{-3}$  and 4, respectively.

| Tasks | RTE                          | WSC                          | WiC                          | BoolQ                        |
|-------|------------------------------|------------------------------|------------------------------|------------------------------|
| PT    | 1.89                         | 1.68                         | 1.77                         | 1.45                         |
| +RG   | 0.81                         | 0.78                         | 0.91                         | 0.56                         |
| +A2T  | 0.45                         | 0.43                         | <b>0.39</b> <sub>↓1.38</sub> | 0.59                         |
| +RM   | 0.68                         | 0.95                         | 0.87                         | 0.65                         |
| +PGD  | <b>0.35</b> <sub>↓1.54</sub> | <b>0.31</b> <sub>↓1.37</sub> | 0.47                         | <b>0.43</b> <sub>↓1.02</sub> |

Table 6: The variance of the scores on the dev sets of RTE, COPA and BoolQ from the FewGLUE benchmark. We compute it on 5 runs with different random seeds (other hyper-parameter are the same). We employ bold font to denote the smallest deviation in each task and blue font to denote the decrease when compared to PT.

881 the ablation of our proposed PTP+A2T training  
882 algorithm.

883 Figure 4 shows the ablation of RG perturbations  
884 on WiC (Pilehvar and Camacho-Collados, 2018),  
885 RTE (Dagan et al., 2005), and COPA (Gordon  
886 et al., 2012) datasets under fully-supervised learn-  
887 ing setting while Figure 5 presents the ablation of  
888 our PTP+PGD training method on WiC (Pilehvar  
889 and Camacho-Collados, 2018), BoolQ (Clark et al.,  
890 2019) and WSC (Levesque et al., 2012) datasets  
891 under few-shot learning settings. We show the  
892 ablation of our proposed PTP+RM algorithm as  
893 Figure 6 on FewGLUE benchmark.

894 All experiments are conducted on servers with  
895 RTX A6000 GPUs, each having 48GB of memory.



| MultiRC       | E=1  | E=5                        | E=10 | E=20 |
|---------------|------|----------------------------|------|------|
| $\sigma=1e-2$ | 75.2 | 74.3 $\downarrow$ 0.7      | 75.8 | 74.7 |
| $\sigma=1e-3$ | 77.3 | <b>78.1</b> $\uparrow$ 3.1 | 77.1 | 76.6 |
| $\sigma=1e-4$ | 76.9 | 77.2                       | 76.8 | 75.5 |

Table 7: Results of different  $\sigma$  in RG perturbation (MultiRC task). The baseline method is PT2 and its performance is 75.0. E denotes number of embeddings that are perturbed. We mark the best and the worst.

| COPA          | t=1  | 2    | 3                          | 4    | 5                     |
|---------------|------|------|----------------------------|------|-----------------------|
| $\alpha=1e-2$ | 70.3 | 72.1 | 72.0                       | 71.0 | 69.0 $\downarrow$ 4.0 |
| $1e-3$        | 73.4 | 70.8 | 72.9                       | 73.5 | 73.2                  |
| $1e-4$        | 73.1 | 73.4 | <b>74.7</b> $\uparrow$ 1.7 | 73.8 | 72.5                  |

Table 8: Results of PTP+PGD on fully-supervised COPA dataset. The baseline method is PT2 (Liu et al., 2022), whose accuracy is 73.0. The backbone LM employed is BERT-large.  $\alpha$  is the perturbation size while  $t$  is PGD iterations (see Eq. (8)).

| Dataset      | LM      | 1     | 2     | 3            | 4     | 5     | 6     | 7     | 8     | 9     | 10    |
|--------------|---------|-------|-------|--------------|-------|-------|-------|-------|-------|-------|-------|
| BoolQ (Full) | BERT    | -0.12 | +0.50 | <b>+1.57</b> | +0.65 | +1.35 | +1.14 | +1.12 | +0.24 | +0.25 | -0.38 |
|              | RoBERTa | +0.31 | +0.76 | <b>+1.10</b> | +0.88 | +0.69 | -0.36 | +0.19 | -0.43 | -0.67 | -0.36 |

Table 9: The results of different  $i$  in Eq. (7) (number of [MASK] randomly inserted into input sequence as perturbation). We select BoolQ with fully-supervised settings to report. The reported increase or decrease is compared to the baseline method PT2.




Figure 6: The results of different  $i$  in Eq. (7) (number of [MASK] randomly inserted into input sequence as perturbation). We select MultiRC and RTE tasks with Few-shot setting to report.

| Cosine Sim    | 0.2   | 0.4   | 0.6   | 0.8   |
|---------------|-------|-------|-------|-------|
| BoolQ(Full)   | -0.83 | -0.46 | +0.61 | +0.52 |
| BoolQ(Few)    | -0.96 | -0.13 | +0.78 | +0.84 |
| MultiRC(Full) | -0.77 | -0.35 | +0.25 | +0.71 |
| MultiRC(Few)  | -0.59 | -0.54 | +1.03 | +0.68 |

Table 10: The results of different minimum cosine similarity in A2T perturbation. Full: fully-supervised learning setting. Few: few-shot learning setting.