

# UNIPELT: A Unified Framework for Parameter-Efficient Language Model Tuning

Yuning Mao<sup>1\*</sup>, Lambert Mathias<sup>2</sup>, Rui Hou<sup>2</sup>, Amjad Almahairi<sup>2</sup>,  
Hao Ma<sup>2</sup>, Jiawei Han<sup>1</sup>, Wen-tau Yih<sup>2</sup>, Madian Khabsa<sup>2</sup>

<sup>1</sup>University of Illinois Urbana-Champaign {yuningm2, hanj}@illinois.edu

<sup>2</sup>Meta AI {mathiasl, rayhou, aalmah, haom, scottyih, mkhabasa}@fb.com

## Abstract

Recent parameter-efficient language model tuning (PELT) methods manage to match the performance of fine-tuning with much fewer trainable parameters and perform especially well when training data is limited. However, different PELT methods may perform rather differently on the same task, making it non-trivial to select the most appropriate method for a specific task, especially considering the fast-growing number of new PELT methods and tasks. In light of model diversity and the difficulty of model selection, we propose a unified framework, UNIPELT, which incorporates different PELT methods as submodules and learns to activate the ones that best suit the current data or task setup via gating mechanism. On the GLUE benchmark, UNIPELT consistently achieves 1~4% gains compared to the best individual PELT method that it incorporates and outperforms fine-tuning under different setups. Moreover, UNIPELT generally surpasses the upper bound that takes the best performance of all its submodules used individually on each task, indicating that a mixture of multiple PELT methods may be inherently more effective than single methods.<sup>1</sup>

## 1 Introduction

As pre-trained language models (PLMs) (Devlin et al., 2019) grow larger and larger (Brown et al., 2020), it becomes increasingly infeasible to perform conventional fine-tuning, where separate replicas of the model parameters are modified per single task. To solve the issue, there has recently been a surge of studies on parameter-efficient language model tuning (PELT), namely how to effectively tune the PLMs with fewer trainable parameters.

Existing PELT research generally aims at achieving performance comparable to fine-tuning with

\*Work was done during internship at Meta AI.

<sup>1</sup>Our code can be found at <https://github.com/morningmoni/UniPELT>.

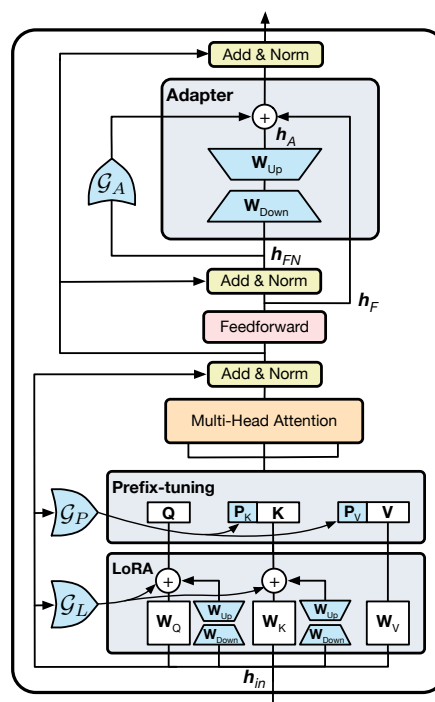


Figure 1: Illustration of UNIPELT, which subsumes existing PELT methods as submodules and controls them via gating mechanism  $\mathcal{G}$ . Different (combinations of) submodules can be activated for different samples. The trainable parameters are shown in blue.

as few trainable parameters as possible, which has seen significant progress – the task-specific trainable parameters used in most recent approaches (Lester et al., 2021; Guo et al., 2021) are almost negligible compared to the total parameters of the PLM (<1%). A more challenging yet less studied problem is whether one can achieve better performance than fine-tuning with fewer parameters. Recent studies (He et al., 2021; Li and Liang, 2021; Karimi Mahabadi et al., 2021b) find that some PELT methods are more effective than fine-tuning on certain tasks when training data is limited, possibly due to the reduced risk of overfitting. However, as found in our experiments (Table 1), different PELT methods exhibit diverse characteristics and perform rather differently on the same task, which

makes it nontrivial to select the most appropriate method for a specific task, especially considering the fast-growing number of new PELT methods and tasks (Ding and Hu, 2021).

In light of the diverse performance of PELT methods and the cost of selecting the best method, we propose a unified PELT framework, named UNIPELT, which incorporates different PELT methods as submodules and learns to dynamically activate the (combination of) submodules that best suit the current data or task setup. As a result, model selection is no longer needed and consistently better performance is achieved under different setups. The activation of each submodule in UNIPELT is controlled by *gating mechanism*, which learns to favor (assign more weight to) the submodules that positively contribute to a given task. In addition, since the number of parameters introduced by each submodule is generally small, combining multiple methods leads to negligible losses in model efficiency.

We select four representative PELT methods for our study – adapter (Houlsby et al., 2019), prefix-tuning (Li and Liang, 2021), LoRA (Hu et al., 2021), and BitFit (Ben Zaken et al., 2021), which largely cover the major categories of PELT methods. We perform two sets of analysis that carefully examines (i) the characteristics of individual PELT methods and (ii) their effectiveness when coordinated by UNIPELT under various setups.<sup>2</sup>

Extensive experiments on the GLUE benchmark (Wang et al., 2019), with 32 setups (8 tasks  $\times$  4 data sizes) and 1,000+ runs, not only reveal the diverse behavior of PELT methods, but also show that UNIPELT is more effective and robust than using each method alone in various task and data setups. Specifically, UNIPELT consistently improves the best submodule that it incorporates by 1–4 points and even outperforms fine-tuning, achieving the best average performance on the GLUE benchmark under different setups. Moreover, UNIPELT generally surpasses the upper bound that takes the best performance of all its submodules used individually on each task, which suggests that UNIPELT maintains (near) optimal performance under different setups. The fact that UNIPELT outperforms the upper bound also implies that a mixture of PELT methods involving different parts of the PLM architecture may be inherently more effective than

individual methods.

**Contributions.** (1) We conduct a comprehensive study of representative PELT methods and carefully examine their differences and commonalities in terms of performance and characteristics. (2) We propose a unified PELT framework that can incorporate existing methods as submodules and automatically learn to activate the appropriate submodules for a given task. (3) Our proposed framework achieves better average performance than fine-tuning and the PELT methods that it incorporates under various setups, often performing the best and never the worst at per-task level, exhibiting superior effectiveness and robustness with negligible losses in model efficiency.

## 2 Preliminaries

### 2.1 PELT Methods without Additional Parameters

PLMs can be used as feature extractors where only the top layers or prediction head are fine-tuned without additional parameters (Lee et al., 2019). However, such fine-tuning approaches generally lead to degenerate model performance that is much worse than fine-tuning all parameters (Lee et al., 2019; Pfeiffer et al., 2021). A recent method BitFit (Ben Zaken et al., 2021) only tunes the bias terms of the PLM and is shown to achieve performance comparable to fine-tuning on certain tasks when training data is limited. Therefore, we select BitFit as the representative of this category for analysis.

### 2.2 PELT Methods with Additional Parameters

Alternatively, one may fix the entire PLM and introduce a small number of new trainable parameters. Notable examples in this category include adapter (Houlsby et al., 2019) and its extensions (Pfeiffer et al., 2021; Karimi Mahabadi et al., 2021b), prefix-tuning (Li and Liang, 2021) and its extensions (Lester et al., 2021), and additive methods (Guo et al., 2021; Hu et al., 2021).

Next, we will briefly describe these methods to facilitate the introduction of our proposed framework. An illustration is shown in Fig. 1 for better understanding.

**Adapter.** Adapter (Houlsby et al., 2019) adds a trainable *bottleneck layer* after the feedforward network in each Transformer layer of the PLM. A bottleneck layer consists of a down+up projection pair that shrinks and recovers the size of token hidden

<sup>2</sup>BitFit is not included in UNIPELT as it typically performs the worst in our preliminary experiments.

states. Mathematically, if we denote the output of the feedforward network after residual connection and layer normalization as  $\mathbf{h}_{FN}$  with hidden size  $D_{\text{hidden}}$  and bottleneck size  $D_{\text{mid}}$ , then the output of a bottleneck layer  $\mathbf{h}_A$  is:

$$\mathbf{h}_A = \mathbf{W}_{\text{up}}^T \phi(\mathbf{W}_{\text{down}}^T \mathbf{h}_{FN}), \quad (1)$$

where  $\mathbf{W}_{\text{down}} \in \mathbb{R}^{D_{\text{hidden}} \times D_{\text{mid}}}$ ,  $\mathbf{W}_{\text{up}} \in \mathbb{R}^{D_{\text{mid}} \times D_{\text{hidden}}}$ ,  $\phi$  is a nonlinear activation function, and the bias terms are omitted for brevity. The parameters in layer normalization and the final prediction head sometimes are also fine-tuned depending on the specific adapter variants.

Adapter has shown to be on par with fine-tuning and sometimes exhibits better effectiveness in the low-resource setting (He et al., 2021). Later studies extend adapter to multi-lingual (Pfeiffer et al., 2020b) and multi-task (Karimi Mahabadi et al., 2021b) settings, or further reduce its trainable parameters (Karimi Mahabadi et al., 2021a), which can be easily incorporated into UNIPELT as a replacement of the vanilla adapter.

**Prefix-tuning.** Prefix-tuning (Li and Liang, 2021) prepends a number of task-specific trainable vectors to the input of multi-head attention in *each* Transformer layer, which the original tokens can attend to as if they were virtual tokens. Specifically, we denote the original sequence length  $L_0$ , the number of trainable vectors (*i.e.*, prefix length)  $L$ , and the Transformer layer input  $\mathbf{h}_{\text{in}} \in \mathbb{R}^{D_{\text{hidden}} \times L_0}$ . First, three linear projections  $\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V \in \mathbb{R}^{D_{\text{hidden}} \times D_{\text{hidden}}}$  transform  $\mathbf{h}_{\text{in}}$  into Query  $\mathbf{Q}$ , Key  $\mathbf{K}$ , and Value  $\mathbf{V}$ . Then, two prefix matrices  $\mathbf{P}_K$  and  $\mathbf{P}_V \in \mathbb{R}^{D_{\text{hidden}} \times L}$  are prepended to  $\mathbf{K}$  and  $\mathbf{V}$ . To stabilize optimization, the prefix matrix  $\mathbf{P}$  is reparameterized by a feedforward network:

$$\mathbf{P}' = \mathbf{W}_{\text{up}}^T \phi(\mathbf{W}_{\text{down}}^T \mathbf{P}), \quad (2)$$

where  $\mathbf{W}_{\text{down}} \in \mathbb{R}^{D_{\text{hidden}} \times D_{\text{mid}}}$ ,  $\mathbf{W}_{\text{up}} \in \mathbb{R}^{D_{\text{mid}} \times 2N_{\text{layer}} D_{\text{hidden}}}$ , and  $N_{\text{layer}}$  denotes the number of Transformer layers. The parameters of this network can be discarded after training, and only  $2N_{\text{layer}}$  prefix matrices  $\in \mathbb{R}^{D_{\text{hidden}} \times L}$  are needed (2 matrices for each layer).

Prefix-tuning is originally evaluated on natural language generation and we adapt it to understanding tasks. A follow-up method named prompt-tuning (Lester et al., 2021) further reduces task-specific parameters by limiting the prefix to the first layer but only performs competitively with

very large model sizes (billions of total parameters), and is thus not considered in our study. Note that prefix-tuning (or prompt-tuning) is different from prompt-based *fine-tuning* methods (Schick and Schütze, 2021; Gao et al., 2021) (see App. A for specific differences).

**Additive Methods.** Additive PELT methods treat the model parameters after fine-tuning as an addition of the pre-trained parameters  $\theta_{\text{pre-trained}}$  and task-specific differences  $\delta_{\text{task}}$ , where  $\theta_{\text{pre-trained}}$  is fixed and a new (sub)set of model parameters are added on top:  $\theta_{\text{task}} = \theta_{\text{pre-trained}} + \delta_{\text{task}}$ . There are various ways to parameterize  $\delta_{\text{task}}$ , leading to different additive methods such as LoRA (Hu et al., 2021), diff pruning (Guo et al., 2021), and side-tuning (Zhang et al., 2020). We take LoRA as a representative and incorporate it into UNIPELT. Other methods are conceptually similar and can be incorporated in the same fashion.

LoRA introduces trainable low-rank matrices and combines them with the original matrices in the multi-head attention. Specifically, two matrices  $\mathbf{W}_{\text{down}} \in \mathbb{R}^{D_{\text{hidden}} \times D_{\text{mid}}}$  and  $\mathbf{W}_{\text{up}} \in \mathbb{R}^{D_{\text{mid}} \times D_{\text{hidden}}}$  are added for the query and key projections along with the original matrix  $\mathbf{W}_Q$  and  $\mathbf{W}_K \in \mathbb{R}^{D_{\text{hidden}} \times D_{\text{hidden}}}$ :

$$\mathbf{Q} = (\mathbf{W}_Q^T + \alpha \mathbf{W}_{\text{up}}^T \mathbf{W}_{\text{down}}^T) \mathbf{h}_{\text{in}}, \quad (3)$$

where  $\alpha$  is a fixed scalar hyperparameter for scaling the task-specific differences. The form of the trainable matrices in LoRA is quite similar to those in adapter or prefix-tuning, but there is no activation function  $\phi$  in between.

## 3 Unifying PELT Methods

### 3.1 Task Formulation

Given a large PLM  $\mathcal{M}$  with size  $|\mathcal{M}|$  that cannot be fine-tuned directly due to computational or storage cost, suppose that we have a list of PELT methods  $\{m_i\}$ , the trainable parameters of which are negligible (*i.e.*,  $\sum_i |m_i| \ll |\mathcal{M}|$ ), our goal is to design a unified PELT framework that incorporates  $\{m_i\}$  as submodules and learns to dynamically activate (up-weight) different submodules when appropriate under different scenarios, such that one could achieve satisfactory results in terms of both model effectiveness and robustness without the hassle of permuting all the method  $\times$  task  $\times$  data combinations.

### 3.2 Proposed Method

**Motivation & Intuition.** During the analysis of individual PELT methods, we observe that different PELT methods exhibit diverse characteristics and perform rather differently on the same task. For example, prefix-tuning generally performs well on natural language inference tasks regardless of the size of training data. Also, as can be seen in Fig. 1 and Sec. 2, different PELT methods often involve different parts of the PLM architecture (e.g., before multi-head attention for prefix-tuning and after feedforward layer for adapter), making it feasible to combine multiple PELT methods without (directly) interfering with each other.

In light of the two observations above, we propose a unified PELT framework, UNIPELT, which takes a hybrid approach by incorporating multiple PELT methods as submodules. At a high level, UNIPELT improves over single PELT methods due to two factors. First, UNIPELT learns to activate (upweight) the submodules that best suit the current task or specific data sample and deactivate (downweight) the rest. Second, we find that UNIPELT generally performs better than taking the best performance of all its submodules used individually on each task, suggesting that there could be some compounding effects that lead to better model effectiveness when multiple PELT methods (that modify different parts of the PLM) are used.

Next, we will introduce how different PELT methods can be incorporated into UNIPELT via gating mechanism.

**Gating Mechanism.** To achieve fine-grained control of submodule (de)activation, we add a trainable gate  $\mathcal{G}_{m_i}$  for each submodule  $m_i \in \{A, P, L\}$  in every Transformer layer (see Fig. 1). The letters A, P, L stand for Adapter, Prefix-tuning, and LoRA, respectively. Intuitively, if  $m_i$  is useful for a given data  $\times$  task setup (or a particular instance), the gate output for  $m_i$  would be higher such that  $m_i$  plays a more important role. The actual interplay of submodules, however, is more complicated given the interdependency of the submodules and the compounding effects of multiple layers.

Specifically, for adapter, there is a residual connection between the feedforward network and the adapter submodule that sums the adapter input (before normalization)  $\mathbf{h}_F$  and output  $\mathbf{h}_A$  as its final output:  $\mathbf{h}'_A = \mathbf{h}_A + \mathbf{h}_F$ . We design a gating function  $\mathcal{G}_A \in (0, 1)$  that estimates the importance of adapter by its direct input  $\mathbf{h}_{FN}$  using a feedforward

network with sigmoid activation and then scales its output:  $\mathbf{h}'_A = \mathcal{G}_A \mathbf{h}_A + \mathbf{h}_F$ . The adapter submodule is effectively bypassed if  $\mathcal{G}_A \approx 0$ .

Similarly, for prefix-tuning, we design a gating function  $\mathcal{G}_P \in (0, 1)$  that is applied to the prefix vectors ( $\mathbf{P}_K$  and  $\mathbf{P}_V$ ) with the representation of the original tokens ( $\mathbf{K}$  and  $\mathbf{V}$ ) intact. In this way, the impact of the prefix would be diminished if the gate output of the prefix-tuning submodule is low.<sup>3</sup> The gating function  $\mathcal{G}_P$  is estimated by the Transformer layer input  $\mathbf{h}_{in}$  with another feedforward network.

As for LoRA, we note that there is already a *constant* scaling factor  $\alpha$  in its original design that resembles the purpose of our gating mechanism. We thus simply make the factor learnable per layer by a third feedforward network that takes  $\mathbf{h}_{in}$  as input instead of specifying a constant manually:  $\theta_{task} = \theta_{pre-trained} + \mathcal{G}_L \delta_{task}$ .

Despite the seeming simplicity of UNIPELT, we note that it is nontrivial for a unified approach to work well under different scenarios. Naively combining different PELT methods as a hybrid approach could lead to mixed or worse performance than using individual methods, as observed in both our experiments and prior studies (Hu et al., 2021).

## 4 Experiments

We conduct extensive experiments with 8 tasks  $\times$  4 data sizes  $\times$  7 methods  $\times$  5 runs per setup, along with additional analysis for particular methods, resulting in 1,000+ runs in total.

### 4.1 Experiment Setup

**Task Setup.** We conduct experiments on the General Language Understanding Evaluation (GLUE) benchmark (Wang et al., 2019), which involves four types of natural language understanding tasks including linguistic acceptability (CoLA), sentiment analysis (SST-2), similarity and paraphrase tasks (MRPC, STS-B, QQP), and natural language inference (MNLI, QNLI, RTE). We exclude the WNLI dataset following prior studies (Houlsby et al., 2019; Devlin et al., 2019).

**Data Setup.** We mainly consider a low-resource setting where training data is limited and the performance of different methods varies much. We sample a small subset of the training set for each task with size  $K = \{100, 500, 1000\}$ . As it is infeasible to submit considerable runs to the GLUE

<sup>3</sup>Prefix-tuning cannot be fully eliminated as adapter or LoRA due to the softmax operation in multi-head attention.

Method	SST-2	MRPC	CoLA	RTE	QNLI	STS-B	MNLI	QQP	Avg.
[ $K = 100$ ] Test Performance									
Fine-tuning	79.61 <sub>4.25</sub>	<u>81.81</u> <sub>0.35</sub>	<u>16.56</u> <sub>4.34</sub>	<u>55.88</u> <sub>1.64</sub>	69.25 <sub>5.94</sub>	74.07 <sub>6.51</sub>	<u>42.56</u> <sub>3.43</sub>	60.41 <sub>6.42</sub>	60.02 <sub>1.84</sub>
BitFit	62.94 <sub>4.85</sub>	81.09 <sub>0.17</sub>	2.71 <sub>1.57</sub>	47.65 <sub>3.56</sub>	42.46 <sub>1.37</sub>	54.53 <sub>0.56</sub>	38.16 <sub>0.53</sub>	59.56 <sub>0.39</sub>	48.64 <sub>0.78</sub>
Adapter	80.48 <sub>2.94</sub>	81.40 <sub>0.19</sub>	2.02 <sub>4.04</sub>	52.78 <sub>0.27</sub>	72.25 <sub>0.49</sub>	77.32 <sub>1.54</sub>	38.81 <sub>3.64</sub>	60.88 <sub>4.00</sub>	58.24 <sub>0.99</sub>
Prefix-tuning	60.87 <sub>12.47</sub>	81.22 <sub>0.00</sub>	0.00 <sub>0.00</sub>	<b>55.96</b> <sub>2.00</sub>	71.91 <sub>2.69</sub>	57.69 <sub>0.02</sub>	40.58 <sub>2.49</sub>	15.68 <sub>0.12</sub>	47.99 <sub>1.77</sub>
→ $L = 50$	79.52 <sub>1.21</sub>	81.22 <sub>0.00</sub>	5.19 <sub>8.62</sub>	49.24 <sub>2.08</sub>	66.33 <sub>2.45</sub>	7.15 <sub>10.37</sub>	33.66 <sub>2.21</sub>	58.32 <sub>3.18</sub>	47.56 <sub>1.37</sub>
LoRA	<u>81.56</u> <sub>0.94</sub>	81.66 <sub>0.81</sub>	13.31 <sub>10.00</sub>	55.02 <sub>1.75</sub>	<b>73.52</b> <sub>1.20</sub>	49.35 <sub>21.87</sub>	39.60 <sub>4.98</sub>	0.09 <sub>0.02</sub>	49.26 <sub>2.19</sub>
UNIPELT (AP)	77.22 <sub>3.75</sub>	<b>81.86</b> <sub>0.70</sub>	14.42 <sub>10.24</sub>	55.52 <sub>2.16</sub>	72.26 <sub>0.89</sub>	<u>79.14</u> <sub>1.97</sub>	<b>42.59</b> <sub>1.20</sub>	<b>63.41</b> <sub>1.44</sub>	<u>60.80</u> <sub>1.53</sub>
UNIPELT (APL)	<b>82.36</b> <sub>0.86</sub>	81.71 <sub>0.72</sub>	<b>23.62</b> <sub>8.83</sub>	55.45 <sub>1.28</sub>	<u>73.19</u> <sub>0.93</sub>	<b>79.37</b> <sub>1.07</sub>	42.30 <sub>1.88</sub>	<u>62.70</u> <sub>2.55</sub>	<b>62.59</b> <sub>1.44</sub>
[ $K = 500$ ] Test Performance									
Fine-tuning	<b>85.67</b> <sub>0.97</sub>	<u>83.34</u> <sub>0.55</sub>	36.47 <sub>2.69</sub>	59.64 <sub>1.10</sub>	77.30 <sub>0.49</sub>	<u>84.96</u> <sub>1.19</sub>	55.84 <sub>0.85</sub>	<u>68.23</u> <sub>1.39</sub>	68.93 <sub>0.65</sub>
BitFit	83.44 <sub>0.63</sub>	82.16 <sub>0.37</sub>	3.32 <sub>2.59</sub>	61.88 <sub>2.75</sub>	69.15 <sub>9.91</sub>	76.30 <sub>0.36</sub>	40.82 <sub>3.30</sub>	65.29 <sub>3.66</sub>	60.30 <sub>1.91</sub>
Adapter	84.54 <sub>1.37</sub>	82.53 <sub>0.36</sub>	38.65 <sub>3.97</sub>	59.35 <sub>3.09</sub>	77.39 <sub>0.84</sub>	83.52 <sub>0.33</sub>	50.04 <sub>1.72</sub>	68.12 <sub>0.95</sub>	68.02 <sub>0.77</sub>
Prefix-tuning	83.65 <sub>0.69</sub>	82.96 <sub>1.63</sub>	38.16 <sub>2.25</sub>	63.18 <sub>2.70</sub>	78.50 <sub>1.12</sub>	79.75 <sub>1.49</sub>	<b>58.06</b> <sub>1.04</sub>	54.34 <sub>25.91</sub>	67.32 <sub>3.42</sub>
LoRA	84.98 <sub>1.10</sub>	82.53 <sub>0.70</sub>	<b>39.86</b> <sub>2.71</sub>	63.03 <sub>2.57</sub>	<b>79.46</b> <sub>0.66</sub>	65.05 <sub>26.31</sub>	<u>56.54</u> <sub>2.05</sub>	55.46 <sub>27.74</sub>	65.86 <sub>4.18</sub>
UNIPELT (AP)	84.84 <sub>0.28</sub>	83.25 <sub>0.51</sub>	<u>39.84</u> <sub>5.01</sub>	<u>63.32</u> <sub>1.72</sub>	78.36 <sub>1.06</sub>	84.53 <sub>0.48</sub>	56.08 <sub>3.26</sub>	68.14 <sub>1.39</sub>	<u>69.79</u> <sub>1.02</sub>
UNIPELT (APL)	84.91 <sub>1.41</sub>	<b>83.56</b> <sub>0.59</sub>	39.81 <sub>2.55</sub>	<b>64.12</b> <sub>2.45</sub>	<u>79.28</u> <sub>0.63</sub>	<b>85.26</b> <sub>0.70</sub>	54.07 <sub>3.74</sub>	<b>68.87</b> <sub>0.41</sub>	<b>69.98</b> <sub>0.42</sub>
[ $K = 1000$ ] Test Performance									
Fine-tuning	<u>86.54</u> <sub>1.01</sub>	84.87 <sub>0.64</sub>	43.26 <sub>2.60</sub>	62.31 <sub>2.10</sub>	79.03 <sub>1.11</sub>	86.39 <sub>0.34</sub>	61.95 <sub>1.20</sub>	<b>71.09</b> <sub>0.77</sub>	71.93 <sub>0.37</sub>
BitFit	83.99 <sub>0.39</sub>	83.95 <sub>0.81</sub>	22.44 <sub>17.10</sub>	62.89 <sub>1.40</sub>	77.43 <sub>0.53</sub>	79.04 <sub>0.61</sub>	52.87 <sub>0.72</sub>	69.50 <sub>0.16</sub>	66.51 <sub>2.22</sub>
Adapter	85.60 <sub>0.63</sub>	84.49 <sub>0.60</sub>	42.33 <sub>1.98</sub>	61.81 <sub>1.57</sub>	79.68 <sub>0.23</sub>	85.52 <sub>0.29</sub>	57.86 <sub>2.44</sub>	70.32 <sub>0.71</sub>	70.95 <sub>0.55</sub>
Prefix-tuning	85.09 <sub>0.99</sub>	83.66 <sub>1.82</sub>	44.07 <sub>2.90</sub>	<b>66.71</b> <sub>2.72</sub>	80.34 <sub>0.70</sub>	82.38 <sub>1.25</sub>	<b>63.59</b> <sub>1.12</sub>	68.58 <sub>0.35</sub>	71.81 <sub>0.52</sub>
LoRA	86.26 <sub>1.22</sub>	<u>86.04</u> <sub>0.99</sub>	<b>45.50</b> <sub>1.11</sub>	<u>65.63</u> <sub>2.11</sub>	<u>81.00</u> <sub>0.98</sub>	81.56 <sub>1.97</sub>	61.32 <sub>1.65</sub>	70.89 <sub>0.81</sub>	72.28 <sub>0.69</sub>
UNIPELT (AP)	86.17 <sub>0.37</sub>	85.86 <sub>1.05</sub>	44.33 <sub>3.55</sub>	64.91 <sub>1.92</sub>	80.65 <sub>0.57</sub>	<u>86.82</u> <sub>0.23</sub>	62.17 <sub>0.99</sub>	69.95 <sub>0.90</sub>	<u>72.61</u> <sub>0.53</sub>
UNIPELT (APL)	<b>87.06</b> <sub>0.81</sub>	<b>86.65</b> <sub>1.10</sub>	<u>45.44</u> <sub>1.97</sub>	65.49 <sub>1.92</sub>	<b>81.22</b> <sub>0.51</sub>	<b>87.10</b> <sub>0.21</sub>	<u>62.49</u> <sub>0.94</sub>	<u>70.99</u> <sub>0.95</sub>	<b>73.31</b> <sub>0.52</sub>

Table 1: Results on the GLUE benchmark with  $K = \{100, 500, 1000\}$  training samples. The evaluation metrics are Matthew’s Correlation for CoLA, F1 for MRPC and QQP, Spearman’s correlation for STS-B, and accuracy for the rest. For MNLI, we evaluate on the matched dataset. We report average performance on five random seeds with standard deviation as the subscript. **Best** and 2nd best methods under each setup are bold and underlined.

leaderboard (2 submissions/day), we take 1,000 samples on the training set as the development set to select the best checkpoint and use the original development set as the test set. To reduce variance, we shuffle the data with 5 random seeds and report the average performance. Additionally, we consider a high-resource setting where the whole training set is used and the best performance on the GLUE development set is reported.

**Compared Methods.** We mainly compare UNIPELT with fine-tuning and four representative PELT methods: adapter (Houlsby et al., 2019), prefix-tuning (Li and Liang, 2021), BitFit (Ben Zaken et al., 2021), and LoRA (Hu et al., 2021). For completeness, we consider two model variants UNIPELT (AP) and UNIPELT (APL), which incorporate 2 and 3 PELT methods, respectively.

**Implementation Details.** We use BERT<sub>base</sub> (Devlin et al., 2019) as the base model in the experiments. Consistent results are observed in our preliminary experiments with BART<sub>large</sub> (Lewis et al., 2020) (provided in App. C). We implement and evaluate all the methods in the same codebase to ensure a fair comparison. We largely follow the

default hyperparameters of different methods and keep them the same on all the tasks for generalizability. We set the prefix length  $L = 10$ , adapter bottleneck size  $D_{\text{mid}} = 48$ , LoRA rank  $D_{\text{mid}} = 8$  if not specified otherwise.<sup>4</sup> More implementation and hyperparameter details can be found in App. B.

## 4.2 Analysis of Individual PELT Methods

In Table 1, we show the performance of different methods on the GLUE benchmark with various sizes of training data. The results on the development sets are generally consistent with the test sets and provided in App. D. Although the average performance of different methods over 8 tasks is sometimes similar, the differences between tasks are quite significant under certain setups and can be as large as 5~9 points on a specific task (e.g., STS-B and MNLI,  $K = 500$ ) even when excluding cases where some methods fail to learn effectively (e.g., prefix-tuning on QQP,  $K = 100$ ).

<sup>4</sup>While these hyperparameters may lead to differences in trainable parameters, we keep them for analysis as they are used by the official implementation. Also, we observe that more trainable parameters do not guarantee better results.

Next, we will analyze and examine each individual PELT method more closely.

**Analysis of Adapter.** The performance of adapter is relatively stable – there is no significantly better or worse result than fine-tuning consistent on different tasks or sizes of training data. In general, adapter is slightly worse than fine-tuning in most cases. We do not observe that adapter consistently outperforms fine-tuning in the low-resource setting as in He et al. (2021), possibly because they tune model hyperparameters on each task, which could be computationally prohibitive when there are considerable tasks. For example, they choose the bottleneck size  $D_{\text{mid}}$  from  $\{64, 128, 256\}$ , while  $D_{\text{mid}} = 48$  is fixed across tasks in our experiments. Also, we only add one adapter in each Transformer layer instead of two following Pfeiffer et al. (2021). These two differences result in 62.4%~90.5% fewer parameters than the adapter used in He et al. (2021).

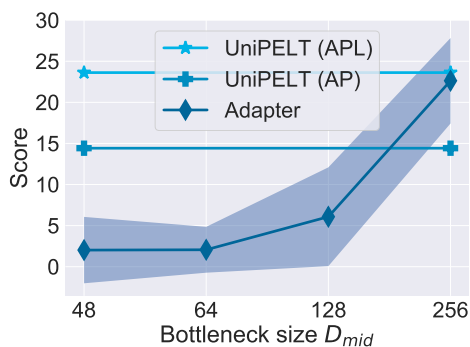


Figure 2: Performance changes when the bottleneck size of adapter is increased (on CoLA,  $K = 100$ ).

To further study the effect of bottleneck size  $D_{\text{mid}}$  in adapter, we increase  $D_{\text{mid}}$  and re-evaluate adapter on a setup that it performs poorly (CoLA,  $K = 100$ ). As shown in Fig. 2, the performance of adapter is increased gradually and becomes significantly better only when  $D_{\text{mid}} = 256$ , which involves  $5.3\times$  trainable parameters than the adapter used originally ( $D_{\text{mid}} = 48$ ),  $4.3\times$  than UNIPELT (AP), and  $3.4\times$  than UNIPELT (APL), suggesting that a larger bottleneck size could be beneficial when adapter learns ineffectively.

On the other hand, there are certain tasks (*e.g.*, STS-B) that adapter largely outperforms competitive methods such as prefix-tuning and LoRA regardless of the size of training data, suggesting that one should favor adapter over other PELT methods under certain scenarios as well.

**Analysis of Prefix-tuning.** Prefix-tuning performs

poorly with  $K = \{100, 500\}$  and becomes on par with fine-tuning when  $K$  reaches 1000. We also observe that prefix-tuning fails to learn effectively on certain tasks when the training data is limited (*e.g.*,  $K = 100$  on SST-2 and  $K = 500$  on QQP), leading to unsatisfactory performance and (or) large variance across different runs. Similar phenomena have been observed in a concurrent study (Gu et al., 2021) on few-shot prompt-tuning.

To ensure that the poor performance of prefix-tuning is not due to its fewer trainable parameters (based on its default setting), we further increase the prefix length to  $L = 50$  such that its trainable parameters are comparable to adapter, and re-evaluate prefix-tuning on all 8 tasks with  $K = 100$ . For the 4 tasks where prefix-tuning ( $L = 10$ ) performs poorly (SST2, CoLA, STS-B, and QQP), while its performance is significantly improved on 3 tasks, it also performs significantly worse on the other task (STS-B), which suggests that training instability in the low-resource regime is still an issue for prefix-tuning even with more trainable parameters.<sup>5</sup> Besides, prefix-tuning ( $L = 50$ ) still lags behind adapter or UNIPELT (AP) on 3 of the 4 tasks. Furthermore, the average performance of prefix-tuning ( $L = 50$ ) on 8 tasks is even slightly worse than with  $L = 10$ , which indicates that increasing prefix length may not be a panacea for all the scenarios. A larger  $L$  also leads to significant training/inference slowdown due to the costly multi-head attention. More broadly, such results suggest that using more trainable parameters does not guarantee better performance.

On the bright side, prefix-tuning performs well on certain tasks such as natural language inference (RTE and MNLI) with various sizes of training data, which suggests that one should also prefer prefix-tuning in certain cases.

**Analysis of BitFit & LoRA.** Tuning only the bias terms of the model does not lead to very satisfactory results in our experiments – BitFit never performs the best and generally performs the worst in different data and task setups. Therefore, we do not consider BitFit in the following experiments and exclude BitFit as a submodule of UNIPELT. As for LoRA, there are a few setups where LoRA fails to learn effectively as well, such as STS-B and QQP ( $K = \{100, 500\}$ ), leading to high variance across runs. Apart from that, LoRA performs

<sup>5</sup>Tuning other hyperparameters like learning rate does not appear to alleviate the issue either.

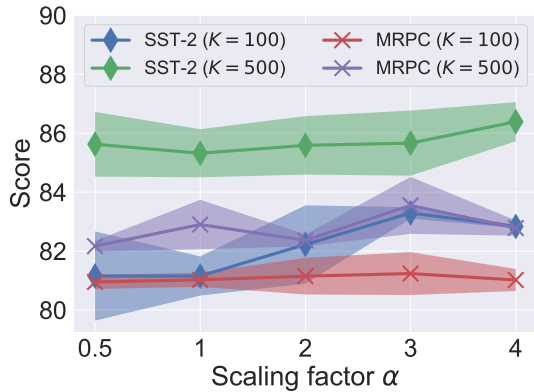


Figure 3: Performance comparison of various scaling factors for LoRA on  $2 \times 2$  task and data setups.

quite competitively despite using fewer trainable parameters than methods like adapter, especially when  $K = 1000$ , achieving the best or 2nd best performance on 4 of 8 tasks.

As LoRA has a scaling factor  $\alpha$  that can be seen as a static gating function under our formulation, we further investigate its importance by evaluating LoRA with different  $\alpha$ . As shown in Fig. 3, LoRA is quite sensitive to the scaling factor and there seems to be no single optimal value that works well across multiple task and data setups. Such findings suggest that gating is critical and motivate us to use more fine-grained and dynamic control for UNIPELT. Besides, we observe that increasing  $\alpha$  consistently results in faster convergence, possibly because the trainable parameters would receive larger gradient updates with a larger  $\alpha$ .

### 4.3 Analysis of UNIPELT

Next, we will turn to our proposed framework UNIPELT, which incorporates multiple existing PELT methods as submodules.

**Low-Resource Performance.** Overall, UNIPELT (APL) and UNIPELT (AP) consistently achieve the best and second best average performance on both the development and test sets regardless of the number of training samples. The gains are generally 1~4% over the submodule that performs the best (when used individually). Such results demonstrate the advantages of our hybrid approach regarding model effectiveness and generalizability.

At the per-task level, UNIPELT (APL) and UNIPELT (AP) perform the best or second best on 7/6/7 of 8 tasks when trained with 100/500/1,000 samples, and never perform the worst in any setup. When comparing the two variants, UNIPELT

(APL) outperforms UNIPELT (AP) on 4/6/8 of 8 tasks when trained with 100/500/1,000 samples. Such results indicate that UNIPELT is quite robust and performs reliably under different scenarios. The improvements of UNIPELT over its submodules are generally larger when having fewer training samples, suggesting that UNIPELT performs especially well in the low-resource regime. In particular, on the tasks where other PELT methods fail to learn effectively such as CoLA and QQP ( $K = 100$ ), UNIPELT manages to achieve performance better than fine-tuning.

**UNIPELT vs. Upper Bound.** In Table 2, we show the comparison of UNIPELT and the upper bound that takes the best performance of its submodules on each task. We observe that both UNIPELT (AP) and UNIPELT (APL) perform similarly or even better than their upper bound, which suggests that UNIPELT successfully learns to leverage different submodules and maintains (near) optimal performance under different setups. The fact that UNIPELT can outperform the upper bound also hints that a mixture of PELT methods (involving different parts of the PLM) might be inherently more effective than single methods (with a limited scope of the PLM architecture).

$K$	$\max(\{A, P\})$	UNIPELT	$\max(\{A, P, L\})$	UNIPELT
100	58.86	<b>60.80</b>	60.60	<b>62.59</b>
500	69.69	<b>69.79</b>	<b>70.02</b>	69.98
1000	72.58	<b>72.61</b>	73.19	<b>73.31</b>

Table 2: Comparison of average test performance between UNIPELT and the upper bound that takes the best performance of its submodules on each task.

**High-Resource Performance.** In Table 3, we list the performance of different methods when all training samples are used. UNIPELT again achieves the best overall performance. The gains are not as significant as in the low-resource setting, which is somewhat expected as existing PELT methods typically perform on par with fine-tuning given abundant training data and the potential of improvement is not as high. That said, the performance of UNIPELT is still the best or 2nd best on all 8 tasks, and generally comparable to the best submodule used individually on each task. Besides, simply combining multiple PELT methods without gating does not work well in the high-resource setting – although UNIPELT-NoGate never performs the worst in each task, its average performance is unsatisfactory (-0.89 vs. UNIPELT).

Method	SST-2	MRPC	CoLA	RTE	QNLI	STS-B	MNLI	QQP	Avg.
[ $K = all$ ] Best Performance on GLUE Dev									
Fine-tuning	91.63	<u>90.94</u>	<b>62.08</b>	66.43	89.95	<b>89.76</b>	83.23	<b>87.35</b>	82.67
Adapter	<b>91.86</b>	89.86	61.51	71.84	<u>90.55</u>	88.63	83.14	86.78	83.02
Prefix-tuning	90.94	<b>91.29</b>	55.37	<b>76.90</b>	90.39	87.19	81.15	83.30	82.07
LoRA	91.51	90.03	60.47	71.48	89.93	85.65	82.51	85.98	82.20
UNIPELT (AP)	<b>91.86</b>	90.28	61.15	71.84	<b>90.77</b>	88.86	<u>83.41</u>	86.74	<u>83.12</u>
-NoGate	<u>91.74</u>	90.18	58.63	71.12	90.30	88.76	81.58	85.53	82.23
UNIPELT (APL)	91.51	<u>90.94</u>	<u>61.53</u>	<u>73.65</u>	90.50	<u>88.93</u>	<b>83.89</b>	<u>87.12</u>	<b>83.50</b>

Table 3: Results on the GLUE benchmark when all training samples are used.

Method	#Param.	Time <sub>T</sub>	Time <sub>I</sub>
Fine-tuning	110M (100%)	100%	100%
BitFit	103K (0.09%)	65%	102%
Prefix-tuning	184K (0.17%)	56%	114%
LoRA	295K (0.27%)	53%	105%
Adapter	895K (0.81%)	55%	107%
UNIPELT (AP)	1.1M (0.99%)	55%	118%
UNIPELT (APL)	1.4M (1.26%)	67%	127%

Table 4: Number of trainable parameters and Training/Inference time relative to fine-tuning.

#### 4.4 Efficiency of PELT Methods

We benchmark the efficiency of PELT methods and list in Table 4 their number of trainable parameters and training/inference time relative to fine-tuning.

**Parameter Efficiency.** As the trainable parameters in PELT methods are almost negligible, combining multiple methods does not lead to significant losses in parameter efficiency. UNIPELT still has few trainable parameters compared to fine-tuning (0.99%~1.26%). The parameters can be further reduced if one uses more parameter-efficient variants (*e.g.*, Karimi Mahabadi et al. (2021a)), which can be easily swapped with the vanilla version used in our current framework. Also, note that more trainable parameters do not always lead to better performance, as shown in our experiments and prior studies (He et al., 2021; Pfeiffer et al., 2021).

**Training and Inference Efficiency.** Due to parameter efficiency, all PELT methods train 30%~50% faster than fine-tuning and incorporating multiple PELT methods into UNIPELT does not suffer from slower training. On the other hand, the inference time of PELT methods is generally longer since they involve more FLOPs. UNIPELT has a slightly larger inference overhead (4%~11% compared to its slowest submodule), which we argue is insignificant since larger models that may achieve similar performance gains (*e.g.*, BERT<sub>large</sub>) need

around 300% inference time (Wolf et al., 2020).

## 5 Related Work

**Parameter-Efficient Tuning of PLMs.** As it is increasingly infeasible to train and store full copies of large PLMs for various downstream tasks, how to efficiently tune the PLMs with few trainable parameters becomes critical. Existing PELT methods can be largely divided into two categories based on whether new trainable parameters are introduced. Specifically, one may either train a subset of the model parameters such as the prediction head (Lee et al., 2019) and bias terms (Ben Zaken et al., 2021), or introduce task-specific parameters to different parts of the PLM such as before multi-head attention (Li and Liang, 2021) or after feedforward layer (Houlsby et al., 2019). As the number of PELT methods keeps increasing, the purpose of UNIPELT is to better understand and leverage the distinctions of various methods instead of proposing yet another method.

**Mixture-of-Experts.** UNIPELT is also related to approaches that involve a high-capacity network and activate (upweight) different parts of the network given different inputs. One notable example is Mixture-of-Experts (MoE) (Jacobs et al., 1991; Shazeer et al., 2017), which maintains a set of experts (neural networks) and one or more trainable gates that select a combination of the experts specific to each input. Despite being conceptually similar, UNIPELT is different from MoE: the submodules in UNIPELT are not combined explicitly by summation like MoE but in sequential order and affect each other implicitly. Moreover, the “experts” are diverse in UNIPELT while usually homogeneous or identical in MoE methods.

## 6 Conclusion

In this paper, we present a comprehensive study of representative parameter-efficient language model



tuning (PELT) methods and propose a unified framework, which incorporates different PELT methods as submodules and learns to activate the most appropriate submodules for a given task or data setup. Our proposed framework consistently outperforms conventional fine-tuning as well as the submodules that it incorporates under different setups, and generally surpasses the upper bound that takes the best performance of each submodule used individually on each task. Our findings suggest that a mixture of multiple PELT methods that involve different parts of the PLM may be favorable regarding both model effectiveness and robustness. For future work, we will try to better understand the discrepancy of various PELT methods in different scenarios. We also plan to investigate a multi-task setting where multiple submodules can be activated and cooperate at the task level.

## Acknowledgements

We thank Xiang Lisa Li, Hai Ye, Rabeeh Karimi Mahabadi, Junxian He, Yiqing Xie, Yaqing Wang, and Liyuan Liu for helpful discussions and feedback. We thank anonymous reviewers for valuable comments and suggestions.

## References

- Elad Ben Zaken, Shauli Ravfogel, and Yoav Goldberg. 2021. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. *arXiv e-prints*, pages arXiv–2106.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- 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.
- Ning Ding and Shengding Hu. 2021. Must-read papers on prompt-based tuning for pre-trained language models. <https://github.com/thunlp/PromptPapers>.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. **Making pre-trained language models better few-shot learners**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3816–3830, Online. Association for Computational Linguistics.
- Yuxian Gu, Xu Han, Zhiyuan Liu, and Minlie Huang. 2021. Ppt: Pre-trained prompt tuning for few-shot learning. *arXiv preprint arXiv:2109.04332*.
- Demi Guo, Alexander Rush, and Yoon Kim. 2021. **Parameter-efficient transfer learning with diff pruning**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4884–4896, Online. Association for Computational Linguistics.
- Ruidan He, Linlin Liu, Hai Ye, Qingyu Tan, Bosheng Ding, Liying Cheng, Jiawei Low, Lidong Bing, and Luo Si. 2021. **On the effectiveness of adapter-based tuning for pretrained language model adaptation**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2208–2222, Online. Association for Computational Linguistics.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pages 2790–2799. PMLR.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. 1991. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87.
- Rabeeh Karimi Mahabadi, James Henderson, and Sebastian Ruder. 2021a. Compacter: Efficient low-rank hypercomplex adapter layers. *arXiv preprint arXiv:2106.04647*.
- Rabeeh Karimi Mahabadi, Sebastian Ruder, Mostafa Dehghani, and James Henderson. 2021b. **Parameter-efficient multi-task fine-tuning for transformers via shared hypernetworks**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 565–576, Online. Association for Computational Linguistics.
- Jaesun Lee, Raphael Tang, and Jimmy Lin. 2019. What would elsa do? freezing layers during transformer fine-tuning. *arXiv preprint arXiv:1911.03090*.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*.

- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Xiang Lisa Li and Percy Liang. 2021. [Prefix-tuning: Optimizing continuous prompts for generation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4582–4597, Online. Association for Computational Linguistics.
- Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. 2021. [AdapterFusion: Non-destructive task composition for transfer learning](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 487–503, Online. Association for Computational Linguistics.
- Jonas Pfeiffer, Andreas Rücklé, Clifton Poth, Aishwarya Kamath, Ivan Vulić, Sebastian Ruder, Kyunghyun Cho, and Iryna Gurevych. 2020a. [Adapterhub: A framework for adapting transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP 2020): Systems Demonstrations*, pages 46–54, Online. Association for Computational Linguistics.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 2020b. [MAD-X: An Adapter-Based Framework for Multi-Task Cross-Lingual Transfer](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7654–7673, Online. Association for Computational Linguistics.
- Timo Schick and Hinrich Schütze. 2021. [Exploiting cloze-questions for few-shot text classification and natural language inference](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 255–269, Online. Association for Computational Linguistics.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2019. Glue: A multi-task benchmark and analysis platform for natural language understanding. In *International Conference on Learning Representations*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. How to benchmark transformer models. <https://huggingface.co/transformers/benchmarks.html>.
- Jeffrey O Zhang, Alexander Sax, Amir Zamir, Leonidas Guibas, and Jitendra Malik. 2020. Side-tuning: A baseline for network adaptation via additive side networks. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 698–714. Springer.

## A Prefix-tuning vs. Prompt-based Fine-tuning

We note that prefix-tuning (or prompt-tuning) is different from prompt-based *fine-tuning* methods (Schick and Schütze, 2021; Gao et al., 2021) in many ways: (1) Prompt-based fine-tuning is not parameter-efficient as it updates all model parameters while prefix-tuning only updates the prefix matrix  $P$ . (2) The prompts are only used in model input for prompt-based fine-tuning but added to every Transformer layer in prefix-tuning (stored as different vectors). (3) Prompt-based fine-tuning typically leverages carefully designed *natural language* prompts while prefix-tuning uses continuous prompts (virtual tokens).

## B Implementation Details

**Data Preparation.** We shuffle the training set with seed  $s$ , take the first  $K$  samples as the new training set, and the next 1,000 samples as the development set. We use  $s = \{111, 222, 333, 444, 555\}$  as the data seeds and the same seed ( $s = 42$ ) for model training. We also conduct another set of preliminary experiments by fixing the data and using 5 different random seeds for model training, the results of which are similar (Table 5).

**Hyperparameters.** We adopt AdapterHub (Pfeiffer et al., 2020a), a library based on HuggingFace Transformers (Wolf et al., 2019), as our codebase. We largely follow the recommended hyperparameters used in different methods for a fair comparison. We set the input length to 128 and the training batch size to 16. We set the number of epochs to 50 and adopt early stopping with a patience of 10 non-increasing epochs. We set the learning rate of fine-tuning and adapter to  $2e-5$  and  $1e-4$  according to the findings in prior studies (Pfeiffer et al., 2020a; He et al., 2021). For prefix-tuning and UNIPeLT, as they are not previously evaluated on NLU tasks, we tune their learning rates from  $\{1e-4, 2e-4, 5e-4\}$  on the development set and set to  $2e-4$  and  $5e-4$ , respectively. For BitFit and LoRA, we choose the learning rates commonly used in their own experiments ( $1e-4$  and  $5e-4$ , respectively). We set  $\alpha = 2$  and  $r = 8$  in LoRA according to its official scripts.

## C BART Results

In our preliminary experiments, we also evaluated UNIPeLT on BART<sub>large</sub> (Lewis et al., 2020). We show the results of fine-tuning, adapter, prefix-

Setup	Fine-tuning	Adapter	Prefix-tuning	UNIPeLT (AP)
Model seed	78.39 <sub>2.92</sub>	77.12 <sub>0.50</sub>	73.16 <sub>2.89</sub>	<b>78.66</b> <sub>0.24</sub>
Data seed	77.55 <sub>2.94</sub>	76.87 <sub>0.55</sub>	71.90 <sub>2.47</sub>	<b>79.02</b> <sub>0.44</sub>

Table 5: Average performance with  $K = 1000$  on the GLUE benchmark with BART<sub>large</sub> as the base model. Results are averaged over 5 runs by changing the model or data seeds.

tuning, and UNIPeLT (AP) in Table 5. 1000 training examples are used and the average best performance on the GLUE development set is reported (excluding QQP). The results are largely consistent with those on BERT<sub>base</sub>. UNIPeLT again achieves the best performance with notably smaller variance.

## D Detailed Performance

In Table 6, we list the detailed results on both development and test sets of the GLUE benchmark. The observations and findings are largely consistent on the two evaluation splits.

