Sparse Structure Search for Parameter-Efficient Tuning

Anonymous Author(s)
Affiliation
Address
email

Abstract

Adapting large pre-trained models (PTMs) through fine-tuning imposes prohibitive computational and storage burdens. Recent studies of parameter-efficient tuning (PET) find that only optimizing a small portion of parameters conditioned on PTMs could yield on-par performance compared to conventional fine-tuning. Generally, PET methods exquisitely design parameter-efficient modules (PET modules) which could be applied to arbitrary fine-grained positions inside PTMs. However, the effectiveness of these fine-grained positions largely relies on sophisticated manual designation, thereby usually producing sub-optimal results. In contrast to the manual designation, we explore constructing PET modules in an automatic manner. We automatically **S**earch for the **S**parse **S**tructure of **P**arameter-**E**fficient **T**uning (S^3PET). Based on a unified framework of various PET methods, S^3PET conducts the differentiable PET structure search through bi-level optimization and proposes shifted global sigmoid method to explicitly control the number of trainable parameters. Extensive experiments show that S^3PET surpasses manual and random structures with less trainable parameters. The searched structures preserve more than 99% fine-tuning performance with 0.01% trainable parameters. Moreover, the advantage of S^3PET is amplified with extremely low trainable parameters budgets (0.0009%$\sim$0.01%). The searched structures are transferable and explainable, providing suggestions and guidance for the future design of PET methods.

1 Introduction

Increasingly large pre-trained models (PTMs) [17, 27, 30, 31, 11] building upon Transformers [36] have been emerging and achieving state-of-the-art results on a variety of downstream tasks. Despite the blessing of effectiveness, these big models also bring the curse of prohibitive costs on computation and storage during the adaptation because of the gradient computation of the whole model and the giant size of the fine-tuned checkpoint.

To alleviate such costs, studies of parameter-efficient tuning (PET) [14, 28, 42, 15, 25, 8, 20, 18] have been developed, which only train a small portion of PTMs and keep the vast majority of parameters frozen. Studies have verified that parameter-efficient tuning could achieve competitive performance compared to conventional fine-tuning with very few trainable parameters, resulting in considerable savings in model adaptation costs. Generally, these approaches manually design parameter-efficient modules (PET modules) to complete model adaptation. For example, adapter-based methods [14, 28, 25] inject two newly-introduced feed-forward layers to Transformers and only fine-tune 0.5%-8% parameters to yield promising results; BitFit [42] only fine-tunes the bias terms (0.04% - 0.1%

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

parameters) within Transformers; LoRA [15] inserts trainable rank decomposition matrices to each layer of Transformers and is successfully adopted on GPT-3 [2] with 175 billion parameters.

While early research focused on how to design practically effective PET modules, more recent research has advanced the understanding of parameter-efficient tuning more deeply. He et al. [12] bridge connections among different approaches to form a unified framework. And Ding et al. [6] indicate that the combination of different trainable modules could bring different levels of gain on downstream tasks. The above empirical evidence implies that there may exist *an optimal mixture of PET modules* that is more effective than manually designed structures. In fact, considering the fine-grained structure inside PTMs, the positions where the PET modules could be applied are numerous, but not all PET modules at all positions contribute equally to the task performance. How to find *the optimal structure of PET modules* and remove the redundancy in the trainable parameters is essential for a more efficient adaption method. Predictably, such optimal structure is difficult to construct artificially and may vary with specific tasks and models. Therefore, we propose to automatically search the optimal structure that contains a mixture of PET modules at diverse positions inside PTMs. Also importantly, the structure should be sparse to ensure the parameter efficiency.

We present **S**parse **S**tructure **S**earch for **P**arameter **E**fficient **T**uning (S^3PET) to automatically search such optimal trainable structure, which could flexibly control the number of trainable parameters according to practical requirements. The searching process and the optimization of S^3PET is guided by performance to ensure the effectiveness on specific tasks. Moreover, the structures change automatically to suit the preset limitation of the number of trainable parameters. In contrast, heuristically designed structures are usually coarse-grained and independent of performance and budget, making them neither optimal nor flexible to adjust the number of trainable parameters.

In terms of the specific methodology, we firstly construct a unified search space by applying probabilistic gating controlled by structural parameters to all potential PET modules. Then, we develop a framework of differentiable PET structure search by treating the problem as a constrained neural architecture search problem. In our framework, the structural parameters are updated via bi-level optimization [22]. Unlike the traditional neural architecture search that learns from scratch, we implement the first neural structure search based on a pre-defined backbone and under the parameter efficient tuning scenario. To search under a pre-defined budget of trainable parameters, we develop a *shifted global sigmoid* to explicitly control the number of activated PET modules in the searching phase.

We conducted extensive experiments to study the effectiveness of S^3PET. Firstly, the experiments show that with 0.01% parameters, we are able to recover 99% and 98% fine-tuning performance on GLUE [37] and SuperGLUE [38], respectively. Secondly, the searched structure surpasses the human-designed structures considerably while consuming less ($\sim 1/5$) trainable parameters. Moreover, the advantage enlarges when the number of trainable parameters is minimal (0.0009%$\sim$ 0.01%). Furthermore, the searched structures are *transferable* across tasks, which significantly strengthens the usefulness of the searched structures. Apart from the performance boost, we visualize and explain the searched structures, which is beneficial to the future design of new PET methods.

2 Related Work

Parameter Efficient Tuning (PET). Our work is related to the studies of parameter efficient tuning (PET) for pre-trained models [6]. Generally, PET only optimizes a small portion of parameters and leaves the vast majority of parameters untouched for the adaptation to downstream tasks. Pioneer work select parts of the PTMs to be trainable [35, 10, 9]. Adapter [14] is one of the earliest methods that apply the concept of parameter-wise efficiency to pre-trained language models, which inserts linear neural modules to every Transformer layer and achieves on-par results to full fine-tuning. As the PTMs scaling in recent years, PET is valued for its efficiency in computing and storage. This has spawned not only empirical studies [28, 13] and variants on the adapter [28, 25, 34], but also a range of other approaches. Prefix tuning [20] prepends embeddings to the hidden states of the Transformer model, and prompt tuning [18] further simplifies the strategy and only prepends such embeddings to the input layer. There are also approaches which specify some of the parameters inside PTMs that can be trainable to achieve good results, such as Masking [43], BitFit [42], DiffPruning [8], etc.

LoRA [15] assumes that the change in model weights is intrinsically low-rank after fine-tuning, and uses trainable rank-decomposition matrices for model adaptation. In addition to specific methods, some studies have comprehensively investigated PET methods. He et al. [12] models multiple methods in a unified manner, Ding et al. [6] provides a theoretical discussion and comprehensive empirical study of these methods. Our work proposes to automatically search for trainable structures in the context of PET, which is a different perspective from all the aforementioned work. In terms of the structure of PET, AdapterDrop [33] explores dropping a fraction of Adapter modules based on manual trials. A concurrent work [26] learns switches on Adapter modules to select the beneficial adapter modules. However, it is not optimized under a preset number of trainable parameters. They are also both limited to adapter-based methods. On the contrary, our proposed method can search within a mixture of almost all PET modules under a constrained trainable parameter budget.

Neural Architecture Search (NAS). Our work conducts structure search in the scope of PET, which is related to the Neural Architecture Search algorithms. A line of NAS algorithms uses Reinforcement Learning or Evolutionary Algorithms to explore the best structure with reward from training the structure from scratch [44, 45, 32, 29], which usually consumes prohibitive computation resources. Another line of NAS algorithms [22, 21, 5] approaches the problem with gradient-based optimization. DARTS [22] relaxes the discrete structure using continuous structural parameters, which are optimized with gradient-based optimization. DARTS achieves competitive performances with much fewer computational resources. We take inspirations from DARTS in optimizing the structural parameters of S^{3}PET. We are also the first to conduct NAS conditioned on a pre-trained backbone model. We also take inspiration from the NAS algorithms with binary gates [3, 40].

3 Method

In this section, we firstly introduce the preliminaries of pre-trained model adaptation, transformer architecture, and the parameter-efficient tuning. Then we introduce our method S^{3}PET in detail.

3.1 Preliminaries

Pretrained Model Adaptation. The recent prevalent pre-train then fine-tune paradigm in deep learning takes advantage of a pre-trained model $\mathcal{M}$ with parameters Θ and continues to optimize Θ on a downstream task $\mathcal{D} = \{\mathcal{D}_{\text{train}}, \mathcal{D}_{\text{val}}, \mathcal{D}_{\text{test}}\}$ under an objective function $\mathcal{L}$. In fine-tuning, all the parameters of the pre-trained model are optimized using the train split to minimize $\mathcal{L}$, i.e.,

$$\min_{\Theta} \mathcal{L}(M(\Theta), \mathcal{D}_{\text{train}}). \tag{1}$$

Transformer Architecture. The pre-trained models typically adopt the Transformer model [36] as their backbone. The Transformer model is composed of multiple stacked Transformer layers that processes the hidden state sequentially through different computation modules, such as Self-Attention module (SelfAttn), Cross-Attention module (CrossAttn), Feed-Forward module (FFN), and Layer Normalization module (LN), etc., and details of each module are in Appendix A. The computation process in the Transformer can be abstracted by a sequence of transformations of the hidden representation. In each computation step, the input hidden representation $\mathbf{H}^{\text{in}} \in \mathbb{R}^{s \times d_1}$ is transformed into an output hidden representation $\mathbf{H}^{\text{out}} \in \mathbb{R}^{s \times d_2}$, where s is the sequence length of the input and d_1, d_2 are the hidden dimensions,

$$\mathbf{H}^{\text{out}} = m(\mathbf{H}^{\text{in}}). \tag{2}$$

Parameter-Efficient Tuning (PET). PET methods only train a small portion of parameters conditioned on the backbone PTMs to improve the adaptation efficiency [14, 28, 42, 15, 25, 8, 20, 18]. Although the specific forms of the various PET modules are substantially different, He et al. [12] unify them as modifications Δ of the hidden state [1],

$$\mathbf{H}^{\text{out}} = m(\mathbf{H}^{\text{in}}) + \Delta. \tag{3}$$

The formulas of some PET methods under the unified view are listed in Table 1. The PET modules can be applied to extensive positions on the backbone PTMs, which are listed in the rightmost

[1] We use a little bit more flexible notation than [12], which takes into account the frozen backbone module m and thus can distinguish the PET modules that take either $\mathbf{H}^{\text{out}}$ or $\mathbf{H}^{\text{in}}$ as their input.

Table 1: Different PET methods are the specializations of the unified view (Equation (3)) and can be applied to extensive positions on the PTMs.

Method	Transformation	Δ	Potential Positions
LoRA [15]	$\mathbf{H}^{\text{out}} = \mathbf{H}^{\text{in}}(\mathbf{W} + \mathbf{AB})$	$\mathbf{H}_0\mathbf{AB}$	Weight matrices
Adapter [14]	$\mathbf{H}^{\text{out}} = m(\mathbf{H}^{\text{in}}) + f(m(\mathbf{H}^{\text{in}})\mathbf{W}_{\text{down}})\mathbf{W}_{\text{up}}$	$f(m(\mathbf{H}^{\text{in}})\mathbf{W}_{\text{down}})\mathbf{W}_{\text{up}}$	After any modules
Parallel Adapter [12]	$\mathbf{H}^{\text{out}} = m(\mathbf{H}^{\text{in}}) + f(\mathbf{H}^{\text{in}}\mathbf{W}_{\text{down}})\mathbf{W}_{\text{up}}$	$f(\mathbf{H}^{\text{in}}\mathbf{W}_{\text{down}})\mathbf{W}_{\text{up}}$	Between any two modules
BitFit [42]	$\mathbf{H}^{\text{out}} = m(\mathbf{H}^{\text{in}}) + \mathbf{b}_\delta$	$\mathbf{b}_\delta$	Linear layers
LNFit [2]	$\mathbf{H}^{\text{out}} = \frac{\mathbf{H}^{\text{in}}}{\text{Var}(\mathbf{H}^{\text{in}})}(\mathbf{s} + \mathbf{s}_\delta) + \mathbf{b}$	$\frac{\mathbf{H}^{\text{in}}}{\text{Var}(\mathbf{H}^{\text{in}})}\mathbf{s}_\delta$	Layer Normalization modules

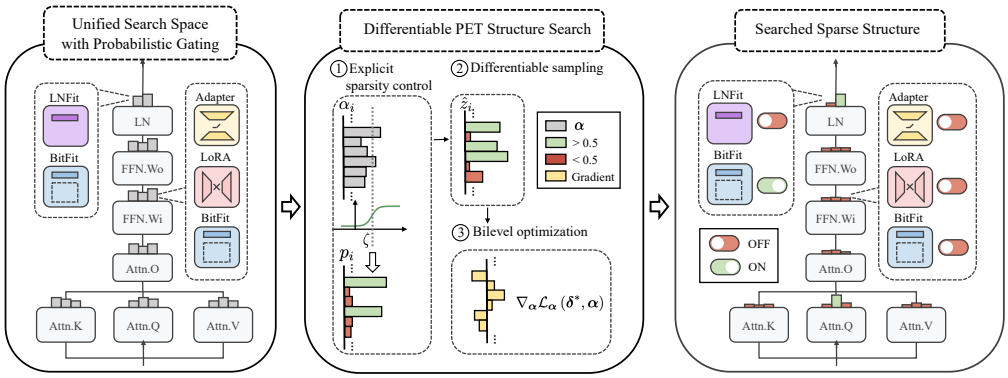

Figure 1: The framework of S^{3}PET. We propose a unified search space with probabilistic gating to enable search among a mixture of PET methods. We find the optimal sparse structure using the differentiable PET structure search and explicit sparsity control.

column in Table 1. In training, we freeze all the parameters in the backbone module m, i.e., Θ, and set parameters introduced in computing Δ, denoted by $\boldsymbol{\delta}$, as the only trainable parameters. Therefore, the adaptation objective in PET is

$$\min_{\boldsymbol{\delta}} \mathcal{L}(\mathcal{M}(\Theta, \boldsymbol{\delta}), \mathcal{D}_{\text{train}}). \tag{4}$$

For the convenience of notation, we simplify Equation (4) into

$$\min_{\boldsymbol{\delta}} \mathcal{L}_{\text{train}}(\boldsymbol{\delta}). \tag{5}$$

3.2 Sparse Structure Search for Parameter-Efficient Tuning

Our target is to search for the optimal structure of PET constrained by a pre-defined and limited trainable parameter budget $\mathcal{B}$. To achieve this, we design Sparse Structure Search for Parameter-Efficient Tuning (S^{3}PET), which is driven by three essential components: a *unified search space with probabilistic gating*, an efficient *differentiable PET structure search* algorithm, and an *explicit sparsity control algorithm using shifted global sigmoid*.

Unified Search Space with Probabilistic Gating. The potential positions in the backbone models to which PET modules could be applied are extensive, especially when we consider a mixture of different types of PET modules through the unified view (Equation (3)). However, not all positions contribute to the task performance equally, and only a fraction of positions should be *activated* to avoid redundancy of the trainable parameters. To this end, we design a *probabilistic gating* mechanism over all possible PET modules' positions. Specifically, for each PET module that computes Δ_i for the hidden representation, we *activate* the modification Δ_i with probability $p_i \in [0, 1]$,

$$\mathbf{H}_i^{\text{out}} = m(\mathbf{H}_i^{\text{in}}) + z_i \Delta_i, \tag{6}$$

where $z_i \in \{0, 1\} \sim B(1, p_i)$ is a random sample from Bernoulli distribution.

Differentiable PET Structure Search. Finding the optimal structure from a search space with abundant potential positions is challenging due to the compositionality of activated positions. Furthermore,

[2]LNFit only trains the variance vector in the Layer Normalization module of the PTMs, which is inspired by Frankle et al. [7] who only train the Batch Normalization module in Convolutional Neural Networks.

directly comparing the fully trained model with each PET structure is intractable. In our work, we propose to *optimize the gating probability p_i with gradient-based optimization*. To make the sampling process differentiable, we use the Binary Concrete Distribution [24, 16] as a soft and differentiable approximation to the Bernoulli distribution,

$$\hat{z}_i = \sigma\left(\frac{1}{\beta}\log\frac{up_i}{(1-u)(1-p_i)}\right), \tag{7}$$

where $u \sim U(0,1)$ is a random sample from the uniform distribution in $[0,1]$ and β is the temperature to control the sharpness of the distribution $\hat{z}_i$ distribution [3]. Similar distributions are used in learning sparse networks [23] or pruning dense networks [39]. By replacing the hard sample z_i with the soft approximation $\hat{z}_i$, we can back-propagate through p_i in training. However, directly optimizing p_i in the probability space $[0,1]$ may lead to numerical instability, therefore, we parameterize it with structural parameters $\alpha_i \in \mathbb{R}$, i.e., $p_i = g(\alpha_i)$. And we denote all the structure parameters as $\boldsymbol{\alpha}$.

We optimize $\boldsymbol{\alpha}$ through bi-level optimization [1, 22], i.e., optimizing $\boldsymbol{\alpha}$ conditioned on the optimized parameters $\boldsymbol{\delta}^*$ of the PET modules. The inner and outer level of optimization are conducted on separate splits of training data, denoted by $\mathcal{D}_{\boldsymbol{\delta}}, \mathcal{D}_{\boldsymbol{\alpha}}$, which is analogous to validating structures trained on $\mathcal{D}_{\boldsymbol{\delta}}$ using a different split $\mathcal{D}_{\boldsymbol{\alpha}}$ to avoid over-fitting to $\mathcal{D}_{\boldsymbol{\delta}}$. Thus, the optimization objective is

$$\min_{\boldsymbol{\alpha}} \mathcal{L}_{\boldsymbol{\alpha}}(\mathcal{M}(\Theta, \boldsymbol{\delta}^*, \boldsymbol{\alpha})), \tag{8}$$

$$s.t.\ \boldsymbol{\delta}^* = \operatorname{argmin}_{\boldsymbol{\delta}} \mathcal{L}_{\boldsymbol{\delta}}(\mathcal{M}(\Theta, \boldsymbol{\delta}, \boldsymbol{\alpha})). \tag{9}$$

Following DARTS [22], we make approximations to the gradient of the structural parameters by applying chain rule and taking finite difference approximations [4],

$$\nabla_{\boldsymbol{\alpha}}\mathcal{L}_{\boldsymbol{\alpha}}\left(\boldsymbol{\delta}^*, \boldsymbol{\alpha}\right) \tag{10}$$

$$\approx \nabla_{\boldsymbol{\alpha}}\mathcal{L}_{\boldsymbol{\alpha}}\left(\boldsymbol{\delta} - \xi\nabla_{\boldsymbol{\delta}}\mathcal{L}_{\boldsymbol{\delta}}(\boldsymbol{\delta}, \boldsymbol{\alpha}), \boldsymbol{\alpha}\right) \tag{11}$$

$$\approx \nabla_{\alpha}\mathcal{L}_{\boldsymbol{\alpha}}\left(\boldsymbol{\delta}', \boldsymbol{\alpha}\right) - \xi\nabla^2_{\boldsymbol{\alpha},\boldsymbol{\delta}}\mathcal{L}_{\boldsymbol{\delta}}(\boldsymbol{\delta}, \boldsymbol{\alpha})\nabla_{\boldsymbol{\delta}'}\mathcal{L}_{\boldsymbol{\alpha}}\left(\boldsymbol{\delta}', \boldsymbol{\alpha}\right) \tag{12}$$

$$\approx \nabla_{\boldsymbol{\alpha}}\mathcal{L}_{\boldsymbol{\alpha}}(\boldsymbol{\delta}', \boldsymbol{\alpha}) - \xi\frac{\nabla_{\boldsymbol{\alpha}}\mathcal{L}_{\boldsymbol{\delta}}(\boldsymbol{\delta}^+, \boldsymbol{\alpha}) - \nabla_{\boldsymbol{\alpha}}\mathcal{L}_{\boldsymbol{\delta}}(\boldsymbol{\delta}^-, \boldsymbol{\alpha})}{2\epsilon}, \tag{13}$$

where the optimal $\boldsymbol{\delta}^*$ is approximated by the parameters of one-step update $\boldsymbol{\delta}' = \boldsymbol{\delta} - \xi\nabla_{\boldsymbol{\delta}}\mathcal{L}_{\boldsymbol{\delta}}(\boldsymbol{\delta}, \boldsymbol{\alpha})$. ξ is the learning rate of parameters $\boldsymbol{\delta}$, and ϵ is a small scalar used in the finite difference approximation.

Explicit Sparsity Control with Shifted Global Sigmoid. Most of the PET modules in the search space are redundant and contribute little to the performance. However, the search algorithm may not be aware of the sparsity target and degenerate to greedily adding more PET modules. As opposed to previous sparse network learning methods [23, 8] which punish the dense structures with L_0 regularization, we explicitly control the sparsity of structure at the target level during the search through a *shifted global sigmoid* parameterization. (See Section 4.9 for comparing the two methods),

$$p_i = \tilde{p}_i\frac{\sum_i \mathrm{Detach}(\tilde{p}_i)}{\sum_i \tilde{p}_i}, \tag{14}$$

$$\text{where} \quad \tilde{p}_i = \mathrm{Sigmoid}(\frac{\alpha_i - \zeta}{\tau}). \tag{15}$$

The $\mathrm{Detach}(\cdot)$ operator turns a parameter that requires gradient into a scalar that is free from gradient computation. Equation (14) doesn't change the value of $\tilde{p}_i$, but it enforces the competition among different positions and PET modules, which is similar to $\mathrm{Softmax}$ operation (See Appendix C.2 for details).

In Equation (15), ζ is a scalar. Increasing ζ's value will monotonically reduce p_i to 0 while keeping p_i in $[0,1]$. So the expected number of trainable parameters $\mathbb{E}[N]$ is a monotonic function w.r.t. ζ,

$$\mathbb{E}[N] = \mathbb{E}\left[\sum_i \mathbb{I}(z_i = 1)|\boldsymbol{\delta}_i|\right] \approx \mathbb{E}\left[\sum_i \mathbb{I}(\hat{z}_i > 0.5)|\boldsymbol{\delta}_i|\right] = \sum_i p_i|\boldsymbol{\delta}_i|, \tag{16}$$

[3] The distribution of $\hat{z}_i$ has the property that $P(\hat{z}_i > 0.5) = p_i$, and when β approximate 0, the distribution of $\hat{z}_i$ converges to $B(1, p_i)$ (See Appendix C.1), which makes it a suitable surrogate for Bernoulli distribution.

[4] We use the same $\hat{z}_i$ sample to compute $\nabla_{\boldsymbol{\alpha}}\mathcal{L}_{\boldsymbol{\delta}}(\boldsymbol{\delta}^+, \boldsymbol{\alpha})$, $\nabla_{\boldsymbol{\alpha}}\mathcal{L}_{\boldsymbol{\delta}}(\boldsymbol{\delta}^-, \boldsymbol{\alpha})$, $\nabla_{\boldsymbol{\delta}}\mathcal{L}_{\boldsymbol{\delta}}(\boldsymbol{\delta}, \boldsymbol{\alpha})$, and $\nabla_{\boldsymbol{\alpha}}\mathcal{L}_{\boldsymbol{\alpha}}\left(\boldsymbol{\delta}', \boldsymbol{\alpha}\right)$.

where the $|\boldsymbol{\delta}_i|$ is the number of parameters introduced in computing Δ_i. Thus, we can dynamically adjust ζ to make $\mathbb{E}[N]$ approach $\mathcal{B}$ via monotonic optimization,

$$\zeta^* = \text{argmin}_\zeta(\sum_i p_i|\boldsymbol{\delta}_i| - \mathcal{B}), \text{ where } \sum_i p_i|\boldsymbol{\delta}_i| \leq \mathcal{B}. \tag{17}$$

Evaluation of the Searched Structure. To determine the final structure of PET, instead of sampling from p_i, we choose the set of positions where the sum of p_i is the highest while still being within the budget $\mathcal{B}$. This deterministic algorithm reduces the variance of the final structures. After obtaining the final structure, we re-initialize and re-train the parameters in the PET modules to converge on $\mathcal{D}_{\text{train}}$.

Algorithm 1 Algorithm of S^{3}PET

Initialize all PET modules in the search space, and initialize $\boldsymbol{\alpha}$.
while *not converged* **do**
1. Calculate ζ, p_i, and sample $\hat{z}_i$.
2. Compute the each loss terms by forward and backward propagation.
3. Update $\boldsymbol{\alpha}$ according to Equation (13).
4. Update $\boldsymbol{\delta}$ using $\nabla_{\boldsymbol{\delta}}\mathcal{L}_{\boldsymbol{\delta}}(\boldsymbol{\delta}, \boldsymbol{\alpha})$.

end while
Determine and evaluate the final structure.

4 Experiments

4.1 Datasets and PTMs

We apply S^{3}PET to multitask benchmarks GLUE [37] and SuperGLUE [38] following previous works. All datasets are downloaded from the HuggingFace Datasets [19]. Since the test splits of these datasets are held officially and invisible to the researchers, we conduct random splits from either train set or validation set to make the new train, validation, and test splits, which is critical to ensure fair evaluations according to Chen et al. [4]. We repeat 4 times using different random seeds for experiments in Table 2, and 8 times for experiments in Figure 2. The details are in Appendix B. We use the T5$_{\text{large}}$ model (703M parameters) as the backbone PTMs.

4.2 Baselines

We compare S^{3}PET with several widely used baselines (See Appendix B for details).

Fine-tune. Traditional fine-tuning trains all parameters in the PTMs.

LoRA. We apply LoRA linear layer to the Self-Attention's query modules and value modules as Hu et al. [15] suggest. We include two rank levels ($r = 8$ and $r = 1$) in our experiments.

Adapter. We adopt the first adapter method proposed by Houlsby et al. [14]. Their method requires more parameters than the other methods but achieves good empirical results.

Low Rank Adapter (Adapter-LR). We adopt the Low Rank Adapter as an efficient variant of the adapter-based method. It is proposed in [25] as a simple but effective baseline. The rank is set to 1.

BitFit. BitFit proposes to only adapt the bias layer in the model. We adopt the same setting as Zaken et al. [42] that tunes the bias inside all linear modules, and the Layer Normalization layer[5].

LNFit. We train the variance vector of all Layer Normalization layers', including the Layer Normalization after the whole transformer encoder.

We do not include Prompt Tuning [18] as our baseline because it takes much longer steps to converge and doesn't achieve competitive performance on T5$_{\text{large}}$ [18].

4.3 S^{3}PET Search Spaces and Budgets

The search space has a noteworthy influence on the performance of S^{3}PET. In our experiment, we define two kinds of search space.

[5] Although T5 has no bias in linear modules, we can treat it as bias vectors with zero initialization.

Table 2: Results on GLUE [37] benchmark (above) and SuperGLUE [38] benchmark (below). Green and blue represent the best and second best scores, respectively, among the methods in our search space. The first three rows represent the results of fine-tuning and other PET methods, which are not used in our search space due to high trainable parameter ratios. On SuperGLUE tasks, since the results on COPA vary dramatically ($\pm$26.00), the average results of SuperGLUE become easily dominated by the results on COPA. Therefore we also report the average results that exclude COPA ($\text{AVG}_{-\text{COPA}}$). The widths of the yellow rectangles are proportional to the trainable parameter ratios.

GLUE									
Parameter Ratios	Method	CoLA	SST2	MRPC	QQP	STSB	MNLI	QNLI	AVG
10000%%	Fine-tune	62.25 ± 3.96	95.87 ± 0.42	91.86 ± 1.19	89.50 ± 0.22	91.86 ± 0.46	89.61 ± 0.30	94.22 ± 0.35	87.88
65.33%%	Adapter	59.03 ± 3.06	95.90 ± 0.29	93.02 ± 0.28	88.39 ± 0.06	91.77 ± 0.25	89.53 ± 0.07	94.17 ± 0.19	87.40
21.32%%	LoRA(r=8)	58.43 ± 4.16	95.79 ± 0.27	92.21 ± 0.88	88.35 ± 0.25	91.78 ± 0.31	89.38 ± 0.32	94.14 ± 0.12	87.15
Methods in the Search Space									
8.13%%	BitFit	56.98 ± 3.89	96.24 ± 0.33	92.16 ± 0.68	88.12 ± 0.07	91.59 ± 0.08	89.10 ± 0.09	94.07 ± 0.21	86.90
4.12%%	Adapter-LR	56.78 ± 4.80	95.90 ± 0.14	92.76 ± 0.67	88.08 ± 0.13	91.26 ± 0.31	89.30 ± 0.14	93.94 ± 0.07	86.86
2.67%%	LoRA(r=1)	56.77 ± 2.29	95.81 ± 0.27	92.45 ± 1.00	88.08 ± 0.11	91.54 ± 0.33	89.16 ± 0.17	94.10 ± 0.05	86.84
1.70%%	LNFit	56.15 ± 4.06	95.81 ± 0.20	91.71 ± 0.39	88.17 ± 0.10	91.37 ± 0.24	89.11 ± 0.09	93.99 ± 0.20	86.62
1.39%%	S^3PET-M	59.34 ± 4.75	95.84 ± 0.14	92.13 ± 2.09	88.04 ± 0.23	91.58 ± 0.25	89.14 ± 0.13	94.12 ± 0.12	**87.17**
1.39%%	S^3PET-L	56.71 ± 3.03	95.93 ± 0.15	93.27 ± 1.39	88.14 ± 0.08	91.58 ± 0.49	88.81 ± 0.44	93.95 ± 0.11	**86.91**
0.35%%	S^3PET-M	54.56 ± 3.66	95.93 ± 0.24	92.14 ± 1.10	88.02 ± 0.20	91.38 ± 0.34	89.04 ± 0.25	93.93 ± 0.14	86.43

SuperGLUE										
Parameter Ratios	Method	BoolQ	CB	COPA	MultiRC	ReCORD	RTE	WIC	AVG	$\text{AVG}_{-\text{COPA}}$
10000%%	Fine-tune	86.67 ± 0.21	96.43 ± 2.92	73.50 ± 5.26	76.65 ± 1.01	85.03 ± 0.67	88.49 ± 2.12	73.12 ± 1.71	82.84	84.40
65.33%%	Adapter	85.98 ± 0.68	94.64 ± 6.19	63.00 ± 7.75	77.60 ± 0.84	85.96 ± 0.37	89.21 ± 2.94	71.63 ± 0.90	81.15	84.17
21.32%%	LoRA(r=8)	85.06 ± 0.70	91.96 ± 3.42	51.00 ± 4.16	76.94 ± 1.16	85.84 ± 0.21	87.05 ± 0.59	72.10 ± 1.31	78.56	83.16
Methods in the Search Space										
8.13%%	BitFit	85.02 ± 0.48	89.29 ± 2.92	75.00 ± 8.08	75.79 ± 1.15	85.85 ± 0.32	86.15 ± 1.48	72.34 ± 1.61	81.35	82.41
4.12%%	Adapter-LR	84.53 ± 0.37	84.82 ± 8.44	49.50 ± 6.81	76.67 ± 1.37	86.04 ± 0.09	85.61 ± 2.42	71.39 ± 0.70	76.94	81.51
2.67%%	LoRA(r=1)	85.60 ± 0.45	84.82 ± 1.79	67.50 ± 5.00	76.71 ± 1.05	85.95 ± 0.36	86.87 ± 1.08	71.32 ± 1.29	79.82	81.88
1.70%%	LNFit	84.07 ± 0.50	82.14 ± 2.92	49.00 ± 1.15	75.52 ± 1.16	86.14 ± 0.11	86.69 ± 1.81	69.28 ± 1.49	76.12	80.64
1.39%%	S^3PET-M	84.92 ± 0.68	92.86 ± 2.92	70.50 ± 3.79	76.38 ± 0.92	86.10 ± 0.11	86.69 ± 1.90	71.63 ± 1.07	81.30	**83.10**
1.39%%	S^3PET-L	85.00 ± 0.67	90.18 ± 6.10	60.00 ± 9.52	76.17 ± 1.41	86.02 ± 0.15	85.79 ± 1.36	71.63 ± 1.39	79.26	**82.46**
0.35%%	S^3PET-M	83.56 ± 0.53	87.50 ± 4.61	54.00 ± 4.32	76.09 ± 0.97	86.10 ± 0.26	85.79 ± 1.89	68.42 ± 1.89	77.35	81.24

Mix. The first search space considers a mixture of LoRA, Adapter-LR, Bitfit, and LNFIT modules. LoRA can be applied to any linear modules in the transformer block, including the query(Q), key(K), value(V), output(O) sub-modules of the attention module(ATTN), and the two sub layer W1 and W2 in Feed Forward modules(FFN). The Adapter-LR can theoretically be applied to any position in the computational graph. However, to avoid overcomplicating the search space, we limit it to the outputs of the ATTN module and the FFN modules. For BitFit, the potential applied positions are all the linear modules (Q, K, V, O, W1, W2) and the Layer Normalization modules (LN). For LNFit, the potential positions are all the LN modules. In this search space, there are 916 potential PET modules to be selected and the total number of candidate structures is 2^{916} if we do not consider the budget constraint. We denote the structures searched on this search space as **S^3PET-M**.

LoRA. We narrow down the search space into a single type of PET module. We choose LoRA as an example. The potential positions are the same as the LoRA modules in Mix search space. There are 288 potential positions in total. We denote the structures searched on this search space as **S^3PET-L**.

We also explore different numbers of trainable parameters. Experiments in Table 2 are conducted on 1.39%% and 0.35%% trainable parameters ratios. More sparsity levels are tested in section 4.5.

4.4 Results on GLUE and SuperGLUE

Table 2 shows the performance of different methods on GLUE and SuperGLUE tasks. Comparing S^3PET with the manual structures within the search space, we find that S^3PET-M (1.39%%) achieves the highest average score on GLUE and SuperGLUE (without COPA) despite using the least number of trainable parameters ($\sim 1/5$ compared to BitFit). S^3PET-M (0.35%%) also surpasses Adapter-LR, LoRA (r=1), LNFit using approximately $1/12, 1/8, 1/5$ trainable parameters, respectively. Narrowing the search space from Mix to LoRA leads to a moderate decrease in performance, which justifies the need to search among a mixture of PET modules. It may also hint that the combination of different PET modules could lead to stronger performance. However, even though the performance of S^3PET-L is not optimal, it compares favorably to LoRA(r=1), which also shows that the human designed structures, though benefit from applying PET modules uniformly on the PTMs, is sub-optimal. Compared with fine-tuning, S^3PET-M (1.39%%) preserves 99.2%

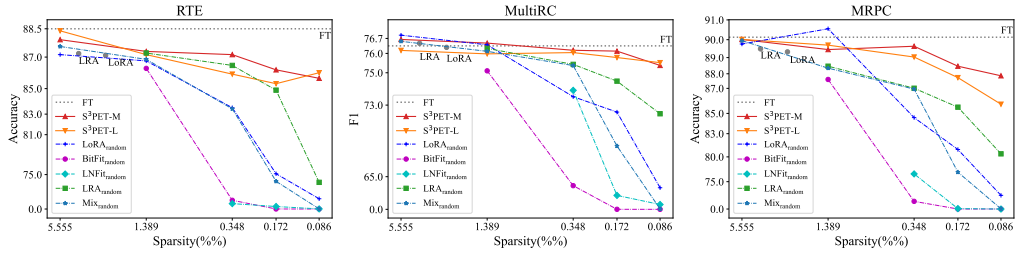

(a) RTE (b) MultiRC (c) MRPC

Figure 2: Performances under different trainable parameters ratios. The x-axis represents the ratio of the number of trainable parameters to the backbone PTM's parameters. The scaled y-axis represents the scores. The accuracy of fine-tuning is in gray horizontal line. The result of LoRA (r=1) and Low Rank Adapter are plotted in grey dot.

and 98.1% performance on GLUE and SuperGLUE, respectively. In fact, we must emphasize that S^{3}PET is orthogonal to specific PET modules. The performance of S^{3}PET can benefit from the future invention of better PET modules, thus potentially achieving comparable or even superior performance to fine-tuning with extremely limited trainable parameters.

4.5 Performance under Different Sparsity Levels

To explore the limit of trainable parameter reduction, we train different methods with decreasing sparsity levels from 5.6%% to 0.086%%. To apply baseline methods on target numbers of trainable parameters, we randomly sample a set of potential positions in their corresponding search space to reach the target sparsity level. In Figure 2, we demonstrate the results on three datasets, RTE, MultiRC and MRPC. We can see that S^{3}PET-M and S^{3}PET-L have considerable advantages in extremely low trainable parameter budgets. For example, S^{3}PET-M trains only 0.086%% parameters whereas recovering 96.8%,98.7%,97.5% of the FT performances on RTE, MultiRC, MRPC, respectively. With 5.6%% trainable parameters on MultiRC and MPRC, all the methods saturate to FT performance, proving the feasibility of removing redundant parameters with S^{3}PET.

4.6 Transferability of the Searched Structures

Another essential characteristic of S^{3}PET is the transferability of the searched structure. In Table 4, we split the GLUE benchmark into source datasets and target datasets. We search on the source dataset (Mix search space) and train the searched structure on the target datasets. We can see that the searched structures are highly transferable, even surpassing the structures direct searched on the target datasets sometimes. The transferability guarantees the reusability of the searched structures.

4.7 Efficiency of the Search Process

Although S^{3}PET focuses on the parameter-efficiency of the searched structures, we also analyze the searching efficiency in Table 3. Generally speaking, the search for an optimal structure consumes 5$\sim$8 times training time and 2 times GPU memory (Due to bi-level optimization). However, it is affordable compared to manually designing different structures and running numerous evaluations.

4.8 Visualization and Explanations of the Search Structures

To understand the searched structures, we draw the heat maps of the p_i on different datasets in Appendix D.2. We find obvious patterns and similarities in most datasets. Therefore, we average the p_i across datasets to see the overall pattern of the searched structure. Figure 3 shows the heatmap of p_i of S^{3}PET-M. We can see that (1) The BitFit modules in the Self-Attention modules and Cross-Attention modules in the higher decoder layers are highly preferred, proving that the BitFit modules are simple and effective. This observation is also beyond the intuition of human experts, as most previous work ignores the contribution of training or applying PET methods to Cross-Attention modules; (2) The last layers of the encoder and decoder are emphasized, which is close to the traditional use of PTMs

Table 3: The computational resources in the searching phase and re-training phase, Computation time, memory consumption are listed.

Dataset	Search	Re-train	Ratio
Time/min			
RTE	148.6	30.0	5.0
STSB	139.3	26.3	5.3
CoLA	145.6	17.0	8.6
Memory/GB			
RTE	27.7	16.6	1.7
STSB	28.9	10.6	2.7
CoLA	16.1	8.9	1.8

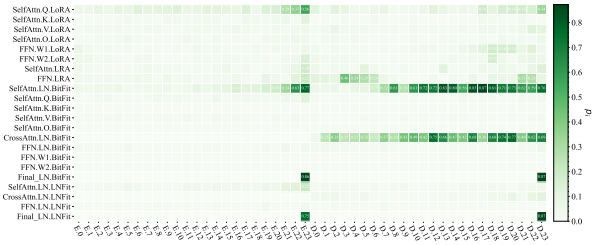

Figure 3: Visualization of p_i of S^{3}PET-M. The numbers on the squares are the average of p_i across all datasets and all seeds. The deeper the color is, the more activated is the PET module. The x-axis represents different layers of PTMs (E denotes Encoder, and D denotes Decoder), and the y-axis represents different positions (modules) of PTMs.

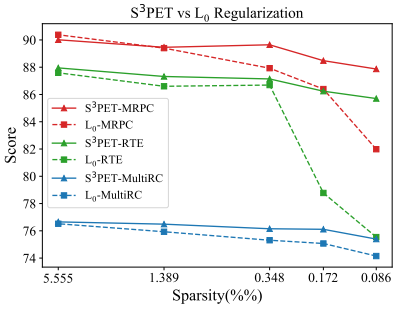

Figure 4: Comparing *shifted global sigmoid* to L_0 regularization. Performance on different datasets is in different colors. *Shifted global sigmoid* and L_0 *regularization* are in solid lines and dotted lines, respectively.

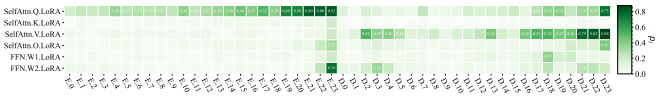

Figure 5: Visualization of p_i of PET modules in S^{3}PET-L.

Table 4: Structure transfer from source datasets and target datasets. The target datasets are in the row name, and the source datasets are in the column names. "No transfer" means the structure is searched on the target dataset.

Source	Target Datasets			
	STSB	QQP	QNLI	CoLA
No transfer	91.58 ± 0.25	88.03 ± 0.23	94.11 ± 0.12	59.34 ± 4.75
MRPC	91.63 ± 0.31	88.16 ± 0.08	93.96 ± 0.06	56.41 ± 3.81
MNLI	91.39 ± 0.67	88.06 ± 0.08	94.13 ± 0.10	56.38 ± 3.98
SST2	91.37 ± 0.23	88.02 ± 0.10	94.14 ± 0.16	55.58 ± 4.13

as feature extractors by training only the last layer; (3) We also observe that BitFit modules tend to be distributed approximately evenly across the higher layers (See Appendix D.2 for details). Figure 5 shows the p_i of S^{3}PET-L. The trend of choosing higher layers still exists. Interestingly, the query sub-modules are prioritized in the encoder, while the value sub-modules are stressed in the decoder.

4.9 Ablation Study

To explicitly control sparsity, we propose *shifted global sigmoid*, which differs from the L_0 regularization used in previous work [23]. We compare the results of regularization using *shifted global sigmoid* and L_0 on three datasets. From Figure 4, it is clear that *shifted global sigmoid* has an advantage over L_0 regularization at almost all sparsity, and the advantage increases with increasing sparsity.

5 Conclusion

In this paper, we propose Sparse Structure Search for Parameter Efficient Tuning (S^{3}PET), which conducts differentiable PET structure search with explicit sparsity control in a unified search space of a mixture of various PET modules. Experiments demonstrate the effectiveness of S^{3}PET to find the optimal structure of PET modules and push the limit of trainable parameter reduction. For future works, there are open questions that are worth investigating. (1) Better search spaces or better PET modules could be designed to further explore the potential of structure search. (2) The current NAS algorithms are not tailored for the scenario where a pre-trained backbone model exists. Therefore, more specialized search algorithms could be developed for PET structure search.

References

[1] G Anandalingam and Terry L Friesz. Hierarchical optimization: An introduction. *Annals of Operations Research*, 34(1):1–11, 1992.

[2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

[3] Han Cai, Ligeng Zhu, and Song Han. Proxylessnas: Direct neural architecture search on target task and hardware. *arXiv preprint arXiv:1812.00332*, 2018.

[4] Guanzheng Chen, Fangyu Liu, Zaiqiao Meng, and Shangsong Liang. Revisiting parameter-efficient tuning: Are we really there yet? *arXiv preprint arXiv:2202.07962*, 2022.

[5] Xin Chen, Lingxi Xie, Jun Wu, and Qi Tian. Progressive differentiable architecture search: Bridging the depth gap between search and evaluation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1294–1303, 2019.

[6] Ning Ding, Yujia Qin, Guang Yang, Fuchao Wei, Zonghan Yang, Yusheng Su, Shengding Hu, Yulin Chen, Chi-Min Chan, Weize Chen, et al. Delta tuning: A comprehensive study of parameter efficient methods for pre-trained language models. *arXiv preprint arXiv:2203.06904*, 2022.

[7] Jonathan Frankle, David J Schwab, and Ari S Morcos. Training batchnorm and only batchnorm: On the expressive power of random features in cnns. *arXiv preprint arXiv:2003.00152*, 2020.

[8] Demi Guo, Alexander Rush, and Yoon Kim. 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, 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.378. URL `https://aclanthology.org/2021.acl-long.378`.

[9] Yunhui Guo, Honghui Shi, Abhishek Kumar, Kristen Grauman, Tajana Rosing, and Rogerio Feris. Spottune: transfer learning through adaptive fine-tuning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4805–4814, 2019.

[10] Yunhui Guo, Yandong Li, Liqiang Wang, and Tajana Rosing. Adafilter: Adaptive filter fine-tuning for deep transfer learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 4060–4066, 2020.

[11] Xu Han, Zhengyan Zhang, Ning Ding, Yuxian Gu, Xiao Liu, Yuqi Huo, Jiezhong Qiu, Liang Zhang, Wentao Han, Minlie Huang, Qin Jin, Yanyan Lan, Yang Liu, Zhiyuan Liu, Zhiwu Lu, Xipeng Qiu, Ruihua Song, Jie Tang, Ji-Rong Wen, Jinhui Yuan, Wayne Xin Zhao, and Jun Zhu. Pre-trained models: Past, present and future. *AI Open*, 2021. ISSN 2666-6510. doi: https://doi.org/10.1016/j.aiopen.2021.08.002. URL `https://www.sciencedirect.com/science/article/pii/S2666651021000231`.

[12] Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Towards a unified view of parameter-efficient transfer learning. In *International Conference on Learning Representations*, 2022. URL `https://openreview.net/forum?id=0RDcd5Axok`.

[13] Ruidan He, Linlin Liu, Hai Ye, Qingyu Tan, Bosheng Ding, Liying Cheng, Jia-Wei Low, Lidong Bing, and Luo Si. On the effectiveness of adapter-based tuning for pretrained language model adaptation. *arXiv preprint arXiv:2106.03164*, 2021.

[14] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for NLP. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799. PMLR, 2019. URL `http://proceedings.mlr.press/v97/houlsby19a.html`.

[15] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *ArXiv preprint*, abs/2106.09685, 2021. URL `https://arxiv.org/abs/2106.09685`.

[16] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.

[17] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186, 2019.

[18] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *ArXiv preprint*, abs/2104.08691, 2021. URL `https://arxiv.org/abs/2104.08691`.

[19] Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, et al. Datasets: A community library for natural language processing. *arXiv preprint arXiv:2109.02846*, 2021.

[20] Xiang Lisa Li and Percy Liang. 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, 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.353. URL `https://aclanthology.org/2021.acl-long.353`.

[21] Hanwen Liang, Shifeng Zhang, Jiacheng Sun, Xingqiu He, Weiran Huang, Kechen Zhuang, and Zhenguo Li. Darts+: Improved differentiable architecture search with early stopping. *arXiv preprint arXiv:1909.06035*, 2019.

[22] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055*, 2018.

[23] Christos Louizos, Max Welling, and Diederik P Kingma. Learning sparse neural networks through l_0 regularization. *arXiv preprint arXiv:1712.01312*, 2017.

[24] Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.

[25] Rabeeh Karimi Mahabadi, James Henderson, and Sebastian Ruder. Compacter: Efficient low-rank hypercomplex adapter layers. *ArXiv preprint*, abs/2106.04647, 2021. URL `https://arxiv.org/abs/2106.04647`.

[26] Nafise Sadat Moosavi, Quentin Delfosse, Kristian Kersting, and Iryna Gurevych. Adaptable adapters. *arXiv preprint arXiv:2205.01549*, 2022.

[27] Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. In *Proceedings of NAACL-HLT*, pages 2227–2237, 2018.

[28] Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. Adapterfusion: Non-destructive task composition for transfer learning. *arXiv preprint arXiv:2005.00247*, 2020.

[29] Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, and Jeff Dean. Efficient neural architecture search via parameters sharing. In *International conference on machine learning*, pages 4095–4104. PMLR, 2018.

[30] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. 2018.

[31] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21:1–67, 2020.

[32] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. In *Proceedings of the aaai conference on artificial intelligence*, volume 33, pages 4780–4789, 2019.

[33] Andreas Rücklé, Gregor Geigle, Max Glockner, Tilman Beck, Jonas Pfeiffer, Nils Reimers, and Iryna Gurevych. AdapterDrop: On the efficiency of adapters in transformers. In *Proceedings of EMNLP*, pages 7930–7946, 2021. URL `https://aclanthology.org/2021.emnlp-main.626`.

[34] Yi-Lin Sung, Jaemin Cho, and Mohit Bansal. Vl-adapter: Parameter-efficient transfer learning for vision-and-language tasks. *arXiv preprint arXiv:2112.06825*, 2021.

[35] Nima Tajbakhsh, Jae Y Shin, Suryakanth R Gurudu, R Todd Hurst, Christopher B Kendall, Michael B Gotway, and Jianming Liang. Convolutional neural networks for medical image analysis: Full training or fine tuning? *IEEE transactions on medical imaging*, 35(5):1299–1312, 2016.

[36] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[37] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*, 2018.

[38] Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. Superglue: A stickier benchmark for general-purpose language understanding systems. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL `https://proceedings.neurips.cc/paper/2019/file/4496bf24afe7fab6f046bf4923da8de6-Paper.pdf`.

[39] Ziheng Wang, Jeremy Wohlwend, and Tao Lei. Structured pruning of large language models. *arXiv preprint arXiv:1910.04732*, 2019.

[40] Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Yanghan Wang, Fei Sun, Yiming Wu, Yuandong Tian, Peter Vajda, Yangqing Jia, and Kurt Keutzer. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10734–10742, 2019.

[41] Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tieyan Liu. On layer normalization in the transformer architecture. In *International Conference on Machine Learning*, pages 10524–10533. PMLR, 2020.

[42] Elad Ben Zaken, Shauli Ravfogel, and Yoav Goldberg. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. *ArXiv preprint*, abs/2106.10199, 2021. URL `https://arxiv.org/abs/2106.10199`.

[43] Mengjie Zhao, Tao Lin, Fei Mi, Martin Jaggi, and Hinrich Schütze. Masking as an efficient alternative to finetuning for pretrained language models. *arXiv preprint arXiv:2004.12406*, 2020.

[44] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016.

[45] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8697–8710, 2018.