# Scattered or Connected? An Optimized Parameter-efficient **Tuning Approach for Information Retrieval**

Xinyu Ma

CAS Key Lab of Network Data Science and Technology, ICT, CAS University of Chinese Academy of Sciences Beijing, China maxinyu17g@ict.ac.cn

Jiafeng Guo\* CAS Key Lab of Network Data Science and Technology, ICT, CAS University of Chinese Academy of Sciences Beijing, China guojiafeng@ict.ac.cn

Yixing Fan CAS Key Lab of Network Data Science and Technology, ICT, CAS University of Chinese Academy of Sciences Beijing, China fanyixing@ict.ac.cn

**Ruqing Zhang** CAS Key Lab of Network Data Science and Technology, ICT, CAS University of Chinese Academy of Sciences Beijing, China zhangruqing@ict.ac.cn

## ABSTRACT

Pre-training and fine-tuning have achieved significant advances in the information retrieval (IR). A typical approach is to fine-tune all the parameters of large-scale pre-trained models (PTMs) on downstream tasks. As the model size and the number of tasks increase greatly, such approach becomes less feasible and prohibitively expensive. Recently, a variety of parameter-efficient tuning methods have been proposed in natural language processing (NLP) that only fine-tune a small number of parameters while still attaining strong performance. Yet there has been little effort to explore parameterefficient tuning for IR.

In this work, we first conduct a comprehensive study of existing parameter-efficient tuning methods at both the retrieval and reranking stages. Unlike the promising results in NLP, we find that these methods cannot achieve comparable performance to full finetuning at both stages when updating less than 1% of the original model parameters. More importantly, we find that the existing methods are just parameter-efficient, but not learning-efficient as they suffer from unstable training and slow convergence. To analyze the underlying reason, we conduct a theoretical analysis and show that the separation of the inserted trainable modules makes the optimization difficult. To alleviate this issue, we propose to inject additional modules alongside the pre-trained models (PTMs) to make the original scattered modules connected. In this way, all the trainable modules can form a pathway to smooth the loss surface

CIKM '22, October 17-21, 2022, Atlanta, GA, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9236-5/22/10...\$15.00

https://doi.org/10.1145/3511808.3557445

Xueqi Cheng CAS Key Lab of Network Data Science and Technology, ICT, CAS University of Chinese Academy of Sciences Beijing, China cxq@ict.ac.cn

and thus help stabilize the training process. Experiments at both retrieval and re-ranking stages show that our method outperforms existing parameter-efficient methods significantly, and achieves comparable or even better performance over full fine-tuning.

## **CCS CONCEPTS**

• Information systems → Retrieval models and ranking.

## **KEYWORDS**

Information Retrieval, Dense Retrieval, Parameter-efficient Tuning

#### **ACM Reference Format:**

Xinyu Ma, Jiafeng Guo, Ruqing Zhang, Yixing Fan, and Xueqi Cheng. 2022. Scattered or Connected? An Optimized Parameter-efficient Tuning Approach for Information Retrieval. In Proceedings of the 31st ACM International Conference on Information and Knowledge Management (CIKM '22), October 17-21, 2022, Atlanta, GA, USA. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3511808.3557445

#### **INTRODUCTION** 1

"Pre-training and fine-tuning" has became the prevalent paradigm in the natural language processing (NLP) [1, 31]. The success of Transformer-based pre-trained models (PTMs) in the NLP has also attracted attention in the information retrieval (IR) community [6, 19]. Many researchers have applied the popular PTMs, e.g., BERT [4] and RoBERTa [21], into the multi-stage search pipeline [28, 35], including the first-stage retrieval and the re-ranking stage. The first-stage retrieval aims to return a subset of candidate documents efficiently, and the re-ranking stage attempts to re-rank those candidates accurately. Studies have shown that leveraging the existing PTMs can benefit both the retrieval and re-ranking stages significantly [14, 15, 28, 38].

The mainstream approach to adapt large-scale PTMs to the downstream tasks is via full fine-tuning, which updates all the parameters of the PTMs. Though effective, this fine-tuning approach has drawbacks on its parameter efficiency. Firstly, every downstream task

<sup>\*</sup>Jiafeng Guo is the corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

needs a separate copy of fine-tuned model parameters, containing as many parameters as in the original PTMs. This is prohibitively expensive when serving models that perform a wide range of tasks. Secondly, larger models are usually trained every few months with the ever increasing size ranging from millions [4] to hundreds of billions [5] or even trillions of trainable parameters [7]. As the model size and the number of tasks grow, re-training all model parameters becomes less feasible and raises critical deployment challenges.

To alleviate this issue, a surge of development of parameterefficient tuning methods have been proposed in NLP, which update only a small number of extra parameters while keeping the original PTMs parameters frozen [8, 11, 12, 18, 33, 37]. The representative methods include addition-based such as Adapter [11] and prefixtuning [18], specification-based such as Bitfit [37], and low-rank adaption like LoRA [12]. Most of these methods are injected to the PTMs in an inside manner where the extra tunable modules are scattered in the sub-layers of the Transformer. In essence, the inside modules have a big impact to the final output due to their interaction with the original PTMs. These methods have been reported to achieve comparable performance over full fine-tuning on NLP tasks, with only updating less than 1% of the original parameters.

Yet there has been little effort to adopt parameter-efficient tuning to the IR scenario. The most related work in this direction focused on the re-ranking stage [13], where prefix-tuning and LoRA are delicately leveraged. Their experimental results demonstrate that these two methods generally perform on par with or even outperform the full fine-tuning by tuning less than 1% of the original model parameters. However, the retrieval stage remains less well studied. Besides, the proposed mechanism is designed for the bi-encoder and unsuitable for the cross-encoder, resulting in the limitation of its flexibility. In addition, their experimental results are on small test sets, which may not be representative enough.

In this work, we first conduct a comprehensive study of several representative parameter-efficient tuning methods for both the retrieval and re-ranking stage. The first research question is: can existing methods perform as well in IR as in NLP? The results show that: (1) These methods lag behind full fine-tuning on both stages by tuning less than 1% of orignal parameters, which is different from the findings observed from the small IR datasets [13]. (2) Existing parameter-efficient tuning methods suffer from unstable training and slow convergence. That is, these methods are just parameter-efficient, but not learning-efficient.

This phenomenon raises the second research question: why the standard setup of parameter-efficient tuning methods falls short in IR? To analyze the underlying reason, we conduct a theoretical analysis and find the potential reason is that the separation of the inserted trainable modules results in a discrepancy between the ideal optimization direction and the actual update direction. Specifically, the computation of the optimization direction depends on the parameters of the whole model (including the PTMs and injected modules), while the actual gradient update only performs on the injected modules. Such discrepancy makes the optimization difficult, which may hurt the performance.

The above analysis leads to the third research question: can we design a parameter-efficient tuning approach to stabilize the training process? Inspired by the skip connection [9] in deep learning, we propose to insert extra modules in an aside manner beyond

the inside manner. The key idea is that extra modules injected alongside the PTMs could make the original scattered modules connected. In this way, all the trainable modules can form a pathway to smooth the loss surface and thus help stabilize the training process. In this work, we carefully design three insertion ways of the aside module. By combing the inside and aside modules, our method can well inherit their advantages, i.e., smoothed loss of the aside modules and big impact of the inside modules. Note that our method can combine most of the parameter-efficient methods and is able to serve both the retrieval and re-ranking stage, and both cross-encoder and bi-encoder. Experiments at both retrieval and re-ranking stages show that our method is significantly better than existing parameter-efficient tuning methods. With tuning less than 1% of the original parameters, our method can achieve comparable performance over full fine-tuning. With tuning 6.7% of the original parameters, our method is able to outperform the full fine-tuning on most tasks.

## 2 PRELIMINARY

In this section, we give a brief description of the ranking problem in IR, the Transformer architecture, as well as several representative parameter-efficient tuning methods.

#### 2.1 Problem Statement

To balance the search efficiency and effectiveness, modern search systems typically employ a multi-stage ranking pipeline in practice, including the first-stage retrieval stage and the re-ranking stage [6].

2.1.1 Dense Retrieval. For the retrieval stage, the model needs to recall a small set of documents from a large-scale corpus efficiently. Dense retrieval models usually employ a representation-based architecture (i.e., bi-encoder) to encode queries and documents into low-dimensional representations independently [14, 22]. Simple similarity functions like dot-product are adopted to compute the relevance score with the dense representations.

Without the loss of generality, the retrieval function with the representation-based architecture can be formulated as follows:

$$rel(q,d) = f(\phi_{PTM}(q), \varphi_{PTM}(d)), \tag{1}$$

where  $\phi_{PTM}$  and  $\varphi_{PTM}$  are query and document encoders, and f is the similarity function.

2.1.2 *Re-ranking.* At the re-ranking stage, the interaction-focused model is widely adopted to produce more accurate ranking list [23, 24, 26]. The relevance score is usually computed by a feed-forward neural network at the top of PTMs where queries and documents are concatenated together as the input to the model.

Without loss of generality, the re-ranking function with the interaction-based architecture could be abstracted as:

$$rel(q,d) = f(\eta_{PTM}(q,d))$$
(2)

where  $\eta_{PTM}$  is the interaction function based on PTMs, and f is the scoring function based on the interaction features. Even though the representation-based models can also be applied to the re-ranking stage, studies have shown that they are less effective than the interaction-based models [25, 30].

Scattered or Connected? An Optimized Parameter-efficient Tuning Approach for Information Retrieval



Figure 1: Illustration of a Transformer layer and several representative parameter-efficient tuning methods. Note that MAM Adapter uses a parallel adapter on FFN sub-layer and prefix-tuning on self-attention sub-layer.

## 2.2 Transformer

Transformer is the dominant model architecture for PTMs. Specifically, a Transformer layer [34] contains a self-attention sub-layer, a feed-forward neural network sub-layer, and residual connection followed by layer normalization.

2.2.1 Self-Attention. The input hidden states are firstly transformed to three vectors, i.e., queries, keys, and values, *m* times independently where *m* is the number of heads. Then a dot-product function is applied on queries and keys to compute attention weights for each head, and then a weighted sum operation is performed on the values. Given the hidden state  $h \in \mathbb{R}^{n \times d}$ , the *i*-th attention is computed as:

Attention<sub>i</sub>(**h**) = 
$$\sum_{m} \operatorname{softmax}\left(\frac{W_{i}^{q} \mathbf{h} \cdot W_{i}^{k} \mathbf{h}}{\sqrt{d/m}}\right) W_{i}^{v} \mathbf{h},$$
 (3)

where  $W_i^q$ ,  $W_i^k$ ,  $W_i^v \in \mathbb{R}^{d/m \times m}$  are the learned transformation matrices for queries, keys and values.

Finally, the output of the multi-head attention is computed as a concatenation of the output vectors of all the heads

$$MH(\mathbf{h}) = Concat(Attention_1(\mathbf{h}), \dots, Attention_n(\mathbf{h}))W^o, \quad (4)$$

where  $W^o \in \mathbb{R}^{d \times d}$  is the projection matrix.

2.2.2 *Feed-forward Neural Network.* The feed-forward network is a position-wise fully connected feed-forward network (FFN), which is applied to each position separately and identically,

$$FFN(\mathbf{h}) = ReLU(\mathbf{h}W_1 + b_1)W_2 + b_2 \tag{5}$$

where  $W_1 \in \mathbb{R}^{d \times 4d}$ ,  $W_2 \in \mathbb{R}^{4d \times d}$ ,  $b_1$ , and  $b_2$  are learned bias terms.

Each of the two sub-layers, i.e., the self-attention sub-layer and the FFN sub-layer, employ a residual connection followed by layer normalization (RCLN) to compute the final output

$$RCLN(h) = LayerNorm(SubLayer(h) + h),$$
 (6)

Method	Insertion position	Number of parameters
Bitfit	-	$11 \times d$
Prefix-tuning	attn	$2 \times l \times d$
Adapter	attn/ffn	$4 \times r \times d$
MAM Adapter	attn/ffn	$2 \times r \times d + 2 \times l \times d$
LoRA	attn	$4 \times r \times d$

Table 1: Number of parameters used at each layer for different methods. Note that for Bitfit, there are 8 bias terms in each transformer layer and 1 bias term in embedding layer.

where LayerNorm( $\cdot$ ) is layer normalization and SubLayer representations Eq. (4) and Eq. (5).

## 2.3 Parameter-efficient Tuning Methods

We introduce five representative parameter-efficient tuning methods as illustrated in Figure 1. These methods can be categorized to three groups, i.e., Addition-based, Specification-based and Lowrank adaption.

2.3.1 Addition-based. Addition-based methods introduce extra parameters by inserting small neural modules such as Adapter [11], or trainable tokens such as prefix-tuning [18]. Only these additional parameters are tuned while the original parameters of PTMs are kept frozen. Besides adapter and prefix-tuning, we also consider the recently proposed Mix-And-Match Adapter (MAM Adapter) [8].

Prefix-tuning extends the prompt-tuning [16] by prepending *m* trainable prefix (token) vectors to the keys and values of the self-attention at every layer. In detail, two sets of newly initialized prefix vectors P<sup>k</sup><sub>i</sub>, P<sup>v</sup><sub>i</sub> ∈ ℝ<sup>l×d</sup> are concatenated with the original key vector and value vector in the self-attention:

$$concat(P_i^k, W_i^k \mathbf{h}), concat(P_i^v, W_i^v \mathbf{h}).$$
 (7)

• Adapter injects two small modules after the self-attention sublayer and the FFN sub-layer sequentially. The adapter module consists of a down-projection, an up-projection and a nonlinear function between them.

$$Adapter(\mathbf{h}) = \mathbf{h} + f(\mathbf{h}W_{down})W_{up},\tag{8}$$

where **h** is the output from a sub-layer,  $W_{down} \in \mathbb{R}^{d \times r}, W_{up} \in \mathbb{R}^{r \times d}$ , and *f* is ReLU.

 MAM Adapter adds prefix-tuning in the self-attention (i.e., Eq. 7) and inserts a parallel adapter module at the FFN side:

$$h = Adapter(\mathbf{h}) + FFN(\mathbf{h}) \tag{9}$$

2.3.2 *Specification-based*. Specification-based methods only tune certain parameters in the original model.

 Bitfit [37] is a very simple method that only trains the bias vectors of the original PTMs and keeps the rest frozen.

2.3.3 Low-rank adaptation. This type of method hypothesizes that the change of weights during model optimizing has a low intrinsic rank. Thus, learning a low-rank decomposition matrix for a frozen pre-trained weight matrix can approximate its weight updates, i.e., a fine-tuned pre-trained weight matrix.

 LoRA trains rank decomposition matrices, which is a downproject and a up-projection, for the dense layer to approximate

Table 2: Comparison between full fine-tuning and various parameter-efficient tuning methods using bi-encoder architecture at the retrieval stage. Best results are marked bold. Note that adding 6.7% params (l = 400) for prefix-tuning increases excessive computational cost to document-based tasks which is unacceptable, we thus only experiment with adding 3.6% params.

Method	#Params	MARCO Passage		TREC2019 Passage		MARCO Doc		TREC2019 Doc	
initia di			R@1000	nDCG@10	R@100	MRR@100	R@100	nDCG@10	R@100
Full fine-tuning	100%	0.316	0.949	0.600	0.715	0.312	0.801	0.462	0.409
Bitfit Prefix-tuning Adapter MAM Adapter LoRA	0.09% 0.5% (l=32) 0.5% (r=16) 0.5% (r=16,l=16) 0.5% (r=16)	0.262 0.294 0.304 0.304 0.302	$\begin{array}{c} 0.921 \\ 0.939 \\ 0.941 \\ 0.944 \\ 0.943 \end{array}$	0.562 0.596 <b>0.606</b> <b>0.609</b> <b>0.608</b>	0.677 0.692 0.696 0.712 0.707	$\begin{array}{c} 0.264 \\ 0.266 \\ 0.255 \\ 0.280 \\ 0.271 \end{array}$	0.785 0.782 0.770 0.799 0.794	$\begin{array}{c} 0.437 \\ 0.423 \\ 0.418 \\ 0.458 \\ 0.417 \end{array}$	0.345 0.326 0.370 0.381 0.376
Prefix-tuning Adapter MAM Adapter LoRA	3.6% (l=200) 6.7% (r=200) 6.7% (r=200,l=200) 6.7% (r=200)	$0.304 \\ 0.316 \\ 0.314 \\ 0.316$	0.943 0.946 0.947 0.946	0.580 0.587 <b>0.616</b> 0.597	0.702 0.687 <b>0.720</b> 0.715	0.265 0.270 0.283 0.279	0.775 0.785 0.792 0.794	$\begin{array}{c} 0.395 \\ 0.433 \\ 0.438 \\ 0.417 \end{array}$	$0.376 \\ 0.400 \\ 0.402 \\ 0.379$

Table 3: Comparison between full fine-tuning and various parameter-efficient tuning methods using cross-encoder architecture at the re-ranking stage. Best results are marked bold.

Method	#Params	MARCO Passage		TREC2019 Passage		MARCO Doc		TREC2019 Doc	
"Turuno		MRR@10	MRR@100	nDCG@10	nDCG100	MRR@10	MRR@100	nDCG@10	nDCG@100
Full fine-tuning	100%	0.376	0.383	0.738	0.637	0.404	0.408	0.657	0.536
Bitfit Prefix-tuning Adapter MAM Adapter LoRA	0.09% 0.5% (l=32) 0.5% (r=16) 0.5% (r=16,l=16) 0.5% (r=16)	0.325 0.355 0.366 0.365 0.363	0.334 0.363 0.371 0.373 0.372	0.562 0.705 0.714 0.717 0.720	0.483 0.626 0.626 0.629 0.635	0.364 0.387 0.397 0.390 0.386	0.357 0.381 0.392 0.395 0.392	0.630 0.640 0.653 0.632 0.637	0.531 0.530 0.534 0.531 0.529
Prefix-tuning Adapter MAM Adapter LoRA	3.6% (l=200) 6.7% (r=200) 6.7% (r=200,l=200) 6.7% (r=200)	0.363 0.373 0.369 0.370	0.371 0.381 0.380 0.378	0.722 0.735 0.731 0.730	0.632 0.637 0.633 0.631	0.384 0.402 0.397 0.401	0.389 0.407 0.402 0.396	$0.640 \\ 0.631 \\ 0.630 \\ 0.647$	0.532 0.528 0.528 0.530

the weight updates. Specifically, LoRA adds the low-rank matrices to the query and value projection matrices  $(W^q, W^v)$  in the self-attention. Taking  $W^q$  as an exmaple:

$$\mathbf{h} = \mathbf{h} \cdot W^{q} + \Delta W = \mathbf{h} \cdot W^{q} + s \cdot \mathbf{h} \cdot W_{down} W_{up}, \qquad (10)$$

where *s* is a tunable scalar hyperparameter,  $W_{down} \in \mathbb{R}^{d \times r}$ , and  $W_{up} \in \mathbb{R}^{r \times d}$ .

We also present the number of parameters used by these methods in Table 1. Based on this, we can change the number of tunable prefixes l or the hidden size r of the inserted module to control the total number of tunable parameters for fair comparisons.

#### **3 A COMPREHENSIVE STUDY**

In this section, we conduct a comprehensive study of the above introduced parameter-efficient tuning methods at both the retrieval and re-ranking stages. We first analyze the overall experimental performance of existing methods. Then, we provide some empirical observations and a following theoretical analysis.

#### 3.1 Overall Performance

We conduct experiments on four large-scale standard benchmarks, including MS MARCO passage ranking datasets (MARCO Dev Passage) [27], MS MARCO document ranking datasets (MARCO Dev Doc) [27], TREC 2019 Deep Learning Track passage ranking task (TREC2019 Passage) [3], and TREC 2019 Deep Learning Track document ranking task (TREC2019 Doc) [3]. The detailed experimental setting can be found in Section 5. Table 2 and Table 3 show the results at the retrieval stage and the re-ranking stage, respectively.

For the retrieval stage, we have the following observations: (1) Unlike the promising results in NLP, all representative methods cannot achieve a comparable performance over full fine-tuning with less than 1% of the model parameters on all datasets. Note that our full fine-tuning baseline is strong as we use mulitple negatives for each query in a mini-batch. (2) With tuning 6% of the original model parameters, these methods achieve comparable performance to the full fine-tuning baseline, but still underperform these baselines on the MARCO Passage and TREC2019 Passage. On MARCO Doc and TREC2019 Doc, they are still very lower than full fine-tuning since we train the dense retrieval models with BM25 negatives which is a little weaker. Existing works like ANCE [35] and ADORE [38] always use the checkpoint trained on the MARCO Passage as the starting point for MARCO Doc. But this makes for an unfair comparison since the parameter-efficient tuning methods would have different starting points. We will do further comparisons by training the dense retrieval models with hard negatives in Section 6.3. (3) Among these parameter-efficient tuning methods, Bitfit performs worst, while LoRA, Adapter and MAM Adapter are more effective than prefix-tuning. Prefix-tuning increases computational cost as it prepends additional trainable tokens in the hidden layer.



Figure 2: Top: The retrieval performance of various parameter-efficient tuning methods using different learning rates on MARCO Passage. Bottom: The loss value of full fine-tuning and the two best performing parameter-efficient methods (i.e., Adapter and LoRA) over training steps.

For the re-ranking stage, we can see that: (1) The relative order of different parameter-efficient tuning methods at this stage is quite consistent with that at the retrieval stage. (2) Our finding is not consistent with Jung et al. [13] in which they found that prefixtuning and LoRA are able to outperform the full fine-tuning on small datasets including Robust04 and ClueWeb09, and nonstandard MARCO document ranking dataset with less than 1% of model parameters. In our experiments, with more strong baselines (i.e., training cross-encoder with several negatives in a mini-batch), these parameter-efficient tuning methods cannot outperform the full finetuning on standard large-scale datasets.

## 3.2 Empirical Observation

Besides the above overall performance, we provide some empirical observations about the training and convergence of existing methods. We find that these methods are very sensitive to hyperparameters, such as learning rate. We also noticed that the loss value at the early training stage is very high which seems to be hard to converge. We take the MARCO Passage as an example and other datasets have the same observation. We show the results at the retrieval stage with different learning rates, i.e., ranging from 4e-5 to 1e-4. The detailed experimental setting can be found in Section 5.

As shown in Figure 2, the performance of these parameterefficient tuning methods varies wildly with different learning rates. A low learning rate always performs worse than a high learning rate in terms of both MRR@10 and R@100 metrics. The results implicate that low learning rates may not find a good optimization direction and are difficult to skip the local optima, leading to slower convergence. Then, we take a look at the loss curve of these methods in the early training stage. As shown at the bottom of Figure 2, we can see that the two best performing parameter-efficient methods Adapter and LoRA have a higher loss value compared to full fine-tuning and their loss values fluctuate wildly, ranging from about 50 to 10. The unstable training process and slow convergence indicate that these methods may be essentially hard to optimize.

## 3.3 Theoretical Analysis

We theoretically show why the standard setting of existing parameterefficient tuning methods is not learning-efficient.

For Transformer-based PTMs, each layer contains a multi-head attention layer (MH), a FFN layer and two RCLN functions. As we introduced in Section 2.3, most of the parameter-efficient tuning methods inject modules to the FFN sub-layer and MH sub-layer in an inside manner, e.g., an Adapter following the MH and another following the FFN.

For simplicity, this can be treated as some form of modifications to FFN and MH, i.e.,  $f(MH(\cdot)), g(FFN(\cdot))$ . Thus, these parameter-tuning methods can be formulated as follows:

$$h = RCLN(g(FFN(RCLN(f(MH(x)))))),$$
(11)

where f, g are the inserted module, such as the Adapter module in Eq. 8, the LoRA module in Eq. 10, the prefix module in Eq. 7 and the MAM Adapter module in Eq. 9.

During training, the goal is to minimize the loss over every training example x = (q, d). The model parameters  $\Theta$  of step *t* are optimized by gradient descent methods (GD):

$$\Theta_{t+1} = \Theta_t - \eta \nabla (J(x, y; \Theta_t), \tag{12}$$

where  $\eta$  is the learning rate, y is the label, and J is the loss function. For simplicity, we will omit the complex loss functions used in the ranking task here. According to the chain rule, the gradient on step t in Eq. 11 is computed as follows:

$$\frac{dJ}{dx} = \frac{dJ}{dRCLN_t} \frac{dRCLN_t}{dg_t} \frac{dg_t}{dFFN_t} \frac{dFFN_t}{dRCLN_t} \frac{dRCLN_t}{df_t} \frac{df_t}{dMH_t} \frac{dMH_t}{dx_t}.$$
(13)

Since the parameters of  $RCLN_t$ ,  $FFN_t$  and  $MH_t$  are kept frozen, only  $f_t$ ,  $g_t$  are updated to:

$$g_{t+1} = g_t - \eta \nabla(g) = g_t - \eta \frac{dJ(x, y; \Theta)}{dg} = g_t - \eta \frac{dJ}{dRCLN_t} \frac{dRCLN_t}{dg_t},$$
  

$$f_{t+1} = f_t - \eta \nabla(f) = f_t - \eta \frac{dJ(x, y; \Theta)}{df}$$
  

$$= g_t - \eta \frac{dJ}{dRCLN_t} \frac{dRCLN_t}{dg_t} \frac{dg_t}{dFFN_t} \frac{dFFN_t}{dRCLN_t} \frac{dRCLN_t}{df_t},$$
(14)

As we can see, the gradients of f, g are computed based on the frozen parameters including  $RCLN_t$ ,  $FFN_t$  and  $MH_t$ . The ideal gradient descent direction is  $\nabla(\Theta_t)$  including all parameters, but the actual gradient update direction is only  $\nabla(f_t, g_t)$ ). So there remains a discrepancy between the ideal optimization direction and the actual update direction

$$\delta = \nabla(\Theta_t) - \nabla(f_t, q_t)).$$

The MH (or FFN) may be the main contribution module to the gradient of the  $\nabla(\Theta_t)$ , that is, updating the parameters of MH (or FFN) may greatly decrease the loss value of the input batch. But only  $f_t, g_t$  are updated, so this can explain why the loss value of parameter-efficient tuning methods varies wildly during training, since  $f_t, g_t$  may contribute little to the gradient of the  $\nabla(\Theta_t)$ . Therefore, the separation of these inserted trainable parameters leads to the discrepancy problem which can make the optimization difficult and hurt the performance.



Figure 3: The aside module and three variants with different insertion ways. In this way, all the extra inserted modules can form a pathway. *l* denotes the number of Transformer.

## **4 OUR METHOD**

Our analysis shows that these scattered modules in an inside manner lead to unsmooth transferring of updatable gradients. Inspired by skip connection [9], beyond the inside manner, we propose to inject additional modules alongside the PTMs to create a pathway for updatable gradients. Specifically, in this way, the scattered modules can be directly connected throughout the whole PTMs. Formally, we denote this type of module as the *aside* module and the module which is injected into the model as the *inside* module. Without the effect of the original frozen model parameters, these aside modules can create an unimpeded path to make the updatable gradients flow fluently.

The aside module is denoted as z(x). So, according to Eq. (14),  $\nabla(z)$  is better than  $\nabla(f)$  and  $\nabla(g)$  since its gradient is only based on the final loss J and itself, i.e.,  $\nabla(z) = \frac{dJ}{dz}$ , and don't have to multiply the gradients of RCLN, FFN and MH. In this way, the aside module is updated without the barrier of frozen model parameters and thus mitigates the optimization discrepancy.

Although the inside module suffers from the optimization discrepancy, it's more expressive and has a bigger impact on the final output than the aside module since its output will be transformed by the next following complex modules like MH and FFN. To leverage the merits of these two kinds of modules, we propose to combine the inside module and the aside module for better performance.

**The Inside Module** We can adopt any parameter-efficient tuning methods which inject new parameters into PTMs, as our inside module. In our pilot experiments, we find that Adapter-based and LoRA perform best across all parameter-efficient tuning methods. So in our main experiments, we employ Adapter as our inside module, and we also conduct experiments with LoRA in Section 6.2. We leave the study of choosing or designing the inside module for future work.

**The Aside Module** As shown in Figure 3, our proposed aside module is a bottleneck architecture (BN) containing a down-projection, a nonlinear function and an up-projection. Compared to Adapter, BN has no residual connection, and compared to LoRA, BN adds a nonlinear function. As depicted in Figure 3, we investigate three ways of inserting BN to the model along with the inside modules, i.e., outside the sub-layer, outside the layer, and outside the model.

We denote these three Inside and Aside (IAA) structures as IAA-S, IAA-L, and IAA-M respectively.

- **IAA-S** inserts two BN modules outside the two sub-layer in Transformer, i.e., FFN sub-layer and MH sub-layer. Note that the output of each BN is added to the output of the residual connection and layer normalization. The number of parameters is computed as  $4 \times r \times d \times l$ .
- IAA-L inserts one BN module outside the Transformer layer and there are *l* modules for a whole PTMs. The number of parameters is computed as  $2 \times r \times d \times l$ .
- **IAA-M** inserts one BN modules outside the whole PTMs. It takes the output of the embedding layer and adds its output to the final output of the whole model. The number of parameters is computed as  $2 \times r \times d$ .

In order to fairly compare these structures, we can control the hidden size *r* to keep the number of parameters in each structure the same. As the BN gets farther from the original Transformer, it can have a larger hidden size to have the same number of parameters.

## **5 EXPERIMENTAL SETTINGS**

In this section, we introduce our experimental settings, including datasets, baseline methods, evaluation metrics, and training details.

#### 5.1 Datasets

We conduct our experiments on 4 standard ranking datasets, including MS MARCO passage ranking datasets (MARCO Passage) [27], MS MARCO document ranking datasets (MARCO Doc) [27], TREC 2019 Deep Learning Track passage ranking task (TREC2019 Passage) [3], and TREC 2019 Deep Learning Track document ranking task (TREC2019 Doc) [3]. MARCO Passage contains 0.5 million training queries, 6 thousand dev queries and 8.8 million passages. MARCO Doc contains 0.4 million training queries, 5 thousand dev queries and 3 million documents. For these two MARCO datasets, we report the performance on dev set following existing work [23, 28, 35, 38]. The two TREC2019 datasets share the same training set and document collection with their corresponding MARCO datasets, but they have a fine-grained test set containing 200 queries. Scattered or Connected? An Optimized Parameter-efficient Tuning Approach for Information Retrieval

Table 4: Comparisons between IAA and the baselines at the retrieval stage. Two-tailed t-tests demonstrate the improvements of IAA over baselines are statistically significant ( $p \le 0.05$ ). \* indicate significant improvements over full fine-tuning. † indicate significant improvements over best parameter-efficient tuning methods (PET) at the same setting.

Method	#Params	MARCO Passage		TREC2019 Passage		MARCO Doc		TREC2019 Doc	
	"I didino	MRR@10	R@1000	nDCG@10	R@100	MRR@100	R@100	nDCG@10	R@100
Full fine-tuning	100%	0.316	0.949	0.600	0.715	0.312	0.801	0.462	0.409
Best PET	0.5%	0.304	0.944	0.609	0.712	0.280	0.799	0.458	0.381
IAA-S Adapter	0.5% (r=8,ar=8)	$0.312^{\dagger}$	0.941	0.605	0.719	0.285	0.785	0.454	0.384
IAA-L Adapter	0.5% (r=12,ar=12)	$0.314^{\dagger}$	0.943	$0.615^{+}$	$0.735^{*}$	0.292	0.792	0.446	0.391
IAA-M Adapter	0.5% (r=15,ar=24)	0.309	0.941	0.602	0.721	0.287	0.782	0.449	0.385
Best PET	6.7%	0.316	0.946	0.616	0.720	0.283	0.792	0.438	0.402
IAA-S Adapter	6.7% (r=100,ar=100)	0.324	0.947	0.581	0.719	0.290	0.798	0.441	0.398
IAA-L Adapter	6.7% (r=50,ar=300)	0.327 <sup>†*</sup>	0.951	$0.617^{*}$	$0.735^\dagger$	$0.295^{\dagger}$	0.795	0.439	0.395
IAA-M Adapter	6.7% (r=185,ar=960)	0.321	0.948	0.592	0.710	0.285	0.793	0.437	0.402

## 5.2 Baselines

We use the BERT-base model as the backbone, and use the crossencoder architecture and the bi-encoder architecture for re-ranking and dense retrieval, respectively.

Our baseline includes the full fine-tuning and 5 representative parameter-efficient tuning methods as we introduced in Section 2.3, including BitFit [37], prefix-tuning [18], Adapter [11], MAM Adapter [8] and LoRA [12]. The recent proposed Semi-Siamese method [13] is applied to the prefix-tuning and LoRA in their experiment. The Semi-Siamese prefix-tuning (SS prefix), besides the common prefix, uses some specific prefixes for the query and the document respectively to model their distinct characteristics. The Semi-Siamese LoRA (SS LoRA) use the same query weight matrices and different value weight matrices for the query and the document

## 5.3 Evaluation Metrics

We report the official metrics of these four benchmarks. For the MARCO Passage, we report the Mean Reciprocal Rank at 10 (MRR@10) and recall at 1000 (R@1000). For the MARCO Doc, we report the MRR@100 and R@100. For TREC2019 Passage, we report normalized discounted cumulative gain at 10 (NDCG@10), and R@1000 while for TREC2019 Doc, we report NDCG@10 and R@100.

## 5.4 Training and Optimization

For the cross-encoder model which is used for the re-ranking stage, the query and the document are concatenated into a single sequence to input to the model. We truncate the sequence to the first 128 tokens and 512 tokens for passage datasets and document datasets, respectively. We use cross-entropy pairwise loss and pair 5 negative examples for each query in a mini-batch. We use the official top-k candidates as the negatives. We use a batch size of 72 and 36 for passage datasets and document datasets, respectively. We train 5 epochs for all methods and choose the best checkpoint. The only difference between full fine-tuning and other baselines is that we set different learning rates. For full fine-tuning, we use a learning rate of 2e-5. For all other parameter-efficient tuning methods, we use a learning rate of 1e-4.

For the bi-encoder model which is used for dense retrieval, the query and the document are encoded separately. We set the maximum length of the query to 32, the passage to 128, and the document to 512. We use the official top-k candidates for the passage retrieval task and use BM25 top-k candidates retrieved by anserini [36] for document retrieval task [36]. Training dense retrieval models with official top-k candidates on MARCO Doc results in bad performance. We pair 7 negative examples for each query on passage retrieval and 1 negative example on document retrieval. We use a batch size of 64 and 44 for passage datasets and document datasets, respectively. We train 3 epochs, and 6 epochs for passage datasets and document datasets, respectively. For full fine-tuning, we use a learning rate of 2e-5. For all parameter-efficient tuning methods, we use a learning rate of 1e-4. For all experiments, we use the Adam optimizer with a linear warm-up over the first 10% steps.

## 6 EMPIRICAL RESULTS

In this section, we report and analyze the experimental results to demonstrate the effectiveness of the proposed method. We target the following research questions:

- **RQ1:** How does our method perform compared with full finetuning and other parameter-efficient tuning methods?
- **RQ2:** How does our method perform compared with the Semi-Siamese bi-encoder neural models on the re-ranking stage?
- RQ3: How does our method perform compared with advanced dense retrieval models when training with hard negatives?
- **RQ4**: How does the hidden size of the aside module affect the performance?
- **RQ5:** How does the connected modules affect the optimization process?

## 6.1 Main Results

To answer **RQ1**, we compare three variants of IAA Adapter with full fine-tuning and the best parameter-efficient tuning methods on four standard large-scale datasets. Table 4 and Table 5 show the results at the retrieval stage and the re-ranking stage, respectively.

We first look at the results at the retrieval stage: (1) Our best IAA model with tuning less than 1% of the model parameters achieve a comparable performance over full fine-tuning, and is significantly better than the best PET on some datasets like MARCO Passage. (2) By tuning 6.7% of the model parameters, our best model could outperform the full fine-tuning baseline on two passage retrieval datasets. On MARCO Passage, it's also significantly better than

nificant improvements over best parameter-efficient tuning methods (PET) at the same setting.									
Method	#Params	MARCO Passage		TREC2019 Passage		MARCO Doc		TREC2019 Doc	
Method	"I di di li li	MRR@10	MRR@100	nDCG@10	nDCG100	MRR@10	MRR@100	nDCG@10	nDCG@100
Full fine-tuning	100%	0.376	0.383	0.738	0.637	0.404	0.408	0.657	0.536
Best PET	0.5%	0.366	0.371	0.720	0.635	0.397	0.392	0.653	0.534
IAA-S Adapter	0.5% (r=8,ar=8)	0.371	0.377	$0.731^{\dagger}$	0.632	0.395	0.393	0.655	0.533
IAA-L Adapter IAA-M Adapter	0.5% (r=12,ar=12) 0.5% (r=15,ar=24)	0.373 <sup>†</sup> 0.369	0.379 <sup>†</sup> 0.373	$0.732^{\dagger} \\ 0.725$	0.633 0.630	0.399 0.393	0.403 <sup>†</sup> 0.391	0.656 0.652	0.537 0.531
Best PET	6.7%	0.373	0.381	0.735	0.637	0.402	0.407	0.647	0.530
IAA-S Adapter	6.7% (r=100 ar=100)	$0.382^{\dagger}$	0.385	0.742	0.635	0.408	0.412	0.651	0.535

0.740

0.739

0.639

0.636

0.412

0.404

0.392\*\*

0.384

Table 5: Comparisons between IAA and the baselines on the re-ranking stage. Two-tailed t-tests demonstrate the improvements of IAA over baselines are statistically significant ( $p \le 0.05$ ). \* indicate significant improvements over full fine-tuning. † indicate significant improvements over best parameter-efficient tuning methods (PET) at the same setting.

full fine-tuning in terms of MRR@10. This demonstrates that by introducing the connected aside module, our method is able to improve the performance. On two document retrieval tasks, we find that our methods cannot outperform the full fine-tuning baseline indicting training with BM25 negatives is not enough for bi-encoder on document retrieval. We leave this for further study. (3) Compare the three insertion structures, we find that IAA-L which injects the aside module outside the layer performs best. One possible reason is that IAA-S which injects the aside module outside the sub-layer has a smaller hidden size of the inside module than IAA-L which may limit its capacity. For IAA-M, although it have bigger hidden size for the aside module, its representative power is not as good as IAA-L since the output of each aside module in IAA-L can be transformed by the original parameters.

IAA-L Adapter 6.7% (r=50,ar=300) 0.385\*

IAA-M Adapter 6.7% (r=185,ar=960) 0.379

We then look at the re-ranking stage, and we find that the performance trend on the re-ranking stage is consistent with the retrieval stage: (1) Our method is significantly better than the best parameterefficient tuning methods in terms of MRR@10 on MARCO Passage and MARCO Doc. (2) All types of IAA can outperform the full finetuning baseline by tuning 6.7% of the model parameters, indicating the effectiveness of IAA. (3) Unlike the poor performance on document retrieval tasks, IAA could outperform the full fine-tuning with 6.7% of the model parameters. It demonstrates that applying parameter-efficient tuning methods on cross-encoder perform better than on the bi-encoder.

## 6.2 Comparison with Semi-Siamese Bi-encoder Baseline

To answer **RQ2**, we compare our method with the recently proposed Semi-Siamese methods, i.e., SS prefix-tuning and SS LoRA [13]. These two methods can only apply to the bi-encoder architecture and they leverage this method at the re-ranking stage. As our method is a general method, thus we utilize a bi-encoder architecture on the re-ranking stage for a fair comparison. We also use IAA-S LoRA which uses LoRA as the inside module to compare with SS LoRA. Experiments are conducted on MARCO Passage with only tuning 0.5% of the model parameters. The results are shown in the table 6. We can see the SS prefix-tuning performs worst and this is consistent with our previous findings where prefix-tuning is not as effective as LoRA and Adapter-based. Our methods including

MARCO Passage	MARCO Doc	
MRR@10	MRR@100	
0.342	0.375	
0.351	0.383	
0.366 <sup>†*</sup>	0.391 <sup>†</sup> *	
<b>0.367</b> <sup>†*</sup>	$0.389^{*}$	
	MARCO Passage MRR@10 0.342 0.351 0.366 <sup>†</sup> * <b>0.367</b> <sup>†</sup> *	

0.414

0.410

0.657

0.649

0.538

0.529

Table 6: Performance comparison with Semi-Siamese using a bi-encoder architecture on the re-ranking stage. Two-tailed t-tests demonstrate the improvements are statistically significant (\*,  $\dagger$  indicates  $p \le 0.05$  over SS prefix-tuning and SS LoRA, respectively).

IAA-L LoRA and IAA-L Adapter, are significantly better than the two baseline methods indicating the aside module is more useful and effective for model training.

## 6.3 Comparison with Advanced Dense Retrieval Models by Training with Hard Negatives

To answer **RQ3**, we train the dense retrieval models using hard negatives for the parameter-efficient tuning methods. Following STAR [38], we mine the static hard negatives using BM25 warm-up checkpoint and train the dense retrieval model on hard negatives for another 2-3 epochs. As shown in Table 7, we can observe that by training with hard negatives, parameter-efficient tuning methods achieve comparable performance over some advanced dense retrieval models such as ANCE, and ADORE. And our proposed IAA-L Adapter can still outperform full fine-tuning baseline and is significantly better than other parameter-efficient tuning methods such as Adapter and LoRA. We could see that compared with RocketQA which utilizes several training techniques like crossbatch training, denoising false negatives, and data augmentation, all parameter-efficient tuning methods are still far behind it.

## 6.4 Impact of Hidden Size of the Aside Module

To answer **RQ4**, we conduct an analysis to investigate the impact of the hidden size of the aside model. We experiment on MARCO Passage under the dense retrieval setting. We vary the hidden size of the aside model but still keep the total number of tuning parameters Scattered or Connected? An Optimized Parameter-efficient Tuning Approach for Information Retrieval

Table 7: Comparison with advanced dense retrieval models
by training PET with hard negatives on the MARCO Passage
Best results are marked bold.

Model	MRR@10	R@1000
ANCE[35]	0.330	0.959
TCT-ColBERT[20]	0.335	0.964
TAS-B[10]	0.343	0.976
ADORE+STAR[38]	0.347	-
RoctetQA [32]	0.367	-
full fine-tuning	0.341	0.961
Adapter	0.334	0.953
MAM Adapter	0.332	0.959
LoRA	0.331	0.957
IAA-L Adapter	0.343	0.971

fixed. That is, the larger the aside module, the smaller the inside module and vice versa.

As shown in Figure 4, we can see that different sizes have a big impact on the performance. When the hidden size of the aside model is 0, it degrades to the skip connection which only connects the input and output of the original model. In this setting, IAA-S performs better than IAA-L and IAA-M indicating that a fine-grained skip connection is better than a coarser-grained. When the hidden size of the inside module is 0, it becomes a totally parallel aside module. We can see that the aside module underperforms the inside module with skip connection, i.e., the hidden size of the aside module is 0. This verifies our hypothesis that the inside module is more expressive and has a larger capacity than the aside module. One possible reason is that the output of the inside module will be transformed by the next following complex Transformer modules like multi-head attention.

### 6.5 Convergence Analysis

To answer **RQ5**, we visualize the training loss Adapter and IAA-L Adapter on MARCO Passage at the retrieval stage. As shown in Figure 5, IAA-L Adapter has a lower loss value than Adapter and also converges faster than Adapter. This demonstrates that by adding the aside module, IAA-L Adapter could alleviate the optimization discrepancy problem which is caused by the separation of the trainable modules. One possible reason is that the aside module eases optimization and accelerates training convergence by smoothing the loss surface. This has been verified by the [17] which says skip connections could promote flat minimizers and prevent the transition to chaotic behavior.

## 7 RELATED WORK

In this section, we briefly review the fine-tuning approaches for PTMs in IR. Fully fine-tuning large PTMs like BERT [4] is the widely used approach in IR, since it achieve strong performance at both the retrieval stage [2, 14, 35, 38] and the re-ranking stage [23, 28]. Another approach is the feature-based as used in ELMo [29]. The pre-trained representations input to task-specific architectures as features. CEDR [26] has investigated this approach in several TREC datasets and found the performance of feature-based degrades greatly compared with the fully fine-tuning. Jung et al. [13] firstly



Figure 4: The impact of the hidden size of the aside module.



Figure 5: The loss value over training steps.

apply prefix-tuning and LoRA to the re-ranking stage. They found that the two kinds of parameter-efficient methods can outperform the full fine-tuning on small test data. But with more strong baselines, our findings are not consistent with theirs and we propose a more universal method which can be applied to various architectures and parameter-efficient tuning methods.

## 8 CONCLUSION

In this paper, we conduct comprehensive empirical studies of parameterefficient tuning methods in IR scenarios, at both the retrieval stage and the re-ranking stage. On four standard large-scale benchmarks, we find that these methods are unable to outperform or even achieve a comparable performance over full fine-tuning with tuning less than 1% of original model parameters. Through mathematical analysis, we certify the reason is that the separation of the trainable parameters results in a discrepancy between the ideal optimization direction and the actual update direction. We thus introduce the aside module to help to stabilize the optimization process. Experiments show that our method is significantly better than existing methods and could outperform the full fine-tuning on most tasks by tuning 6.7% of original model parameters. In future work, we would study their ability of domain adaptation in IR.

## ACKNOWLEDGMENTS

This work was funded by the National Natural Science Foundation of China (NSFC) under Grants No. 62006218 and 61902381, the Youth Innovation Promotion Association CAS under Grants No. 20144310, and 2021100, the Young Elite Scientist Sponsorship Program by CAST under Grants No. YESS20200121, and the Lenovo-CAS Joint Lab Youth Scientist Project. CIKM '22, October 17-21, 2022, Atlanta, GA, USA

#### REFERENCES

- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. On the opportunities and risks of foundation models. arXiv preprint arXiv:2108.07258 (2021).
- [2] Jiangui Chen, Ruqing Zhang, J. Guo, Yixing Fan, and Xueqi Cheng. 2022. GERE: Generative Evidence Retrieval for Fact Verification. Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (2022).
- [3] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Fernando Campos, and Ellen M. Voorhees. 2020. Overview of the TREC 2020 Deep Learning Track. *ArXiv* abs/2102.07662 (2020).
- [4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. In Proceedings of the 2019 Conference of the North. Association for Computational Linguistics. https://doi.org/10.18653/v1/n19-1423
- [5] Tom B. Brown .et.al. 2020. Language Models are Few-Shot Learners. In Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual. https://proceedings.neurips.cc/paper/2020/hash/ 1457c0dbfbc4967418bfb8ac142f64a-Abstract.html
- [6] Yixing Fan, Xiaohui Xie, Yinqiong Cai, Jia Chen, Xinyu Ma, Xiangsheng Li, Ruqing Zhang, Jiafeng Guo, and Yiqun Liu. 2021. Pre-training Methods in Information Retrieval. CoRR abs/2111.13853 (2021). arXiv:2111.13853 https://arxiv.org/abs/ 2111.13853
- [7] William Fedus, Barret Zoph, and Noam Shazeer. 2021. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. arXiv preprint arXiv:2101.03961 (2021).
- [8] Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. 2021. Towards a unified view of parameter-efficient transfer learning. *ICLR 2022* (2021).
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition. 770–778.
- [10] Sebastian Hofstätter, Sheng-Chieh Lin, Jheng-Hong Yang, Jimmy Lin, and Allan Hanbury. 2021. Efficiently Teaching an Effective Dense Retriever with Balanced Topic Aware Sampling. In Proceedings of the 44th International ACM SI-GIR Conference on Research and Development in Information Retrieval. ACM. https://doi.org/10.1145/3404835.3462891
- [11] 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*. PMLR, 2790–2799.
- [12] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, and Weizhu Chen. 2022. LoRA: Low-Rank Adaptation of Large Language Models. *ICLR* (2022).
- [13] Euna Jung, Jaekeol Choi, and Wonjong Rhee. 2022. Semi-Siamese Bi-encoder Neural Ranking Model Using Lightweight Fine-Tuning. In Proceedings of the ACM Web Conference 2022. ACM. https://doi.org/10.1145/3485447.3511978
- [14] Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen tau Yih. 2020. Dense Passage Retrieval for Open-Domain Question Answering. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP). Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.emnlp-main.550
- [15] Omar Khattab and Matei Zaharia. 2020. ColBERT: Efficient and Effective Passage Search via Contextualized Late Interaction over BERT. In Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval. ACM. https://doi.org/10.1145/3397271.3401075
- [16] Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The Power of Scale for Parameter-Efficient Prompt Tuning. In Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing. Association for Computational Linguistics. https://doi.org/10.18653/v1/2021.emnlp-main.243
- [17] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. 2018. Visualizing the loss landscape of neural nets. Advances in neural information processing systems 31 (2018).
- [18] 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). Association for Computational Linguistics. https://doi.org/10.18653/v1/2021.acl-long.353
- [19] Jimmy Lin, Rodrigo Nogueira, and Andrew Yates. 2021. Pretrained Transformers for Text Ranking: BERT and Beyond. Synthesis Lectures on Human Language Technologies 14, 4 (Oct. 2021), 1–325. https://doi.org/10.2200/ s01123ed1v01y202108hlt053
- [20] Sheng-Chieh Lin, Jheng-Hong Yang, and Jimmy Lin. 2021. In-Batch Negatives for Knowledge Distillation with Tightly-Coupled Teachers for Dense Retrieval. In Proceedings of the 6th Workshop on Representation Learning for NLP (RepL4NLP-2021). Association for Computational Linguistics. https://doi.org/10.18653/v1/

2021.repl4nlp-1.17

- [21] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692 (2019).
- [22] Xinyu Ma, J. Guo, Ruqing Zhang, Yixing Fan, and Xueqi Cheng. 2022. Pre-train a Discriminative Text Encoder for Dense Retrieval via Contrastive Span Prediction. Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (2022).
- [23] Xinyu Ma, Jiafeng Guo, Ruqing Zhang, Yixing Fan, Xiang Ji, and Xueqi Cheng. 2021. PROP: Pre-training with Representative Words Prediction for Ad-hoc Retrieval. In Proceedings of the 14th ACM International Conference on Web Search and Data Mining. ACM. https://doi.org/10.1145/3437963.3441777
- [24] Xinyu Ma, Jiafeng Guo, Ruqing Zhang, Yixing Fan, Yingyan Li, and Xueqi Cheng. 2021. B-PROP: Bootstrapped Pre-Training with Representative Words Prediction for Ad-Hoc Retrieval. In Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (Virtual Event, Canada) (SIGIR '21). Association for Computing Machinery, New York, NY, USA, 1513–1522. https://doi.org/10.1145/3404835.3462869
- [25] Sean MacAvaney, Franco Maria Nardini, Raffaele Perego, Nicola Tonellotto, Nazli Goharian, and Ophir Frieder. 2020. Efficient Document Re-Ranking for Transformers by Precomputing Term Representations. In Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval. ACM. https://doi.org/10.1145/3397271.3401093
- [26] Sean MacAvaney, Andrew Yates, Arman Cohan, and Nazli Goharian. 2019. CEDR. In Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval. ACM. https://doi.org/10.1145/3331184. 3331317
- [27] Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. MS MARCO: A Human Generated MAchine Reading COmprehension Dataset. In Proceedings of the Workshop on Cognitive Computation: Integrating neural and symbolic approaches 2016 co-located with the 30th Annual Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain, December 9, 2016 (CEUR Workshop Proceedings, Vol. 1773).
- [28] Rodrigo Nogueira and Kyunghyun Cho. 2019. Passage Re-ranking with BERT. arXiv preprint arXiv:1901.04085 (2019).
- [29] Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep Contextualized Word Representations. In Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers). Association for Computational Linguistics. https://doi.org/10.18653/v1/n18-1202
- [30] Yifan Qiao, Chenyan Xiong, Zhenghao Liu, and Zhiyuan Liu. 2019. Understanding the Behaviors of BERT in Ranking. arXiv preprint arXiv:1904.07531 (2019).
- [31] XiPeng Qiu, TianXiang Sun, YiGe Xu, YunFan Shao, Ning Dai, and XuanJing Huang. 2020. Pre-trained models for natural language processing: A survey. *Science China Technological Sciences* 63, 10 (Sept. 2020), 1872–1897. https://doi. org/10.1007/s11431-020-1647-3
- [32] Yingqi Qu, Yuchen Ding, Jing Liu, Kai Liu, Ruiyang Ren, Wayne Xin Zhao, Daxiang Dong, Hua Wu, and Haifeng Wang. 2021. RocketQA: An Optimized Training Approach to Dense Passage Retrieval for Open-Domain Question Answering. In Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. Association for Computational Linguistics. https://doi.org/10.18653/v1/2021.naacl-main.466
- [33] 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. Association for Computational Linguistics. https: //doi.org/10.18653/v1/2021.eacl-main.20
- [34] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. Advances in neural information processing systems 30 (2017).
- [35] Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul N. Bennett, Junaid Ahmed, and Arnold Overwijk. 2021. Approximate Nearest Neighbor Negative Contrastive Learning for Dense Text Retrieval. In 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021. OpenReview.net. https://openreview.net/forum?id=zeFrfgyZln
- [36] Peilin Yang, Hui Fang, and Jimmy Lin. 2017. Anserini. In Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval. ACM. https://doi.org/10.1145/3077136.3080721
- [37] Elad Ben Zaken, Shauli Ravfogel, and Yoav Goldberg. 2022. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. ACL (2022).
- [38] Jingtao Zhan, Jiaxin Mao, Yiqun Liu, Jiafeng Guo, Min Zhang, and Shaoping Ma. 2021. Optimizing Dense Retrieval Model Training with Hard Negatives. In Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval. ACM. https://doi.org/10.1145/3404835. 3462880