

# Delta-LoRA: Fine-Tuning High-Rank Parameters with the Delta of Low-Rank Matrices

Anonymous ACL submission

## Abstract

In this paper, we present **Delta-LoRA**, which is a novel parameter-efficient approach to fine-tune large language models (LLMs). In contrast to LoRA and other low-rank adaptation methods such as AdaLoRA, Delta-LoRA not only updates the low-rank matrices  $A$  and  $B$ , but also propagate the learning to the pre-trained weights  $W$  via updates utilizing the delta of the product of two low-rank matrices ( $A^{(t+1)}B^{(t+1)} - A^{(t)}B^{(t)}$ ). Such a strategy effectively addresses the limitation that the incremental update of low-rank matrices is inadequate for learning representations capable for downstream tasks. Moreover, as the update of  $W$  does not need to compute the gradients of  $W$  and store their momentums, Delta-LoRA shares comparable memory requirements and computational costs with LoRA. Extensive experiments show that Delta-LoRA significantly outperforms existing low-rank adaptation methods. We further support these results with comprehensive analyses that underscore the effectiveness of Delta-LoRA.

## 1 Introduction

Large Language Models (LLMs) recently have attracted considerable attention due to their remarkable performance across a broad spectrum of downstream tasks. Diverging from conventional Transformers characterized by a scale of millions of parameters, modern LLMs typically scale up to billions of parameters, endowing them with notable advantages such as emergent capabilities and robust generalization as detailed in (Bubeck et al., 2023). Fine-tuning such highly capable LLMs on downstream tasks (Raffel et al., 2020; Devlin et al., 2019; Radford et al., 2019; He et al., 2021; Liu et al., 2019; Brown et al., 2020) has consequently become a mainstream paradigm to reduce the training time required for individual tasks, yet with superior performance compared with other methods (Lester et al., 2021; Li and Liang, 2021;

Houlsby et al., 2019; Zhang et al., 2023a).

However, fine-tuning a LLM with all the learnable parameters (Full Fine-tuning) requires multiple GPUs with high memory demand (Dettmers et al., 2023; Hu et al., 2022), which is unattainable for many companies and research institutions. Full fine-tuning poses exceptional challenges to researchers: with massive parameter size, LLMs already demand more storage space than regular models; Further training exaggerates the GPU memory requirement because common optimizers such as AdamW (Loshchilov and Hutter, 2019) often maintain several copies of the model parameters, which is 2-3 times of memory overhead.

To this end, a series of methods have been proposed (Valipour et al., 2023; Zhang et al., 2022; Li and Liang, 2021; Liu et al., 2022a; Lv et al., 2023; Dettmers et al., 2023; Liu et al., 2022b; Zaken et al., 2021; Pfeiffer et al., 2021; Guo et al., 2021; Zhou et al., 2023; Zhang et al., 2023b; Houlsby et al., 2019; Wang et al., 2022; Xu et al., 2024; Guo et al., 2024; Kopiczko et al., 2024; Lialin et al., 2024; Liu et al., 2023a; Li et al., 2024; Wen and Chaudhuri, 2024; Chen et al., 2024; Kopiczko et al., 2023) to reduce memory overhead at the training stage. Some even accelerate the fine-tuning process with only less than 1% trainable parameters. Among these methods, LoRA (Hu et al., 2022) is the most attractive for its stable performance on broad downstream tasks (Ding et al., 2023), no observed overfitting, as well as no extra memory and computation cost at inference.

While LoRA and its successors (Zhang et al., 2022; Valipour et al., 2023) have indeed exhibited superior performance in comparison to alternative approaches within the realm of Parameter Efficient Fine-Tuning (PEFT), a substantial **performance gap** persists when compared to the full fine-tuning, as highlighted in most scenarios (Ding et al., 2023). This discrepancy is attributed to the inherent limitation of updating only a fraction of the model’s

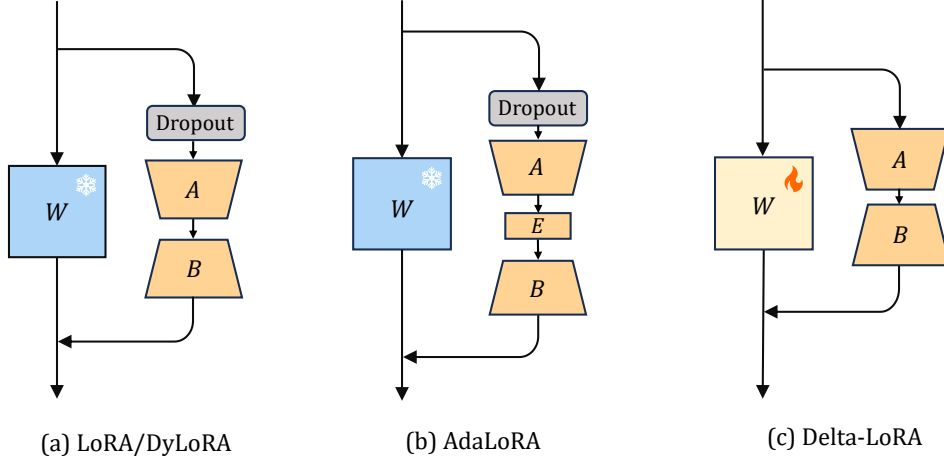


Figure 1: An overview of the proposed **Delta-LoRA** structure, compared to **LoRA**, **DyLoRA** and **AdaLoRA**. Note that **DyLoRA** and **LoRA** basically share the same architecture.  $W$  is the pre-trained weight which is frozen (signified by blue) when performing efficient-parameter fine-tuning in (a) and (b). Orange trapezoids  $A$ ,  $B$  and  $E$  denote the trainable parameters. In our proposed Delta-LoRA, the light orange rectangle means that pre-trained weights can be updated via the delta. Note that our proposed Delta-LoRA removes the Dropout layer to ensure reasonable delta for pre-trained matrix.

parameters, rendering it inadequate to fit the intricacies presented in the training data.

To bridge this gap, a reasonable strategy is to introduce more parameters into the optimization process. In this paper, we introduce Delta-LoRA as shown in Fig. 1, a novel PEFT approach that simultaneously updates the pre-trained matrix and two low-rank matrices while maintaining the same memory consumption as the original LoRA. Specifically, the pre-trained matrix  $W$  is updated with the delta of the product of two low-rank matrices in two consecutive iterations ( $\Delta AB = A^{(t+1)}B^{(t+1)} - A^{(t)}B^{(t)}$ ), while two low-rank matrices are updated by the AdamW optimizer automatically. This is based on the mathematical property that  $\frac{\partial \mathcal{L}}{\partial W} = \frac{\partial \mathcal{L}}{\partial AB}$  and  $\Delta AB$  is a surrogate to direct the update of  $W$  (see Sec. 3 for details). Since we neither store the gradient of  $W$  nor use the optimizer to update the pre-trained matrix, the proposed method thus does not yield any extra memory overhead. This strategic integration effectively mitigates the sub-optimal representation learning stemming from only updating the two low-rank matrices. Moreover, our approach aligns the update direction of the pre-trained weights with that of the incremental update matrix. Furthermore, we discard the Dropout layer in low-rank branches to obtain a more reasonable delta for  $W$ , in order to ensure  $\frac{\partial \mathcal{L}}{\partial W} = \frac{\partial \mathcal{L}}{\partial AB}$ . The advantages of our proposed method are conspicuous: including the pre-trained weights in the optimization pro-

cess engenders a broader integration of parameters, thereby enhancing the potential for learning intricate representations.

The main contributions of this paper can be summarized as:

- We introduce Delta-LoRA, a novel PEFT method that simultaneously updates the full weight matrix and two low-rank matrices. Delta-LoRA leverages the delta of the product of  $A$  and  $B$  to update the pre-trained weights and thus prevent storing the first and the second-order momentums in the optimizer.
- We analyze the gradient flow of Delta-LoRA and show that the Dropout layer in the low-rank branch makes  $\frac{\partial \mathcal{L}}{\partial W} \neq \frac{\partial \mathcal{L}}{\partial AB}$ . Thus, we remove the Dropout layer in our proposed Delta-LoRA to get reasonable delta for  $W$ .
- We conduct comprehensive experiments to show that Delta-LoRA has consistent gains on a broad range of NLP tasks. Additionally, we provide thorough explanations to analyze its superiority and the value contributed by each component.

## 2 Related Works

With the ever-growing parameter scale in current Transformer-based models, fine-tuning such a large language model (LLM) requires considerable number of GPUs equipped with high memory capacity.

This is mainly due to the fact that common optimizers such as AdamW (Loshchilov and Hutter, 2019) requires maintaining three times of extra parameter size (gradients, first-order and second-order momentums). To bridge this gap, a series of Parameter-Efficient Fine-Tuning (PEFT) methods have been proposed (Hu et al., 2022; Liu et al., 2022b; Shin et al., 2020; Houlsby et al., 2019). The Adapter (Houlsby et al., 2019) introduces lightweight trainable parameters between pre-trained layers while keeping the pre-trained weights fixed. Prompt-Tuning (Lester et al., 2021) aims to optimize the prompt to achieve comparable performance with fine-tuning for specific task, while Prefix-Tuning optimizes for trainable prefixes and prepends these trainable parameters to each hidden state (Li and Liang, 2021). Despite the notable performance achievements, these methods inevitably introduce extra overhead at the inference stage.

Hu et al. (2022) proposed LoRA to utilize the multiplication of two low-rank matrices to model the incremental update of a full-rank matrix. LoRA merges the incremental updates to pre-trained weights after training, thereby avoiding any extra computation overhead during inference. Furthermore, it stands out as one of the most effective PEFT techniques according to Ding et al. (2023); Xu et al. (2023)’s evaluation. Subsequent to its inception, a series of enhanced methods building upon LoRA was proposed. Notably, G-LoRA (Chavan et al., 2023) leverages a generalized prompt module to fine-tune pre-trained weights resulting in better representations for computer vision tasks. DyLoRA (Valipour et al., 2023) aims to adjust the rank of two lightweight matrices after the training stage. Differing from the conventional approach of maintaining a static rank during training, DyLoRA introduces rank variations to its blocks. AdaLoRA (Zhang et al., 2022) emphasizes the disparate importance attributed to distinct weight parameters. This technique intelligently allocates the parameter budget across weight matrices based on their respective importance scores. Additionally, Q-LoRA (Dettmers et al., 2023) was proposed to further reduce the average memory footprint by quantizing the pre-trained model with 4-bit NormalFloat. This quantization approach not only preserves the model’s efficacy but also effectively alleviates the resource-intensive nature of LLM training and addresses a pertinent concern.

### 3 Methodology

This section introduces the novel fine-tuning approach termed as Delta-LoRA. Delta-LoRA encompasses two pivotal designs as shown in Figure 1 and Figure 2: (i) It simultaneously updates the full weight matrix ( $\mathbf{W}$ ) alongside the two low-rank adaptation matrices ( $\mathbf{A}$  and  $\mathbf{B}$ ), utilizing the delta ( $\mathbf{A}^{(t+1)}\mathbf{B}^{(t+1)} - \mathbf{A}^{(t)}\mathbf{B}^{(t)}$ ) resulting from incremental updates to refine the pre-trained weights ( $\mathbf{W}$ ); (ii) The Dropout layer as originally integrated within the conventional LoRA module, is excluded in Delta-LoRA. This omission stems from the realization that its presence violates the required assumption  $\frac{\partial \mathcal{L}}{\partial \mathbf{W}} = \frac{\partial \mathcal{L}}{\partial \mathbf{AB}}$ .

#### 3.1 Update the Delta of Low-rank Matrices on Pre-trained Weights

For an input  $x$  and its corresponding hidden state  $\mathbf{h}$ , LoRA optimizes two low-rank matrices  $\mathbf{A}$  and  $\mathbf{B}$  to learn an incremental update  $\mathbf{AB}$  for the pre-trained and fixed weight matrix  $\mathbf{W}$ . Different from previous methods, we argue that  $\mathbf{W}$  also needs to be updated. In this way, we can introduce more learnable parameters to the optimization process for higher learning capability. However, acquiring the normalized gradients (i.e. the gradients after normalization in optimizer) to fine-tune the weight matrix  $\mathbf{W}$  is non-trivial, since the optimizer such as AdamW must maintain at least three extra copies of the parameters (i.e. gradients as well as the first-order and the second-order moments of gradients) in GPU memory. Intriguingly, we note that the gradients of the loss  $\mathcal{L}$  with respect to matrices  $\mathbf{AB}$  and  $\mathbf{W}$  are precisely identical, under the presumption that the LoRA module exclusively retains matrices  $\mathbf{A}$  and  $\mathbf{B}$ , while disregarding the Dropout layer. This correspondence can be formally represented as:

$$\begin{aligned} \mathbf{g}_{\mathbf{W}} &= \frac{\partial \mathcal{L}}{\partial \mathbf{h}_{i+1}} \cdot \frac{\partial \mathbf{h}_{i+1}^\top}{\partial \mathbf{W}} = \frac{\partial \mathcal{L}}{\partial \mathbf{h}_{i+1}} \cdot \mathbf{h}_i^\top, \\ \mathbf{g}_{\mathbf{AB}} &= \frac{\partial \mathcal{L}}{\partial \mathbf{h}_{i+1}} \cdot \frac{\partial \mathbf{h}_{i+1}^\top}{\partial \mathbf{AB}} = \frac{\partial \mathcal{L}}{\partial \mathbf{h}_{i+1}} \cdot \mathbf{h}_i^\top, \quad (1) \\ &\implies \mathbf{g}_{\mathbf{W}} = \mathbf{g}_{\mathbf{AB}}, \end{aligned}$$

where  $\mathbf{h}_{i+1} = \mathbf{W}\mathbf{h}_i + \mathbf{AB}\mathbf{h}_i$ ,  $\mathbf{h}_i$  and  $\mathbf{h}_{i+1}$  are the outputs of the  $i$ -th layer and the  $i+1$ -th layer respectively.  $\mathbf{AB}$  is the matrix product of the adaptation matrices  $\mathbf{A}$  and  $\mathbf{B}$ ,  $\mathcal{L}$  is the loss function, while  $\mathbf{g}_{\mathbf{W}+\mathbf{AB}}$ ,  $\mathbf{g}_{\mathbf{W}}$  and  $\mathbf{g}_{\mathbf{AB}}$  denote the gradients of  $\frac{\partial \mathcal{L}}{\partial (\mathbf{W}+\mathbf{AB})}$ ,  $\frac{\partial \mathcal{L}}{\partial \mathbf{W}}$ , and  $\frac{\partial \mathcal{L}}{\partial \mathbf{AB}}$  respectively.

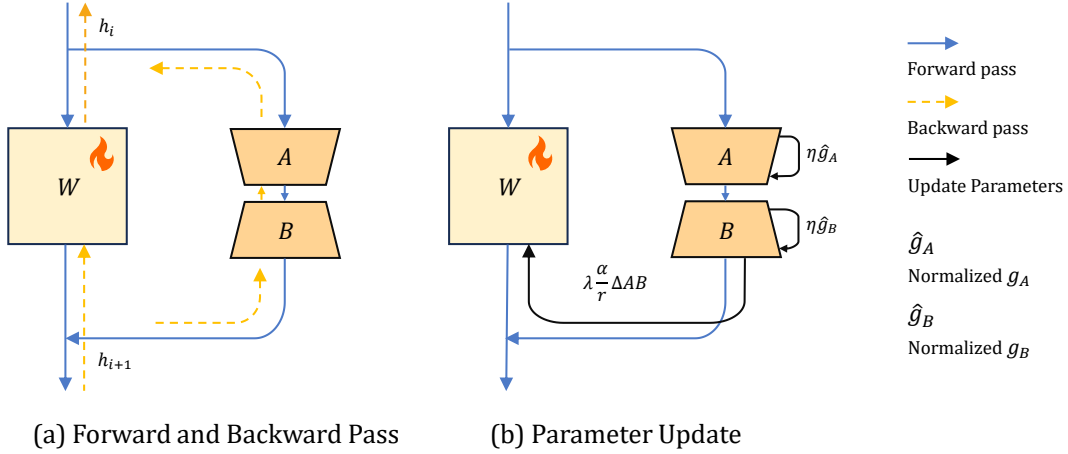


Figure 2: The framework of our proposed Delta-LoRA. The blue arrows represent forward pass while yellow dashed arrows denote backward propagation. The black solid arrows in (b) represent the process of updating the low-rank adaptation matrices  $A$  and  $B$  with normalized gradients  $\hat{g}_A$  and  $\hat{g}_B$  multiplied by the learning rate  $\eta$ , as well as updating the pre-trained weights  $W$  with the delta matrix  $\Delta AB$  multiplied by the update ratio  $\lambda$ .

Equation 1 inspires us to use  $g_{AB}$  to assimilate  $g_W$  when learning the parameter updates for weight matrix  $W$ . Unfortunately, we are only able to obtain the gradients  $g_A$  and  $g_B$  rather than  $g_W$  during the back-propagation process. Furthermore, the computation of the gradients for  $AB$  is as expensive as for the matrix  $W$ , since both matrices share the same dimensions of  $d \times k$ , consequently entailing an equivalent GPU memory overhead.

Considering a typical optimization process, the model updates its parameters by applying the gradient descent:  $W^{(t+1)} = W^{(t)} - \eta g_W$ , with the parameter update denoted as  $\Delta W = -\eta g_W$ , using the learning rate  $\eta$ . Similarly, we regard  $-\Delta AB$  as the gradients for  $AB$  and utilize this matrix as a substitute for  $g_W$  according to Equation 1. Here, we can compute  $\Delta AB$  as:

$$\begin{aligned} \Delta AB &= A^{(t+1)} B^{(t+1)} - A^{(t)} B^{(t)} \\ &= \eta A^{(t)} g_B + \eta g_A B^{(t)} - \eta^2 g_A g_B, \end{aligned} \quad (2)$$

where  $A^{(t)}$ ,  $B^{(t)}$  and  $W^{(t)}$  are the weights of  $A$ ,  $B$  and  $W$  at the  $t$ -th step respectively,  $A^{(t+1)} = A^{(t)} - \eta g_A$ ,  $B^{(t+1)} = B^{(t)} - \eta g_B$  and  $\eta$  is the learning rate. To be precise,  $-\Delta AB$  does not equate directly to  $g_{AB}$  and  $g_W$  as elaborated in Appendix A.4.2. Nonetheless,  $\Delta AB$  has the capability to symbolize the genuine directions of update for the matrix  $AB$ . Based on this assumption, it is reasonable to employ  $-\Delta AB$  as the gradient for directing the update of  $W$ .

Therefore, during the training phase we introduce the matrix  $\Delta AB$  to update the pre-trained

weights  $W$  in the following manner:

$$\begin{aligned} W^{(t+1)} &= W^{(t)} + \lambda \cdot \frac{\alpha}{r} \cdot \Delta AB, \\ \text{where } \Delta AB &= A^{(t+1)} B^{(t+1)} - A^{(t)} B^{(t)}, \end{aligned} \quad (3)$$

where  $\lambda$  represents the hyper-parameter to trade off the update ratio of  $AB$  and the pre-trained weights  $W$ . The parameter updates for  $W$  commence after  $K$  training iterations. The procedural details of the algorithm are illustrated in Algorithm 1.

### 3.2 The structure of our Delta-LoRA

Both LoRA and its successor AdaLoRA put a Dropout layer before two low-rank matrices  $A$  and  $B$ . However, this arrangement results in a disparity between the gradient matrices  $g_W$  and  $g_{AB}$  (or the matrix  $g_{ABE}$  in the context of AdaLoRA). The derivation of this disparity can be shown as:

$$g_W = \frac{\partial \mathcal{L}}{\partial h_{i+1}} \cdot h_i^\top \neq g_{AB} = \frac{\partial \mathcal{L}}{\partial h_{i+1}} \cdot \text{Drop}(h_i)^\top, \quad (4)$$

where  $\text{Drop}(\cdot)$  denotes the Dropout layer which leads to  $g_W \neq g_{AB}$ . A reasonable choice is to remove the Dropout layer in the low-rank module and activate the Dropout layer between pre-trained layers if overfitting problem occurs. This modification also brings additional benefits: (1) it can alleviate under-fitting to some extent, thereby enhancing the learned representations of the networks. The rationale behind this improvement lies in the fact that LoRA and its successors formulate low-rank updates for pre-trained weights, involving less than



Table 1: The evaluation results of our proposed Delta-LoRA and other existing methods on E2E NLG Challenge dataset. † indicates fine-tuning all layers except embedding layer. ‡ indicates only fine-tuning weights for query and value. ¶ means we choose different settings with AdaLoRA: we only tune  $W_Q$  and  $W_V$  instead of all layers. The best results of Fine-Tuning methods are underlined. The best results of PEFT methods are **boldfaced**.

Method	Trainable Parameters	Extra Updatable Parameters	BLEU	NIST	METEOR	ROUGE-L	CIDEr
Full Fine-Tuning	354.92M	✗	69.70±0.20	8.81±0.05	46.21±0.17	71.88±0.30	2.45±0.03
Fine-Tuning†	305.84M	✗	69.54±0.24	8.79±0.03	46.12±0.06	71.65±0.28	2.42±0.02
Fine-Tuning‡	48M	✗	<u>69.81±0.09</u>	<u>8.82±0.02</u>	46.16±0.17	71.85±0.10	2.44±0.04
LoRA (repr.)	0.375M	✗	69.13±0.49	8.79±0.01	45.82±0.30	71.41±0.30	2.40±0.05
DyLoRA	0.375M	✗	67.53±0.39	8.57±0.07	45.09±0.89	70.79±0.25	2.32±0.06
AdaLoRA¶	0.375M	✗	68.47±0.38	8.69±0.10	45.40±1.13	70.81±0.15	2.39±0.04
Delta-LoRA (Ours)	0.375M	✓ 48M	<b>70.90±0.09</b>	<b>8.91±0.02</b>	<b>46.42±0.09</b>	<b>72.11±0.11</b>	<b>2.53±0.01</b>

Table 2: The evaluation results of our proposed Delta-LoRA and other existing methods on WebNLG Challenge 2017 dataset. † indicates fine-tuning all layers except embedding layer. ‡ indicates only fine-tuning weights for query and value. ¶ means we choose different settings with AdaLoRA: we only tune  $W_Q$  and  $W_V$  instead of all layers. The best results of Fine-Tuning methods are underlined. The best results of PEFT methods are **boldfaced**.

Method	Trainable Parameters	Extra Updatable Parameters	BLEU↑			METEOR↑			TER↓		
			S	U	A	S	U	A	S	U	A
Full Fine-Tuning	354.92M	✗	61.35±0.15	45.64±0.64	54.38±0.14	0.44±0.00	0.38±0.00	0.41±0.00	0.36±0.01	0.50±0.03	0.41±0.01
Fine-Tuning†	305.84M	✗	63.20±0.47	47.17±0.72	55.94±0.03	0.45±0.00	<u>0.39±0.00</u>	0.42±0.00	0.34±0.00	0.49±0.01	0.41±0.01
Fine-Tuning‡	48M	✗	<u>64.40±0.21</u>	<u>47.81±0.35</u>	<u>56.66±0.59</u>	<u>0.45±0.01</u>	0.38±0.01	<u>0.43±0.01</u>	<u>0.34±0.01</u>	<u>0.47±0.01</u>	<u>0.40±0.00</u>
LoRA (repr.)	0.375M	✗	61.78±0.52	47.01±0.70	55.07±0.03	0.44±0.00	0.38±0.01	0.41±0.01	0.35±0.01	0.49±0.01	0.41±0.01
DyLoRA	0.375M	✗	59.60±1.10	45.57±0.40	53.25±0.44	0.43±0.01	0.37±0.01	0.40±0.01	0.37±0.01	0.48±0.01	0.42±0.01
AdaLoRA¶	0.375M	✗	57.37±1.38	44.26±0.17	51.42±0.85	0.42±0.01	0.37±0.00	0.40±0.01	0.39±0.01	0.49±0.00	0.44±0.01
Delta-LoRA (Ours)	0.375M	✓ 48M	<b>62.56±0.45</b>	<b>47.94±0.36</b>	<b>55.99±0.04</b>	<b>0.46±0.01</b>	<b>0.40±0.01</b>	<b>0.43±0.01</b>	<b>0.34±0.01</b>	<b>0.48±0.01</b>	<b>0.40±0.01</b>

1% of the complete parameters. However, relying solely on such a small fraction of parameters may not bestow an adequate representation capacity in most cases; (2) This alteration also yields memory-saving benefits. By negating the requirement to store intermediate features, the model curtails the memory consumption. Consequently, there is a reduction in activation memory employed during the back-propagation process.

## 4 Experiments

We evaluate our proposed model fine-tuning method Delta-LoRA with RoBERTa (Liu et al., 2019), GPT-2 (Radford et al., 2019) and BART (Lewis et al., 2019) on a broad set of datasets. Specifically, we train (1) RoBERTa on GLUE benchmark which consists of 8 NLP understanding tasks; (2) GPT-2 on E2E Challenge and WebNLG Challenge 2017 following the setting of Hu et al. (2022); and (3) BART on XSum dataset by using the setting provided by Zhang et al. (2022). See Appendix A.6 for more training details on the datasets. The setups and detailed introductions of baseline methods are shown in Appendix

A.1. We use *PyTorch* to implement our experiments and download the pre-trained weights as well as configuration files from *HuggingFace* (Wolf et al., 2019).

### 4.1 Natural Language Generation

**Models and Datasets.** We use GPT2-Medium to verify the effectiveness of our Delta-LoRA on two datasets for data-to-text tasks, including the E2E NLG Challenge (Puzikov and Gurevych, 2018) and WebNLG Challenge 2017 (Gardent et al., 2017). The E2E NLG Challenge dataset comprises 42,000 samples for training, 4,600 for validation, and 4,600 for testing purposes. In contrast, the WebNLG Challenge 2017 consists of 21,855 training samples across nine categories, expanding to a total of 14 categories in the test set. For the text summarization task, we employed BART-Large (Lewis et al., 2019) to evaluate the efficacy of our method using the XSum dataset (Narayan et al., 2018). The XSum dataset is composed of 204,045 training samples, 11,332 validation samples, and 11,332 test samples. We also use LLaMA-7B (Touvron et al., 2023), a popular pre-trained

Table 3: The evaluation results of our proposed Delta-LoRA and other existing methods on GLUE benchmark. We report the overall (matched and mismatched) accuracy for MNLI, Matthew’s correlation for CoLA, Pearson correlation for STS-B, and accuracy for other tasks. † indicates fine-tuning all layers except the embedding layer. ‡ indicates only fine-tuning weights for query and value. ¶ means we choose different settings with AdaLoRA: we only tune  $W_Q$  and  $W_V$  instead of all layers. The best results of Fine-Tuning methods are underlined. The best results of PEFT methods are **boldfaced**.

Method	Trainable Parameters	Extra Updatable Parameters	MNLI	SST-2	MRPC	CoLA	QNLI	QQP	RTE	STS-B	AVG
Full Fine-Tuning	118.87M	✗	87.41±0.10	94.07±0.17	87.99±0.25	64.42±0.30	92.87±0.14	91.23±0.64	84.71±0.55	90.75±0.19	86.68
Fine-Tuning†	82.05M	✗	<u>87.61±0.03</u>	94.22±0.18	<u>89.87±0.37</u>	62.90±0.27	92.72±0.16	<u>91.62±0.54</u>	<u>86.28±0.36</u>	90.64±0.39	<u>86.98</u>
Fine-Tuning‡	13.5M	✗	87.34±0.26	<u>95.02±0.06</u>	89.05±0.14	61.00±0.17	92.64±0.12	91.36±0.47	85.19±0.36	90.08±0.71	86.46
LoRA	0.28M	✗	87.32±0.10	94.61±0.12	89.70±0.65	63.28±0.75	92.91±0.10	90.63±0.05	86.28±0.36	91.51±0.06	87.03
DyLoRA	0.28M	✗	86.60±0.27	94.34±0.24	88.89±0.51	61.51±0.36	92.36±0.16	90.37±0.33	84.35±0.55	91.24±0.20	86.20
AdaLoRA¶	0.28M	✗	87.35±0.04	94.49±0.12	90.11±0.14	61.68±0.37	92.98±0.12	90.24±0.13	85.80±0.56	91.31±0.16	86.74
Delta-LoRA (Ours)	0.28M	✓13.5M	<b>87.62±0.21</b>	<b>95.29±0.23</b>	<b>90.60±0.14</b>	<b>64.64±0.86</b>	<b>93.09±0.15</b>	<b>91.01±0.06</b>	<b>87.00±0.36</b>	<b>91.61±0.04</b>	<b>87.60</b>

large language model with 7 Billion parameters, to fine-tune on Alpaca dataset (Taori et al., 2023).

Table 4: The evaluation results of Delta-LoRA with LLaMA-7B on the Instruction-Tuning dataset provided by Stanford Alpaca (Taori et al., 2023). We use GPT-4 to choose from *a. LoRA*, *b. Delta-LoRA* or *c. Both LoRA and Delta-LoRA* to decide the text from which method is better.

Both	LoRA	Delta-LoRA	Total
886	10	104	1,000

**Implementation Details.** In order to compare with LoRA and its successors fairly, we adopt the model setups from LoRA to implement our Delta-LoRA and three PEFT methods. We only learn the low-rank incremental update for  $W_Q$  and  $W_V$  in MHA module. For data-to-text datasets, we use the same training configurations as adopted by LoRA, including the number of training epochs, batch size and etc. We use update ratio  $\lambda = 2$  and set start steps  $K = 500$  for Delta-LoRA. More details about Delta-LoRA are listed in the Appendix A.6. For the text-summarization task, we use the implementation of AdaLoRA and adopt the same training configurations. We set the update ratio  $\lambda = 0.5$  and the start steps  $K = 1000$  for Delta-LoRA.

**Experimental Results.** Table 1 shows the results for E2E Challenge dataset on 5 evaluation metrics, demonstrating that our method achieves state-of-the-art performance over 3 baselines and a set of fine-tuning methods. For the BLEU and ROUGE-L metrics, our method obtains 1.77 and 0.7 performance gains compared with LoRA, with 0.12, 0.6 and 0.13 improvement on NIST, METEOR and CIDEr respectively. Table 2 demonstrates that Delta-LoRA outperforms baselines on BLEU score for WebNLG Challenge 2017 dataset, with

0.78, 0.93 and 0.92 improvement on Seen, Unseen and All test data, respectively. Additionally, for the METEOR and TER evaluation metrics, Delta-LoRA also achieves state-of-the-art performance, with 0.02 and 0.01 improvement over LoRA on all data. For the text-summarization task, the test results are shown in Table 5, which demonstrates that our method achieves state-of-the-art results across 3 parameter-efficient methods on 4 evaluation metrics. To fairly evaluate our method, we utilized LLaMA-7B and compared it with LoRA. We employed GPT-4 to generate 1,000 questions and presented these questions to the parameter-efficient fine-tuned LLaMA-7B. Subsequently, we leveraged GPT-4 to compare the texts generated by LoRA-tuned and Delta-LoRA-tuned LLMs. Additional details can be found in Appendix A.2. According to the findings in Table 4, Delta-LoRA establishes state-of-the-art performance in the evaluation of Language Models (LLMs). GPT-4 identified 104 samples generated by Delta-LoRA as superior to LoRA, while only 10 samples generated by LoRA exhibited higher quality than Delta-LoRA. This underscores Delta-LoRA’s effectiveness even when utilized within models containing billions of parameters.

## 4.2 Natural Language Understanding

**Models and Datasets.** We use RoBERTa-base (Liu et al., 2019) to evaluate the performance of our proposed method, prior works and three fine-tuning methods. We choose the GLUE benchmark consisting of 8 datasets (Wang et al., 2019), including classification tasks, similarity and paraphrase tasks and natural language inference tasks.

**Implementation Details.** We use RoBERTa-base with 118M parameters to conduct our experiments and to compare our method with the base-

Table 5: The evaluation results of our proposed Delta-LoRA and other existing methods on XSum dataset. † indicates fine-tuning all layers except the embedding layer. ‡ indicates only fine-tuning weights for query and value. ¶ means we choose different settings with AdaLoRA: we only tune  $W_Q$  and  $W_V$  instead of all layers. The best results of Fine-Tuning methods are underlined. The best results of PEFT methods are **boldfaced**.

Method	Trainable Parameters	Extra Updatable Parameters	Rouge-1	Rouge-2	Rouge-L	Rouge-Sum
Full Fine-Tuning	387.5M	✗	<u>45.18±0.17</u>	<u>22.08±0.08</u>	<u>37.19±0.05</u>	<u>37.2±0.05</u>
Fine-Tuning†	338.4M	✗	44.87±0.23	21.94±0.16	36.74±0.20	36.76±0.20
Fine-Tuning‡	72M	✗	44.67±0.27	21.29±0.14	36.23±0.12	36.24±0.13
LoRA	0.56M	✗	43.17±0.12	20.01±0.14	34.93±0.20	34.93±0.20
DyLoRA	0.56M	✗	42.08±0.32	19.01±0.27	33.83±0.32	33.85±0.33
AdaLoRA¶	0.56M	✗	42.9±0.03	19.75±0.07	34.66±0.07	34.66±0.07
Delta-LoRA (Ours)	0.56M	✓72M	<b>43.45±0.05</b>	<b>20.21±0.12</b>	<b>35.21±0.05</b>	<b>35.21±0.05</b>

lines. We mostly adopt the same training configurations of LoRA, more details can get from Appendix A.6. We set the rank to 8 and the target rank to 6 for AdaLoRA and choose the rest of hyper-parameters according to the characteristics of different tasks. For Delta-LoRA, we set the update ratio  $\lambda$  to 0.5 and choose different start steps  $K$  according to warmup steps used in individual tasks.

**Experimental Results.** We compare our method with prior PEFT works. According to Table 3, our method outperforms existing methods on all 8 tasks in GLUE benchmark. Among these tasks, our method demonstrates significant improvement on SST-2, CoLA and RTE. This is mainly due to the fact that these datasets contain less training data, which hinders the model’s capacity to effectively acquire a robust representation when using prior fine-tuning methods. Delta-LoRA also achieves decent performance on the rest of the datasets, including MNLI, MRPC, QNLI as well STS-B, which proves that our method is stable and reliable across different settings.

### 4.3 Comprehensive Understanding of Delta-LoRA

**The Extra Updatable Parameters.** We introduce the concept of extra updatable parameters to point out the superiority of Delta-LoRA. For most PEFT methods, they can only adjust the low-rank adapters, such as Adapter(Houlsby et al., 2019) and LoRA(Hu et al., 2022). Thus, they don’t have any extra parameters to update, which means their extra updatable parameters are 0. However, our Delta-LoRA can achieve the purpose of updating the  $W$  matrix without increasing the GPU memory consumption, which means its extra updatable parameters are the parameter number of  $W$ .

**Ablation study.** To better understand the contri-

buton of our modified LoRA module (i.e. Delta-LoRA module) and the effectiveness of our update algorithm, we conduct studies on E2E Challenge dataset with GPT2-medium. As shown in Table 6, only updating the pre-trained matrices with delta of low-rank update can indeed achieve performance improvement, while further discarding the dropout in Delta-LoRA module obtains the best performance. This observation confirms the indispensable role played by each component within our proposed methodology. We have devised an experiment to further differentiate whether the performance enhancement stems from the inherent characteristics of our method rather than solely from the substantial update magnitude. According to our algorithm, we update the parameters of both pre-trained and low-rank matrices, which can arise the doubt of whether the improvement is caused by updating larger  $\Delta AB$  on the weights instead of introducing more parameters into the optimization process. To answer this question, we design an experiment with results shown in Table 7 to prove the effectiveness of our method. We scale the learning rate of LoRA from  $2e-4$  to  $6e-4$  making sure that  $W + AB$  can be updated with  $3 \times \Delta AB$ , which is equivalent to Delta-LoRA when  $\lambda$  is set to 2. We find that even by updating with  $3 \times \Delta AB$  on  $AB$ , the performance is still not comparable with Delta-LoRA. This experiment further proves that introducing more parameters into the optimization process can force to learn better representation.

**The cosine similarity between fine-tuned and the pre-trained parameters to measure learning effects.** We conduct a comparative analysis of three methods including Fine-Tuning‡, LoRA and Delta-LoRA, in order to elucidate the reasons behind Delta-LoRA’s superior performance. We use the last checkpoint trained on E2E Challenge dataset

Table 6: The ablation study of our proposed Delta-LoRA on E2E Challenge dataset demonstrates the importance of each component. The best results are **boldfaced**.

Method	Trainable Parameters	Extra Updatable Parameters	BLEU	NIST	METEOR	ROUGE-L	CIDEr
LoRA (repr.)	0.375M	✗	69.60	8.78	45.61	71.12	2.45
Delta-LoRA + LoRA Module	0.375M	✓48M	70.29	8.88	46.38	71.88	2.51
Delta-LoRA	0.375M	✓48M	<b>70.84</b>	<b>8.91</b>	<b>46.47</b>	<b>72.24</b>	<b>2.53</b>

Table 7: The ablation study of our proposed Delta-LoRA to eliminate the impact of hyper-parameter  $\lambda$  on E2E Challenge dataset. The best results are **boldfaced**.

Method	Learning Rate	$\lambda$	BLEU	NIST	METEOR	ROUGE-L	CIDEr
LoRA (repr.)	2e-4	-	69.60	8.78	45.61	71.12	2.45
LoRA (repr.)	6e-4	-	69.63	8.79	45.70	71.55	2.39
Delta-LoRA	2e-4	2	<b>70.84</b>	<b>8.91</b>	<b>46.47</b>	<b>72.24</b>	<b>2.53</b>

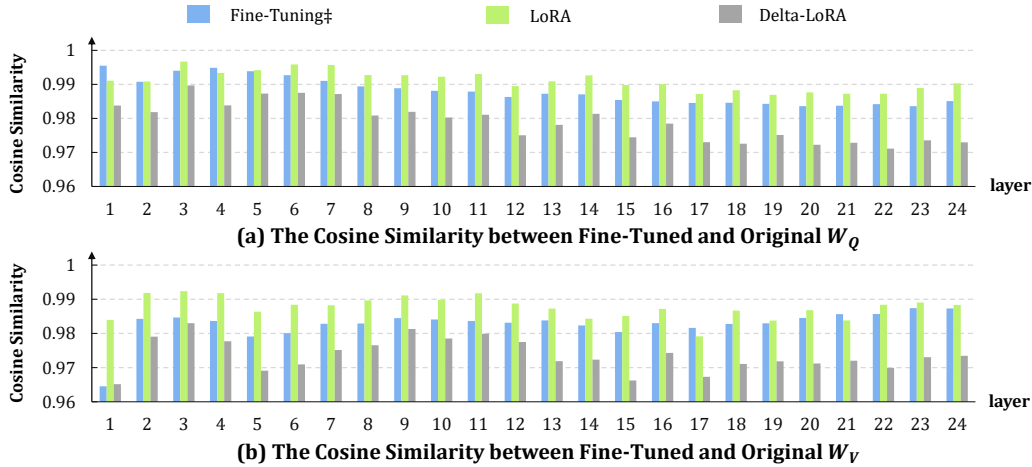


Figure 3: The comparison of Fine-Tuning $\ddagger$ , LoRA as well as Delta-LoRA for the cosine similarity between the fine-tuned parameters and the original pre-trained parameters in each transformer block. *Higher value means higher similarity.*

to give understanding. As depicted in Figure 3, it is evident that LoRA exhibits the highest similarity across the majority of transformer blocks. This observation suggests that LoRA primarily modifies the matrix  $W^* = W + AB$  within a limited range. Nevertheless, Delta-LoRA showcases the lowest cosine similarity, underscoring that our approach induces the most significant modifications to the final matrix  $W^*$ . Due to this property, our approach can effectively stimulate the model to acquire better representations, leading to state-of-the-art performance across all four PEFT methods. This observation further aligns with the evaluation results in Table 1: Delta-LoRA achieves the best performance among the three methods, whereas LoRA is slightly worse than Fine-Tuning $\ddagger$ .

## 5 Conclusion

In this paper, we have introduced Delta-LoRA, a novel method to simultaneously update the full weight matrix and two low-rank matrices. Delta-LoRA leverages the delta ( $A^{(t+1)}B^{(t+1)} - A^{(t)}B^{(t)}$ ) to update the pre-trained weights ( $W$ ). In this way, we introduce more learnable parameters into the optimization process such that the model can learn a better representation with comparable memory cost as LoRA. Meanwhile, we identify the Dropout layer in the low-rank branch to be unnecessary according to the gradient flow. We also provide thorough analysis of our method to understand its effectiveness and robustness. Extensive experiments on a broad range of NLP tasks are conducted to empirically verify the effectiveness of our Delta-LoRA.



## 6 Limitations

This work has several limitations. First, it involves a two-step process to compose the delta of two low-rank matrices into a gradient matrix, which is somewhat time-consuming. Second, since this work was conducted in 2023, we could not compare it with subsequent advancements in the field. Third, our evaluation was limited to GPT-2 Medium, BART-Large, RoBERTa-Base, and LLaMA-7B models. Future work will focus on experimenting with other large language models.

## References

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *NeurIPS*.

Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. 2023. *Sparks of artificial general intelligence: Early experiments with gpt-4*.

Arnav Chavan, Zhuang Liu, Deepak Gupta, Eric Xing, and Zhiqiang Shen. 2023. One-for-all: Generalized lora for parameter-efficient fine-tuning. *arXiv preprint arXiv:2306.07967*.

Yukang Chen, Shengju Qian, Haotian Tang, Xin Lai, Zhijian Liu, Song Han, and Jiaya Jia. 2024. LongLoRA: Efficient fine-tuning of long-context large language models. In *ICLR*.

Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. Qlora: Efficient finetuning of quantized llms. *arXiv preprint arXiv:2305.14314*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*.

Ning Ding, Yujia Qin, Guang Yang, Fuchao Wei, Zonghan Yang, Yusheng Su, Shengding Hu, Yulin Chen, Chi-Min Chan, Weize Chen, Jing Yi, Weilin Zhao, Zhiyuan Liu, Hai-Tao Zheng, Jianfei Chen, Yang Liu, Jie Tang, Juanzi Li, and Maosong Sun. 2023. Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nature Machine Intelligence*.

Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017. The WebNLG challenge: Generating text from RDF data. In *Proceedings of the 10th International Conference on Natural Language Generation*.

Demi Guo, Alexander M Rush, and Yoon Kim. 2021. Parameter-efficient transfer learning with diff pruning. In *ACL*.

Han Guo, Philip Greengard, Eric Xing, and Yoon Kim. 2024. LQ-LoRA: Low-rank plus quantized matrix decomposition for efficient language model finetuning. In *ICLR*.

Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021. DeBERTaV3: Improving DeBERTa using Electra-style pre-training with gradient-disentangled embedding sharing. *arXiv preprint arXiv:2111.09543*.

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 *ICML*.

Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-rank adaptation of large language models. In *ICLR*.

Shaoyi Huang, Dongkuan Xu, Ian Yen, Yijue Wang, Sung-En Chang, Bingbing Li, Shiyang Chen, Mimi Xie, Sanguthevar Rajasekaran, Hang Liu, and Caiwen Ding. 2022. Sparse progressive distillation: Resolving overfitting under pretrain-and-finetune paradigm. In *ACL*.

Dawid Jan Kopiczko, Tijmen Blankevoort, and Yuki M Asano. 2024. ELoRA: Efficient low-rank adaptation with random matrices. In *ICLR*.

Dawid Jan Kopiczko, Tijmen Blankevoort, and Yuki Markus Asano. 2023. Vera: Vector-based random matrix adaptation. *arXiv preprint arXiv:2310.11454*.

Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. In *EMNLP*.

Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.

Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. In *ACL*.

Yixiao Li, Yifan Yu, Chen Liang, Nikos Karampatziakis, Pengcheng He, Weizhu Chen, and Tuo Zhao. 2024. Loftq: LoRA-fine-tuning-aware quantization for large language models. In *ICLR*.

618	Vladislav Lialin, Sherin Muckatira, Namrata Shiva-	Taylor Shin, Yasaman Razeghi, Robert L Logan IV,	672
619	gunde, and Anna Rumshisky. 2024. ReloRA: High-	Eric Wallace, and Sameer Singh. 2020. Autoprompt:	673
620	rank training through low-rank updates. In <i>ICLR</i> .	Eliciting knowledge from language models with au-	674
621	Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mo-	tomatically generated prompts. In <i>EMNLP</i> .	675
622	hta, Tenghao Huang, Mohit Bansal, and Colin A Raf-	Chuanqi Tan, Fuchun Sun, Tao Kong, Wenchang Zhang,	676
623	fel. 2022a. Few-shot parameter-efficient fine-tuning	Chao Yang, and Chunfang Liu. 2018. A survey on	677
624	is better and cheaper than in-context learning. In	deep transfer learning. In <i>ICANN</i> .	678
625	<i>NeurIPS</i> .		
626	Weiyang Liu, Zeju Qiu, Yao Feng, Yuliang Xiu, Yux-	Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann	679
627	uan Xue, Longhui Yu, Haiwen Feng, Zhen Liu,	Dubois, Xuechen Li, Carlos Guestrin, Percy Liang,	680
628	Juyeon Heo, Songyou Peng, et al. 2023a. Parameter-	and Tatsunori B. Hashimoto. 2023. Stanford alpaca:	681
629	efficient orthogonal finetuning via butterfly factoriza-	An instruction-following llama model. <a href="https://github.com/tatsu-lab/stanford_alpaca">https://</a>	682
630	tion. <i>arXiv preprint arXiv:2311.06243</i> .	<a href="https://github.com/tatsu-lab/stanford_alpaca">github.com/tatsu-lab/stanford_alpaca</a> .	683
631	Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Tam, Zhengx-	Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier	684
632	iao Du, Zhilin Yang, and Jie Tang. 2022b. P-tuning:	Martinet, Marie-Anne Lachaux, Timothée Lacroix,	685
633	Prompt tuning can be comparable to fine-tuning	Baptiste Rozière, Naman Goyal, Eric Hambro,	686
634	across scales and tasks. In <i>ACL</i> .	Faisal Azhar, et al. 2023. Llama: Open and effi-	687
635	Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang,	cient foundation language models. <i>arXiv preprint</i>	688
636	Ruo Chen Xu, and Chenguang Zhu. 2023b. Gpte-	<i>arXiv:2302.13971</i> .	689
637	val: Nlg evaluation using gpt-4 with better human	Mojtaba Valipour, Mehdi Rezagholizadeh, Ivan	690
638	alignment. <i>arXiv preprint arXiv:2303.16634</i> .	Kobyzev, and Ali Ghodsi. 2023. Dylora: Parameter-	691
639	Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Man-	efficient tuning of pre-trained models using dynamic	692
640	dar Joshi, Danqi Chen, Omer Levy, Mike Lewis,	search-free low-rank adaptation. In <i>EACL</i> .	693
641	Luke Zettlemoyer, and Veselin Stoyanov. 2019.	Alex Wang, Amanpreet Singh, Julian Michael, Felix	694
642	Roberta: A robustly optimized bert pretraining ap-	Hill, Omer Levy, and Samuel R. Bowman. 2019.	695
643	proach. <i>arXiv preprint arXiv:1907.11692</i> .	GLUE: A multi-task benchmark and analysis plat-	696
644	Ilya Loshchilov and Frank Hutter. 2019. Decoupled	form for natural language understanding. In <i>ICLR</i> .	697
645	weight decay regularization. In <i>ICLR</i> .	Yaqing Wang, Subhabrata Mukherjee, Xiaodong Liu,	698
646	Kai Lv, Yuqing Yang, Tengxiao Liu, Qinghui Gao,	Jing Gao, Ahmed Hassan Awadallah, and Jian-	699
647	Qipeng Guo, and Xipeng Qiu. 2023. <a href="#">Full parameter</a>	feng Gao. 2022. Adamix: Mixture-of-adapter for	700
648	<a href="#">fine-tuning for large language models with limited</a>	parameter-efficient tuning of large language models.	701
649	<a href="#">resources</a> .	In <i>EMNLP</i> .	702
650	Shashi Narayan, Shay B Cohen, and Mirella Lap-	Yeming Wen and Swarat Chaudhuri. 2024. Batched	703
651	ata. 2018. Don't give me the details, just the	low-rank adaptation of foundation models. In <i>ICLR</i> .	704
652	summary! topic-aware convolutional neural net-	Thomas Wolf, Lysandre Debut, Victor Sanh, Julien	705
653	works for extreme summarization. <i>arXiv preprint</i>	Chaumond, Clement Delangue, Anthony Moi, Pier-	706
654	<i>arXiv:1808.08745</i> .	eric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz,	707
655	Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé,	et al. 2019. Huggingface's transformers: State-of-	708
656	Kyunghyun Cho, and Iryna Gurevych. 2021.	the-art natural language processing. <i>arXiv preprint</i>	709
657	Adapterfusion: Non-destructive task composition for	<i>arXiv:1910.03771</i> .	710
658	transfer learning. In <i>EACL</i> .	Lingling Xu, Haoran Xie, Si-Zhao Joe Qin, Xiaohui	711
659	Yevgeniy Puzikov and Iryna Gurevych. 2018. E2E	Tao, and Fu Lee Wang. 2023. Parameter-efficient	712
660	NLG challenge: Neural models vs. templates. In	fine-tuning methods for pretrained language models:	713
661	<i>Proceedings of the 11th International Conference on</i>	A critical review and assessment. <i>arXiv preprint</i>	714
662	<i>Natural Language Generation</i> .	<i>arXiv:2312.12148</i> .	715
663	Alec Radford, Jeffrey Wu, Rewon Child, David Luan,	Yuhui Xu, Lingxi Xie, Xiaotao Gu, Xin Chen,	716
664	Dario Amodei, Ilya Sutskever, et al. 2019. Language	Heng Chang, Hengheng Zhang, Zhengsu Chen, XI-	717
665	models are unsupervised multitask learners. <i>OpenAI</i>	AOPENG ZHANG, and Qi Tian. 2024. QA-loRA:	718
666	<i>blog</i> , 1(8):9.	Quantization-aware low-rank adaptation of large lan-	719
667	Colin Raffel, Noam Shazeer, Adam Roberts, Katherine	guage models. In <i>ICLR</i> .	720
668	Lee, Sharan Narang, Michael Matena, Yanqi Zhou,	Elad Ben Zaken, Shauli Ravfogel, and Yoav Gold-	721
669	Wei Li, Peter J Liu, et al. 2020. Exploring the limits	berg. 2021. Bitfit: Simple parameter-efficient	722
670	of transfer learning with a unified text-to-text trans-	fine-tuning for transformer-based masked language-	723
671	former. <i>JMLR</i> , 21(140):1–67.	models. <i>arXiv preprint arXiv:2106.10199</i> .	724

725 Longteng Zhang, Lin Zhang, Shaohuai Shi, Xiaowen  
726 Chu, and Bo Li. 2023a. [Lora-fa: Memory-efficient](#)  
727 [low-rank adaptation for large language models fine-](#)  
728 [tuning.](#)

729 Longteng Zhang, Lin Zhang, Shaohuai Shi, Xiaowen  
730 Chu, and Bo Li. 2023b. Lora-fa: Memory-efficient  
731 low-rank adaptation for large language models fine-  
732 tuning. *arXiv preprint arXiv:2308.03303*.

733 Qingru Zhang, Minshuo Chen, Alexander Bukharin,  
734 Pengcheng He, Yu Cheng, Weizhu Chen, and  
735 Tuo Zhao. 2022. Adaptive budget allocation for  
736 parameter-efficient fine-tuning. In *ICLR*.

737 Han Zhou, Xingchen Wan, Ivan Vulić, and Anna Korho-  
738 nen. 2023. Autopeft: Automatic configuration search  
739 for parameter-efficient fine-tuning. *arXiv preprint*  
740 *arXiv:2301.12132*.

## A Appendix

### A.1 Baselines

We compare our proposed method Delta-LoRA with Fine-Tuning and prior works of LoRA, AdaLoRA, and DyLoRA. For PEFT methods, we only train the incremental updates for  $W_V$  and  $W_Q$ , following the setup as used in LoRA’s paper. For Fine-Tuning methods, we use two extra training paradigms: (1) freeze the embedding and train all the other parameters as Fine-Tuning †; (2) train  $W_V$  and  $W_Q$  only as Fine-Tuning ‡.

**Fine-Tuning.** In the past few years, fine-tuning has become the mainstream paradigm for both NLP and CV tasks. However, fine-tuning full parameters is subject to potential drawbacks including overfitting and training instability (Huang et al., 2022). Therefore, freezing a subset of network layers and fine-tuning the rest has become a popular choice (Tan et al., 2018). In our experiments, we compare with full fine-tuning, fine-tuning with embedding layers frozen (Fine-tuning †) and fine-tuning query and value matrices only (Fine-tuning ‡).

**LoRA** (Hu et al., 2022) uses multiplication of two low-rank matrices to learn the incremental updates with reduced GPU memory cost. We follow their setups to reproduce experimental results for fair comparison.

**DyLoRA** (Valipour et al., 2023) randomly chooses a rank  $r$  for LoRA modules during learning.

**AdaLoRA** (Zhang et al., 2022) focuses on the challenge of determining the optimal rank for incremental updates. It employs an adaptive approach to singular value pruning, tailoring the rank selection to the magnitude of each singular value. Consequently, distinct ranks are employed for different layers.

### A.2 The Comparison between LoRA and Delta-LoRA with LLaMA-7B

#### A.2.1 Training and Inference Arguments Used in Our Method and Baseline

We choose LLaMA-7B to evaluate our method and LoRA. Here, we set the learning rate  $\gamma = 1e-4$ , batch size to 128,  $r = 8$ ,  $\alpha = 16$ , and training epochs to 3 for both two methods. Following the LoRA’s paper, we only tune  $W_Q$  and  $W_V$ . For Delta-LoRA, we choose start steps  $K = 100$  and  $\lambda = 0.25$ . When inference, we set the no\_repeat\_ngram\_size = 10, temperature = 0 and beam size = 4 to get a certain answer.

### A.2.2 The Evaluation for Our Method and Baselines

Current LLMs obtain the training data from the Internet, which may unintentionally cause data leakage. Therefore, using the mainstream NLP datasets to evaluate the effectiveness of Large Language Model is not reasonable and wisdom. Inspired by evaluation approach proposed by Liu et al. (2023b), we decided to use GPT-4 to judge the text generated by which method is accurate. First, we ask GPT-4 to generate 1,000 different questions. Second, we use the LLaMA-7B trained by two methods to generate the texts. Finally, we ask GPT-4 to give decision to tell us which text is accurate. It can choose from three options: *a. Choice 1* (LoRA generates accurate text), *b. Choice 2* (Delta-LoRA generates accurate text) and *c. Both Choice 1 and 2* (Both LoRA and Delta-LoRA generate accurate texts). The prompt we used for evaluation:

*Help me to determine which text is accurate for the given instruction and question. The answer can be chosen from a. Choice 1 is accurate, b. Choice 2 is accurate or c. both Choice 1 and 2 are accurate. Give me a certain answer and this is a choice question. Please don’t give reasons and the answer must be shorter than 20 words.*

*Question: ""*

*(Choice 1): ""*

*(Choice 2): ""*

### A.3 Algorithm of Delta-LoRA

Our Delta-LoRA can be found in Algorithm 1. Compared to LoRA, we added a step to update the pre-trained  $W$  without any extra GPU memory consumption.

### A.4 A Further Understanding of Delta-LoRA

#### A.4.1 The Differences between LoRA and Delta-LoRA

There are some fundamental differences between LoRA and Delta-LoRA.

- Given  $W + AB$ ,  $W$  is fixed in LoRA, but  $W$  will be updated in our Delta-LoRA. This is **the largest difference between LoRA and Delta-LoRA**. This modification can yield more training differences between LoRA and Delta-LoRA in the next few training steps.

- $\text{Rank}(\Delta W_{\text{Delta-LoRA}}) = \text{Rank}(W^{(T)} - W^{(0)} + AB) > \text{Rank}(\Delta W_{\text{LoRA}}) = \text{Rank}(AB)$ . The rank of the learned incremental weight matrix in our Delta-LoRA is larger than that in the original



---

**Algorithm 1: Delta-LoRA**


---

**Input:** Learning rate  $\eta$ ; weight decay  $\beta$ ; total training iterations  $T$ ; low rank  $r$ ; scale factor  $\alpha$ ; start steps  $K$ ; update ratio  $\lambda$ .

$\mathbf{A}$  is initialized by Kaiming Initialization,  $\mathbf{B} = \mathbf{0}$  and  $\mathbf{W}$  is initialized with pre-trained weights.

**for**  $t = 0, \dots, T - 1$  **do**

    Sample a mini-batch and compute gradients for  $\{\mathbf{A}, \mathbf{B}\}$  in each Delta-LoRA module.

    Update the first and second moments maintained by the optimizer with the computed gradients, and get the normalized gradients  $\hat{g}_A$  and  $\hat{g}_B$ .

$\mathbf{A}^{(t+1)} \leftarrow \mathbf{A}^{(t)} - \eta \hat{g}_A - \eta \beta \mathbf{A}^{(t)}$

$\mathbf{B}^{(t+1)} \leftarrow \mathbf{B}^{(t)} - \eta \hat{g}_B - \eta \beta \mathbf{B}^{(t)}$

**if**  $t > K$  **do**

$\mathbf{W}^{(t+1)} \leftarrow \mathbf{W}^{(t)} + \lambda \cdot \frac{\alpha}{r} \cdot (\mathbf{A}^{(t+1)} \mathbf{B}^{(t+1)} - \mathbf{A}^{(t)} \mathbf{B}^{(t)})$

**end if**

**end for**

**Output:** the fine-tuned parameters  $\{\mathbf{W}^{(T)}, \mathbf{A}^{(T)}, \mathbf{B}^{(T)}\}$

---

LoRA.

• The gradient flow is different between LoRA and Delta-LoRA. Suppose that we have  $\mathbf{W} \in R^{m \times n}$ ,  $\mathbf{A} \in R^{m \times r}$  and  $\mathbf{B} \in R^{r \times n}$ , where  $r \leq \min(m, n)$ . For LoRA, it keeps  $\mathbf{W}$  frozen, so that  $\mathbf{W}^* = \mathbf{W}^{(0)} + \frac{\alpha}{r} \cdot \mathbf{A}^{(t)} \mathbf{B}^{(t)}$ . For Delta-LoRA, it updates all matrices, and has  $\mathbf{W}^* = \mathbf{W}^{(t)} + \frac{\alpha}{r} \cdot \mathbf{A}^{(t)} \mathbf{B}^{(t)}$ .

So, we have the following equation according to Figure 4:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{W}^{(t)}} = \left( \frac{\partial \mathbf{h}_{i+3}}{\partial \mathbf{W}^{(t)}} \right)^\top \frac{\partial \mathcal{L}}{\partial \mathbf{h}_{i+4}} = \mathbf{h}_i^\top \frac{\partial \mathcal{L}}{\partial \mathbf{h}_{i+4}}$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \mathbf{A}^{(t)}} &= \left( \frac{\partial \mathbf{h}_{i+1}}{\partial \mathbf{A}^{(t)}} \right)^\top \frac{\partial \mathcal{L}}{\partial \mathbf{h}_{i+4}} \cdot \frac{\partial \mathbf{h}_{i+4}}{\partial \mathbf{h}_{i+2}} \cdot \frac{\partial \mathbf{h}_{i+2}}{\partial \mathbf{h}_{i+1}} \\ &= \mathbf{h}_i^\top \frac{\partial \mathcal{L}}{\partial \mathbf{h}_{i+4}} \frac{\partial \mathbf{h}_{i+2}}{\partial \mathbf{h}_{i+1}} = \mathbf{h}_i^\top \frac{\partial \mathcal{L}}{\partial \mathbf{h}_{i+4}} \mathbf{B}^{(t)\top} \\ &= \partial \mathcal{L} \frac{\partial \mathcal{L}}{\partial \mathbf{W}^{(t)} \cdot \mathbf{B}^{(t)\top}} \end{aligned}$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \mathbf{B}^{(t)}} &= \left( \frac{\partial \mathbf{h}_{i+1}}{\partial \mathbf{B}^{(t)}} \right)^\top \cdot \frac{\partial \mathcal{L}}{\partial \mathbf{h}_{i+4}} \cdot \frac{\partial \mathbf{h}_{i+4}}{\partial \mathbf{h}_{i+2}} \\ &= \mathbf{h}_{i+2}^\top \cdot \frac{\partial \mathcal{L}}{\partial \mathbf{h}_{i+4}} = (\mathbf{A}^{(t)})^\top \cdot \mathbf{h}_i^\top \cdot \frac{\partial \mathcal{L}}{\partial \mathbf{h}_{i+4}} \\ &= (\mathbf{A}^{(t)})^\top \cdot \frac{\partial \mathcal{L}}{\partial \mathbf{W}^{(t)}} \end{aligned}$$

Here, we provide the back-propagation process of LoRA:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \mathbf{A}^{(t+1)}} &= \frac{\partial \mathcal{L}}{\partial \mathbf{W}^{(t+1)}} \cdot (\mathbf{B}^{(t+1)})^\top \\ &= \frac{\partial \mathcal{L}}{\partial (\mathbf{W}^{(0)} + (\mathbf{A}^{(t)} + \Delta \mathbf{A}^{(t)}) (\mathbf{B}^{(t)} + \Delta \mathbf{B}^{(t)})}} \cdot (\mathbf{B}^{(t)} + \Delta \mathbf{B}^{(t)})^\top \end{aligned}$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \mathbf{B}^{(t+1)}} &= (\mathbf{A}^{(t+1)})^\top \cdot \frac{\partial \mathcal{L}}{\partial \mathbf{W}^{(t+1)}} \\ &= (\mathbf{A}^{(t)} + \Delta \mathbf{A}^{(t)})^\top \cdot \frac{\partial \mathcal{L}}{\partial (\mathbf{W}^{(0)} + (\mathbf{A}^{(t)} + \Delta \mathbf{A}^{(t)}) (\mathbf{B}^{(t)} + \Delta \mathbf{B}^{(t)})}} \end{aligned}$$

This is the back-propagation process of Delta-LoRA:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \mathbf{A}^{(t+1)}} &= \frac{\partial \mathcal{L}}{\partial \mathbf{W}^{(t+1)}} \cdot (\mathbf{B}^{(t+1)})^\top \\ &= \frac{\partial \mathcal{L}}{\partial ((\mathbf{W}^{(t)} + \lambda \Delta \mathbf{A}^{(t)}) \mathbf{B}^{(t)} + (\mathbf{A}^{(t)} + \Delta \mathbf{A}^{(t)}) (\mathbf{B}^{(t)} + \Delta \mathbf{B}^{(t)})}} \cdot (\mathbf{B}^{(t)} + \Delta \mathbf{B}^{(t)})^\top \\ \frac{\partial \mathcal{L}}{\partial \mathbf{B}^{(t+1)}} &= (\mathbf{A}^{(t+1)})^\top \cdot \frac{\partial \mathcal{L}}{\partial \mathbf{W}^{(t+1)}} \\ &= (\mathbf{A}^{(t)} + \Delta \mathbf{A}^{(t)})^\top \cdot \frac{\partial \mathcal{L}}{\partial ((\mathbf{W}^{(t)} + \lambda \Delta \mathbf{A}^{(t)}) \mathbf{B}^{(t)} + (\mathbf{A}^{(t)} + \Delta \mathbf{A}^{(t)}) (\mathbf{B}^{(t)} + \Delta \mathbf{B}^{(t)})}} \end{aligned}$$

#### A.4.2 The Expansion of $\Delta \mathbf{AB}$

In the real training process, we need to consider a variety of training arguments, such as optimizer and the regularization for  $\Delta \mathbf{AB}$ . Suppose that we use the AdamW (Loshchilov and Hutter, 2019) and  $L_2$  regularization, the  $\Delta \mathbf{AB}$  can be expanded in




Figure 4: The backward propagation of Delta-LoRA.

the following equation:

$$\begin{aligned}
\Delta \mathbf{A} \mathbf{B} &= \mathbf{A}^{(t+1)} \mathbf{B}^{(t+1)} - \mathbf{A}^{(t)} \mathbf{B}^{(t)} \\
&= (\mathbf{A}^{(t)} - \eta \hat{g}_A - \eta \beta \mathbf{A}^{(t)}) \cdot (\mathbf{B}^{(t)} \\
&\quad - \eta \hat{g}_B - \eta \beta \mathbf{B}^{(t)}) - \mathbf{A}^{(t)} \mathbf{B}^{(t)} \\
&= \mathbf{A}^{(t)} \mathbf{B}^{(t)} - \eta \mathbf{A}^{(t)} \hat{g}_B - \eta \beta \mathbf{A}^{(t)} \mathbf{B}^{(t)} \\
&\quad - \eta \hat{g}_A \mathbf{B}^{(t)} + \eta^2 \hat{g}_A \hat{g}_B + \eta^2 \beta \hat{g}_A \mathbf{B}^{(t)} \\
&\quad - \eta \beta \mathbf{A}^{(t)} \mathbf{B}^{(t)} + \eta^2 \beta \mathbf{A}^{(t)} \hat{g}_B \\
&\quad + \eta^2 \beta^2 \mathbf{A}^{(t)} \mathbf{B}^{(t)} - \mathbf{A}^{(t)} \mathbf{B}^{(t)} \\
&= -\eta \mathbf{A}^{(t)} \hat{g}_B - \eta \beta \mathbf{A}^{(t)} \mathbf{B}^{(t)} - \eta \hat{g}_A \mathbf{B}^{(t)} \\
&\quad + \eta^2 \hat{g}_A \hat{g}_B + \eta^2 \beta \hat{g}_A \mathbf{B}^{(t)} - \eta \beta \mathbf{A}^{(t)} \mathbf{B}^{(t)} \\
&\quad + \eta^2 \beta \mathbf{A}^{(t)} \hat{g}_B + \eta^2 \beta^2 \mathbf{A}^{(t)} \mathbf{B}^{(t)} \\
&\approx -\eta \mathbf{A}^{(t)} \hat{g}_B - \eta \hat{g}_A \mathbf{B}^{(t)}
\end{aligned} \tag{5}$$

where  $\eta$  is the learning rate,  $\beta$  is weight decay. What's more, for pre-trained weight  $\mathbf{W}$ ,  $\Delta \mathbf{W} = \eta \hat{g}_W + \eta \beta \mathbf{W}^{(t)}$ . As a consequence,  $\Delta \mathbf{A} \mathbf{B}$  is not equal to  $\Delta \mathbf{W}$  in the training process.

### A.5 The Parameter Sensitivity Study

**Parameter Sensitivity.** Here, we explore the hyper-parameter  $K$  in Algorithm 1 and  $\lambda$  in Equation 3. For the hyper-parameter  $K$ , we select it from 0 to 1000 with the interval of 100. From Table 9, we find that our Delta-LoRA could not bring in any improvement before  $K = 400$ , and it will keep a relatively good performance when  $K$  is larger than 500. What is more, we choose different numbers for  $\lambda$ , ranging from 0 to 5. According to Table 8, the 5 metrics rise rapidly after  $\lambda = 0$  and reach best at  $\lambda = 2$ , while the performance has small drops on 5 evaluation scores if  $\lambda$  is chosen from 3 to 5.

### A.6 Hyper-Parameter Used in Our Experiments

We report the hyper-parameter that used in our experiments. Table 10 and Table 11 show the hyper-parameter that we used for the training and evaluation on E2E Challenge and WebNLG Challenge 2017 dataset. The Table 12 and Table 13 are the training and evaluation hyper parameter for XSum dataset, and the Table 14 consists of hyper-parameters for 8 datasets in GLUE benchmark.

Table 8: The parameter sensitivity study of update ratio  $\lambda$  for our proposed Delta-LoRA on E2E Challenge dataset. The best results are **boldfaced**.

$\lambda$	BLEU	NIST	METEOR	ROUGE-L	CIDEr
0	68.94	8.73	45.27	70.81	2.41
1	69.77	8.81	45.99	71.58	2.46
2	<b>70.84</b>	<b>8.91</b>	<b>46.47</b>	<b>72.24</b>	<b>2.53</b>
3	70.14	8.84	46.39	71.45	2.45
4	70.03	8.83	46.21	71.56	2.47
5	70.13	8.85	46.35	71.72	2.48

Table 9: The parameter sensitivity study of start steps  $K$  for our proposed Delta-LoRA on E2E Challenge dataset. The best results are **boldfaced**.

$K$	BLEU	NIST	METEOR	ROUGE-L	CIDEr
0	69.10	8.75	45.54	71.31	2.41
100	69.97	8.84	46.07	71.40	2.46
200	69.72	8.83	45.82	71.41	2.43
300	69.73	8.86	45.98	71.09	2.46
400	70.18	8.89	46.30	71.66	2.49
500	70.84	8.91	46.47	<b>72.24</b>	<b>2.53</b>
600	70.38	8.86	46.38	71.70	2.47
700	70.61	8.89	46.43	72.13	2.51
800	70.70	8.89	46.30	71.97	2.51
900	<b>71.00</b>	<b>8.92</b>	<b>46.47</b>	72.04	2.52
1000	70.87	8.89	46.31	72.06	2.50

Table 10: The training hyper-parameter used for E2E Challenge and WebNLG Challenge 2017 dataset.

Hyper-Parameter	E2E Challenge	WebNLG Challenge 2017
Learning Rate $\eta$	2e-4	2e-4
Batch Size	8	8
Number of Epochs	5	5
Weight Decay $\beta$	0.01	0.01
Resid_pdrop	0	0.09
Attn_pdrop	0	0.09
Embd_pdrop	0	0
Label Smooth	0	0
Start Steps $K$	500	500
Update Ratio $\lambda$	2	5
Rank $r$	4	4
Alpha $\alpha$	32	32
Trainable Matrices	$\mathbf{W}_Q, \mathbf{W}_V$	$\mathbf{W}_Q, \mathbf{W}_V$
LR Scheduler	Linear	Linear
Warmup Steps	500	500

Table 11: The hyper-parameter for evaluation used for E2E Challenge and WebNLG Challenge 2017 dataset.

Hyper-Parameter	E2E Challenge	WebNLG Challenge 2017
Beam Size	10	5
Penalty	0.8	1.0
No Repeat Ngram Size	4	4

Table 12: The training hyper-parameter used for XSum dataset.

Hyper-Parameter	Xsum
Learning Rate $\eta$	2e-4
Batch Size	64
Number of Epochs	25
Weight Decay $\beta$	0
Activation Dropout	0
Dropout	0
Classifier Dropout	0
Start Steps $K$	1000
Update Ratio $\lambda$	0.5
Rank $r$	4
Alpha $\alpha$	32
Trainable Matrices	$\mathbf{W}_Q, \mathbf{W}_V$
LR Scheduler	Linear
Warmup Steps	3000

Table 13: The hyper-parameter for evaluation used for XSum dataset.

Hyper-Parameter	Xsum
Beam Size	8
Penalty	1.0
No Repeat N-gram Size	4

Table 14: The training hyper-parameters of our proposed Delta-LoRA on GLUE benchmark. We adopt the most of hyper-parameters in LoRA’s paper and implement our method based on the codes given by LoRA’s repository.

Hyper-Parameter	MNLI	SST-2	MRPC	CoLA	QNLI	QQP	RTE	STS-B
Learning Rate $\eta$	4e-4	5e-4	5e-4	6e-4	3e-4	6e-4	4e-4	4e-4
Batch Size	128	128	128	64	256	128	128	128
Number of Epochs	30	60	30	80	25	25	80	40
Weight Decay $\beta$	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
Max Sequence Length	256	256	256	256	256	256	512	256
Start Steps $K$	2000	400	10	200	600	400	200	200
Rank $r$	8	8	8	8	8	8	8	8
Alpha $\alpha$	16	16	16	16	16	16	16	16
LR Scheduler	Linear	Linear	Linear	Linear	Linear	Linear	Linear	Linear
Update Ratio $\lambda$	0.5	0.5	0.5	1	1	0.5	0.5	0.5
Trainable Matrices	$\mathbf{W}_Q, \mathbf{W}_V$	$\mathbf{W}_Q, \mathbf{W}_V$	$\mathbf{W}_Q, \mathbf{W}_V$	$\mathbf{W}_Q, \mathbf{W}_V$	$\mathbf{W}_Q, \mathbf{W}_V$	$\mathbf{W}_Q, \mathbf{W}_V$	$\mathbf{W}_Q, \mathbf{W}_V$	$\mathbf{W}_Q, \mathbf{W}_V$
Warmup Ratio	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06
Evaluation Metrics	Accuracy	Accuracy	Accuracy	Matthews Correlation	Accuracy	Accuracy	Accuracy	Pearson