

000 001 002 003 004 005 DUAL LoRA: ENHANCING LoRA WITH MAGNITUDE 006 AND DIRECTION UPDATES 007 008 009

010 **Anonymous authors**
011 Paper under double-blind review
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026

ABSTRACT

027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053
Low-rank adaptation (LoRA) is one of the most popular methods among parameter-efficient fine-tuning (PEFT) methods to adapt pre-trained large language models (LLMs) to specific downstream tasks. However, the model trained based on LoRA often has an unsatisfactory performance due to its low-rank assumption. In this paper, we propose a novel method called Dual LoRA to improve the performance by incorporating an inductive bias into the original LoRA. Specifically, we separate low-rank matrices into two groups: the magnitude group to control whether or not and how far we should update a parameter and the direction group to decide whether this parameter should move forward or backward, to better simulate the parameter updating process of the full fine-tuning based on gradient-based optimization algorithms. We show that this can be simply achieved by adding a ReLU function to the magnitude group and a sign function to the direction group. We conduct several experiments over a wide range of NLP tasks, including natural language generation (NLG), understanding (NLU), and commonsense reasoning datasets on GPT-2, RoBERTa, DeBERTa, and LLaMA-1/2/3 as baseline models. The results show that we consistently outperform LoRA and its state-of-the-art variants with the same number of trainable parameters.

1 INTRODUCTION

Large language models (LLMs) have shown promising results on almost all natural language processing (NLP) tasks (Touvron et al., 2023a; Achiam et al., 2023) and other multi-modal tasks (Liu et al., 2024a), by adapting a well trained LLM to different downstream applications. Full fine-tuning (FFT) is a straightforward way to achieve this goal, but it requires tremendous computational resources and time to complete the fine-tuning process. Thus, parameter-efficient fine-tuning (PEFT) which updates a small fraction (less than 2%) of parameters has attracted more and more attention due to its low memory and time requirements.

Traditional PEFT methods include adapter tuning (Hu et al., 2023) which adds trainable tiny modules to adapt to downstream tasks, prompt tuning (Peng et al., 2024) that inserts learnable prompt vectors to the existing input, and low-rank adaptation (LoRA) (Hu et al., 2021a) which updates the original parameters by adding low-rank matrices. Among them, LoRA surpasses other methods by achieving better performance without generating additional inference costs.

Many follow-ups manage to improve the fine-tuning performance of LoRA. LoRA+ (Hayou et al., 2024) uses different learning rates to update low-rank matrices A and B and enhance the performance with a well-chosen learning rate ratio. DoRA (Liu et al., 2024b) decomposes the original weight matrix into a normalized matrix and its corresponding norm and applies the original LoRA to the normalized matrix. FLoRA (Si et al., 2024) generates LoRA to high dimensional space and inserts a low-rank core matrix into the original LoRA matrices to improve its performance. MoRA (Jiang et al., 2024) replaces the low-rank matrices with a square matrix to achieve high-rank updating and applies a compress layer and a decompress layer to maintain a roughly similar number of trainable parameters. However, they share a common drawback: as the trainable parameters are much fewer than those of FFT, updating them without incorporating prior knowledge will inevitably result in unsatisfactory model accuracy.

Thus, in this paper we introduce an inductive bias into the original LoRA method, *i.e.*, to simulate the parameter updating process of FFT, which utilizes gradient-based optimization algorithms.

054 Specifically, we divide the low-rank matrices into two groups: the magnitude group, which controls
 055 whether and to what extent a parameter should be updated; and the direction group, which deter-
 056 mines the direction of the update—whether it should be positive or negative. The whole fine-tuning
 057 process can be treated as adjusting the sign and magnitude of each element in the update matrix and
 058 adding them back to the original parameters to gradually achieve the optimal solution. We conduct
 059 experiments to validate the effectiveness of our method over a wide range of NLP tasks including
 060 natural language generation (NLG), understanding (NLU), and commonsense reasoning to make a
 061 fair comparison with state-of-the-art methods. The evaluation results on different LLM models such
 062 as GPT-2, RoBERTa, DeBERTa, LLaMA-7B/13B, LLaMA2-7B, LLaMA3-8B, and LLaMA3-70B-
 063 Instruct show that we can achieve consistent improvements over these SOTA methods by using the
 064 same number of training parameters.

065 The contributions of our method are summarized as follows:

- 067 We introduce Dual LoRA, a novel method that replaces the original low-rank matrices in
 068 LoRA with two groups of parameters: a magnitude group and a direction group to sep-
 069 arately determine the amplitude and sign of the update to the original parameters in the
 070 LLMs. This can be treated as incorporating an inductive bias into the original LoRA to
 071 better learn the parameter updating process of FFT, which can improve the performance.
- 072 Dual LoRA consistently outperforms state-of-the-art methods on a wide range of NLP
 073 tasks across various baseline models with different sizes (from 7B to up to 70B), which
 074 demonstrates the effectiveness of our method.

076 2 RELATED WORKS

079 In this section, we first introduce different parameter-efficient fine-tuning (PEFT) methods, followed
 080 by a deeper dive into the LoRA series methods.

082 2.1 PEFT METHODS IN LLMs

084 **Prefix tuning** is the first kind of methods (Li & Liang, 2021; Liu et al., 2022; Zhang et al., 2024)
 085 in PEFT. It was first proposed by Li *et.al.* (Li & Liang, 2021), which was a lightweight alternative
 086 to FFT that kept LLM parameters frozen and only optimized a sequence of continuous task-specific
 087 vectors called prefix. Dynamic prefix-tuning (Liu et al., 2022) proposed a generative template-based
 088 event extraction method with dynamic prefixes by integrating context information with type-specific
 089 prefixes to learn a context-specific prefix for each context. Selective prefix-tuning (Zhang et al.,
 090 2024) showed that prefix tokens carried context-specific information and enhanced their special-
 091 ization can improve model performance. Thus, they integrated a selective mechanism inspired by
 092 selective self-attention and introduced selective loss to encourage diversity in prefix tokens.

093 **Prompt tuning** is the second kind of PEFT method that added trainable embeddings to original word
 094 embeddings and learned these soft prompts through back-propagation and tuned them to incorporate
 095 signals from any number of labeled examples (Lester et al., 2021). P-Tuning v2 (Liu et al., 2021)
 096 empirically found that properly optimized prompt tuning can be universally effective across a wide
 097 range of model scales and NLU tasks, which increased the capacity of continuous prompts and
 098 closed the gap to FFT. Knowledgeable Prompt-tuning (Hu et al., 2021b) improved and stabilized the
 099 original prompt-tuning method by expanding the label word space of the verbalizer with external
 100 knowledge bases and refining it with PLM before predicting.

101 **Representation fine-tuning (ReFT)** aims to train interventions that manipulate model repres-
 102 entations to steer model behaviors on downstream tasks at inference time. ReFT (Wu et al., 2024)
 103 introduced a family of ReFT methods that operated on a frozen base model and learned task-specific
 104 interventions on hidden representations.

105 Although the aforementioned methods improved the performance of LLMs in downstream tasks,
 106 they suffered the problem that the original architecture of the baseline model needed to be changed
 107 and the inference speed was slowed down. Compared to them, LoRA-based methods had exactly
 the same inference latency to the baseline LLMs.

108
109

2.2 LORA-BASED METHODS

110

LoRA (Hu et al., 2021a) assumed that only a small number of task-specific parameters needed to be tuned to fit the downstream tasks and updated the weights with two low-rank matrices. These matrices can be merged back into the original weights during inference to avoid additional computational costs. LoRA+ (Hayou et al., 2024) argued that LoRA led to sub-optimal results, and the problem can be corrected by setting different learning rates for the low-rank matrices A and B with a fixed learning rate ratio. MoRA (Jiang et al., 2024) believed that the low-rank updating mechanism limited the ability of LLMs and used a square matrix to achieve high-rank updating with the same number of trainable parameters. Two non-parameter operators were used to reduce the input dimension and increase the output dimension of this square matrix. DoRA (Liu et al., 2024b) decomposed the pre-trained weight into magnitude and direction for fine-tuning, and employed original LoRA for direction component update to accelerate the training process.

120

The methods mentioned above can improve the performance of downstream tasks. However, their performance is still unsatisfactory because of the low-rank assumption (Hu et al., 2021a; Hayou et al., 2024; Liu et al., 2024b). Although MoRA (Jiang et al., 2024) attempted to address this issue by using a high-rank matrix, its rank and the number of trainable parameters remained significantly lower than those in FFT. Thus, it is difficult to achieve satisfactory model performance without incorporating prior knowledge into the training process.

126

Note that both DoRA and our method have magnitude and direction groups, but the meaning behind them is totally different. The direction and magnitude in DoRA can be treated as a normalized weight matrix and its corresponding norm. In our method, we are trying to simulate the parameter updating process of FFT which utilizes gradient-based optimization algorithms. Thus, the direction and magnitude control the sign and to what extent a parameter should be updated.

131

Another family of methods aim to modify the gradient calculation and backward propagation process of training, such as GaLore (Zhao et al., 2024), FLoRA (Hao et al., 2024) and GaRare (Liu et al.). These methods are orthogonal to the proposed Dual LoRA which only focuses on the architecture and forward pass modification, and a detailed discussion falls outside the scope of this paper.

136

3 METHOD

138

In this section, we first introduce the preliminaries of LoRA and optimization methods. Then, we give a thorough analysis of our proposed Dual LoRA and explain its advantage over previous LoRA-based methods.

142

3.1 LOW-RANK ADAPTATION (LORA)

144

Given a pre-trained weight matrix $W_0 \in \mathbb{R}^{d \times k}$, LoRA (Hu et al., 2021a) assumes that a low “intrinsic rank” is enough during adaptation on downstream tasks and constrains the updated matrix with a low-rank decomposition:

147

$$W' = W_0 + \Delta W = W_0 + \frac{\alpha}{r} \cdot BA, \quad (1)$$

150

where $B \in \mathbb{R}^{d \times r}$ and $A \in \mathbb{R}^{r \times k}$ are two low-rank matrices with rank $r \ll \min(d, k)$, α is a fixed hyper-parameter to control the influence of the low-rank matrices, and W' is the final weight matrix after fine-tuning.

153

Given an original forward pass $h = W_0 x$ with an input x , the modified forward pass can be expressed as:

155

$$h = W_0 x + \Delta W x = (W_0 + \frac{\alpha}{r} \cdot BA)x. \quad (2)$$

158

Note that the usage of LoRA does not affect the inference speed since the low-rank matrices A and B can be merged back into the original weight W_0 , and the dimension of the final weight matrix W' is the same as the pre-trained weight matrix W_0 . Since the trainable low-rank matrices have fewer parameters (less than 2%) compared to the original matrices, LoRA usually has insufficient performance.

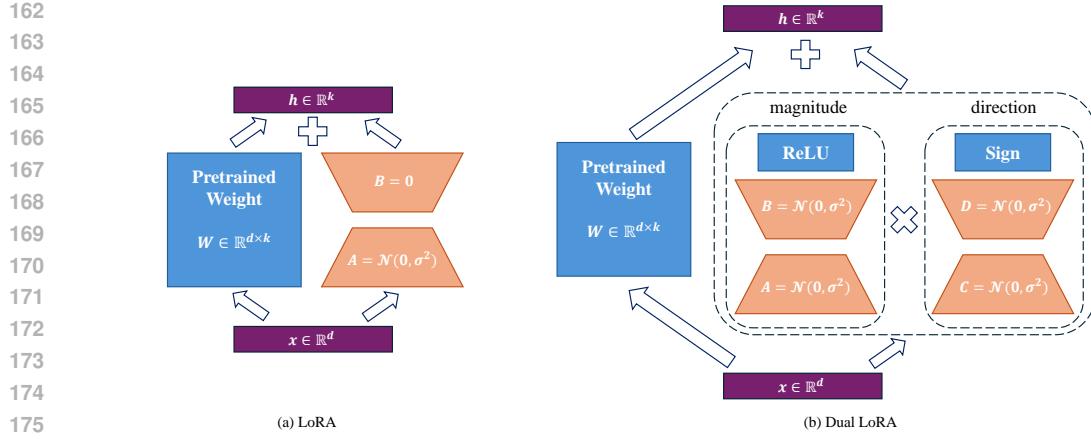


Figure 1: The architecture of the original LoRA and our proposed Dual LoRA. The low-rank update matrices are separated into the magnitude group and the direction group.

3.2 OPTIMIZATION METHODS

Given a loss function $\ell(\hat{y}, y)$ which measures the cost between the prediction \hat{y} and the ground-truth label y , we can choose a family \mathcal{F} of functions $f_w(x)$ with learnable weight w and input x , and seek the function $f \in \mathcal{F}$ to minimize the loss $\ell(f_w(x), y)$ averaged on the input examples:

$$E_n(f_w) = \frac{1}{n} \sum_{i=1}^n \ell(f_w(x_i), y_i). \quad (3)$$

In order to minimize the empirical risk $E_n(f_w)$, a global optimum weight w^* needs to be found step by step using a series of optimization methods. Specifically, we have:

$$w_{t+1} = w_t + \gamma \Delta w, \quad (4)$$

where γ is the learning rate and w_{t+1} is expected to converge to the global optimum w^* as the training proceed.

To achieve this, different optimization methods leverage different ways to compute Δw . For example, gradient descent (Bottou, 2010) uses $\Delta w = \frac{1}{n} \sum_{i=1}^n \nabla_w \ell(f(x_i), y_i)$ to compute the update, and Adam (Kingma, 2014) uses $\Delta w = \hat{m}_t / (\sqrt{\hat{v}_t} + \epsilon)$ where \hat{m}_t and \hat{v}_t are the first-moment estimate and second-moment estimate, and $\epsilon = 10^{-8}$.

Both FFT and LoRA fine-tune the model based on the optimization methods mentioned above. However, FFT assumes Δw is a full-rank matrix while LoRA decomposes Δw into two low-rank matrices and trains them without any other prior knowledge, which is the main reason that causes the performance drop.

3.3 DUAL LoRA

Note that the update matrix Δw can always be decomposed into magnitude and direction regardless of the optimization method used. Learning these components separately can be treated as adding an inductive bias into the original LoRA, aiding in facilitating the search for the optimal solution within the solution space.

Instead of using two low-rank matrices, we use four low-rank matrices and separate them into a magnitude group and a direction group in Dual LoRA, as shown in Fig. 1.

Magnitude group. Given two low-rank matrices $A \in \mathbb{R}^{r_1 \times k}$ and $B \in \mathbb{R}^{d \times r_1}$, the magnitude group can be computed as:

$$W_m = \text{ReLU}(BA), \quad (5)$$

216 which has two effects. Firstly, non-negative outputs can be treated as learning the magnitude of the
 217 update during the training process. Secondly, we can easily freeze some of the elements that are
 218 already well-trained for the downstream tasks in the original weight matrix by learning the output
 219 elements of BA to be negative and filter them out with ReLU function, which is hard for previous
 220 LoRA-based methods to achieve such a goal.

221 **Direction group.** Given two low-rank matrices $C \in \mathbb{R}^{r_2 \times k}$ and $D \in \mathbb{R}^{d \times r_2}$, the direction group
 222 can be computed as:

$$223 \quad W_d = \text{Sign}(DC), \quad (6)$$

225 where $\text{Sign}(\cdot)$ is an element-wise operation that outputs $+1$ for positive input and -1 otherwise.
 226 Note that the gradient of the sign function is zero almost everywhere, and backward propagation
 227 cannot be applied during training. Thus, given $x_b = \text{Sign}(x)$, the straight-through estimator (STE)
 228 method (Bengio et al., 2013) is introduced to compute its gradient as:

$$229 \quad \frac{\partial \mathcal{L}}{\partial x} = \text{Clip}\left(\frac{\partial \mathcal{L}}{\partial x_b}, -1, 1\right), \quad (7)$$

231 in which \mathcal{L} is the corresponding loss function for a downstream task and:

$$232 \quad \text{Clip}(x, -1, 1) = \begin{cases} -1, & \text{if } x < -1, \\ 233 \quad x, & \text{if } -1 \leq x < 1, \\ 234 \quad 1, & \text{otherwise.} \end{cases} \quad (8)$$

235 The direction group can control the sign of each element in the update matrix, which is a two-
 236 way direction to decide whether the element in the original weight matrix should move forward or
 237 backward.

238 **Overall update.** Given a pre-trained weight matrix W_0 , the overall update of our Dual LoRA can
 239 be expressed as:

$$241 \quad W' = W_0 + \Delta W = W_0 + \frac{\alpha}{\sqrt{r_1 r_2}} W_m \odot W_d, \quad (9)$$

243 where \odot represents an element-wise product (Hadamard product) between two matrices. Similarly,
 244 given the original forward pass $h = W_0 x$, the modified forward pass is:

$$245 \quad h = W_0 x + \Delta W x = (W_0 + \frac{\alpha}{\sqrt{r_1 r_2}} W_m \odot W_d) x, \quad (10)$$

247 which does not affect the inference process as long as we merge ΔW into W_0 .

248 **Initialization.** LoRA uses random Gaussian initialization for A and zero for B to make sure the
 249 update matrix is zero at the beginning of training, as shown in Fig. 1(a). In Dual LoRA, however,
 250 none of the low-rank matrices in the magnitude group should be initialized with zero. Otherwise,
 251 either all trainable parameters are dead, or we cannot achieve the goal that the update matrix is zero
 252 due to the $\text{ReLU}(\cdot)$ function and $\text{Sign}(\cdot)$ function.

253 Specifically, given

$$255 \quad \Delta W = \frac{\alpha}{\sqrt{r_1 r_2}} \text{ReLU}(BA) \odot \text{Sign}(DC), \quad (11)$$

256 we can compute the gradient of the loss function \mathcal{L} with respect to four low-rank matrices as:

$$258 \quad \frac{\partial \mathcal{L}}{\partial A} = \frac{\partial \mathcal{L}}{\partial \Delta W} \cdot \frac{\alpha}{\sqrt{r_1 r_2}} B^\top \cdot \text{Sign}(DC) \cdot \mathbb{1}_{BA > 0}, \quad \frac{\partial \mathcal{L}}{\partial B} = \frac{\partial \mathcal{L}}{\partial \Delta W} \cdot \frac{\alpha}{\sqrt{r_1 r_2}} \text{Sign}(DC) \cdot \mathbb{1}_{BA > 0} \cdot A^\top, \\ 259 \quad \frac{\partial \mathcal{L}}{\partial C} = \text{Clip}\left(\frac{\partial \mathcal{L}}{\partial \Delta W}, -1, 1\right) \cdot \frac{\alpha}{\sqrt{r_1 r_2}} D^\top \cdot \text{ReLU}(BA), \quad \frac{\partial \mathcal{L}}{\partial D} = \text{Clip}\left(\frac{\partial \mathcal{L}}{\partial \Delta W}, -1, 1\right) \cdot \frac{\alpha}{\sqrt{r_1 r_2}} \text{ReLU}(BA) \cdot C^\top, \\ 260 \quad 261 \quad 262 \quad 263 \quad (12)$$

where $\mathbb{1}$ is the indicator function.

264 It is easy to know that when setting $A = 0$ or $B = 0$, we will have $\mathbb{1}_{BA > 0} = 0$ and $\text{ReLU}(BA) = 0$
 265 and all four gradients in Eq. 12 are zeros which will cause the training process to be dead. Setting
 266 $C = 0$ or $D = 0$ will not result in such a problem, but it cannot achieve the goal that the update
 267 matrix Eq. 11 is zero at the beginning of training since $\text{Sign}(x)$ always outputs $+1$ or -1 depending
 268 on the input. Thus, during the experiments, we use random Gaussian initialization for all four low-
 269 rank matrices and apply a warm-up strategy for the first few training steps to make sure that $\Delta W = 0$
 at the start.

270 Table 1: The results of the proposed Dual LoRA and other competitors with LLaMA-7B/13B,
 271 LLaMA2-7B, LLaMA3-8B and LLaMA3-70B-Instruct on commonsense reasoning datasets. For
 272 all matrices, higher is better.

275 Model	276 Methods	277 Trainable 278 Param. (%)	279 Commonsense Reasoning Datasets								
			280 BoolQ	281 PIQA	282 SIQA	283 HellaS	284 WinoG	285 ARC-e	286 ARC-c	287 OBQA	288 Avg.
289 L-7B	Adapter-P	3.54	67.9	76.4	78.8	69.8	78.9	73.7	57.3	75.2	72.2
	LoRA ($r = 64$)	1.64	67.3	79.0	76.3	76.6	78.8	74.5	59.3	77.4	73.6
	DoRA ($r = 32$)	0.84	68.7	83.3	79.4	85.5	81.3	80.8	66.0	78.8	78.0
	DoRA ($r = 64$)	1.65	68.9	82.1	77.4	75.9	80.0	80.0	64.8	81.0	76.3
	Dual LoRA ($r = 32$)	1.64	70.0	83.2	79.5	87.3	83.0	81.6	65.2	81.0	78.9
290 L-13B	LoRA ($r = 32$)	0.67	71.6	83.4	80.0	89.9	84.2	81.2	67.7	80.8	79.9
	DoRA ($r = 16$)	0.35	71.7	84.2	80.6	90.5	85.2	83.1	68.4	80.4	80.5
	DoRA ($r = 32$)	0.68	72.4	84.9	81.2	91.5	83.7	84.6	68.9	81.6	81.1
	Dual LoRA ($r = 16$)	0.67	72.5	84.2	79.9	92.7	83.8	84.8	72.4	83.2	81.7
291 L2-7B	LoRA ($r = 16$)	0.41	70.4	82.9	79.0	81.3	81.5	82.4	69.2	80.4	78.4
	LoRA ($r = 32$)	0.83	68.9	82.2	78.1	86.9	81.2	79.3	65.4	78.4	77.6
	DoRA ($r = 16$)	0.43	63.5	82.8	79.5	90.6	82.4	83.9	69.9	81.8	79.3
	DoRA ($r = 32$)	0.84	72.2	83.5	80.3	89.0	82.5	84.1	69.5	80.4	80.2
	Dual LoRA ($r = 16$)	0.83	72.3	83.3	79.8	89.8	84.6	84.8	70.2	82.8	81.0
292 L3-8B	RandLoRA	0.70	76.3	88.1	80.3	95.7	86.1	90.4	80.9	87.0	85.6
	LoRA ($r = 16$)	0.35	71.7	86.8	79.5	93.9	84.4	87.4	76.3	84.2	83.0
	LoRA ($r = 32$)	0.70	71.2	85.1	79.3	92.1	82.6	85.2	70.1	81.4	80.9
	DoRA ($r = 16$)	0.35	75.1	87.8	80.8	95.6	86.3	90.4	80.0	85.6	85.2
	DoRA ($r = 32$)	0.71	71.7	88.0	80.2	95.5	86.6	90.7	78.4	85.0	84.5
	Dual LoRA ($r = 16$)	0.70	75.5	89.2	81.4	95.8	86.0	90.5	81.1	86.6	85.8
293 L3-70B	LoRA ($r = 16$)	0.197	78.6	92.8	83.4	92.7	92.6	97.5	91.7	94.4	90.5
	DoRA ($r = 16$)	0.202	78.4	93.0	83.8	96.5	92.3	97.6	92.3	94.6	91.1
	Dual LoRA ($r = 8$)	0.197	81.4	94.0	84.4	97.9	93.6	97.3	91.0	95.2	91.9

300 4 EXPERIMENTS

301 In this section, we evaluate the effectiveness of the proposed Dual LoRA on various NLP tasks. We
 302 compare our methods with other PEFT competitors by fine-tuning LLaMA-7B/13B, LLaMA2-7B,
 303 LLaMA3-8B, and LLaMA3-70B-Instruct models on a series of commonsense reasoning datasets.
 304 Then, we explore the ability of our method on the neural language understanding (NLU) dataset
 305 GLUE by fine-tuning RoBERTa base/large and DeBERTa XXL. Furthermore, we conduct exper-
 306 iments on neural language generation (NLG) datasets including E2E NLG Challenge, DART and
 307 WebNLG using GPT2_M and GPT2_L as backbones (see Appendix A). All experiments above show
 308 that Dual LoRA can surpass other LoRA-based methods with the same or fewer trainable parameters
 309 and achieve state-of-the-art results. Finally, we analyze our method further by performing a series
 310 of ablation studies. In the following experiments, we set the rank of the magnitude group and the
 311 direction group as the same, *i.e.*, $r_1 = r_2 = r$ unless specified.

312 **Competitors.** We compare Dual LoRA with a series of baseline methods including LoRA-based
 313 methods (LoRA (Hu et al., 2021a), LoRA+ (Hayou et al., 2024), GaLore (Zhao et al., 2024),
 314 GaRare (Liu et al.), Delta-LoRA (Zi et al., 2023), CorDA (Yang et al., 2024), VeRA (Kopitzko
 315 et al., 2024), RandLoRA (Albert et al., 2025), and DoRA (Liu et al., 2024b)) and other PEFT meth-
 316 ods (efficient adapter design with LayerNorm (Adapter-L) (Lin et al., 2020), parallel adapter tuning
 317 (Adapter-P) (He et al., 2021) and prefix-layer tuning (Prefix) (Li & Liang, 2021)).

319 320 4.1 COMMONSENSE REASONING

321 **Datasets and baseline models.** We evaluate Dual LoRA and different PEFT methods on the com-
 322 monsense reasoning task which is composed of eight different sub-tasks including BoolQ (Clark
 323 et al., 2019), PIQA (Bisk et al., 2020), Social IQa (Sap et al., 2019), HellaSwag (Zellers et al.,

324 Table 2: The results of the proposed Dual LoRA and other competitors with RoBERTa base/large
 325 and DeBERTa XXL on GLUE datasets. For all matrices, higher is better.
 326

328 Model	329 Methods	330 Trainable 331 Param. (M)	332 GLUE								
			MNLI	SST-2	MRPC	CoLA	QNLI	QQP	RTE	STS-B	Avg.
333 RoB _{base}	FFT	125.0	87.6	94.8	90.2	63.6	92.8	91.9	78.7	91.2	86.4
	GaLore ($r = 8$)	0.3	87.2	94.4	92.0	61.8	92.3	91.2	79.1	90.8	85.9
	GaRare ($r = 8$)	0.3	87.2	94.4	91.5	61.1	92.3	90.9	79.3	90.3	85.9
	Delta-LoRA ($r = 8$)	0.3	87.5	95.1	90.2	63.8	93.1	90.9	87.0	91.6	87.4
	CorDA ($r = 128$)	21	-	93.1	89.7	59.6	91.5	-	88.1	90.2	-
	VeRA	0.3	-	91.9	88.4	59.9	90.5	-	74.9	90.4	-
	RandLoRA	0.7	-	92.2	88.0	59.4	91.3	-	74.7	90.3	-
	LoRA ($r = 8$)	0.3	87.0	94.6	89.2	60.9	92.9	90.7	92.0	91.1	86.1
	LoRA ($r = 16$)	0.6	87.0	95.1	89.0	63.9	93.0	91.2	83.4	91.1	86.7
	LoRA+ ($r = 16$)	0.6	87.8	95.2	90.4	65.9	92.6	91.2	82.3	91.4	87.1
334 RoB _{large}	DoRA ($r = 16$)	0.6	87.7	95.3	87.8	64.8	92.6	90.8	82.2	90.8	86.5
	Dual LoRA ($r = 8$)	0.6	87.8	95.8	91.7	67.8	93.3	90.7	88.1	91.7	88.3
335 DeB _{XXL}	FFT	355.0	90.2	96.4	90.9	68.0	94.7	92.2	86.6	92.4	88.9
	GaLore ($r = 16$)	1.6	90.8	96.1	91.7	68.3	95.7	91.9	87.0	92.5	89.3
	GaRare ($r = 16$)	1.6	91.3	96.2	91.7	67.9	94.6	91.8	87.4	92.3	89.2
	VeRA	0.3	-	95.8	89.3	65.3	94.1	-	81.6	91.8	-
	RandLoRA	1.8	-	95.5	90.1	67.4	94.1	-	84.5	91.4	-
	LoRA ($r = 8$)	0.8	90.2	95.6	89.5	63.8	94.5	91.5	88.8	92.5	88.3
	LoRA ($r = 16$)	1.6	90.2	95.9	90.9	66.0	94.4	91.6	87.4	92.3	88.6
	LoRA+ ($r = 16$)	1.6	90.3	96.3	91.4	68.7	94.7	91.6	88.8	92.5	89.3
	DoRA ($r = 16$)	1.6	90.5	96.2	89.7	68.5	92.6	91.5	89.2	92.3	88.8
	Dual LoRA ($r = 8$)	1.6	90.5	96.4	91.9	70.2	95.1	91.2	89.5	92.6	89.7

349 2019), WinoGrande (Sakaguchi et al., 2021), ARC-easy/challange (Clark et al., 2018) and Open-
 350 BookQA (Mihaylov et al., 2018). Similarly to DoRA, we merge the training sets from all sub-tasks
 351 to get the final training set and perform evaluations on their own testing datasets for each task.

352 For the baseline models, we use LLaMA-7B/13B (Touvron et al., 2023a), LLaMA2-7B (Touvron
 353 et al., 2023b), LLaMA3-8B (Dubey et al., 2024), and LLaMA3-70B-Instruct (Dubey et al., 2024).
 354 We halve the rank of our low-rank matrices to ensure that the same number of trainable parameters
 355 are used compared to other LoRA-based methods. We tune the learning rate for our method, and all
 356 other training hyper-parameters are kept unchanged as in DoRA in order to make a fair comparison.
 357 We train one epoch for LLaMA3-70B-instruct, and three epochs for other baseline models.

358 **Results.** The results in Tab. 1 show that we can consistently outperform DoRA with less trainable
 359 parameters on all of the baseline models. For example, Dual LoRA enhances the average accuracy
 360 by 0.9%/0.6% compared to the previous best result on LLaMA-7B/13B. The performance gains are
 361 still notable on LLaMA2-7B, LLaMA3-8B and LLaMA3-70B-Instruct, which are 0.8%, 0.2% and
 362 0.8%.

371 372 4.2 NEURAL LANGUAGE UNDERSTANDING (NLU)

373 **Datasets and baseline models.** We evaluate our method on a widely used natural language under-
 374 standing dataset GLUE. It consists of eight different datasets includes MNLI (Williams et al., 2017),
 375 SST-2 (Socher et al., 2013), MRPC (Dolan & Brockett, 2005), CoLA (Warstadt, 2019), QNLI (Ra-
 376 jpurkar et al., 2018), QQP, RTE, and STS-B (Cer et al., 2017). The diversity makes the GELU
 377 benchmark a robust dataset for evaluating LLMs on NLU tasks.

378 For the baseline models, we use RoBERTa base/large (Liu, 2019) and DeBERTa XXL (He et al.,
 379 2020) as pretrained baseline models from the HuggingFace Transformers library (Wolf et al., 2020).
 380 Similarly to LoRA, we initialize the model to the LoRA-adapted MNLI checkpoint for MRPC, RTE,
 381 and STSB rather than the pre-trained baseline model. All other training parameters are the same as
 382 LoRA except for the learning rate.

383 **Results.** As shown in Tab. 2, the proposed Dual LoRA shows state-of-the-art results on all three
 384 baseline models. For example, on the small model RoBERTa base we can defeat previous methods
 385 LoRA, LoRA+, and DoRA by 1.6%, 1.2%, and 1.8% average accuracy. Similarly, on medium-sized
 386 model RoBERTa large, our Dual LoRA surpasses LoRA, LoRA+, and DoRA by 1.1%, 0.4%, and
 387 0.9%. On DeBERTa XXL model with over 1500M total parameters, Dual LoRA can still exceed
 388 LoRA, LoRA+, and DoRA by 1.3%, 1.1%, and 1.1%. Note that we can even surpass the FFT meth-
 389 ods on these baseline models by 1.9%, 0.8%, and 0.5%, which shows the priority of the proposed
 390 method.

391 4.3 ABLATION STUDY

392 We conduct several ablation studies to further verify the effectiveness of the proposed method.

393 **Dealing with the sign function.** In the previous section and experiments, we use the straight-
 394 through estimator (STE) method (Bengio et al., 2013) to compute the gradient of the sign function.
 395 Note that the sign function is a standard operator in the area of binary neural networks (BNNs), and
 396 there are many studies on dealing with the forward and backward passes of the sign function. For
 397 example, XNOR-Net (Rastegari et al., 2016) scales the weights after binarized:

$$400 \quad \text{Forward: } x_b = \text{Sign}(x) \times \mathbf{E}_F(|x|), \quad \text{Backward: } \frac{\partial \mathcal{L}}{\partial x} = \frac{\partial \mathcal{L}}{\partial x_b}, \quad (13)$$

401 where $\mathbf{E}_F(|x|)$ is the mean of the absolute value of each output channel of weights. Dorefa-
 402 Net (Zhou et al., 2016) uses a constant scalar to scale all of the weights instead of doing channel-wise
 403 scaling:

$$404 \quad \text{Forward: } x_b = \text{Sign}(x) \times \mathbf{E}(|x|), \quad \text{Backward: } \frac{\partial \mathcal{L}}{\partial x} = \frac{\partial \mathcal{L}}{\partial x_b}. \quad (14)$$

405 The experimental results of using different meth-
 406 ods to deal with the sign function are shown
 407 in Tab. 3. The original STE method performs
 408 best among different methods. This conclusion
 409 is different from that in BNNs. We analyze that
 410 this is because both XNOR-Net (Rastegari et al.,
 411 2016) and Dorefa-Net (Zhou et al., 2016) modify
 412 the forward pass of sign function by adding per-
 413 channel scales or a constant scale, which contam-
 414 inate the direction group and make it unable to
 415 focus on giving the correct binary outputs.

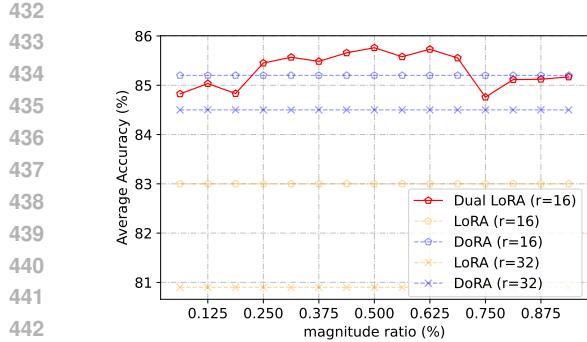
416 Table 3: Different methods are used to deal
 417 with sign function. The experiments are con-
 418 ducted on LLaMA-7B and the commonsense
 419 reasoning dataset.

Method	Trainable Params (%)	Avg.
STE (ours)	1.64	78.9
XNOR-Net	1.64	77.9
Dorefa-Net	1.64	78.1

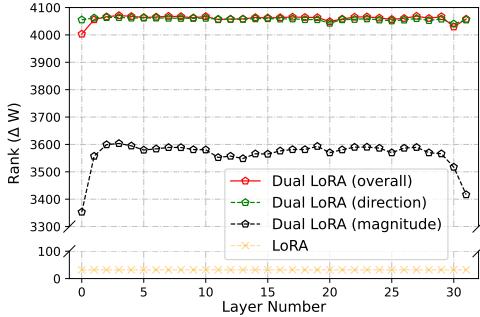
420 **The influence of r_1 and r_2 .** As shown in Sec. 3.3,
 421 r_1 and r_2 are the ranks of the low-rank matrices
 422 in the magnitude and direction groups, respectively. In previous experiments, we set $r_1 = r_2 = r$ in
 423 default to avoid introducing new hyper-parameters compared to other LoRA-based methods. Thus,
 424 in this section we dive deeper into the influence of r_1 and r_2 .

425 Specifically, we keep the total trainable parameters unchanged by setting $r_1 + r_2 = 2r$, and adjust
 426 the ratio of parameters in the magnitude group and direction group, which are controlled by r_1 and
 427 r_2 , respectively. We conduct experiments on the commonsense reasoning dataset with LLaMA3-
 428 8B as our baseline model and set 15 different ratios. Other training hyper-parameters are kept the
 429 same as in the previous experiments. The results are shown in Fig. 2. We can see that the proposed
 430 Dual LoRA can consistently outperform LoRA and DoRA when having a roughly balanced param-
 431 eter ratio between the magnitude group and direction group (from 25% to 70%), which shows the
 432 robustness of our method.

433 More ablation studies are shown in Appendix B.



432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
Figure 2: Average accuracy on the common- sense reasoning datasets using LLaMA3-8B as the baseline model with $r_1 = \{2, 4, \dots, 30\}$ and $r_2 = 32 - r_1$ in the experiments. The proposed Dual LoRA, the blue/orange lines represent DoRA/LoRA with different ranks.



432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
Figure 3: The average rank of ΔW for LoRA, magnitude group of Dual LoRA, direction group of Dual LoRA, and the overall Dual LoRA. The experiments are conducted on LLaMA2-7B.

5 ANALYSIS OF THE RANK OF THE UPDATE MATRIX

In this section, we give an analyze of the rank of the update matrix Eq. 11 to further show the priority of our method. Note that for a given matrix $X \in \mathbb{R}^{m \times n}$, $\text{Rank}(X) \leq \min(m, n)$ always holds true. Thus, in the original LoRA, given $A \in \mathbb{R}^{r \times k}$ and $B \in \mathbb{R}^{d \times r}$ with $r \ll \min(d, k)$, the rank of the update matrix $\Delta W = \frac{\alpha}{r} \cdot BA$ is upper bounded by:

$$\text{Rank}(\Delta W) = \text{Rank}(BA) \leq \min(\text{Rank}(A), \text{Rank}(B)) \leq r. \quad (15)$$

In the proposed Dual LoRA, however, we found that the rank of the update matrix can achieve a higher upper bound. Specifically, given two low-rank matrices $A \in \mathbb{R}^{r_1 \times k}$ and $B \in \mathbb{R}^{d \times r_1}$ in the magnitude group and two low-rank matrices $C \in \mathbb{R}^{r_2 \times k}$ and $D \in \mathbb{R}^{d \times r_2}$ in the direction group with $r_1, r_2 \ll \min(d, k)$, the rank of the update matrix $\Delta W' = \frac{\alpha}{\sqrt{r_1 r_2}} \text{ReLU}(BA) \odot \text{Sign}(DC)$ is:

$$\begin{aligned} \text{Rank}(\Delta W') &= \text{Rank}(\text{ReLU}(BA) \odot \text{Sign}(DC)) \\ &\leq \text{Rank}(\text{ReLU}(BA)) \times \text{Rank}(\text{Sign}(DC)) \\ &\leq \min(k, d)^2. \end{aligned} \quad (16)$$

Note that the $\text{ReLU}(\cdot)$ and $\text{Sign}(\cdot)$ operations break the low-rank limitation of the original input matrix and derive output matrices with high rank.

In Fig. 3, we explicitly show the average rank of the update matrix in LoRA and the proposed Dual LoRA (magnitude group, direction group, and overall) over different layers. The experiments are conducted on LLaMA2-7B. We can see that the update matrix and the direction group almost achieve full rank (4096). The magnitude group has a relatively lower rank but is still much larger than that in LoRA. The results show the priority of our method from the perspective of matrix rank.

6 CONCLUSION

Original LoRA and its followers fine-tune the model without incorporating any prior knowledge and share a common drawback: as the trainable parameters are limited, the model accuracy is unsatisfactory. In this paper, we propose a new LoRA-based method called Dual LoRA, which incorporates an inductive bias into the original LoRA and improve the performance by introducing four low-rank matrices and separating them into the magnitude group and the direction group. The former controls the amplitude and whether or not we should update a parameter, and the latter decides whether or not this parameter should be updated in a positive or negative direction. Parameters in two groups are combined together to simulate the parameter updating process of FFT with gradient-based optimization methods. Experimental results on a wide range of NLP tasks and baseline models show that our Dual LoRA can consistently outperform LoRA and other state-of-the-art methods such as LoRA+ and DoRA with the same or less number of trainable parameters.

486 REFERENCES
487

488 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Ale-
489 man, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical
490 report. *arXiv preprint arXiv:2303.08774*, 2023.

491 Paul Albert, Frederic Z Zhang, Hemanth Saratchandran, Cristian Rodriguez-Opazo, Anton van den
492 Hengel, and Ehsan Abbasnejad. Randlora: Full-rank parameter-efficient fine-tuning of large mod-
493 els. *International Conference on Learning Representations*, 2025.

494

495 Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients
496 through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.

497

498 Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, et al. Piqa: Reasoning about physical com-
499 monsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*,
500 volume 34, pp. 7432–7439, 2020.

501

502 Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of
503 COMPSTAT'2010: 19th International Conference on Computational StatisticsParis France, Au-
gust 22-27, 2010 Keynote, Invited and Contributed Papers*, pp. 177–186. Springer, 2010.

504

505 Daniel Cer, Mona Diab, Eneko Agirre, Inigo Lopez-Gazpio, and Lucia Specia. Semeval-2017 task
506 1: Semantic textual similarity-multilingual and cross-lingual focused evaluation. *arXiv preprint
507 arXiv:1708.00055*, 2017.

508

509 Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina
510 Toutanova. Boolq: Exploring the surprising difficulty of natural yes/no questions. *arXiv preprint
arXiv:1905.10044*, 2019.

511

512 Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and
513 Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge.
514 *arXiv preprint arXiv:1803.05457*, 2018.

515

516 Bill Dolan and Chris Brockett. Automatically constructing a corpus of sentential paraphrases. In
517 *Third international workshop on paraphrasing (IWP2005)*, 2005.

518

519 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha
520 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models.
521 *arXiv preprint arXiv:2407.21783*, 2024.

522

523 Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. The webnlg
challenge: Generating text from rdf data. In *10th International Conference on Natural Language
Generation*, pp. 124–133. ACL Anthology, 2017.

524

525 Yongchang Hao, Yanshuai Cao, and Lili Mou. Flora: Low-rank adapters are secretly gradient
526 compressors. *International conference on machine learning*, 2024.

527

528 Soufiane Hayou, Nikhil Ghosh, and Bin Yu. Lora+: Efficient low rank adaptation of large models.
529 *arXiv preprint arXiv:2402.12354*, 2024.

530

531 Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Towards a
unified view of parameter-efficient transfer learning. *arXiv preprint arXiv:2110.04366*, 2021.

532

533 Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. Deberta: Decoding-enhanced bert
with disentangled attention. *arXiv preprint arXiv:2006.03654*, 2020.

534

535 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,
536 and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint
537 arXiv:2106.09685*, 2021a.

538

539 Shengding Hu, Ning Ding, Huadong Wang, Zhiyuan Liu, Jingang Wang, Juanzi Li, Wei Wu, and
Maosong Sun. Knowledgeable prompt-tuning: Incorporating knowledge into prompt verbalizer
for text classification. *arXiv preprint arXiv:2108.02035*, 2021b.

540 Zhiqiang Hu, Lei Wang, Yihuai Lan, Wanyu Xu, Ee-Peng Lim, Lidong Bing, Xing Xu, Soujanya
 541 Poria, and Roy Ka-Wei Lee. Llm-adapters: An adapter family for parameter-efficient fine-tuning
 542 of large language models. *arXiv preprint arXiv:2304.01933*, 2023.

543

544 Ting Jiang, Shaohan Huang, Shengyue Luo, Zihan Zhang, Haizhen Huang, Furu Wei, Weiwei Deng,
 545 Feng Sun, Qi Zhang, Deqing Wang, et al. Mora: High-rank updating for parameter-efficient fine-
 546 tuning. *arXiv preprint arXiv:2405.12130*, 2024.

547 Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*,
 548 2014.

549 Dawid J Kopiczko, Tijmen Blankevoort, and Yuki M Asano. Vera: Vector-based random matrix
 550 adaptation. *International Conference on Learning Representations*, 2024.

551

552 Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt
 553 tuning. *arXiv preprint arXiv:2104.08691*, 2021.

554

555 Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv
 556 preprint arXiv:2101.00190*, 2021.

557 Zhaojiang Lin, Andrea Madotto, and Pascale Fung. Exploring versatile generative language model
 558 via parameter-efficient transfer learning. *arXiv preprint arXiv:2004.03829*, 2020.

559

560 Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances
 561 in neural information processing systems*, 36, 2024a.

562 Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-
 563 Ting Cheng, and Min-Hung Chen. Dora: Weight-decomposed low-rank adaptation. *arXiv
 564 preprint arXiv:2402.09353*, 2024b.

565

566 Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Lam Tam, Zhengxiao Du, Zhilin Yang, and Jie Tang. P-
 567 tuning v2: Prompt tuning can be comparable to fine-tuning universally across scales and tasks.
 568 *arXiv preprint arXiv:2110.07602*, 2021.

569

570 Xiao Liu, Heyan Huang, Ge Shi, and Bo Wang. Dynamic prefix-tuning for generative template-
 571 based event extraction. *arXiv preprint arXiv:2205.06166*, 2022.

572

573 Xu-Hui Liu, Yali Du, Jun Wang, and Yang Yu. On the optimization landscape of low rank adaptation
 574 methods for large language models. In *International Conference on Learning Representations*.

575

576 Yinhan Liu. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint
 577 arXiv:1907.11692*, 364, 2019.

578

579 Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct
 580 electricity? a new dataset for open book question answering. *arXiv preprint arXiv:1809.02789*,
 581 2018.

582

583 Linyong Nan, Dragomir Radev, Rui Zhang, Amrit Rau, Abhinand Sivaprasad, Chiachun Hsieh,
 584 Xiangru Tang, Aadit Vyas, Neha Verma, Pranav Krishna, et al. Dart: Open-domain structured
 585 data record to text generation. *arXiv preprint arXiv:2007.02871*, 2020.

586

587 Jekaterina Novikova, Ondřej Dušek, and Verena Rieser. The e2e dataset: New challenges for end-
 588 to-end generation. *arXiv preprint arXiv:1706.09254*, 2017.

589

590 Cheng Peng, XI Yang, Kaleb E Smith, Zehao Yu, Aokun Chen, Jiang Bian, and Yonghui Wu. Model
 591 tuning or prompt tuning? a study of large language models for clinical concept and relation
 592 extraction. *Journal of biomedical informatics*, 153:104630, 2024.

593

594 Pranav Rajpurkar, Robin Jia, and Percy Liang. Know what you don't know: Unanswerable questions
 595 for squad. *arXiv preprint arXiv:1806.03822*, 2018.

596

597 Mohammad Rastegari, Vicente Ordóñez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet
 598 classification using binary convolutional neural networks. In *European conference on computer
 599 vision*, pp. 525–542. Springer, 2016.

594 Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adver-
 595 sarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106, 2021.
 596

597 Maarten Sap, Hannah Rashkin, Derek Chen, Ronan LeBras, and Yejin Choi. Socialqa: Common-
 598 sense reasoning about social interactions. *arXiv preprint arXiv:1904.09728*, 2019.

599 Chongjie Si, Xuehui Wang, Xue Yang, Zhengqin Xu, Qingyun Li, Jifeng Dai, Yu Qiao, Xi-
 600 aokang Yang, and Wei Shen. Flora: Low-rank core space for n-dimension. *arXiv preprint*
 601 *arXiv:2405.14739*, 2024.

602

603 Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng,
 604 and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment
 605 treebank. In *Proceedings of the 2013 conference on empirical methods in natural language pro-*
 606 *cessing*, pp. 1631–1642, 2013.

607

608 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée
 609 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and
 610 efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023a.

611 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Niko-
 612 lay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open founda-
 613 tion and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023b.

614

615 A Warstadt. Neural network acceptability judgments. *arXiv preprint arXiv:1805.12471*, 2019.

616

617 Adina Williams, Nikita Nangia, and Samuel R Bowman. A broad-coverage challenge corpus for
 618 sentence understanding through inference. *arXiv preprint arXiv:1704.05426*, 2017.

619 Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi,
 620 Pierrick Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. Transformers: State-of-the-art
 621 natural language processing. In *Proceedings of the 2020 conference on empirical methods in*
 622 *natural language processing: system demonstrations*, pp. 38–45, 2020.

623

624 Zhengxuan Wu, Aryaman Arora, Zheng Wang, Atticus Geiger, Dan Jurafsky, Christopher D Man-
 625 ning, and Christopher Potts. Reft: Representation finetuning for language models. *arXiv preprint*
 626 *arXiv:2404.03592*, 2024.

627

628 Yibo Yang, Xiaojie Li, Zhongzhu Zhou, Shuaiwen Song, Jianlong Wu, Liqiang Nie, and Bernard
 629 Ghanem. Corda: Context-oriented decomposition adaptation of large language models for task-
 630 aware parameter-efficient fine-tuning. *Advances in Neural Information Processing Systems*, 37:
 631 71768–71791, 2024.

632

633 Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a ma-
 634 chine really finish your sentence? *arXiv preprint arXiv:1905.07830*, 2019.

635

636 Hongyi Zhang, Zuchao Li, Ping Wang, and Hai Zhao. Selective prefix tuning for pre-trained lan-
 637 guage models. In *Findings of the Association for Computational Linguistics ACL 2024*, pp. 2806–
 2813, 2024.

638

639 Jiawei Zhao, Zhenyu Zhang, Beidi Chen, Zhangyang Wang, Anima Anandkumar, and Yuandong
 640 Tian. Galore: Memory-efficient ILM training by gradient low-rank projection. In *International*
 641 *conference on machine learning*. PMLR, 2024.

642

643 Shuchang Zhou, Yuxin Wu, Zekun Ni, Xinyu Zhou, He Wen, and Yuheng Zou. Dorefa-net: Train-
 644 ing low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint*
 645 *arXiv:1606.06160*, 2016.

646

647 Bojia Zi, Xianbiao Qi, Lingzhi Wang, Jianan Wang, Kam-Fai Wong, and Lei Zhang. Delta-
 648 lora: Fine-tuning high-rank parameters with the delta of low-rank matrices. *arXiv preprint*
 649 *arXiv:2309.02411*, 2023.

648 A EXPERIMENTS ON NEURAL LANGUAGE GENERATION (NLG)
649

650 **Datasets.** We conduct experiments on E2E NLG Challenge (Novikova et al., 2017), DART (Nan
651 et al., 2020) and WebNLG (Gardent et al., 2017) datasets. The E2E dataset consists of approximately
652 42,000 training data, 4,600 validation data, and 4,600 test data. Each sample is composed of a
653 sequence of slot-value pairs (x, y) and a corresponding natural language reference text. DART is
654 an open-domain data-to-text dataset with around 82,000 samples, each sample is structured as a
655 sequence of entity-relation-entity triple. WebNLG has a total of 22,000 examples from 14 different
656 categories, each sample is structured as a sequence of subject-property-object triple.
657

658 Table 4: The results of the proposed Dual LoRA and other competitors with GPT2_M and GPT2_L
659 on the E2E NLG Challenge dataset. For all matrices, higher is better.
660

661 Model	662 Methods	663 Trainable Parameters (M)	664 E2E NLG Challenge				
			665 BLEU	666 NIST	667 MET	668 ROUGE-L	669 CIDEr
664 GPT2_M	FFT	354.92	68.2	8.62	46.2	71.0	2.47
	Adapter-L	11.09	68.9	8.71	46.1	71.3	2.47
	Prefix	0.35	69.7	8.81	46.1	71.4	2.49
	LoRA ($r = 4$)	0.35	68.9	8.69	46.4	71.4	2.52
	LoRA ($r = 8$)	0.70	69.9	8.77	46.8	71.7	2.50
	LoRA+ ($r = 8$)	0.70	70.2	8.81	46.6	71.6	2.53
	DoRA ($r = 8$)	0.71	69.5	8.75	46.4	71.4	2.52
671 GPT2_L	Dual LoRA ($r = 4$)	0.70	70.6	8.86	46.9	72.4	2.56
	FFT	774.03	68.5	8.78	46.0	69.9	2.45
	Adapter-L	23.00	69.1	8.68	46.3	71.4	2.49
	Prefix	0.77	70.3	8.85	46.2	71.7	2.47
	LoRA ($r = 4$)	0.77	70.3	8.85	46.8	71.9	2.52
	LoRA ($r = 8$)	1.54	70.0	8.80	46.8	71.7	2.54
	LoRA+ ($r = 8$)	1.54	70.0	8.83	46.8	71.9	2.53
672 GPT2_L	DoRA ($r = 8$)	1.56	69.8	8.78	46.6	71.6	2.52
	Dual LoRA ($r = 4$)	1.54	70.6	8.86	47.1	72.5	2.54

673 Table 5: The results of the proposed Dual LoRA and other competitors with GPT2_M and GPT2_L
674 on DART and WebNLG datasets. The up arrow indicates that the higher is better, and the down
675 arrow indicates that the lower is better.
676

677 Model	678 Methods	679 Trainable Parameters (M)	680 DART			681 WebNLG		
			682 BLEU↑	683 MET↑	684 TER↓	685 BLEU↑	686 MET↑	687 TER↓
688 GPT2_M	LoRA ($r = 4$)	0.35	47.4	0.36	0.47	55.0	0.37	0.39
	LoRA ($r = 8$)	0.70	47.5	0.36	0.47	55.6	0.38	0.39
	LoRA+ ($r = 8$)	0.70	47.6	0.36	0.47	56.1	0.38	0.39
	DoRA ($r = 8$)	0.71	47.0	0.36	0.48	53.5	0.36	0.40
	Dual LoRA ($r = 4$)	0.70	48.3	0.36	0.47	56.6	0.38	0.38
692 GPT2_L	LoRA ($r = 4$)	0.35	47.7	0.36	0.47	57.1	0.37	0.38
	LoRA ($r = 8$)	0.70	47.5	0.36	0.47	57.6	0.38	0.38
	LoRA+ ($r = 8$)	0.70	47.7	0.36	0.47	57.5	0.38	0.38
	DoRA ($r = 8$)	0.71	47.2	0.36	0.47	57.4	0.39	0.38
	Dual LoRA ($r = 4$)	0.70	48.4	0.36	0.47	57.7	0.39	0.38

693 **Results.** As in previous experiments, we reduce our rank to half of the other LoRA-based methods
694 to ensure the same number of trainable parameters. All other hyper-parameters are the same as
695 LoRA, except that we tune the learning rate. As the results shown in Tab. 4 and Tab. 5, Dual LoRA
696 can consistently outperform all other competitors on three different datasets with baseline model
697 GPT2_M and GPT2_L.

702 B MORE ABLATION STUDIES 703

704 In this section, we conduct more ablation studies on the proposed Dual LoRA.
705

706 B.1 DIFFERENT FORMS OF THE UPDATE MATRIX 707

708 Specifically, we investigate the following settings:
709

- 710 • Setting 1: Remove $\text{ReLU}(\cdot)$ function in Dual LoRA, which means $\Delta W = \frac{\alpha}{\sqrt{r_1 r_2}} (BA) \odot$
711 $\text{Sign}(DC)$.
- 712 • Setting 2: Remove $\text{Sign}(\cdot)$ function in Dual LoRA, which means $\Delta W =$
713 $\frac{\alpha}{\sqrt{r_1 r_2}} \text{ReLU}(BA) \odot (DC)$.
- 714 • Setting 3: Replace the output of the direction group with a random-initialized binary matrix
715 and fix this matrix during training. Given W_b as the random initialized binary matrix, we
716 have $\Delta W = \frac{\alpha}{r_1} \text{ReLU}(BA) \odot W_b$.
717

718 Table 6: The results of the previous settings on the commonsense reasoning dataset, with LLaMA-
719 7B as the base model.
720

722 Methods	723 Trainable 724 Param. (%)	725 Commonsense Reasoning Datasets								
		726 BoolQ	727 PIQA	728 SIQA	HellaS	729 WinoG	730 ARC-e	731 ARC-c	732 OBQA	733 Avg.
Dual LoRA ($r = 32$)	1.64	70.0	83.2	79.5	87.3	83.0	81.6	65.2	81.0	78.9
Setting 1 ($r = 32$)	1.64	70.3	82.9	78.8	85.4	83.3	82.1	65.2	79.4	78.4
Setting 2 ($r = 32$)	1.64	70.6	80.8	78.9	84.7	81.8	80.4	65.2	78.8	77.7
Setting 3 ($r = 32$)	0.84	60.4	0.2	1.7	0.2	0.2	0.5	0	0	7.9
Setting 3 ($r = 64$)	1.64	36.3	36.2	6.6	0	14.3	4.9	11.5	0	13.7

734 We conduct experiments on LLaMA-7B and perform an evaluation on the commonsense reasoning
735 dataset. The results are shown in Tab. 6. We can see that removing $\text{ReLU}(\cdot)$ and $\text{Sign}(\cdot)$ functions
736 cause marginal performance drop, and using a random initialized binary matrix to replace the
737 direction group severely degrades the performance, which shows the importance of the proposed
738 architecture.
739

740 B.2 DIFFERENT ACTIVATION FUNCTIONS FOR THE MAGNITUDE GROUP 741

742 In the main paper, we use $\text{ReLU}(\cdot)$ function for the magnitude group. There are other functions that
743 can keep the activation greater than zero, such as $\text{Abs}(x) = |x|$ and $\text{Sigmoid}(x) = 1/(1 + e^{-x})$.
744 We compare them in Tab. 7 on the commonsense reasoning dataset with LLaMA-7B model.
745

746 Table 7: The results of different activation functions for the magnitude group on the commonsense
747 reasoning dataset, with LLaMA-7B as the base model.
748

749 Activation	750 Commonsense Reasoning Datasets								
	751 BoolQ	752 PIQA	753 SIQA	HellaS	754 WinoG	755 ARC-e	756 ARC-c	757 OBQA	Avg.
ReLU (ours)	70.0	83.2	79.5	87.3	83.0	81.6	65.2	81.0	78.9
Abs	68.6	82.9	77.9	74.3	74.0	81.6	65.8	79.6	75.6
Sigmoid	66.5	80.3	78.5	77.8	79.4	75.4	60.9	76.6	74.4

758 We can see that the original setting with $\text{ReLU}(\cdot)$ function performs the best. This is because
759 neither the $\text{Abs}(\cdot)$ function nor the $\text{Sigmoid}(\cdot)$ function can flexibly give zero output. Note that zero
760 output means that we can easily freeze some of the elements that are already well-trained for the
761 downstream tasks in the original weight matrix, which is the advantage of the $\text{ReLU}(\cdot)$ function.
762