# BONE: BLOCK-AFFINE ADAPTATION OF LARGE LAN-GUAGE MODELS

Anonymous authors

Paper under double-blind review

#### ABSTRACT

Low-Rank Adaptation (LoRA) has achieved remarkable training results by freezing the original weights and training only low-rank matrices, establishing itself as the predominant fine-tuning method for LLMs. Many LoRA variants have emerged, yet they lack a design tailored to the characteristics of LLM weights and fail to leverage the original weights effectively. To address the sparsity of LLM weights, and drawing inspiration from GQA and MQA, we propose Block-Affine Adaptation (Bone), a novel PEFT technique distinct from LoRA. By dividing the original weights into multiple subspaces that share a single matrix for weight updates, Bone simplifies the process by requiring the trainable matrix to be initialized to zero, eliminating the need for complex initialization as in some LoRA variants. Compared to LoRA, Bone significantly reduces memory usage and achieves faster computation. Evaluation of both NLU and NLG tasks demonstrates that Bone substantially outperforms LoRA and its variants. Inspired by Pissa, we propose a new theory called "Weight Guide" to better utilize the information embedded in the original weights. This approach extracts valuable information through a linear transformation of the original weight matrix using a trainable matrix. To validate the effectiveness of "Weight Guide" we combined it with Bone to create a new structure called Block-Affine Transformation (Bat), and ablation experiments confirmed the effectiveness of "Weight Guide".

028 029

031

004

010 011

012

013

014

015

016

017

018

019

021

025

026

027

#### 1 INTRODUCTION

033 Large models have been integrated into various industries, revolutionizing many traditional technologies Radford et al. (2019); Raffel et al. (2020). However, general-purpose large models often struggle to meet the needs of all downstream tasks, making it necessary to fine-tune base models for specific scenarios. Full-scale fine-tuning of large models is computationally costly; for example, finetuning the LLaMA2-7B Touvron et al. (2023) model with bfloat16 Wang & Kanwar (2019) 037 precision requires around 60GB of VRAM. In contrast, PEFT (Parameter-Efficient Fine-Tuning)Xu et al. (2023) techniques could reduce the VRAM requirement to fit into a 24GB VRAM GPU. As a result, numerous PEFT techniques and quantization methods have emerged to reduce the training 040 costs of large models. LoRA (Low-Rank Adaptation) Hu et al. (2021) has become one of the most 041 popular PEFT methods due to its small tunable parameter size, its effectiveness, and the possibility 042 of zero inference overhead after finetuning. 043

LoRA significantly reduces memory usage by freezing the original weights W and updating two 044 low-rank matrices A and B. Typically, either A or B is initialized to zero, ensuring that the initial state of LoRA is consistent with the pre-trained model, The figure 2 illustrates the structure visu-046 alization. However, extensive experiments Ding et al. (2023); Liu et al. (2024b); Biderman et al. 047 (2024) have shown that LoRA's convergence is significantly slower compared to full fine-tuning. 048 This slow convergence is likely due to the small gradients caused by the zero initialization of either Aor B. To address this issue, researchers have proposed several LoRA variants, such as LoRA+ Hayou et al. (2024), PISSA Meng et al. (2024), and LoRA-GAWang et al. (2024). Despite their excellent 051 performance, these LoRA variants inevitably introduce complexity into the fine-tuning process and reduce structural flexibility. For example, LoRA+ requires manual adjustment of different learning 052 rates for A and B; PISSA necessitates performing SVD decomposition on W at the beginning of training, which can be time-consuming when the model parameters are large. As research on LoRA becomes increasingly saturated, it is essential to explore new directions for PEFT techniques and effectively leverage the unique characteristics of LLMs.

In this work, we first designed Block-Affine Adaptation (Bone), which divides the original weights into multiple subspaces that share a single trainable matrix for updates, making it more efficient than LoRA's two low-rank matrices. Secondly, to address the limitations of LoRA and fully utilize the information in the original weights, we propose a new theory, "Weight Guide". "Weight Guide" enables feature extraction by applying a simple linear transformation to the original weights. Finally, to validate the effectiveness of "Weight Guide", we combined Bone with "Weight Guide" to develop Block-Affine Transformation (Bat).

063 Our extensive evaluation shows that Bone not only excels in both Natural Language Understand-064 ing (NLU) and Natural Language Generation (NLG) tasks but also significantly outperforms LoRA 065 and its variants. Additionally, Bone retains the advantages of LoRA, such as ease of use and no 066 additional computational overhead for the model. In NLG tasks, Bone demonstrates superior per-067 formance, with evaluation metrics surpassing even the strong LoRA variant Pissa across the board. 068 To verify the feasibility of Bone, we conducted experiments on two different LLM architectures 069 (LLaMA2 Xu et al. (2023), RWKV6 Peng et al. (2024)). As shown in Figure 1, Bone achieves the fastest convergence, and the results on the test set demonstrate its superior data fitting and general-071 ization capabilities (Table 2). As a completely new structure distinct from LoRA, Bone offers clear improvements in computational efficiency and memory savings, as detailed in table 7. 072

### 073 Our contributions can be summarized as follows

074

075

076

077

078

079

081

**1.** We propose a novel PEFT technique called Block-Affine Adaptation (Bone). Bone outperforms LoRA and its variants across various fine-tuning tasks. Additionally, Bone is more memory-efficient and computationally faster compared to LoRA.

**2.** To effectively leverage the implicit information in the original weights, we propose the "Weight Guide" theory. By integrating this theory with Bone, we design Block-Affine Transformation (Bat), demonstrating the effectiveness of "Weight Guide".



Figure 1: The left image shows the loss curve for LLaMA2-7B fine-tuned on the MetaMathQA dataset, with the first 100 steps highlighted for closer observation. Comparing the loss curves reveals that Bone demonstrates superior fitting ability across various architectures and parameter settings. Additionally, Bone exhibits a rapid decrease in loss within the first 100 steps, highlighting its effectiveness. The left image is a demonstration of the training state.

107

### 108 2 RELATED WORKS

109 110

The PEFT (Parameter-Efficient Fine-Tuning) techniques are diverse and include approaches like adapter tuningHoulsby et al. (2019); He et al. (2022); Wang et al. (2022); Pfeiffer et al. (2020), prefix tuningLiu et al. (2023); Li & Liang (2021), prompt tuningBrown (2020); Liu et al. (2023); Lester et al. (2021); Razdaibiedina et al. (2023); Li & Liang (2021), LoRAHu et al. (2021); Meng et al. (2024); Wang et al. (2024); Si et al. (2024), and layer-freezing methods such as LISA.

The adapter method does not require fine-tuning all the parameters of the pre-trained model. Instead, it introduces a small number of task-specific parameters to store knowledge related to that task, thereby reducing the computational demands of model fine-tuning. Prefix tuning is a lightweight fine-tuning method for generative tasks. It adds a continuous, task-specific vector sequence, called a prefix, to the input. Unlike prompts, prefixes are entirely composed of free parameters and do not correspond to actual tokens. Compared to traditional fine-tuning, prefix tuning only optimizes a set of prefixes related to specific tasks, without any change to the original model.

Prompt tuning defines a unique prompt for each task, prepending it to the input data while freezing the pre-trained model during training. LoRA allows us to indirectly train certain dense layers in a neural network by optimizing low-rank decomposition matrices that adapt these layers while keeping the pre-trained weights unchanged. This approach significantly addresses the inference latency introduced by adapters and the fitting challenges of prefix tuning.

127 In LoRA, the adapter matrices A and B are updated with the same learning rate, but using the same 128 rate for both may not effectively learn the features. LoRA+ extended this method by introducing in-129 dependent learning rates for matrices A and B with a fixed ratio, improving the method's efficiency. 130 The DoRA Liu et al. (2024b) method combines weight decomposition to achieve learning capabili-131 ties similar to full fine-tuning without sacrificing LoRA's inference efficiency. PiSSA optimizes the 132 compact parameter space by representing the matrices in the model as the product of two trainable 133 matrices, augmented with a residual matrix for error correction. Using Singular Value Decomposition (SVD), PiSSA initializes the dominant singular values and vectors to train these matrices, 134 while keeping the residual matrix static during fine-tuning. OLoRA Büyükakyüz (2024) leverages 135 QR decomposition to initialize the adaptation matrices during the fine-tuning process, ensuring that 136 these matrices are orthogonal. This orthogonal initialization helps maintain the stability of the pa-137 rameter space during optimization. LoRA-GA and PiSSA are similar in form, but they differ in that 138 LoRA-GA initializes A and B by computing the initial gradient, thereby closely approximating full 139 fine-tuning. 140

141 142

143

3 Method

- 144 3.1 MOTIVATION
- 146 The primary motivations are as follows:

Firstly, current research on LoRA has reached a saturation point, underscoring the need for a new direction in Parameter-Efficient Fine-Tuning (PEFT) techniques.

- Secondly, existing PEFT methods lack a thorough exploration and effective utilization of the original
   model weights.
- 152

154

**153 3.2 BONE: BLOCK-AFFINE UPDATE MATRICES** 

To design a novel PEFT structure that rivals LoRA, we conducted an in-depth study of dense matrices. Current theories suggest that LLM weights exhibit sparse properties. For example, LoRA utilizes low-rank matrices to update dense matrices, while techniques like GQA Ainslie et al. (2023) and MQA Shazeer (2019) employ shared mechanisms to reduce parameter overhead with remarkable effectiveness.

Inspired by these approaches, we hypothesize that dividing LLM weights into multiple subspaces
 allows different subspaces to share a single low-rank matrix for updates. As illustrated in the figure
 we name this structure Block-Affine Adaptation (Bone).

It is evident that the ways to partition LLM weights are diverse, as illustrated in the figure. Different partitioning methods influence the shape of the Bone matrix, the computation logic, and ultimately, the processing speed. In this work, we propose the Best Bone structure. For a pre-trained weight matrix  $W \in \mathbb{R}^{d \times k}$ , we update it using a single matrix:  $W + \Delta W = W + expand(B)$ , where  $B \in \mathbb{R}^{r \times d}$ , and the rank  $r \ll k$ . This design not only retains the simplicity and ease of use characteristic of LoRA but also achieves higher computational efficiency. The calculation is expressed as follows:

$$Y = Wx + \Delta Wx = Wx + sum(x)B \tag{1}$$



Figure 2: Bone requires training only a single low-rank matrix initialized to zero. Best Bone represents the optimal design structure of Bone. Subsequent experiments default to using Best Bone.

185 186 187

188

179

181

182

183

3.3 "WEIGHT GUIDE"

189 It is well-known that LLMs have extremely large parameters, with complex interactions among 190 their internal weights, but most LoRA variants utilize original weight information minimally, with 191 only a few leveraging it during initialization to enhance performance. To increase the utilization 192 of this information and promote internal interactions among weights, we propose a novel concept called "Weight Guide" This approach ensures that the trainable matrices are consistently guided and 193 constrained by the original weights throughout every step of the training process. As illustrated in 194 Figure 1b, LoRA variants interact with the original weights only once during initialization to extract 195 essential components. In contrast, the Bone structure is guided (constrained) by the original weights 196 after each update, significantly increasing the new weights' utilization of the original weights. As 197 the training steps increase, the influence of the original weights becomes more pronounced. This is evident from the loss curves, which show that the differences between Bone and other methods 199 grow larger in the later stages of training. 200

So, how can we extract the critical information from the original weights? A straightforward ap-201 proach is to apply a learnable matrix directly to the original weights using the Hadamard prod-202 uct. This method essentially assigns a weight to each element, allowing the learnable matrix to 203 autonomously determine the importance of individual elements. Although this method is simple 204 and straightforward, the elements within the weights remain isolated, making it impossible to learn 205 interactions between them. Therefore, we replace the Hadamard product with matrix multiplication, 206 essentially applying a simple linear transformation. This approach links the entire weight matrix 207 space, providing a better chance to discover the optimal solution. The corresponding formula is 208 shown below:

$$\Delta W_{n,n} = W_{n,n} \odot bone_{n,n}$$
  

$$\Delta W_{n,n} = W_{n,n} \otimes bone_{n,n}$$
(2)

210 211

209

212 3.4 BAT: BLOCK-AFFINE TRANSFORMATION

213

In the previous section, we introduced Bone, which is both simple and efficient. However, applying

In the previous section, we introduced Bone, which is both simple and efficient. However, applying the same updates to different subspaces of the weight matrix  $W \in \mathbb{R}^{d \times k}$  is clearly counterintuitive. To allow for more flexibility in the update directions of different subspaces, we combine Bone with <sup>216</sup> "Weight Guide", proposing a new method called Block-Affine Transformation (Bat). To balance <sup>217</sup> computational efficiency, we divide the original weights into smaller subspaces compared to Bone. <sup>218</sup> Here, we name the trainable matrix  $bone \in \mathbb{R}^{r \times d}$ , where its rank r is replaced with block size b. <sup>219</sup> The formula for Bone-col is as follows:

$$W_{k/b,d/b,b,b} = Reshape(W_{d,k})$$
  

$$bone_{d/b,b,b} = Reshape(bone_{b,d})$$
  

$$\Delta W_{d,k} = Reshape(W_{k/b,d/b,b,b} \otimes bone_{d/b,b,b} + bone_{d/b,b,b})$$
(3)

Similar to Bone, Bat can also adopt various grouping methods, as illustrated in the figure 3. In Section 6, we compared the performance of different partitioning strategies for the structure. The formula for Bone-row is as follows:

$$W_{d/b,k/b,b,b} = Reshape(W_{d,k})$$

$$bone_{k/b,b,b} = Reshape(bone_{b,k})$$

$$\Delta W_{d,k} = Reshape(W_{d/b,k/b,b,b} \otimes bone_{k/b,b,b} + bone_{k/b,b,b})$$
(4)



Figure 3: A comparison between the visualizations of LoRA and Bone structures reveals that when n = m in W, setting bone\_b to  $2 \times \text{lora}$  resures that the trainable parameters of both structures are equal. The figure illustrates three different grouping methods for the Bone structure. Grouping by rows and columns allows for seamless adaptation to any LLM structure, making it easier to adapt to different LLM architectures.

#### 4 EXPERIMENTS

In this section, we evaluate the performance of Bone on various benchmark datasets. Initially, we assess Natural Language Understanding (NLU) capabilities using a subset of the GLUE dataset with the robert-base model. Subsequently, we evaluated the Natural Language Generation (NLG) capabilities by fine-tuning the LLM.

The experiments were conducted on 4×NVIDIA 4090 24G GPUs.

4.1 EXPERIMENTS ON NATURAL LANGUAGE UNDERSTANDING

Models and Datasets We fine-tune the RoBERTa-base model on several datasets from the GLUE benchmark, including MNLI, SST-2, CoLA, QNLI, and MRPC. Performance is evaluated on the development set using accuracy as the primary metric.

Implementation Details The experimental hyperparameter settings were aligned with those in
 the LoRA repository, but training was conducted using a single 4090 GPU. Each experiment is
 conducted with 3 different random seeds, and the average performance is reported.

Results As shown in Table 1, Bone demonstrates outstanding performance, particularly on the
 CoLA dataset, where it exhibits significantly faster convergence and superior data-fitting capabilities, far surpassing LoRA and Pissa.

Method	Trainable	MNLI	SST-2	CoLA	QNLI	MRPC
LoRA	0.236%	$85.63_{\pm 0.01}$	$94.03_{\pm0.02}$	$62.40_{\pm 0.71}$	$91.37_{\pm 0.97}$	$87.98_{\pm 0.23}$
Pissa	0.236%	$85.72_{\pm 0.40}$	$93.64_{\pm 0.13}$	$67.28_{\pm 0.59}$	$91.40_{\pm 0.54}$	$88.11_{\pm 0.24}$
Bone	0.236%	$85.71_{\pm 0.32}$	$93.60_{\pm 0.07}$	$72.86_{\pm 3.13}$	$91.43_{\pm 0.76}$	$88.14_{\pm 0.60}$

Table 1: The results of fine-tuning RoBERTa-base using Bone and various LoRA variants were
 compared on a subset of the GLUE benchmark.

#### 4.2 EXPERIMENT ON NATURAL LANGUAGE GENERATION

**Models and Datasets** To verify the generalizability of Bone, we conducted more comprehensive experiments on LLM. we conducted 3 more task finetuning experiments on LLM: *math, code,* and *chat.* 

1. *Math*: We trained our model on a 395k subset of MetaMathQA Yu et al. (2023), a dataset boot-strapped from other math instruction tuning datasets like GSM8K Cobbe et al. (2021) and MATH Yu et al. (2023), with higher complexity and diversity.

289 2. *Code*: We train our model on a 100k subset of CodeFeedback Zheng et al. (2024b), a high290 quality code instruction dataset, removing explanations after code blocks. The model is tested on
291 HumanEval Chen et al. (2021).

292 3. *Chat*: We train our model on a 70k subset of WizardLM-Evol-Instruct Xu et al. (2024). We test
293 our model on the MT-Bench dataset Zheng et al. (2024a), which consists of 80 multi-turn questions
294 designed to assess LLMs on multiple aspects. We used GPT-40 to judge the quality of responses, as
295 shown in lm-sys/FastChat.

**Implementation Details** The hyperparameter settings for this experiment were kept equal, while the train steps were adjusted according to the specific fine-tuning datasets used. It is worth noting that the weights of LLaMA2-7B are not fully symmetric, making it impossible to perfectly align the trainable parameters when comparing Bone and LoRA. To address this, we set the rank r of LoRA to 36 and the rank r of Bone to 64, ensuring that Bone uses fewer parameters than LoRA to demonstrate its superiority. Each experiment is conducted with 2 different random seeds, and the average performance is reported.

**Result** The results, as shown in Table 2 and Figure 1a, demonstrate that Bone outperforms other PEFT methods in terms of convergence speed, data fitting, and generalization capabilities. Bone demonstrates outstanding performance across three different tasks. On LLaMA2-7B, Bone achieves results that surpass Pissa, despite using fewer parameters than LoRA and its variants. On RWKV6-7B, Bone and LoRA have the same number of trainable parameters, yet Bone consistently outperforms LoRA and its variants across all tasks.

4.3 EFFECT OF RANK r

This subsection explores the upper limits of the Bone structure by varying the rank r in the Bone matrix. Comparative experiments were conducted by fine-tuning LLaMA2-7B on the MetaMathQA dataset and validating on GSM8K and Math benchmarks. The test results, as shown in Table 3, demonstrate that the fine-tuning performance improves as the value of b increases. Notably, when r = 16, the Bone structure, with only one-quarter of the trainable parameters compared to PiSSA, surpasses PiSSA's performance on the GSM8k benchmark. However, its performance on the Math benchmark is only 3.73. The GSM8K score surpasses that of PiSSA, but the Math score is significantly lower, indicating The size of r impacts the model's ability to understand unseen data. Based on this observation, we hypothesize that when the rank is too small, it significantly limits the model's generalization ability.

Model	Strategy	Trainable	GSM8K	Math	HumanEval	MT-Bench
	LoRA	89.9M	40.75	5.22	17.68	3.73
	OLoRA	89.9M	42.93	6.51	21.12	4.03
Llama2-7B	PiSSA	89.9M	43.89	6.92	22.25	4.11
	Bone	87.0M	48.16	8.58	24.08	4.31
	LoRA	88.1M	38.13	6.06	-	-
RWKV 6-7B	PiSSA	88.1M	40.48	6.12	-	-
	Bone	88.1M	41.73	6.52	-	-
	LoRA	89.1M	65.17	15.82	39.02	-
Mistral-7B	PiSSA	89.1M	67.01	18.13	40.85	-
	Bone	88.1M	66.94	18.85	41.76	-

Table 2: We fine-tuned LLMs using Bone and various LoRA variants, and evaluated performance
 on GSM8k, Math, HumanEval, and MT-Bench.

Table 3: Comparing different values of rank (r)

Model	rank	Trainable	GSM8K	Math
	16	21.7M	45.90	3.77
Llama 2 7P	32	43.5M	46.18	7.43
Liailia2-7D	64	87.0M	48.16	8.58
	128	174.0M	53.49	10.08

4.4 BONE VS BAT

To validate the effectiveness of the "Weight Guide", we fine-tuned LLaMA2-7B and RWKV6-7B using both Bone and Bat on the MetaMathQA dataset and evaluated their performance on Math and GSM8K. As shown in the table, Bat, equipped with "Weight Guide", achieves significant improvements in performance metrics compared to Bone.

Model	Strategy	Trainable	GSM8K	Math
Llama2-7B	Bone	87.0M	48.16	8.58
	Bat	87.0M	<b>49.36</b>	<b>8.88</b>
RWKV6-7B	Bone	55.1M	41.73	6.52
	Bat	55.1M	<b>42.76</b>	<b>6.60</b>

Table 4: Comparing Bone, Bat on math tasks

4.5 ABLATION EXPERIMENTS ON DIFFERENT GROUPING METHODS FOR BAT

In this subsection, we explore the impact of different grouping methods in the Bat structure on model fine-tuning performance. Due to structural differences in the weight matrix W, the Bat-free grouping requires manual configuration, which is inconvenient. Therefore, this subsection only compares row-wise and column-wise grouping, both of which can be easily extended to any structure. We fine-tuned LLaMA2-7B on the MetaMathQA dataset and validated the results on GSM8k and Math. The results are shown in Table 5. Since in LLaMA2-7B, the dimension of gate\_proj in the MLP part is (4096, 11008), this leads to an asymmetry between row-wise and column-wise grouping in the Bat structure, making it difficult to align parameter counts. Although Bat-row uses 15M fewer parameters than Bat-col, it still delivers excellent performance. However, this discrepancy in



Figure 4: Training loss curves of Bone with different rank r on the MetaMathQA dataset.

parameter counts makes it challenging to accurately evaluate the differences between the groupingmethods.

To explore the differences between the two grouping methods and the effect of parameter count, we added a comparative experiment with RWKV6-3B, as RWKV6's symmetrical structure ensures that the trainable parameter count is the same whether using row-wise or column-wise grouping during Bat fine-tuning. This allows for a fairer comparison between Bat-row and Bat-col. The experimental results, shown in Table 5, indicate that the difference between the two is minimal, with both performing well. Therefore, we believe that Bat can effectively fit data regardless of the grouping method used. The key factor influencing Bat's performance remains the block size. However, this doesn't imply that different grouping methods are meaningless. As more LLMs begin to use techniques like GQA and MLA Liu et al. (2024a) to reduce KV cache overhead Dai et al. (2024); Lee et al. (2024); Shazeer (2019), the main weight matrices become smaller, and Bat will need to adjust its grouping or employ other techniques to adapt to these new technologies. 

Tuble 5. Comparing Dut 10%, Dut cor on much tuble	Table 5:	Comparing	Bat-row,	Bat-col	on math tasks	
---	----------	-----------	----------	---------	---------------	--

Model	Strategy	Trainable	GSM8K	Math
Llama2-7B	Bat-row	72.8M	45.76	7.82
	Bat-col	87.0M	<b>49.36</b>	<b>8.88</b>
RWKV6-3B	Bat-row	55.1M	25.93	3.12
	Bat-col	55.1M	25.25	3.09

#### 4.6 ABLATION EXPERIMENTS ON THE SPECIFIC IMPLEMENTATION OF "WEIGHT GUIDE"

Evaluating the best computation method for promoting internal weight feature fusion on RWKV6-3B: As seen in the table 6, Bone-Hadamard is significantly weaker compared to matrix multiplication.

The formula for Bone-Hadamard is as follows:

$$W_{k/b,d/b,b,b} = Reshape(W_{d,k})$$
  

$$bone_{d/b,b,b} = Reshape(bone_{b,d})$$
  

$$\Delta W_{d,k} = Reshape(W_{k/b,d/b,b,b} \odot bone_{d/b,b,b} + bone_{d/b,b,b})$$
(5)

#### 5 RESOURCE AND EFFICIENCY

Table 7 compares the training resources and token throughput required for fine-tuning RWKV6 using LoRA, Bone, and Bat on a single 4090 GPU. The specific fine-tuning settings are as follows: batch size = 1, context length (ctx\_len) = 512.

	Model	Strategy	Trainable	GSM8K	Math
	RWKV6-3B	Bat	55.1M	25.93	3.32
		Bat-Hadamard	55.1M	22.44	2.67

Table 6: Block-Affine on math tasks

The results show that Bone has the highest computational efficiency, being nearly 10% faster than LoRA while also being more memory-efficient. However, Bat incurs significantly higher memory usage due to large intermediate values and is slower in comparison. 

At the end of the table, we provide the actual resource costs for fine-tuning RWKV6 on the Meta-MathQA dataset using 4 NVIDIA 4090 GPUs, with checkpoint techniques applied.

Therefore, a key focus of our future work will be improving the Bat operator to enhance token throughput and reduce memory usage.

Model	Strategy	Trainable	GPU Memory	Token throughput
RWKV-3B	LoRA	55.1M	12074 MB	3.62 kt/s
	Bone	55.1M	<b>11052</b> MB	3.99 kt/s
	Bat	55.1M	22978 MB	2.16 kt/s
RWKV-3B	LoRA(use checkpoint)	55.1M	4*15328 MB	15.6 kt/s
	Bone(use checkpoint)	55.1M	4 * 15304 MB	<b>16.0</b> kt/s
	Bat(use checkpoint)	55.1M	4*15305 MB	14.2 kt/s

#### CONCLUSION

This work, inspired by GQA and MQA, leverages the sparsity of LLM weights to design the Block-Affine Adaptation (Bone) structure. In Bone, the original weights are divided into multiple subspaces, all of which share a single low-rank matrix initialized to zero for updates. Extensive exper-iments demonstrate that Bone consistently outperforms LoRA and its variants across various tasks, while also offering superior computational efficiency. To break the limitations of LoRA and ef-fectively utilize the information from the original weights, we propose the "Weight Guide" theory, which enables significant improvements through simple linear transformations. 

By integrating these innovations, we introduce a new structure called Block-Affine Transformation (Bat). Bat not only validates the effectiveness of "Weight Guide" but also addresses the limitation of Bone, where identical updates are applied to all subspaces. Experimental results show that Bat surpasses Bone on multiple tasks, although with a trade-off in computational efficiency. 

Bone brings new possibilities to existing LLM PEFT techniques. Instead of focusing solely on optimizing LoRA, we should shift our attention to innovative PEFT methods that are better suited to the architecture of LLMs.

#### FUTURE WORK

Bone and "Weight Guide" are merely a starting point, offering a foundation for researchers to ex-plore additional branches inspired by Bat. There are many directions for future work. 

1. Can Bone adapt to multimodal tasks and other complex scenarios? 

2. Can "Weight Guide" be implemented using alternative computation methods? 

486	We welcome the community to provide additional suggestions and conduct further tests
487	we welcome the community to provide additional suggestions and conduct further tests.
488	
489	
490	
491	
492	
493	
494	
495	
496	
497	
498	
499	
500	
501	
502	
503	
504	
505	
506	
507	
508	
509	
510	
511	
512	
513	
514	
515	
516	
517	
518	
519	
520	
521	
522	
523	
524	
525	
526	
527	
528	
529	
530	
531	
532	
533	
534	
535	
536	
537	
538	
539	

## 540 REFERENCES

574

575

576

579

586

592

- Joshua Ainslie, James Lee-Thorp, Michiel de Jong, Yury Zemlyanskiy, Federico Lebrón, and Sumit
   Sanghai. Gqa: Training generalized multi-query transformer models from multi-head check points. *arXiv preprint arXiv:2305.13245*, 2023.
- Dan Biderman, Jose Gonzalez Ortiz, Jacob Portes, Mansheej Paul, Philip Greengard, Connor Jennings, Daniel King, Sam Havens, Vitaliy Chiley, Jonathan Frankle, et al. Lora learns less and forgets less. *arXiv preprint arXiv:2405.09673*, 2024.
- Tom B Brown. Language models are few-shot learners. *arXiv preprint ArXiv:2005.14165*, 2020.
- Kerim Büyükakyüz. Olora: Orthonormal low-rank adaptation of large language models. *arXiv* preprint arXiv:2406.01775, 2024.
- 553 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared 554 Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, 555 Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, 556 Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex 558 Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, 559 Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob Mc-561 Grew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large 562 language models trained on code, 2021. 563
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Jincheng Dai, Zhuowei Huang, Haiyun Jiang, Chen Chen, Deng Cai, Wei Bi, and Shuming Shi.
   Sequence can secretly tell you what to discard. *arXiv preprint arXiv:2404.15949*, 2024.
- Ning Ding, Yujia Qin, Guang Yang, Fuchao Wei, Zonghan Yang, Yusheng Su, Shengding Hu, Yulin
  Chen, Chi-Min Chan, Weize Chen, et al. Parameter-efficient fine-tuning of large-scale pre-trained
  language models. *Nature Machine Intelligence*, 5(3):220–235, 2023.
  - Soufiane Hayou, Nikhil Ghosh, and Bin Yu. Lora+: Efficient low rank adaptation of large models. arXiv preprint arXiv:2402.12354, 2024.
- Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Towards a
   unified view of parameter-efficient transfer learning, 2022.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *International conference on machine learning*, pp. 2790–2799. PMLR, 2019.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- Wonbeom Lee, Jungi Lee, Junghwan Seo, and Jaewoong Sim. {InfiniGen}: Efficient generative inference of large language models with dynamic {KV} cache management. In *18th USENIX Symposium on Operating Systems Design and Implementation (OSDI 24)*, pp. 155–172, 2024.
- Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning, 2021.
- 593 Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv* preprint arXiv:2101.00190, 2021.

594 Aixin Liu, Bei Feng, Bin Wang, Bingxuan Wang, Bo Liu, Chenggang Zhao, Chengqi Dengr, Chong 595 Ruan, Damai Dai, Daya Guo, et al. Deepseek-v2: A strong, economical, and efficient mixture-596 of-experts language model. arXiv preprint arXiv:2405.04434, 2024a. 597 Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-598 Ting Cheng, and Min-Hung Chen. Dora: Weight-decomposed low-rank adaptation. arXiv preprint arXiv:2402.09353, 2024b. 600 601 Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. Gpt understands, too. AI Open, 2023. 602 603 Fanxu Meng, Zhaohui Wang, and Muhan Zhang. Pissa: Principal singular values and singular 604 vectors adaptation of large language models. arXiv preprint arXiv:2404.02948, 2024. 605 Bo Peng, Daniel Goldstein, Quentin Anthony, Alon Albalak, Eric Alcaide, Stella Biderman, Eugene 606 Cheah, Teddy Ferdinan, Haowen Hou, Przemysław Kazienko, et al. Eagle and finch: Rwkv with 607 matrix-valued states and dynamic recurrence. arXiv preprint arXiv:2404.05892, 2024. 608 609 Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. Adapter-610 fusion: Non-destructive task composition for transfer learning. arXiv preprint arXiv:2005.00247, 611 2020. 612 Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language 613 models are unsupervised multitask learners. OpenAI blog, 1(8):9, 2019. 614 615 Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi 616 Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text 617 transformer. Journal of machine learning research, 21(140):1–67, 2020. 618 Anastasia Razdaibiedina, Yuning Mao, Rui Hou, Madian Khabsa, Mike Lewis, Jimmy Ba, and 619 Amjad Almahairi. Residual prompt tuning: Improving prompt tuning with residual reparameter-620 ization, 2023. 621 Noam Shazeer. Fast transformer decoding: One write-head is all you need. arXiv preprint 622 arXiv:1911.02150, 2019. 623 624 Chongjie Si, Xuehui Wang, Xue Yang, Zhengqin Xu, Qingyun Li, Jifeng Dai, Yu Qiao, Xi-625 aokang Yang, and Wei Shen. Flora: Low-rank core space for n-dimension. arXiv preprint 626 arXiv:2405.14739, 2024. 627 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Niko-628 lay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open founda-629 tion and fine-tuned chat models. arXiv preprint arXiv:2307.09288, 2023. 630 Shaowen Wang, Linxi Yu, and Jian Li. Lora-ga: Low-rank adaptation with gradient approximation. 631 arXiv preprint arXiv:2407.05000, 2024. 632 633 Shibo Wang and Pankaj Kanwar. Bfloat16: The secret to high performance on cloud tpus. Google 634 Cloud Blog, 4, 2019. 635 Yaqing Wang, Sahaj Agarwal, Subhabrata Mukherjee, Xiaodong Liu, Jing Gao, Ahmed Hassan 636 Awadallah, and Jianfeng Gao. Adamix: Mixture-of-adaptations for parameter-efficient model 637 tuning. arXiv preprint arXiv:2205.12410, 2022. 638 639 Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, Qingwei 640 Lin, and Daxin Jiang. Wizardlm: Empowering large pre-trained language models to follow com-641 plex instructions. In The Twelfth International Conference on Learning Representations, 2024. 642 Lingling Xu, Haoran Xie, Si-Zhao Joe Qin, Xiaohui Tao, and Fu Lee Wang. Parameter-efficient 643 fine-tuning methods for pretrained language models: A critical review and assessment. arXiv 644 preprint arXiv:2312.12148, 2023. 645 Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T Kwok, Zhen-646 guo Li, Adrian Weller, and Weiyang Liu. Metamath: Bootstrap your own mathematical questions 647 for large language models. arXiv preprint arXiv:2309.12284, 2023.

648 649 650	Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. <i>Advances in Neural Information Processing Systems</i> , 36, 2024a.
652	Tianyu Zheng, Ge Zhang, Tianhao Shen, Xueling Liu, Bill Yuchen Lin, Jie Fu, Wenhu Chen, and Xiang Yue. Opencode interpreter: Integrating code generation with execution and refinement
653	arXiv preprint arXiv:2402 14658 2024b
654	unity proprim unity.2702.17050, 20270.
655	
656	
657	
658	
659	
660	
661	
662	
663	
664	
665	
667	
669	
660	
670	
671	
672	
673	
674	
675	
676	
677	
678	
679	
680	
681	
682	
683	
684	
685	
686	
687	
688	
689	
690	
691	
692	
693	
694	
695	
696	
697	
698	
699	
700	
701	