

Calibrating and Rotating: A Unified Framework for Weight Conditioning in PEFT

Da Chang^{1,2,4}, Peng Xue^{1,2,4}, Yu Li³, Yongxiang Liu^{1*}, Pengxiang Xu¹, Shixun Zhang¹

¹Pengcheng Laboratory

²Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences

³George Washington University

⁴University of Chinese Academy of Sciences

{changda24,xuepeng24}@mails.uca.ac.cn, yul@gwu.edu, {liuyx,xupx,zhangshx}@pcl.ac.cn

Abstract

Parameter-Efficient Fine-Tuning (PEFT) methods are crucial for adapting large pre-trained models. Among these, LoRA is considered a foundational approach. Building on this, the influential DoRA method enhances performance by decomposing weight updates into magnitude and direction. However, its underlying mechanism remains unclear, and it introduces significant computational overhead. In this work, we first identify that DoRA’s success stems from its capacity to increase the singular value entropy of the weight update matrix, which promotes a more uniform update distribution akin to full fine-tuning. We then reformulate DoRA into a mathematically equivalent and more efficient matrix form, revealing it as a learnable weight conditioning method. Based on this insight, we propose a unified framework for designing advanced PEFT methods by exploring two orthogonal dimensions: the architectural placement and the transformation type of the conditioning matrix. Within this framework, we introduce two novel methods: (1) **Pre-Diag**, which applies a diagonal conditioning matrix before the LoRA update to efficiently calibrate the pre-trained weights, thereby enhancing performance while reducing training time; and (2) **Skewed Orthogonal Rotation Adaptation (SORA)**, which employs a parameter-efficient orthogonal rotation to perform a more powerful, norm-preserving transformation of the feature space. Extensive experiments on natural language understanding and generation tasks demonstrate that our proposed methods achieve superior performance and efficiency compared to both LoRA and DoRA.

Code — <https://github.com/MaeChd/SORA>

Full paper — <https://arxiv.org/abs/2511.00051>

Introduction

Large-scale Pre-trained Models constitute the bedrock of modern artificial intelligence, yet their substantial fine-tuning costs drive research into Parameter-Efficient Fine-Tuning (PEFT) (Ding et al. 2023). PEFT techniques adapt these models with minimal computational and storage overhead through diverse strategies, ranging from inserting small, trainable adapter modules into the model’s architecture (Houlsby et al. 2019) to methods like prompt-tuning, which only modify the input embeddings (Lester, Al-Rfou, and Constant 2021; Li

and Liang 2021). Among these, LoRA (Hu et al. 2022) is one of the most widely adopted PEFT techniques, which operates by adding a low-rank update \mathbf{BA} to the pre-trained weights \mathbf{W}_{pre} (i.e., $\mathbf{W} = \mathbf{W}_{\text{pre}} + s \cdot \mathbf{BA}$).

The simplicity and effectiveness of LoRA inspire a continuous stream of research, leading to numerous variants that improve upon the original method along several key axes. These advancements primarily focus on: (i) dynamic rank allocation during training, as exemplified by AdaLoRA (Zhang et al. 2023) and ReLoRA (Lialin et al. 2023); (ii) improved initialization of low-rank matrices from pre-trained weights, proposed in methods like PiSSA (Meng, Wang, and Zhang 2024) and LoRA-GA (Wang, Yu, and Li 2024); and (iii) refined optimization strategies, such as those introduced in LoRA+ (Hayou, Ghosh, and Yu 2024) and LoRA-RITE (Yen et al. 2024).

Our work connects to the thread of principled weight preparation, with a focus on DoRA (Weight-Decomposed Low-Rank Adaptation) (Liu et al. 2024). As an improvement upon LoRA, DoRA pushes this philosophy further by decoupling a pre-trained weight matrix \mathbf{W} into its magnitude and direction. It then applies the low-rank update exclusively to the directional component, leaving the magnitude fixed. This process is formulated as $\mathbf{W} = \mathbf{m} \frac{\mathbf{W}_{\text{pre}} + s\mathbf{BA}}{\|\mathbf{W}_{\text{pre}} + s\mathbf{BA}\|_c}$.

DoRA exhibits superior performance to LoRA on numerous tasks, demonstrating the potential of this decomposition. However, this success is accompanied by a significant drawback: the column-wise norm computation ($\|\cdot\|_c$) required during training introduces substantial computational overhead. This overhead hinders its practical application, especially for larger models and longer sequences. Concurrently, a more fundamental question remains unanswered:

What is the true driving force behind DoRA’s performance improvement?

To answer this question, we conduct an in-depth empirical analysis. We find that while prior work points to stable rank (Lion et al. 2025), its explanatory power is limited in many practical tasks. Instead, we find that DoRA consistently increases the **Singular Value Entropy** of the weight update matrix. A higher entropy signifies a broader distribution of update energy across features, which prevents over-

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

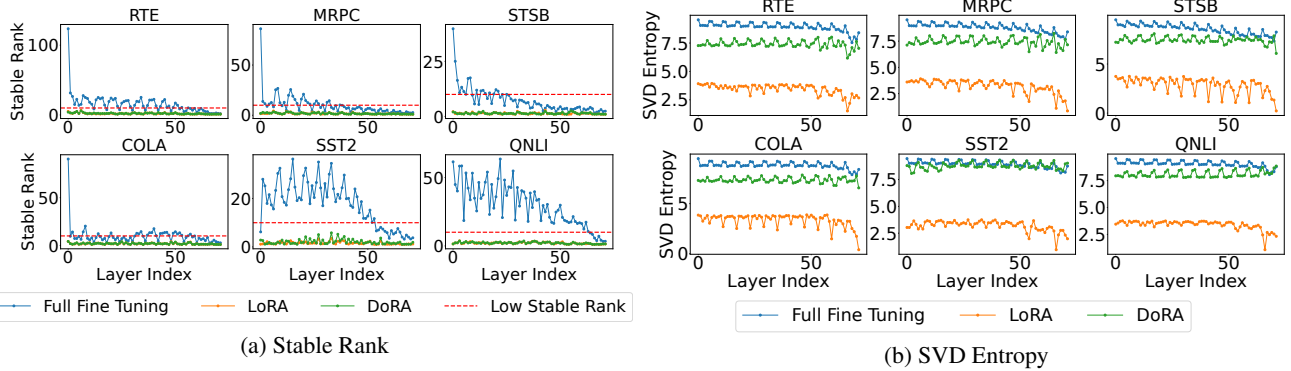


Figure 1: (a) DeBERTaV3-Base on GLUE Benchmark: Stable Rank $\|\Delta\mathbf{W}\|_F^2/\|\Delta\mathbf{W}\|_2^2$ Across Layers; (b) DeBERTaV3-Base on GLUE Benchmark: SVD Entropy $H(\sigma) = -\sum_i p_i \log p_i$ Across Layers, with Comparisons of Full Fine-Tuning, LoRA, and DoRA.

specialization and emulates the desirable behavior of full fine-tuning (Liu et al. 2025). This suggests DoRA’s success stems from indirectly optimizing the weight update’s singular value spectrum. To validate this and address DoRA’s inefficiency, we reframe it as a more efficient **Weight Conditioning** method (Saratchandran, Wang, and Lucey 2024). This perspective reveals DoRA’s mechanism as modulating the main weight matrix with a lightweight, learnable matrix. This discovery opens a broader design space, for which we propose a unified weight conditioning framework to create more powerful and efficient PEFT methods.

Within this framework, we introduce two novel methods: (i) **Pre-Diag** which efficiently calibrates pre-trained weights to improve performance while reducing training time, and (ii) **Skewed Orthogonal Rotation Adaptation (SORA)**, a powerful, norm-preserving transformation that uses parameter-efficient orthogonal rotations to align the feature space, building on Orthogonal Fine-Tuning (OFT) (Qiu et al. 2023).

Our contributions can be summarized as follows:

1. We are the first to connect the success of DoRA to an increase in singular value entropy and, based on this, reframe it as an efficient weight conditioning method.
2. We propose a unified weight conditioning framework, from which we derive two novel methods: Pre-Diag for efficient weight calibration and SORA for powerful rotational adaptation of the feature space.
3. We demonstrate through extensive experiments that our proposed methods outperform existing state-of-the-art PEFT methods in terms of both performance and efficiency.

Analysis and Insights

In this section, we first demonstrates the inadequacy of stable rank for explaining performance differences among fine-tuning methods. We then introduce singular value entropy as a more discerning metric, providing both empirical and theoretical evidence that DoRA attains a higher entropy than LoRA. Ultimately, we uncover DoRA’s core mechanism by

reformulating it as a form of weight conditioning, which enables it to modify the full singular value spectrum and thus enhance update uniformity.

The Inadequacy of Stable Rank

To understand the performance discrepancies among different fine-tuning methods, we first analyze the stable rank of the parameter update matrix, $\Delta\mathbf{W} = \mathbf{W} - \mathbf{W}_{\text{pre}}$, where \mathbf{W} represents the final weights and \mathbf{W}_{pre} the initial pre-trained weights. The stable rank is defined as $\|\Delta\mathbf{W}\|_F^2/\|\Delta\mathbf{W}\|_2^2$ (Rudelson and Vershynin 2007). Our choice of this metric is motivated by prior work that has linked it to the diversity of update directions (Lion et al. 2025) and the expressive capacity of low-rank models (Zeng and Lee 2023; Kalajdziewski 2023). Our initial hypothesis was that the stable rank could quantitatively explain why methods like Full Fine-Tuning (FFT) and DoRA outperform LoRA.

However, our layer-wise analysis reveals that the aggregated, model-wide stable rank is an insufficient and potentially misleading metric. As shown in Figure 1a, the stable rank of FFT is highly volatile across layers, soaring in some while dropping to levels comparable to the consistently low ranks of LoRA and DoRA in others. This volatility indicates that FFT’s high global rank is driven by only a subset of layers. Consequently, stable rank fails to robustly distinguish the learning dynamics of these methods or explain their performance differences.

SVD Entropy as a More Faithful Explainer

The limitations of stable rank motivate our search for a more discerning metric. We therefore employ singular value entropy, defined as the Shannon entropy of the normalized singular value spectrum: $H(\sigma) = -\sum_i p_i \log p_i$, where $p_i = \sigma_i^2/\sum_j \sigma_j^2$ (Roy and Vetterli 2007; Alter, Brown, and Botstein 2000). Intuitively, this metric measures how uniformly the update’s energy is distributed across its singular directions. A higher entropy indicates that the update is more uniformly distributed, rather than being dominated by a few

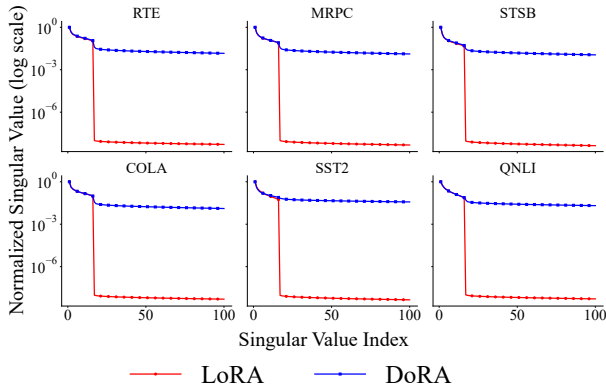


Figure 2: Distribution of average singular values after layer-wise normalization for LoRA and DoRA on GLUE tasks with DeBERTaV3-Base.

high-magnitude singular values. This uniformity prevents the update from concentrating on only a few dominant features.

In stark contrast to the ambiguous results from stable rank, Figure 1b shows that layer-wise singular value entropy reveals a clear and consistent hierarchy: FFT consistently maintains the highest entropy, followed closely by DoRA, with LoRA lagging significantly behind.

To provide theoretical support for this empirical finding, we analyze a simplified model inspired by the step-like structure of the singular value distributions observed in Figure 2. We model the normalized singular values of LoRA and DoRA as two-step and three-step distributions, respectively. This leads to the following proposition:

Theorem 1. Let $\Delta \mathbf{W}_{LoRA}$ be a LoRA update matrix of rank r . Its set of normalized singular values is $\Sigma_{LoRA} = \{1, \underbrace{\alpha, \dots, \alpha}_{r-1}\}$. Let $\Delta \mathbf{W}_{DoRA}$ be a DoRA update matrix of rank $s > r$. Its set of normalized singular values is $\Sigma_{DoRA} = \{1, \underbrace{\beta, \dots, \beta}_{r-1}, \underbrace{\gamma, \dots, \gamma}_{s-r}\}$. The singular value entropy

$H(\sigma(\Delta \mathbf{W}))$ is the Shannon entropy of the probability distribution $\{p_i\}$ formed by the squared normalized singular values, where $p_i = \sigma_i^2 / \sum_j \sigma_j^2$. Given the constraint that the sum of squares of minor singular values is preserved, $(r-1)\alpha^2 = (r-1)\beta^2 + (s-r)\gamma^2$, and assuming $\alpha > \gamma > 0$ and $\beta > 0$. Then, the singular value entropy of the DoRA matrix is strictly greater than that of the LoRA matrix:

$$H(\sigma(\Delta \mathbf{W}_{DoRA})) > H(\sigma(\Delta \mathbf{W}_{LoRA}))$$

A detailed proof is provided in the Appendix 1.

Remark 1. This proposition formally demonstrates that by activating more singular directions and distributing the update energy across them (even if the newly activated values γ are small), the overall entropy of the update matrix increases. This reinforces our central finding: the ability to increase singular value entropy appears to be the primary driving force behind DoRA’s performance gains over LoRA.

Unveiling DoRA: Weight Conditioning

This finding raises a crucial question: **how does the DoRA formulation mechanically increase singular value entropy?** The original DoRA update rule (Liu et al. 2024) decomposes the weight matrix \mathbf{W} into a learnable magnitude \mathbf{m} and a direction component:

$$\mathbf{W} = \mathbf{m} \frac{\mathbf{W}_{pre} + s \cdot \mathbf{BA}}{\|\mathbf{W}_{pre} + s \cdot \mathbf{BA}\|_c},$$

where $\|\cdot\|_c$ denotes the column norm.

We can algebraically reformulate this operation as a form of weight conditioning (Saratchandran, Wang, and Lucey 2024). This reformulation reveals the effective weight change $\Delta \mathbf{W}_{DoRA} = \mathbf{W} - \mathbf{W}_{pre}$ to be:

$$\Delta \mathbf{W}_{DoRA} = \mathbf{W}_{pre}(\mathbf{D} - \mathbf{I}) + s \cdot \mathbf{BAD},$$

where $\mathbf{D} = \text{Diag} \left(\frac{\mathbf{m}}{\|\mathbf{W}_{pre} + s \cdot \mathbf{BA}\|_c} \right)$.

This reformulation offers two critical insights. Computationally, it replaces the costly column-wise norm with an efficient matrix multiplication. More importantly, it reveals DoRA’s core mechanism: unlike the strictly low-rank LoRA update, the DoRA update $\Delta \mathbf{W}_{DoRA}$ incorporates a higher-rank term, $\mathbf{W}_{pre}(\mathbf{D} - \mathbf{I})$.

This higher-rank component is key. For singular value indices $i > r$, the LoRA update is zero (i.e., $\sigma_i(\Delta \mathbf{W}_{LoRA}) = 0$). In contrast, Weyl’s theorem (Horn and Johnson 2012) shows that DoRA maintains non-zero singular values in these dimensions:

$$\sigma_i(\Delta \mathbf{W}_{DoRA}) \leq \sigma_1(\mathbf{W}_{pre}(\mathbf{D} - \mathbf{I})) + \sigma_i(\mathbf{BAD}) \quad \text{for } i > r.$$

This ability to modify the full singular value spectrum allows DoRA to create a richer, more uniform update distribution, thereby increasing its entropy, as illustrated in Figure 2.

Framework and Methods

Our analysis reveals that DoRA’s success stems from its capacity to increase the singular value entropy of the update matrix. This effect is achieved through an implicit form of weight conditioning, where a diagonal matrix \mathbf{D} modulates the weights. Based on this insight, we propose a Unified Framework for Weight Conditioning to move beyond ad-hoc improvements and systematically design novel PEFT methods.

The framework’s core idea is to enhance the standard LoRA update, $\Delta \mathbf{W} = s\mathbf{BA}$, with a learnable conditioning matrix. We identify two fundamental and orthogonal design axes:

1. **Architectural Placement:** Where should the conditioning matrix be applied? Should it act on the entire weight after the LoRA update as in DoRA, or only on the pre-trained weights?
2. **Transformation Type :** What transformation should the conditioning matrix perform? Should it be a diagonal matrix \mathbf{D} that only performs axis-aligned scaling as implicitly done by DoRA, or a more powerful orthogonal matrix \mathbf{P} capable of performing rotation?

By systematically exploring these two axes, we develop principled and efficient methods that explicitly optimize the update spectrum.

Method	Formulation	Condition	Placement	Key idea
LoRA (Hu et al. 2022)	$\mathbf{W}_{\text{pre}} + \mathbf{BA}$	Identity Matrix	N/A	Low-Rank Adaption
DoRA (Liu et al. 2024)	$(\mathbf{W}_{\text{pre}} + \mathbf{BA})\mathbf{D}$	Diagonal Matrix	Post	Weight-Decomposed
Pre-Diag (ours)	$\mathbf{W}_{\text{pre}}\mathbf{D} + \mathbf{BA}$	Diagonal Matrix	Pre	Calibration
SORA (ours)	$(\mathbf{W}_{\text{pre}}\mathbf{D} + \mathbf{BA})\mathbf{P}$	Orthogonal Matrix	Pre & Post	Calibration & Rotation

Table 1: Comparison of properties between Fully Fine-Tuning, LoRA, DoRA and our’s Pre-Diag, SORA.

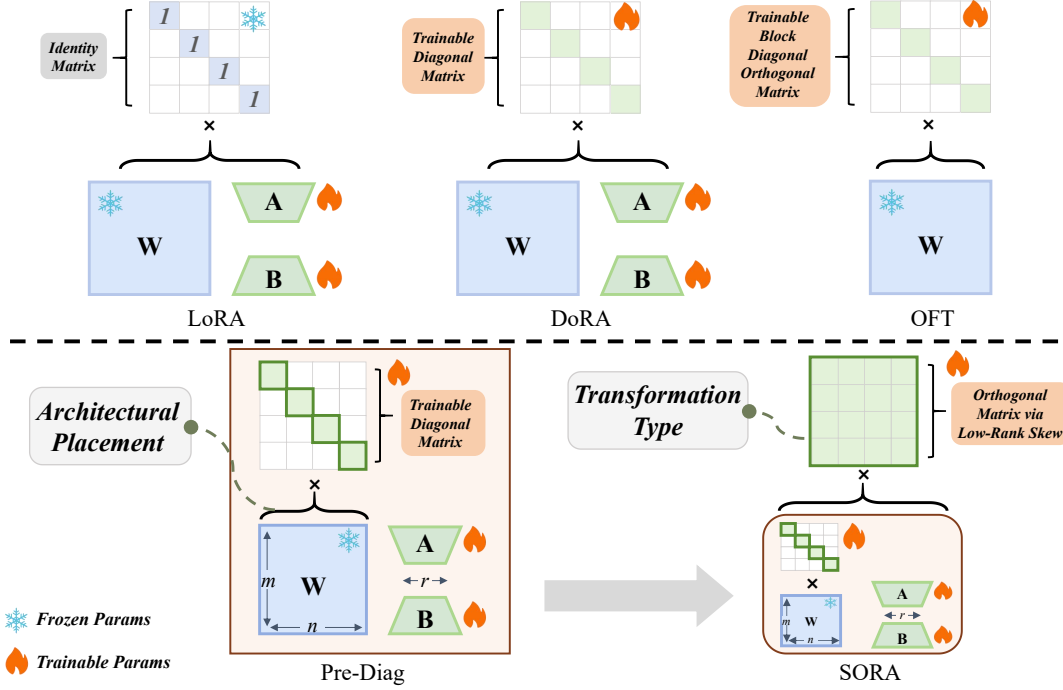


Figure 3: Taking LoRA as the baseline, we can regard DoRA and OFT as weight-conditioned methods built upon it. Our method Pre-Diag, in turn, adjusts pre-trained weights by using a trainable diagonal matrix based on DoRA; SORA, on the other hand, replaces the conditional matrix in DoRA with a trainable low-parameter orthogonal matrix to automatically rotate input features.

Placement: Pre-Diag for Efficient Calibration

We first explore the architectural placement of the conditioning matrix. Recall from previous section that DoRA’s update can be expressed as $\Delta\mathbf{W}_{\text{DoRA}} = \mathbf{W}_{\text{pre}}(\mathbf{D} - \mathbf{I}) + s \cdot \mathbf{BA}$. In this formulation, the conditioning matrix \mathbf{D} is coupled with both the pre-trained weights \mathbf{W}_{pre} and the LoRA update \mathbf{BA} , creating a complex gradient path.

To simplify this, we propose Pre-Diag, which applies the conditioning directly and solely to the pre-trained weights. Its formulation is:

$$\mathbf{W} = \mathbf{W}_{\text{pre}}\mathbf{D} + s \cdot \mathbf{BA}.$$

The resulting update matrix is $\Delta\mathbf{W}_{\text{Pre-Diag}} = \mathbf{W}_{\text{pre}}(\mathbf{D} - \mathbf{I}) + s \cdot \mathbf{BA}$.

The motivation for this design is to decouple two distinct goals: (1) using the diagonal matrix \mathbf{D} to efficiently calibrate the scale of the pre-trained weights \mathbf{W}_{pre} , and (2) using the LoRA update \mathbf{BA} to learn new, task-specific features. This cleaner, more direct structure simplifies the gradient computation and, as we show in experiments, leads to faster training.

Pre-Diag thus serves as a more direct and efficient method for spectrum optimization, achieving high singular value entropy with reduced computational overhead.

Transformation: SORA for Rotational Adaptation

We further enhance the Pre-Diag, $\mathbf{W} = \mathbf{W}_{\text{pre}}\mathbf{D} + s \cdot \mathbf{BA}$, by introducing a learnable transformation. This model implicitly ends with an identity matrix: $(\mathbf{W}_{\text{pre}}\mathbf{D} + s \cdot \mathbf{BA})\mathbf{I}$. We propose replacing this static identity with a dynamic, parameter-efficient rotation matrix \mathbf{P} .

This leads to our method, **SORA** (Skewed Orthogonal Rotation Adaptation), which applies this rotation to refine feature interactions:

$$\mathbf{W}' = (\mathbf{W}_{\text{pre}}\mathbf{D} + s \cdot \mathbf{BA})\mathbf{P}.$$

To maintain efficiency, the orthogonal matrix \mathbf{P} is constructed as the exponential of a low-rank skew-symmetric matrix $\mathbf{S}_{\mathbf{P}}$, where $\mathbf{S}_{\mathbf{P}}$ is parameterized by two small matrices $\mathbf{D}_{\mathbf{P}}, \mathbf{C}_{\mathbf{P}} \in \mathbb{R}^{n \times r_{\mathbf{P}}}$ ($r_{\mathbf{P}} \ll n$).

As the direct computation of the matrix exponential, $\mathbf{P} = \exp(s_{\mathbf{P}} \cdot \mathbf{S}_{\mathbf{P}})$, is prohibitively expensive, we employ a highly

accurate first-order Taylor approximation for small rotations:

$$\mathbf{P} \approx \mathbf{I} + s_P \cdot \mathbf{S}_P = \mathbf{I} + s_P \cdot (\mathbf{D}_P \mathbf{C}_P^\top - \mathbf{C}_P \mathbf{D}_P^\top).$$

This approximation makes \mathbf{P} near-orthogonal while drastically reducing computational demands. (Singla and Feizi 2021) first introduced the skew-orthogonal strategy into CNNs and provided a bounding for the approximation of matrix exponentials. Similarly, in SORA, for the orthogonal matrix \mathbf{P} , we have the following theorem:

Theorem 2. For Skew-Symmetric \mathbf{J} , we have the inequality:

$$\|\exp(\mathbf{J}) - T_k(\mathbf{J})\|_2 \leq \frac{\|\mathbf{J}\|_2^k}{k!} \quad \text{where} \quad T_k(\mathbf{J}) = \sum_{i=0}^{k-1} \frac{\mathbf{J}^i}{i!}.$$

Then, Let $s_P = \frac{\epsilon}{2\|\mathbf{D}_P\|_F\|\mathbf{C}_P\|_F + \epsilon}$, we have the inequality:

$$\|\exp(\mathbf{P}) - T_1(\mathbf{P})\|_2 \leq \epsilon.$$

A detailed proof is provided in the Appendix 2.

Remark 2. This implies that when $\|\mathbf{P}\|_F$ is extremely small, $\exp(\mathbf{P})$ can be approximated with high precision, and this can be achieved by controlling s_P , thereby significantly reducing the time complexity involved in solving $\exp(\mathbf{P})$.

To efficiently compute $\mathbf{W}_{PD}\mathbf{P}$ where $\mathbf{W}_{PD} = \mathbf{W}_{pre}\mathbf{D} + s \cdot \mathbf{B}\mathbf{A}$, we optimize the matrix multiplication. A naive calculation of $\mathbf{W}_{PD}(\mathbf{D}_P \mathbf{C}_P^\top - \mathbf{C}_P \mathbf{D}_P^\top)$ has $\mathcal{O}(mn^2)$ complexity. By reordering the operations as $(\mathbf{W}_{PD}\mathbf{D}_P)\mathbf{C}_P^\top - (\mathbf{W}_{PD}\mathbf{C}_P)\mathbf{D}_P^\top$, we reduce the complexity to $\mathcal{O}(mnr_P)$, ensuring SORA is both powerful and computationally efficient.

Figure 3 and Table 1 detail the differences between our proposed **Pre-Diag** and **SORA** and existing methods.

Experiment

In this section, we evaluate the performance of our proposed methods, Pre-Diag and SORA.

Setup All experiments were conducted on a setup consisting of four NVIDIA RTX 4090 (24GB) GPUs and eight Ascend 910C (64GB) NPUs. Without additional specifications, we use $r_P = 1$ by default in our experiments. Detailed experimental settings are provided in Appendix 3.

Natural language understanding To test general language understanding, we fine-tuned DeBERTaV3-Base (He et al. 2020) on the GLUE benchmark (Wang et al. 2018). We test three learning rates ($\{1e-3, 8e-4, 4e-4\}$), each with three random seeds for robustness. Table 2 shows aggregated results comparing our Pre-Diag and SORA with LoRA and DoRA, reporting the best average performance. The findings indicate that our methods achieve superior average performance on the benchmark.

Commonsense reasoning For commonsense reasoning, we conduct several experiments. First, we train LLaMA3-8B (Dubey et al. 2024) on the Commonsense15k dataset and test its performance on a commonsense benchmark. The experimental results are presented in Table 3. SORA achieves the optimal average score across all tasks with different ranks, which effectively confirms the effectiveness of SORA.

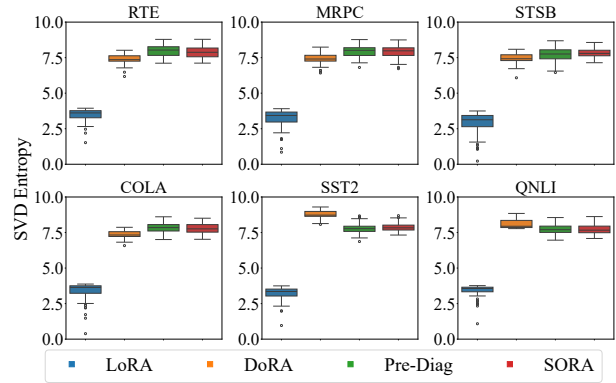


Figure 4: SVD Entropy of DeBERTaV3-Base Fine-Tuned on GLUE tasks: Comparing LoRA, DoRA, Our Pre-Diag, and SORA.




Figure 5: Comparison of inference speed (steps per second) and training speed (steps per second) among LoRA, DoRA, and our proposed method on the PIQA, ARC-c, and OBQA tasks, using the LLaMA3-8B model.

Mathematical reasoning For mathematical reasoning, we conduct several experiments. First, we train Gemma-7B (Gemma Team et al. 2024) on the MetaMathQA14k dataset (Yu et al. 2024) and test its performance on a mathematical benchmark. The experimental results are presented in Table 4. SORA achieves the optimal average score across all tasks, which effectively confirms the effectiveness of SORA.

Update Diversity and Efficiency

We evaluate our methods on two fronts: update diversity and computational efficiency. (i) Figure 4 shows the SVD entropy of weight updates on GLUE tasks. Both Pre-Diag and SORA consistently achieve higher entropy than LoRA, and surpass DoRA on most tasks, indicating a more uniform update distribution. (ii) We benchmark training and inference speeds on LLaMA3-8B in Figure 5. While LoRA is the fastest, DoRA introduces substantial overhead, our methods offer a compelling trade-off. Compared to DoRA, Pre-Diag and SORA significantly accelerate both training (by 50.60% and 36.96%) and inference (by 65.09% and 51.27%). These results demonstrate that our proposed methods successfully balance enhanced update diversity with high computational efficiency.

Method	Trainable Params Rate	RTE 2.5k	MRPC 3.7k	STS-B 7k	CoLA 8.5k	SST-2 67k	QNLI 105k	QQP 364k	Avg.
LoRA	1.42%	86.63 \pm 0.63	91.25 \pm 0.79	90.44 \pm 0.36	69.71 \pm 0.51	96.21 \pm 0.00	93.94 \pm 0.10	91.80 \pm 0.60	88.57
DoRA	1.46%	86.83 \pm 1.02	92.15 \pm 0.73	90.34 \pm 0.41	69.90 \pm 0.45	96.25 \pm 0.07	94.06 \pm 0.27	92.38 \pm 0.02	88.84
Pre-Diag	1.46%	87.24\pm1.67	91.17\pm0.25	90.50\pm0.46	71.19\pm0.49	96.10\pm0.20	94.04\pm0.14	92.21\pm0.03	88.98
SORA	1.55%	87.90\pm0.42	91.82\pm1.02	90.79\pm0.57	70.80\pm1.18	96.36\pm0.12	94.30\pm0.23	92.38\pm0.06	89.19

Table 2: Comparison of best average results of fine-tuning DeBERTaV3-base model on GLUE benchmark. We report the Matthew’s (Pearson) correlation for CoLA (STS-B), and accuracy for other tasks

Rank	Method	BoolQ	PIQA	SIQA	HeSw	WiGr	ARC-e	ARC-c	OBQA	Avg.
4	LoRA	68.96	75.46	73.18	79.40	71.19	84.09	69.54	78.20	75.00
	DoRA	69.72	75.35	72.31	78.36	73.88	83.92	69.45	78.20	75.15
	Pre-Diag	69.02	77.09	73.08	79.27	71.59	80.98	68.94	74.20	74.27
	SORA	69.63	76.28	72.62	78.92	71.59	84.43	70.90	77.80	75.27
8	LoRA	71.13	74.76	72.52	81.49	74.19	84.22	70.31	77.60	75.78
	DoRA	70.86	75.90	72.72	80.85	74.90	84.89	70.73	78.80	76.08
	Pre-Diag	70.06	77.09	73.08	81.68	72.61	84.68	70.56	79.40	76.15
	SORA	70.76	76.66	73.03	81.49	73.40	84.89	70.73	78.80	76.22
16	LoRA	70.09	76.01	73.59	81.08	73.88	84.60	70.90	78.80	76.12
	DoRA	71.77	76.33	72.76	82.27	74.19	83.84	68.77	79.60	76.19
	Pre-Diag	70.24	76.71	73.44	81.26	74.82	84.64	71.16	79.40	76.46
	SORA	71.56	76.71	73.69	82.82	76.40	85.43	70.14	79.00	76.97
32	LoRA	71.87	77.97	73.85	83.46	76.40	85.98	71.33	80.20	77.63
	DoRA	71.35	78.56	73.29	83.70	75.92	86.03	71.67	80.40	77.62
	Pre-Diag	71.83	78.07	73.03	83.33	75.37	85.48	72.44	81.20	77.59
	SORA	72.32	78.67	73.95	83.77	76.87	86.53	71.93	80.40	78.06

Table 3: Single-task commonsense reasoning performance on LLaMA3-8B. Models are trained on commonsense15k, comparing LoRA, DoRA, and our method across a range of ranks.

Method	Trainable Params Rate	GSM8K	MultiArith	AQuA	SVAMP	AddSub	SingleEq	Avg.
OFT	0.58%	73.09	99.00	38.19	75.00	86.10	94.69	77.68
LoRA+	0.40%	74.83	<u>98.67</u>	36.58	75.10	85.32	93.11	77.27
LoRA	0.40%	74.07	97.50	37.80	75.90	85.08	95.85	77.70
DoRA	0.41%	73.77	97.17	<u>38.34</u>	77.8	88.40	93.92	78.23
SORA	0.47%	74.98	98.50	40.49	<u>76.30</u>	<u>87.63</u>	95.87	78.96

Table 4: Single-task mathematical reasoning performance on Gemma-7B. Models are trained on MetaMathQA14k and evaluated on various benchmarks. All methods use a rank of 16 ($r_p = 2$ for SORA).

Ablation Study

Component Analysis To validate our architectural choices, we ablate SORA along two orthogonal axes: whether the weight conditioning is placed before or after the LoRA update, and whether its transformation is diagonal or orthogonal. We benchmark four models: the LoRA baseline, the post-diagonal model DoRA, our pre-diagonal model Pre-Diag, and our final model SORA, which combines a pre-diagonal with a post-orthogonal transform. This comparison isolates the performance contribution of each component, empirically justifying the benefits of pre-conditioning and introducing an orthogonal transformation. From Table 5, we find that Pre-Ortho leads to worse performance with adjusted LoRA

weights BA, indicating rotation transformation is more suitable for global use.

Hyperparameter Sensitivity On the RTE task, we sweep $r_P \in \{1, 2, 4, 8, 16\}$ and observe that increasing r_P beyond 1 yields at most marginal accuracy gains while monotonically increasing inference cost and memory footprint. In particular, $r_P = 1$ achieves the most favorable accuracy–efficiency trade-off, matching or slightly outperforming higher-rank settings under comparable budgets. The summary in Table 6 corroborates this trend, indicating that the minimal rank is already sufficient for this task and that larger ranks exhibit diminishing returns.

Model	Placement		Type		Avg. GLUE Score
	Pre	Post	Diag (D)	Ortho (P)	
<i>Baseline Experiments</i>					
LoRA					88.57
DoRA		✓	✓		88.84
Pre-Diag (<i>ours</i>)	✓		✓		88.98
<i>New Ablation Experiments</i>					
Pre-Ortho	✓			✓	87.76
Post-Ortho		✓		✓	88.93
SORA (<i>ours</i>)	✓	✓	✓	✓	89.23

Table 5: Ablation study of the weight conditioning framework. Pre/Post refers to the placement of the conditioning matrix. **D** and **P** denote diagonal and orthogonal transformations, respectively. SORA (*ours*) applies Pre-Diag calibration followed by Post-Ortho rotation.

Rotational Rank (r_P)	Trainable Params	Avg. RTE Score	Training Speed (step/s)
1	1.55%	87.90	2.17
2	1.64%	87.72	1.92
4	1.81%	88.44	1.82
8	2.16%	88.44	1.75
16	2.85%	89.53	1.72

Table 6: Sensitivity analysis of rotational rank r_p in SORA.

Related Work

Parameter-Efficient Fine-Tuning (PEFT) methods reduce the significant computational and storage costs associated with adapting large pre-trained models. Among various approaches, such as adapter modules (Houlsby et al. 2019) and prompt-tuning (Lester, Al-Rfou, and Constant 2021), methods based on low-rank adaptation are particularly prominent for their effectiveness and efficiency.

LoRA (Hu et al. 2022) is a cornerstone of this approach. It approximates weight updates as the product of two low-rank matrices, based on the assumption that these updates occupy a low-dimensional manifold. This method is also provably capable of approximating any target transformer model under mild assumptions (Zeng and Lee 2023). This elegant mechanism inspires a wave of research that unfolds along three main threads.

The first thread focuses on expressiveness and optimisation. Methods like AdaLoRA (Zhang et al. 2023) and ReLoRA (Lialin et al. 2023) replace LoRA’s static rank with dynamic allocation schedules to place capacity where it is most needed. Others enrich the update structure beyond a simple matrix product; for instance, BTT (Potapczynski et al. 2024) and CoMERA (Yang et al. 2024) introduce tensor factorisations while maintaining linear complexity. The training process itself is also refined. LoRA+ (Hayou, Ghosh, and Yu 2024) advocates for distinct learning rates for the two factors, while LORO (Zhanfeng Mo 2025) and ScaleAdam (Zhang and Pilanci 2024) reframe the optimisation on Riemannian manifolds to align updates with the parameter space’s intrinsic curvature.

A second thread pursues efficiency under tight resource constraints. QLoRA (Dettmers et al. 2023) enables fine-tuning on a single consumer GPU by quantising the model backbone to 4-bits while keeping LoRA activations in higher precision. Complementary efforts further shrink memory footprints by compressing the adapters or sharing low-rank factors across layers and tasks, as seen in methods like LoQT and VeRA (Loeschcke et al. 2024; Kopiczko, Blankevoort, and Asano 2024).

The third thread involves the principled initialisation of weights. PiSSA (Meng, Wang, and Zhang 2024) and OLoRA (Büyükakyüz 2024) initialise the low-rank factors using singular vectors or orthogonal bases from the pre-trained weights, aligning the initial update with salient directions. DoRA (Liu et al. 2024) extends this by decomposing each weight into magnitude and direction, applying LoRA only to the directional component. Such targeted adaptation yields significant gains and highlights the value of structuring the update, not just constraining its rank.

Beyond LoRA-centric innovations, other structured PEFT methods explore different constraints. One class of methods, including HRA (Yuan, Liu, and Xu 2024) and HOFT (Arcas et al. 2025), constructs complex orthogonal matrices from simple low-rank operations like Householder transformations. Another evolutionary line of structured methods seeks to balance expressive power, cost, and efficiency. This line evolves from the restrictive block-diagonal form of OFT (Qiu et al. 2023), to the denser but more expensive butterfly structure of BOFT, and finally to the efficient group-and-shuffle structure of GSOF (Gorbunov et al. 2024). Our method, SORA, also operates in this space. Inspired by the Cayley transform (Trockman and Kolter 2021), SORA constructs orthogonal matrices using a first-order Taylor expansion of the matrix exponential, an approximation also explored in (Singla and Feizi 2021). As a future direction, SORA can be integrated with orthogonalisation techniques to enhance performance.

Conclusion

This paper identifies singular value entropy as the key mechanism behind DoRA’s success leading to a unified weight conditioning framework for PEFT. Our framework yields two novel methods: (i) **Pre-Diag**, which effectively calibrates pre-trained weights by applying diagonal matrix conditioning prior to the LoRA update, enhancing both training speed and performance; and (ii) **SORA**, which introduces a parameter-efficient orthogonal rotation to perform a more powerful, norm-preserving transformation of the feature space, further improving model performance. Extensive experiments demonstrate that both methods achieve superior performance and computational efficiency over LoRA and DoRA. Ultimately, this research advocates for a paradigm shift in PEFT, moving from simple low-rank updates to more principled, structured optimizations of the weight space.

Acknowledgments

This research was supported by Guangdong S&T Programme (No. 2024B0101010003). The computing resources were provided by Pengcheng Cloud Brain.

References

- Alter, O.; Brown, P. O.; and Botstein, D. 2000. Singular value decomposition for genome-wide expression data processing and modeling. *Proceedings of the National Academy of Sciences*, 97(18): 10101–10106.
- Arcas, A. M.; Sanchís, A.; Civera, J.; and Juan-Císcar, A. 2025. HOFT: Householder Orthogonal Fine-tuning. *ArXiv*, abs/2505.16531.
- Büyükakyüz, K. 2024. OLoRA: Orthonormal Low-Rank Adaptation of Large Language Models. *ArXiv*, abs/2406.01775.
- Dettmers, T.; Pagnoni, A.; Holtzman, A.; and Zettlemoyer, L. 2023. QLoRA: Efficient Finetuning of Quantized LLMs. In *Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023*.
- Ding, N.; Qin, Y.; Yang, G.; Wei, F.; Yang, Z.; Su, Y.; Hu, S.; Chen, Y.; Chan, C.-M.; Chen, W.; et al. 2023. Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nature Machine Intelligence*, 5(3): 220–235.
- Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; and et al, A. L. 2024. The Llama 3 Herd of Models. *ArXiv*, abs/2407.21783.
- Gemma Team, T. M.; Hardin, C.; Dadashi, R.; Bhupatiraju, S.; Sifre, L.; Rivière, M.; Kale, M. S.; Love, J.; Tafti, P.; Hussenot, L.; and et al. 2024. Gemma.
- Gorbunov, M.; Yudin, N.; Soboleva, V.; Alanov, A.; Naumov, A.; and Rakhuba, M. 2024. Group and Shuffle: Efficient Structured Orthogonal Parametrization. In *Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024*.
- Hayou, S.; Ghosh, N.; and Yu, B. 2024. LoRA+: Efficient Low Rank Adaptation of Large Models. In *Forty-first International Conference on Machine Learning, ICML 2024*.
- He, P.; Liu, X.; Gao, J.; and Chen, W. 2020. DeBERTa: Decoding-enhanced BERT with Disentangled Attention. In *9th International Conference on Learning Representations, ICLR 2021*.
- Horn, R. A.; and Johnson, C. R. 2012. *Matrix analysis*. Cambridge university press.
- Houlsby, N.; Giurgiu, A.; Jastrzebski, S.; Morrone, B.; De Laroussilhe, Q.; Gesmundo, A.; Attariyan, M.; and Gelly, S. 2019. Parameter-efficient transfer learning for NLP. In *International conference on machine learning*, 2790–2799. PMLR.
- Hu, J. E.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *The Tenth International Conference on Learning Representations, ICLR 2022*.
- Kalajdzievski, D. 2023. A Rank Stabilization Scaling Factor for Fine-Tuning with LoRA. *ArXiv*, abs/2312.03732.
- Kopiczko, D. J.; Blankevoort, T.; and Asano, Y. M. 2024. VeRA: Vector-based Random Matrix Adaptation. In *The Twelfth International Conference on Learning Representations, ICLR 2024*.
- Lester, B.; Al-Rfou, R.; and Constant, N. 2021. The Power of Scale for Parameter-Efficient Prompt Tuning. In *Conference on Empirical Methods in Natural Language Processing*.
- Li, X. L.; and Liang, P. 2021. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021*, 4582–4597. Association for Computational Linguistics.
- Lialin, V.; Shivagunde, N.; Muckatira, S.; and Rumshisky, A. 2023. ReLoRA: High-Rank Training Through Low-Rank Updates. In *International Conference on Learning Representations*.
- Lion, K.; Zhang, L.; Li, B.; and He, N. 2025. PoLAR: Polar-Decomposed Low-Rank Adapter Representation. *ArXiv*, abs/2506.03133.
- Liu, J.; Su, J.; Yao, X.; Jiang, Z.; Lai, G.; Du, Y.; Qin, Y.; Xu, W.; Lu, E.; Yan, J.; Chen, Y.; Zheng, H.; Liu, Y.; Liu, S.; Yin, B.; He, W.; Zhu, H.; Wang, Y.; Wang, J.; Dong, M.; Zhang, Z.; Kang, Y.; Zhang, H.; Xu, X.; Zhang, Y.; Wu, Y.; Zhou, X.; and Yang, Z. 2025. Muon is Scalable for LLM Training. *ArXiv*, abs/2502.16982.
- Liu, S.-Y.; Wang, C.-Y.; Yin, H.; Molchanov, P.; Wang, Y.-C. F.; Cheng, K.-T.; and Chen, M.-H. 2024. DoRA: Weight-Decomposed Low-Rank Adaptation. In *Forty-first International Conference on Machine Learning, ICML 2024*.
- Loeschcke, S.; Toftrup, M. B.; Kastoryano, M. J.; Belongie, S. J.; and Snæbjarnarson, V. 2024. LoQT: Low-Rank Adapters for Quantized Pretraining. In *Neural Information Processing Systems*.
- Meng, F.; Wang, Z.; and Zhang, M. 2024. PiSSA: Principal Singular Values and Singular Vectors Adaptation of Large Language Models. In *Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024*.
- Potapczynski, A.; Qiu, S.; Finzi, M.; Ferri, C.; Chen, Z.; Goldblum, M.; Bruss, B.; Sa, C. D.; and Wilson, A. G. 2024. Searching for Efficient Linear Layers over a Continuous Space of Structured Matrices. In *Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024*.
- Qiu, Z.; Yu, Liu, W.; Feng, H.; Xue, Y.; Feng, Y.; Liu, Z.; Zhang, D.; Weller, A.; and Scholkopf, B. 2023. Controlling Text-to-Image Diffusion by Orthogonal Finetuning. In *Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023*.
- Roy, O.; and Vetterli, M. 2007. The effective rank: A measure of effective dimensionality. In *2007 15th European Signal Processing Conference*, 606–610.
- Rudelson, M.; and Vershynin, R. 2007. Sampling from large matrices: An approach through geometric functional analysis. *Journal of the ACM (JACM)*, 54(4): 21–es.
- Saratchandran, H.; Wang, T. X.; and Lucey, S. 2024. Weight Conditioning for Smooth Optimization of Neural Networks. In *European Conference on Computer Vision*.
- Singla, S.; and Feizi, S. 2021. Skew Orthogonal Convolutions. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021*, volume 139 of *Proceedings of Machine Learning Research*, 9756–9766. PMLR.

Trockman, A.; and Kolter, J. Z. 2021. Orthogonalizing Convolutional Layers with the Cayley Transform. In *9th International Conference on Learning Representations, ICLR 2021*.

Wang, A.; Singh, A.; Michael, J.; Hill, F.; Levy, O.; and Bowman, S. R. 2018. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. In *BlackboxNLP@EMNLP*.

Wang, S.; Yu, L.; and Li, J. 2024. LoRA-GA: Low-Rank Adaptation with Gradient Approximation. In *Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024*.

Yang, Z.; Choudhary, S.; Xie, X.; Gao, C.; Kunzmann, S.; and Zhang, Z. 2024. CoMERA: Computing- and Memory-Efficient Training via Rank-Adaptive Tensor Optimization. In *Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024*.

Yen, J.-N.; Si, S.; Meng, Z.; Yu, F.; Duvvuri, S. S.; Dhillon, I. S.; Hsieh, C.-J.; and Kumar, S. 2024. LoRA Done RITE: Robust Invariant Transformation Equilibration for LoRA Optimization. *arXiv preprint arXiv:2410.20625*.

Yu, L.; Jiang, W.; Shi, H.; Yu, J.; Liu, Z.; Zhang, Y.; Kwok, J. T.; Li, Z.; Weller, A.; and Liu, W. 2024. Metamath: Bootstrap your own mathematical questions for large language models. In *The Twelfth International Conference on Learning Representations, ICLR 2024*.

Yuan, S.; Liu, H.; and Xu, H. 2024. Bridging The Gap between Low-rank and Orthogonal Adaptation via Householder Reflection Adaptation. In *Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024*.

Zeng, Y.; and Lee, K. 2023. The Expressive Power of Low-Rank Adaptation. *ArXiv*, abs/2310.17513.

Zhanfeng Mo, S. J. P., Long-kai Huang. 2025. Parameter and Memory Efficient Pretraining via Low-rank Riemannian Optimization. In *The Thirteenth International Conference on Learning Representations*.

Zhang, F.; and Pilanci, M. 2024. Riemannian Preconditioned LoRA for Fine-Tuning Foundation Models. In *Forty-first International Conference on Machine Learning, ICML 2024*.

Zhang, Q.; Chen, M.; Bukharin, A.; Karampatziakis, N.; He, P.; Cheng, Y.; Chen, W.; and Zhao, T. 2023. AdaLoRA: Adaptive Budget Allocation for Parameter-Efficient Fine-Tuning.