

Efficiently Seeking Flat Minima for Better Generalization in Fine-Tuning Large Language Models and Beyond

Jiaxin Deng^{1*}, Qingcheng Zhu^{2*}, Junbiao Pang^{1†}, Linlin Yang³, Zhongqian Fu⁴, Baochang Zhang^{5‡}

¹School of Information Science and Technology, Beijing University of Technology, Beijing, China

²School of Electronic Information Engineering, Beihang University, Beijing, China

³State Key Laboratory of Media Convergence and Communication, Communication University of China, Beijing, China

⁴Huawei Noah’s Ark Lab, China

⁵Hangzhou Research Institute, School of Artificial Intelligence, Beihang University, China

dengjiaxin@emails.bjut.edu.cn, zhuqc@buaa.edu.cn, junbiao_pang@bjut.edu.cn, lyang@cuc.edu.cn,

fuzhongqian@huawei.com, bczhang@buaa.edu.cn

Abstract

Little research explores the correlation between the expressive ability and generalization ability of the low-rank adaptation (LoRA). Sharpness-Aware Minimization (SAM) improves model generalization for both Convolutional Neural Networks (CNNs) and Transformers by encouraging convergence to locally flat minima. However, the connection between sharpness and generalization has not been fully explored for LoRA due to the lack of tools to either empirically seek flat minima or develop theoretical methods. In this work, we propose Flat Minima LoRA (FMLoRA) and its efficient version, *i.e.*, EFMLoRA, to seek flat minima for LoRA. Concretely, we theoretically demonstrate that perturbations in the full parameter space can be transferred to the low-rank subspace. This approach eliminates the potential interference introduced by perturbations across multiple matrices in the low-rank subspace. Our extensive experiments on large language models and vision-language models demonstrate that EFMLoRA achieves optimization efficiency comparable to that of LoRA while simultaneously attaining comparable or even better performance. For example, on the GLUE dataset with RoBERTa-large, EFMLoRA outperforms LoRA and full fine-tuning by 1.0% and 0.5% on average, respectively. On vision-language models, *e.g.*, Qwen-VL-Chat, there are performance improvements of 1.5% and 1.0% on the SQA and VizWiz datasets, respectively. These empirical results also verify that the generalization of LoRA is closely related to sharpness, which is omitted by previous methods.

Introduction

Parameter-Efficient Fine-Tuning (PEFT) methods only update a small subset of parameters, *e.g.*, adapters (Hu et al. 2022) or prompt weights (Li and Liang 2021) for Large language models (LLMs) with substantially lower memory and computational costs. Specifically, Low-Rank Adaptation (LoRA) (Hu et al. 2022) stands out for achieving per-

*These authors contributed equally.

†Corresponding authors.

‡Project leader.

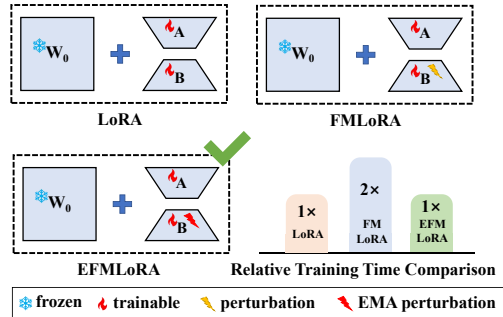


Figure 1: Comparison of Methods: LoRA, FMLoRA, and EFMLoRA.

formance comparable to full fine-tuning (FT) while being considerably more efficient.

Many works have been proposed to enhance the performance of LoRA by introducing more dedicated budgets for rank allocation (Zhang et al. 2023b), decomposing optimization for direction and magnitude updates (Liu et al. 2024), or designing better initialization strategies for LoRA parameters (Meng, Wang, and Zhang 2024), etc. These studies demonstrate the significant potential to improve LoRA performance. However, most existing approaches fail to effectively address bias inheritance, where LLMs may propagate and amplify their inherent biases, significantly impacting model performance and robustness on downstream tasks (Li et al. 2025). Therefore, a natural question is: how to model and understand the generalization of LoRA for various LLMs and beyond, *e.g.*, vision-language models?

It is widely believed that a flatter loss landscape can lead to better generalization performance (Hochreiter and Schmidhuber 1994) (Hochreiter and Schmidhuber 1997). For instance, Foret et al. proposed Sharpness-Aware Minimization (SAM) (Foret et al. 2021), which seeks parameter regions where the training loss remains uniformly flat. SAM and its variants have demonstrated State-Of-The-Art (SOTA) performances across various applications, such as

classification (Kwon et al. 2021), transfer learning (Zhuang et al. 2022), domain generalization (Dong et al. 2024) and federated learning (Dai et al. 2023).

To the best of our knowledge, compared to theoretical analysis, *e.g.*, (Neyshabur et al. 2017), empirically connecting sharpness and generalization ability of LoRA is a practical approach, *e.g.*, (Andriushchenko et al. 2023). For the second line of research, a naive approach is to combine SAM with LoRA. However, if perturbations in SAM are applied simultaneously to two low-rank subspaces of LoRA, they may change the maximum loss within the neighborhood of LoRA’s full parameter space (Dinh et al. 2017a); besides, SAM incurs a computational cost twice that of Stochastic Gradient Descent (SGD) (Deng et al. 2024). The key question in the second line of research is how to efficiently find flat minima in LoRA, aiming to better understand the connection between sharpness and generalization.

In this paper, we propose a novel PEFT method, FM-LoRA, that promotes convergence toward flatter minima. Specifically, we theoretically uncover that perturbations in the full parameter space can be equivalently reparameterized as perturbations within the low-rank space. In addition, we propose EFMLoRA to accelerate FMLoRA by an Exponential Moving Average (EMA) strategy. We validate that EFMLoRA improves generalization performance on downstream tasks while maintaining computational efficiency comparable to that of LoRA. Fig. 1 compares three methods: LoRA, FMLoRA, and EFMLoRA. We conducted comprehensive experiments on diverse tasks (fine-tuning, few-shot learning) and various model types (RoBERTa (Liu et al. 2019), GPT-2 (Radford et al. 2019), CLIP (Zanella and Ben Ayed 2024), Qwen-VL-Chat (Bai et al. 2023)) and scales. We find that EFMLoRA achieves model accuracy very close to, or even surpass both full fine-tuning and LoRA across many tasks. Our main contribution can be summarized as follows:

- We propose FMLoRA, a novel PEFT training method that integrates SAM into the LoRA framework. Furthermore, EFMLoRA provides an efficient tool for empirically understanding the connection between sharpness and generalization in LLMs and beyond. We empirically show that reducing sharpness is highly correlated with improved generalization in PEFT tasks, which has been rarely explored in PEFT studies before.
- We conduct comprehensive experiments on LLMs (*e.g.*, RoBERTa, GPT-2) and vision-language models (*e.g.*, CLIP, Qwen-VL-Chat) across various tasks including fine-tuning and few-shot learning. Results show that EFMLoRA achieves optimization efficiency comparable to that of LoRA while simultaneously attaining comparable or even better performance.

Related Works

Low-rank Adaption

Hu et al. proposed LoRA (Hu et al. 2022) as a PEFT method that introduced low-rank adapters into each layer of a pre-trained model. Recent advancements in LoRA can be

broadly categorized into two directions: 1) advanced architectures and 2) optimization methods. In the first research line, for example, LoraHub (Huang et al. 2023) trained multiple adapters and strategically combined them based on the domain during inference. LoRA-FA (Zhang et al. 2023a) chose to freeze the projection-down weight of \mathbf{A} and update the projection-up weight of \mathbf{B} in each LoRA layer. DoRA (Liu et al. 2024) improved LoRA by incorporating a learnable magnitude vector to re-scale the normalized product of low-rank matrices. HydraLoRA (Tian et al. 2024) extended the LoRA framework with an asymmetric architecture that shared a common \mathbf{A} matrix for efficiency while dynamically assigning samples to multiple \mathbf{B} matrices via a MoE mechanism. In the second line, for example, LoRA+ (Hayou, Ghosh, and Yu 2024) applied different learning rates to the two low-rank matrices. Additionally, Galore (Zhao et al. 2024) employed SVD to compress the gradients and its first and second momentum of full training into a low-rank space, thereby reducing the memory footprint during pre-training and fine-tuning. Recently, Li et al. (Li et al. 2024a) proposed combining SAM with LoRA for better generalization, but they used random perturbation. Our method belongs to the second research line. Different from (Li et al. 2024a), our method transfers the perturbation from the full parameter space to a single low-rank parameter space without changing the maximum perturbed loss, avoiding misalignment with SAM’s training behavior.

Sharpness and Generalization Ability

Research on the relationship between sharpness and generalization could be traced back to (Hochreiter and Schmidhuber 1997). Following the observation by (Keskar et al. 2017) that larger batch sizes tended to increase sharpness and generalization error. (Jastrzkebski et al. 2017) extended this by finding a correlation between the sharpness and the ratio of learning rate to batch size. (Dinh et al. 2017b) showed that one can easily construct networks with good generalization but with arbitrary large sharpness by reparameterization. (Jiang et al. 2020) performed a large-scale empirical study on various generalization measures and showed that sharpness-based measures have the highest correlation with generalization. Theoretical understandings on the generalization error using sharpness-related measures were provided in (Neyshabur et al. 2017), (Wang and Mao 2022). Collectively, these studies justified the goal of seeking flatter minima to improve generalization. However, to the best of our knowledge, the correlation between sharpness and generalization for LoRA has barely been discussed due to the lack of theoretical understanding or efficient tools for empirical analysis. Our method provides an efficient tool for empirical analysis in this domain.

Recap of SAM

Foret et al. (Foret et al. 2021) proposed the SAM to enhance model generalization as follows:

$$\min_{\mathbf{w}} [(\max_{\|\boldsymbol{\varepsilon}\| \leq \rho} L(\mathbf{w} + \boldsymbol{\varepsilon}) - L(\mathbf{w})) + L(\mathbf{w}) + \lambda \|\mathbf{w}\|_2^2], \quad (1)$$

where \mathbf{w} represents the weights of the network, $\boldsymbol{\varepsilon}$ represents the perturbation of weights \mathbf{w} in a Euclidean ball with the

radius ρ ($\rho > 0$), $L(\cdot)$ is the loss function, and $\lambda\|\mathbf{w}\|_2^2$ is a standard L2 regularization term.

SAM utilizes Taylor expansion to search for the maximum perturbed loss ($\max_{\|\varepsilon\| \leq \rho} L(\mathbf{w} + \varepsilon)$) in local parameter space as follows:

$$\arg \max_{\|\varepsilon\| \leq \rho} L(\mathbf{w} + \varepsilon) \approx \arg \max_{\|\varepsilon\| \leq \rho} \varepsilon^\top \nabla_{\mathbf{w}} L(\mathbf{w}). \quad (2)$$

By solving Eq. (2), SAM obtains the perturbation as follows:

$$\hat{\varepsilon} = \rho \nabla_{\mathbf{w}} L(\mathbf{w}) / \|\nabla_{\mathbf{w}} L(\mathbf{w})\|. \quad (3)$$

Substituting the perturbation $\hat{\varepsilon}$ back into Eq. (1), we then have:

$$\begin{aligned} \nabla_{\mathbf{w}} \max_{\|\varepsilon\| \leq \rho} L(\mathbf{w} + \varepsilon) &\approx \nabla_{\mathbf{w}} L(\mathbf{w} + \hat{\varepsilon}(\mathbf{w})) \\ &= \nabla_{\mathbf{w}} L(\mathbf{w})|_{\mathbf{w} + \hat{\varepsilon}(\mathbf{w})} + \frac{d\hat{\varepsilon}(\mathbf{w})}{d\mathbf{w}} \nabla_{\mathbf{w}} L(\mathbf{w})|_{\mathbf{w} + \hat{\varepsilon}(\mathbf{w})}. \end{aligned} \quad (4)$$

By dropping the second-order terms in Eq.(4), SAM calculates the gradient at $\mathbf{w} + \hat{\varepsilon}$ as follows:

$$\nabla_{\mathbf{w}} \max_{\|\varepsilon\| \leq \rho} L(\mathbf{w} + \varepsilon) \approx \nabla_{\mathbf{w}} L(\mathbf{w})|_{\mathbf{w} + \hat{\varepsilon}}. \quad (5)$$

Finally, SAM uses the gradients from Eq. (5) for optimization.

SAM Variants

Recently, SAM variants could be broadly categorized into three groups: 1) studies on the perturbation radius ρ in SAM, 2) studies that speed up the optimization process of SAM, and 3) redefinitions of sharpness in SAM. For the first direction, Kwon et al. (Kwon et al. 2021) proposed Adaptive SAM (ASAM), which adapted the perturbation radius in a scale-aware manner, allowing SAM to be effectively applied to scale-invariant neural networks. For the second group, Kim et al. (Kim et al. 2023) introduced a multi-step ascent approach to improve SAM. Li et al. (Li et al. 2024b) introduced Friendly SAM (F-SAM), which improved generalization by removing the detrimental influence of the full gradient component and instead utilizing batch-specific gradients to guide optimization more effectively. For the third group, Zhuang et al. (Zhuang et al. 2022) pointed out that SAM did not always favor flat minima. Consequently, they proposed GSAM, which minimized the surrogate gap and the perturbed loss to better encourage flatness. Zhang et al. introduced the first-order flatness (Zhang et al. 2023c), which assessed the maximal gradient norm within a perturbation radius. Consequently, they proposed GAM which explicitly seeks minima characterized by uniformly small curvature.

Method

SAM on LoRA

LoRA achieves parameter efficiency by modeling the low-rank decomposed weight (Li et al. 2022). Specifically, the weight change for each layer $\mathbf{W}_0 \in \mathbb{R}^{n \times m}$ is represented as $\Delta \mathbf{W} = s\mathbf{B}\mathbf{A}$, where s is a scaling factor, $\mathbf{B} \in \mathbb{R}^{n \times r}$,

$\mathbf{A} \in \mathbb{R}^{r \times m}$, with $\text{rank } r \ll \min(n, m)$. Given an input \mathbf{x} , the forward is as follows:

$$\mathbf{y} = \mathbf{W}_0 \mathbf{x} + \Delta \mathbf{W} \mathbf{x} = (\mathbf{W}_0 + s\mathbf{B}\mathbf{A})\mathbf{x}, \quad (6)$$

where matrix \mathbf{A} is typically initialized by the Kaiming’s method (He et al. 2015), \mathbf{B} is set to zeros. \mathbf{W}_0 remains unchanged during fine-tuning, while \mathbf{B} and \mathbf{A} are trained. During inference, $\Delta \mathbf{W}$ is merged into \mathbf{W}_0 .

If SAM is naively combined with LoRA, the optimization loss can be rewritten as follows:

$$\min_{\mathbf{A}, \mathbf{B}} \max_{\substack{\|\mathbf{E}^{\mathbf{A}}\|_F \leq \rho, \\ \|\mathbf{E}^{\mathbf{B}}\|_F \leq \rho}} L(\mathbf{W}_0 + s(\mathbf{B} + \mathbf{E}^{\mathbf{B}})(\mathbf{A} + \mathbf{E}^{\mathbf{A}})), \quad (7)$$

where $\mathbf{E}^{\mathbf{B}} \in \mathbb{R}^{n \times r}$ and $\mathbf{E}^{\mathbf{A}} \in \mathbb{R}^{r \times m}$ represent the perturbations applied to the parameters \mathbf{B} and \mathbf{A} , respectively, and ρ is the radius of perturbations. There are two key challenges:

- Two separate perturbations in two low-rank subspaces interfere with each other, leading to an inconsistency between the maximum loss obtained when perturbing in the low-rank subspaces and the maximum loss obtained when perturbing in the full parameter space.
- SAM requires computing gradients twice per iteration, resulting in approximately twice the computational cost compared to LoRA.

FMLoRA

To deal with the first challenge, we propose to re-parameterize the perturbation from the full parameter space to a single low-rank parameter space. Concretely, the loss in the full parameter space can be formulated as follows:

$$\min_{\mathbf{A}, \mathbf{B}} \max_{\|\mathbf{E}^{\mathbf{W}}\|_F \leq \rho} L(\mathbf{W}_0 + s\mathbf{B}\mathbf{A} + \mathbf{E}^{\mathbf{W}}). \quad (8)$$

To solve the minimax problem in Eq. (8), it is necessary to first find optimal $\hat{\mathbf{E}}^{\mathbf{W}} \in \mathbb{R}^{n \times m}$. Analogous to SAM, we approximate the optimal perturbation $\hat{\mathbf{E}}^{\mathbf{W}}$ to maximize $L(\mathbf{W} + \mathbf{E}^{\mathbf{W}})$ where $\mathbf{W} = \mathbf{W}_0 + s\mathbf{B}\mathbf{A}$ as follows:

$$\hat{\varepsilon}^{\mathbf{w}} = \rho \text{sign}(\mathbf{g}^{\mathbf{w}}) \frac{\mathbf{g}^{\mathbf{w}}}{\|\mathbf{g}^{\mathbf{w}}\|}, \quad (9)$$

where $\mathbf{g}^{\mathbf{w}} = \text{Vector}(\nabla L_{\mathbf{W}}(\mathbf{W}))$ and $\hat{\varepsilon}^{\mathbf{w}} = \text{Vector}(\hat{\mathbf{E}}^{\mathbf{W}})$, in which the $\text{Vector}(\cdot)$ function represents a vectorized operation. However, the solution for $\hat{\mathbf{E}}^{\mathbf{W}}$ explicitly depends on the gradient of the matrix \mathbf{W} . That is, the form of solution in Eq. (9) is undesirable since $\nabla L_{\mathbf{W}}(\mathbf{W})$ is unknown during LoRA optimization.

In this paper, we propose to approximate the unknown gradient $\nabla L_{\mathbf{W}}(\mathbf{W})$ using standard LoRA gradients, which can be computed in two ways:

$$(1) \quad \nabla L_{\mathbf{W}}(\mathbf{W}) = \frac{1}{s} \nabla L_{\mathbf{B}}(\mathbf{W}_0 + s\mathbf{B}\mathbf{A})(\mathbf{A}^\top)^+, \quad (10)$$

$$(2) \quad \nabla L_{\mathbf{W}}(\mathbf{W}) = \frac{1}{s} (\mathbf{B}^\top)^+ \nabla L_{\mathbf{A}}(\mathbf{W}_0 + s\mathbf{B}\mathbf{A}), \quad (11)$$

where $(\mathbf{A}^\top)^+$ and $(\mathbf{B}^\top)^+$ represent the pseudo-inverse of \mathbf{A}^\top and \mathbf{B}^\top , respectively. The accuracy of the pseudo-inverse depends on the condition number of matrix. A

smaller condition number leads to a more accurate pseudo-inverse. Matrices with lower condition numbers are better suited for stable representation. In LoRA, we found that the condition number is typically low, around 3.

To obtain a more accurate estimate of the gradient of the full weights, we combine the above two approaches to compute $\nabla L_{\mathbf{W}}(\mathbf{W})$ as follows:

$$\begin{aligned} \overline{\nabla L_{\mathbf{W}}(\mathbf{W})} &= 0.5 \left(\frac{1}{s} \nabla L_{\mathbf{B}}(\mathbf{W}_0 + s\mathbf{B}\mathbf{A})(\mathbf{A}^\top)^\dagger \right. \\ &\quad \left. + \frac{1}{s} (\mathbf{B}^\top)^\dagger \nabla L_{\mathbf{A}}(\mathbf{W}_0 + s\mathbf{B}\mathbf{A}) \right). \end{aligned} \quad (12)$$

Let $\bar{\mathbf{g}}^{\mathbf{W}} = \text{Vector}(\overline{\nabla L_{\mathbf{W}}(\mathbf{W})})$. Then the perturbation in Eq. (9) could be rewritten as follows:

$$\bar{\mathbf{E}}^{\mathbf{W}} = \text{Matrix}(\rho \text{sign}(\bar{\mathbf{g}}^{\mathbf{W}}) \frac{\bar{\mathbf{g}}^{\mathbf{W}}}{\|\bar{\mathbf{g}}^{\mathbf{W}}\|}), \quad (13)$$

where $\text{Matrix}(\cdot)$ denotes the operation that converts a vector into a matrix. We transfer the perturbation from the full parameter space to a single low-rank parameter space without changing the maximum loss in the local region of the parameters. We apply no perturbation to matrix \mathbf{A} , *i.e.*, $\mathbf{E}^{\mathbf{A}} = \mathbf{0}$, and ensure that the loss under perturbations in the low-rank subspace in Eq. (7) matches the inner maximum loss in Eq. (8), as follows:

$$\begin{aligned} L(\mathbf{W}_0 + s(\mathbf{B} + \mathbf{E}^{\mathbf{B}})\mathbf{A}) \\ = \max_{\|\mathbf{E}^{\mathbf{W}}\|_F \leq \rho} L(\mathbf{W}_0 + s\mathbf{B}\mathbf{A} + \mathbf{E}^{\mathbf{W}}). \end{aligned} \quad (14)$$

Substituting $\bar{\mathbf{E}}^{\mathbf{W}}$ into Eq. (14), we obtain:

$$\mathbf{E}^{\mathbf{B}} \approx \frac{1}{s} \bar{\mathbf{E}}^{\mathbf{W}} \mathbf{A}^\dagger, \quad (15)$$

where \mathbf{A}^\dagger is the pseudo-inverse of \mathbf{A} . An alternative approach is to transfer the perturbation to matrix \mathbf{A} . Following the observations from HydraLoRA (Tian et al. 2024), matrix \mathbf{A} shows high parameter similarity across heads, likely due to initialization, making it capture domain-common features, while matrix \mathbf{B} remains distinct and domain-specific. Since different tasks require different perturbations, we adopt the approach of transferring the perturbation to the matrix \mathbf{B} , as expressed in Eq. (14). The detailed derivation of Eq. (10) and the pseudo-algorithm for FMLoRA are provided in the supplementary file.

Balancedness of FMLoRA. Balancedness is well-appreciated in domains such as matrix factorization/sensing (Ge, Jin, and Zheng 2017) (Du, Hu, and Lee 2018). It is also observed that balanced neural networks are easier to optimize relative to unbalanced ones (Neyshabur, Salakhutdinov, and Srebro 2015). Recently, Balancedness $B_t := \frac{1}{2}(\|\mathbf{x}_t\|^2 - \|\mathbf{y}_t\|^2)$ (where \mathbf{x}_t and \mathbf{y}_t are variables) turns out to be an intriguing alternative to sharpness on the scale-invariant problem (Li, Zhang, and He 2024).

To investigate the balancedness of our proposed method, we express the update process of FMLoRA analogously to

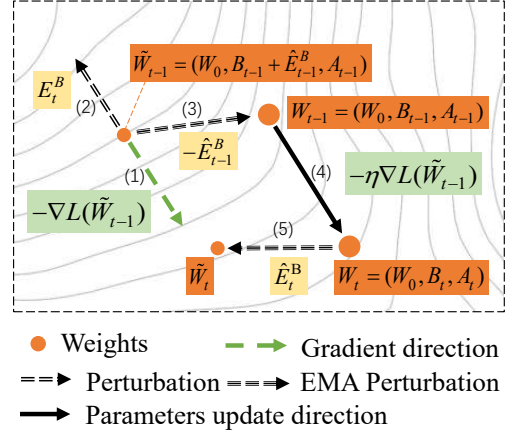


Figure 2: Parameter update process for EFMLoRA.

Eq.(4) in (Li, Zhang, and He 2024) as follows:

$$\begin{aligned} \tilde{\mathbf{x}}_t &= \mathbf{x}_t + \rho \frac{1}{s} \frac{\mathbf{G}_t}{\|\mathbf{G}_t\|_F} \mathbf{y}_t^\dagger, \quad \tilde{\mathbf{y}}_t = \mathbf{y}_t, \\ \mathbf{g}_{\tilde{\mathbf{x}}_t} &= \tilde{\mathbf{G}}_t \tilde{\mathbf{y}}_t, \quad \mathbf{g}_{\tilde{\mathbf{y}}_t} = \tilde{\mathbf{G}}_t^\top \tilde{\mathbf{x}}_t, \\ \mathbf{x}_{t+1} &= \mathbf{x}_t - \eta \mathbf{g}_{\tilde{\mathbf{x}}_t}, \quad \mathbf{y}_{t+1} = \mathbf{y}_t - \eta \mathbf{g}_{\tilde{\mathbf{y}}_t}, \end{aligned} \quad (16)$$

where $\mathbf{x}_t = \text{Vector}(\mathbf{B}_t)$, $\mathbf{y}_t = \text{Vector}(\mathbf{A}_t)$, $\mathbf{G}_t = \nabla L(\mathbf{x}_t \mathbf{y}_t^\top)$ is the gradient of the full parameter space at the original parameter point, $\tilde{\mathbf{G}}_t = \nabla L(\tilde{\mathbf{x}}_t \tilde{\mathbf{y}}_t^\top)$ is the gradient of the full parameter space at the perturbed parameter point, and \mathbf{y}_t^\dagger is the pseudo inverse of \mathbf{y}_t .

Theorem 1 Let $B_t := \frac{1}{2}(\|\mathbf{x}_t\|^2 - \|\mathbf{y}_t\|^2)$. For the learning rate $\eta \Rightarrow 0$, the limiting flow of FMLoRA guarantees that:

$$\left| \frac{1}{2} \frac{d(\|\mathbf{x}_t\|^2 - \|\mathbf{y}_t\|^2)}{dt} \right| \leq \left| \rho \frac{1}{s} \frac{1}{\|\mathbf{y}_t\|} \|\mathbf{g}_{\tilde{\mathbf{x}}_t}\| \right|. \quad (17)$$

Theorem 1 indicates that the balancedness of FMLoRA is influenced by the perturbation range ρ , the norm of the gradient at the perturbed point, the ℓ_2 -norm of \mathbf{y}_t , and the scale constraint of LoRA. To ensure that the balancedness of FMLoRA gradually decreases during training, we reduce ρ progressively. In addition, the norm of the gradient with respect to \mathbf{y}_t at the perturbed point also decreases due to the weight decay. The ℓ_2 -norm of \mathbf{y}_t is bounded within a certain range, these factors collectively contribute to the reduction in the balancedness of FMLoRA.

Efficient FMLoRA

The optimization processes of FMLoRA also require two gradient computations per iteration. To enhance optimization efficiency, we propose Efficient FMLoRA (EFMLoRA), which estimates the subsequent perturbation $\mathbf{E}^{\mathbf{B}}$ in Eq. (15) by maintaining an Exponential Moving Average (EMA) of previous perturbations as follows:

$$\hat{\mathbf{E}}_t^{\mathbf{B}} = (1 - \beta) \hat{\mathbf{E}}_{t-1}^{\mathbf{B}} + \beta \mathbf{E}_t^{\mathbf{B}}, \quad (18)$$

where $\beta \in (0, 1)$ is the momentum coefficient that determines the update rate of the exponential moving average.

$\mathbf{E}_t^{\mathbf{B}}$ is the perturbation on matrix \mathbf{B}_t at t -th iteration, $\hat{\mathbf{E}}_t^{\mathbf{B}}$ is the EMA perturbation at t -th iteration. Fig. 2 illustrates the parameter update process of EFMLoRA: (1) Calculate the gradient at the perturbed point $(\mathbf{W}_0, \mathbf{B}_{t-1} + \hat{\mathbf{E}}_{t-1}^{\mathbf{B}}, \mathbf{A}_{t-1})$. (2) Calculate the perturbation $\mathbf{E}_t^{\mathbf{B}} = \frac{1}{s} \hat{\mathbf{E}}_t^{\mathbf{B}} \mathbf{A}_{t-1}^+$. (3) Return to the original parameter point $(\mathbf{W}_0, \mathbf{B}_{t-1}, \mathbf{A}_{t-1})$. (4) Update the parameters to $(\mathbf{W}_0, \mathbf{B}_t, \mathbf{A}_t)$. (5) Calculate the EMA perturbation by Eq.(18) and update the parameters to the next perturbed point $(\mathbf{W}_0, \mathbf{B}_t + \hat{\mathbf{E}}_t^{\mathbf{B}}, \mathbf{A}_t)$. During this optimization process, each optimization step requires only a single forward and backward. The algorithmic pseudocode is provided in the supplementary file.

To theoretically analyze the error of EFMLoRA, some necessary assumptions are listed below, all of which are common and standard when analyzing SAM optimization (Du et al. 2022) (Zhuang et al. 2022).

Assumption 1 (Smooth) $L(\mathbf{w})$ is τ -Lipschitz smooth in \mathbf{w} , i.e., $\|\nabla L(\mathbf{w}) - \nabla L(\mathbf{v})\| \leq \tau \|\mathbf{w} - \mathbf{v}\|$.

Assumption 2 (Bounded gradients). By the assumption that an upper bound exists on the gradient of each mini-batch. There exists $G > 0$ for each mini-batch such that $\mathbb{E}[\|\nabla L(\mathbf{w})\|] \leq G$.

Assumption 3 (Bounded variance of stochastic gradients). Given the training set \mathbf{D} and a mini-batch $\mathbf{B} \in \mathbf{D}$. There exists $\sigma \geq 0$, the variance of stochastic gradient $L_{\mathbf{B}}(\mathbf{w})$ is bounded by $\mathbb{E}[\|\nabla L_{\mathbf{B}}(\mathbf{w}) - \nabla L_{\mathbf{D}}(\mathbf{w})\|^2] \leq \sigma^2$.

Assumption 4 (Convex) We assume that the loss function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ is convex and twice differentiable over an open domain. That is, for all $x, y \in \text{dom}(f)$, it satisfies: $f(y) \geq f(x) + \nabla f(x)^\top (y - x)$.

This convexity assumption is reasonable in the fine-tuning stage, as the model is typically close to a local minimum and the loss landscape is approximately convex in a local neighborhood (Jang, Lee, and Ryu 2024).

Theorem 2 [EMA perturbation approximate perturbation of SAM due to the convex of the loss landscape] Assume that during fine-tuning, the solution is already close to a local minimum and the local loss function is convex. Let the model weights at i -th iteration be \mathbf{w}_t . Under Assumptions 1, 2, and 3, let $\rho_t = \frac{\rho_0}{\sqrt{t}}$, the error between the sharpness calculated using the EMA perturbation (S^{EMA}) and that calculated using the original SAM perturbation (S^{SAM}) is bounded as follows:

$$\begin{aligned} & \left| \underbrace{[L(\mathbf{w}_t + \hat{\mathbf{e}}_{t-1}) - L(\mathbf{w}_t)]}_{S^{\text{EMA}}} - \underbrace{[L(\mathbf{w}_t + \tilde{\mathbf{e}}_t) - L(\mathbf{w}_t)]}_{S^{\text{SAM}}} \right| \\ & \leq \left(\left(1 + (1 - \beta)^{t-1}\right) \tau \rho_0 + G + \sigma^2 \right) \\ & \quad \cdot \left(\left(1 + (1 - \beta)^{t-1}\right) \rho_0 + \frac{\rho_0}{\sqrt{t}} \right). \end{aligned} \quad (19)$$

Theorem 2 demonstrates that as t increases, the difference between S^{EMA} and S^{SAM} gradually decreases. The perturbation estimated by the EMA can effectively approximate the original SAM perturbation.

Memory and Time Complexity

LoRA reduces the number of trainable parameters by decomposing weight updates as $\Delta \mathbf{W} \approx \mathbf{B}\mathbf{A}$, where $\mathbf{B} \in \mathbb{R}^{n \times r}$ and $\mathbf{A} \in \mathbb{R}^{r \times m}$ with $r \ll \min(n, m)$. Both FMLoRA and EFMLoRA retain this parameter efficiency:

$$\begin{aligned} P_{\text{LoRA}} &= P_{\text{FMLoRA}} = P_{\text{EFMLoRA}} \\ &= O(nr + rm) \ll O(nm). \end{aligned} \quad (20)$$

However, FMLoRA and EFMLoRA introduce additional memory overhead. Specifically, FMLoRA temporarily stores the original values of \mathbf{B} and \mathbf{A} , as well as the gradients of \mathbf{A} . The memory usage of FMLoRA is calibrated as follows:

$$M_{\text{FMLoRA}} = M_{\text{LoRA}} + O(1.5 \times (nr + rm)), \quad (21)$$

where M_{LoRA} indicates the memory required by LoRA. The memory of EFMLoRA needs to maintain the EMA perturbation on \mathbf{B} as follows:

$$M_{\text{EFMLoRA}} = M_{\text{LoRA}} + O(2 \times (nr + rm)). \quad (22)$$

Notably, modern optimizers like AdamW already require $O(2 \times (nr + rm))$ memory for momentum and second-moment statistics when applied to LoRA.

For time complexity, suppose that the time complexity of optimizing the model with LoRA is $O(T)$, which mainly includes the time for forward and backward. Theoretically, the time complexity of FMLoRA is approximately as follows:

$$T_{\text{FMLoRA}} \approx O(2T) = 2 \times T_{\text{LoRA}}. \quad (23)$$

In contrast, the time complexity of EFMLoRA can be approximated as follows:

$$T_{\text{EFMLoRA}} \approx O(T) = T_{\text{LoRA}}. \quad (24)$$

We implement QR decomposition by Householder transformations, with time complexity of $O(r^2n)$ for an $r \times n$ matrix, e.g., r is rank, n is the input dimension in LORA.

Experiments and Discussions

The best and second-best results are highlighted in bold and underline, respectively. Additional experimental details are provided in the supplementary file.

Experiments on Large Language Models

Few-shot with RoBERTa-large. We first consider few-shot learning with EFMLoRA. Following the setup of (Li, Zhang, and He 2024), we adopt RoBERTa-large—a 355M-parameter language model—as the backbone. The results in Table 1 show that FMLoRA outperforms all other methods with the highest average score (83.1), particularly excelling on SST-2, SNLI, and MNLI. EFMLoRA follows closely with an average score of 82.3. It consistently surpasses baseline LoRA (+2.3), LoRA-SAM (+1.0), and both BAR variants. These results highlight its superior generalization ability under distribution shift and limited supervision. We conjecture that the performance gap between SAM and EFMLoRA comes from EFMLoRA eliminating the mutual interference between perturbations in the two low-rank subspaces.

RoBERTa	SST-2	SST-5	SNLI	MNLI	RTE	TREC	avg.↑
Zero-Shot*	79.0	35.5	50.2	48.8	51.4	32.0	49.5
LoRA*	91.1 \pm 0.8	52.3 \pm 2.9	84.3 \pm 0.3	78.1 \pm 1.3	77.5 \pm 2.3	96.6 \pm 1.0	80.0
LoRA-SAM*	92.2 \pm 0.4	54.2 \pm 2.0	85.5 \pm 0.7	78.7 \pm 1.0	80.6 \pm 4.3	96.7 \pm 0.2	81.3
LoRA-oBAR*	91.5 \pm 0.9	54.5 \pm 2.7	84.9 \pm 0.5	78.3 \pm 2.2	79.7 \pm 2.0	96.7 \pm 0.5	80.9
LoRA-nBAR*	91.4 \pm 0.5	55.0 \pm 2.0	84.9 \pm 1.4	78.1 \pm 0.2	81.0 \pm 1.0	96.7 \pm 1.0	81.2
FMLoRA	95.1 \pm 0.5	54.4 \pm 1.3	86.4 \pm 0.8	82.7 \pm 1.0	82.7 \pm 1.2	96.7 \pm 0.2	83.1
EFMLoRA	91.9 \pm 1.7	54.7 \pm 1.6	85.7 \pm 0.7	82.1 \pm 0.6	82.8 \pm 0.2	96.8 \pm 0.4	<u>82.3</u>

Table 1: Experiments on few-shot RoBERTa (355M). Results marked with * are taken from (Li, Zhang, and He 2024).

RoBERTa	SST2	STS-B	RTE	QQP	QNLI	MRPC	MNLI	CoLA	avg.↑
FT [†]	96.4	92.4	86.6	92.2	94.7	90.9	90.2	68.0	88.9
Adapter [†]	96.6	91.9	80.1	91.7	94.8	89.7	-	67.8	-
LoRA*	95.8	92.4	88.2	91.4	94.7	89.6	90.6	64.8	88.4
LoRA-oBAR*	<u>96.0</u>	<u>92.6</u>	88.7	<u>91.6</u>	94.8	<u>90.3</u>	90.6	65.1	88.7
LoRA-nBAR*	<u>96.0</u>	<u>92.6</u>	89.2	<u>91.6</u>	94.7	<u>90.3</u>	90.8	<u>65.6</u>	88.9
EFMLoRA	<u>96.3</u> \pm 0.2	92.7 \pm 0.1	89.3 \pm 0.6	<u>91.6</u> \pm 0.1	94.8 \pm 0.1	91.5 \pm 0.4	<u>90.7</u> \pm 0.1	68.0 \pm 1.2	89.4

Table 2: Experiments on finetuning RoBERTa (355M). Results marked with † are taken from (Hu et al. 2022), and those with * are taken from (Li, Zhang, and He 2024).

Fine-tuning with RoBERTa-large. We apply EFM-LoRA to finetune RoBERTa-large. Our implementation follows (Hu et al. 2022), using the same hyperparameters as those in its GitHub repository. The results can be found in Table 2. We observe that EFMLoRA achieves the highest scores on all datasets, and achieves the highest accuracy on average over these datasets. Specifically, on average over these datasets, EFMLoRA surpasses standard LoRA with a margin of 1.0. Additionally, EFMLoRA even achieves better performance than full fine-tuning on some datasets. This superior performance may be attributed to overfitting in full fine-tuning, where optimizing all model parameters can lead to overfitting on the training data, thus reducing the model’s generalization to the test set. This effect is particularly pronounced on small datasets, such as MRPC, which contains only 3.7k training data.

Fine-tuning with GPT-2. Having shown that FMLoRA is effective for NLU tasks, we now explore whether EFM-LoRA can improve LoRA in NLG models like GPT-2 Medium and Large (Radford et al. 2019). To enable a direct comparison, we adopt the experimental setup of (Li and Liang 2021) with minimal deviation. Table 3 demonstrates the effectiveness of EFMLoRA on the E2E NLG Challenge (Novikova, Dušek, and Rieser 2017) with GPT-2 Medium and Large models. Compared with existing PEFT methods such as Adapter and LoRA, EFMLoRA consistently achieves superior performance across all metrics. Notably, it achieves this improvement without increasing the number of trainable parameters, maintaining the same efficiency as standard LoRA.

Experiments on Vision Language Models

Few-shot with CLIP. Recent advances in few-shot adaptation of Vision-Language Models (VLMs) have signifi-

cantly enhanced their generalization. CLIP-LoRA (Zanella and Ben Ayed 2024) explores the application of LoRA in this few-shot VLM setting. In our work, we also apply FMLoRA and EFMLoRA to VLMs to evaluate their effectiveness. For a fair comparison, our experimental setup follows that of CLIP-LoRA. We consider five datasets for fine-grained classification of satellite imagery (EuroSAT (Helber et al. 2019), Ox-fordPets (Parkhi et al. 2012), Flower102 (Nilsback and Zisserman 2008), Caltech101 (Fei-Fei, Fergus, and Perona 2004), DTD (Cimpoi et al. 2014)). These datasets offer a thorough benchmarking framework for evaluating few-shot visual classification tasks. Table 4 demonstrates that FMLoRA and EFMLoRA outperformed Adapter and LoRA in most settings. In the low-data regimes (1-shot and 4-shot), EFMLoRA shows clear advantages. These results highlight the effectiveness of EFMLoRA in improving generalization in few-shot adaptation of vision-language models.

Fine-tuning with Qwen-VL-Chat. Qwen-VL-Chat (Bai et al. 2023) is a multimodal conversational large language model capable of understanding both images and text. We apply EFMLoRA to fine-tune Qwen-VL-Chat, following the same experimental setup as in (Zhou et al. 2024). Table 5 presents the results on the ScienceQA (Lu et al. 2022) and VizWiz (Gurari et al. 2018) datasets. The results in Table 5 demonstrate that the perturbation size ρ significantly influences the performance of EFMLoRA when fine-tuning Qwen-VL-Chat. By tuning ρ , EFMLoRA adapts to different tasks, enabling improved generalization—achieving higher accuracy than LoRA. Specifically, a larger ρ (e.g., $\rho = 0.2$) yields the best accuracy on ScienceQA, while a smaller ρ (e.g., $\rho = 0.05$) performs better on VizWiz. This suggests that different tasks benefit from different levels of perturbation. Therefore, selecting an appropriate ρ based on the task

Model & Method	# Trainable Parameters	E2E NLG Challenge				
		BLEU \uparrow	NIST \uparrow	MET \uparrow	ROUGE-L \uparrow	CIDEr \uparrow
GPT-2 M (FT) \dagger	354.92M	68.2	8.62	46.2	71.0	2.47
GPT-2 M (Adapter ^L) \dagger	0.37M	66.3	8.41	45.0	69.8	2.40
GPT-2 M (LoRA)	0.35M	69.2	<u>8.72</u>	<u>46.5</u>	<u>71.5</u>	<u>2.51</u>
GPT-2 M (FMLoRA)	0.35M	<u>69.2</u>	<u>8.72</u>	46.6	<u>71.5</u>	<u>2.51</u>
GPT-2 M (EFMLoRA)	0.35M	69.7	8.77	46.6	71.7	2.53
GPT-2 L (FT) \dagger	774.03M	68.5	8.78	46.0	69.9	2.45
GPT-2 L (Adapter ^L) \dagger	0.88M	69.1	8.68	<u>46.3</u>	71.4	2.49
GPT-2 L (LoRA)	0.77M	69.9	8.82	46.8	71.8	2.53
GPT-2 L (FMLoRA)	0.77M	<u>70.0</u>	<u>8.83</u>	46.8	71.8	2.53
GPT-2 L (EFMLoRA)	0.77M	70.2	8.84	46.8	71.8	<u>2.52</u>

Table 3: GPT-2 medium (M) and large (L) with different adaptation methods on the E2E NLG Challenge. Results marked with \dagger are taken from (Hu et al. 2022).

Shots	Method	Eur. \uparrow	Pets \uparrow	Flo. \uparrow	Cal. \uparrow	DTD \uparrow
0	CLIP	47.5	89.1	71.4	92.9	43.6
1	Adapter	49.3	89.0	71.3	92.0	44.2
	LoRA	72.3	<u>92.3</u>	83.2	93.7	54.3
	FMLoRA	<u>72.6</u>	92.8	<u>82.8</u>	94.5	54.9
	EFMLoRA	78.3	92.8	81.0	<u>93.9</u>	<u>54.6</u>
4	Adapter	51.2	90.8	73.1	94.0	46.1
	LoRA	84.9	91.0	93.7	<u>95.2</u>	63.8
	FMLoRA	90.0	93.1	94.9	95.6	65.7
	EFMLoRA	<u>87.6</u>	<u>91.1</u>	<u>94.0</u>	95.6	<u>65.0</u>
16	Adapter	71.4	92.3	92.9	94.9	59.4
	LoRA	<u>92.1</u>	<u>92.4</u>	98.0	96.4	<u>72.0</u>
	FMLoRA	92.2	93.4	98.5	<u>96.5</u>	72.7
	EFMLoRA	91.6	91.5	<u>98.1</u>	96.6	71.9

Table 4: Detailed results for five datasets with CLIP-Adapter, CLIP-LoRA and EFMLoRA.

characteristics is crucial for achieving optimal fine-tuning performance on multimodal large language models.

Runtime and Memory Consumption

The results in Table 6 confirm the theoretical time complexity analysis. As expected, FMLoRA has approximately double the runtime of LoRA (2.1 \times on both GPT-2 Medium and Large), consistent with its theoretical complexity of $O(2T)$ due to two forward and backward passes for sharpness optimization. In contrast, EFMLoRA operates with near-LoRA efficiency, requiring only 1.1 \times and 1.2 \times more time on GPT-2 Medium and Large, respectively. This supports the theoretical claim that EFMLoRA maintains a time complexity close to $O(T)$ while benefiting from sharpness-aware optimization. In addition, EFMLoRA maintains a memory usage almost identical to that of LoRA, with only negligible increases (less than 0.4 GB across both model scales). These results demonstrate that EFMLoRA achieves near-LoRA efficiency in both memory and runtime.

Method	ρ	SQA \uparrow	VizWiz \uparrow
LoRA	-	<u>90.1</u>	50.69
EFMLoRA	0.05	90.0	51.7
	0.1	90.0	50.6
	0.2	91.6	<u>51.0</u>
	0.6	89.6	50.0

Table 5: EFMLoRA Fine-Tuning Results on Qwen-VL-Chat with different ρ .

Methods	GPT-2 Medium		GPT-2 Large	
	Memory \downarrow	Time \downarrow	Memory \downarrow	Time \downarrow
LoRA	23.6	4.30	23.2	8.45
FMLoRA	<u>24.0</u>	9.10	<u>23.6</u>	17.47
EFMLoRA	<u>24.0</u>	<u>4.80</u>	23.2	<u>10.00</u>

Table 6: Runtime (Hour) and memory (GB) of LoRA, FMLoRA and EFMLoRA on fine-tuning GPT-2 Medium/Large.

Conclusion

In this work, we propose FMLoRA, a novel PEFT method that integrates sharpness-aware optimization into the LoRA framework to promote convergence toward flatter minima. We theoretically demonstrate that perturbations in the full parameter space can be equivalently represented within the low-rank subspace. To improve computational efficiency, we introduce EFMLoRA, which leverages an exponential moving average to approximate perturbations, significantly reducing runtime overhead while maintaining effectiveness. Extensive experiments across various large language and vision-language models demonstrate that EFMLoRA achieves comparable or even superior generalization performance to full fine-tuning and LoRA. Our results emphasize the importance of reducing sharpness to improve generalization in PEFT methods, offering valuable insights and practical tools for future research on the link between sharpness and generalization in LLMs and beyond.

Acknowledgments

This research was supported by National Natural Science Foundation of China under Grant NO. 62576018, the Zhejiang Provincial Natural Science Foundation of China under Grant No. LD24F020007, Beijing Natural Science Foundation L223024 and L244043, “One Thousand Plan” projects in Jiangxi Province Jxsq2023102268, Taiyuan City “Double hundred Research action” 2024TYJB0127. The authors gratefully acknowledge the financial support for this work by the National Key R&D Program of China (Grant Nos. 2019YFF0301401).

References

- Andriushchenko, M.; Croce, F.; Müller, M.; Hein, M.; and Flammarion, N. 2023. A modern look at the relationship between sharpness and generalization. In *Proceedings of the International Conference on Machine Learning (ICML)*, 840–902.
- Bai, J.; Bai, S.; Yang, S.; Wang, S.; Tan, S.; Wang, P.; Lin, J.; Zhou, C.; and Zhou, J. 2023. Qwen-VL: A versatile vision-language model for understanding, localization, text reading, and beyond. In <https://api.semanticscholar.org/CorpusID:261101015>.
- Cimpoi, M.; Maji, S.; Kokkinos, I.; Mohamed, S.; and Vedaldi, A. 2014. Describing textures in the wild. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 3606–3613.
- Dai, R.; Yang, X.; Sun, Y.; Shen, L.; Tian, X.; Wang, M.; and Zhang, Y. 2023. FedGAMMA: Federated learning with global sharpness-aware minimization. *IEEE Transactions on Neural Networks and Learning Systems*, 1–14. Doi:10.1109/TNNLS.2023.3304453.
- Deng, J.; Pang, J.; Zhang, B.; and Wang, T. 2024. Effective gradient sample size via variation estimation for accelerating sharpness aware minimization. *arXiv preprint arXiv:2403.08821*.
- Dinh, L.; Pascanu, R.; Bengio, S.; and Bengio, Y. 2017a. Sharp minima can generalize for deep nets. In *Proceedings of the International Conference on Machine Learning (ICML)*, 1019–1028. PMLR.
- Dinh, L.; Pascanu, R.; Bengio, S.; and Bengio, Y. 2017b. Sharp minima can generalize for deep nets. In *Proceedings of the International Conference on Machine Learning (ICML)*, 1019–1028. PMLR.
- Dong, M.; Yang, Y.; Zeng, K.; Wang, Q.; and Shen, T. 2024. Implicit sharpness-aware minimization for domain generalization. *Remote Sensing*, 16(16): 2877.
- Du, J.; Yan, H.; Feng, J.; Zhou, J. T.; Zhen, L.; Goh, R. S. M.; and Tan, V. Y. 2022. Efficient sharpness-aware minimization for improved training of neural networks.
- Du, S. S.; Hu, W.; and Lee, J. D. 2018. Algorithmic regularization in learning deep homogeneous models: Layers are automatically balanced. *Advances in neural information processing systems*, 31.
- Fei-Fei, L.; Fergus, R.; and Perona, P. 2004. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 178–178. IEEE.
- Foret, P.; Kleiner, A.; Mobahi, H.; and Neyshabur, B. 2021. Sharpness-aware minimization for efficiently improving generalization. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Ge, R.; Jin, C.; and Zheng, Y. 2017. No spurious local minima in nonconvex low rank problems: A unified geometric analysis. In *Proceedings of the International Conference on Machine Learning (ICML)*, 1233–1242. PMLR.
- Gurari, D.; Li, Q.; Stangl, A. J.; Guo, A.; Lin, C.; Grauman, K.; Luo, J.; and Bigham, J. P. 2018. Vizwiz grand challenge: Answering visual questions from blind people. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 3608–3617.
- Hayou, S.; Ghosh, N.; and Yu, B. 2024. Lora+: Efficient low rank adaptation of large models. *arXiv preprint arXiv:2402.12354*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 1026–1034.
- Helber, P.; Bischke, B.; Dengel, A.; and Borth, D. 2019. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7): 2217–2226.
- Hochreiter, S.; and Schmidhuber, J. 1994. Simplifying neural nets by discovering flat minima. *Advances in neural information processing systems*, 7.
- Hochreiter, S.; and Schmidhuber, J. 1997. Flat minima. *Neural computation*, 9(1): 1–42.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Huang, C.; Liu, Q.; Lin, B. Y.; Pang, T.; Du, C.; and Lin, M. 2023. Lorahub: Efficient cross-task generalization via dynamic lora composition. *arXiv preprint arXiv:2307.13269*.
- Jang, U.; Lee, J. D.; and Ryu, E. K. 2024. LoRA training in the NTK regime has no spurious local minima. In *Proceedings of the International Conference on Machine Learning (ICML)*, 21306–21328.
- Jastrzebski, S.; Kenton, Z.; Arpit, D.; Ballas, N.; Fischer, A.; Bengio, Y.; and Storkey, A. 2017. Three factors influencing minima in sgd. *arXiv preprint arXiv:1711.04623*.
- Jiang, Y.; Neyshabur, B.; Mobahi, H.; Krishnan, D.; and Bengio, S. 2020. Fantastic generalization measures and where to find them.
- Keskar, N. S.; Mudigere, D.; Nocedal, J.; Smelyanskiy, M.; and Tang, P. T. P. 2017. On large-batch training for deep learning: Generalization gap and sharp minima. *Proceedings of the International Conference on Learning Representations (ICLR)*.

- Kim, H.; Park, J.; Choi, Y.; Lee, W.; and Lee, J. 2023. Exploring the effect of multi-step ascent in sharpness-aware minimization. *arXiv preprint arXiv:2302.10181*.
- Kwon, J.; Kim, J.; Park, H.; and Choi, I. K. 2021. Asam: Adaptive sharpness-aware minimization for scale-invariant learning of deep neural networks. In *Proceedings of the International Conference on Machine Learning (ICML)*, 5905–5914. PMLR.
- Li, B.; Zhang, L.; and He, N. 2024. Implicit regularization of sharpness-aware minimization for scale-invariant problems. *Advances in Neural Information Processing Systems*, 37: 44444–44478.
- Li, M.; Chen, H.; Wang, Y.; Zhu, T.; Zhang, W.; Zhu, K.; Wong, K.-F.; and Wang, J. 2025. Understanding and mitigating the bias inheritance in LLM-based data augmentation on downstream tasks. *arXiv preprint arXiv:2502.04419*.
- Li, T.; He, Z.; Li, Y.; Wang, Y.; Shang, L.; and Huang, X. 2024a. Flat-LoRA: Low-rank adaptation over a flat loss landscape. *arXiv preprint arXiv:2409.14396*.
- Li, T.; Tan, L.; Huang, Z.; Tao, Q.; Liu, Y.; and Huang, X. 2022. Low dimensional trajectory hypothesis is true: Dnns can be trained in tiny subspaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3): 3411–3420.
- Li, T.; Zhou, P.; He, Z.; Cheng, X.; and Huang, X. 2024b. Friendly sharpness-aware minimization. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 5631–5640.
- 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*. Association for Computational Linguistics.
- 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 *Proceedings of the International Conference on Machine Learning (ICML)*.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Lu, P.; Mishra, S.; Xia, T.; Qiu, L.; Chang, K.-W.; Zhu, S.-C.; Tafjord, O.; Clark, P.; and Kalyan, A. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35: 2507–2521.
- Meng, F.; Wang, Z.; and Zhang, M. 2024. Pissa: Principal singular values and singular vectors adaptation of large language models. *Advances in Neural Information Processing Systems*, 37: 121038–121072.
- Neyshabur, B.; Bhojanapalli, S.; McAllester, D.; and Srebro, N. 2017. Exploring generalization in deep learning. *Advances in Neural Information Processing Systems (NeurIPS)*, 30.
- Neyshabur, B.; Salakhutdinov, R. R.; and Srebro, N. 2015. Path-sgd: Path-normalized optimization in deep neural networks. *Advances in neural information processing systems*, 28.
- Nilsback, M.-E.; and Zisserman, A. 2008. Automated flower classification over a large number of classes. In *2008 Sixth Indian conference on computer vision, graphics & image processing*, 722–729. IEEE.
- Novikova, J.; Dušek, O.; and Rieser, V. 2017. The E2E dataset: New challenges for end-to-end generation. *arXiv preprint arXiv:1706.09254*.
- Parkhi, O. M.; Vedaldi, A.; Zisserman, A.; and Jawahar, C. 2012. Cats and dogs. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 3498–3505. IEEE.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Tian, C.; Shi, Z.; Guo, Z.; Li, L.; and Xu, C.-Z. 2024. Hydralora: An asymmetric lora architecture for efficient fine-tuning. *Advances in Neural Information Processing Systems*, 37: 9565–9584.
- Wang, Z.; and Mao, Y. 2022. On the generalization of models trained with SGD: Information-theoretic bounds and implications. In *International Conference on Learning Representations*.
- Zanella, M.; and Ben Ayed, I. 2024. Low-rank few-shot adaptation of vision-language models. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 1593–1603.
- Zhang, L.; Zhang, L.; Shi, S.; Chu, X.; and Li, B. 2023a. Lora-fa: Memory-efficient low-rank adaptation for large language models fine-tuning. *arXiv preprint arXiv:2308.03303*.
- Zhang, Q.; Chen, M.; Bukharin, A.; He, P.; Cheng, Y.; Chen, W.; and Zhao, T. 2023b. Adaptive budget allocation for parameter-efficient fine-tuning. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Zhang, X.; Xu, R.; Yu, H.; Zou, H.; and Cui, P. 2023c. Gradient norm aware minimization seeks first-order flatness and improves generalization. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 20247–20257.
- Zhao, J.; Zhang, Z.; Chen, B.; Wang, Z.; Anandkumar, A.; and Tian, Y. 2024. Galore: Memory-efficient llm training by gradient low-rank projection. *arXiv preprint arXiv:2403.03507*.
- Zhou, X.; He, J.; Ke, Y.; Zhu, G.; Gutiérrez-Basulto, V.; and Pan, J. Z. 2024. An empirical study on parameter-efficient fine-tuning for multimodal large language models. *arXiv preprint arXiv:2406.05130*.
- Zhuang, J.; Gong, B.; Yuan, L.; Cui, Y.; Adam, H.; Dvornik, N.; Tatikonda, S.; Duncan, J.; and Liu, T. 2022. Surrogate gap minimization improves sharpness-aware training.