

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

6-2024

Friendly sharpness-aware minimization

Tao LI

Pan ZHOU

Singapore Management University, panzhou@smu.edu.sg

Zhengbao HE

Xinwen CHENG

Xiaolin HUANG

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Theory and Algorithms Commons](#)

Citation

LI, Tao; ZHOU, Pan; HE, Zhengbao; CHENG, Xinwen; and HUANG, Xiaolin. Friendly sharpness-aware minimization. (2024). *Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition Conference (CVPR), Seattle, 2024 June 17-21*. 5631-5640.

Available at: https://ink.library.smu.edu.sg/sis_research/9018

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

Friendly Sharpness-Aware Minimization

Tao Li¹ Pan Zhou² Zhengbao He¹ Xinwen Cheng¹ Xiaolin Huang¹

¹Department of Automation, Shanghai Jiao Tong University, Shanghai, China

²Singapore Management University, Singapore

{li.tao, lstefanie, xinwencheng, xiaolinhuang}@sjtu.edu.cn panzhou@smu.edu.sg

Abstract

Sharpness-Aware Minimization (SAM) has been instrumental in improving deep neural network training by minimizing both training loss and loss sharpness. Despite the practical success, the mechanisms behind SAM’s generalization enhancements remain elusive, limiting its progress in deep learning optimization. In this work, we investigate SAM’s core components for generalization improvement and introduce “Friendly-SAM” (F-SAM) to further enhance SAM’s generalization. Our investigation reveals the key role of batch-specific stochastic gradient noise within the adversarial perturbation, i.e., the current minibatch gradient, which significantly influences SAM’s generalization performance. By decomposing the adversarial perturbation in SAM into full gradient and stochastic gradient noise components, we discover that relying solely on the full gradient component degrades generalization while excluding it leads to improved performance. The possible reason lies in the full gradient component’s increase in sharpness loss for the entire dataset, creating inconsistencies with the subsequent sharpness minimization step solely on the current minibatch data. Inspired by these insights, F-SAM aims to mitigate the negative effects of the full gradient. It removes the full gradient estimated by an exponentially moving average (EMA) of historical stochastic gradients, and then leverages stochastic gradient noise for improved generalization. Moreover, we provide theoretical validation for the EMA approximation and prove the convergence of F-SAM on non-convex problems. Extensive experiments demonstrate the superior generalization performance and robustness of F-SAM over vanilla SAM. Code is available at <https://github.com/nblt/F-SAM>.

1. Introduction

Deep neural networks (DNNs) have achieved remarkable performance in various vision and language processing tasks [1–3]. A critical factor contributing to their success is the choice of optimization algorithms [4–9], designed to

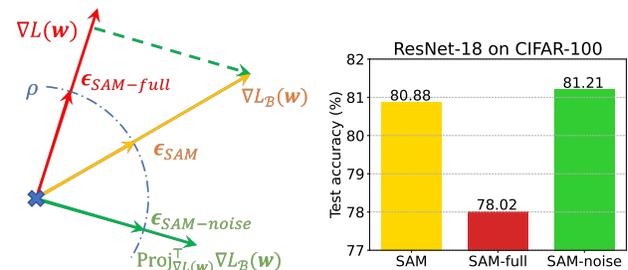


Figure 1. Investigation on SAM’s adversarial perturbation direction. We decompose the minibatch gradient $\nabla_{L_B}(\mathbf{w})$ into two components: the full gradient component and the remaining batch-specific stochastic gradient noise. Solely using the full gradient component leads to a dramatic generalization degradation, while only using the noise component enhances the generalization.

efficiently optimize DNN model parameters. To achieve better performance, it is often desirable for an optimizer to converge to flat minima, characterized by uniformly low loss values, as opposed to sharp minima with abrupt loss changes, as the former typically results in better generalization [10–13]. However, the complex landscape of over-parameterized DNNs, featuring numerous sharp minima, presents challenges for optimizers in real applications.

To address this challenge, Sharpness-Aware Minimization (SAM) [14] is proposed to simultaneously minimize the training loss and the loss sharpness, enabling it to identify flat minima associated with improved generalization performance. For each training iteration, given the minibatch training loss $L_B(\mathbf{w})$ parameterized by \mathbf{w} , SAM achieves this by adversarially computing an adversarial perturbation ϵ_s to maximize the training loss $L_B(\mathbf{w} + \epsilon_s)$. It then minimizes the loss of this perturbed objective via one updating step of a base optimizer, such as SGD [4]. In practice, to efficiently compute ϵ_s , SAM takes a linear approximation of the loss objective, and employs the minibatch gradient $\nabla L_B(\mathbf{w})$ as the accent direction for searching $\epsilon_s = \rho \cdot \text{norm}(\nabla L_B(\mathbf{w}))$, where $\text{norm}(\mathbf{x}) = \mathbf{x}/\|\mathbf{x}\|_2$ and $\rho > 0$ is a radius. This approach enables SAM to seek models situated in neighborhoods characterized by consis-

tently low loss values, resulting in better generalization when used in conjunction with base optimizers such as SGD and Adam [5].

Despite SAM’s practical success, the underlying mechanism responsible for its empirical generalization improvements remains limited [15, 16]. This open challenge hinders developing new and more advanced deep learning optimization algorithms in a principle way. To address this issue, we undertake a comprehensive exploration of SAM’s core components contributing to its generalization improvement. Subsequently, we introduce a new variant of SAM, which offers a simple yet effective approach to further enhance the generalization performance in conjunction with popular base optimizers.

First, our investigations reveal that the batch-specific stochastic gradient noise present in the minibatch gradient $\nabla L_B(\mathbf{w})$ plays a crucial role in SAM’s generalization performance. Specifically, as shown in Fig. 1, we decompose the perturbation direction $\nabla L_B(\mathbf{w})$ into two orthogonal components: the full gradient component $\text{Proj}_{\nabla L(\mathbf{w})} \nabla L_B(\mathbf{w})$, representing the projection of $\nabla L_B(\mathbf{w})$ onto the full gradient direction $\nabla L(\mathbf{w})$, and the stochastic gradient noise component $\text{Proj}_{\nabla L(\mathbf{w})}^\perp \nabla L_B(\mathbf{w})$, denoting the residual projection. Surprisingly, we empirically find that as shown in Fig. 1, using only the full gradient component for perturbation significantly degrades SAM’s generalization, which is also observed by [15]. Conversely, excluding the full gradient component leads to improved generalization performance. This observation suggests that the effectiveness of SAM primarily arises from the presence of the stochastic gradient noise component within the minibatch gradient $\nabla L_B(\mathbf{w})$. The reason behind this phenomenon may be that the full gradient component in the perturbation increases the sharpness loss of the entire dataset, leading to inconsistencies with the subsequent sharpness minimization step which only uses the current minibatch data to minimize training loss value and its sharpness. See more discussion in Sec. 4.

Next, building on this insight, we introduce a straightforward yet highly effective modification to SAM, referred to as “Friendly-SAM” or F-SAM. F-SAM mitigates the undesirable effects of the full gradient component within the minibatch gradient and leverages the beneficial role of batch-specific stochastic gradient noise to further enhance SAM’s generalization performance. Since removing the full gradient component in the perturbation step can relieve the adversarial perturbation impact on other data points except for current minibatch data, the adversarial perturbation in F-SAM is more “friendly” to other data points compared with vanilla SAM, which helps the adversarial perturbation and the following minimization perform on the same minibatch data and improves sharpness minimization consistency. Moreover, such friendliness enables F-SAM to be

less sensitive to the choice of perturbation radius as discussed in Sec. 6.3. For efficiency, F-SAM approximates the computationally expensive full gradient through an exponentially moving average (EMA) of the historical stochastic gradients and computes the stochastic gradient noise as the perturbation direction. Moreover, we also provide theoretical validation of the EMA approximation to the full gradient and prove the convergence of F-SAM on non-convex problems. Our main contribution can be summarized as follows:

- We take an in-depth investigation into the key component in SAM’s adversarial perturbation and identify that the stochastic gradient noise plays the most significant role.
- We propose a novel variant of SAM called F-SAM by eliminating the undesired full gradient component and harnessing the beneficial stochastic gradient noise for adversarial perturbation to enhance generalization.
- Extensive experimental results demonstrate that F-SAM improves SAM’s generalization, aids in training, and exhibits better robustness across different perturbation radii.

2. Related Work

Sharp Minima and Generalization. The connection between the flatness of local minima and generalization has received a rich body of studies [10–13, 17–19]. Recently, many works have tried to improve the model generalization by seeking flat minima [14, 17, 20–25]. For example, Chaudhari et al. [17] proposed Entropy-SGD that actively searches for flat regions by minimizing local entropy. Wen et al. [20] designed SmoothOut framework to smooth out the sharp minima and obtain generalization improvements. Notably, sharpness-aware minimization (SAM) [14] provides a generic training scheme for seeking flat minima by formulating the optimization as a min-max problem and encourage parameters sitting in neighborhoods with uniformly low loss, achieving state-of-the-art generalization improvements across various tasks. Later, a line of works improves SAM from the perspective of the neighborhood’s geometric measure [26, 27], surrogate loss function [28], friendly adversary [29], and training efficiency [30–34].

Understanding SAM. Despite the empirical success of SAM across various tasks, a deep understanding of its generalization performance is still limited. Foret et al. [14] explained the success of SAM via using a PAC-Bayes generalization bound. However, they did not analyze the key components of SAM that contributes to its generalization improvement. Andriushchenko et al. [15] investigated the implicit bias of SAM for diagonal linear networks. Wen et al. [35] demonstrated that SAM minimizes the top eigenvalue of Hessian in the full-batch setting and thus improves the flatness of the minima. Chen et al. [16] studied SAM on the non-smooth convolutional ReLU networks, and explained its success because of its ability to prevent noise

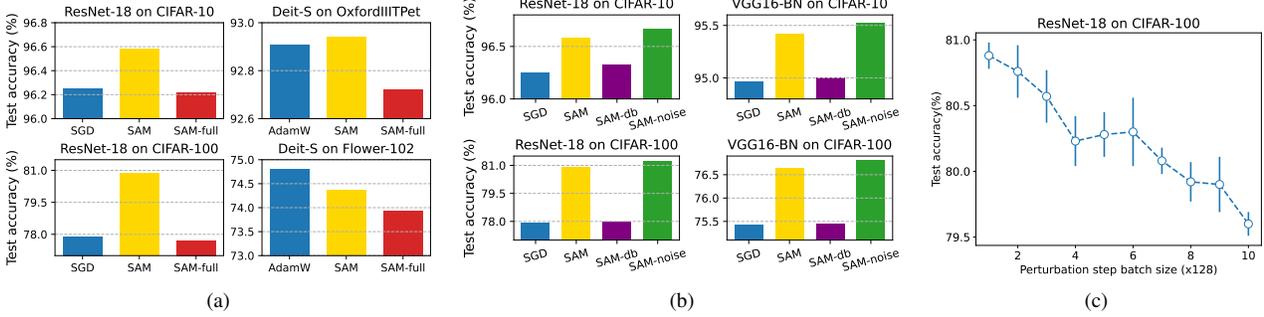


Figure 2. Performance comparison of different versions of SAM with SGD/AdamW as its base optimizer. In (a), SAM-full denotes the version of SAM using the full gradient component as the perturbation. In (b), SAM-db represents the SAM using different random-selected data batch for perturbation and its following minimization step. (c) compares SAM using different minibatch size to compute the perturbation but always fixing minibatch size of 128 for the following minimization step.

learning. While previous works primarily focus on simplified networks or ideal loss objectives, we aim to deepen the understanding of SAM by undertaking a deep investigation into the key components that contribute to its practical generalization improvement.

3. Preliminaries

SAM. Let $f(\mathbf{x}; \mathbf{w})$ parameterized by $\mathbf{w} \in \mathbb{R}^d$ be a neural network, and $L(f(\mathbf{x}_i; \mathbf{w}), \mathbf{y}_i)$ ($L_i(\mathbf{w})$ for short) denote the loss to measure the discrepancy between the prediction $f(\mathbf{x}_i; \mathbf{w})$ and the ground-truth label \mathbf{y}_i . Given a dataset $\mathcal{S} = (\mathbf{x}_i, \mathbf{y}_i)_{i=1}^n$ i.i.d. drawn from a data distribution \mathcal{D} , the empirical training loss is often defined as

$$L(\mathbf{w}) = \frac{1}{n} \sum_{i=1}^n L_i(\mathbf{w}). \quad (1)$$

To solve this problem, one often uses SGD or Adam to optimize it. To further improve the generalization performance, SAM [14] aims to minimize the worst-case loss within a defined neighborhood for guiding the training process towards flat minima. The objective function of SAM is given by:

$$L^{\text{SAM}}(\mathbf{w}) = \max_{\|\epsilon\|_2 \leq \rho} L(\mathbf{w} + \epsilon), \quad (2)$$

where ρ denotes the neighborhood radius. Eqn. (2) shows that SAM seeks to minimize the maximum loss over the neighborhood surrounding the current weight \mathbf{w} . By doing so, SAM encourages the optimization process to converge to flat minima which often enjoys better generalization performance [10–13].

To efficiently optimize $L^{\text{SAM}}(\mathbf{w})$, SAM first approximates Eqn. (2) via first-order expansion and compute the adversarial perturbation ϵ_s as follows:

$$\epsilon_s \approx \arg \max_{\|\epsilon\|_2 \leq \rho} \epsilon^\top \nabla_{\mathbf{w}} L(\mathbf{w}) = \rho \cdot \text{norm}(\nabla_{\mathbf{w}} L(\mathbf{w})). \quad (3)$$

where $\text{norm}(\mathbf{x}) = \mathbf{x} / \|\mathbf{x}\|_2$. Subsequently, one can compute the gradient at the perturbed point $\mathbf{w} + \epsilon_s$, and then use the updating step of the base SGD optimizer to update

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \gamma \nabla L_{\mathcal{B}}(\mathbf{w})|_{\mathbf{w}_t + \epsilon_t}. \quad (4)$$

Other base optimizers, e.g., Adam, can also be used to update the model parameters in Eqn. (4).

Assumptions. Before delving into our analysis, we first make some standard assumptions in stochastic optimization [15, 34, 36, 37] that will be used in our theoretical analysis.

Assumption 1 (β -smoothness). Assume the loss function $L : \mathbb{R}^d \mapsto \mathbb{R}$ to be β -smooth. There exists $\beta > 0$ such that

$$\|\nabla L(\mathbf{w}) - \nabla L(\mathbf{v})\|_2 \leq \beta \|\mathbf{w} - \mathbf{v}\|_2, \quad \forall \mathbf{w}, \mathbf{v} \in \mathbb{R}^d. \quad (5)$$

Assumption 2 (Bounded variance). There exists a constant $M > 0$ for any data batch \mathcal{B} such that

$$\mathbb{E} [\|\nabla L_{\mathcal{B}}(\mathbf{w}) - \nabla L(\mathbf{w})\|_2^2] \leq M, \quad \forall \mathbf{w} \in \mathbb{R}^d. \quad (6)$$

Assumption 3 (Bounded gradient). There exists $G > 0$ for any data batch \mathcal{B} such that

$$\mathbb{E} [\|\nabla L_{\mathcal{B}}(\mathbf{w})\|_2] \leq G, \quad \forall \mathbf{w} \in \mathbb{R}^d. \quad (7)$$

4. Empirical Analysis of SAM

Here we conduct a set of experiments to identify the effective components in the adversarial perturbation that contributes to SAM’s improved generalization performance. To this end, we decompose SAM’s perturbation direction $\nabla L_{\mathcal{B}}(\mathbf{w})$ in Eqn. (3) into two orthogonal components. The first one is the full gradient component $\text{Proj}_{\nabla L(\mathbf{w})} \nabla L_{\mathcal{B}}(\mathbf{w})$ which is the projection onto the full gradient direction, i.e.,

$$\text{Proj}_{\nabla L(\mathbf{w})} \nabla L_{\mathcal{B}}(\mathbf{w}) = \cos(\nabla L(\mathbf{w}), \nabla L_{\mathcal{B}}(\mathbf{w})) \nabla L(\mathbf{w}),$$

where $\cos(\cdot, \cdot)$ denotes the cosine similarity function. Another one is the residual projection, i.e., the batch-specific stochastic gradient noise component which is defined as

$$\text{Proj}_{\nabla L(\mathbf{w})}^{\perp} \nabla L_{\mathcal{B}}(\mathbf{w}) = \nabla L_{\mathcal{B}}(\mathbf{w}) - \text{Proj}_{\nabla L(\mathbf{w})} \nabla L_{\mathcal{B}}(\mathbf{w}). \quad (8)$$

In the following, we will investigate the effects of these two orthogonal components to the performance of SAM.

4.1. Effect of Full Gradient Component

Here we investigate the impact of the full gradient component on the generalization performance of SAM. Specifically, for each training iteration, we set the perturbation ϵ_s as $\epsilon_s = \rho \cdot \text{norm}(\text{Proj}_{\nabla L(\mathbf{w})} \nabla L_{\mathcal{B}}(\mathbf{w}))$ where $\text{norm}(\mathbf{x}) = \mathbf{x} / \|\mathbf{x}\|_2$, and then follow the minimization step in vanilla SAM to minimize the perturbed training loss. For brevity, this modified SAM version is called ‘‘SAM-full’’ since it uses the full gradient direction as the perturbation direction. Then we follow the standard training settings of SAM and compare the performance of SGD/AdamW, SAM, and SAM-full in training from scratch and transfer learning tasks. See more training details in Appendix A5.

Fig. 2a (a) summarizes the results. One can observe that SAM significantly outperforms the base optimizer in most scenarios, except for the Flower-102 dataset [38], where the training data are very limited (see Sec. 6.2 for more discussion). However, as shown in Fig. 2a (a), SAM-full indeed impairs the performance of base optimizers, such as SGD base optimizer on the CIFAR-10, CIFAR-100 [39], and OxfordIIITPet datasets [40], and even severely hurt the performance of AdamW base optimizer on the Flower-102 dataset, of which the training sample number is small. This implies that 1) utilizing the full gradient component $\text{Proj}_{\nabla L(\mathbf{w})} \nabla L_{\mathcal{B}}(\mathbf{w})$ as the perturbation direction actually does not improve generalization and can even have negative effects; 2) the effectiveness of SAM primarily arises from the presence of stochastic gradient noise component $\text{Proj}_{\nabla L(\mathbf{w})}^{\perp} \nabla L_{\mathcal{B}}(\mathbf{w})$ within the minibatch gradient $\nabla L_{\mathcal{B}}(\mathbf{w})$ in adversarial perturbation.

To provide more evidence, we further modify the perturbation step in SAM by increasing the minibatch size in Eqn. (3) to compute the adversarial perturbation. As illustrated in Fig. 2c, we use a minibatch \mathcal{B}' with size $k \times 128$ to compute the perturbation in SAM by following Eqn. (3), where k is selected from $\{1, 2, 3, \dots, 10\}$. However, we maintain the use of a minibatch \mathcal{B} of size 128 for the second minimization step of standard SAM in Eqn. (4). Here we ensure the sample set \mathcal{B} in the second minimization belong to the sample set \mathcal{B}' in the first perturbation step, namely, $\mathcal{B} \subseteq \mathcal{B}'$, which allows to better observe the effects of large minibatch to the performance of SAM. For the perturbation $\epsilon_s = \rho \cdot \text{norm} \nabla L_{\mathcal{B}'}(\mathbf{w})$, its direction gradually aligns more closely with the full gradient direction along with increment

of k . Consequently, as the minibatch size grows, the contribution of the full gradient component in the perturbation step becomes stronger, while the residual stochastic gradient noise component weakens. This analysis helps to reveal the perturbation effects of the full gradient direction to SAM’s performance.

By observing Fig. 2c, one can find that when minibatch size of \mathcal{B}' in the first adversarial perturbation step increases, the classification accuracy of SAM gradually decreases when using ResNet18 on CIFAR-100 dataset. See more similar experimental results in Appendix A5. All these results together indicate that strengthening the full gradient component in the perturbation step indeed impairs the generalization performance of vanilla SAM.

4.2. Effect of Stochastic Gradient Noise Component

To delve into the role of the batch-specific stochastic gradient noise in the generalization improvement of SAM, we first replace the stochastic gradient noise component of the current minibatch data with another random minibatch. That is, for each training iteration, we use two different minibatch data of the same size to do the first adversarial perturbation step (3) and the second minimization step (4). We term this modified version ‘‘SAM-db’’. As illustrated in Fig. 2b, we observe that for the four training settings, SAM-db with SGD as its base optimizer often does not make significant improvements in generalization compared to SGD. Note that the adversarial perturbation in SAM-db contains a similar full gradient component as vanilla SAM (in terms of expectation), but the key difference lies in replacing the stochastic gradient noise component from the current minibatch data with another random minibatch. However, this substitution ultimately hinders generalization. Therefore, we conclude that the stochastic gradient noise associated with the minibatch in decent step in perturbation plays a pivotal role in improving the generalization of SAM.

To further solidify this observation, we modify SAM by setting the adversarial perturbation ϵ_s in vanilla SAM as $\epsilon_s = \rho \cdot \text{norm}(\text{Proj}_{\nabla L(\mathbf{w})}^{\perp} \nabla L_{\mathcal{B}}(\mathbf{w}))$. Note, in this modification, we compute the full gradient for each iteration and then follow Eqn. (8) to calculate the stochastic gradient noise $\text{Proj}_{\nabla L(\mathbf{w})}^{\perp} \nabla L_{\mathcal{B}}(\mathbf{w})$. We refer to this modified version as ‘‘SAM-noise’’. As illustrated in Fig. 2b, one can observe that SAM-noise not only restores the generalization performance of SAM but even exhibits a notable enhancement. This result highlights that exclusively utilizing the stochastic gradient noise component as SAM’s perturbation can further enhance its generalization ability, thereby inspiring our F-SAM algorithm in the next section.

5. The F-SAM Algorithm

In this section, we will present our proposed F-SAM algorithm. We begin by introducing the efficient estimation of

the stochastic gradient noise component. We then describe the algorithmic steps of our proposed F-SAM. Finally, we formulate the loss objective of our F-SAM algorithm and highlight the distinctions from vanilla SAM.

5.1. Estimation of Stochastic Gradient Noise Component

From the empirical results in Sec. 4, we know that the full gradient component $\text{Proj}_{\nabla L(\mathbf{w})} \nabla L_{\mathcal{B}}(\mathbf{w})$ in the perturbation impairs the performance of SAM, while the batch-specific stochastic gradient noise component $\text{Proj}_{\nabla L(\mathbf{w})}^{\top} \nabla L_{\mathcal{B}}(\mathbf{w})$ is essential to improve performance. Accordingly, it is natural to remove the full gradient component and only use the stochastic gradient noise component.

To this end, we first reformulate the batch-specific stochastic gradient noise component in Eqn. (8) as

$$\text{Proj}_{\nabla L(\mathbf{w})}^{\top} \nabla L_{\mathcal{B}}(\mathbf{w}) = \nabla L_{\mathcal{B}}(\mathbf{w}) - \sigma \nabla L(\mathbf{w}). \quad (9)$$

where $\sigma = \cos(\nabla L(\mathbf{w}), \nabla L_{\mathcal{B}}(\mathbf{w}))$. From Eqn. (9), one can observe that to compute $\text{Proj}_{\nabla L(\mathbf{w})}^{\top} \nabla L_{\mathcal{B}}(\mathbf{w})$, we need to compute the minibatch gradient $\nabla L_{\mathcal{B}}(\mathbf{w})$ and also the full gradient $\nabla L(\mathbf{w})$ at each training iteration. For $\nabla L_{\mathcal{B}}(\mathbf{w})$, it is also computed in vanilla SAM and thus does not bring extra computation overhead. Regarding the full gradient $\nabla L(\mathbf{w})$, it is computed on the whole dataset, and thus is computationally prohibitive in practice.

To address this issue, we resort to estimating $\nabla L(\mathbf{w})$ with an exponentially moving average (EMA) which accumulates the historical minibatch gradients as follows:

$$\mathbf{m}_t = \lambda \mathbf{m}_{t-1} + (1 - \lambda) \nabla L_{\mathcal{B}_t}(\mathbf{w}_t), \quad (10)$$

where $\lambda > 0$ is a hyper-parameter. This approach enables us to approximate the full gradient $\nabla L(\mathbf{w}_t)$ with minimal additional computational overhead. We also prove that \mathbf{m}_t is a good estimation to the full gradient $\nabla L(\mathbf{w}_t)$ as shown in following Theorem 1 with its proof in Appendix A1.

Theorem 1. *Suppose Assumption 1, 2, and 3 hold. Assume that SAM uses SGD as the base optimizer with a learning rate γ to update the model parameter in Eqn. (4). Then by setting $\lambda = 1 - C\gamma^{2/3}$, after $T > C'\gamma^{-2/3}$ training iterations, with probability $1 - \delta$, we have*

$$\Phi_T = \|\mathbf{m}_T - \nabla L(\mathbf{w}_T)\|_2 \leq \mathcal{O}\left(\gamma^{\frac{1}{3}} \beta^{\frac{1}{3}} G^{\frac{1}{3}} M^{\frac{1}{3}} \log\left(\frac{1}{\delta}\right)\right),$$

where C and C' are two universal constants.

Theorem 1 shows that the error bound Φ_t between the full gradient $\nabla L(\mathbf{w}_t)$ and its EMA estimation \mathbf{m}_t is at the order of $\mathcal{O}(\gamma^{1/3})$. On the non-convex problem, learning rate γ is often set as $\mathcal{O}(1/\sqrt{T})$ to ensure convergence as shown in our Theorem 2 in Section 5.4. This implies that the error bound Φ_t is $\mathcal{O}(T^{-1/6})$ and thus is small since the training

iteration T is often large. So \mathbf{m}_t is a good estimation the full gradient $\nabla L(\mathbf{w}_t)$.

In this way, the stochastic gradient noise component in Eqn. (8) is approximated as

$$\text{Proj}_{\nabla L(\mathbf{w}_t)}^{\top} \nabla L_{\mathcal{B}}(\mathbf{w}_t) \approx \widetilde{\text{Proj}}_{\nabla L(\mathbf{w}_t)}^{\top} \nabla L_{\mathcal{B}}(\mathbf{w}_t), \quad (11)$$

where $\widetilde{\text{Proj}}_{\nabla L(\mathbf{w}_t)}^{\top} \nabla L_{\mathcal{B}}(\mathbf{w}_t) = \nabla L_{\mathcal{B}}(\mathbf{w}_t) - \sigma_t \mathbf{m}_t$ and $\sigma_t = \cos(\mathbf{m}_t, \nabla L_{\mathcal{B}}(\mathbf{w}_t))$. In the real network training, the training frameworks, e.g., PyTorch, often uses propagation to update \mathbf{m}_t and $\nabla L_{\mathcal{B}}(\mathbf{w}_t)$ from layer to layer. So one needs to wait until the training framework finishes the updating of all layers, and then compute σ_t which is not efficient, especially for modern over-parameterized networks. Moreover, both \mathbf{m}_t and $\nabla L_{\mathcal{B}}(\mathbf{w}_t)$ are high dimensional, their cosine computation can be unstable due to the possible noises in the approximation \mathbf{m}_t . So to improve the training efficiency and stability, we set σ_t as a constant σ , and obtain

$$\widetilde{\text{Proj}}_{\nabla L(\mathbf{w}_t)}^{\top} \nabla L_{\mathcal{B}}(\mathbf{w}_t) = \nabla L_{\mathcal{B}}(\mathbf{w}_t) - \sigma \mathbf{m}_t. \quad (12)$$

The experimental results in Section 6 show that this practical setting achieves satisfactory performance. In theory, we also prove that this practical setting is validated as shown in the following Lemma 1 with its proof in Appendix A2.

Lemma 1. *Assuming \mathbf{m}_t is a unbiased estimator to the full gradient $\nabla L(\mathbf{w}_t)$. With $\sigma = 1$, we have*

$$\mathbb{E} \left\langle \widetilde{\text{Proj}}_{\nabla L(\mathbf{w}_t)}^{\top} \nabla L_{\mathcal{B}}(\mathbf{w}_t), \nabla L(\mathbf{w}_t) \right\rangle = 0. \quad (13)$$

Lemma 1 shows that $\widetilde{\text{Proj}}_{\nabla L(\mathbf{w}_t)}^{\top} \nabla L_{\mathcal{B}}(\mathbf{w}_t)$ is orthogonal to the full gradient $\nabla L(\mathbf{w}_t)$ in expectation when $\sigma = 1$. So using $\widetilde{\text{Proj}}_{\nabla L(\mathbf{w}_t)}^{\top} \nabla L_{\mathcal{B}}(\mathbf{w}_t)$ as the perturbation direction can remove the effects of the full gradient $\nabla L(\mathbf{w}_t)$. As Theorem 1 shows that \mathbf{m}_t is good estimation to full gradient $\nabla L(\mathbf{w}_t)$, the assumption in Lemma 1 is not restrictive.

5.2. Algorithmic Steps of F-SAM

After estimating the stochastic gradient noise component in Eqn. (12), one can use it to replace the adversarial perturbation direction in vanilla SAM, and straightforwardly develop our proposed F-SAM. Specifically, for the t -th training iteration, following the algorithmic steps in vanilla SAM, we first define the perturbation ϵ as follows:

$$\epsilon_t = \rho \cdot \text{norm}\left(\widetilde{\text{Proj}}_{\nabla L(\mathbf{w}_t)}^{\top} \nabla L_{\mathcal{B}}(\mathbf{w}_t)\right), \quad (14)$$

where ρ is the radius. Then, following SAM, we use SGD as its base optimizer, and updates the model parameter \mathbf{w}_t by using Eqn. (4) in Section 3. We summarize the algorithmic steps of F-SAM with SGD as base optimizer in Algorithm 1. One can also use other base optimizers, e.g., AdamW, to update the model parameters in Eqn. (4).

Algorithm 1: F-SAM algorithm

Input: Loss function $L(\mathbf{w})$, training datasets $\mathcal{S} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$, minibatch size b , neighborhood size ρ , learning rate γ , momentum factor λ , projection constant σ

Output: Trained weight \mathbf{w}

```

1 Initialize  $\mathbf{w}_0, t \leftarrow 0, \mathbf{m}_{-1} = \mathbf{0}$ ;
2 while not converged do
3   Sample a minibatch data  $\mathcal{B}$  of size  $b$  from  $\mathcal{S}$ ;
4    $\mathbf{g}_t = \nabla L_{\mathcal{B}}(\mathbf{w}_t)$ ;
5    $\mathbf{m}_t = \lambda \mathbf{m}_{t-1} + (1 - \lambda) \mathbf{g}_t$ ;
6   Compute adversarial adversarial perturbation:
7      $\boldsymbol{\epsilon}_t = \rho \frac{\mathbf{d}_t}{\|\mathbf{d}_t\|}$  where  $\mathbf{d}_t = \mathbf{g}_t - \sigma \mathbf{m}_t$ ;
8   Compute the gradient approximation  $\mathbf{g}$ :
9      $\mathbf{g}'_t \leftarrow \nabla L_{\mathcal{B}}(\mathbf{w}_t + \boldsymbol{\epsilon}_t)$ ;
10  Update  $\mathbf{w}$  using gradient descent:
11     $\mathbf{w}_{t+1} \leftarrow \mathbf{w}_t - \gamma \mathbf{g}'_t$ ;
12     $t \leftarrow t + 1$ ;
13 return  $\mathbf{w}_t$ .
```

5.3. A Friendly Perspective of the Loss Objective

Here we provide an intuitive understanding on the difference between SAM and F-SAM. For a minibatch data \mathcal{B} , SAM solves the maximization problem to seek adversarial perturbation: $\boldsymbol{\epsilon}_s^{\text{SAM}} = \arg \max_{\|\boldsymbol{\epsilon}\|_2 \leq \rho} L_{\mathcal{B}}(\mathbf{w} + \boldsymbol{\epsilon})$. In contrast, by observing $\widetilde{\text{Proj}}_{\nabla L(\mathbf{w}_t)}^{\top} \nabla L_{\mathcal{B}}(\mathbf{w}_t)$ in Eqn. (12), F-SAM indeed computes the adversarial perturbation by maximizing the current minibatch loss while minimizing the loss on the entire dataset:

$$\boldsymbol{\epsilon}_s^{\text{F-SAM}} = \arg \max_{\|\boldsymbol{\epsilon}\|_2 \leq \rho} L_{\mathcal{B}}(\mathbf{w} + \boldsymbol{\epsilon}) - \sigma L_{\mathcal{D}}(\mathbf{w} + \boldsymbol{\epsilon}).$$

Then combining the following minimization problem, one can write F-SAM’s bi-level optimization problem as

$$\begin{aligned} \min_{\mathbf{w}} \quad & \mathbb{E}_{\mathcal{B}} [L_{\mathcal{B}}(\mathbf{w} + \boldsymbol{\epsilon}_s)] \\ \text{s.t.} \quad & \boldsymbol{\epsilon}_s = \arg \max_{\|\boldsymbol{\epsilon}\|_2 \leq \rho} L_{\mathcal{B}}(\mathbf{w} + \boldsymbol{\epsilon}) - \sigma L_{\mathcal{D}}(\mathbf{w} + \boldsymbol{\epsilon}). \end{aligned}$$

The insight behind is that F-SAM aims to find an adversarial perturbation $\boldsymbol{\epsilon}_s$ that increases the loss sharpness of the current minibatch data while minimizing the impact on the loss sharpness of the other data points as much as possible. Then F-SAM minimizes the loss value and sharpness of current minibatch data. So the adversarial perturbation in F-SAM is “friendly” to other data points, yielding consistent sharpness minimization. Moreover, this friendliness enables high robustness of F-SAM to the perturbation radius as shown in Sec. 6.3. But for vanilla SAM, its adversarial perturbation increases the loss sharpness of current minibatch data and

other data points but minimizes the loss value and sharpness of only current minibatch data. So in vanilla SAM, the second minimization step indeed does not consider the loss sharpness increment caused by its first adversarial perturbation step. This inconsistency may impair the sharpness optimization and thus limit the generalization.

5.4. Convergence Analysis

In this section, we analyze the convergence properties of the F-SAM algorithm under non-convex setting. We follow basic assumptions that are standard in convergence analysis of stochastic optimization [15, 36, 37].

Theorem 2. *Assume Assumption 1 and 2 hold. Assume that SAM uses SGD as the base optimizer with a learning rate γ to update the model parameter in Eqn. (4). By setting learning rate $\gamma = \frac{\gamma_0}{\sqrt{T}} \leq 1/\beta$ and the perturbation radius $\rho_t = \frac{\rho_0}{\sqrt{t}}$, we have*

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} \|\nabla L(\mathbf{w}_t)\|^2 \leq \frac{2\Delta}{\gamma_0 \sqrt{T}} + \frac{\Theta}{\sqrt{T}} + \frac{\Pi \log T}{\sqrt{T}}, \quad (15)$$

where $\Delta = \mathbb{E}[L(\mathbf{w}_0) - L(\mathbf{w}^*)]$ with the optimal solution \mathbf{w}^* to $L(\mathbf{w})$, $\Theta = 2\beta M \gamma_0 + \rho_0^2 \beta^3 \gamma_0$, and $\Pi = \rho_0^2 \beta^2$.

See the proof in Appendix A3. For non-convex stochastic optimization, Theorem 2 shows that F-SAM has the convergence rate $\mathcal{O}(\log T/\sqrt{T})$ and share the same convergence speed as SAM. But F-SAM enjoys better generalization performance than SAM as shown in Section 6.

6. Numerical Experiments

Here we test F-SAM on various tasks and network architectures under two popular settings: 1) training from scratch, and 2) transfer learning by fine-tuning pretrained models. Moreover, ablation study and additional experiments show more insights on F-SAM.

6.1. Training From Scratch

CIFAR. We first evaluate over CIFAR-10/100 datasets [39]. The network architectures we considered include VGG16-BN [41], ResNet-18 [42], WRN-28-10 [43] and PyramidNet-110 [44]. The first three models are trained for 200 epochs while for PyramidNet-110 we train for 300 epochs. We set the initial learning rate as 0.05 with a cosine learning rate schedule. The momentum and weight decay are set to 0.9 and 0.0005 for SGD, respectively. SAM and variants adopt the same setting except that the weight decay is set to 0.001 following [29, 45]. We apply standard random horizontal flipping, cropping, normalization, and cutout augmentation [46]. For SAM, we set the perturbation radius ρ as 0.1 and 0.2 for CIFAR-10 and CIFAR-100 [29, 45]. For ASAM, we adopt the recommend ρ as 2 and

CIFAR-10	SGD	ASAM	FisherSAM	SAM	F-SAM (ours)
VGG16-BN	94.96 \pm 0.15	95.57 \pm 0.05	95.55 \pm 0.08	95.42 \pm 0.05	95.62 \pm 0.11
ResNet-18	96.25 \pm 0.06	96.63 \pm 0.15	96.72 \pm 0.03	96.58 \pm 0.10	96.75 \pm 0.09
WRN-28-10	97.08 \pm 0.16	97.64 \pm 0.13	97.46 \pm 0.18	97.32 \pm 0.11	97.53 \pm 0.11
PyramidNet-110	97.63 \pm 0.09	97.82 \pm 0.07	97.64 \pm 0.09	97.70 \pm 0.10	97.84 \pm 0.05

Table 1. Test accuracy (%) comparison of various neural networks on CIFAR-10.

CIFAR-100	SGD	ASAM	FisherSAM	SAM	F-SAM (ours)
VGG16-BN	75.43 \pm 0.19	76.27 \pm 0.35	76.90 \pm 0.37	76.63 \pm 0.20	77.08 \pm 0.17
ResNet-18	77.90 \pm 0.07	81.68 \pm 0.12	80.99 \pm 0.13	80.88 \pm 0.10	81.29 \pm 0.12
WRN-28-10	81.71 \pm 0.13	84.99 \pm 0.22	84.91 \pm 0.07	84.88 \pm 0.10	85.16 \pm 0.07
PyramidNet-110	84.65 \pm 0.11	86.47 \pm 0.09	86.53 \pm 0.07	86.55 \pm 0.08	86.70 \pm 0.14

Table 2. Test accuracy (%) comparison of various neural networks on CIFAR-100.

4 for CIFAR-10 and CIFAR-100 [26]. For FisherSAM, we adopt the recommended $\rho = 0.1$. For F-SAM, we use the same ρ as SAM, and tune $\lambda = \{0.6, 0.9, 0.95\}$ although all parameters outperforms SAM. We find that $\lambda = 0.9$ works best for WRN-28-10 and PyramidNet-110, while $\lambda = 0.6$ achieves the best for others. We simply set $\sigma = 1$ for all the following experiments. Experiments are repeated over 3 independent trials.

CIFAR-10	ASAM	F-ASAM (ours)
ResNet-18	96.63 \pm 0.15	96.77 \pm 0.11
WRN-28-10	97.64 \pm 0.13	97.73 \pm 0.04
CIFAR-100	ASAM	F-ASAM (ours)
ResNet-18	81.68 \pm 0.12	81.82 \pm 0.14
WRN-28-10	84.99 \pm 0.22	85.24 \pm 0.16

Table 3. Results on integration with ASAM.

In Tab. 1 and 2, we observe that F-SAM achieves 0.1 to 0.2 accuracy gain on CIFAR-10 and 0.2 to 0.4 on CIFAR-100 in all tested scenarios. Moreover, F-SAM outperforms SAM’s adaptive variants in most cases, and in Tab. 3, we integrate our friendly approach to ASAM (see Appendix A4) and again observe consistent improvement.

ImageNet. Next, we investigate the performance of F-SAM on larger scale datasets by training ResNet-50 [42] on ImageNet [47] from scratch. We set the training epochs to 90 with batch size 128, weight decay 0.0001, and momentum 0.9. We set the initial learning rate as 0.05 with a cosine schedule [29]. We apply basic data preprocessing and augmentation [48]. We set $\rho = 0.075$ for all three SAM and variants and $\lambda = 0.95$ for F-SAM. In Tab. 4, we observe that F-SAM outperforms SAM and ASAM and offers an accuracy gain of 0.21 over SAM. This validates the effectiveness of F-SAM on large-scale problems.

ImageNet	SGD	ASAM	SAM	F-SAM
ResNet-50	76.62 \pm 0.12	77.10 \pm 0.14	77.14 \pm 0.16	77.35 \pm 0.12

Table 4. Test accuracy (%) comparison on ImageNet.

6.2. Transfer Learning

One of the fascinating training pipelines for modern DNNs is transfer learning, i.e., first training a model on a large datasets and then easily and quickly adapting to novel target datasets by fine-tuning. In this subsection, we evaluate the performance of transfer learning for F-SAM. We use a Deit-small model [49] pre-trained on ImageNet (with public available checkpoints¹). We use AdamW [6] as base optimizer and train the model for 10 epochs with batch size 128, weight decay 10^{-5} and initial learning rate of 10^{-4} . We adopt $\rho = 0.075$ for SAM and F-SAM. The results are shown in Tab. 5. We observe that F-SAM consistently outperforms SAM and AdamW by an accuracy gain of 0.1 to 0.7. An intriguing observation is that on small Flowers102 dataset [38] (with 1020 images), AdamW even outperforms SAM. We attribute this to the fact that when the dataset is small, the full gradient component within the minibatch gradient is more prominent, and its detrimental impact on convergence and generalization becomes more evident.

Datasets	AdamW	SAM	F-SAM (ours)
CIFAR-10	98.10 \pm 0.10	98.27 \pm 0.05	98.43 \pm 0.07
CIFAR-100	88.44 \pm 0.10	89.10 \pm 0.11	89.49 \pm 0.12
Stanford Cars	74.78 \pm 0.09	75.39 \pm 0.05	75.82 \pm 0.14
OxfordIIITPet	92.42 \pm 0.43	92.70 \pm 0.26	92.90 \pm 0.23
Flowers102	74.82 \pm 0.36	74.38 \pm 0.24	75.15 \pm 0.35

Table 5. Results on transfer learning by fine-tuning.

6.3. Additional Studies

6.3.1 Robustness to Label Noise

Since previous works have shown that SAM is robust to label noise, in this subsection, we also test the performance of F-SAM in the presence of symmetric label noise by random flipping. The training settings are the same as in Sec. 6.1. From the results in Tab. 6, we observe that F-SAM consistently improves the performance from SAM, confirming

¹<https://github.com/facebookresearch/deit>

its improved generalization. Notably, such improvement is more obvious when the noise rate is large (i.e. 60%-80%). This is perhaps because when the data label is wrong, the harmfulness of increasing the sharpness of other batch data is more prominent, even posing challenges to convergence. On severe noise ratio of 80%, we find that vanilla SAM even suffers a collapse with the original setting, which is also reported by Foret et al. [14], and our F-SAM relieves it, achieving a remarkable accuracy gain of 28%.

Rates	20%	60%	70%	80%
SGD	87.05±0.06	52.21±0.23	39.31±0.13	27.86±0.53
SAM	95.27±0.09	90.08±0.09	84.89±0.10	31.73±3.37
F-SAM	95.42±0.08	90.47±0.05	86.48±0.07	59.39±1.77

Table 6. Results under label noise on CIFAR-10 with ResNet-18.

6.3.2 Robustness to Perturbation Radius

One deficiency of SAM is its sensitivity to the perturbation radius, which may require choosing different radii for optimal performance on different datasets. Especially, an excessively large perturbation radius can result in a sharp decrease in generalization performance. We attribute a portion of this over-sensitivity to the impact of the adversarial perturbation derived from the current minibatch data on the other data samples. Such an impact can be more obvious when the perturbation radius increases as the magnitude of the full gradient component grows correspondingly. In our F-SAM, such impact on other data samples is eliminated as much as possible, and thus, we can expect that it is more robust to the choices of the perturbation radius. In Fig. 3, we plot the performance of SAM and F-SAM under different perturbation radii. We can observe that F-SAM is much less sensitive to ρ than SAM. Especially on larger perturbation radius, the performance gain of F-SAM over SAM is more significant. For example, when ρ is set to 2x larger than the optimal on CIFAR100, SAM’s performance suffers a sharp drop from 80.88% to 78.83%, but F-SAM still maintains a good performance of 80.30%. We further compare the training curves of SAM and F-SAM under different perturbation radii in Appendix A6, showing that F-SAM greatly facilitates training on large radius. This confirms F-SAM’s improved robustness to different perturbation radii.

6.3.3 Results on Different Batch Sizes

As the main claim of this paper is that the full gradient component existing in SAM’s adversarial perturbation does not contribute to generalization and may have detrimental effects. As the training batch size increases, the corresponding full gradient component existing in the minibatch gradient can be strengthened. Thus, the improvement of F-SAM from removing the full gradient component could be more obvious. In Fig. 4, we compare the performance of SGD, SAM, and our F-SAM under batch sizes from

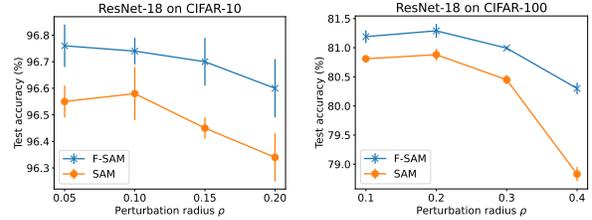


Figure 3. Results under different perturbation radii ρ .

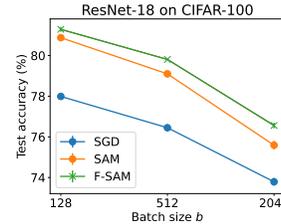


Figure 4. Performance comparison with different batch sizes.

{128, 512, 2048} while keeping other hyper-parameters unchanged. We observe that the performance gain of F-SAM over SAM increases significantly as the batch size increases. This further validates our findings and supports the effectiveness of removing the full gradient component.

Connection to m -sharpness. We would like to relate our removing the full gradient component to the intriguing m -sharpness in SAM, which is defined as the search for an adversarial perturbation that maximizes the sum of losses over sub-batches of m training points [14]. This phenomenon emphasizes the importance of using a small m to effectively improve generalization [14, 15, 50]. It is worth noting that by utilizing a small sub-batch, one implicitly suppresses the effects of the full gradient component. Hence, our findings can provide new insights into understanding m -sharpness.

7. Conclusion

In this paper, we conduct an in-depth investigation into the core components of SAM’s generalization. By decomposing the minibatch gradient into two orthogonal components, we discover that the full gradient component in adversarial perturbation contributes minimally to generalization and may even have detrimental effects, while the stochastic gradient noise component plays a crucial role in enhancing generalization. We then propose a new variant of SAM called F-SAM to eliminate the undesirable effects of the full gradient component. Extensive experiments across various tasks demonstrate that F-SAM significantly improves the robustness and generalization performance of SAM.

Acknowledgement This work was supported by National Key Research Development Project (2023YFF1104202), National Natural Science Foundation of China (62376155), Shanghai Municipal Science and Technology Research Program (22511105600) and Major Project (2021SHZDZX0102). Pan Zhou was supported by the Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier 1 grant.

References

- [1] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. **1**
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [3] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (CVPR)*, 2021. **1**
- [4] H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(3):400–407, 1951. **1**
- [5] Diederik P. Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. **2**
- [6] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. **7**
- [7] Juntang Zhuang, Tommy Tang, Yifan Ding, Sekhar C Tatikonda, Nicha Dvornek, Xenophon Papademetris, and James Duncan. Adabelief optimizer: Adapting stepsizes by the belief in observed gradients. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [8] Xingyu Xie, Pan Zhou, Huan Li, Zhouchen Lin, and Shuicheng Yan. Adan: Adaptive nesterov momentum algorithm for faster optimizing deep models. *arXiv preprint arXiv:2208.06677*, 2022.
- [9] Pan Zhou, Xingyu Xie, and YAN Shuicheng. Win: Weight-decay-integrated nesterov acceleration for adaptive gradient algorithms. In *International Conference on Learning Representations (ICLR)*, 2023. **1**
- [10] Sepp Hochreiter and Jürgen Schmidhuber. Flat minima. *Neural computation*, 1997. **1, 2, 3**
- [11] Laurent Dinh, Razvan Pascanu, Samy Bengio, and Yoshua Bengio. Sharp minima can generalize for deep nets. In *International Conference on Machine Learning (ICML)*, 2017.
- [12] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [13] Pan Zhou, Jiashi Feng, Chao Ma, Caiming Xiong, Steven Chu Hong Hoi, et al. Towards theoretically understanding why sgd generalizes better than adam in deep learning. *Advances in Neural Information Processing Systems*, 33:21285–21296, 2020. **1, 2, 3**
- [14] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. In *International Conference on Learning Representations (ICLR)*, 2020. **1, 2, 3, 8**
- [15] Maksym Andriushchenko and Nicolas Flammarion. Towards understanding sharpness-aware minimization. In *International Conference on Machine Learning (ICML)*, 2022. **2, 3, 6, 8**
- [16] Zixiang Chen, Junkai Zhang, Yiwen Kou, Xiangning Chen, Cho-Jui Hsieh, and Quanquan Gu. Why does sharpness-aware minimization generalize better than sgd? *arXiv preprint arXiv:2310.07269*, 2023. **2**
- [17] Pratik Chaudhari, Anna Choromanska, Stefano Soatto, Yann LeCun, Carlo Baldassi, Christian Borgs, Jennifer Chayes, Levent Sagun, and Riccardo Zecchina. Entropy-sgd: Biasing gradient descent into wide valleys. In *International Conference on Learning Representations (ICLR)*, 2017. **2**
- [18] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. In *International Conference on Learning Representations (ICLR)*, 2017.
- [19] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. *arXiv preprint arXiv:1803.05407*, 2018. **2**
- [20] Wei Wen, Yandan Wang, Feng Yan, Cong Xu, Chunpeng Wu, Yiran Chen, and Hai Li. Smoothout: Smoothing out sharp minima to improve generalization in deep learning. *arXiv preprint arXiv:1805.07898*, 2018. **2**
- [21] Yusuke Tsuzuku, Issei Sato, and Masashi Sugiyama. Normalized flat minima: Exploring scale invariant definition of flat minima for neural networks using pac-bayesian analysis. In *International Conference on Machine Learning*, pages 9636–9647. PMLR, 2020.
- [22] Yaowei Zheng, Richong Zhang, and Yongyi Mao. Regularizing neural networks via adversarial model perturbation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [23] Devansh Bisla, Jing Wang, and Anna Choromanska. Low-pass filtering sgd for recovering flat optima in the deep learning optimization landscape. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2022.
- [24] Tao Li, Weihao Yan, Zehao Lei, Yingwen Wu, Kun Fang, Ming Yang, and Xiaolin Huang. Efficient generalization improvement guided by random weight perturbation. *arXiv preprint arXiv:2211.11489*, 2022.
- [25] Tao Li, Weihao Yan, Qinghua Tao, Zehao Lei, Yingwen Wu, Kun Fang, Mingzhen He, and Xiaolin Huang. Revisiting random weight perturbation for efficiently improving generalization. In *OPT 2023: Optimization for Machine Learning*, 2023. **2**
- [26] Jungmin Kwon, Jeongseop Kim, Hyunseo Park, and In Kwon Choi. Asam: Adaptive sharpness-aware minimization for scale-invariant learning of deep neural networks. In *International Conference on Machine Learning (ICML)*, 2021. **2, 7**
- [27] Minyoung Kim, Da Li, Shell X Hu, and Timothy Hospedales. Fisher sam: Information geometry and sharpness aware minimisation. In *International Conference on Machine Learning (ICML)*, 2022. **2**

- [28] Juntang Zhuang, Boqing Gong, Liangzhe Yuan, Yin Cui, Hartwig Adam, Nicha Dvornek, Sekhar Tatikonda, James Duncan, and Ting Liu. Surrogate gap minimization improves sharpness-aware training. In *International Conference on Learning Representations (ICLR)*, 2022. 2
- [29] Bingcong Li and Georgios B Giannakis. Enhancing sharpness-aware optimization through variance suppression. *arXiv preprint arXiv:2309.15639*, 2023. 2, 6, 7
- [30] Jiawei Du, Hanshu Yan, Jiashi Feng, Joey Tianyi Zhou, Liangli Zhen, Rick Siow Mong Goh, and Vincent YF Tan. Efficient sharpness-aware minimization for improved training of neural networks. In *International Conference on Learning Representations (ICLR)*, 2022. 2
- [31] Yong Liu, Siqi Mai, Xiangning Chen, Cho-Jui Hsieh, and Yang You. Towards efficient and scalable sharpness-aware minimization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (CVPR)*, 2022.
- [32] Jiawei Du, Daquan Zhou, Jiashi Feng, Vincent YF Tan, and Joey Tianyi Zhou. Sharpness-aware training for free. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [33] Yang Zhao, Hao Zhang, and Xiuyuan Hu. SS-SAM: Stochastic scheduled sharpness-aware minimization for efficiently training deep neural networks. *arXiv preprint arXiv:2203.09962*, 2022.
- [34] Weisen Jiang, Hansi Yang, Yu Zhang, and James Kwok. An adaptive policy to employ sharpness-aware minimization. In *International Conference on Learning Representations (ICLR)*, 2023. 2, 3
- [35] Kaiyue Wen, Tengyu Ma, and Zhiyuan Li. How sharpness-aware minimization minimizes sharpness? In *International Conference on Learning Representations (ICLR)*, 2023. 2
- [36] Saeed Ghadimi and Guanghui Lan. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 2013. 3, 6
- [37] Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the polyak-lojasiewicz condition. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 2016. 3, 6
- [38] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *Indian conference on computer vision, graphics & image processing*. IEEE, 2008. 4, 7
- [39] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. *Technical Report*, 2009. 4, 6
- [40] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *IEEE conference on computer vision and pattern recognition (ICCV)*, 2012. 4
- [41] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 6
- [42] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6, 7
- [43] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In Richard C. Wilson, Edwin R. Hancock, and William A. P. Smith, editors, *British Machine Vision Conference (BMVC)*, 2016. 6
- [44] Dongyoon Han, Jiwhan Kim, and Junmo Kim. Deep pyramidal residual networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 6
- [45] Peng Mi, Li Shen, Tianhe Ren, Yiyi Zhou, Xiaoshuai Sun, Rongrong Ji, and Dacheng Tao. Make sharpness-aware minimization stronger: A sparsified perturbation approach. *arXiv preprint arXiv:2210.05177*, 2022. 6
- [46] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017. 6
- [47] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 7
- [48] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 7
- [49] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning (ICML)*, 2021. 7
- [50] Kayhan Behdin, Qingquan Song, Aman Gupta, David Durfee, Ayan Acharya, Sathiya Keerthi, and Rahul Mazumder. Improved deep neural network generalization using m-sharpness-aware minimization. *arXiv preprint arXiv:2212.04343*, 2022. 8