
A Universal Class of Sharpness-Aware Minimization Algorithms

Behrooz Tahmasebi¹ Ashkan Soleymani² Dara Bahri³ Stefanie Jegelka^{4,1} Patrick Jaillet²

Abstract

Recently, there has been a surge in interest in developing optimization algorithms for overparameterized models as achieving generalization is believed to require algorithms with suitable biases. This interest centers on minimizing sharpness of the original loss function; the Sharpness-Aware Minimization (SAM) algorithm has proven effective. However, most literature only considers a few sharpness measures, such as the maximum eigenvalue or trace of the training loss Hessian, which may not yield meaningful insights for non-convex optimization scenarios like neural networks. Additionally, many sharpness measures are sensitive to parameter invariances in neural networks, magnifying significantly under rescaling parameters. Motivated by these challenges, we introduce a new class of sharpness measures in this paper, leading to new sharpness-aware objective functions. We prove that these measures are *universally expressive*, allowing any function of the training loss Hessian matrix to be represented by appropriate hyperparameters. Furthermore, we show that the proposed objective functions explicitly bias towards minimizing their corresponding sharpness measures, and how they allow meaningful applications to models with parameter invariances (such as scale-invariances). Finally, as instances of our proposed general framework, we present *Frob-SAM* and *Det-SAM*, which are specifically designed to minimize the Frobenius norm and the determinant of the Hessian of the training loss, respectively. We also demonstrate the advantages of our general framework through extensive experiments.

1. Introduction

Understanding the generalization capabilities of overparameterized networks is a fundamental, yet unsolved challenge in deep learning. It is postulated that achieving near-zero training loss alone may be insufficient as there exist many instances where global minima fail to exhibit satisfactory generalization performance. To this end, a dominant observation asserts that the characteristics of the loss landscape play a pivotal role in determining which parameters have low training loss while also exhibiting generalization capabilities.

A recently proposed approach to consider the geometric aspects of the loss landscape, with the aim of achieving generalization, entails the avoidance of sharp minima. For example, the celebrated Sharpness-Aware Minimization (SAM) algorithm has shown enhancements in generalization across many practical tasks (Foret et al., 2021). While the concept of sharpness lacks a precise definition in a general sense, people often introduce various measures to quantify it in practice (Dinh et al., 2017). Many sharpness measures in the literature rely on the second-order derivative characteristics of the training loss function, such as the trace or the operator norm of the Hessian matrix (Chaudhari et al., 2019; Keskar et al., 2017).

Nevertheless, traditional methodologies for quantifying sharpness may not suffice to ensure generalization, given the intricate geometry of the loss landscape, which may necessitate different regularization techniques. Moreover, many existing sharpness measures fail to encapsulate the genuine essence of sharpness in deep neural networks because the Hessian matrix no longer maintains positive semi-definiteness. Furthermore, neural networks exhibit parameter invariances, wherein different parameterizations can yield identical functions — such as scaling invariances in ReLU networks. Consequently, an effective measure of sharpness should remain invariant in the face of such parameter variations. Unfortunately, conventional approaches for quantifying sharpness frequently fall short in addressing this phenomenon.

Therefore, a fundamental question arises: how can one succinctly represent all measures of sharpness within a compact parameterized framework that also enables meaningful applications to models with parameter invariances? This

¹MIT CSAIL ²MIT LIDS ³Google DeepMind ⁴TU Munich.
Correspondence to: Behrooz Tahmasebi <bzt@mit.edu>.

Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024. Copyright 2024 by the author(s).

question holds significance in applications as it allows *learning/designing the regularization* in cases where information about the geometry of the loss landscape or parameter invariances is provided, either empirically or through assumption. To the best of our knowledge, this question has remained fairly unexplored in the deep learning literature.

In this paper, we characterize *all* sharpness measures (i.e., functions of the Hessian of the training loss) through an average-based parameterized representation. We prove that by changing the (hyper)parameters, the provided representation spans all the sharpness measures as a function of the Hessian matrix. In other words, it is provably a *universal representation*. We also provide quantitative theoretical results on the complexity of the sharpness representation as a function of the data dimension.

Moreover, attached to any representation of sharpness, we provide a new loss function, and we prove that the new loss function is *biased* toward minimizing its corresponding sharpness measure. Since the parameterized representation reduces to SAM (i.e., worst-direction) and average-direction sharpness measures (Wen et al., 2023a) in special cases, it can be considered as a generalized (hyper)parameterized sharpness-aware minimization algorithm. This generalizes the recent study of the explicit bias of a few sharpness-aware minimization algorithms (Wen et al., 2023a) to a comprehensive class of objectives. Furthermore, this allows us to readily design algorithms with *any* bias of interest, while to the best of our knowledge, only algorithms with biases towards minimizing the trace, operator norm of the Hessian matrix, and a few other sharpness measures are known in the literature. As instances of our proposed general algorithm, we present *Frob-SAM* and *Det-SAM*, two new sharpness-aware minimization algorithms that are specifically designed to minimize the Frobenius norm and the determinant of the Hessian of the training loss function, respectively.

An interesting feature of the given representation is that it provides a systematic way to construct sharpness measures respecting parameter invariances, e.g., scale-invariance in neural networks. As a specific example, we provide a class of loss functions and algorithms that are invariant under parameter scaling. Note that the classical sharpness measures, such as trace or the operator norm of the Hessian matrix, are not invariant to rescaling or group actions.

In the experiments, we explore extensively two specific choices of these algorithms: (1) *Frob-SAM*: an algorithm biased toward minimizing the Frobenius norm of the Hessian, a meaningful sharpness notion for non-convex optimization problems and (2) *Det-SAM*: an algorithm biased toward minimizing the determinant of the Hessian, a scale-invariant sharpness measure. We demonstrate the advantages of these two cases through an extensive series of experiments.

In short, in this paper we make the following contributions:

- We propose a new class of sharpness measures, as function of the training loss Hessian. We prove that the new representation is *universally expressive*, meaning that it covers all sharpness measures of the Hessian as its (hyper)parameters change.
- Along with each sharpness measure we provide an optimization objective and prove that the new objective is explicitly *biased* toward minimizing the corresponding sharpness measure.
- The structure of the proposed method allows meaningful applications to models with parameter invariances, as it provides a class of objective functions for any desired type of parameter invariance.
- We introduce two fundamental illustrative examples of our proposed general representation and the corresponding algorithms: *Frob-SAM* and *Det-SAM*. *Frob-SAM* is geared towards minimizing the Frobenius norm of the Hessian matrix, providing a meaningful and natural solution to the definition problem of sharpness for non-convex optimization problems. Conversely, *Det-SAM* is focused on minimizing the determinant¹ of Hessian, addressing scale-invariant issues related to parameterization.

2. Related Work

Foret et al. (2021) recently proposed the Sharpness-Aware Minimization (SAM) algorithm to avoid sharp minima. The SAM objective has connections to a similar robust optimization problem that was suggested for the study of adversarial attacks in deep learning (Madry et al., 2018). Besides SAM, Nitanda et al. (2023) show how parameter averaging for SGD is biased toward flatter minima. Label noise SGD also prefers flat minima (Damian et al., 2021). Woodworth et al. (2020) studied the role of sharpness in overparameterization from a kernel perspective. See (Wang et al., 2023) for the applications of flat minima for domain generalization (see also (Cha et al., 2021)). For applications of SAM in large language models, see (Bahri et al., 2022) (also (Zhong et al., 2022)), and (Qu et al., 2022; Shi et al., 2023) for federated learning). Besides those applications, Wen et al. (2023b) prove that current sharpness minimization algorithms sometimes fail to generalize for non-generalizing flattest models.

The (implicit) bias of many optimization algorithms and architectures has been studied, from the Gradient Descent (GD) (Ji & Telgarsky, 2019b; Soudry et al., 2018) to mirror descent (Gunasekar et al., 2018a; Azizan & Hassibi, 2019); see also (Gunasekar et al., 2018b) for linear convolutional networks, and (Lawrence et al., 2022) for equivariant net-

¹To be more precise, the product of non-zero eigenvalues.

works. Ji & Telgarsky (2019a) observed that linear neural networks are biased toward weight alignment for different layers (see (Le & Jegelka, 2022) for non-linear networks). Andriushchenko & Flammarion (2022) study implicit bias of SAM for diagonal linear networks, and Wen et al. (2023a) find the explicit bias of the Gaussian averaging method and other SAM variants.

The role of scale-invariance in generalization in deep learning is emphasized in (Neyshtabur et al., 2017). Dinh et al. (2017) point out that parameter invariances can lead to the different parameterization of the same function, making the definition of flatness challenging; see also (Andriushchenko et al., 2023b) for a recent study. This motivates the study of sharpness measures that are invariant to such reparametrizations.

There have been a few attempts to address reparametrization problems with sharpness measures recently. Kwon et al. (2021) proposed to adaptively calculate the sharpness in a normalized ball around the loss function to achieve scale invariance. However, their method is limited to scaling problems. Kim et al. (2022) took a step further and introduced a new SAM algorithm by capturing the neighborhood of the parameters in an ellipsoid induced by the Fisher information. This way, the neighborhood becomes invariant with respect to the parameter invariances in the network. Jang et al. (2022) defined an information geometric sharpness measure by investigating the eigenspaces of Fisher Information Matrix (FIM) of distribution parameterized by neural networks. They proved scale-invariance properties for their notion. Even though Kim et al. (2022) and Jang et al. (2022) enjoy some parameter invariance properties, (1) in practice, their methods are limited to classification tasks because of FIM calculation, (2) the underlying explicit biasing of their algorithms remains a mystery and is not guaranteed.²

3. Background

3.1. Setting

Consider a standard learning setup with a labeled dataset \mathcal{S} , and a training loss function $L : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$, where $L(x)$ denotes the training loss over \mathcal{S} computed for the parameters $x \in \mathbb{R}^d$. The main objective in Empirical Risk Minimization (ERM) is to minimize the training loss $L(x)$ over the feasibility set $\mathcal{X} \subseteq \mathbb{R}^d$. However, achieving parameters satisfying $L(x) \approx 0$ in overparameterized models is often straightforward. This is because in contrast to other models, in overparameterized models, there are *many* global minima, i.e., the set $\Gamma := \{x \in \mathcal{X} : L(x) = 0\}$ is a manifold – it is called the *zero-loss manifold* in the literature. Moreover, in practical scenarios, it is noteworthy that not all global

²Please refer to Appendix A for a more detailed overview of related work.

minima exhibit favorable generalization capabilities (Foret et al., 2021).

3.2. Background on SAM

It is hypothesized that the avoidance of sharp minima can enhance generalization performance (Hochreiter & Schmidhuber, 1997; Keskar et al., 2017; Izmailov et al., 2018). However, it should be noted that the concept of sharpness encompasses a multitude of distinct definitions in practical contexts. The Sharpness-Aware Minimization (SAM) algorithm (Foret et al., 2021) suggests minimizing the training loss function over a small ball around the parameters:

$$\min_{x \in \mathcal{X}} \left\{ L_{\text{SAM}}(x) := \max_{\|v\|_2 \leq 1} L(x + \rho v) \right\},$$

where $\rho \in \mathbb{R}_{\geq 0}$ is the *perturbation parameter*. Note that L_{SAM} can be decomposed into two terms:

$$L_{\text{SAM}}(x) = \underbrace{L(x)}_{\text{empirical loss}} + \underbrace{\max_{\|v\|_2 \leq 1} \{L(x + \rho v) - L(x)\}}_{\text{sharpness}}.$$

Foret et al. (2021) also suggest alternative average-based sharpness-aware objectives to use PAC bounds on the generalization error of overparameterized models; we follow the definition in (Wen et al., 2023b):

$$\begin{aligned} L_{\text{AVG}}(x) &:= \mathbb{E}_{v \sim \mathcal{N}(0, I)} \left[L\left(x + \frac{\rho v}{\|v\|_2}\right) \right] \\ &= \underbrace{L(x)}_{\text{empirical loss}} + \underbrace{\mathbb{E}_{v \sim \mathcal{N}(0, I)} \left[L\left(x + \frac{\rho v}{\|v\|_2}\right) - L(x) \right]}_{\text{sharpness}}. \end{aligned}$$

Wen et al. (2023b) recently proved that minimizing $L_{\text{SAM}}(x)$ will lead to global minima (i.e., $L(x) \approx 0$) with small $\lambda_{\max}(\nabla^2 L(x))$. In other words, SAM is (explicitly) biased towards minimizing $\lambda_{\max}(\nabla^2 L(x))$. Moreover, they show that using $L_{\text{AVG}}(x)$ biases towards minimizing $\frac{1}{d} \text{tr}(\nabla^2 L(x))$. This means that SAM measures the sharpness of a global minimum by $\lambda_{\max}(\nabla^2 L(x))$, while the average-based objective uses $\frac{1}{d} \text{tr}(\nabla^2 L(x))$ to evaluate it.

In the next examples, we argue how both sharpness measures above fail to define a meaningful notion for overparameterized models. In Example 2, a special case of problem with parameter invariances, i.e., under parameter rescalings is discussed.

Example 1. *The sharpness measures $\lambda_{\max}(\nabla^2 L(x))$, $\lambda_{\min}(\nabla^2 L(x))$ and $\text{tr}(\nabla^2 L(x))$ are conceptually meaningful when the objective function $L(x)$ is convex, therefore λ_i s are nonnegative. However, the Loss landscape of neural networks is highly nonconvex, and as a result, λ_i can be potentially negative. Consider the toy non-convex example of $L(x_1, x_2) = \frac{1}{2}(x_1^2 - x_2^2)$,*

$$\nabla^2 L = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \quad (1)$$

For all $x_1, x_2 \in \mathbb{R}$, we know that $\text{tr}(\nabla^2 L) = 1 + (-1) = 0$, which in the Trace measure of sharpness it suggests that all the points $x_1, x_2 \in \mathbb{R}$ are equally flat. Are these sharpness notions really capturing the intended concepts? For a better illustration, consider the plot of this function provided in Figure 1. This problem extends to other existing notions.

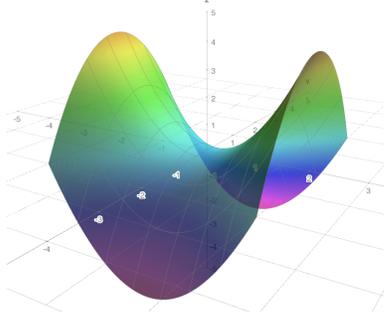


Figure 1. The loss landscape of the non-convex objective function $L(x_1, x_2) = \frac{1}{2}(x_1^2 - x_2^2)$. This examples shows how existing sharpness measures fall short of capturing sharpness meaning in non-convex settings. In particular, for all points $(x_1, x_2) \in \mathbb{R}^2$, $\text{tr}(\nabla^2 L(x_1, x_2)) = 1 + (-1) = 0$.

Example 2. Consider the loss function $L(x_1, x_2) = x_1^2 x_2^2 - 2x_1 x_2 + 1$ with two parameters $x_1, x_2 \in \mathbb{R}$. It is scale-invariant, i.e., $L(kx_1, \frac{x_2}{k}) = L(x_1, x_2)$ for all $k \neq 0$. Indeed, the zero-loss manifold $\Gamma = \{(x_1, x_2) : x_1 x_2 = 1\}$ contains infinitely many global minima. Straightforward calculation shows $\nabla^2 L(x_1, x_2) = \begin{pmatrix} 2x_2^2 & 4x_1 x_2 - 2 \\ 4x_1 x_2 - 2 & 2x_1^2 \end{pmatrix}$. Thus, we have $\frac{1}{2} \text{tr}(\nabla^2 L(x_1, x_2)) = x_1^2 + x_2^2$. After rescaling, we get $\frac{1}{2} \text{tr}(\nabla^2 L(x_1, x_2)) \Big|_{(kx_1, k^{-1}x_2)} = k^2 x_1^2 + \frac{x_2^2}{k^2} \neq \frac{1}{2} \text{tr}(\nabla^2 L(x_1, x_2))$. Therefore, as a sharpness measure, $\text{tr}(\nabla^2 L(x_1, x_2))$ is not scale-invariant. The problem magnifies in the limit: $\lim_{k \rightarrow \infty} \text{tr}(\nabla^2 L(x_1, x_2)) \Big|_{(kx_1, k^{-1}x_2)} = \infty$. Similar problems exist for $\lambda_{\max}(\nabla^2 L(x_1, x_2))$. However, $\det(\nabla^2 L(x_1, x_2))$ is scale-invariant; we have $\det(\nabla^2 L(x_1, x_2)) \Big|_{(kx_1, k^{-1}x_2)} = \det(\nabla^2 L(x_1, x_2))$ for all $k \neq 0$.

Note that neural networks are often scale-invariant, e.g., linear networks or ReLU networks after scaling up the parameters of one hidden layer and scaling down the parameters of another hidden layer encode the same functions.

4. A New Class of Sharpness Measures

To define a new class of sharpness measures, we take a closer look at the average-based sharpness-aware objective $L_{\text{AVG}}(x)$; using its Taylor expansion (Wen et al., 2023b), we

have

$$\begin{aligned} L_{\text{AVG}}(x) &= \mathbb{E}_{v \sim \mathcal{N}(0, I)} \left[L\left(x + \frac{\rho v}{\|v\|_2}\right) \right] \\ &\approx L(x) + \rho \mathbb{E}_{v \sim \mathcal{N}(0, I)} \left[\langle \nabla L(x), \frac{v}{\|v\|_2} \rangle \right] \\ &\quad + \rho^2 \mathbb{E}_{v \sim \mathcal{N}(0, I)} \left[\frac{v^t \nabla^2 L(x) v}{\|v\|_2^2} \right] \\ &= L(x) + \rho^2 \frac{\text{tr}(\nabla^2 L(x))}{d}. \end{aligned}$$

This intuitively tells us that for a small perturbation parameter ρ , the leading term in the objective function is the training loss $L(x)$, and after we get close to the zero-loss manifold Γ , the leading term becomes $\frac{1}{d} \text{tr}(\nabla^2 L(x))$, which is exactly the explicit bias of the average-based sharpness-aware minimization objective. This motivates us to define the following parameterized sharpness measure.

Definition 1 ((ϕ, ψ, μ) -sharpness measure). For any continuous functions $\phi, \psi : \mathbb{R} \rightarrow \mathbb{R}$ and any (Borel) measure μ on \mathbb{R}^d , the (ϕ, ψ, μ) -sharpness measure $S(x; \phi, \psi, \mu)$ is defined as

$$S(x; \phi, \psi, \mu) := \phi \left(\int \psi \left(\frac{1}{2} v^t \nabla^2 L(x) v \right) d\mu(v) \right). \quad (2)$$

Similarly, one can consider continuous functions $\psi : \mathbb{R} \rightarrow \mathbb{R}^m$ and $\phi : \mathbb{R}^m \rightarrow \mathbb{R}$, for some positive integer $m \geq 1$, and (Borel) measures $\mu_\ell, \ell \in [m]$, and define

$$S(x; \phi, \psi, \mu) := \phi \left(\int \psi_1 \left(\frac{1}{2} v^t \nabla^2 L(x) v \right) d\mu_1(v), \quad (3)$$

$$\int \psi_2 \left(\frac{1}{2} v^t \nabla^2 L(x) v \right) d\mu_2(v), \quad (4)$$

$$\dots, \quad (5)$$

$$\int \psi_m \left(\frac{1}{2} v^t \nabla^2 L(x) v \right) d\mu_m(v) \Big), \quad (6)$$

where we use $\mu := \mu_1 \otimes \mu_2 \otimes \dots \otimes \mu_m$ for the sake of brevity in our notation, and $\psi = (\psi_1, \psi_2, \dots, \psi_m)^t$.

We specify several examples of hyperparameters (ϕ, ψ, μ) in Table 1, which shows how (ϕ, ψ, μ) -sharpness measures can represent various notions of sharpness, as a function of the training loss Hessian matrix.

5. Expressive Power and Universality

In this section, we prove that the proposed class of sharpness measures is *universal*. In other words, for any continuous function $S : \mathbb{R}^d \rightarrow \mathbb{R}$, we specify continuous functions ϕ, ψ and a (Borel) probability measure³ μ on \mathbb{R}^d such that $S(\lambda_1, \lambda_2, \dots, \lambda_d) = S(x; \phi, \psi, \mu)$, where $\lambda_i, i \in [d]$, are the eigenvalues of the Hessian matrix $\nabla^2 L(x)$.

³We indeed prove that Borel probability measures (as a subset of arbitrary Borel measures) are enough to achieve universality.

Table 1. Various (ϕ, ψ, μ) -sharpness measures (Appendix B; $\lambda_i, i \in [d]$, are eigenvalues of Hessian).

Hyperparameters				$S(x; \phi, \psi, \mu)$ (or bias)	Reference
$\phi(t)$	$\psi(t)$	m	μ		
t	t	1	Uniform(\mathbb{S}^{d-1})	$\frac{1}{2d} \text{tr}(\nabla^2 L(x)) = \frac{1}{2d} \sum_{i=1}^d \lambda_i$	(Wen et al., 2023a)
$2 \max(\{\cdot\})$	(t, t, \dots)	∞	$\otimes_{\ v\ _2=1} \delta_v$	$\max_{i \in [d]} \lambda_i$	(Wen et al., 2023a)
$2(t_2 - t_1^2)$	(t, t^2)	2	$\mathcal{N}(0, I_d) \otimes \mathcal{N}(0, I_d)$	$\sum_{i=1}^d \lambda_i^2 = \ \nabla^2 L(x)\ _F$	This paper
$(2\pi)^d/t^2$	$\exp(-t)$	1	Lebesgue measure on \mathbb{R}^d	$\det(\nabla^2 L(x)) = \prod_{i=1}^d \lambda_i$	This paper
t	t^n	1	Uniform(\mathbb{S}^{d-1})	$q(\lambda_1, \lambda_2, \dots, \lambda_d)$	This paper
$1/t^2$	$\exp(\sigma t)$	1	$\mathcal{N}(0, I_d)$	$\prod_{i=1}^d (1 - \sigma \lambda_i)$	This paper

* $q(\lambda_1, \lambda_2, \dots, \lambda_d)$ is a specific homogeneous polynomial of degree n ; see Equation (22).

Theorem 1 (Universality of the (ϕ, ψ, μ) -sharpness measures for functions of Hessian eigenvalues). *Let $\mathcal{A} \subseteq \mathbb{R}^d$ be a compact set. For any continuous function $S : \mathcal{A} \rightarrow \mathbb{R}$, there exist a product (Borel) probability measure μ , a positive integer $m \leq d$, and continuous functions $\phi : \mathbb{R}^m \rightarrow \mathbb{R}$ and $\psi : \mathbb{R} \rightarrow \mathbb{R}^m$, such that $S(\lambda_1, \lambda_2, \dots, \lambda_d) = S(x; \phi, \psi, \mu)$ for any $x \in \mathcal{A}$, where $\lambda_i, i \in [d]$, are the eigenvalues of the Hessian matrix $\nabla^2 L(x)$.*

We present the proof of Theorem 1 in Appendix C.

Note that to achieve universality, we need the functions ϕ, ψ to be of dimension $m = d$. However, as one can see in Table 1, many celebrated sharpness measures can indeed be represented using only small m . We believe that practically small hyperparameter m is enough, as it is motivated from the measures in Table 1.

While we proved the universality of the proposed class of sharpness measures for continuous functions of the Hessian eigenvalues, one may be interested in measuring sharpness with more information about the loss Hessian (e.g., the eigenvectors of the loss Hessian). The following theorem proves the universality for this class of arbitrary functions.

Theorem 2 (Universality of the (ϕ, ψ, μ) -sharpness measures for arbitrary functions of Hessian). *For any continuous function $S : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}$, there exist a positive integer $m \leq d(d+1)/2$, (Borel) probability measures $\mu_\ell, \ell \in [m]$, and continuous functions $\phi : \mathbb{R}^m \rightarrow \mathbb{R}$ and $\psi : \mathbb{R} \rightarrow \mathbb{R}^m$, such that $S(\nabla^2 L(x)) = S(x; \phi, \psi, \mu)$ for any $x \in \mathbb{R}^d$, where $\mu := \mu_1 \otimes \mu_2 \otimes \dots \otimes \mu_m$ is a product probability measure.*

We present the proof of Theorem 2 in Appendix D.

Note that arbitrary functions of the Hessian matrix can be quite hard to compute, e.g., consider the permanent of the Hessian matrix. Moreover, the dimension m must be quite large to allow us to prove the universality in overparameterized models (for d of considerable size), since the generality bound scales as $\mathcal{O}(d^2)$. Nevertheless, in practice, only small m allows to cover many interesting cases.

6. Explicit Bias

Now that we defined a flexible set of sharpness measures and we proved that it is universally expressive, the following question arises: how can one achieve $S(x; \phi, \psi, \mu)$ as the explicit bias of an objective function that only relies on the zeroth-order information about the training loss, similar to $L_{\text{SAM}}(x)$ and $L_{\text{AVG}}(x)$? To answer this question, we introduce the (ϕ, ψ, μ) -sharpness-aware loss function as follows.

Definition 2. *The (ϕ, ψ, μ) -sharpness-aware loss function*

$$\begin{aligned}
L_{(\phi, \psi, \mu)}(x) &:= \underbrace{L(x)}_{\text{empirical loss}} \\
&+ \underbrace{\rho^2 \phi \left(\int \psi \left(\frac{1}{\rho^2} (L(x + \rho v) - L(x)) \right) d\mu(v) \right)}_{:= R_\rho(x) \text{ sharpness}} \\
&= L(x) + \rho^2 R_\rho(x),
\end{aligned}$$

where ρ is the perturbation parameter and $R_\rho(x)$ denotes the sharpness regularizer.

Extending this definition to the cases with $m > 1$ is straightforward.

In the above definition, the new regularizer $R_\rho(x)$ is an approximation of the sharpness measure $S(x; \phi, \psi, \mu)$ as $\rho \rightarrow 0^+$. As a result, it is expected that minimizing $L_{(\phi, \psi, \mu)}(x)$ lead to minimizing the training loss as well as the sharpness measure $S(x; \phi, \psi, \mu)$. The next theorem formalizes this intuitive observation via characterizing the explicit bias of minimizing the sharpness-aware loss function $L_{(\phi, \psi, \mu)}(x)$.

Theorem 3 (Explicit bias of the (ϕ, ψ, μ) -sharpness-aware loss function). *Given a triplet (ϕ, ψ, μ) , $m \geq 1$, and a training loss function $L : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$, assume that:*

- $L(x)$ is third-order continuously differentiable and satisfies the following upper bound

$$\max_{i, j, k \in \{1, 2, 3\}} |\partial_i \partial_j \partial_k L(v)| = \mathcal{O}(\|v\|^{-1}), \quad (7)$$

for $v \in \mathbb{R}^d$ as $\|v\|_2 \rightarrow \infty$.

- The two functions ϕ, ψ are continuously differentiable.
- For some $C > \max_{x \in \mathcal{X}} \max_{i \in [d]} |\lambda_i(\nabla^2 L(x))|$, we have $\int \|v\|_2^2 \tilde{\psi}_i(v) d\mu(v) < \infty$, $i \in [m]$, where

$$\tilde{\psi}_i(v) := \max_{|t| \leq \|v\|_2} |\psi'_i(Ct^2)|. \quad (8)$$

Then, there exists an open neighborhood $U \supseteq \Gamma$, where Γ is the zero-loss manifold, for connected U and Γ , such that if for some $u \in U$, one has

$$L(u) + \rho^2 R_\rho(u) - \inf_{x \in U} (L(x) + \rho^2 R_\rho(x)) \leq \Delta \rho^2, \quad (9)$$

with some optimality gap $\Delta > 0$, then

$$L(u) \leq \underbrace{\inf_{x \in U} L(x)}_{=0} + (\Delta + o_\rho(1))\rho^2, \quad (10)$$

and also

$$S(u; \phi, \psi, \mu) \leq \inf_{x \in \Gamma} S(x; \phi, \psi, \mu) + \Delta + o_\rho(1). \quad (11)$$

We present the proof of Theorem 3 in Appendix E.

The above theorem shows how using the new objective function $L_{(\phi, \psi, \mu)}(x)$ leads to explicitly biased optimization algorithms towards minimizing the sharpness measure $S(x; \phi, \psi, \mu)$ over the zero-loss manifold Γ . Indeed, it proves that if we are close to the zero-loss manifold (i.e., $u \in U$ for some open neighborhood $U \supseteq \Gamma$), and also $L_{(\phi, \psi, \mu)}(u)$ is close to its global minimum over U , then (1) the training loss function $L(u)$ is close to zero, and (2) the corresponding sharpness measure $S(u; \phi, \psi, \mu)$ is close to its global minimum over the zero-loss manifold, with respect to an optimality gap Δ .

7. Invariant Sharpness-Aware Minimization

For which hyperparameters (ϕ, ψ, μ) is the corresponding sharpness measure scale-invariant? The following theorem answers this question.

Theorem 4 (Scale-invariant (ϕ, ψ, μ) -sharpness measures). *Consider a scale-invariant loss function $L(x)$ and let μ be a Borel measure of the form*

$$d\mu(x) = f\left(\prod_{i=1}^d x_i\right) \prod_{i=1}^d dx_i, \quad (12)$$

where $f : \mathbb{R} \rightarrow \mathbb{R}$ is a measurable function⁴. Then, for any continuous functions ϕ, ψ , the corresponding sharpness

⁴In Lemma 1, we show that any scale-invariant measure is of this form.

measure $S(x; \phi, \psi, \mu)$ is scale-invariant; this means that $S(x; \phi, \psi, \mu) = S(Dx; \phi, \psi, \mu)$ for any diagonal matrix $D \in \mathbb{R}^{d \times d}$ with $\det(D) = 1$.

We present the proof of Theorem 4 in Appendix F.

Example 3. Note that $\det(\nabla^2 L(x))$ is a scale-invariant sharpness measure; for any diagonal matrix $D \in \mathbb{R}^{d \times d}$ with $\det(D) = 1$,

$$\begin{aligned} \det(\nabla^2 L(x)) \Big|_{Dx} &= \det(D^{-1} \nabla^2 L(x) D^{-1}) \\ &= \det(D^{-1})^2 \det(\nabla^2 L(x)) = \det(\nabla^2 L(x)). \end{aligned}$$

Note that Theorem 4 also supports the scale-invariance of the determinant; the Lebesgue measure satisfies the condition in Theorem 4 with $f \equiv 1$, and we have the representation of the determinant in Table 1.

While in Theorem 4 we only considered scale-invariances, one can generalize it to a general class of parameter invariances in the following theorem.

Theorem 5 (General parameter-invariant (ϕ, ψ, μ) -sharpness measures). *Let G be a group acting by matrices on \mathbb{R}^d , and assume that $L(x)$ is invariant with respect to the action of G . Then, for any G -invariant (Borel) measure μ , and any continuous functions ϕ, ψ , the corresponding sharpness measure $S(x; \phi, \psi, \mu)$ is G -invariant; this means that $S(x; \phi, \psi, \mu) = S(A_g x; \phi, \psi, \mu)$ for any matrix $A_g \in \mathbb{R}^{d \times d}$ corresponding to the action of an element $g \in G$.*

The proof of this theorem is analogous to Theorem 4 and is deferred to Appendix G. Thus, the strategy to create sharpness measures invariant to any group action G is simply to choose a group action invariant measure μ . Now, for any family of choices of functions ϕ and ψ , we obtain a family of G -invariant sharpness measures. Consequently, there is a family of G -invariant Sharpness-Aware Minimization algorithms, as explained in Section 8.

8. (ϕ, ψ, μ) -Sharpness-Aware Minimization Algorithm

In this section, we present the pseudocode for the (ϕ, ψ, μ) -Sharpness-Aware Minimization Algorithm (see Algorithm 1). For simplicity, we present the algorithm for the full-batch gradient descent, and assume that $m = 1$. Extending it to the mini-batch case with $m > 1$ is straightforward (see Algorithm 3). The idea is to apply (stochastic) gradient decent or other optimization algorithms on the (ϕ, ψ, μ) -sharpness-aware loss function defined in Definition 2,

$$L_{(\phi, \psi, \mu)} = L(x) + \rho^2 R_\rho(x).$$

Algorithm 1 (ϕ, ψ, μ) -Sharpness-Aware Minimization Algorithm (with $m = 1$)

Input: The triplet (ϕ, ψ, μ) , Training loss $L(x)$, Step size η , Perturbation parameter ρ , Number of samples n ,

Output: Model parameters x_t trained with (ϕ, ψ, μ) -Sharpness-Aware Minimization Algorithm

Initialization: $x \leftarrow x_0$ and $t \leftarrow 0$

while 1 do

 Sample $v_1, v_2, \dots, v_n \stackrel{\text{i.i.d.}}{\sim} \mu$

 Compute the following:

$$g_t = \nabla L(x_t) + \phi' \left(\sum_{i=1}^n \frac{1}{n} \psi \left(\frac{1}{\rho^2} (L(x_t + \rho v_i) - L(x_t)) \right) \right) \\ \times \sum_{i=1}^n \frac{1}{n} \left\{ \psi' \left(\frac{1}{\rho^2} (L(x_t + \rho v_i) - L(x_t)) \right) \times \left(\nabla L(x_t + \rho v_i) - \nabla L(x_t) \right) \right\}.$$

 Update the parameters: $x_{t+1} = x_t - \eta g_t$

$t \leftarrow t + 1$

end while

However, calculating the sharpness term $R_\rho(x)$ directly is analytically hard to do because of the integration with respect to the probability measure μ . Hence, we propose to estimate the inner integration at each iteration with i.i.d. random variables $\nu_1, \nu_2, \dots, \nu_n \sim \mu$ as perturbations, i.e.,

$$\tilde{R}_\rho(x) := \phi \left(\frac{1}{n} \sum_{i=1}^n \psi \left(\frac{1}{\rho^2} (L(x + \rho v_i) - L(x)) \right) \right).$$

When ϕ satisfies continuity conditions, for large enough n , the estimator $\tilde{R}_\rho(x)$ will converge to $R_\rho(x)$. Now, we calculate the gradients of $L(x) + \rho^2 \tilde{R}_\rho(x)$. By chain rule,

$$\rho^2 \nabla \tilde{R}_\rho(x) = \phi' \left(\sum_{i=1}^n \frac{1}{n} \psi \left(\frac{1}{\rho^2} (L(x_t + \rho v_i) - L(x_t)) \right) \right) \\ \times \sum_{i=1}^n \frac{1}{n} \left\{ \psi' \left(\frac{1}{\rho^2} (L(x_t + \rho v_i) - L(x_t)) \right) \right. \\ \left. \times \left(\nabla L(x_t + \rho v_i) - \nabla L(x_t) \right) \right\},$$

which leads to Algorithm 1.

Our algorithm needs $n + 1$ gradient evaluations per iteration, which for $n = 1$ matches the SAM algorithm (Foret et al., 2021). In practice, small values for n demonstrate the expected results, therefore, the computational overhead of our algorithm is not a barrier.

Note that to recover the original SAM algorithm, one can set the function ϕ, ψ to identity, $m = 1$, and choose μ to be the single-point measure on $\nabla L(x_t) / \|\nabla L(x_t)\|_2$ with $n = 1$ sample for each t .

Moreover, even though to prove universality, we only used probability measures, we proposed a compact representation of determinant with Lebesgue measure with $m = 1$ in Table 1 and Theorem 4. However, integrals with respect to Lebesgue measure cannot be estimated via sampling and we need to truncate the integral to integration over a large hypercube; this allows us to use Algorithm 1 for the scale-invariant sharpness measures. Also, this approximation

achieves non-zero sharpness in cases that the Hessian matrix is not full-rank (which happens in overparametrized models), as it gets the product of non-zero eigenvalues. We use this approximation in the next section to implement the method.

9. Frobenius SAM and Determinant SAM

To be more concrete, we specify Algorithm 3 (for arbitrary m) to the case with the Frobenius norm regularizes (with $m = 2$), which we call the *Frob-SAM* algorithm. Note that to achieve this, one needs to specify $\phi(t_1, t_2) = 2(t_2 - t_1^2)$ and $\psi(t) = (t, t^2)$. Furthermore, since we only need to collect samples from the Gaussian distribution to get the Frobenius norm bias (see Table 1), we can use the same samples to estimate both integrals for the functions $\psi_1(t) = t$ and $\psi_2(t) = t^2$. Replacing these assumptions into the formula given in Algorithm 3, we get the following update rule:

$$g_t = \nabla L(x_t) \\ + 4 \sum_{i=1}^n \frac{1}{n \rho^2} \left\{ (L(x_t + \rho v_i) - L(x_t)) \right. \\ \left. \times \left(\nabla L(x_t + \rho v_i) - \nabla L(x_t) \right) \right\} \\ - 4 \left\{ \sum_{i=1}^n \frac{1}{n \rho} (L(x_t + \rho v_i) - L(x_t)) \right\} \\ \times \left\{ \sum_{j=1}^n \frac{1}{n \rho} \left(\nabla L(x_t + \rho v_j) - \nabla L(x_t) \right) \right\}.$$

If we take a closer look at this, we observe that

$$g_t = \nabla L(x_t) + \frac{4}{\rho^2} \widehat{\text{cov}} \left((L(x_t + \rho v) - L(x_t)), \right. \\ \left. (\nabla L(x_t + \rho v) - \nabla L(x_t)) \right),$$

where $\widehat{\text{cov}}$ denotes the (biased) empirical cross-covariance between the scalar random variable $L(x_t + \rho v) - L(x_t)$ and

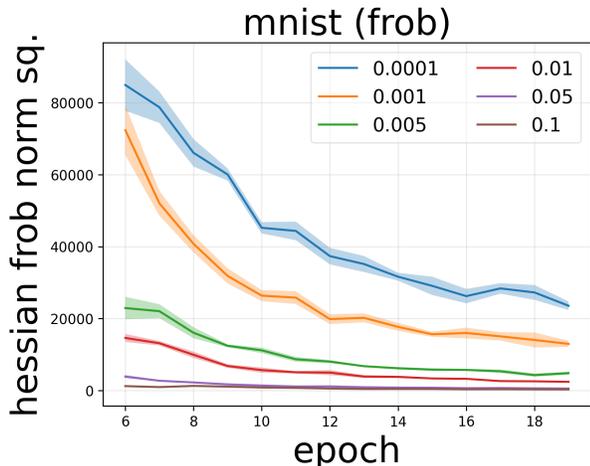


Figure 2. Sharpness measure while training on MNIST for different amounts of regularization λ . When training with Frob-SAM ($\rho = 0.01$, $n = 2$), the Frobenius norm of the Hessian decreases throughout training and larger λ results in lower values of the norm. One standard error bar is shaded. All λ 's achieve over 94% final test accuracy.

the vector-valued random variable $\nabla L(x_t + \rho v) - \nabla L(x_t)$, for $v \sim \mathcal{N}(0, I_d)$. Since the covariance is not sensitive to the means of random variables/vectors, we can further simplify the update rule to

$$g_t = \nabla L(x_t) + \frac{4}{\rho^2} \widehat{\text{cov}}\left(L(x_t + \rho v), \nabla L(x_t + \rho v)\right).$$

We can further replace the unbiased estimator of the cross-covariance instead of $\widehat{\text{cov}}$ which leads to Algorithm 2.

Moreover, achieving *Det-SAM* is also similar to *Frob-SAM*, but the only difficulty is that it involves computing an integral with respect to the Lebesgue measure which can be challenging (Table 1). To address this issue, we instead sample a point from the hypercube $[-t, t]^d$ for a hyperparameter $t \in \mathbb{R}$ to approximate the Lebesgue measure.

10. Experiments

The goal of our experiments is twofold. Firstly, we validate Theorem 3 by showing that minimization of the sharpness-aware loss defined in Definition 2 and codified in Algorithm 1 has the explicit bias of minimizing the sharpness measure. Secondly, we show our method is useful in practice by evaluating it on benchmark vision tasks. Our code is available at https://github.com/dbahri/universal_sam.

10.1. Setup

We evaluate on three vision datasets: CIFAR10, CIFAR100, and SVHN. Furthermore, we study how our *explicit* bias may be helpful in settings that generally benefit from regularization – specifically, when *training data is limited* and when *training labels are noisy*. For the former, we artificially sub-sample each original dataset, keeping only the first 10% of training samples, and we denote these sub-sampled datasets with “-S” (i.e. CIFAR10-S). For the latter, we choose a random 20% of training samples to corrupt, and we corrupt these samples by flipping their labels to a different label chosen uniformly at random over the remaining classes. We denote these datasets by “-C” (i.e. CIFAR10-C).

Full experimental details are deferred to Appendix I; we summarize them here. We train ResNet18 (He et al., 2016) on the datasets using momentum-SGD and a multi-step learning rate schedule. We run each experiment under four different random seeds. We evaluate three sharpness measures and the following other baselines:

Frob-SAM / Trace-SAM / Det-SAM. These correspond to instances of our algorithm where the measure is, respectively, the Frobenius norm, trace, and determinant of the Hessian. For Det-SAM, we set t , half the edge width of the approximating hypercube, to 0.01 with little tuning.

SAM. We set ρ to 0.05 / 0.1 / 0.05 for CIFAR10 / CIFAR100 / SVHN, following Foret et al. (2021).

Adaptive SAM (ASAM). Kwon et al. (2021) proposes a modification of SAM that is scale-invariant. We set ρ to 0.5 / 1 / 0.5 and η to 0.01 / 0.1 / 0.01 for CIFAR10 / CIFAR100 / SVHN.

Sparse SAM (SSAM). Mi et al. (2022) speeds up and improves the performance of SAM by only perturbing important parameters, as determined via Fisher information and sparse dynamic training. We use SSAM-F with ρ set to 0.1 / 0.2 / 0.1 for CIFAR10 / CIFAR100 / SVHN. 50% sparsity is used with 16 samples.

Furthermore, to demonstrate that minimization of our loss has the explicit bias we expect, we train a simple 6-layer ReLU network with 128 hidden units on MNIST using momentum-SGD (with momentum 0.9 and learning rate 0.001) for 20 epochs. We estimate the Frobenius norm of the Hessian (see Appendix I for details) throughout training for Frob-SAM, under a range of regularization strengths λ .

10.2. Results

MNIST results are shown in Figure 2. We see that for Frob-SAM, minimization of the loss leads to reduction in the corresponding sharpness measure and the reduction expectedly scales proportionally with regularization strength λ .

Algorithm 2 *Frob-SAM*

Input: Training loss $L(x)$, Step size η , Perturbation parameter ρ , Number of samples n ,

Output: Model parameters x_t trained with Frobenius SAM

Initialization: $x \leftarrow x_0$ and $t \leftarrow 0$

while 1 do

Sample $v_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_d)$ for any $i \in [n]$

Compute the following:

$$g_t = \nabla L(x_t) + 4 \sum_{i=1}^n \frac{1}{(n-1)\rho^2} L(x_t + \rho v_i) \nabla L(x_t + \rho v_i) - 4 \sum_{i=1}^n \frac{1}{(n-1)\rho} L(x_t + \rho v_i) \times \sum_{i=1}^n \frac{1}{n\rho} \nabla L(x_t + \rho v_i).$$

Update the parameters: $x_{t+1} = x_t - \eta g_t$

$t \leftarrow t + 1$

end while

	CIFAR10	CIFAR100	SVHN	CIFAR10-S	CIFAR100-S	SVHN-S
Frob	94.96 \pm 0.05	77.16 \pm 0.16	96.26 \pm 0.05	74.49 \pm 1.15	38.30 \pm 0.89	89.23 \pm 0.32
Trace	95.04 \pm 0.07	77.61 \pm 0.08	96.37 \pm 0.04	74.46 \pm 0.39	37.50 \pm 0.78	89.90 \pm 0.32
Det	95.10 \pm 0.03	77.64 \pm 0.12	96.32 \pm 0.05	74.22 \pm 0.80	37.48 \pm 0.42	89.85 \pm 0.13
SSAM	95.78 \pm 0.05	78.42 \pm 0.19	96.49 \pm 0.04	73.39 \pm 0.48	35.33 \pm 0.55	89.87 \pm 0.27
ASAM	95.49 \pm 0.09	78.91 \pm 0.06	96.18 \pm 0.05	74.80 \pm 0.64	37.40 \pm 0.66	89.21 \pm 0.11
SAM	95.54 \pm 0.06	78.62 \pm 0.02	96.45 \pm 0.03	72.96 \pm 0.77	36.87 \pm 0.47	89.80 \pm 0.19
SGD	94.81 \pm 0.07	77.00 \pm 0.11	96.06 \pm 0.08	74.98 \pm 0.92	37.64 \pm 0.88	89.89 \pm 0.27

Table 2. Final test accuracy and standard errors for the full and 10% sub-sampled datasets. Entries within two standard errors of the best are bolded. Our biases particularly help when data is limited. For example, Frob-SAM achieves over 1.5% higher test accuracy than SAM on CIFAR10-S.

	CIFAR10-C	CIFAR100-C	SVHN-C
Frob	87.57 \pm 0.10	65.73 \pm 0.22	90.11 \pm 0.20
Trace	87.20 \pm 0.12	65.29 \pm 0.12	91.01 \pm 0.44
Det	83.61 \pm 0.06	64.73 \pm 0.29	90.56 \pm 0.09
SSAM	86.24 \pm 0.08	67.06 \pm 0.14	90.92 \pm 0.12
ASAM	85.88 \pm 0.12	67.78 \pm 0.04	89.54 \pm 0.12
SAM	85.42 \pm 0.08	66.12 \pm 0.13	90.66 \pm 0.16
SGD	82.78 \pm 0.20	64.16 \pm 0.22	87.78 \pm 0.16

Table 3. Final test accuracy and standard errors when the labels of 20% of the training examples are corrupted. Entries within two standard errors of the best are bolded. We see that our bias can provide a boost in this scenario. For example, Frob-SAM achieves over 2% higher accuracy than SAM on CIFAR10-C.

Results for the main datasets are shown in Table 2 and Table 3. We find that our method nearly always outperforms SGD and is at or sometimes above par with SAM and alternatives, especially in the noisy label and limited data settings. For example, for noisy label CIFAR10 (CIFAR10-C), Frob-SAM achieves nearly 5% and 2% higher final test accuracy than SGD and SAM respectively. Our findings suggest that the explicit biases we propose can be practically useful, especially in the noisy label and limited data scenarios, though it is often unclear what the *best* bias for a particular task is, a priori.

11. Conclusion

In this paper, we introduce a new family of sharpness measures and demonstrate how this parameterized representation can generate many meaningful sharpness notions (Table 1). These measures are indeed universally expressive. Furthermore, in Theorem 3, we illustrate how the corresponding zeroth-order objective function for each sharpness measure is explicitly biased towards minimizing the sharpness of the training loss. Moreover, in Theorem 4, we prove how specific (Borel) measures can lead to scale-invariant sharpness measures, such as the determinant of the Hessian matrix. We conclude the paper with a series of numerical experiments showcasing the efficacy of the proposed loss function on various practical datasets. Given the broad class of sharpness measures we proposed, an interesting future direction is to evaluate in practice which sharpness measure/algorithm performs the best for different dataset. Another interesting attitude is, given the universal expressivity of our algorithm, to meta-learn the sharpness measures on the top of the model. This could potentially yield the identification of optimal, intricate sharpness measures tailored to specific datasets, thereby contributing to unraveling hidden secrets of overparameterized models.

Acknowledgments

The authors appreciate Joshua Robinson for his insightful comments and valuable suggestions. BT and SJ are supported by the Office of Naval Research award N00014-20-1-2023 (MURI ML-SCOPE), NSF award CCF-2112665 (TILOS AI Institute), NSF award 2134108, and the Alexander von Humboldt Foundation. AS and PJ are supported by the National Research Foundation Singapore and DSO National Laboratories under the AI Singapore Programme (AISG Award No: AISG2-RP-2020-018).

Impact Statement

This paper introduces a universal class of sharpness measures with associated algorithms. While these algorithms may improve the generalization of neural networks, the primary aim is to enhance our understanding of the loss landscape, investigate the impact of overparameterization, and clarify the reasons behind the generalization performance of overparameterized models. Consequently, no direct social or ethical concerns are anticipated in relation to this paper.

References

- Agarwala, A. and Dauphin, Y. SAM operates far from home: eigenvalue regularization as a dynamical phenomenon. In *Int. Conference on Machine Learning (ICML)*, 2023.
- Andriushchenko, M. and Flammarion, N. Towards understanding sharpness-aware minimization. In *Int. Conference on Machine Learning (ICML)*, 2022.
- Andriushchenko, M., Bahri, D., Mobahi, H., and Flammarion, N. Sharpness-aware minimization leads to low-rank features. *arXiv preprint arXiv:2305.16292*, 2023a.
- Andriushchenko, M., Croce, F., Müller, M., Hein, M., and Flammarion, N. A modern look at the relationship between sharpness and generalization. In *Int. Conference on Machine Learning (ICML)*, 2023b.
- Arora, S., Li, Z., and Panigrahi, A. Understanding gradient descent on the edge of stability in deep learning. In *Int. Conference on Machine Learning (ICML)*, 2022.
- Azizan, N. and Hassibi, B. Stochastic gradient/mirror descent: Minimax optimality and implicit regularization. In *Int. Conference on Learning Representations (ICLR)*, 2019.
- Bahri, D., Mobahi, H., and Tay, Y. Sharpness-aware minimization improves language model generalization. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, 2022.
- Bartlett, P. L., Long, P. M., and Bousquet, O. The dynamics of sharpness-aware minimization: Bouncing across ravines and drifting towards wide minima. *arXiv preprint arXiv:2210.01513*, 2022.
- Behdin, K. and Mazumder, R. On statistical properties of sharpness-aware minimization: Provable guarantees. *arXiv preprint arXiv:2302.11836*, 2023.
- Behdin, K., Song, Q., Gupta, A., Durfee, D., Acharya, A., Keerthi, S., and Mazumder, R. Improved deep neural network generalization using m-sharpness-aware minimization. *arXiv preprint arXiv:2212.04343*, 2022.
- Blanc, G., Gupta, N., Valiant, G., and Valiant, P. Implicit regularization for deep neural networks driven by an ornstein-uhlenbeck like process. In *Conference on Learning Theory (COLT)*, 2020.
- Cha, J., Chun, S., Lee, K., Cho, H.-C., Park, S., Lee, Y., and Park, S. Swad: Domain generalization by seeking flat minima. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Chaudhari, P., Choromanska, A., Soatto, S., Lecun, Y., Baldassi, C., Borgs, C., Chayes, J., Sagun, L., and Zecchina, R. Entropy-sgd: Biasing gradient descent into wide valleys. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12):124018, 2019.
- Compagnoni, E. M., Biggio, L., Orvieto, A., Proske, F. N., Kersting, H., and Lucchi, A. An SDE for modeling sam: Theory and insights. In *Int. Conference on Machine Learning (ICML)*, 2023.
- Damian, A., Ma, T., and Lee, J. D. Label noise SGD provably prefers flat global minimizers. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Dinh, L., Pascanu, R., Bengio, S., and Bengio, Y. Sharp minima can generalize for deep nets. In *Int. Conference on Machine Learning (ICML)*, 2017.
- Du, J., Zhou, D., Feng, J., Tan, V., and Zhou, J. T. Sharpness-aware training for free. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- Foret, P., Kleiner, A., Mobahi, H., and Neyshabur, B. Sharpness-aware minimization for efficiently improving generalization. In *Int. Conference on Learning Representations (ICLR)*, 2021.
- Gunasekar, S., Lee, J., Soudry, D., and Srebro, N. Characterizing implicit bias in terms of optimization geometry. In *Int. Conference on Machine Learning (ICML)*, 2018a.
- Gunasekar, S., Lee, J. D., Soudry, D., and Srebro, N. Implicit bias of gradient descent on linear convolutional

- networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018b.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Hochreiter, S. and Schmidhuber, J. Flat minima. *Neural computation*, 9(1):1–42, 1997.
- Izmailov, P., Podoprikin, D., Gariipov, T., Vetrov, D., and Wilson, A. G. Averaging weights leads to wider optima and better generalization. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2018.
- Jang, C., Lee, S., Park, F., and Noh, Y.-K. A reparametrization-invariant sharpness measure based on information geometry. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- Ji, Z. and Telgarsky, M. Gradient descent aligns the layers of deep linear networks. In *Int. Conference on Learning Representations (ICLR)*, 2019a.
- Ji, Z. and Telgarsky, M. The implicit bias of gradient descent on nonseparable data. In *Conference on Learning Theory (COLT)*, 2019b.
- Jiang, Y., Neyshabur, B., Mobahi, H., Krishnan, D., and Bengio, S. Fantastic generalization measures and where to find them. In *Int. Conference on Learning Representations (ICLR)*, 2020.
- Kaddour, J., Liu, L., Silva, R., and Kusner, M. J. When do flat minima optimizers work? In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- Keskar, N. S., Mudigere, D., Nocedal, J., Smelyanskiy, M., and Tang, P. T. P. On large-batch training for deep learning: Generalization gap and sharp minima. In *Int. Conference on Learning Representations (ICLR)*, 2017.
- Kim, H., Park, J., Choi, Y., Lee, W., and Lee, J. Exploring the effect of multi-step ascent in sharpness-aware minimization. *arXiv preprint arXiv:2302.10181*, 2023.
- Kim, M., Li, D., Hu, S. X., and Hospedales, T. Fisher sam: Information geometry and sharpness aware minimisation. In *Int. Conference on Machine Learning (ICML)*, 2022.
- Kwon, J., Kim, J., Park, H., and Choi, I. K. ASAM: Adaptive sharpness-aware minimization for scale-invariant learning of deep neural networks. In *Int. Conference on Machine Learning (ICML)*, 2021.
- Lawrence, H., Georgiev, K., Dienes, A., and Kiani, B. T. Implicit bias of linear equivariant networks. In *Int. Conference on Machine Learning (ICML)*, 2022.
- Le, T. and Jegelka, S. Training invariances and the low-rank phenomenon: beyond linear networks. In *Int. Conference on Learning Representations (ICLR)*, 2022.
- Li, Z., Wang, T., and Arora, S. What happens after sgd reaches zero loss?—a mathematical framework. In *Int. Conference on Learning Representations (ICLR)*, 2022a.
- Li, Z., Wang, T., and Yu, D. Fast mixing of stochastic gradient descent with normalization and weight decay. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022b.
- Liang, T., Poggio, T., Rakhlin, A., and Stokes, J. Fisher-rao metric, geometry, and complexity of neural networks. In *Int. Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019.
- Lim, D., Robinson, J., Zhao, L., Smidt, T., Sra, S., Maron, H., and Jegelka, S. Sign and basis invariant networks for spectral graph representation learning. In *Int. Conference on Learning Representations (ICLR)*, 2023.
- Liu, Y., Mai, S., Chen, X., Hsieh, C.-J., and You, Y. Towards efficient and scalable sharpness-aware minimization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022a.
- Liu, Y., Mai, S., Cheng, M., Chen, X., Hsieh, C.-J., and You, Y. Random sharpness-aware minimization. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022b.
- Long, P. M. and Bartlett, P. L. Sharpness-aware minimization and the edge of stability. *arXiv preprint arXiv:2309.12488*, 2023.
- Lu, P., Kobzyev, I., Rezagholizadeh, M., Rashid, A., Ghodsi, A., and Langlais, P. Improving generalization of pre-trained language models via stochastic weight averaging. *arXiv preprint arXiv:2212.05956*, 2022.
- Lyu, K., Li, Z., and Arora, S. Understanding the generalization benefit of normalization layers: Sharpness reduction. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. In *Int. Conference on Learning Representations (ICLR)*, 2018.
- Mi, P., Shen, L., Ren, T., Zhou, Y., Sun, X., Ji, R., and Tao, D. Make sharpness-aware minimization stronger: A sparsified perturbation approach. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

- Neyshabur, B., Bhojanapalli, S., McAllester, D., and Srebro, N. Exploring generalization in deep learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- Nitanda, A., Kikuchi, R., and Maeda, S. Parameter averaging for sgd stabilizes the implicit bias towards flat regions. *arXiv preprint arXiv:2302.09376*, 2023.
- Qu, Z., Li, X., Duan, R., Liu, Y., Tang, B., and Lu, Z. Generalized federated learning via sharpness aware minimization. In *Int. Conference on Machine Learning (ICML)*, 2022.
- Shi, Y., Liu, Y., Wei, K., Shen, L., Wang, X., and Tao, D. Make landscape flatter in differentially private federated learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- Soudry, D., Hoffer, E., Nacson, M. S., Gunasekar, S., and Srebro, N. The implicit bias of gradient descent on separable data. *Journal of Machine Learning Research*, 2018.
- Sun, H., Shen, L., Zhong, Q., Ding, L., Chen, S., Sun, J., Li, J., Sun, G., and Tao, D. AdaSAM: Boosting sharpness-aware minimization with adaptive learning rate and momentum for training deep neural networks. *arXiv preprint arXiv:2303.00565*, 2023.
- Wang, P., Zhang, Z., Lei, Z., and Zhang, L. Sharpness-aware gradient matching for domain generalization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- Wen, K., Ma, T., and Li, Z. How does sharpness-aware minimization minimize sharpness? In *Int. Conference on Learning Representations (ICLR)*, 2023a.
- Wen, K., Ma, T., and Li, Z. Sharpness minimization algorithms do not only minimize sharpness to achieve better generalization. *arXiv preprint arXiv:2307.11007*, 2023b.
- Woodworth, B., Gunasekar, S., Lee, J. D., Moroshko, E., Savarese, P., Golan, I., Soudry, D., and Srebro, N. Kernel and rich regimes in overparametrized models. In *Conference on Learning Theory (COLT)*, 2020.
- Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How powerful are graph neural networks? In *Int. Conference on Learning Representations (ICLR)*, 2019.
- Yao, Z., Gholami, A., Keutzer, K., and Mahoney, M. W. Pyhessian: Neural networks through the lens of the hessian. In *2020 IEEE international conference on big data (Big data)*, pp. 581–590. IEEE, 2020.
- Zaheer, M., Kottur, S., Ravanbakhsh, S., Póczos, B., Salakhutdinov, R. R., and Smola, A. J. Deep sets. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- Zhao, Y., Zhang, H., and Hu, X. Randomized sharpness-aware training for boosting computational efficiency in deep learning. *arXiv preprint arXiv:2203.09962*, 2022.
- Zhong, Q., Ding, L., Shen, L., Mi, P., Liu, J., Du, B., and Tao, D. Improving sharpness-aware minimization with fisher mask for better generalization on language models. *arXiv preprint arXiv:2210.05497*, 2022.
- Zhu, X., Wang, Z., Wang, X., Zhou, M., and Ge, R. Understanding edge-of-stability training dynamics with a minimalist example. In *Int. Conference on Learning Representations (ICLR)*, 2023.
- Zhuang, J., Gong, B., Yuan, L., Cui, Y., Adam, H., Dvornik, N., Tatikonda, S., Duncan, J., and Liu, T. Surrogate gap minimization improves sharpness-aware training. In *Int. Conference on Learning Representations (ICLR)*, 2022.

A. Additional Related Work

SAM can provide a strong regularization of the eigenvalues throughout the learning trajectory (Agarwala & Dauphin, 2023). Bartlett et al. (2022) show that the dynamics of SAM similar to GD on the spectral norm of Hessian. Compagnoni et al. (2023) propose an SDE for modeling SAM, while Behdin & Mazumder (2023) study the statistical benefits of SAM (see also (Li et al., 2022a) for a general framework for the dynamics of SGD around the zero-loss manifold). It is shown that SAM can reduce the feature rank (i.e., allowing learning low-rank features) (Andriushchenko et al., 2023a). Blanc et al. (2020) proved that SGD is implicitly biased toward minimizing the trace of Hessian.

Kim et al. (2023) proposes a multi-step ascent approach to improve SAM, while Mi et al. (2022) suggested sparsification of SAM. Zhuang et al. (2022) improve SAM by changing the directions in the ascent step; their method is called Surrogate Gap Guided Sharpness-Aware Minimization (GSAM) (see also (Behdin et al., 2022)). Random smoothing-based SAM (R-SAM) is another SAM variant that is proposed to reduce its computational complexity (Liu et al., 2022b) (see also (Du et al., 2022; Liu et al., 2022a; Zhao et al., 2022; Sun et al., 2023) for more). Adaptive SAM (ASAM) is proposed for applying SAM on scale-invariant neural networks and has shown generalization benefits (Kwon et al., 2021). Li et al. (2022b) also prove that scale-invariant loss functions allow faster mixing in function spaces for neural networks. Lyu et al. (2022) show how normalization can make GD reduce the sharpness via a continuous sharpness-reduction flow. Liang et al. (2019) propose a capacity measure based on information geometry for parameter invariances in overparameterized models (for more on information geometry, see (Kim et al., 2022; Jang et al., 2022)). Jiang et al. (2020) empirically compare different complexity measures for overparameterized models. Keskar et al. (2017) show how a large batch yields sharp minima but a small batch achieves flat minima.

Stochastic Weight Averaging (SWA) is another way to improve generalization and it relies on finding wider minima by averaging multiple points along the trajectory of SGD (Izmailov et al., 2018). (see e.g., (Lu et al., 2022) which uses this method for language models). See (Kaddour et al., 2022) for the empirical comparison between two popular flat-minima optimization approaches: SWA and SAM.

Learning with group invariant architectures has recently gained a lot of interest due to its applications in physics and biology; see e.g., deep sets (Zaheer et al., 2017), Graph Neural Networks (GNNs) (Xu et al., 2019), and also sign-flips for spectral data (Lim et al., 2023). These architectures are all owing their practical success to their specific parameter invariance.

It is also worth mentioning that neural networks trained with large learning rates often generalize better (the edge-of-stability regime); see (Arora et al., 2022; Zhu et al., 2023; Long & Bartlett, 2023) for the theoretical understanding of this phenomenon.

B. Examples of (ϕ, ψ, μ) -Sharpness Measures

In this section, we prove various notions of sharpness can be achieved using the proposed approach in this paper (Table 1). For the last row of Table 1, we refer the reader to the proof of Theorem 1.

- **Trace.** Let $\phi(t) = \psi(t) = t$, and note that

$$S(x; \phi, \psi, \mu) = \int \frac{1}{2} v^t \nabla^2 L(x) v d\mu(v) \quad (13)$$

$$= \frac{1}{2} \mathbb{E}_{v \sim \mu} [v^t \nabla^2 L(x) v], \quad (14)$$

where μ is the uniform distribution over the $(d-1)$ -sphere $\mathbb{S}^{(d-1)} := \{x \in \mathbb{R}^d : \|x\|_2 = 1\}$. Denote the entries of $\nabla^2 L(x)$ as $(\nabla^2 L(x))_{i,j}$. Then, by the linearity of expectation

$$\mathbb{E}_{v \sim \mu} [v^t \nabla^2 L(x) v] = \sum_{i,j=1}^d (\nabla^2 L(x))_{i,j} \mathbb{E}[v_i v_j] = \sum_{i=1}^d \frac{1}{d} (\nabla^2 L(x))_{i,i} = \frac{1}{d} \text{tr}(\nabla^2 L(x)), \quad (15)$$

since $\mathbb{E}[v_i v_j] = \frac{1}{d} \delta_{i,j}$, where $\delta_{i,j}$ denotes the Kronecker delta function.

- **Determinant.** To achieve the determinant, we choose $\phi(t) = (2\pi)^d / t^2$ and $\psi(t) = \exp(-t)$. Then,

$$S(x; \phi, \psi, \mu) = (2\pi)^d \left(\int \exp\left(-\frac{1}{2} v^t \nabla^2 L(x) v\right) dv \right)^{-2}, \quad (16)$$

where dv denotes the Lebesgue measure. However, using the multivariate Gaussian integral, we have

$$\int \exp\left(-\frac{1}{2}v^t \nabla^2 L(x)v\right) dv = (2\pi)^{d/2} \det(\nabla^2 L(x))^{-1/2}. \quad (17)$$

Replacing this into the definition of $S(x; \phi, \psi, \mu)$ gives the desired result.

- **Polynomials of eigenvalues.** First assume that $\psi(t) = t^n$ for some $n \geq 0$. Then, for any function $\phi(t)$,

$$S(x; \phi, \psi, \mu) = \phi\left(\int \left(\frac{1}{2}v^t \nabla^2 L(x)v\right)^n d\mu(v)\right) \quad (18)$$

$$= \phi\left(\mathbb{E}_{v \sim \mu}\left[\left(\frac{1}{2}v^t \nabla^2 L(x)v\right)^n\right]\right), \quad (19)$$

where μ is the uniform distribution over the $(d-1)$ -sphere $\mathbb{S}^{(d-1)}$. Since $\nabla^2 L(x)$ is a symmetric matrix, we can find an orthogonal matrix Q such that $\nabla^2 L(x) = Q^t D Q$, where D is a diagonal matrix with diagonal entries $\lambda_1, \lambda_2, \dots, \lambda_d$. Now we write $(v^t \nabla^2 L(x)v)^n = (v^t Q^t D Q v)^n$. But Qv is distributed uniformly over the $(d-1)$ -sphere $\mathbb{S}^{(d-1)}$, similar to v . Thus, we conclude

$$S(x; \phi, \psi, \mu) = \phi\left(\mathbb{E}_{v \sim \mu}\left[\left(\frac{1}{2}v^t \nabla^2 L(x)v\right)^n\right]\right) \quad (20)$$

$$= \phi\left(\mathbb{E}_{v \sim \mu}\left[\left(\frac{1}{2}\sum_{i=1}^d \lambda_i v_i^2\right)^n\right]\right). \quad (21)$$

Define

$$q(\lambda_1, \lambda_2, \dots, \lambda_d) := \mathbb{E}_{v \sim \mu}\left[\left(\sum_{i=1}^d \frac{1}{2}\lambda_i v_i^2\right)^n\right], \quad (22)$$

which is clearly a polynomial function (by the linearity of expectation).

Note that the above computation is still valid if we replace the uniform distribution on hypersphere with the Gaussian multivariate distribution with identity covariance $\mathcal{N}(0, I_d)$. Indeed, let us compute this polynomial for $n = 2$ with Gaussian distribution. Note that

$$q(\lambda_1, \lambda_2, \dots, \lambda_d) = \mathbb{E}_{v \sim \mu}\left[\left(\sum_{i=1}^d \frac{1}{2}\lambda_i v_i^2\right)^2\right] = \frac{1}{4}\sum_{i=1}^d \lambda_i^2 \mathbb{E}[Z^4] + \sum_{i \neq j} \frac{1}{4}\lambda_i \lambda_j (\mathbb{E}[Z^2])^2 \quad (23)$$

$$= \frac{3}{4}\sum_{i=1}^d \lambda_i^2 + \frac{1}{4}\sum_{i \neq j} \lambda_i \lambda_j, \quad (24)$$

where Z is a zero-mean Gaussian random variable with unit variance, and note that $E[Z^2] = 1$ and $E[Z^4] = 3$.

Now if we take $m = 2$, and $\psi(t) = (t, t^2)$, with $\mu = \mathcal{N}(0, I_d) \otimes \mathcal{N}(0, I_d)$, we have that

$$\int \psi\left(\frac{1}{2}v^t \nabla^2 L(x)v\right) d\mu(v) = \left(\frac{1}{2}\sum_{i=1}^d \lambda_i, \frac{3}{4}\sum_{i=1}^d \lambda_i^2 + \frac{1}{4}\sum_{i \neq j} \lambda_i \lambda_j\right). \quad (25)$$

Finally, by taking $\phi(t_1, t_2) = 2(t_2 - t_1^2)$, we obtain

$$\phi\left(\int \psi\left(\frac{1}{2}v^t \nabla^2 L(x)v\right) d\mu(v)\right) = \sum_{i=1}^d \lambda_i^2. \quad (26)$$

C. Proof of Theorem 1

Theorem 1 (Universality of the (ϕ, ψ, μ) -sharpness measures for functions of Hessian eigenvalues). *Let $\mathcal{A} \subseteq \mathbb{R}^d$ be a compact set. For any continuous function $S : \mathcal{A} \rightarrow \mathbb{R}$, there exist a product (Borel) probability measure μ , a positive integer $m \leq d$, and continuous functions $\phi : \mathbb{R}^m \rightarrow \mathbb{R}$ and $\psi : \mathbb{R} \rightarrow \mathbb{R}^m$, such that $S(\lambda_1, \lambda_2, \dots, \lambda_d) = S(x; \phi, \psi, \mu)$ for any $x \in \mathcal{A}$, where $\lambda_i, i \in [d]$, are the eigenvalues of the Hessian matrix $\nabla^2 L(x)$.*

Proof. We explicitly construct the (Borel) probability measure μ and the function $\psi : \mathbb{R}^m \rightarrow \mathbb{R}$. Let us take $m = d$ to prove the universality theorem, while we believe lower m should be enough for specific practical sharpness measures.

Indeed, let us consider μ to be the multivariate Gaussian probability measure with identity covariance matrix. Also, define a (parameterized) function $\psi_\sigma(t) = \exp(\sigma t)$, for some σ to be set later. We are interested to compute the following quantity:

$$\int \psi_\sigma\left(\frac{1}{2}v^t \nabla^2 L(x)v\right) d\mu(v). \quad (27)$$

Note that μ is the standard Gaussian probability measure on \mathbb{R}^d , and specifically, it's invariant under the action of the orthogonal matrices. Indeed, there exists an orthogonal matrix Q such that $\nabla^2 L(x) = Q^t \Lambda Q$, where Λ is a diagonal matrix with diagonal entries $\lambda_1, \lambda_2, \dots, \lambda_d$. Observe that $u := Qv$ is also distributed according to the Gaussian probability distribution with the identity covariance matrix on \mathbb{R}^d , similar to v . Now we write

$$\int \psi_\sigma\left(\frac{1}{2}v^t \nabla^2 L(x)v\right) d\mu(v) = \int \psi_\sigma\left(\frac{1}{2}v^t Q^t \Lambda Q v\right) d\mu(v) \quad (28)$$

$$= \int \psi_\sigma\left(\frac{1}{2}u^t \Lambda u\right) d\mu(u) \quad (29)$$

$$= \int \psi_\sigma\left(\frac{1}{2} \sum_{i=1}^d \lambda_i u_i^2\right) d\mu(u) \quad (30)$$

$$= \int (2\pi)^{-d/2} \psi_\sigma\left(\frac{1}{2} \sum_{i=1}^d \lambda_i u_i^2\right) \exp\left(-\frac{1}{2} \sum_{i=1}^d u_i^2\right) du \quad (31)$$

$$= \int (2\pi)^{-d/2} \exp\left(\frac{1}{2} \sum_{i=1}^d \sigma \lambda_i u_i^2\right) \exp\left(-\frac{1}{2} \sum_{i=1}^d u_i^2\right) du \quad (32)$$

$$= \int (2\pi)^{-d/2} \exp\left(\sum_{i=1}^d \frac{1}{2}(\sigma \lambda_i - 1)u_i^2\right) du, \quad (33)$$

where du denotes the Lebesgue measure on \mathbb{R}^d . Now to compute the integral, note that $du = du_1 \times du_2 \times \dots \times du_d$ is a product measure and the integrand also takes on a product form; thus,

$$\int \psi_\sigma\left(\frac{1}{2}v^t \nabla^2 L(x)v\right) d\mu(v) = \int (2\pi)^{-d/2} \exp\left(\sum_{i=1}^d \frac{1}{2}(\sigma \lambda_i - 1)u_i^2\right) du \quad (34)$$

$$= \prod_{i=1}^d \int (2\pi)^{-1/2} \exp\left(\frac{1}{2}(\sigma \lambda_i - 1)u_i^2\right) du_i \quad (35)$$

$$\stackrel{(a)}{=} \prod_{i=1}^d \frac{1}{\sqrt{1 - \sigma \lambda_i}} \underbrace{\int \sqrt{\frac{1 - \sigma \lambda_i}{2\pi}} \exp\left(\frac{1}{2}(\sigma \lambda_i - 1)u_i^2\right) du_i}_{=1} \quad (36)$$

$$= \prod_{i=1}^d \frac{1}{\sqrt{1 - \sigma \lambda_i}}, \quad (37)$$

where (a) holds by the Gaussian integral identities.

Note that to calculate the integral above, we assumed that $1 - \sigma \lambda_i > 0$ for any $i \in [d]$. This is equivalent to having

$$\max_{\lambda_i < 0} \lambda_i^{-1} < \sigma < \min_{\lambda_i > 0} \lambda_i^{-1}. \quad (38)$$

Now, due to the assumption in the theorem, we study the target sharpness measure $S(\lambda_1, \lambda_2, \dots, \lambda_d)$ only on a compact domain $\mathcal{A} \subseteq \mathbb{R}^d$, and this means that there exists an open interval $I = (-\epsilon, \epsilon)$ with

$$\epsilon := \min_{(\lambda_1, \lambda_2, \dots, \lambda_d) \in \mathcal{A}} \min_i |\lambda_i|^{-1}, \quad (39)$$

such that the above integral is well-defined and finite for all $\sigma \in I$.

Let us define the function $\tilde{\phi} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ as follows:

$$\tilde{\phi}(t_1, t_2, \dots, t_d) := (t_1^{-2}, t_2^{-2}, \dots, t_d^{-2}). \quad (40)$$

Finally, consider the following polynomial in one variable with degree d :

$$p(x) := (1 - \lambda_1 x) \times (1 - \lambda_2 x) \times \dots \times (1 - \lambda_d x). \quad (41)$$

Claim 1. For any $\sigma_i \in I$, $i \in [d]$, we have

$$\tilde{\phi} \left(\int \psi_{\sigma_1} \left(\frac{1}{2} v^t \nabla^2 L(x) v \right) d\mu(v), \int \psi_{\sigma_2} \left(\frac{1}{2} v^t \nabla^2 L(x) v \right) d\mu(v), \dots, \int \psi_{\sigma_d} \left(\frac{1}{2} v^t \nabla^2 L(x) v \right) d\mu(v) \right) \quad (42)$$

$$= (p(\sigma_1), p(\sigma_2), \dots, p(\sigma_d)). \quad (43)$$

The above claim simply follows from the integral we calculated before.

Now we are ready to complete the proof. Choose arbitrary non-zero distinct $\sigma_i \in I$, $i \in [d]$, and note that having access to $(p(\sigma_1), p(\sigma_2), \dots, p(\sigma_d))$ is enough to recover all the eigenvalues. Indeed, assume that $p(x) = p_0 + p_1 x + p_2 x^2 + \dots + p_d x^d$ and note that $p(0) = 1$ by definition. Let also $V(0, \sigma_1, \sigma_2, \dots, \sigma_d) \in \mathbb{R}^{d \times d}$ denote a Vandermonde matrix of order $d + 1$, which is provably invertible by definition, and note that

$$(p(\sigma_1), p(\sigma_2), \dots, p(\sigma_d))^t = V(0, \sigma_1, \sigma_2, \dots, \sigma_d) \times (p_0, p_1, \dots, p_{d+1})^t \quad (44)$$

$$\implies (p_0, p_1, \dots, p_{d+1})^t = V(0, \sigma_1, \sigma_2, \dots, \sigma_d)^{-1} \times (p(\sigma_1), p(\sigma_2), \dots, p(\sigma_d))^t. \quad (45)$$

Indeed, this shows that having access to the vector $(p(\sigma_1), p(\sigma_2), \dots, p(\sigma_d))^t$ is enough to reconstruct the polynomial $p(x) = p_0 + p_1 x + \dots + p_d x^d$. Having access to this polynomial is equivalent to having access to its roots, so one can find a continuous function $\phi_1 : \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that

$$(\lambda_1, \lambda_2, \dots, \lambda_d)^t \quad (46)$$

$$= \phi_1 \circ \tilde{\phi} \left(\int \psi_{\sigma_1} \left(\frac{1}{2} v^t \nabla^2 L(x) v \right) d\mu(v), \int \psi_{\sigma_2} \left(\frac{1}{2} v^t \nabla^2 L(x) v \right) d\mu(v), \dots, \int \psi_{\sigma_d} \left(\frac{1}{2} v^t \nabla^2 L(x) v \right) d\mu(v) \right). \quad (47)$$

Since the sharpness measure $S(\lambda_1, \lambda_2, \dots, \lambda_d)$ is a continuous function of its coordinates, we conclude that

$$S(\lambda_1, \lambda_2, \dots, \lambda_d) \quad (48)$$

$$= S \circ \phi_1 \circ \tilde{\phi} \left(\int \psi_{\sigma_1} \left(\frac{1}{2} v^t \nabla^2 L(x) v \right) d\mu(v), \int \psi_{\sigma_2} \left(\frac{1}{2} v^t \nabla^2 L(x) v \right) d\mu(v), \dots, \int \psi_{\sigma_d} \left(\frac{1}{2} v^t \nabla^2 L(x) v \right) d\mu(v) \right). \quad (49)$$

Now to complete the proof, we define a continuous function $\psi : \mathbb{R} \rightarrow \mathbb{R}^d$ as $\psi = (\psi_{\sigma_1}, \psi_{\sigma_2}, \dots, \psi_{\sigma_d})^t$, and a continuous function $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ as $\phi = S \circ \phi_1 \circ \tilde{\phi}$, and observe that $S(\lambda_1, \lambda_2, \dots, \lambda_d) = S(x; \phi, \psi, \mu)$ for any $x \in \mathcal{A}$. This completes the proof. \square

D. Proof of Theorem 2

Theorem 2 (Universality of the (ϕ, ψ, μ) -sharpness measures for arbitrary functions of Hessian). *For any continuous function $S : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}$, there exist a positive integer $m \leq d(d + 1)/2$, (Borel) probability measures μ_ℓ , $\ell \in [m]$, and continuous functions $\phi : \mathbb{R}^m \rightarrow \mathbb{R}$ and $\psi : \mathbb{R} \rightarrow \mathbb{R}^m$, such that $S(\nabla^2 L(x)) = S(x; \phi, \psi, \mu)$ for any $x \in \mathbb{R}^d$, where $\mu := \mu_1 \otimes \mu_2 \otimes \dots \otimes \mu_m$ is a product probability measure.*

Proof. We explicitly construct a set of functions/probability measures to achieve the desired representation. Indeed, let's take $m = d(d+1)/2$ and consider the following Dirac measures: $\mu_i = \delta_{e_i}, i \in [d]$, and also $\mu_{ij} = \delta_{e_i+e_j}$, for any $i, j \in [d]$ such that $i < j$. Here, e_i denotes the unit vector in the i th coordinate in \mathbb{R}^d . Now, note that we have

$$(\nabla^2 L(x))_{i,i} = \int v^t \nabla^2 L(x) v d\mu_i(v) \quad (50)$$

for any $i \in [d]$, and

$$2(\nabla^2 L(x))_{i,j} + (\nabla^2 L(x))_{i,i} + (\nabla^2 L(x))_{j,j} = \int v^t \nabla^2 L(x) v d\mu_{i,j}(v), \quad (51)$$

for any $i, j \in [d]$ such that $i < j$. The above system of linear equations has clearly a unique solution, as the Hessian matrix is symmetric. This means that, similar to the proof of Theorem 1, one can find continuous functions $\psi : \mathbb{R} \rightarrow \mathbb{R}^m$ and $\phi : \mathbb{R}^m \rightarrow \mathbb{R}$, along with m constructed probability measures such that $S(\nabla^2 L(x)) = S(x; \phi, \psi, \mu)$ for any $x \in \mathbb{R}^d$, where $\mu = \mu_1 \otimes \mu_2 \otimes \dots \otimes \mu_m$ is a product probability measure. \square

E. Proof of Theorem 3

Theorem 3 (Explicit bias of the (ϕ, ψ, μ) -sharpness-aware loss function). *Given a triplet (ϕ, ψ, μ) , $m \geq 1$, and a training loss function $L : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$, assume that:*

- $L(x)$ is third-order continuously differentiable and satisfies the following upper bound

$$\max_{i,j,k \in \{1,2,3\}} |\partial_i \partial_j \partial_k L(v)| = O(\|v\|^{-1}), \quad (7)$$

for $v \in \mathbb{R}^d$ as $\|v\|_2 \rightarrow \infty$.

- The two functions ϕ, ψ are continuously differentiable.
- For some $C > \max_{x \in \mathcal{X}} \max_{i \in [d]} |\lambda_i(\nabla^2 L(x))|$, we have $\int \|v\|_2^2 \tilde{\psi}_i(v) d\mu(v) < \infty, i \in [m]$, where

$$\tilde{\psi}_i(v) := \max_{|t| \leq \|v\|_2} |\psi'_i(Ct^2)|. \quad (8)$$

Then, there exists an open neighborhood $U \supseteq \Gamma$, where Γ is the zero-loss manifold, for connected U and Γ , such that if for some $u \in U$, one has

$$L(u) + \rho^2 R_\rho(u) - \inf_{x \in U} (L(x) + \rho^2 R_\rho(x)) \leq \Delta \rho^2, \quad (9)$$

with some optimally gap $\Delta > 0$, then

$$L(u) \leq \underbrace{\inf_{x \in U} L(x)}_{=0} + (\Delta + o_\rho(1))\rho^2, \quad (10)$$

and also

$$S(u; \phi, \psi, \mu) \leq \inf_{x \in \Gamma} S(x; \phi, \psi, \mu) + \Delta + o_\rho(1). \quad (11)$$

Proof. For simplicity, we assume that $m = 1$. The general proof for $m > 1$ follows with a similar argument to this special case. Define an open set $U \subseteq \mathbb{R}^d$ as follows:

$$U := \{x \in \mathbb{R}^d : \|\nabla L(x)\|_2 < \rho^2\}. \quad (52)$$

Note that this set contain the zero-loss manifold, i.e., $\Gamma \subseteq U$. We study the behavior of the loss function on this open set.

Let us denote the sharpness term in the loss function $L_{(\phi, \psi, \mu)}(x)$ by $R_\rho(x)$:

$$L_{(\phi, \psi, \mu)}(x) = L(x) + \rho^2 R_\rho(x) \quad (53)$$

$$= L(x) + \rho^2 \phi \left(\int \psi \left(\frac{1}{\rho^2} (L(x + \rho v) - L(x)) \right) d\mu(v) \right). \quad (54)$$

We first study the convergence of $R_\rho(x)$ to the corresponding sharpness measure $S(x; \phi, \psi, \mu)$. Fix any point $x \in U$ and note that using Taylor's theorem for the function $L : \mathbb{R}^d \rightarrow \mathbb{R}$ and for any $v \in \mathbb{R}^d$, one has

$$\frac{1}{\rho^2} (L(x + \rho v) - L(x)) = \frac{1}{\rho^2} \left(\rho \langle \nabla L(x), v \rangle + \frac{1}{2} \rho^2 v^t \nabla^2 L(x) v + O_x(\rho^3 \|v\|_2^3 \times \|v\|_2^{-1}) \right) \quad (55)$$

$$= \rho^{-1} \langle \nabla L(x), v \rangle + \frac{1}{2} v^t \nabla^2 L(x) v + O_x(\rho \|v\|_2^2 \times \|v\|_2^{-1}), \quad (56)$$

where in above we used the fact that $L(x)$ is third-order continuously differentiable and its third-order derivative satisfies the estimate

$$\max_{i,j,k \in \{1,2,3\}} |\partial_i \partial_j \partial_k L(v)| = O(\|v\|^{-1}), \quad (57)$$

for $v \in \mathbb{R}^d$ as $\|v\|_2 \rightarrow \infty$.

Note that using the assumption $x \in U$, we have that

$$|\rho^{-1} \langle \nabla L(x), v \rangle| \leq \rho^{-1} \|\nabla L(x)\|_2 \|v\|_2 < \rho \|v\|_2. \quad (58)$$

Thus, we have

$$\frac{1}{\rho^2} (L(x + \rho v) - L(x)) = \frac{1}{2} v^t \nabla^2 L(x) v + O_x(\rho (\|v\|_2^2 + \|v\|_2)). \quad (59)$$

Note that we study the above approximation only for $x \in U$, and for small enough ρ , we know that U is a precompact set. Therefore, we drop the dependence on x in the error term above.

Now using the above approximation, we have

$$\int \psi \left(\frac{1}{\rho^2} (L(x + \rho v) - L(x)) \right) d\mu(v) = \int \psi \left(\frac{1}{2} v^t \nabla^2 L(x) v + O(\rho (\|v\|_2^2 + \|v\|_2)) \right) d\mu(v). \quad (60)$$

Let us use Taylor's theorem for the function ψ and write

$$\psi \left(\frac{1}{2} v^t \nabla^2 L(x) v + O(\rho (\|v\|_2^2 + \|v\|_2)) \right) = \psi \left(\frac{1}{2} v^t \nabla^2 L(x) v \right) + \rho \times O(\tilde{\psi}(v) (\|v\|_2^2 + \|v\|_2)), \quad (61)$$

where

$$\tilde{\psi}(v) := \max_{|t| \leq \|v\|_2} |\psi'(Ct^2)|, \quad (62)$$

and C is a constant, and it's big enough to absorb the quadratic growth of $\nabla^2 L(x)$; i.e.,

$$C > \max_{x \in \mathcal{X}} \max_{i \in [d]} |\lambda_i(\nabla^2 L(x))|. \quad (63)$$

Therefore, we conclude that

$$\int \psi \left(\frac{1}{\rho^2} (L(x + \rho v) - L(x)) \right) d\mu(v) = \int \psi \left(\frac{1}{2} v^t \nabla^2 L(x) v \right) d\mu(v) + \rho \times O \left(\int \tilde{\psi}(v) (\|v\|_2^2 + \|v\|_2) d\mu(v) \right) \quad (64)$$

$$= \int \psi \left(\frac{1}{2} v^t \nabla^2 L(x) v \right) d\mu(v) + O(\rho), \quad (65)$$

by the assumption. This allows us to conclude that

$$R_\rho(x) = \phi \left(\int \psi \left(\frac{1}{\rho^2} (L(x + \rho v) - L(x)) \right) d\mu(v) \right) \quad (66)$$

$$= \phi \left(\int \psi \left(\frac{1}{2} v^t \nabla^2 L(x) v \right) d\mu(v) \right) + O(\rho) \quad (67)$$

$$= S(x; \phi, \psi, \mu) + O(\rho), \quad (68)$$

again, by assuming that

$$\max_{x \in \mathcal{X}} \phi' \left(\int \psi \left(\frac{1}{2} v^t \nabla^2 L(x) v \right) d\mu(v) \right) < \infty, \quad (69)$$

which holds by the compactness of \mathcal{X} , and also using $U \subseteq \mathcal{X}$.

Now according to the assumption, for some $u \in U$, we have

$$L(u) + \rho^2 R_\rho(u) - \inf_{x \in U} (L(x) + \rho^2 R_\rho(x)) \leq \Delta \rho^2, \quad (70)$$

for some optimally gap Δ . Using the following proven approximation

$$R_\rho(x) = S(x; \phi, \psi, \mu) + O(\rho), \quad (71)$$

we conclude that

$$L(u) + \rho^2 S(u; \phi, \psi, \mu) - \inf_{x \in U} (L(x) + \rho^2 S(x; \phi, \psi, \mu)) \leq (\Delta + O(\rho)) \rho^2. \quad (72)$$

Now by proof by contradiction, assume that

$$L(u) \geq \inf_{x \in U} L(x) + (\Delta + \delta) \rho^2, \quad (73)$$

for some strictly positive δ , as $\rho \rightarrow 0^+$. Note that $\inf_{x \in U} L(x) = 0$ as $\Gamma \subseteq U$. Thus, we can conclude that

$$\rho^2 S(u; \phi, \psi, \mu) + (\Delta + \delta) \rho^2 \leq L(u) + \rho^2 S(u; \phi, \psi, \mu) \quad (74)$$

$$\leq \inf_{x \in U} (L(x) + \rho^2 S(x; \phi, \psi, \mu)) + (\Delta + O(\rho)) \rho^2 \quad (75)$$

$$\leq \rho^2 \inf_{x \in \Gamma} S(x; \phi, \psi, \mu) + (\Delta + O(\rho)) \rho^2, \quad (76)$$

since $\Gamma \subseteq U$. This shows that

$$S(u; \phi, \psi, \mu) \leq \inf_{x \in \Gamma} S(x; \phi, \psi, \mu) - \delta + O(\rho). \quad (77)$$

This must hold for as $\rho \rightarrow 0^+$. However, as $\rho \rightarrow 0^+$, we have that $U_\rho \rightarrow \Gamma$. This means that

$$S(u; \phi, \psi, \mu) \leq \inf_{x \in \Gamma} S(x; \phi, \psi, \mu) - \delta, \quad (78)$$

for some $u \in \Gamma$, which is a contradiction. This shows that

$$L(u) \leq \inf_{x \in U} L(x) + (\Delta + o(1)) \rho^2. \quad (79)$$

Also, to prove the next part of the theorem, for any $u \in U$ satisfying the assumptions, similarly we can show

$$\rho^2 S(u; \phi, \psi, \mu) \leq L(u) + \rho^2 S(u; \phi, \psi, \mu) \quad (80)$$

$$= \inf_{x \in U} (L(x) + \rho^2 S(x; \phi, \psi, \mu)) + (\Delta + O(\rho)) \rho^2 \quad (81)$$

$$\leq \rho^2 \inf_{x \in \Gamma} S(x; \phi, \psi, \mu) + (\Delta + O(\rho)) \rho^2, \quad (82)$$

which implies that

$$S(u; \phi, \psi, \mu) \leq \inf_{x \in \Gamma} S(x; \phi, \psi, \mu) + (\Delta + O(\rho)). \quad (83)$$

The proof is thus complete. □

F. Proof of Theorem 4

Theorem 4 (Scale-invariant (ϕ, ψ, μ) -sharpness measures). *Consider a scale-invariant loss function $L(x)$ and let μ be a Borel measure of the form*

$$d\mu(x) = f\left(\prod_{i=1}^d x_i\right) \prod_{i=1}^d dx_i, \quad (12)$$

where $f : \mathbb{R} \rightarrow \mathbb{R}$ is a measurable function⁵. Then, for any continuous functions ϕ, ψ , the corresponding sharpness measure $S(x; \phi, \psi, \mu)$ is scale-invariant; this means that $S(x; \phi, \psi, \mu) = S(Dx; \phi, \psi, \mu)$ for any diagonal matrix $D \in \mathbb{R}^{d \times d}$ with $\det(D) = 1$.

Proof. Let $D \in \mathbb{R}^{d \times d}$ be an arbitrary diagonal matrix. Then,

$$S(Dx; \phi, \psi, \mu) = \phi\left(\int \psi\left(\frac{1}{2}v^t \nabla^2 L(x)\Big|_{Dx} v\right) d\mu(v)\right). \quad (84)$$

But note that by the assumption

$$L(x) = L(Dx) \implies \nabla^2 L(x) = D^t \nabla^2 L(x)\Big|_{Dx} D. \quad (85)$$

Therefore,

$$S(Dx; \phi, \psi, \mu) = \phi\left(\int \psi\left(\frac{1}{2}v^t D^{-1} \nabla^2 L(x) D^{-1} v\right) d\mu(v)\right). \quad (86)$$

Now define a new variable $u := D^{-1}v$. Then,

$$d\mu(v) = f\left(\prod_{i=1}^d v_i\right) \prod_{i=1}^d dv_i = f\left(\prod_{i=1}^d D_{i,i} \prod_{i=1}^d u_i\right) \prod_{i=1}^d D_{i,i} \prod_{i=1}^d du_i \quad (87)$$

$$= f(\det(D) \prod_{i=1}^d u_i) \det(D) \prod_{i=1}^d du_i \quad (88)$$

$$= f\left(\prod_{i=1}^d u_i\right) \prod_{i=1}^d du_i. \quad (89)$$

Therefore, we conclude that

$$S(Dx; \phi, \psi, \mu) = \phi\left(\int \psi\left(\frac{1}{2}u^t \nabla^2 L(x) u\right) d\mu(u)\right) = S(x; \phi, \psi, \mu), \quad (90)$$

and this completes the proof. □

Lemma 1. *For any scale-invariant measure μ that is absolutely continuous with respect to the Lebesgue measure on \mathbb{R}^d , one has*

$$d\mu(x) = f\left(\prod_{i=1}^d x_i\right) \prod_{i=1}^d dx_i.$$

Proof. Any measure μ that is absolutely continuous with respect to the Lebesgue measure on \mathbb{R}^d can be written as

$$d\mu(x_1, x_2, \dots, x_d) = \tilde{f}(x_1, x_2, \dots, x_d) \prod_{i=1}^d dx_i.$$

⁵In Lemma 1, we show that any scale-invariant measure is of this form.

For any non-zero choice of x_1, x_2, \dots, x_d , by scale invariance property of $d\mu$,

$$\begin{aligned} d\mu(x_1, x_2, \dots, x_d) &= d\mu(1, 1, \dots, 1, \prod_{i=1}^d x_i) \\ \Rightarrow \tilde{f}(x_1, x_2, \dots, x_d) \prod_{i=1}^d dx_i &= \tilde{f}(1, 1, \dots, 1, \prod_{i=1}^d x_i) \prod_{i=1}^d dx_i \\ \Rightarrow \tilde{f}(x_1, x_2, \dots, x_d) &= \tilde{f}(1, 1, \dots, 1, \prod_{i=1}^d x_i). \end{aligned}$$

Now, it suffices to choose $f(\prod_{i=1}^d x_i) := \tilde{f}(1, 1, \dots, 1, \prod_{i=1}^d x_i)$ almost surely and the proof is concluded. \square

G. Proof of Theorem 5

Theorem 5 (General parameter-invariant (ϕ, ψ, μ) -sharpness measures). *Let G be a group acting by matrices on \mathbb{R}^d , and assume that $L(x)$ is invariant with respect to the action of G . Then, for any G -invariant (Borel) measure μ , and any continuous functions ϕ, ψ , the corresponding sharpness measure $S(x; \phi, \psi, \mu)$ is G -invariant; this means that $S(x; \phi, \psi, \mu) = S(A_g x; \phi, \psi, \mu)$ for any matrix $A_g \in \mathbb{R}^{d \times d}$ corresponding to the action of an element $g \in G$.*

Proof. We start by evaluating $S(A_g x; \phi, \psi, \mu)$.

$$S(Dx; \phi, \psi, \mu) = \phi \left(\int \psi \left(\frac{1}{2} v^t \nabla^2 L(x) \Big|_{A_g x} v \right) d\mu(v) \right). \quad (91)$$

But again here, note that by the assumption

$$L(x) = L(A_g x) \implies \nabla^2 L(x) = A_g^t \nabla^2 L(x) \Big|_{A_g x} A_g. \quad (92)$$

Therefore,

$$S(A_g x; \phi, \psi, \mu) = \phi \left(\int \psi \left(\frac{1}{2} v^t A_g^{-1} \nabla^2 L(x) A_g^{-1} v \right) d\mu(v) \right). \quad (93)$$

Now define a new variable $u := A_g^{-1} v$. Therefore, we conclude that

$$S(A_g x; \phi, \psi, \mu) = \phi \left(\int \psi \left(\frac{1}{2} u^t \nabla^2 L(x) u \right) d\mu(u) \right) = S(x; \phi, \psi, \mu), \quad (94)$$

and this completes the proof. \square

H. (ϕ, ψ, μ) -Sharpness-Aware Minimization Algorithm

To propose an algorithm for the general case (i.e., arbitrary m), we compute the gradient of

$$\tilde{R}_\rho(x) := \phi \left(\frac{1}{n} \sum_{i=1}^n \psi_1 \left(\frac{1}{\rho^2} (L(x + \rho v_{i,1}) - L(x)) \right) \right), \quad (95)$$

$$\frac{1}{n} \sum_{i=1}^n \psi_2 \left(\frac{1}{\rho^2} (L(x + \rho v_{i,2}) - L(x)) \right), \quad (96)$$

$$\dots \quad (97)$$

$$\frac{1}{n} \sum_{i=1}^n \psi_m \left(\frac{1}{\rho^2} (L(x + \rho v_{i,m}) - L(x)) \right), \quad (98)$$

Algorithm 3 (ϕ, ψ, μ) -Sharpness-Aware Minimization Algorithm (with arbitrary m)

Input: The triplet (ϕ, ψ, μ) , Training loss $L(x)$, Step size η , Perturbation parameter ρ , Number of samples n ,

Output: Model parameters x_t trained with (ϕ, ψ, μ) -Sharpness-Aware Minimization Algorithm

Initialization: $x \leftarrow x_0$ and $t \leftarrow 0$

while 1 do

Sample $v_{i,\ell} \stackrel{\text{i.i.d.}}{\sim} \mu_\ell$, for any $i \in [n]$ and $\ell \in [m]$

Compute the following:

$$g_t = \nabla L(x_t) + \sum_{\ell=1}^m \partial_\ell \phi \left(\sum_{i=1}^n \frac{1}{n} \psi_\ell \left(\frac{1}{\rho^2} (L(x_t + \rho v_{i,\ell}) - L(x_t)) \right) \right) \\ \times \sum_{i=1}^n \frac{1}{n} \left\{ \psi'_\ell \left(\frac{1}{\rho^2} (L(x_t + \rho v_{i,\ell}) - L(x_t)) \right) \times \left(\nabla L(x_t + \rho v_{i,\ell}) - \nabla L(x_t) \right) \right\}.$$

Update the parameters: $x_{t+1} = x_t - \eta g_t$

$t \leftarrow t + 1$

end while

	CIFAR10	CIFAR100	SVHN	CIFAR10-S	CIFAR100-S	SVHN-S	CIFAR10-C	CIFAR100-C	SVHN-C
Frob-SAM	5e-3	1e-4	5e-3	1e-4	1e-4	1e-3	0.01	0.01	1e-3
Trace-SAM	1	0.1	0.1	0.01	0.01	1	1	0.1	0.01
Det-SAM	1	0.1	0.1	0.01	1	1	0.01	0.1	0.1

Table 4. Choice of regularizer weight λ for different datasets.

	CIFAR10	CIFAR100	SVHN	CIFAR10-S	CIFAR100-S	SVHN-S	CIFAR10-C	CIFAR100-C	SVHN-C
Frob-SAM	5e-3	5e-3	5e-3	1e-2	5e-3	5e-3	1e-2	5e-3	1e-2

Table 5. Choice of ρ for Frob-SAM for different datasets.

where $v_{i,\ell} \stackrel{\text{i.i.d.}}{\sim} \mu_\ell$ for each $\ell \in [m]$. Note that $\psi = (\psi_1, \psi_2, \dots, \psi_m)^t$ for some scalar functions $\psi_\ell, \ell \in [m]$. Let $\partial_\ell \phi$ denote partial derivatives of the function $\phi : \mathbb{R}^m \rightarrow \mathbb{R}$, for any $\ell \in [m]$. Then,

$$\rho^2 \nabla \tilde{R}_\rho(x) = \sum_{\ell=1}^m \partial_\ell \phi \left(\sum_{i=1}^n \frac{1}{n} \psi_\ell \left(\frac{1}{\rho^2} (L(x_t + \rho v_{i,\ell}) - L(x_t)) \right) \right) \quad (99)$$

$$\times \sum_{i=1}^n \frac{1}{n} \left\{ \psi'_\ell \left(\frac{1}{\rho^2} (L(x_t + \rho v_{i,\ell}) - L(x_t)) \right) \times \left(\nabla L(x_t + \rho v_{i,\ell}) - \nabla L(x_t) \right) \right\}, \quad (100)$$

and this leads to Algorithm 3.

I. Experiments

I.1. Experimental Details

We now describe experimental details that were omitted from the main text.

For CIFAR10 and CIFAR100, we apply random crops and random horizontal flips. We use a momentum term of 0.9 for all datasets and a weight decay of 5e-4 for CIFAR10 and SVHN and 1e-3 for CIFAR100. We use batch size 128 and train for 200 epochs. We use a multi-step schedule where the learning rate is initially 0.1 and decays by a multiplicative factor of 0.1 every 50 epochs. We run each experiment with four different random seeds to access statistical significance. We use 1280 training examples and 100 noise samples to estimate the Frobenius norm and trace via Hessian-vector products. We set ρ to 1.0 for Det-SAM, 0.01 for Trace-SAM, and sweep it in $\{0.005, 0.01\}$ for Frob-SAM. For Det-SAM and Trace-SAM we sweep λ in $\{0.01, 0.1, 1.0\}$ and set $n = 1$. For Frob-SAM, we sweep λ in $\{0.0001, 0.001, 0.005, 0.01, 0.05, 0.1\}$ and set $n = 2$. The hyper-parameters selected for each setting is given in Table 4 and Table 5.

We use the PyHessian library (Yao et al., 2020) to estimate the trace of the Hessian. Adapting this library, we estimate the Frobenius norm (squared) as $\frac{1}{k} \sum_{i=1}^k \|H z_i\|_2^2$, where H is the Hessian matrix and $z_i \sim \mathcal{N}(0, I_d)$.

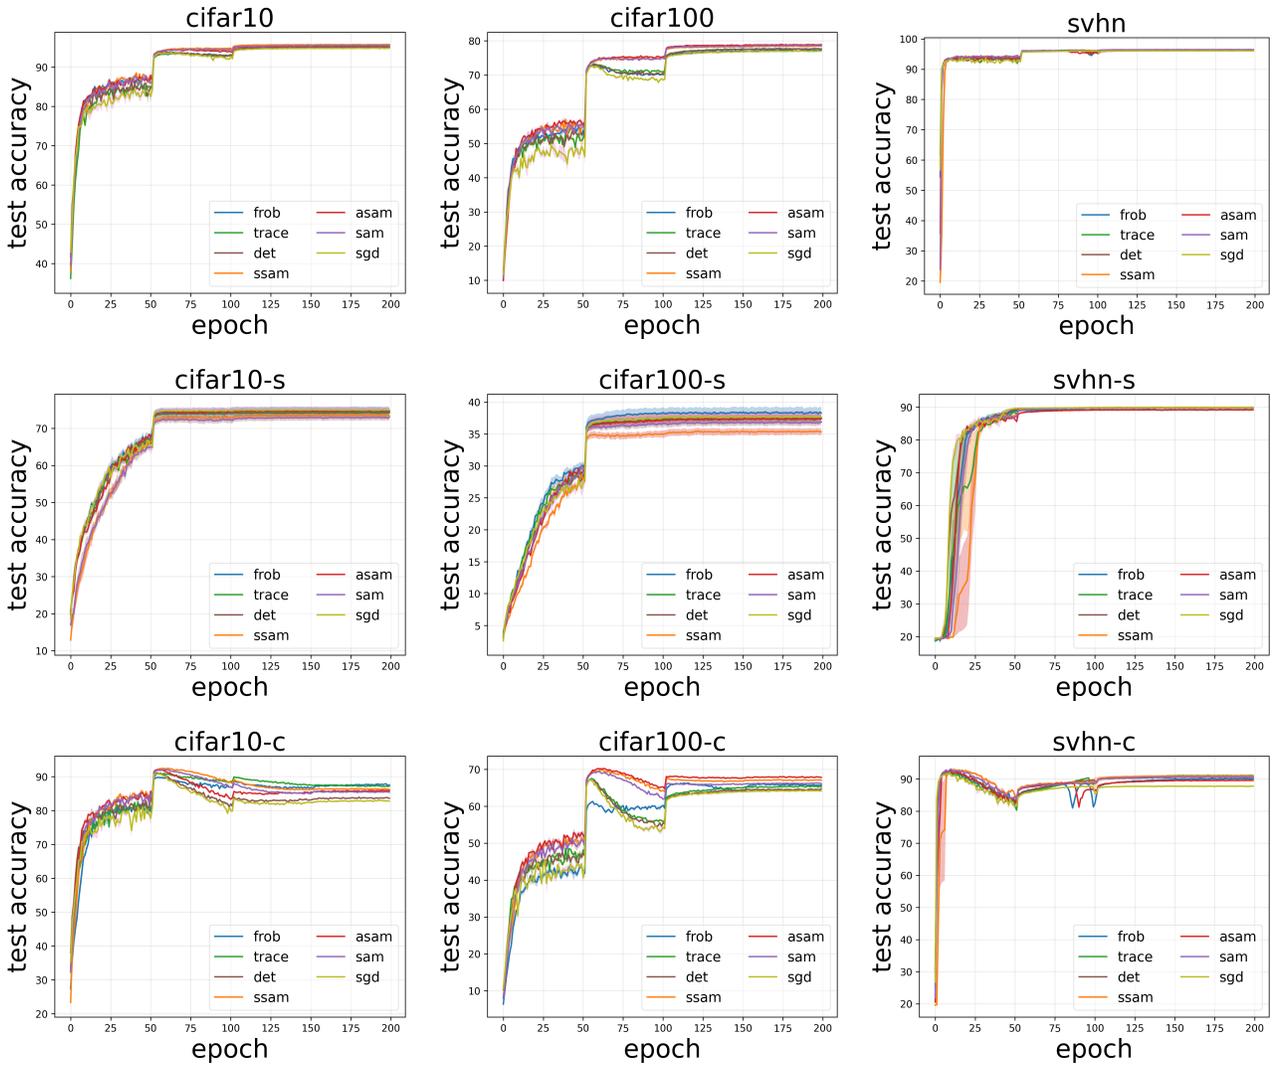


Figure 3. Training plots for all settings. One standard error is shaded.

I.2. Additional Plots

All training plots are shown in Figure 3.