

TOWARDS ROBUST OUT-OF-DISTRIBUTION GENERALIZATION BOUNDS VIA SHARPNESS

Yingtian Zou, Kenji Kawaguchi, Yingnan Liu, Jiashuo Liu*, Mong-Li Lee, Wynne Hsu

School of Computing, National University of Singapore

Institute of Data Science, National University of Singapore

*Department of Computer Science & Technology, Tsinghua University, China

ABSTRACT

Generalizing to out-of-distribution (OOD) data or unseen domain, termed OOD generalization, still lacks appropriate theoretical guarantees. Canonical OOD bounds focus on different distance measurements between *source* and *target* domains but fail to consider the optimization property of the learned model. As empirically shown in recent work, the *sharpness* of learned minima influences OOD generalization. To bridge this gap between optimization and OOD generalization, we study the effect of sharpness on how a model tolerates data change in domain shift which is usually captured by "robustness" in generalization. In this paper, we give a rigorous connection between sharpness and robustness, which gives better OOD guarantees for robust algorithms. It also provides a theoretical backing for "flat minima leads to better OOD generalization". Overall, we propose a sharpness-based OOD generalization bound by taking robustness into consideration, resulting in a tighter bound than non-robust guarantees. Our findings are supported by the experiments on a ridge regression model, as well as the experiments on deep learning classification tasks.

1 INTRODUCTION

Machine learning systems are typically trained on a given distribution of data and achieve good performance on new, unseen data that follows the same distribution as the training data. Out-of-Distribution (OOD) generalization requires machine learning systems trained in the source domain to generalize to unseen data or target domains with different distributions from the source domain. A myriad of algorithms (Sun & Saenko, 2016; Arjovsky et al., 2019; Sagawa et al., 2019; Koyama & Yamaguchi, 2020; Pezeshki et al., 2021; Ahuja et al., 2021) aim to learn the invariant components along the distribution shifting. Optimization-based methods such as (El Ghaoui & Le Bret, 1997; Duchi & Namkoong, 2018; Liu et al., 2021; Rame et al., 2022) focus on maximizing robustness by optimizing for worst-case error over an uncertainty distribution set. While these methods are sophisticated, they do not always perform better than Empirical Risk Minimization (ERM) when evaluated across different datasets (Gulrajani & Lopez-Paz, 2021; Wiles et al., 2022). This raises the question of how to understand the OOD generalization of algorithms and which criteria should be used to select models that are provably better (Gulrajani & Lopez-Paz, 2021). These questions highlight the need for more theoretical research in the field of OOD generalization (Ye et al., 2021).

To characterize the generalization gap between the source domain and the target domain, a canonical method (Blitzer et al., 2007) from domain adaptation theory decouples this gap into an In-Distribution (ID) generalization and a hypothesis-specific Out-of-Distribution (OOD) distance. However, this distance is based on the notion of VC-dimension (Kifer et al., 2004), resulting in a loose bound due to the large size of the hypothesis class in the modern overparameterized neural networks. Subsequent works improve the bound based on Rademacher Complexity (Du et al., 2017), whereas Germain et al. (2016) improves the bound based on PAC-Bayes. Unlike the present paper, these works did not consider algorithmic robustness, which has natural interpretation and advantages for distribution shifts. In this work, we consider algorithmic robustness to derive the OOD generalization bound. The key idea is to partition the input space into K non-overlapping subspaces such that the error difference in the model's performance between any pair of points in each subspace is bounded by some constant ϵ . Within each subspace, any distributional shift is considered subtle for the robust

model thus leading to less impact on OOD generalization. Figure 1 illustrates this with the two distributions where the target domain has a distributional shift from the source domain. Compared to existing non-robust OOD generalization bounds Zhao et al. (2018), our new generalization error does not depend on hypothesis size, which is more reliable in the overparameterized regime. Our goal is to measure the generalizability of a model by considering how it is robust to this shift and achieves a tighter bound than existing works.

Although robustness captures the tolerance to distributional shift, it is intractable to compute robustness constant ϵ due to the inaccessibility of target distribution. The robustness definition in Xu & Mannor (2012) indicates that the loss landscape induced by the model’s parameters is closely tied to its robustness. To gain a deeper understanding of robustness, we further study the learned model from an optimization perspective. As shown in (Lyu et al., 2022; Petzka et al., 2021), when the loss landscape is "flat", there is a good generalization, which is also observed in OOD settings (Izmailov et al., 2018; Cha et al., 2021). However, the relationship between robustness and this geometric property of the loss landscape, termed *Sharpness*, remains an open question. In this paper, we establish a provable dependence between robustness and sharpness for ReLU random neural network classes. It allows us to replace robustness constant ϵ with the sharpness of a learned model which is only computed from the training dataset that addresses the problem mentioned above. Our result of the interplay between robustness and sharpness can be applied to both ID and OOD generalization bounds. We also show an example to generalize our result beyond our assumption and validate it empirically.

Our main contributions can be summarized as follows:

- We proposed a new framework for Out-of-distribution/ Out-of-domain generalization bounds. In this framework, we use robustness to capture the tolerance of distribution shift which leads to tighter upper bounds generally.
- We reveal the underlying connection between the robustness and sharpness of the loss landscape and use this connection to enrich our robust OOD bounds under one-hidden layer ReLU NNs. This is the first optimization-based bound in Out-of-Distribution/Domain generalization.
- We studied two cases in ridge regression and classification which support and generalize our main theorem well. All the experimental results corroborate our findings well.

2 PRELIMINARY

Notations We use $[n]$ to denote the integers set $\{1, \dots, n\}$. We use $\|\cdot\|_k$ to denote the ℓ_2 -norm (Euclidean norm). In vector form, w_i denotes the i -th instance while the w_j is the j -th element of the vector \mathbf{w} and we use $|w_j|$ for the element-wise absolute value of vector \mathbf{w} . We use n, d for the training set size and input dimension. \mathcal{O} is the Big-O notation.

2.1 PROBLEM FORMULATION

Consider a source domain and a target domain of the OOD generalization problem where we use D_S, D_T to represent the source and target distribution respectively. Let each D be the probability




Figure 1: An example of a target distribution (red) directly translated from a source distribution (blue). The 1D density reflects the marginal distribution. Unlike existing works (left), we divide the distributions into disjoint partitions as a small change in distribution for a robust model is negligible (right). The sharpness of the model will decide the tolerance of change thus affecting the partitions. If two sub-distributions S, T have small shifts such that they fall into the same partition (red grid), their distance measure $d'(S, T)$ by considering robustness will be zero.

measure of sample \mathbf{z} from sample space $Z = X \times Y$ with $X \subseteq \mathbb{R}^d$. In the source domain, we have a training set $S = \{z_i\}_{i=1}^n$, \mathcal{D}_S while the target is to learn a model $f \in \mathcal{F}$ with S and parameters θ where $f : X \rightarrow \mathbb{R}$ generalizes well. Given loss function $\ell : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_+$, which is for short, the expected risk over the source distribution \mathcal{D}_S will be

$$L_S(f) = \mathbb{E}_{z \sim \mathcal{D}_S}[\ell(f(z), y)] = \mathbb{E}_{z \sim \mathcal{D}_S}[\ell(f(x), y)], \quad \hat{L}_S(f) = \frac{1}{n} \sum_{z_i \in S} \ell(z_i).$$

We use $\ell(z)$ as the shorthand. The OOD generalization is to measure between target domain expected risk $L_T(f)$ and the source domain empirical risk $\hat{L}_S(f)$ which involves two parts: (1) In-domain generalization error gap between empirical risk and expected risk $L_S(f)$ in the source domain. (2) Out-of-Domain distance between source and target domains. A model-agnostic example in [Zhao et al. \(2018\)](#) gave the following uniform bound:

Proposition 2.1 (Zhao et al. [Zhao et al. \(2018\)](#) Theorem 2 & 3.2). *With hypothesis class F and pseudo dimension $\text{Pdim}(F) = d'$, unlabeled empirical datasets from source and target distribution \mathcal{D}_S and \mathcal{D}_T with size n each, then with probability at least $1 - \delta$, for all $f \in F$,*

$$L_T(f) \leq \hat{L}_S(f) + \frac{1}{2} d_{\mathcal{F}}(\mathcal{D}_T; \mathcal{D}_S) + O\left(\sqrt{\frac{d'}{n}}\right)$$

where $d_{\mathcal{F}}(\mathcal{D}_T; \mathcal{D}_S) := 2 \sup_{A \in \mathcal{F}} |\mathbb{P}_{\mathcal{D}_T}[A] - \mathbb{P}_{\mathcal{D}_S}[A]|$ and \oplus is the XOR operator. Specific form of $O\left(\sqrt{\frac{d'}{n}}\right)$ is defined in [Appendix C.6](#).

2.2 ALGORITHMIC ROBUSTNESS

Definition 2.2 (Robustness, [Xu & Mannor \(2012\)](#)). A learning model f on training set S is $(K, \epsilon(\cdot))$ -robust, for $K \geq \mathbb{N}$, if Z can be partitioned into K disjoint sets, denoted by $\{C_i\}_{i=1}^K$, such that $\mathcal{D}_S \subseteq S$ we have

$$|\mathbb{E}_{z \in C_i} \ell(f(z)) - \mathbb{E}_{\mathbf{s} \in C_i} \ell(f(\mathbf{s}))| \leq \epsilon(S).$$

This definition captures the robustness of the model in terms of the input. Within each partitioned set C_i , the loss difference between any sample \mathbf{z} belonging to C_i and training sample $\mathbf{s} \in C_i$ will be upper bounded by the robustness constant $\epsilon(S)$. The generalization result given by [Xu & Mannor \(2012\)](#) provides a framework to bound the empirical risk with algorithmic robustness which has been stated in [Appendix C](#). Based on this framework, we are able to reframe the existing OOD generalization theory.

3 MAIN RESULTS

In this section, we propose a new Out-of-Distribution (OOD) generalization bound for robust algorithms that have not been extensively studied yet. We then compare our result to the existing domain shift bound in [Proposition 2.1](#) and discuss its implications for OOD and domain generalization problems by considering algorithmic robustness. To further explain the introduced robustness, we connect it to the *sharpness* of the minimum (a widely concerned geometric property in optimization) by showing a rigorous dependence between robustness and sharpness. This interplay will give us a better understanding of the OOD generalization problem, and meanwhile, provide more information on the final generalization bound. Detailed assumptions are clarified in [Appendix B.1](#).

3.1 ROBUST OOD GENERALIZATION BOUND

The main concern in OOD generalization is to measure the domain shift of a learned model. However, existing methods fail to consider the intrinsic property of the model, such as robustness. [Definition 2.2](#) gives us a new robustness measurement to the model trained on dataset S where the training set S is a collection of *i.i.d.* data pair (\mathbf{x}, y) sampled from source distribution \mathcal{D}_S with size n . The measurement provides an intuition that if a test sample from the target domain is similar to a specific group of training samples, their losses will be similar as well. In other words, the model's robustness is a reflection of its ability to generalize to unseen data. Our first theorem shows that by utilizing the "robustness" measurement, we can more effectively handle domain shifts by setting a tolerance

range for distribution changes. This robustness measurement, therefore, provides a useful tool for addressing OOD generalization.

Theorem 3.1. *Let \hat{D}_T be the empirical distribution of size n drawn from D_T . Assume that the loss ℓ is upper bounded by M . With probability at least $1 - \delta$ (over the choice of the samples S), for every f trained on S satisfying $(K, \epsilon(S))$ -robust, we have*

$$L_T(f) \leq \ell_S(f) + Md_{(\cdot; K)}(S, \hat{D}_T) + 2\epsilon(S) + 3M \sqrt{\frac{2K \ln 2 + 2 \ln(2/\delta)}{n}} \quad (1)$$

where the total variation distance $d_{(\cdot; K)}$ for discrete empirical distributions is defined by:

$$\forall i \in [n], n_i(S) := \#\{z \in S \setminus C_i\}, \quad d_{(\cdot; K)}(S, \hat{D}_T) := \sum_{i=1}^K \frac{n_i(S)}{n} - \frac{n_i(\hat{D}_T)}{n} \quad (2)$$

and $n_i(S), n_i(\hat{D}_T)$ are the number of samples from S and \hat{D}_T that fall into the set C_i , respectively.

Remark. The result can be decomposed into in-domain generalization and out-domain distance $jL_T(f) - L_S(f)j$ (please refer to Lemma C.1). Both of them depend on robustness $\epsilon(S)$.

See proof in Appendix C. The last three terms on the RHS of (1) are *distribution distance*, *robustness constant*, and *error term*, respectively. Unlike traditional distribution measures, we partition the sample space and the distributional shift separately in the K sub-groups instead of measuring it point-wisely. We argue that the $d_{(\cdot; K)}(S, \hat{D}_T)$ can be zero measure if all small changes happen within the same partition where a 2D illustrative case is shown in Figure 1. Under the circumstances, our distribution distance term will be significantly smaller than Proposition 2.1 as the target distribution is essentially a perturbation of the source distribution. As a robust OOD generalization measure, our bound characterizes how robust the learned model is to negligible distributional perturbations. To prevent a bound that expands excessively, we also propose an alternate solution tailored for non-robust algorithms ($K \neq 1$) as follows.

Corollary 3.2. *If $K \neq 1$, the domain shift bound $jL_T(f) - L_S(f)j$ can be replaced to the distribution distance in Proposition 2.1 where*

$$jL_T(f) - L_S(f)j \leq \frac{1}{2}d_{\mathcal{F}}(S; D_T) + \frac{1}{2}d_{\mathcal{F}}(S; \mathcal{D}_T) + O(\sqrt{d'/n}) \quad (3)$$

where the pseudo dimension $\text{Pdim}(F) = d'$.

The proof is in Appendix C.1. As dictated in Theorem 3.1, when $K \neq 1$, the use infinite number of partitions on the data distribution leads to meaningless robustness. However, Corollary C.7 suggests that our bound can be replaced by $d_{\mathcal{F}}(D_S; D_T)$ in the limit of infinite K . This avoids computing a vacuous bound for non-robust algorithms. In summary, Theorem 3.1 presents a novel approach for quantifying distributional shifts by incorporating the concept of robustness. Our framework is particularly beneficial when a robust algorithm is able to adapt to local shifts in the distribution. Additionally, our data-dependent result remains valid and useful in the overparameterized regime, since K does not depend on the model size.

3.2 SHARPNESS AND ROBUSTNESS

Clearly, robustness is inherently tied to the optimization properties of a model, particularly the curvature of the loss landscape. One direct approach to characterize this geometric curvature, referred to as "sharpness," involves analyzing the Hessian matrix (Foret et al., 2020; Cohen et al., 2021). Recent research (Petzka et al., 2021) has shown that the concept of "relative flatness", the *sharpness* in this paper, has a strong correlation with model generalization. However, the impact of relative flatness on OOD generalization remains uncertain, even within the convex setting. To address this problem, we aim to investigate the interplay between robustness and sharpness. With the following definition of sharpness, we endeavor to establish an OOD generalization bound rooted in optimization principles.

Definition 3.3 (Sharpness, Petzka et al. (2021)). For a twice differentiable loss function $L(\mathbf{w}) = \sum_{s \in S} \ell_{\mathbf{w}}(s)$, $\mathbf{w} \in \mathbb{R}^m$ with a sample set S , the sharpness is defined by

$$\kappa(\mathbf{w}, S, A) := \|\mathbf{w}\| \cdot \text{tr}(H_{S;A}(\mathbf{w})) \quad (4)$$

where $H_{S;A}$ is the Hessian matrix of loss $L(\mathbf{w})$ w.r.t. \mathbf{w} with hypothesis set A and input set S .

See proof in Appendix D. From Theorem 3.4 we can see, the robustness coefficient (Stat) is (pointwise) upper bounded by the sharpness of the learned model, as measured by the quantity $\mathcal{S}(\hat{S}; A)$, and the parameter $\epsilon_{\max}(L)$. It should be noted that the parameter $\epsilon_{\max}(L)$ depends on the partition, and as the number of partitions increases, the region $\mathcal{R}(L)$ of the input domain becomes smaller, thereby making sharpness the dominant term in reflecting the model's robustness within the small local area. In Corollary 3.5, we show a stronger connection when the partition satisfies some conditions. Overall, this bound states that the larger the sharpness of the model $\mathcal{S}(\hat{S}; A)$, the larger the upper bound on the robustness parameter $\epsilon_{\max}(L)$. This result aligns with the intuition that a sharper model is more prone to overfitting the training domain and is less robust in the unseen domain. While the dependency is not exact, it still can be regarded as an alternative approach that avoids the explicit computation of the intractable robustness term. By substituting this upper bound for $\mathcal{S}(\hat{S}; A)$ into Theorem 3.1, we derive a sharpness-based OOD generalization bound. This implies that the OOD generalization error will have a high probability of being small if the learned model is sharp enough. Unlike existing works, our generalization bound provides more information about how optimization property influences performance when generalizing to OOD data. It bridges the gap between robustness and sharpness which can also be generalized to non-OOD learning problems. Moreover, we provide a better theoretical grounding for an empirical observation that a sharp minimum improves domain generalization (Cha et al., 2021) by pinpointing a clear dependence on sharpness.

3.3 CASE STUDY

To better demonstrate the relationship between sharpness and robustness, we provide two specific examples: (1) linear ridge regression; (2) two-layer diagonal neural networks for classification.

Example 3.6. In ridge regression models $\mathcal{S}(\hat{S}; A)$ has a reverse relationship to the regularization parameter λ . That is, the more the regularization parameter λ , the more the sharpness of the learned model could be. Following the previous notation, we have $\epsilon_1(\hat{S}; S) = \epsilon_1(\lambda; S) + \epsilon_d$ where ϵ_d has a smaller order than $\epsilon_1(\hat{S}; S)$ for large d (proof refer to Appendix E.1).

As suggested in Ali et al. (2019), let's consider a generic response model $(X; y; \Sigma)$ where $X \in \mathbb{R}^{n \times d}$; $d > n$. The least-square empirical minimizer of the ridge regression problem will be:

$$\hat{\beta} = \arg \min_{\beta} \frac{1}{2n} \|X\beta - y\|^2 + \frac{\lambda}{2} \|\beta\|^2 = (X^T X + \lambda I_n)^{-1} X^T y = M^{-1} X^T y \quad (9)$$

Let S be the training set. It's trivial to get the sharpness of a quadratic loss function where

$$\mathcal{S}(\hat{\beta}; S) = \lambda^{-1} \text{tr}(X^T X + \lambda I_n) = \lambda^{-1} \text{tr}(M) \quad (10)$$

It's obvious that both of the above two equations depend on the same matrix $X^T X + \lambda I_n$. For fixed training samples X , we have $\mathcal{S}(\hat{\beta}; S) = O(\lambda^{-1})$ in the limit of $\lambda \rightarrow 0$. Then it's clear that a higher penalty leads to a sharper minimum. This intuition is rigorously proven in Appendix E.1. According to Theorem 3.1 and Theorem 3.4, a sharper minimum probably associates with lower robustness constant $\mathcal{S}(\hat{S}; A)$. Thus it enjoys a lower OOD generalization error gap. In ridge regression, this phenomenon can be reflected by the regularization coefficient λ . Therefore, in general, the larger λ is, the lower the sharpness $\mathcal{S}(\hat{\beta}; S)$ and variance are. As a consequence, larger λ learns a more robust model resulting in a lower OOD generalization error gap. This idea is later verified in the distributional shift experiments, shown as Figure 2.

Example 3.7. We consider a classification problem using a 2-layer diagonal linear network with exp-loss. The robustness $\mathcal{S}(\hat{S}; A)$ has a similar relationship in Theorem 3.4. Given training set S after iteration $t > T$, $\epsilon_2 > 0$, $\mathcal{S}(\hat{S}; A) \leq \epsilon_2 \sup_{t \geq T} \mathcal{S}(\hat{S}_t; S)$.

In addition to the regression and linear models, we have obtained a similar relationship for 2-layer diagonal linear networks, which are commonly used in the kernel and rich regimes as well as in intermediate settings (Moroshko et al., 2020). Example 3.7 demonstrates that the relationship also holds true when the model is well-trained, even exp-loss does not satisfy the P condition. By extending our theorems to these more complex frameworks, we go beyond our initial assumptions and offer insights into broader applications. Later experiments on non-linear NN also support our statements. However, we still need a unified theorem for general function classes with fewer assumptions.

4 RELATED WORK

Despite various methods (Sun & Saenko, 2016; Sagawa et al., 2019; Shi et al., 2021; Sha eezadeh Abadeh et al., 2015; Li et al., 2018; Cha et al., 2021; Du et al., 2020; Zhang et al., 2022) that have been proposed to overcome the poor generalization brought by unknown distribution shifts, the underlying principles and theories still remain underexplored. As pointed out in Redko et al. (2020); Miller et al. (2021), different tasks that address distributional shifts, such as domain adaptation, OOD, and domain generalization, are collectively referred to as "transfer transductive learning" and share similar generalization theories. In general, the desired generalization bound will be split into In-Distribution/Domain (ID) generalization error and Out-of-Distribution/Domain (OOD) distance. Since Blitzer et al. (2007) establish a VC-dimension-based framework to estimate the domain shift gap by a divergence term, many following works make the effort to improve this term in the following decades, such as Discrepancy (Mansour et al., 2009), Wasserstein measurement (Courty et al., 2017; Shen et al., 2018), Integral Probability Metrics (IPM) (Zhang et al., 2019b; Ye et al., 2021) and χ^2 -divergence (Germain et al., 2016). Among them, new generalization tools like PAC-Bayes, Rademacher Complexity, and Stability are also applied. However, few of them discuss how the sharpness reacts to data distributional shifts.

Beyond this canonical framework, Ye et al. (2021) reformulate the OOD generalization problem and provide a generalization bound using the concepts of "variation" and "informativeness." The causal framework proposed in Peters et al. (2017); Rojas-Carulla et al. (2018) focuses on the impact of interventions on robust optimization over test distributions. However, none of these frameworks consider the optimization process of a model and how it affects OOD generalization. Inspired by previous investigation on the effect of sharpness on ID generalization (Lyu et al., 2022; Petzka et al., 2021), recent work in Cha et al. (2021) found that sharper minima can also improve OOD generalization. Nevertheless, they lack a sufficient theoretical foundation for the relationship between the "sharpness" of a model and OOD generalization, but end with a union bound of Blitzer et al. (2007)'s result. In this paper, we aim to provide a rigorous examination of this relationship.

5 EXPERIMENTS

In light of space constraints, we present only a portion of our experimental results to support the validity of our theorems and findings. For comprehensive results, please refer to the Appendix G

5.1 RIDGE REGRESSION IN DISTRIBUTIONAL SHIFTING

Following Duchi & Namkoong (2021), we investigated the ridge regression on distributional shift. We randomly generate $\mathbf{x} \in \mathbb{R}^d$ in spherical space, and data from the following generating process: $X \stackrel{iid}{\sim} \mathcal{N}(0, 1)$; $y = X \cdot \theta_0$. To simulate distributional shift, we randomly generate a perpendicular unit vector \mathbf{z} to \mathbf{x} . Let $\mathbf{z}_0; \mathbf{z}_1$ be the basis vectors, then shifted ground-truth will be computed from the basis by $\theta = \theta_0 \cos(\alpha) + \theta_1 \sin(\alpha)$. For the source domain, we use \mathbf{x} as our training distribution. We randomly sample 50 data points and train a linear classifier with a gradient descent of 3000 iterations. By minimizing the objective function in (9), we can get the empirical optimum $\hat{\theta}$. Then we gradually shift the distribution by increasing α to get different target domains. Along distribution shifting, the test loss $\mathcal{L}(\hat{\theta}; y)$ will increase. As shown in Figure 2, the test loss will culminate in around $\frac{\pi}{2}$ rads due to the maximum distribution shifting. Comparing different levels of regularization, we found that the larger L2-penalty brings lower OOD generalization error which is shown as darker purple lines. This plot bears out our intuition in the previous section. As stated in the aforementioned case, the sharpness of ridge regression should inversely depend on λ . Correspondingly, we compute sharpness using the definition equation (4) by averaging ten different results. For each trial, we use the same training and test data for every λ . The sharpness of each ridge regressor is shown in the legend of Figure 2. As we can see, larger λ leads to less sharpness.

5.2 SHARPER MINIMUM HURTS OOD GENERALIZATION

In our results, we proved that the upper bound of OOD generalization error involves the sharpness of the trained model. Here we empirically verified our theoretical insight. We follow the experiment setting in DomainBed (Gulrajani & Lopez-Paz, 2021). To easily compute the sharpness, we choose

Figure 2: OOD test losses increase along distributional shifting. The X-axis is the shifting angle and the Y-axis is the test loss of the model which is trained on distribution = 0. Lines are average test losses and shadows are variances of 10 trials. Larger regularization (darker color) causes a lower increase in test loss but smaller sharpness.

Figure 3: The relationship between out-of-domain test accuracy and model sharpness on RotatedMNIST dataset. Here we show 4 different OOD environments [15 ; 30 ; 45 ; 60 rotation] as the OOD test set respectively. Each marker denotes a minimum of an algorithm with a specific seed. The marker style means the models trained in the same environment.

the 4-layer MLP on RotatedMNIST dataset where RotatedMNIST is a rotation of MNIST handwritten digit dataset (LeCun, 1998) with different angles ranging from [15 ; 30 ; 45 ; 60 ; 75]. In this codebase, each environment refers to selecting a domain (a specific rotation angle) as the test domain/OOD test dataset while training on all other domains. After getting the trained model of each environment, we compute the sharpness using all domain training sets based on the implementation of Petzka et al. (2021). To this end, we plot the performances of Empirical Minimization Risk (ERM), SWAD (Cha et al., 2021), Mixup (Yan et al., 2020) and GroupDRO (Sagawa et al., 2019) with 6 seeds of each. Then we measure the sharpness of all these minima. Figure 3 shows the relationship between model sharpness and out-of-domain accuracy. The tendency is clear that at minima give better OOD performances. In general, different environments can not be plotted together due to different training sets. However, we found the middle environments are similar tasks and thus plot them together for a clearer trend. In addition, different algorithms lead to different feature scales which may affect the scale of the sharpness. To address this, we align their scales when putting them together. For more individual results, please refer to Figure 8 in the appendix.

5.3 COMPARISON OF GENERALIZATION BOUNDS

To analyze our generalization bounds, we follow the toy example experiments in Sagawa et al. (2020). In this experiment, the distribution shift terms and generalization error terms can be explicitly computed. Furthermore, their synthetic experiment considers the spurious correlation across distribution shifts which is now a general formulation of OOD generalization (Wald et al., 2021; Aubin et al., 2021; Yao et al., 2022). Consider data $\mathcal{D} = [X_{\text{core}}; X_{\text{spu}}]$ that consist of two features: core feature and spurious feature. The features are generated from the following rule:

$$X_{\text{core}} = \{y \in \{0, 1\}^N \mid y_i = 1; \text{core}\} \cup \{x \in \{0, 1\}^N \mid x_i = 1; \text{spu}\}$$

where $y \in \{0, 1\}^N$ is the label, and $x \in \{0, 1\}^N$ is the spurious attribute. Data with $y = 1$ forms the majority group of size n_{maj} , and data with $y = 0$ forms minority group of size n_{min} . Total number of training points $n = n_{\text{maj}} + n_{\text{min}}$. The spurious correlation probability $\gamma = \frac{n_{\text{maj}}}{n}$ denotes the probability of $y = 1$ in training data. In testing, we always have $\gamma_{\text{test}} = 0.5$. The metric, worst-group error Sagawa et al. (2019) is defined as

$$\text{Err}_{\text{wg}}(w) := \max_{i \in [2]} E_{x, y \sim \mathcal{D}_i} [\ell_0(w; (x; y))]$$

where ℓ_0 is the 0-1 loss in binary classification. Here we compare the robustness of our proposed OOD generalization bound and the baseline in Proposition 2.1. We also give the comparison to other baselines, like PAC-Bayes DA bound in the Appendix G.

Figure 4: Spurious feature synthetic experiment. Each dot represents a trained model. The dash curves are the smoothed function \hat{t} by the test data points. The baseline is Proposition 2(d): the generalization error of the logistic regression models with increasing the model size/correlation probability. (b): concentration error term in domain shift bound (e): comparison of distribution distance bounds (c),(f): comparisons of generalization bounds. Note that model size 500 is the overparameterized regime. The further the correlation probability is 0.5, the greater the distributional shift is.

Along model size We plot the generalization error of the random feature logistic regression along the model size increases in Figure 4(a). In this experiment, we follow the hyperparameter setup of Sagawa et al. (2020) by setting the number of points 500, data dimension $d = 200$ with 100 on each feature. majority fraction $\rho_{\text{maj}} = 0.9$ and noises $\sigma_{\text{spu}}^2 = 1$; $\sigma_{\text{core}}^2 = 100$. The worst-group error turns out to be nearly the same as the model size increases. However, in Figure 4(b), the error term in domain shift bound Proposition 2.1(A) will keep increasing when the model size is increasing. In contrast, our domain shift bound at order \sqrt{r} is independent of the model size which addresses the limitation of their bound. We follow Kawaguchi et al. (2022) to compute an inverse image of the covering in a randomly projected space (see details in appendix). We set the same value $k = 1,000$ in our experiment. Different from the baseline, ϵ is data dependent and leads to a constant concentration error term along with model size increases. Analogously, our OOD generalization bound will not explode as model size increases (shown in Figure 4(c)).

Along distribution shift In addition, we are interested in characterizing OOD generalization when test distribution shifts from train distribution by varying the correlation probability during data generation. As shown in Figure 4(d), when $\rho_{\text{maj}} = 0.5$, there is no distributional shift between training and test data due to no spurious features correlated to training data. Thus, the training and test distributions align closer and closer when $\rho_{\text{maj}} < 0.5$ and increase, resulting in an initial decrease in the test error for the worst-case group. However, as $\rho_{\text{maj}} > 0.5$ and deviates from 0.5, introducing the spurious features, a shift in the distribution occurs. This deviation is likely to impact the worst-case group differently, leading to an increase in the test error. As displayed in Figure 4(e) and Figure 4(f), our distribution distance and generalization bound can capture the distribution shifts but are tighter than the baseline.

6 CONCLUSION

In this paper, we provide a more interpretable and informative theory to understand Out-of-Distribution (OOD) generalization. Based on the notion of robustness, we propose a robust OOD bound that effectively captures the algorithmic robustness in the presence of shifting data distributions. In addition, our in-depth analysis of the relationship between robustness and sharpness further illustrates that sharpness has a negative impact on generalization. Overall, our results advance the understanding of OOD generalization and the principles that govern it.

7 ACKNOWLEDGEMENT

This research/project is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG-GC-2019-001-2A), (AISG Award No: AISG2-TC-2023-010-SGIL) and the Singapore Ministry of Education Academic Research Fund Tier 1 (Award No: T1 251RES2207). We would appreciate the contributions of Fusheng Liu, Ph.D. student at NUS.

REFERENCES

- Kartik Ahuja, Ethan Caballero, Dinghui Zhang, Jean-Christophe Gagnon-Audet, Yoshua Bengio, Ioannis Mitliagkas, and Irina Rish. Invariance principle meets information bottleneck for out-of-distribution generalization. *Advances in Neural Information Processing Systems*, 34:3438–3450, 2021.
- Alnur Ali, J Zico Kolter, and Ryan J Tibshirani. A continuous-time view of early stopping for least squares regression. *The 22nd international conference on artificial intelligence and statistics* pp. 1370–1378. PMLR, 2019.
- Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- Benjamin Aubin, Agnieszka S owik, Martin Arjovsky, Leon Bottou, and David Lopez-Paz. Linear unit-tests for invariance discovery. *arXiv preprint arXiv:2102.10867*, 2021.
- Peter Bandi, Oscar Geessink, Quirine Manson, Marcory Van Dijk, Maschenka Balkenhol, Meyke Hermsen, Babak Ehteshami Bejnordi, Byungjae Lee, Kyunghyun Paeng, Aoxiao Zhong, et al. From detection of individual metastases to classification of lymph node status at the patient level: the camelyon17 challenge. *IEEE Transactions on Medical Imaging*, 37:2018–2028, 2018.
- Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning* 79:151–175, 2010.
- John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman. Learning bounds for domain adaptation. *Advances in neural information processing systems*, 20:205–212, 2007.
- Jean-Christophe Bourin, Eun-Young Lee, and Minghua Lin. Positive matrices partitioned into a small number of hermitian blocks. *Linear Algebra and its Applications*, 438(5):2591–2598, 2013.
- Junbum Cha, Sanghyuk Chun, Kyungjae Lee, Han-Cheol Cho, Seunghyun Park, Yunsung Lee, and Sungrae Park. Swad: Domain generalization by seeking at minimum variance. *Advances in Neural Information Processing Systems*, 34:22405–22418, 2021.
- Jeremy M Cohen, Simran Kaur, Yuanzhi Li, J Zico Kolter, and Ameet Talwalkar. Gradient descent on neural networks typically occurs at the edge of stability. *arXiv preprint arXiv:2103.00065*, 2021.
- Nicolas Courty, Rémi Flamary, Amaury Habrard, and Alain Rakotomamonjy. Joint distribution optimal transportation for domain adaptation. *Advances in Neural Information Processing Systems* 2017.
- Simon S Du, Jayanth Koushik, Aarti Singh, and Barnabás Póczos. Hypothesis transfer learning via transformation functions. *Advances in neural information processing systems*, 30:305–314, 2017.
- Yingjun Du, Xiantong Zhen, Ling Shao, and Cees GM Snoek. Metanorm: Learning to normalize few-shot batches across domains. *International Conference on Learning Representations*, 2020.
- John Duchi and Hongseok Namkoong. Learning models with uniform performance via distributionally robust optimization. *arXiv preprint arXiv:1810.08750*, 2018.
- John C Duchi and Hongseok Namkoong. Learning models with uniform performance via distributionally robust optimization. *The Annals of Statistics*, 49(3):1378–1406, 2021.
- Laurent El Ghaoui and Hervé Lebret. Robust solutions to least-squares problems with uncertain data. *SIAM Journal on matrix analysis and applications*, 19:97–114, 1997.

- Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. *arXiv preprint arXiv:2010.01412*, 2020.
- Pascal Germain, Amaury Habrard, François Laviolette, and Emilie Morvant. A pac-bayesian approach for domain adaptation with specialization to linear classifiers. *International conference on machine learning*, PMLR, 2013.
- Pascal Germain, Amaury Habrard, François Laviolette, and Emilie Morvant. A new pac-bayesian perspective on domain adaptation. *International conference on machine learning*, PMLR, 2016.
- Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. *International Conference on Learning Representations*, 2021.
- 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.
- Kenji Kawaguchi, Zhun Deng, Kyle Luh, and Jiaoyang Huang. Robustness implies generalization via data-dependent generalization bounds. *International Conference on Machine Learning*, PMLR, 2022.
- Daniel Kifer, Shai Ben-David, and Johannes Gehrke. Detecting change in data streams. *DBN*, volume 4, pp. 180–191. Toronto, Canada, 2004.
- Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Bal-subramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, et al. Wilds: A benchmark of in-the-wild distribution shifts. *International Conference on Machine Learning*, pp. 5637–5664. PMLR, 2021.
- Masanori Koyama and Shoichiro Yamaguchi. Out-of-distribution generalization with maximal invariant predictor. 2020.
- Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. *Proceedings of the IEEE international conference on computer vision*, pp. 5542–5550, 2017.
- Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. Learning to generalize: Meta-learning for domain generalization. *Proceedings of the AAAI conference on artificial intelligence*, 2018.
- Jiashuo Liu, Zheyuan Hu, Peng Cui, Bo Li, and Zheyuan Shen. Heterogeneous risk minimization. *International Conference on Machine Learning*, pp. 6804–6814. PMLR, 2021.
- Kaifeng Lyu, Zhiyuan Li, and Sanjeev Arora. Understanding the generalization benefit of normalization layers: Sharpness reduction. *Advances in Neural Information Processing Systems*, 2022.
- Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. Domain adaptation: Learning bounds and algorithms. *arXiv preprint arXiv:0902.3430*, 2009.
- Song Mei and Andrea Montanari. The generalization error of random features regression: Precise asymptotics and the double descent curve. *Communications on Pure and Applied Mathematics*, (4):667–766, 2022.
- John P Miller, Rohan Taori, Aditi Raghunathan, Shiori Sagawa, Pang Wei Koh, Vaishaal Shankar, Percy Liang, Yair Carmon, and Ludwig Schmidt. Accuracy on the line: on the strong correlation between out-of-distribution and in-distribution generalization. *International Conference on Machine Learning*, PMLR, 2021.
- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2018.

- Edward Moroshko, Blake E Woodworth, Suriya Gunasekar, Jason D Lee, Nati Srebro, and Daniel Soudry. Implicit bias in deep linear classification: Initialization scale vs training accuracy in neural information processing systems. *Advances in Neural Information Processing Systems*, 33:22182–22193, 2020.
- Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.
- Henning Petzka, Michael Kamp, Linara Adilova, Cristian Sminchisescu, and Mario Boley. Relative atness and generalization. *Advances in Neural Information Processing Systems*, 34:18420–18432, 2021.
- Mohammad Pezeshki, Oumar Kaba, Yoshua Bengio, Aaron C Courville, Doina Precup, and Guillaume Lajoie. Gradient starvation: A learning proclivity in neural networks. *Advances in Neural Information Processing Systems*, 34:1256–1272, 2021.
- Alexandre Rame, Corentin Dancette, and Matthieu Cord. Fishr: Invariant gradient variances for out-of-distribution generalization. *International Conference on Machine Learning*, pp. 18347–18377. PMLR, 2022.
- Ievgen Redko, Emilie Morvant, Amaury Habrard, Marc Sebban, and Younès Bennani. A survey on domain adaptation theory: learning bounds and theoretical guarantees. preprint arXiv:2004.11829, 2020.
- Herbert Robbins. A remark on stirling's formula. *The American mathematical month*, 62(1):26–29, 1955.
- Mateo Rojas-Carulla, Bernhard Schölkopf, Richard Turner, and Jonas Peters. Invariant models for causal transfer learning. *The Journal of Machine Learning Research*, 2018.
- Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. arXiv preprint arXiv:1911.08731, 2019.
- Shiori Sagawa, Aditi Raghunathan, Pang Wei Koh, and Percy Liang. An investigation of why overparameterization exacerbates spurious correlation. *International Conference on Machine Learning*. PMLR, 2020.
- Soroosh Hashemi Abadeh, Peyman M Mohajerin Esfahani, and Daniel Kuhn. Distributionally robust logistic regression. *Advances in Neural Information Processing Systems*, 2015.
- Jian Shen, Yanru Qu, Weinan Zhang, and Yong Yu. Wasserstein distance guided representation learning for domain adaptation. *Proceedings of the AAAI Conference on Artificial Intelligence* 2018.
- Yuge Shi, Jeffrey Seely, Philip HS Torr, N Siddharth, Awni Hannun, Nicolas Usunier, and Gabriel Synnaeve. Gradient matching for domain generalization. arXiv preprint arXiv:2104.09937, 2021.
- Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *European conference on computer vision*, pp. 443–450. Springer, 2016.
- Yoav Wald, Amir Feder, Daniel Greenfeld, and Uri Shalit. On calibration and out-of-domain generalization. *Advances in neural information processing systems*, 2021.
- Olivia Wiles, Sven Gowal, Florian Stimberg, Sylvestre-Alvise Rebuf, Ira Ktena, Krishnamurthy Dvijotham, and Ali Taylan Cemgil. A fine-grained analysis on distribution shift. *International Conference on Learning Representations*, 2022.
- Huan Xu and Shie Mannor. Robustness and generalization. *Machine learning*, 86(3):391–423, 2012.
- Shen Yan, Huan Song, Nanxiang Li, Lincan Zou, and Liu Ren. Improve unsupervised domain adaptation with mixup training. arXiv preprint arXiv:2001.00677, 2020.
- Huaxiu Yao, Yu Wang, Sai Li, Linjun Zhang, Weixin Liang, James Zou, and Chelsea Finn. Improving out-of-distribution robustness via selective augmentation. arXiv preprint arXiv:2201.00299, 2022.

Haotian Ye, Chuanlong Xie, Tianle Cai, Ruichen Li, Zhenguo Li, and Liwei Wang. Towards a theoretical framework of out-of-distribution generalization. *Advances in Neural Information Processing Systems*, 2021.

Xiao Zhang, Yaodong Yu, Lingxiao Wang, and Quanquan Gu. Learning one-hidden-layer relu networks via gradient descent. *The 22nd international conference on artificial intelligence and statistics*, pp. 1524–1534. PMLR, 2019a.

Xingxuan Zhang, Linjun Zhou, Renzhe Xu, Peng Cui, Zheyang Shen, and Haoxin Liu. Towards unsupervised domain generalization. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4910–4920, 2022.

Yuchen Zhang, Tianle Liu, Mingsheng Long, and Michael Jordan. Bridging theory and algorithm for domain adaptation. *International Conference on Machine Learning*, PMLR, 2019b.

Han Zhao, Shanghang Zhang, Guanhang Wu, José MF Moura, Joao P Costeira, and Geoffrey J Gordon. Adversarial multiple source domain adaptation. *Advances in neural information processing systems* 31, 2018.

Kai Zhong, Zhao Song, Prateek Jain, Peter L Bartlett, and Inderjit S Dhillon. Recovery guarantees for one-hidden-layer neural networks. *International conference on machine learning*, pp. 4140–4149. PMLR, 2017.

A ADDITIONAL EXPERIMENTS

A.1 SHARPNESS VS. OOD GENERALIZATION ON PACS AND WILDS-CAMELYON17

Figure 5: The relationship between out-of-distribution (OOD) test accuracy on the test environment and model sharpness (of last FC layer) on the Wilds-Camelyon17 dataset. Each marker denotes a model trained using ERM with different seed and hyperparameters.

Figure 6: The relationship between OOD test accuracy and model sharpness on PACS dataset. Each marker denotes a model trained using ERM with different seeds and hyperparameters. The marker style shows the out-of-distribution test environment of the model.

To evaluate our theorem more deeply, we examine the relationship between our defined sharpness and OOD generalization error on larger-scale real-world datasets: Wilds-Camelyon17 (Bandi et al. (2018); Koh et al. (2021)) and PACS (Li et al. (2017)). Wilds-Camelyon17 dataset includes 455,954 tumor and normal tissue slide images from five hospitals (environments). One of the hospitals is assigned as the test environment by the dataset publisher. Distribution shift arises from variations in patient population, slide staining, and image acquisition. PACS dataset contains 9,991 images of 7 objects in 4 visual styles (environments): art painting, cartoon, photo, and sketch. Following the common setup in Gulrajani & Lopez-Paz (2021), each environment is used as a test environment in turn. We follow the practice in Petzka et al. (2021) to compute the sharpness using the Hessian matrix from the last Fully-Connected (FC) layer of each model. For the Wilds-Camelyon17 dataset, we test the sharpness of 18 ERM models trained with different random seeds and hyperparameters. Figure 5 shows the result. For the PACS dataset, we run 60 ERM models with different random seeds and hyperparameters for each test environment. To get a clearer correlation, we align the points from 4 environments by their mean performance. Figure 6 shows the result. From the two figures, we can observe a clear correlation between sharpness and out-of-distribution (OOD) accuracy. Sharpness tends to hurt the OOD performance of the model. The result is consistent with what we report in Figure 3. It shows that the correlation between sharpness and OOD accuracy can also be observed on large-scale datasets.

B NOTATIONS AND DEFINITIONS

Notations We use $[n]$ denote the integers set $\{1, \dots, n\}$. $\|\cdot\|_2$ represents ℓ_2 -norm $\|\cdot\|_2$ for short. Without loss of generality, we use $f(\cdot; x; y)$ for the loss function of model on data pair $(x; y)$, which is denoted as $\ell(z)$ and we use $n; d$ for training set size and input dimension. Note that we generally follow the notations in the original papers.

- $L_S; L_T$: expected risk of the source domain and target domain, respectively. The corresponding empirical version will be $\hat{L}_S; \hat{L}_T$
- $\mathcal{C}_i; \mathcal{G}_{i=1}^K$: K partitions on sample space \mathcal{X} and \mathcal{C}_i denotes each partitioned set.
- $D_S; D_T$: distributions of source and target domain. Their sampled dataset will be denoted as $\mathcal{D}_S; \mathcal{D}_T$ accordingly.

- θ : In our setting, $\theta = (w; f; g^m)$ denotes the model parameters where w is the trainable parameters and f, g^m are the random features (we also use $a = [a_1; \dots; a_m]$ for short notation in many places)
- \mathcal{W} : is the minimizer of empirical loss.
- M : upper bound of the loss function.
- $S = \{(x_i; y_i)\}_{i=1}^n = X = [x_1; \dots; x_n]$: training data of size n .

Definition B.1 (Robustness, Xu & Mannor (2012)) A learning algorithm \mathcal{A} on training set S is $(K; \epsilon)$ -robust, for $K \geq 2$, if Z can be partitioned into K disjoint sets, denoted by $C_k, k=1, \dots, K$, such that for all $S \subseteq Z; z \in Z$ we have

$$|\mathcal{A}(S; z) - \mathcal{A}(S; s)| \leq \epsilon \quad \forall S \subseteq Z; z, s \in C_k; |S \cap C_k| \geq \frac{n}{K}$$

Definition B.2 (Hessian Lipschitz continuous) For a twice differentiable function $f: \mathbb{R}^n \rightarrow \mathbb{R}$, it has L -Lipschitz continuous Hessian for domain C_i if

$$\| \nabla^2 f(y) - \nabla^2 f(x) \| \leq L \|y - x\|$$

where $L_i > 0$ depends on input domain C_i and $\| \cdot \|$ is L_2 norm. Then for all K domains $\{C_i\}_{i=1}^K$, let $L := \max_i L_i, i \in [K]$ be the uniform Lipschitz constant, so we have

$$\| \nabla^2 f(y) - \nabla^2 f(x) \| \leq L \|y - x\| \leq L \|y - x\|; \forall (x; y) \in C_i$$

which is uniformly bounded with L .

Lemma B.3 (Hessian Lipschitz Lemma) If f is twice differentiable and has L -Lipschitz continuous Hessian, then

$$|f(y) - f(x) - \nabla f(x) \cdot (y - x) - \frac{1}{2} \nabla^2 f(x) (y - x)(y - x)| \leq \frac{L}{6} \|y - x\|^3$$

Gegenbauer Polynomials

We briefly define Gegenbauer polynomials here whose details can be found in Appendix of Mei & Montanari (2022). First, we denote $\mathbb{S}^{d-1}(r) = \{x \in \mathbb{R}^d : \|x\| = r\}$ as the uniform spherical distribution with the radius r on $d-1$ manifold. Let μ_d be the probability measure on \mathbb{S}^{d-1} and the inner product in functional space $L^2(\mathbb{S}^{d-1}; \mu_d)$ denoted as $\langle \cdot, \cdot \rangle_{L^2}$ and $\| \cdot \|_{L^2}$:

$$\langle f; g \rangle_{L^2} = \int_{\mathbb{S}^{d-1}} f(x)g(x) \mu_d(dx)$$

For any function $f \in L^2(\mathbb{S}^{d-1}; \mu_d)$, where μ_d is the distribution of $x; y \sim \mu_d(x; y \in \mathbb{S}^{d-1}(\bar{r}))$, the orthogonal basis $Q_t^{(d)}$ forms the Gegenbauer polynomial of degree $t \geq 0$, its spherical harmonics coefficients $c_{d;t}(\cdot)$ can be expressed as:

$$c_{d;t}(\cdot) = \int_{\mathbb{S}^{d-1}} f(x) Q_t^{(d)}(\langle \cdot, x \rangle) \mu_d(dx)$$

then the Gegenbauer generating function holds in $L^2(\mathbb{S}^{d-1}; \mu_d)$ sense

$$f(x) = \sum_{k=0}^{\infty} c_{d;t}(\cdot) N_{d;t} Q_t^{(d)}(\langle \cdot, x \rangle)$$

where $N_{d;t}$ is the normalized factor depending on the norm of input.

B.1 ASSUMPTIONS

We discuss and list all assumptions we used in our theorems. The purposes are to offer clarity regarding the specific assumptions required for each theorem and ensure that the assumptions made in our theorems are well-founded and reasonable, reinforcing the validity and reliability of our results.

OOD Generalization

(Setting): Given a full sample space, source and target distributions are two different measures over this whole sample domain. The purpose is to study the robust algorithms in the OOD generalization setting.

(Assumptions): For any sample $z \in Z$, the loss function is bounded $\ell(z) \in [0; M]$. This assumption generally follows the original paper [Xu & Mannor \(2012\)](#). While it is possible to relax this assumption and derive improved bounds, our primary objective is to formulate a framework for robust OOD generalization and establish a clear connection with the optimization properties of the model.

Robustness and Sharpness

(Setting): In order to give a fine-grained analysis, we follow the common choice where a two-layer ReLU Neural Network function class is widely analyzed in most literature, i.e. Neural Tangent Kernel, non-kernel (rich) regime [Moroshko et al. \(2020\)](#) and random feature models. Among them, we select the following random feature models as our function class:

$$f(w; A; x) = \sum_{i=1}^n w_i \sigma(x; a_i) : w_i \in \mathbb{R}; a_i \sim \text{Unif}(S^{d-1}(\sqrt{d})); i \in [m]$$

where n is the hidden size and $A = [a_1; \dots; a_m]$ contains random vectors uniformly distributed on n -dim hypersphere whose surface is a 1 manifold. $\sigma(x) = \max(0, x)$ If $\sigma(x) = \max(0, x)$ denotes the ReLU activation function and $\mathbb{1}$ is the indicator function $w = [w_1; \dots; w_m]^T$ is the trainable parameter. We choose the common loss functions: (1) Homogeneity in regression; (2) (Binary) Cross-Entropy Loss; (3) Negative Log Likelihood (NLL) loss;

(Assumptions):

(i) Let $C_i; i \in [K]$ be any set from whole partition $\bigcup_{i=1}^K C_i$, we assume $z \in C_i$, the loss function $\ell_{w; A}(z)$ satisfies L -Hessian Lipschitz for all $i \in [K]$ (See details in Definition B.2). Note that we only require this assumption to hold within each partition, instead of holding globally. In general, the smoothness and convexity condition is actually equivalent to locally convex which is a weak assumption for most function classes.

(ii) Consider an optimization problem in each partition $C_i; i \in [K]$. Let one of the training points $z_i(A) \in S \setminus C_i$ be the initial point and $d_i(A) \in M_i$ is the corresponding nearest local minima where M_i is the local minima set of partition C_i . For some $\epsilon_i(L) > 0$, we assume $\|z_i(A) - d_i(A)\| \leq \epsilon_i(L)$ holds a.s.. It ensures the hessian $\text{Hess}(\ell_{w; A}(z_i(A)))$ has the lower bound. Similar conditions and estimations can be found in [Zhang et al. \(2019a\)](#).

(iii) To simplify the computation of probability, we assume $z \sim \text{Unif}(S^{d-1}(\sqrt{d}))$ obeys a rotationally invariant distribution.

(iv) For Corollary 3.5, we make additional assumptions that loss function satisfies a bounded condition where the second derivative $\text{Hess}(\ell_{w; A}(z)) \in \text{Unif}(S^{d-1}(\sqrt{d})); j \in \{1, \dots, d\}$ with respect to its argument w should be bounded by $[M_1; M_2]$. Note, we consider the case where the data $z \sim \text{Unif}(S^{d-1}(\sqrt{d}))$ while $m = \text{Poly}(d)$ is to ensure the positive definiteness of $\sum_{i=1}^m a_i a_i^T \in \mathbb{R}^{d \times d}$ almost surely.

C PROOF TO DOMAIN SHIFT

Lemma C.1. Let \hat{D}_T be the empirical distribution of size n drawn from D_T . The loss ℓ is upper bounded by M . With probability at least $1 - \delta$ (over the choice of the samples), for every trained on S , we have

$$L_T(f) - L_S(f) + M d_{(K)}(S; \hat{D}_T) + (S) + 2M \sqrt{\frac{2K \ln 2 + 2 \ln(1/\delta)}{n}} \quad (11)$$

where

$$d_{(K)}(S; \hat{D}_T) := \sum_{i=1}^K \frac{n_i(S)}{n} - \frac{n_i(\hat{D}_T)}{n} \quad (12)$$

and $n_i(S); n_i(\hat{D}_T)$ are the number of samples from S and \hat{D}_T that fall into the set C_i , respectively.

Proof. In the following generalization statement, we use $\ell(f; z)$ to denote the error obtained with input z and hypothesis function f for better illustration. By definition we have,

$$L_T(f) - L_S(f) := E_{z \in D_T}(\ell(f; z)) - E_{z \in S}(\ell(f; z)) \quad (13)$$

Then we make the K partitions for source distribution D_S . Let n_i be the size of collection set of points x fall into the partition C_i where n_i is the i.i.d. multinomial random variable with $(p_S(C_1); \dots; p_S(C_K))$. We use parallel notation for target distribution D_T with S_i^0 $(p_T(C_1); \dots; p_T(C_K))$. Since

$$\begin{aligned} E_{z \in D_S}(\ell(f; z)) &= \sum_{i=1}^K E_{z \in D_S}(\ell(f; z) | z \in C_i) p_S(C_i) \\ E_{z \in D_T}(\ell(f; z)) &= \sum_{i=1}^K E_{z \in D_T}(\ell(f; z) | z \in C_i) p_T(C_i) \end{aligned} \quad (14)$$

and thus we have

$$\begin{aligned} L_T(f) - L_S(f) &= \sum_{i=1}^K (E(\ell(f; z) | z \in C_i) p_T(C_i) - E(\ell(f; z) | z \in C_i) p_S(C_i)) \\ &= \sum_{i=1}^K (E(\ell(f; z) | z \in C_i) (p_T(C_i) - p_S(C_i))) \\ &\quad + \sum_{i=1}^K [E(\ell(f; z) | z \in C_i) - E(\ell(f; z) | z \in C_i)] p_S(C_i) \\ &\quad + \sum_{i=1}^K (E(\ell(f; z) | z \in C_i) (p_T(C_i) - p_S(C_i))) + (S; A) \end{aligned} \quad (15)$$

If we sample empirical distributions $S; \hat{D}_T$ of size n each drawn from D_S and D_T , respectively. $(n_1; \dots; n_K)$ are the i.i.d. random variables belongs to S . We use the parallel notation for target distribution.

$$d_{(K)}(S; \hat{D}_T) := \sum_{i=1}^K \frac{n_i(S)}{n} - \frac{n_i(\hat{D}_T)}{n} \quad (16)$$

Further, we have

$$\begin{aligned} &\sum_{i=1}^K (E(\ell(f; z) | z \in C_i) (p_T(C_i) - p_S(C_i))) - \sum_{i=1}^K (E(\ell(f; z) | z \in C_i) \frac{n_i(\hat{D}_T)}{n} - \frac{n_i(S)}{n}) \\ &= \sum_{i=1}^K (E(\ell(f; z) | z \in C_i) p_T(C_i) \frac{n_i(\hat{D}_T)}{n} - \frac{n_i(\hat{D}_T)}{n} \sum_{i=1}^K (E(\ell(f; z) | z \in C_i) p_S(C_i) \frac{n_i(S)}{n})) \\ &\quad + \sum_{i=1}^K p_T(C_i) \frac{n_i(\hat{D}_T)}{n} - \sum_{i=1}^K p_S(C_i) \frac{n_i(S)}{n} \end{aligned} \quad (17)$$

With Bretaganolle-Huber-Carol inequality we have

$$\sum_{i=1}^K \frac{n_i(S)}{n} - p_S(C_i) \leq \frac{r}{n} \frac{2K \ln 2 + 2 \ln(1/\epsilon)}{n} \quad (18)$$

To integrate these two inequalities, we have

$$\begin{aligned} & \sum_{i=1}^K (E(\ell(f; z^0) | z^0 \sim C_i)) (p_t(C_i) - p_s(C_i)) \\ & \sum_{i=1}^K (E(\ell(f; z^0) | z^0 \sim C_i)) \frac{n_i(\hat{D}_T)}{n} - \frac{n_i(S)}{n} + 2M \sqrt{\frac{2K \ln 2 + 2 \ln(1 - \delta)}{n}} \\ & Md_{(\cdot; K)}(S; \hat{D}_T) + 2M \sqrt{\frac{2K \ln 2 + 2 \ln(1 - \delta)}{n}}. \end{aligned} \quad (19)$$

In summary with probability $1 - \delta$ we have

$$L_T(f) - L_S(f) + Md_{(\cdot; K)}(S; \hat{D}_T) + (S) + 2M \sqrt{\frac{2K \ln 2 + 2 \ln(1 - \delta)}{n}} \quad (20)$$

which completes the proof. \square

With the result of the domain (distribution) shift and the relationship between sharpness and robustness, we can move forward to the final OOD generalization error bound. First, we state the context of ID robustness bound in [Xu & Mannor \(2012\)](#) as follows.

Lemma C.2 (Xu et al. [Xu & Mannor \(2012\)](#)) Assume that for all $h \in \mathcal{H}$ and $z \in \mathcal{Z}$, the loss is upper bounded by M , i.e., $\ell(h; z) \leq M$. If the learning algorithm A is (K, δ) -robust, then for any $\delta > 0$, with probability at least $1 - \delta$ over an iid draw of n samples $S = (z_i)_{i=1}^n$, it holds that:

$$E_z [L(A_S; z)] \leq \frac{1}{n} \sum_{i=1}^n \ell(A_S; z_i) + (S) + M \sqrt{\frac{2K \ln 2 + 2 \ln(1 - \delta)}{n}}$$

As the conclusive results, we briefly prove the following result by summarizing [Lemma C.2](#) and [Lemma C.1](#).

Theorem C.3 (Restatement of [Theorem 3.1](#)) Let \hat{D}_T be the empirical distribution of size n drawn from D_T . The loss is upper bounded by M . With probability at least $1 - \delta$ (over the choice of the samples), for every f trained on S , we have

$$L_T(f) - L_S(f) + Md_{(\cdot; K)}(S; \hat{D}_T) + 2(S) + 3M \sqrt{\frac{2K \ln 2 + 2 \ln(2 - \delta)}{n}}. \quad (21)$$

Proof. Firstly, with [Lemma C.1](#) and probability at least $\frac{\delta}{2}$, we have

$$L_T(f) - L_S(f) + Md_{(\cdot; K)}(\hat{D}_S; \hat{D}_T) + 2M \sqrt{\frac{2K \ln 2 + 2 \ln(2 - \delta)}{n}} + (S)$$

Secondly, with [Lemma C.2](#) ([Xu & Mannor \(2012\)](#) [Theorem 3](#)) and probability at least $\frac{\delta}{2}$, we have

$$L_S(f) - \hat{L}_S(f) \leq (S) + M \sqrt{\frac{2K \ln 2 + 2 \ln(2 - \delta)}{n}}$$

By taking the union bound, we conclude our final result that with probability at least $1 - \delta$

$$L_T(f) - \hat{L}_S(f) + 3M \sqrt{\frac{2K \ln 2 + 2 \ln(2 - \delta)}{n}} + 2(S) + Md_{(\cdot; K)}(S; \hat{D}_T) \quad (22)$$

\square

Here (S) is the robustness constant that we can replace with any sharpness measure.

C.1 PROOF TO COROLLARY 3.2

Definition C.4. $d_{F, F}(D_T; D_S) := 2 \sup_{A(f) \in \mathcal{A}_{F, F}} |\Pr_{D_S}(A(f)) - \Pr_{D_T}(A(f))|$ and F, F is defined as:

$$F, F := \{f \oplus f(x) - f^0(x) : f; f^0 \in F, g\}$$

where \oplus is the XOR operator e.g. $(f^0(x) \oplus f(x))$.

Lemma C.5 (Lemma 2 Zhao et al. (2018)) If $\text{VCdim}(F) = d^0$, then $\text{VCdim}(F, F) \leq 2d^0$.

Proposition C.6 (Zhao et al. (2018)) Let F be a hypothesis class with pseudo dimension $\text{VCdim}(F) = d^0$. If \mathcal{D}_S is the empirical distributions generated with i.i.d. samples from source domain, and \mathcal{D}_T is the empirical distribution on the target domain generated from samples without labels, then with probability at least $1 - \delta$, for all $f \in F$, we have:

$$\begin{aligned} L_T(f) - \mathbb{E} L_S(f) + E &+ \frac{r}{n} \frac{2d^0 \log \frac{en}{d^0}}{n} + \frac{s}{2n} \frac{\log \frac{2}{\delta}}{2n} \\ &+ \frac{1}{2} d_{F, F}(\mathcal{D}_T; \mathcal{D}_S) + 4 \frac{s}{n} \frac{2d^0 \ln(2n) + \ln \frac{4}{\delta}}{n} \end{aligned} \quad (23)$$

(Empirical div Error)

where $E = \mathbb{E} L_S(f^*) + \mathbb{E} L_T(f^*)$ is the total error of best hypothesis over source and target domain.

Proof. With Lemma 4 (Zhao et al., 2018), we have

$$L_T(f) - \mathbb{E} L_S(f) + \frac{1}{2} d_{F, F} + E$$

where

$$E = \inf_{f \in F} L_S(f^0) + L_T(f^0):$$

Lemma 6 (Zhao et al., 2018), which is actually Lemma 1 in (Ben-David et al., 2010), shows the following results

$$d_{F, F}(D_T; D_S) \leq d_{F, F}(\mathcal{D}_T; \mathcal{D}_S) + 4 \frac{s}{n} \frac{\text{VCdim}(F, F) \ln(2n) + \ln \frac{2}{\delta}}{n}:$$

As suggested in Zhao et al. (2018), $\text{VCdim}(F, F)$ is at most $2d^0$. Further, with Theorem 2 (Ben-David et al., 2010), we have at probability at least $\frac{1}{2}$

$$\begin{aligned} L_T(f) - L_S(f) + \frac{1}{2} d_{F, F}(\mathcal{D}_T; \mathcal{D}_S) &+ 4 \frac{s}{n} \frac{\text{VCdim}(F, F) \ln(2n) + \ln \frac{2}{\delta}}{n} + E \\ L_S(f) + \frac{1}{2} d_{F, F}(\mathcal{D}_T; \mathcal{D}_S) &+ 4 \frac{s}{n} \frac{2d^0 \ln(2n) + \ln \frac{2}{\delta}}{n} + E \end{aligned} \quad (24)$$

Using in-domain generalization error Lemma 11.6 (Mohri et al., 2018), with probability at least $\frac{1}{2}$ the result is

$$L_S(f) - \mathbb{E} L_S(f) + M \frac{r}{n} \frac{2d^0 \log \frac{en}{d^0}}{n} + M \frac{s}{2n} \frac{\log \frac{1}{\delta}}{2n}$$

Note in Zhao et al. (2018), $M = 1$ for the normalized regression loss. Combine them all, we conclude the proof. \square

Corollary C.7. If $K \leq 1$; $M = 1$, domain shift bound $|L_T(f) - L_S(f)|$ will be reduced to (Empirical div Error) in Proposition C.6 where

$$|L_T(f) - L_S(f)| \leq \frac{1}{2} d_{F, F}(D_S; D_T) + (\text{Empirical div Error}) \quad (25)$$

Proof. According to Theorem C.3, we have

$$L_T(f) - L_S(f) = \int_{\mathcal{Z}} E(\ell(f; z^0) | z^0 \sim C_i) p_t(C_i) - E(\ell(f; z) | z \sim C_i) p_s(C_i) \quad (26)$$

If $K \geq 1$, let's define a domain $\mathcal{U} := \bigcup_{i=1}^{S_1} C_i$. The equation (26) will be

$$\begin{aligned} L_T(f) - L_S(f) &= \int_{\mathcal{Z}^{02U}} \ell(f; z^0) p_t(z^0) dz - \int_{\mathcal{Z}^{2U}} \ell(f; z) p_s(z) dz \\ &= \int_{z^0 \in D_T} \ell(f; z^0) p_t(z^0) dz - \int_{z \in D_S} \ell(f; z) p_s(z) dz \\ &= E_{z^0 \sim D_T} \ell(f; z^0) - E_{z \sim D_S} \ell(f; z) \end{aligned} \quad (27)$$

In this case, we have,

$$\begin{aligned} &|L_T(f) - L_S(f)| \\ &= \int_{\mathcal{Z}} |E_{z^0 \sim D_T} \ell(f; z^0) - E_{z \sim D_S} \ell(f; z)| \\ &\quad | \Pr_{D_T}(\ell(f; x^0; y^0) > t) - \Pr_{D_S}(\ell(f; x; y) > t) | dt \\ &= \int_0^1 | \Pr_{D_T}(\ell(f; x^0; y^0) > t) - \Pr_{D_S}(\ell(f; x; y) > t) | dt \quad (M = 1) \\ &\leq \sup_{t \in [0;1]} \sup_{f \in \mathcal{F}} | \Pr_{D_T}(\ell(f; x^0; y^0) > t) - \Pr_{D_S}(\ell(f; x; y) > t) | \\ &\leq \sup_{A(f) \in \mathcal{A}_{\mathcal{F}}} | \Pr_{D_T}(A(f)) - \Pr_{D_S}(A(f)) | \\ &= \frac{1}{2} d_{\mathcal{F}}(D_S; D_T) \quad (\text{Empirical div error}) \end{aligned} \quad (28)$$

where $\mathcal{A}_{\mathcal{F}}$ represents a learning algorithm under the hypothesis $\mathcal{F} = \{f(x) = f^0(x) : f \in \mathcal{F}\}$, which completes the proof. \square

D SHARPNESS AND ROBUSTNESS

Lemma D.1 (positive definiteness of Hessian) Let w_{\min} be the minimum value of w_j and $x = \arg \min_{x \in \text{Unif}(S^{d-1}(r))} \ell(f; A; x; y)$. For any $A = (a_1; \dots; a_m); a_i \in \text{Unif}(S^{d-1}(r))$, $\delta_i \in [0, 1]$ denote $\tilde{\lambda}(d; m) = \min_{i=1, \dots, m} (a_i^T a_i) > 0$ be the minimum eigenvalue, where $\mathcal{Q}_t = \prod_{i=1}^m (1 - \delta_i) \prod_{i=1}^m \mathcal{Q}_t^{(d)}(a_i; a_i) = d$ is the polynomial product constant. $M_1 = \text{Poly}(d)$, the hessian $H(x)$ can be lower bound by

$$E_{x \in \text{Unif}(S^{d-1}(r))} H(x) \geq \frac{w_{\min}^2 \tilde{\lambda}(d; m) M_1}{d} I_d \quad (29)$$

Proof. As suggested in Lemma D.6 of (Zhong et al., 2017), we have a similar result to bound the local positive definiteness of Hessian. By previous definition, the Hessian $H(x)$ has a following partial order

$$\begin{aligned} H(x) &= \frac{D_f^2(x; y)}{d} \otimes_{i=1}^n \otimes_{j=1}^n w_i w_j a_i^T a_j^T \ell(a_i^T x) \ell(a_j^T x) A \\ &\geq \frac{M_1}{d} \otimes_{i=1}^n \otimes_{j=1}^n w_i w_j a_i^T a_j^T \ell(a_i^T x) \ell(a_j^T x) A \\ &\geq \frac{w_{\min}^2 M_1}{d} \otimes_{i=1}^n \otimes_{j=1}^n w_i w_j a_i^T a_j^T \ell(a_i^T x) \ell(a_j^T x) A \end{aligned} \quad (30)$$

For the ReLU activation function, we further have

$$\sigma(a \cdot x) = \sigma(a \cdot x) \quad (31)$$

We extend the $L^2([-\bar{d}, \bar{d}]; \mu)$ (where μ is the distribution of $x_1; x_2 = \bar{d}$) by Gegenbauer polynomials that

$$Q_t^{(d)}(x) = \sum_{t=0}^{\infty} N_{d,t} Q_t^{(d)}(x/\bar{d}) \quad (32)$$

Let $A = (a_1; \dots; a_m) \in \mathbb{R}^{m \times d}$. We assume $x \sim \text{Unif}(S^{d-1}(\bar{d}))$. Lemma C.7 in (Mei & Montanari, 2022), suggests that

$$U = E_x \sum_{i,j=1}^m (a_i \cdot x/\bar{d}) (a_j \cdot x/\bar{d}) \in \mathbb{R}^{m \times m} \quad (33)$$

which shows matrix U is a positive definite matrix. Similarly, taking the expectation over terms in RHS of (30) bracket can be rewritten as

$$E_x \sum_{i=1}^n \sum_{j=1}^n a_i a_j \sigma(a_i \cdot x/\bar{d}) \sigma(a_j \cdot x/\bar{d}) = \sum_{i=1}^n \sum_{j=1}^n a_i a_j E_x [(a_i \cdot x/\bar{d}) (a_j \cdot x/\bar{d})] \quad (34)$$

Besides, we have the following property of Gegenbauer polynomials,

- For $x, y \in S^{d-1}(\bar{d})$

$$\int_{L^2(S^{d-1}(\bar{d}); \mu)} Q_j^{(d)}(hx; i) Q_k^{(d)}(hy; i) d\mu = \frac{1}{N_{d;k}} \delta_{jk} Q_k^{(d)}(hx; yi)$$

- For $x, y \in S^{d-1}(\bar{d})$

$$Q_k^{(d)}(hx; yi) = \frac{1}{N_{d;k}} \sum_{i=1}^{N_{d;k}} Y_{ki}^{(d)}(x) Y_{ki}^{(d)}(y)$$

where spherical harmonics $\{Y_{ij}^{(d)}\}_{i,j=1}^{N_{d,i}}$ forms an orthonormal basis which gives the following results

$$\begin{aligned} E_x [(a_i \cdot x/\bar{d}) (a_j \cdot x/\bar{d})] &= \sum_{t=0}^{\infty} \frac{1}{N_{d,t}^2} E_x [Q_t^{(d)}(a_i \cdot x/\bar{d}) Q_t^{(d)}(a_j \cdot x/\bar{d})] \\ &= \sum_{t=0}^{\infty} \frac{1}{N_{d,t}^2} Q_t^{(d)}(a_i; a_j/\bar{d}) = G_{ij} < 1 : \end{aligned} \quad (35)$$

Hence, we have

$$\sum_{i=1}^n \sum_{j=1}^n a_i a_j E_x [(a_i \cdot x/\bar{d}) (a_j \cdot x/\bar{d})] = \sum_{i=1}^n a_i a_i G_{ii} + O(1/d) \text{Var}(a) \quad (36)$$

Since $\text{Poly}(d)$ and $a_{i,2[m]}$ are i.i.d, then $\text{rank} \sum_{i=1}^m a_i a_i G_{ii} = \text{rank} AA^T = d$. Let $\tilde{w}(d; m) = E_{a \sim \text{Poly}(d)} \min_{i=1}^m a_i a_i G_{ii} > 0$ we have

$$E_x H(x) \geq \frac{\tilde{w}_{\min}^2 \tilde{w}(d; m) M_1}{d} I_d \quad (37)$$

□

Lemma D.2. Let $\bigcup_{k=1}^K C_k$ be the whole domain, the notion of $(\mathcal{S}; A)$ -robustness is described by

$$(\mathcal{S}; A), \max_{C_i} \sup_{z; z^0 \in C_i; z \in \mathcal{S}} \mathbb{E}_j \|\nabla_{\mathcal{W}; A}(z) - \nabla_{\mathcal{W}; A}(z_i^0)\|_j$$

Define M_i be the set of global minima \mathcal{M}_i , where

$$M_i = \{z \in C_i \mid f(z; A) = \min_{z \in C_i} f(z; A)\}$$

suppose for some maximum training loss point

$$z_i(A) \in \arg \max_{z \in C_i \setminus \mathcal{S}} f(z; A) - f(z_i(A))$$

there exists $z_i(A) \in M_i$ where

$$z_i(A) = \arg \min_{z \in M_i} \|z - z_i(A)\|_k$$

such that $\|z_i(A) - z_i^0(A)\|_k \leq \frac{\epsilon_i(L)}{L}$ almost surely hold for any $A \in \text{Unif}(\mathcal{S}^d)$ and for any A , $\nabla_{\mathcal{W}; A}(z)$ is L -Hessian Lipschitz continuous. Then $(\mathcal{S}; A)$ can be bounded by

$$(\mathcal{S}; A) \leq \max_{i \in [K]} \frac{\epsilon_i(L)^2}{2L^2} + r^2 \mathbb{E}_{z_i(A)} \|\nabla_{\mathcal{W}; A}(z_i(A))\|_j^2 + \frac{4\epsilon_i(L)}{3}$$

Proof. Let $z \in \mathcal{S}$ be a collection of $(x; y)$ from the training set and z_i^0 denote any collection from the set C_i . We define local minima set M_i (which is the global minima set \mathcal{M}_i). Assume that for some maximum point $z_i(A) \in \arg \max_{z \in C_i \setminus \mathcal{S}} f(z; A)$, there exists $z_i(A) \in M_i$ almost surely for all $A \in \text{Unif}(\mathcal{S}^d)$ such that

$$z_i(A) = \arg \min_{z \in M_i} f(z; A) := f(z; A) \text{ s.t. } \|z_i(A) - z_i^0(A)\|_k \leq \frac{\epsilon_i(L)}{L} \quad (38)$$

By definition, $(\mathcal{S}; A)$ can be rewritten as

$$\begin{aligned} (\mathcal{S}; A) &= \max_{i \in [K]} \sup_{z; z^0 \in C_i; z \in \mathcal{S}} \mathbb{E}_j \|\nabla_{\mathcal{W}; A}(z) - \nabla_{\mathcal{W}; A}(z_i^0)\|_j \\ &= \max_{i \in [K]} \sup_{z \in C_i \setminus \mathcal{S}; z \in M_i} \|\nabla_{\mathcal{W}; A}(z) - \nabla_{\mathcal{W}; A}(z_i(A))\|_j \\ &= \max_{i \in [K]} \|\nabla_{\mathcal{W}; A}(z_i(A)) - \nabla_{\mathcal{W}; A}(z_i(A))\|_j \end{aligned} \quad (39)$$

According to Lemma B.3, we have

$$\begin{aligned} (\mathcal{S}; A) &= \max_{i \in [K]} \|\nabla_{\mathcal{W}; A}(z_i(A)) - \nabla_{\mathcal{W}; A}(z_i(A))\|_j \\ &\stackrel{(i)}{=} \max_{i \in [K]} \mathbb{E}_r \|\nabla_{\mathcal{W}; A}(z_i(A)); z_i(A) - z_i(A)\|_j \\ &\quad + \frac{1}{2} r^2 \mathbb{E}_{z_i(A)} \|\nabla_{\mathcal{W}; A}(z_i(A))\|_j^2 \|z_i(A) - z_i(A)\|_k^2 + \frac{L}{6} \|z_i(A) - z_i(A)\|_k^3 \\ &= \max_{i \in [K]} \frac{1}{2} r^2 \mathbb{E}_{z_i(A)} \|\nabla_{\mathcal{W}; A}(z_i(A))\|_j^2 \|z_i(A) - z_i(A)\|_k^2 + \frac{L}{6} \|z_i(A) - z_i(A)\|_k^3 \\ &\quad + \max_{i \in [K]} \frac{1}{2} r^2 \mathbb{E}_{z_i(A)} \|z_i(A) - z_i(A)\|_k^2 \\ &\quad + \frac{L}{6} \|z_i(A) - z_i(A)\|_k^3 \quad (\text{Cauchy-Schwarz}) \end{aligned} \quad (40)$$

where (i) support by the fact $\nabla_{\mathcal{W}; A}(z) = 0$. With Lipschitz continuous Hessian we have

$$\|r^2 \mathbb{E}_{z_i(A)} \|\nabla_{\mathcal{W}; A}(z_i(A))\|_j^2 \|z_i(A) - z_i(A)\|_k^2 + \frac{L}{6} \|z_i(A) - z_i(A)\|_k^3 \leq r^2 \mathbb{E}_{z_i(A)} \|\nabla_{\mathcal{W}; A}(z_i(A))\|_j^2 \|z_i(A) - z_i(A)\|_k^2 + \frac{L}{6} \|z_i(A) - z_i(A)\|_k^3 \quad (41)$$

Overall, we have

$$\begin{aligned}
 (S; A) & \max_{i \in [K]} \frac{1}{2} \left(r_{w;A}^2(z_i(A)) \|z_i(A)\| + \frac{L}{6} \|z_i(A)\|^3 \right) \\
 & \max_{i \in [K]} \frac{1}{2} \left(kr_{w;A}^2(z_i(A)) \|z_i(A)\| + L \|z_i(A)\|^3 \right) \\
 & + \frac{L}{6} \|z_i(A)\|^3 \\
 & \max_{i \in [K]} \frac{1}{2} \left(kr_{w;A}^2(z_i(A)) \|z_i(A)\| + \frac{i(L)^2}{L^2} + \frac{i(L)^3}{6L^2} \right) \\
 & = \max_{i \in [K]} \frac{i(L)^2}{2L^2} kr_{w;A}^2(z_i(A)) \|z_i(A)\| + \frac{2i(L)^3}{3L^2}
 \end{aligned} \tag{42}$$

which completes the proof. \square

Lemma D.3 (Lemma 2.1 Bourin et al. (2013)) For every matrix in M_{n+m}^+ partitioned into blocks, we have a decomposition

$$\begin{pmatrix} A & X \\ X & B \end{pmatrix} = U \begin{pmatrix} A & 0 \\ 0 & 0 \end{pmatrix} U + V \begin{pmatrix} 0 & 0 \\ 0 & B \end{pmatrix} V$$

for some unitaries $U, V \in M_{n+m}$.

Lemma D.4. Then, given an arbitrary partitioned positive semi-definite matrix,

$$\begin{pmatrix} A & X \\ X & B \end{pmatrix} \preceq_s k_A k_s + k_B k_s$$

for all symmetric norms.

Proof. In lemma D.3 we have

$$\begin{pmatrix} A & X \\ X & B \end{pmatrix} = U \begin{pmatrix} A & 0 \\ 0 & 0 \end{pmatrix} U + V \begin{pmatrix} 0 & 0 \\ 0 & B \end{pmatrix} V$$

for some unitaries $U, V \in M_{n+m}$. The result then follows from the simple fact that symmetric norms are non-decreasing functions of the singular values where $k_s : M \rightarrow \mathbb{R}$, we have

$$f \left(\begin{pmatrix} A & X \\ X & B \end{pmatrix} \right) = f \left(U \begin{pmatrix} A & 0 \\ 0 & 0 \end{pmatrix} U + V \begin{pmatrix} 0 & 0 \\ 0 & B \end{pmatrix} V \right)$$

\square

Lemma D.5. For a $\text{Unif}(S^{d-1}(\frac{1}{\sqrt{d}}))$ and x are some vectors in \mathbb{R}^d with norm $\|x\| = \frac{1}{\sqrt{R(d)}}$, we have

$$P(\langle x; a \rangle^2 \leq k^2) = \min \left\{ 2 \arccos \frac{1}{\sqrt{R(d)}}; 1 - \frac{1}{\sqrt{R(d)}} \exp \left(-\frac{1}{4d} \right) \right\} \tag{43}$$

Proof. We can replace the unit vector a with e by

$$P(\langle x; a \rangle^2 \leq k^2) = P(\langle x; e \rangle^2 \leq 1) \tag{44}$$

Similarly, we can replace x by unit vectors such that

$$P(\langle x; e \rangle^2 \leq 1) = P \left(\langle h; e \rangle^2 \leq \frac{1}{R(d)} \right) \tag{45}$$

Solving $\langle h; e \rangle^2 = \frac{1}{R(d)}$, we get

$$\langle h; e \rangle^2 = \cos^2 \left(\arccos \frac{1}{\sqrt{R(d)}} \right) = \frac{1}{R(d)} \tag{46}$$

In this case, the probability will converge to 1 as $R(d)$ increases. As is known to us, surface area of S^{d-1} equals

$$A_d = r^{d-1} 2^{\frac{d-2}{2}} \quad (47)$$

An area C_d of the spherical cap equals

$$A_d^{\text{cap}}(r) = \int_0^Z A_{d-1}(r \sin \theta) r d\theta = \frac{2^{\frac{(d-1)-2}{2}}}{\frac{d-1}{2}} r^{d-1} \int_0^Z \sin^{d-2} \theta d\theta \quad (48)$$

where $\int_0^{\frac{\pi}{2}} \sin^{n-1} \theta d\theta = \frac{(2(n-1))!}{2^{2(n-1)}(n-1)!}$.

1) When $d = 1$, almost surely we have

$$P((x - e)^2 \leq \epsilon^2) = P(x^2 \leq 1) = 1 \quad (49)$$

2) When $d = 2$, we have S^2 where S^2 is a circle $= 1$ and the probability is the angle between s ; e how much the vectors span within the circle where

$$P(\|s - e\|^2 \leq \frac{1}{R(d)}) = \frac{2 \int_0^R r dr}{R(d)} = \frac{2}{R(d)} \quad (50)$$

3) When $d \geq 3$, the probability equals

$$P(\|s - e\|^2 \leq \frac{1}{R(d)}) = \frac{A_d^{\text{cap}}(r)}{\frac{1}{2}A_d} = 1 - \frac{A_d^{\text{cap}}(r)}{\frac{1}{2}A_d} \quad (51)$$

where $A_d^{\text{cap}}(r)$ is the remaining area of cutting the hyperspherical caps in half of the sphere,

$$\begin{aligned} A_d^{\text{cap}}(r) &= \frac{2^{\frac{(d-1)-2}{2}}}{\frac{d-1}{2}} \int_{\frac{\pi}{2}}^Z \sin^{d-2} \theta d\theta \\ &= \frac{2^{\frac{(d-1)-2}{2}}}{\frac{d-1}{2}} \int_{\frac{\pi}{2}}^Z \sin^{d-2} \theta d\theta \\ &= \frac{2^{\frac{(d-1)-2}{2}}}{\frac{d-1}{2}} (\cos \frac{Z}{2} + \cos \frac{\pi}{2}) \\ &= \frac{2^{\frac{(d-1)-2}{2}}}{\frac{d-1}{2}} \cos \frac{Z}{2} \end{aligned} \quad (52)$$

3-a) If d is even then $\frac{d}{2} = \frac{d}{2} - 1!$, so

$$\frac{2^{\frac{d-1}{2}}}{\frac{d-1}{2}} = \frac{(d-2)!}{2^{d-2} \frac{d-1}{2}} = \frac{2^{d-2}}{2^{d-2} \frac{d-1}{2}} \quad (53)$$

Robbins' bounds (Robbins, 1955) imply that for any positive integer

$$p^{\frac{4^d}{d}} \exp \frac{1}{8d-1} < \frac{2^d}{d} < p^{\frac{4^d}{d}} \exp \frac{1}{8d+1} \quad (54)$$

So we have

$$\begin{aligned} \frac{2A_d^{\text{cap}}(r)}{A_d} &= \frac{2^{\frac{d}{2}}}{\frac{d-1}{2}} = \frac{2^{d-1}}{\frac{d-1}{2}} \cos \frac{Z}{2} \\ &< \frac{2^{d-1}}{\frac{2^{d-2} \exp \frac{1}{4d-9}}}{(d-2)!} \cos \frac{Z}{2} \\ &< \frac{p^{\frac{4^d}{d}}}{\exp \frac{1}{4d-9}} \cos \frac{Z}{2} \end{aligned} \quad (55)$$

So the probability will be

$$P_{hs;ei^2} \frac{1}{R(d)} > 1 - \frac{p^{\frac{2d-4}{4}}}{p} \exp \frac{1}{4d-9} \cos : \quad (56)$$

Suppose $R(d) = kd; k > 1$, we have

$$\lim_{d \rightarrow \infty} \frac{p^{\frac{2d-4}{4}}}{kd \exp \frac{1}{4d-9}} = \frac{r}{k} : \quad (57)$$

3-b) Similarly, if d is odd, then $\frac{d-1}{2} = \frac{d-3}{2}!$, so

$$\frac{\frac{d-1}{2}}{\frac{d}{2}} = \frac{2^{\frac{d-2}{2}} \frac{d-3}{2}!}{(d-2)!} = \frac{2^{\frac{d-2}{2}}}{(d-2)^{\frac{d-2}{2}}} = \frac{d-3}{2} : \quad (58)$$

If $d = 3$ then

$$P_{hs;ei^2} \frac{1}{R(d)} = 1 - \frac{2^{\frac{d-4}{2}}}{2^{\frac{d-4}{2}}} \cos = 1 - \frac{1}{R(d)} : \quad (59)$$

If $d > 3$ then Robbins' bounds imply that

$$\frac{2^{\frac{d-4}{2}}}{2^{\frac{d-4}{2}}} = \frac{2^{\frac{d-4}{2}}}{d^{\frac{d-3}{2}}} > \frac{p^{\frac{2d-4}{4}}}{(d-3)^{\frac{d-3}{2}}} \exp \frac{1}{4d-11} : \quad (60)$$

Thus, the probability will be at least

$$P_{hs;ei^2} \frac{1}{R(d)} > 1 - \frac{p^{\frac{d-2}{2}}}{(d-3)} \exp \frac{1}{4d-11} \cos : \quad (61)$$

To simplify the result, we compare the minimum probability that $d \geq 3$

$$\begin{aligned} P_{hs;ei^2} \frac{1}{R(d)} &> 1 - \frac{p^{\frac{d-2}{2}}}{(d-3)} \exp \frac{1}{4d-11} \cos \\ &= 1 - \frac{p^{\frac{d-2}{2}}}{(d-3)} \exp \frac{1}{4d-11} \cos \\ &= 1 - \frac{p^{\frac{d-2}{2}}}{(d-3)R(d)} \exp \frac{1}{4d-11} \cos \\ &> 1 - \frac{p^{\frac{2d-4}{4}}}{R(d)} \exp \frac{1}{4d-9} \end{aligned} \quad (62)$$

Overall, we have $d \in \mathbb{N}_+$,

$$P_{hs;ei^2} \frac{1}{R(d)} > \min_{w} \frac{2 \arccos \frac{1}{R(d)}}{w}; 1 - \frac{p^{\frac{2d-4}{4}}}{R(d)} \exp \frac{1}{4d-9} : \quad (63)$$

□

D.1 PROOF TO THEOREM 3.4

Proof. Let $A^d = (a_i)_{i \in [d]} \stackrel{i.i.d.}{\sim} \text{Unif}(S^{d-1}(\frac{1}{d}))$. We consider the random ReLU NN function class to be

$$F_{\text{relu}}(A^d) = \{f(w; A; x) = \frac{1}{d} \sum_{i=1}^d \max(w_i, x - a_i) : w_i \in \mathbb{R}; i \in [m]\}$$

where $A = [a_1; \dots; a_m] \in \mathbb{R}^{d \times m}$. The empirical minimizer of the source domain is

$$w = \min_{w \in \mathbb{R}^d} \frac{1}{n} \sum_{x_i, y_i \in S} \max(w; A; x_i; y_i) = \frac{1}{n} \sum_{z_i \in S} \max(w; A(z_i)) : \quad (64)$$

Then with the chain rule, the first derivative of any input \mathbf{w} will be

$$\begin{aligned} \mathbf{r}_x \cdot (f(\mathbf{w}; A; x); y) &= \frac{\partial (f(\mathbf{w}; A; x); y)}{\partial \mathbf{w}} \frac{\partial (f(\mathbf{w}; A; x); y)}{\partial (Ax)} \frac{\partial (Ax)}{\partial \mathbf{w}} \\ &= \frac{D_f^1(x; y)}{d} \sum_{i=1}^n \mathbf{w}_i \mathbf{a}_i \mathbf{a}_i^\top \mathbf{x} \mathbf{0} \\ &= \frac{D_f^1(x; y)}{d} \mathbf{w} \mathbf{w}^\top \mathbf{0}(Ax) \end{aligned} \quad (65)$$

where a short notation $D_f^1(x; y)$ denotes the first order directional derivative of $f(\mathbf{w}; A; x)$ $\mathbf{0}(Ax)$ $\mathbb{R}^{m \times d}$ is the Jacobian matrix w.r.t. input. Apparently, the second order derivative is represented as $D_f^2(x; y)$, thus the Hessian will be

$$\begin{aligned} \mathbf{r}_x \cdot \cdot (f(\mathbf{w}; A; x); y) &= \frac{D_f^2(x; y)}{d} \sum_{i=1}^n \sum_{j=1}^n \mathbf{w}_i \mathbf{w}_j \mathbf{a}_i \mathbf{a}_j^\top \mathbf{x} \mathbf{0}; \mathbf{a}_i^\top \mathbf{x} \mathbf{0} \\ &= \frac{D_f^2(x; y)}{d} \mathbf{0}(Ax) \mathbf{w} \mathbf{w}^\top \mathbf{0}(Ax) \end{aligned} \quad (66)$$

Similarly, we have

$$\mathbf{r}_y \cdot (f(\mathbf{w}; A; x); y) = D_y^2(x; y) (\text{sgn}(y))^2 D_f^2(x; y) \quad (67)$$

where $\text{sgn}(y)$ is the sign function. holds under our choice of the family of loss functions.

1. Homogeneity in regression, i.e. L1, MSE, MAE, Huber Loss, we have $\mathbf{r}_y \cdot (f(\mathbf{w}; A; x); y) = |D_f^2(x; y)|$;
2. (Binary) Cross-Entropy Loss:

$$D_y^2(x; y) = \frac{\partial^2}{\partial y^2} \sum_{c=1}^C y^c \exp(x) = \frac{\partial^2}{\partial y^2} \sum_{c=1}^C \exp(x_c) = \frac{\partial^2}{\partial y^2} \mathbf{0} = 0$$

3. Negative Log Likelihood (NLL) loss $D_y^2(x; y) = 0$.

Besides, as a convex loss function $D_f^2(x; y) \geq 0$. Hence, the range of $D_f^2(x; y)$ will be $[0; D_f^2(x; y)]$.

To combine with robustness, we denote $\mathbf{h}(z; S; A) \in \mathbb{R}^{d+1}$. Therefore, the Hessian will be

$$H(z; S; A) := \frac{\partial^2}{\partial \mathbf{z} \partial \mathbf{z}^\top} (f(\mathbf{w}; A; x); y) = \frac{\partial^2 (f(\mathbf{w}; A; x); y)}{\partial \mathbf{z} \partial \mathbf{z}^\top} \quad (68)$$

With Lemma D.4, the spectral norm of Hessian will be bounded by

$$\|\mathbf{H}(z)\| \leq \mathbf{r}_x \cdot \cdot (f(\mathbf{w}; A; x); y) + \mathbf{r}_y \cdot (f(\mathbf{w}; A; x); y) \quad (69)$$

The first term in (69) can be further bounded by

$$\begin{aligned} D_f^2(x; y) \mathbf{0}(Ax) \mathbf{w} \mathbf{w}^\top \mathbf{0}(Ax) &\leq |D_f^2(x; y)| \mathbf{w} \mathbf{w}^\top \mathbf{0}(Ax) \mathbf{0}(Ax) \\ &= D_f^2(x; y) \|\mathbf{w}\|^2 \mathbf{0}(Ax) \mathbf{0}(Ax) \end{aligned} \quad (70)$$

where the convexity of loss function $\mathbf{0}(x; y; D_f^2(x; y) \geq 0$ supports the last equation. The right term has the facts that

$$\|\mathbf{0}(Ax) \mathbf{0}(Ax)^\top\| \leq \|\mathbf{0}(Ax)\| \|\mathbf{0}(Ax)^\top\| = \text{tr} \mathbf{0}(Ax) \mathbf{0}(Ax)^\top \quad (71)$$

In summary, we have the following inequality:

$$\begin{aligned}
 \mathbb{E} \left[\sum_{i=1}^n D_f^2(x_i; y_i) \right] &\leq \frac{1}{d} \mathbb{E} \left[\sum_{j=1}^n \|w_j\|^2 \text{tr} \left((Ax_j)(Ax_j)^\top \right) \right] + 1 \\
 &= \frac{1}{d} \mathbb{E} \left[\sum_{j=1}^n \|w_j\|^2 \text{tr} \left((Ax_j)(Ax_j)^\top \right) \right] + 1 \\
 &= \frac{1}{d} \mathbb{E} \left[\sum_{j=1}^n \|w_j\|^2 \sum_{i=1}^d (a_i^\top x_j)^2 \right] + 1
 \end{aligned} \tag{72}$$

In Lemma D.2, it depends on some $\epsilon = \epsilon(x_i; y_i) \in \mathcal{S} \setminus \mathcal{C}_i$ that

$$\begin{aligned}
 (\mathcal{S}; A) &= \max_{i \in [K]} \frac{i(L)^2}{2L^2} \mathbb{E} \left[\sum_{j=1}^n D_f^2(x_j; y_j) \right] + \frac{4}{3} \frac{i(L)}{L} \\
 &= \max_{i \in [K]} \frac{i(L)^2}{2L^2} \mathbb{E} \left[\sum_{j=1}^n D_f^2(x_j; y_j) \right] + \frac{4}{3} \frac{i(L)}{L} \\
 &= \frac{\max(L)^2}{2L^2} \mathbb{E} \left[\sum_{j=1}^n D_f^2(x_k; y_k) \right] + \frac{4}{3} \frac{\max(L)}{L} + o_p(1)
 \end{aligned} \tag{73}$$

where $\max(L) = \max_{i \in [K]} i(L)$, then $o_p(1) = O_p \left(\frac{1}{\sqrt{n}} \sum_{j=1}^n \|w_j\|^2 \text{tr} \left((Ax_j)(Ax_j)^\top \right) \right)$ is a smaller order term compared to first term, since $d \rightarrow \infty$. Last equality, the maximum can be taken as we find maximum $(x_k; y_k) \in \mathcal{C}_i \cap \mathcal{S}$. Because $(x_k; y_k) \in \mathcal{S}$ is one of the training samples $(x; y) \in [n]$, there must exist $i \in [0; n]$ that

$$\begin{aligned}
 D_f^2(x_k; y_k) &\leq \frac{1}{d} \sum_{j=1}^n \|w_j\|^2 \text{tr} \left((Ax_j)(Ax_j)^\top \right) \\
 &= \frac{1}{d} \sum_{j=1}^n \|w_j\|^2 \sum_{i=1}^d (a_i^\top x_j)^2
 \end{aligned} \tag{74}$$

Recall that the sharpness of parameters is defined by

$$\begin{aligned}
 (\hat{w}; \mathcal{S}; A) &:= \frac{1}{n} \sum_{j=1}^n \|w_j\|^2 \text{tr} \left[H_{\mathcal{S}; A}(\hat{w}) \right] \\
 &= \frac{1}{n} \sum_{j=1}^n \|w_j\|^2 \text{tr} \left((Ax_j)(Ax_j)^\top \right) \\
 &= \frac{1}{n} \sum_{j=1}^n \|w_j\|^2 \sum_{i=1}^d (a_i^\top x_j)^2
 \end{aligned} \tag{75}$$

Let the $a_i = a_i^\top x_j \sim D(\cdot)$ and the expectation $\mathbb{E} \left[\sum_{i=1}^d (a_i^\top x_j)^2 \right] = d \sigma^2$ where $D(\cdot)$ is some rotationally invariant distribution, i.e. uniform or normal distribution. Under this circumstance, the sample mean of $(a_i^\top x_j)^2$ still obeys the same family distribution $D(\cdot)$. Thus, we have

$$\begin{aligned}
 \mathbb{P} \left(\sum_{j=1}^n \|w_j\|^2 \sum_{i=1}^d (a_i^\top x_j)^2 \geq \frac{1}{d} \sum_{j=1}^n \|w_j\|^2 \sum_{i=1}^d (a_i^\top x_j)^2 \right) &= \mathbb{P} \left(\sum_{j=1}^n \|w_j\|^2 \sum_{i=1}^d (a_i^\top x_j)^2 \geq \frac{1}{d} \sum_{j=1}^n \|w_j\|^2 \sum_{i=1}^d (a_i^\top x_j)^2 \right) \\
 &= \mathbb{P} \left(\sum_{j=1}^n \|w_j\|^2 \sum_{i=1}^d (a_i^\top x_j)^2 \geq \frac{1}{d} \sum_{j=1}^n \|w_j\|^2 \sum_{i=1}^d (a_i^\top x_j)^2 \right) \\
 &= \mathbb{P} \left(\sum_{j=1}^n \|w_j\|^2 \sum_{i=1}^d (a_i^\top x_j)^2 \geq \frac{1}{d} \sum_{j=1}^n \|w_j\|^2 \sum_{i=1}^d (a_i^\top x_j)^2 \right)
 \end{aligned} \tag{76}$$

With Lemma D.5, we have at least a probability at

$$\mathbb{P} \left(\sum_{j=1}^n \|w_j\|^2 \sum_{i=1}^d (a_i^\top x_j)^2 \geq \frac{1}{d} \sum_{j=1}^n \|w_j\|^2 \sum_{i=1}^d (a_i^\top x_j)^2 \right) \geq \min \left(\frac{2}{\pi} \arccos R(d)^{\frac{1}{2}}; 1 - \frac{p}{R(d)} e^{-\frac{1}{4d}} \right) \tag{77}$$

the following inequality holds,

$$\begin{aligned} \mathbb{E}_{\mathcal{S}; A} \left[\frac{\max(L)^2}{2L^2} n^0(\mathbf{w}; \mathcal{S}; A) + \frac{4 \max(L)}{3} + \epsilon \right] \\ = \frac{\max(L)^2}{2L^2} n^0 + O\left(\frac{d}{m}\right) \mathbb{E}_{\mathcal{S}; A} \left[\frac{4}{3} \max(L) \right] \end{aligned} \quad (78)$$

□

D.2 PROOF TO COROLLARY 3.5

Corollary D.6 (Restatement of Corollary 3.5) Let w_{\min} be the minimum value of w_j . Suppose $x \sim \text{Unif}(S^{d-1}(\bar{d}))$ and $j \in [1, d]$. If $m > d$, $\max(L) < \frac{w_{\min}^2 M_1 \sim(d; m)}{2d}$ for any $A = (a_1; \dots; a_m)$; $a_i \sim \text{Unif}(S^{d-1}(\bar{d}))$ $\forall i \in [1, m]$, taking expectation over $x \sim \text{Unif}(S^{d-1}(\bar{d}))$ in S , we have

$$\mathbb{E}_{\mathcal{S}; A} \left[\frac{\max(L)^2}{2L^2} n^0(\mathbf{w}; \mathcal{S}; A) + M_2 \right] \geq \lambda_{\min}(G) \quad (79)$$

where $\lambda_{\min}(G) = \mathbb{E}_{a_i \sim \text{Unif}(S^{d-1}(\bar{d}))} \left(\prod_{i=1}^m a_i a_i^T G_{ii} \right) > 0$ is the minimum eigenvalue and G is product constant of Gegenbauer polynomials

$$G_{ij} = \int_{t=0}^1 \frac{d-1}{2} N_{d,t}^2 Q_t^{(d)}(a_i; a_j) dt$$

Proof. In our main theorem, with some probability, we have the following relation

$$\mathbb{E}_{\mathcal{S}; A} \left[\frac{\max(L)^2}{2L^2} n^0(\mathbf{w}; \mathcal{S}; A) + \frac{4 \max(L)}{3} + d \right] \geq \lambda_{\min}(G)$$

So, we are concerned about the relation between the $\mathbb{E}_{\mathcal{S}; A}$ second term. If $\max(L) < \frac{w_{\min}^2 M_1 \sim(d; m)}{2d}$, we may say the RHS is dominated by sharpness term $\frac{4 \max(L)}{3}$ as well as the main effect is taken by the sharpness. As suggested in Lemma D.1, we have

$$\mathbb{E}_{x \sim \text{Unif}(S^{d-1}(\bar{d}))} H(x) \geq \frac{w_{\min}^2 M_1 \sim(d; m)}{d} \mathbb{I}_d \quad (80)$$

where \mathbb{I}_d is the global minimum over the whole set. As defined in (38), the following condition holds true

$$9z_i(A) \geq 2M_i; \quad k z_i(A) \leq \frac{i(L)}{L}; \quad (81)$$

and with Hessian Lipschitz, the relation is almost surely for arbitrary A

$$\begin{aligned} \mathbb{E}_{z_i(A)} kH(z_i(A)) - H(z_i(A))k &\leq L k z_i(A) - z_i(A)k \leq i(L) \\ &\leq \frac{w_{\min}^2 M_1 \sim(d; m)}{2d} \mathbb{I}_d \\ \mathbb{E}_{A; x} \frac{1}{2} kH(x)k &\geq \frac{1}{2} \mathbb{I}_d \\ \mathbb{E}_{z_i(A)} \frac{1}{2} kH(z_i(A))k &\geq \frac{1}{2} \mathbb{I}_d \end{aligned} \quad (82)$$

Obviously, $kH(z_i(A))k > \frac{1}{2} \mathbb{I}_d$. Following Lemma A.2 of Zhang et al. (2019a), for $z_i(A) \in C_i$, we have a similar result that

$$\lambda_{\min}(H(z_i(A))) \geq \lambda_{\min}(H(z_i(A))) - kH(z_i(A)) - H(z_i(A))k \geq \frac{1}{2} \mathbb{I}_d \quad (83)$$

where λ_{\min} denotes the minimum singular value. With Lemma D.2, we know that

$$\mathbb{E}_{\mathcal{S}; A} \left[\frac{\max(L)^2}{2L^2} n^0(\mathbf{w}; \mathcal{S}; A) + \frac{4 \max(L)}{3} \right] \geq \mathbb{E}_{\mathcal{S}; A} \max_{i \in [K]} \frac{i(L)^2}{2L^2} kH(z_i(A))k + \frac{4 \max(L)}{3} \quad (84)$$

We also have another condition that

$$jD_f^2(x; y)j := \frac{\partial^2 (f(\hat{w}; A; x); y)}{\partial \hat{w}^2} \leq 2 [M_1; M_2]; \forall x; y: \quad (85)$$

Combine all these results, we finally have

$$E_{S;A} (S; A) \leq E_{S;A} \max_{i \in [K]} \frac{(L)^2}{2L^2} \left(1 + \frac{4}{3} \|H(z_i(A))\| \right) + \frac{7 \max(L)^2}{6L^2} \frac{\mathbb{E} \sum_{j=1}^n D_f^2(x_j; y_j)}{d} \frac{\mathbb{E} \|k\|^2}{n} \sum_{j=1}^n \mathbb{E} \|a_j\|^2 \leq M_2 A \quad (86)$$

Recall the definition of $(\hat{w}; S; A)$ in the main theorem that

$$E_{S;A} (\hat{w}; S; A) \leq \mathbb{E}_{f, x, g^n} \|k\|^2 \frac{1}{n} \sum_{j=1}^n \frac{D_f^2(x_j; y_j)}{d} \sum_{i=1}^n \mathbb{E} \|a_i\|^2 \leq 0 \quad (87)$$

Look at the second sum, we have

$$\begin{aligned} \mathbb{E}_{f, x, g^n; f, a, g^m} \sum_{i=1}^n \mathbb{E} \|a_i\|^2 \mathbb{E} \|x_j\|^2 &= \sum_{i=1}^n \mathbb{E}_{x_j} \mathbb{E}_{a_i} \|k\|^2 \mathbb{E} \cos^2(\angle) \mathbb{E} \|a_i\|^2 \leq 0 \\ &= \sum_{i=1}^n \mathbb{E}_{x_j} \mathbb{E}_{a_i} \|k\|^2 \mathbb{E} \|a_i\|^2 \mathbb{E} \cos^2(\angle) \end{aligned} \quad (88)$$

Suppose x and a are i.i.d. from $\text{Uni}(\mathbb{S}^{d-1}(1))$, let $u = [x; a]$, we have a well-known result that

$$\begin{aligned} \mathbb{E}_u [u^2] &= \mathbb{E} [\|x\|^2 + \|a\|^2] \\ &= \mathbb{E} [\|k\|^2 \cos^2(\angle)] \\ &= \mathbb{E} [\cos^2(\angle)] = \frac{1}{d}; \quad x; a \in \mathbb{R}^d \end{aligned} \quad (89)$$

Therefore, in (88),

$$\sum_{i=1}^n \mathbb{E}_{x_j} \mathbb{E}_{a_i} \|k\|^2 \mathbb{E} \|a_i\|^2 \mathbb{E} \cos^2(\angle) = \sum_{i=1}^n \mathbb{E}_{x_j} \mathbb{E}_{a_i} \|k\|^2 \mathbb{E} \|a_i\|^2 \leq 0 \quad (90)$$

and we have (based on proof of main theorem),

$$\begin{aligned} E_{S;A} (S; A) &\leq \frac{7 \max(L)^2}{6L^2} \frac{\mathbb{E} \|k\|^2}{n} \sum_{j=1}^n \frac{D_f^2(x_j; y_j)}{d} \sum_{i=1}^n \mathbb{E}_{x_j} \mathbb{E}_{a_i} \|k\|^2 \mathbb{E} \|a_i\|^2 \leq 0 + M_2 A \\ E_{S;A} \frac{7 \max(L)^2}{6L^2} \sum_{j=1}^n D_f^2(x_j; y_j) &\leq n^0 (\hat{w}; S; A) + M_2 \end{aligned} \quad (91)$$

□

E CASE STUDY

To better illustrate our theorems, we here give two different cases for clearly picturing intuition. The first case is the very basic model, ridge regression. As is known to us, ridge regression provides a straightforward way (by punishing the norm of the weights) to reduce the "variance" of the model in order to avoid overfitting. In this case, this mechanism is equivalent to reducing the model's sharpness.

Example E.1. In ridge regression models, the robustness constant has a reverse relationship to regularization parameter where $\lambda > 0$, the more probably after minimum λ and less sensitivity λ of the learned model could be. Follow the previous notation that $(\hat{w}; S; A)$ denotes the robustness and $(\hat{w}; S)$ is the sharpness on training set, then we have

$$\lambda > 0; c \in (0; n]; (\hat{w}; S; A) \leq c (\hat{w}; S) + \epsilon_d$$

where ϵ_d is a much smaller order than $(\hat{w}; S)$.

E.1 RIDGE REGRESSION

We consider a generic response model as stated in [Ali et al. \(2019\)](#).

$$y_j = (X^T; \mathbf{1})^T$$

ridge regression minimization problem is defined by

$$\min_{\beta} \frac{1}{2n} \|X\beta - y\|^2 + \frac{\lambda}{2} \|\beta\|^2; X \in \mathbb{R}^{n \times d}; n < d: \quad (92)$$

The least-square solution of ridge regression is

$$\hat{\beta} = (X^T X + n\lambda I)^{-1} X^T y: \quad (93)$$

With minimizer $\hat{\beta}$, we now focus on its geometry w.r.t. Let $S = \{(z_j; y_j)\}_{j=1}^n$ be the training set, $(Z; \mathbf{1})$ be a measure space. Consider the bounded sample such that

$$\|z_j\| \leq R; \mathbf{1}^T z_j \leq R: \quad (94)$$

The Z can be partitioned into K disjoint sets $C_i; i \in [K]$. By definition, we have robustness defined by each partition C_i ,

$$\|z_j\| \leq R; z_j \in C_i; \|\hat{\beta}(z_j) - \hat{\beta}(z_0)\| \leq \epsilon(S; A): \quad (95)$$

For this convex function $\ell(\hat{\beta}; z)$, we have the following upper bound in the whole sample domain

$$\begin{aligned} \ell(S) &= \max_{i \in [K]} \sup_{z \in C_i} \ell(\hat{\beta}; z) - \ell(\hat{\beta}; z_0) \\ &= \max_{i \in [K]} \sup_{z \in C_i} \ell(\hat{\beta}; z) - \ell(\hat{\beta}; z_i) + \ell(\hat{\beta}; z_i) - \ell(\hat{\beta}; z_0) \\ &= \sup_{z_j \in S} \ell(\hat{\beta}; z_j) - \ell(\hat{\beta}; z_0) \end{aligned} \quad (96)$$

where the $z_i; z_0$ are the global minimum point in C_i and whole domain $Z = \cup_{i=1}^K C_i$, respectively. z_j is a training data point that has the maximum loss difference from the optimum. Specifically, it as well as the augmented form $\hat{\beta}_+$ can be expressed as

$$z = [x_1; \dots; x_d; y]^T; \hat{\beta}_+ = [\hat{\beta}_1; \dots; \hat{\beta}_d; \mathbf{1}]^T; \in \mathbb{R}^{d+1}:$$

Then the loss difference can be rewritten as

$$\ell_{\hat{\beta}_+}(z) = (\hat{\beta}_+^T z)^2 = (\hat{\beta}_+^T x - y)^2; H(\hat{\beta}_+(z)) \text{ is P.S.D matrix}$$

It is a convex function with regards to z such that

$$\begin{aligned} \ell(S) &= \sup_{z_j \in S} \ell_{\hat{\beta}_+}(z_j) - \ell_{\hat{\beta}_+}(z_0) \\ &= \sup_{z_j \in S} \left[\frac{1}{2} \hat{\beta}_+^T (z_j - z_0) \right]^2 + \frac{1}{2} (z_j - z_0)^T H(\hat{\beta}_+(z_0)) (z_j - z_0) \\ &= \sup_{z_j \in S} \frac{1}{2} (z_j - z_0)^T H(\hat{\beta}_+(z_0)) (z_j - z_0) \\ &= \sup_{z_j \in S} \frac{1}{2} \|H(\hat{\beta}_+(z_0))\| \|z_j - z_0\|^2 \end{aligned} \quad (97)$$

where the second equality is supported by convexity and the third equality is due to $\hat{\beta}_+^T z_0 = 0$. Further, with Lemma D.4, we have

$$\|H(\hat{\beta}_+(z_0))\| \leq \frac{2}{\lambda} \frac{\|\hat{\beta}_+\|}{\|\hat{\beta}_+(z_0)\|} > \frac{2}{\lambda} \frac{\|\hat{\beta}_+(z_j)\|}{\|\hat{\beta}_+(z_0)\|} \frac{3}{5} \|\hat{\beta}_+\| + 1: \quad (98)$$

In (97), we can also bound the norm of the input difference by

$$\|z_j - z\|^2 \leq k^2 \|x_j - x\|^2 + (y_j - y)^2 \quad (99)$$

and then for simplicity, assume $k^2 \|x_j - x\|^2$ we have

$$\begin{aligned} (S) \quad & \sup_{x_j, z \in \mathcal{S}} \frac{1}{2} \|\hat{\beta}^{\lambda_j} - \beta\|^2 + 1 \cdot k^2 \|x_j - x\|^2 + (y_j - y)^2 \\ & \sup_{x_j, z \in \mathcal{S}} \frac{1}{2} \|\hat{\beta}^{\lambda_j}\|^2 + 1 \cdot k^2 \|x_j\|^2 + k^2 \|x\|^2 + (y_j - y)^2 \\ & \sup_{x_j, z \in \mathcal{S}} \frac{1}{2} \|\hat{\beta}^{\lambda_j}\|^2 \cdot k^2 \|x_j\|^2 + k^2 \|x\|^2 + O(d) \\ & \sup_{x_j, z \in \mathcal{S}} \|\hat{\beta}^{\lambda_j}\|^2 \cdot k^2 \|x_j\|^2 + O(d) \end{aligned} \quad (100)$$

where $\|\hat{\beta}^{\lambda_j}\|^2 \cdot k^2 \|x_j\|^2 = O(d^2)$ is the dominate term for large d . Now, let's look at the relation to sharpness. By definition,

$$(\hat{\beta}; S) = k^2 \text{tr} \left(H_{\lambda}(\hat{\beta}; S) \right) = k^2 \text{tr} \left(\frac{X^{\>X}}{n} + I \right) : \quad (101)$$

Since

$$\text{tr} \left(\frac{X^{\>X}}{n} + I \right) = \text{tr} \left(\frac{X^{\>X}}{n} \right) + \text{tr}(I) = \text{tr} \left(\frac{XX^{\>}}{n} \right) + \frac{1}{n} \sum_j X_j^{\>} k^2 + ; \quad (102)$$

so we have

$$(\hat{\beta}; S) = k^2 \frac{1}{n} \sum_j X_j^{\>} k^2 + k^2 : \quad (103)$$

As is known to us, the "variance" of ridge estimator [Ali et al. \(2019\)](#) can be defined by

$$\text{Var}(\hat{\beta}), \text{tr} \hat{\beta}^{\lambda_j} = \text{tr} \left(X^T X + n I \right)^{-1} X^T y y^T X \left(X^T X + n I \right)^{-1} : \quad (104)$$

Note that $\hat{\beta}^{\lambda_j}$ is a PSD with $\text{rank}(\hat{\beta}^{\lambda_j}) = 1$, thus it has only one eigenvalue $\lambda(\hat{\beta}^{\lambda_j}) = k^2 > 0$.

$$\text{Var}(\hat{\beta}) = \text{tr} \hat{\beta}^{\lambda_j} = k^2 = O(d^2) \quad (105)$$

By definition, the covariance matrix $\mathbb{E}[y y^T]$ is a diagonal matrix with entries of σ^2 . Averagely, we have

$$\begin{aligned} \text{tr} \hat{\beta}^{\lambda_j} &= \frac{1}{n} \text{tr} \left(X^{\>X} + n I \right)^{-1} X^{\>X} X^{\>X} \left(X^{\>X} + n I \right)^{-1} \\ &= \frac{1}{n} \text{tr} \left(\frac{X^{\>X}}{n} + I \right)^{-1} \frac{X^{\>X}}{n} \\ &= \frac{1}{n} \sum_{i=1}^{X^d} \frac{\lambda_i(X^{\>X}=n)}{(\lambda_i(X^{\>X}=n) + 1)^2} \end{aligned} \quad (106)$$

where $\frac{X^{\>X}}{n}$ and $\frac{X^T X}{n} + I$ are simultaneously diagonalizable and commutable. Therefore, the greater λ is, the smaller $\text{tr} \hat{\beta}^{\lambda_j}$ is.

Conclusions From our above analysis, we have the following conditions.

- Upper bound of robustness $(S) \sup_{x_j, z \in \mathcal{S}} k^2 \|x_j\|^2 + O(d)$:
- Sharpness expression $(\hat{\beta}; S) = k^2 \frac{1}{n} \sum_{x_j, y_j \in \mathcal{S}} k^2 + ;$

- Variance expression $\text{Var}(\hat{\beta}) = k^2 \sigma^2$:

1) First, let's discuss how inferences sharpness.

$$\begin{aligned} \text{Sharpness}(\hat{\beta}; S) &= k^2 \text{tr}^{-1}(\text{H}(\hat{\beta}; S)) \\ &= \frac{1}{\sum_{x_j, y_j \in S} \frac{1}{n} \sum_{i=1}^n (X^T X + nI)^{-1} X^T X} k^2 + \frac{1}{A} \quad (107) \\ &= O(n^{-1}): \end{aligned}$$

As dictated in above equation, the $\text{Sharpness}(\hat{\beta}; S) = O(n^{-1})$ where sharpness holds an inverse relationship to n .

2) Now, it's trivial to get the relationship between robustness and sharpness by combining the first two points. Because $x_j, y_j \in S$ in supremum is one of the training samples there exists a constant $c < n$ and $\epsilon_d = o(d^2)$ such that

$$\begin{aligned} \text{Robustness}(S) &\leq \frac{c}{n} \sum_j \frac{X^n}{k^2} k^2 k_{x_j}^2 + \epsilon_d \\ &\leq \frac{c}{n} \sum_j \frac{X^n}{k^2} k^2 k_{x_j}^2 + c k^2 + \epsilon_d \quad (108) \\ &= c \text{Sharpness}(\hat{\beta}; S) + \epsilon_d: \end{aligned}$$

This relation is consistent with Theorem 3.4 where the robustness is upper bounded by sharpness $\text{Sharpness}(\hat{\beta}; S)$. Besides, the relation is simpler here, where robustness only depends on sharpness without other coefficients before $\text{Sharpness}(\hat{\beta}; S)$.

3) Finally, the relation between and robustness will be

$$\begin{aligned} \text{Robustness}(S) &\leq \frac{c}{n} \sum_j \frac{X^n}{k^2} \text{Var}(\sum_j k_{x_j}^2) + \epsilon_d = \frac{c}{n} \sum_{i=1}^n \frac{X^d}{(\sum_{i=1}^n (X^T X + nI)^{-1} X^T X)^2} \sum_j k_{x_j}^2 + \epsilon_d \\ &= \frac{2}{n} \sum_{i=1}^n \frac{X^d}{(\sum_{i=1}^n (X^T X + nI)^{-1} X^T X)^2} \sum_j k_{x_j}^2 + \epsilon_d \quad (109) \\ &= O(n^{-2}): \end{aligned}$$

where the $\text{Robustness}(S)$ somehow is the order $O(n^{-2})$ in the limit of n . It's clear to us that the greater n is, the less sensitive (more robust) model we can get. In practical, we show that "over-robust" may hurt the model's performance (we show the detail empirically in Figure 7).

E.2 2-LAYER DIAGONAL LINEAR NETWORK CLASSIFICATION

However, our main theorem assumes the loss function satisfying Polyak- ojasiewicz (P) condition. To extend our result to a more general case, here we study a 2-layer diagonal linear network classification problem whose loss is exponential-based and not satisfied the P condition.

Example E.2. We consider a classification problem using a 2-layer diagonal linear network with exp-loss. The robustness $\text{Robustness}(S; A)$ has a similar relationship in Theorem 3.4. Given training set S after iteration $t > T$, $\epsilon_{C_2} > 0$, $\text{Robustness}(S; A) \leq C_2 \sup_{t > T} \text{Sharpness}(t; S)$.

Given a training set $S = (X; y)$; $X = [x_1; \dots; x_n]$; $X \in \mathbb{R}^{d \times n}$. A depth-2 diagonal linear networks with parameters $u = [u_+; u_-] \in \mathbb{R}^{2d}$ specified by:

$$f(u; x) = \frac{1}{2} (u_+^2 - u_-^2) x_i$$

We consider exponential loss where $\ell(x) = \frac{1}{n} \sum_{i=1}^n \exp(-x_i^T y_i)$ and $y_i \in \{-1, 1\}$. It has the same tail behavior and thus similar asymptotic properties as the logistic or cross-entropy loss. WLOG,

we assume $\|y_i\|_2 = 1$ such that $x_i = y_i x_i$. Suggest by [Moroshko et al. \(2020\)](#), we have

$$k(t) = 2^{-2} \sinh^{-1} \left(\frac{X^T r(t)}{k(t)} \right)$$

where $r(t) = \frac{1}{n} \exp(-X^T r(t))$ and $k(t)k_1 = L(t)$. Note that

$$\frac{d}{dt} k(t) = \frac{4^p}{n} \frac{1}{2(t) + 4^{-41}} X^T \exp(-X^T r(t)) = A(t)X^T r(t) \quad (110)$$

where $A(t) = \text{diag} \left\{ \frac{4^p}{2(t) + 4^{-41}} \right\}$.

Part i, sharpness. First derivative and Hessian can be obtained by

$$\begin{aligned} \nabla L(t) &= (X^T r(t))^T \\ H(L(t)) &= \frac{\partial^2 L(t)}{\partial (t)^2} = \sum_{i=1}^n r_i(t) X_i X_i^T \end{aligned} \quad (111)$$

With the definition of sharpness, we can get

$$\begin{aligned} \kappa(t; S) &= k(t) k^2 \text{tr}(H[L(t)]) \\ &= k(t) k^2 \text{tr} \sum_{i=1}^n r_i(t) X_i X_i^T \\ &= k(t) k^2 \text{tr} \sum_{i=1}^n r_i(t) X_i^T X_i \\ &= k(t) k^2 \sum_{i=1}^n r_i(t) \|x_i\|_2^2 \end{aligned} \quad (112)$$

Part ii, robustness. Now, let's use the same discussion of the robustness constant in the previous case. Follow the previous definition of $\kappa(S; A)$, after some iteration number \bar{t} it is defined by

$$\kappa(S; A) := \sup_{i \in [2, n], t \in T} \sum_{j=1}^n r_j(t) \|r(t)\|_2 \quad (113)$$

where $r_j(t)$ is the (denormalized) point-wise loss of x_j and $r(t)$ denotes the minimum loss of point x . There exists $n^0 < n$, we have

$$\kappa(S; A) = \sup_{t \in T} n^0 (k(t)k_1 \|r(t)\|_2) \sup_{t \in T} n^0 k(t)k_1 \quad (114)$$

Let $k_{\min} k = \min_{i \in [2, n]} \|x_i\|_2$, the above equation

$$\begin{aligned} \kappa(S; A) &= \frac{1}{k_{\min} k^2} \sup_{t \in T} (n^0 k(t)k_1 \|r(t)\|_2) \\ &= \frac{n^0}{k_{\min} k^2} \sup_{t \in T} \sum_{i=1}^n r_i(t) \|x_{\min}\|_2^2 \\ &= \frac{n^0}{k_{\min} k^2} \sup_{t \in T} \sum_{i=1}^n r_i(t) \|x_i\|_2^2 \end{aligned} \quad (115)$$

Part iii, connection. Compare the last part of (112) and (115), we found that robustness and sharpness depend on the same term. Further, we can say that for any $k(t)k^2$ will have the upper bound.

From Lemma 11 of [Moroshko et al. \(2020\)](#), we have the following inequality,

$$k(t)k_1 \leq 2^{-2} \sinh^{-1} \left(\frac{X^T r(t)}{k(t)} \right) \quad (116)$$

where $L(t) = \exp(-X^T r(t))$. Then, we can bound the $k(t)k^2$ via:

$$C_1 k(t)k^2 \leq dk(t)k_1^2 = 4d^{-4} \sinh^{-2} \left(\frac{X^T r(t)}{k(t)} \right) < 1 \quad (117)$$

Note that $C_1 > 0$. In summary, we have

$$\mathcal{L}(S; A) \leq \frac{n^0}{C_1 k_{\min} k} \sup_{t \in T} \mathcal{L}(t; S) \leq C_2 \sup_{t \in T} \mathcal{L}(t; S) \quad (118)$$

where $C_2 = \frac{n^0}{C_1 k_{\min} k} > 0$ is a constant.

Part iv, sanity check— asymptotic. Asymptotically, as $t \rightarrow \infty$, $L(t) \rightarrow 0$ while $k(t)k_2$ will be explode. So if sharpness $\mathcal{L}(t; X) \rightarrow \infty$, then it will fail to imply robustness.

$$\begin{aligned} \mathcal{L}(t; S) &= k(t)k^2 \sum_{i=1}^n r_i(t)kx_i k^2 \\ &= k(t)k^2 \sum_{i=1}^n r_i(t)kx_{\max} k^2 \\ &= k(t)k^2 L(t)kx_{\max} k^2 \\ &= \max_{t \in T} \mathcal{L}(t; S) \end{aligned} \quad (119)$$

Let $\max_{t \in T} \mathcal{L}(t; X)$ be an upper bound of $\mathcal{L}(t; X)$ at any time step t . The dynamics will be

$$\begin{aligned} \frac{d \max_{t \in T} \mathcal{L}(t; S)}{dt} &= r_{\max} \mathcal{L}(t; S) - (t) \\ &= kx_{\max} k^2 \text{tr}(XX^T) - k(t)k^2 L(t) + 2L(t) \mathcal{L}(t) - (t) \\ &= kx_{\max} k^2 \text{tr}(XX^T) - k(t)k^2 (Xr(t))^T A(t)Xr(t) + 2L(t) \mathcal{L}(t) - A(t)Xr(t) \\ &= kx_{\max} k^2 \text{tr}(XX^T) - 2L(t) \mathcal{L}(t) - k(t)k^2 (Xr(t))^T A(t)Xr(t) \end{aligned} \quad (120)$$

As $t \rightarrow \infty$, we have $L(t) \rightarrow 0$, thus it is easy to converge that

$$\lim_{t \rightarrow \infty} \frac{d \max_{t \in T} \mathcal{L}(t; S)}{dt} = kx_{\max} k^2 k(t)k^2 (Xr(t))^T A(t)Xr(t) - kx_{\max} k^2 k(t)k^2 L(t) = 0 \quad (121)$$

As we can see, the dynamics of the derivative of $\max_{t \in T} \mathcal{L}(t; X)$ is decreasing to 0 as $t \rightarrow \infty$ which means the sharpness $\mathcal{L}(t; X)$ will be upper bounded by a converged number $\mathcal{L}(t \rightarrow \infty; X) < C_1 < 1$. So sharpness will not explode.

F COMPARISON TO FEATURE ROBUSTNESS

F.1 THEIR RESULTS

Definition F.1 (Feature robustness Petzka et al. [Petzka et al. \(2020\)](#)): $Y \subseteq \mathbb{R}_+$ denote a loss function, and two positive (small) real numbers ϵ, δ , X, Y a finite sample set, and $A \in \mathbb{R}^{m \times m}$ a matrix. A model $f(x) = (w; A)(x)$ with $(X) \in \mathbb{R}^m$ is called $((\epsilon; S; A); \delta)$ -feature robust, if $E_{\mathcal{F}}(f; S; A) \leq \delta$ for all $0 < \epsilon < \delta$. More generally, for a probability distribution \mathcal{P} on perturbation matrices \mathbb{R}^m , we define

$$E_{\mathcal{F}}(f; S; A) = E_{A \sim \mathcal{P}} E_{\mathcal{F}}(f; S; A);$$

and call the model $((\epsilon; S; A); \delta)$ -feature robust on average over \mathcal{P} if $E_{\mathcal{F}}(f; S; A) \leq \delta$ for $0 < \epsilon < \delta$.

Theorem F.2 (Theorem 5 Petzka et al. [Petzka et al. \(2020\)](#)): Consider a model $f(x; w) = g(w(x))$ as above, a loss function g , and a sample set S , and let $O_m \subseteq \mathbb{R}^{m \times m}$ denote the set of orthogonal matrices. Let ϵ be a positive (small) real number and $w \in \mathbb{R}^d$ denote parameters at a local minimum of the empirical risk on a sample set S . If the labels satisfy that $(A(x_i)) = y(x_i) = y_i$ for all $(x_i; y_i) \in S$ and all $kA k \leq 1$, then $f(x; w)$ is $((\epsilon; S; O_m); \delta)$ -feature robust on average over O_m for $\delta = \frac{\epsilon^2}{2m} (\epsilon) + O(\epsilon^3)$.

Proof sketch

1. With assumption that $\mathbb{E}[f_A(x_i)] = y_i$ for all $(x_i; y_i) \in \mathcal{S}$ and all $k \leq K$, feature perturbation around w is

$$\mathbb{E}_F(f; \mathcal{S}; A) + \mathbb{E}_{\text{emp}}(w; \mathcal{S}) = \mathbb{E}_{\text{emp}}(w + wA; \mathcal{S})$$

2. Since Taylor expansion for local minimum $w = !$ will only remains second order term, thus

$$\mathbb{E}_{\text{emp}}(w + wA; \mathcal{S}) = \mathbb{E}_{\text{emp}}(!; \mathcal{S}) + \frac{1}{2} \sum_{s,t=1}^d \mathbb{E}_{A \sim \mathcal{O}_m} (!_s A) H_{s,t} (!; \mathcal{S}) (!_t A)^T + O(\epsilon^3)$$

3. With basic algebra, one can easily get

$$\begin{aligned} \mathbb{E}_{A \sim \mathcal{O}_m} \mathbb{E}_F(f; \mathcal{S}; A) &= \frac{1}{2} \sum_{s,t=1}^d \mathbb{E}_{A \sim \mathcal{O}_m} (!_s A) H_{s,t} (!_t A)^T + O(\epsilon^3) \\ &= \frac{1}{2m} \sum_{s,t=1}^d \mathbb{E}_{! \sim \mathcal{S}} \text{Tr}(H_{s,t}) + O(\epsilon^3) \\ &= \frac{1}{2m} \text{Tr}(H) + O(\epsilon^3) \end{aligned}$$

F.2 COMPARISON TO OUR RESULTS

We first summarize the commonalities and differences between their results and ours:

- Both of us consider the robustness of the model. But they define the feature robustness while we study the loss robustness [Xu & Mannor \(2012\)](#) which has been studied for many years.
- They consider a non-standard generalization gap by decomposing it into representativeness and the expected deviation of the loss around the sample points. But we strive to integrate sharpness into the generalization guarantees.

For point 1, their defined feature robustness trivially depends on the sharpness. Because sharpness (the curvature information) is just defined by the robust perturbation areas around the desired point. From step 2 in the above proof sketch we can see, the hessian is exactly the second expansion of perturbed expected risk. So we think this definition provides less information about the optimization landscape. In contrast, we consider the loss robustness for two reasons: 1) it is easy to get in practice without finding the orthogonal matrices \mathcal{O}_m . 2) we highlight its dependence on the data manifold.

For point 2, we try to integrate this optimization property (sharpness) into the standard generalization frameworks in order to get a clearer interplay. Unlike feature robustness, the robustness defined by loss function will be easier analyzed in generalization tools, because it's hard and vague to define the "feature" in general. Besides, our result will also benefit the data-dependent bounds [Xu & Mannor \(2012\)](#); [Kawaguchi et al. \(2022\)](#).

G EXPERIMENTS AND DETAILS

G.1 RIDGE REGRESSION WITH DISTRIBUTIONAL SHIFTING

As we stated before, we followed the [Duchi & Namkoong \(2021\)](#) to investigate the ridge regression on distributional shift. We randomly generate $2 \times R^d$ in spherical space, and data from

$$X \stackrel{\text{iid}}{\sim} \mathcal{N}(0; 1); y = X \cdot \theta_0 \quad (122)$$

To simulate distributional shift, we randomly generate a perpendicular unit vector θ_0 . Let $\{e_1, \dots, e_d\}$ be the basis vectors, then shifted ground-truth will be computed from the basis $\theta_0 = \cos(\alpha) e_1 + \sin(\alpha) e_2$. For the source domain, we use θ_0 as our training distribution. We randomly

sample 50 data points and train a linear classifier with a gradient descent of 3000 iterations. Starting from $\theta = 0$, we gradually increase θ to generate different distributional shifts.

From the left panel in Figure 7 we can see that a larger penalty suffers from lower loss increasing when θ ranges from 0 to 2. Since we consider a cycling shift of label space $\theta = 2\pi$ corresponds to the maximum shift thus leading to the highest loss increase. According to our analysis of ridge regression, larger λ means a flatter minimum and more robustness, resulting in a better OOD generalization. This experiment verifies our theoretical analysis. However, it is important to note that too large a coefficient of λ_2 regularization will hurt the performance. As shown in Figure 7 (right panel), the curse of overfitting (indicated by brown colors) appears when $\lambda = 4$.

Figure 7: OOD test losses are increasing along distributional shifting. The X-axis is the shifting angle θ and the Y-axis is the test loss of the model which is trained on distribution $\theta = 0$. Left: The larger regularization λ causes a lower increase in test loss (Darker purple lines have lower test losses). Right: Too large λ_2 penalty coefficient brings poor fitting and thus fails to generalize on both ID and OOD datasets (orange lines). Best viewed in colors.

G.2 ADDITIONAL RESULTS ON SHARPNESS

Here we show more experimental results on RotatedMNIST which is a rotation of MNIST handwritten digit dataset with different angles ranging from 0° ; 15° ; 30° ; 45° ; 60° ; 75° (6 domains/distributions). For each environment, we select a domain as a test domain and train the model on all other domains. Then OOD generalization test accuracy will be reported on the test domain. In our experiments, we run each algorithm with 12 different seeds, thus getting 12 trained models of different minima. Then we compute the sharpness (see Algorithm 1) for all these models and then plot them in Figure 8. For algorithms, we choose Empirical Risk Minimization (ERM), and Stochastic Weight Averaging (SWA) as the plain OOD generalization algorithm which is shown in the first column. In robust optimization algorithms, we choose Group Distributional Robust Optimization (Sagawa et al. (2019)). We also choose CORAL (Sun & Saenko (2016)) as a multi-source domain generalization algorithm. Among these different types of out-of-domain generalization algorithms, we can conclude that sharpness will affect the test accuracy on the OOD dataset.

Experimental configurations are listed in Table 1. For each model, we run 5000 iterations and choose the last model as the test model. To ease the computational burden, we choose the MLP to compute the sharpness.

Algorithms	Optimizer	lr	WD	batch size	MLP size	eta	MMD
ERM(SWA)	Adam	0.001	0	64	265*3	-	-
DRO	Adam	0.001	0	64	265*3	0.01	-
CORAL	Adam	0.001	0	64	265*3	-	1

Table 1: Hyperparameters we use for different DG algorithms in the experiments.

From Figure 8 we can see, the sharpness has an inverse relationship with the out-of-domain generalization performance for every model in each individual environment. To make it clear, we plot similar

tasks from environment 1 to 4 as the last row. Thus, we can see a clearer tendency in all algorithms. It verified our Theorem 3.1. Note that all algorithms have different feature scales. One may need to normalize the results of different algorithms when plotting them together.

Algorithm 1 Pseudocode of model sharpness computation

Require: feature layer $f(x)$, training loss ℓ
Ensure: Sharpness S
 Get Jacobian matrix w.r.t. feature layer $\mathbf{J} = \nabla \ell(f(x), y)$
for each gradient vector \mathbf{J}_i in \mathbf{J}^\top **do**
 Compute Hessian w.r.t. element i, j of $f(x)$ by $\frac{\partial \mathbf{J}_i}{\partial f_i(x)}$
end for
 We store Hessian in the variable \mathbf{H}
 Initialize sharpness $S = 0$
for i in feature layer.shape[0] **do**
 for j in feature layer.shape[0] **do**
 Retrieve the hessian value of i, j element via $h = \mathbf{H}[:, j, i, :]$
 Sharpness $s_{ij} = \text{Trace}(h) / f_{ij}(x)^2$
 $S = S + s_{ij}$
 end for
end for

G.3 COMPARE OUR ROBUST BOUND TO OTHER OOD GENERALIZATION BOUNDS

G.3.1 COMPUTATION OF OUR BOUNDS

First, we follow Kawaguchi et al. (2022) to compute the K in an inverse image of the ϵ -covering in a randomly projected space. The main idea is to partition input space in a projected space with transformation matrix A . The specific steps will be (1) To generate a random matrix A , we i.i.d. sample each entry from the Uniform Distribution $U(0, 1)$. (2) Each row of the random matrix $A \in \mathbb{R}^{3 \times d}$ is then normalized so that $Ax \in [0, 1]^3$, i.e. $A_{ij} = A_{ij} / \sqrt{\sum_{j=1}^d A_{ij}^2}$ (3) After generating a random matrix A , we use the ϵ -covering of the space of $u = Ax$ to define the pre-partition $\{C_i\}_{i=1}^K$.

G.3.2 COMPUTATION OF PAC-BAYES BOUND

We follow the definition to compute expected dis in Germain et al. (2013) where

Definition G.1. Let H be a hypothesis class. For any marginal distributions D_S and D_T over X , any distribution ρ on H , the domain disagreement $\text{dis}(D_S, D_T)$ between D_S and D_T is defined by,

$$\text{dis}(D_S, D_T) \stackrel{\text{def}}{=} \mathbf{E}_{h, h' \sim \rho} [R_{D_T}(h, h') - R_{D_S}(h, h')].$$

Since the $\text{dis}(D_S, D_T)$ is defined as the expected distance, we can compute its empirical version according to their theoretical upper bound as follows.

Proposition G.2 (Germain et al. Germain et al. (2013) Theorem 3). *For any distributions D_S and D_T over X , any set of hypothesis H , any prior distribution π over H , any $\delta \in (0, 1]$, and any real number $\alpha > 0$, with a probability at least $1 - \delta$ over the choice of $S \sim T \sim (D_S \times D_T)^m$, for every ρ on H , we have*

$$\text{dis}(D_S, D_T) \leq \frac{\mathbb{E}_{S, T} [\text{dis}(S, T) + \frac{2\text{KL}(\rho \| \pi) \ln 2}{m\alpha}] + 1}{1 - e^{-2}}$$

. With $\text{dis}(S, T)$, we can then compute the final generalization bound by the following inequality

$$\begin{aligned} & \delta \rho \text{ on } H, R_{P_T}(G) - R_{P_T}(G) - R_{P_S}(G) \\ & + \text{dis}(D_S, D_T) + R_{D_T}(G, G) + R_{D_S}(G, G) \end{aligned}$$

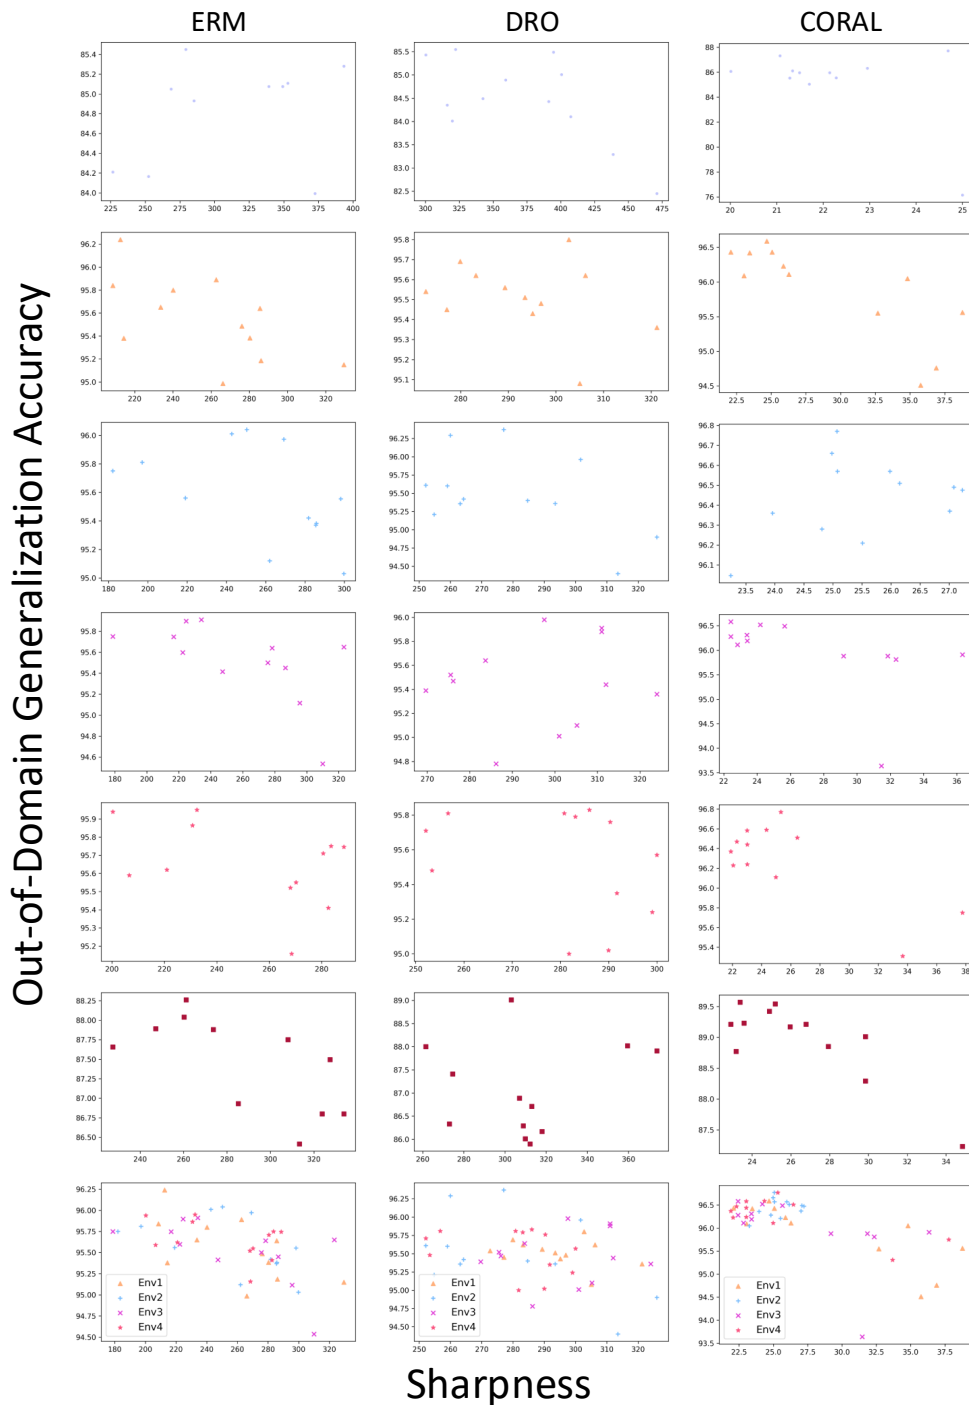


Figure 8: Domain generalization test accuracy on RotatedMNIST. From top to the bottom: environment 0, 1, 2, 3, 4, 5 with angles: $[0^\circ, 15^\circ, 30^\circ, 45^\circ, 60^\circ, 75^\circ]$ and a plot together. Each column shows the 12 runs of each algorithm.

with $\rho_T^* = \operatorname{argmin}_G R_{P_T}(G)$ is the best target posterior, and $R_D(G, G_T) = E_{h \sim P} E_{h' \sim P_T} R_D(h, h')$.

Note that we ignore the expected errors over the best hypothesis by assuming the $R_D(G, G_T) = 0$. We apply the same operation in E^* of Proposition C.6 as well.

G.3.3 COMPARISONS

In this section, we add some additional experiments on comparing to other baselines, i.e. PAC-Bayes bounds [Germain et al. \(2013\)](#). As shown in the first row of Figure 9, our robust framework has a smaller distribution distance in the bound compared to the two baselines when increasing the model size. In the second row, we have similar results in final generalization bounds. From the third and fourth rows we can see, our bound is tighter than baselines when suffering distributional shifts.

