
Normalized Flat Minima: Exploring Scale Invariant Definition of Flat Minima for Neural Networks Using PAC-Bayesian Analysis

Yusuke Tsuzuku^{1,2*} Issei Sato^{1,2} Masashi Sugiyama^{2,1}

Abstract

The notion of flat minima has gained attention as a key metric of the generalization ability of deep learning models. However, current definitions of flatness are known to be sensitive to parameter rescaling. While some previous studies have proposed to rescale flatness metrics using parameter scales to avoid the scale dependence, the normalized metrics lose the direct theoretical connections between flat minima and generalization. We first provide generalization error bounds using existing normalized flatness measures for smooth and stochastic networks using second-order approximation. Using the analysis, we then propose a novel normalized flatness metric. The proposed metric enjoys both direct theoretical connections and better empirical correlation to generalization error.

1. Introduction

Deep learning methods have achieved significant performance improvement in many domains, such as computer vision, language processing, and speech processing (Krizhevsky et al., 2012; Devlin et al., 2019; van den Oord et al., 2018). However, we are still on the way to understand when they perform well on unseen data. Better understanding of the generalization performance of deep learning methods would help improve their performance. Deep learning community has made tremendous effort to understand generalization of neural networks, both theoretically and empirically (Zhang et al., 2017; Neyshabur et al., 2017; Arora et al., 2018).

The notion of “flat minima” has gained attention as a possi-

*Author is currently working at Preferred Networks, Inc.
¹The University of Tokyo, Tokyo, Japan ²RIKEN AIP, Tokyo, Japan. Correspondence to: Yusuke Tsuzuku <tsuzuku@ms.k.u-tokyo.ac.jp>.

ble explanation for the generalization ability of deep learning methods (Hochreiter & Schmidhuber, 1997). Many empirical studies have supported the usefulness of this notion. For example, the notion shed light on the reason for larger generalization gaps in large-batch training (Keskar et al., 2017; Yao et al., 2018), and also inspired various training methods (Hochreiter & Schmidhuber, 1997; Chaudhari et al., 2017; Hoffer et al., 2017). Notably, Jiang et al. (2020) suggested that PAC-Bayes based generalization metrics, which have deep connections with flatness, are effective and promising among many other generalization metrics. As measures of flatness, previous studies proposed the volume of the region in which a network keeps roughly the same loss value (Hochreiter & Schmidhuber, 1997), the maximum loss value around minima (Keskar et al., 2017), and the spectral norm of the Hessian (Yao et al., 2018).

Despite the empirical connections of “flatness” to generalization, current definitions of flatness suffer from unnecessary scale dependence. Dinh et al. (2017) showed that we can arbitrarily change the flatness of the loss-landscape for some networks without changing the functions represented by the networks. Such scale dependence appears in networks with ReLU activation functions or normalization layers such as batch-normalization (Ioffe & Szegedy, 2015) and weight-normalization (Salimans & Kingma, 2016). A major counterargument to the scale dependence is that flatness provides valid generalization error bounds when parameter scales are taken into account (Dziugaite & Roy, 2017; Neyshabur et al., 2017). However, both the flatness and parameter scales have weak correlations with the generalization error in practical settings (Sec. 8).

Previous studies have tried to make flatness metrics invariant with respect to parameter scaling. Li et al. (2018) has successfully visualized connections between empirical performances and flatness hypotheses by scaling the loss-landscape using parameter scales. The success suggests that the normalized loss-landscape better captures networks’ generalization and motivates us to provide their theoretical interpretations. Neyshabur et al. (2017), Achille & Soatto (2018), and Liang et al. (2019) also suggested scaling flatness metrics by parameter scales. While these metrics are

invariant with respect to known scaling issues, they did not provide direct connections between the metrics and generalization error. From the PAC-Bayesian perspective, providing scale-invariant bounds is not straightforward because we need to use the log-uniform distribution as a prior so that the Kullback-Leibler divergence term in PAC-Bayes bounds becomes scale-invariant (Kingma et al., 2015), while the use of the prior invalidates the PAC-Bayesian analysis.

On the way to understand generalization of deep learning and its connections to a normalized loss-landscape, this paper makes two contributions.

- We provide PAC-Bayesian generalization error bounds using scale-invariant loss curvature metrics.
- We propose a novel normalized loss curvature metric that has a close connection to the bounds.

This paper is organized as follows. We start our discussion with the PAC-Bayesian generalization error bound and its connection to flat minima (Sec. 3). Next, we connect a flatness metric normalized by parameter scales per parameter to the PAC-Bayesian analysis (Sec. 4). Unfortunately, per-parameter normalization yields a constant term proportional to the number of parameters in the generalization error bounds, which makes them vacuous for deep learning methods. Thus, we re-analyze the known scale dependence issue (Sec. 5). We identify that we do not need to scale the loss-landscape per-parameter, but we only need to scale it per node. Then, applying the analysis, we improve the method of loss-landscape normalization (Sec. 6). Using the novel normalized flatness definition, we provide generalization error bounds that do not have constant terms that scale with the number of parameters. Empirically, the proposed metric can predict the generalization performance of models more accurately than current unnormalized metrics (Sec. 8).

Remark: We provide the generalization bounds for stochastic networks. Most of them rely on the quadratic approximation of the loss function, but not all of them¹.

2. Related work

The notion of flat minima has its origin in the minimum description length principle (MDL) (Hinton & van Camp, 1993; Rissanen, 1986; Honkela & Valpola, 2004). Dziugaite & Roy (2017) connected flat minima to PAC-Bayesian arguments, which is a generalization of the MDL. They pointed out that the sharpness of local minima is not sufficient for measuring generalization, and that parameter scales need to be taken into account. PAC-Bayesian analysis has also

¹Sec. 6 briefly explains that the stochastic network becomes close to a deterministic network as sample size increases.

been used for generalization analysis outside the context of flat minima (Neyshabur et al., 2018; Letarte et al., 2019). Neyshabur et al. (2018) analyzed the propagation of parameters’ perturbations to bound the generalization error. Given the existence of adversarial examples (Szegedy et al., 2014), however, this approach inevitably provides vacuous bounds. Instead, in this paper, we stick to flat-minima arguments, which will better capture the effect of parameter perturbations.

Since Dinh et al. (2017) pointed out the scale dependences of flat minima, many studies have attempted to resolve the issue. A common approach is scaling sharpness metrics by parameter scales. Neyshabur et al. (2017) pointed out that the scale dependence can be removed by balancing parameter scales and sharpness metrics. However, their proposal requires using data-dependent priors (Guedj, 2019; Parrado-Hernández et al., 2012; Dziugaite & Roy, 2018) or equivalent alternatives, which adds non-trivial costs to the generalization bounds. In this paper, we make the costs explicit in Sec. 3 and later reduce them in Sec. 6. Prior to Neyshabur et al. (2017), Dziugaite & Roy (2017) tried to achieve balance using numerical optimization. While they partially mitigated the scale dependence, they could not completely overcome scale dependence because the prior variance was tied into whole network. Wang et al. (2018) explored a better choice of posterior variance. However, not only their analysis could not remove scale dependence, it was based on a parameter-wise argument, which involved a factor that scales with the number of parameters, making their bound as vacant as naive parameter counting². Our analysis overcomes both problems.

Sharpness metrics normalization has also appeared outside the PAC-Bayesian context. Li et al. (2018) proposed rescaling the loss-landscape by filters’ scales of convolutional layers. Even though they provided useful visualizations and insights, it is unclear how the scaling relates to generalization. Our work bridges the gap between the empirical success and theoretical understandings. Achille & Soatto (2018) proposed a notion of information in weights, which has connections to flat minima and PAC-Bayes. However, their arguments were based on an improper log-uniform prior and did not provide valid generalization error bounds. Even if we avoid the log-uniform prior to use a method described in Neklyudov et al. (2017), the notion does not provide non-vacuous bounds, as we show in Sec. 4. Liang et al. (2019) proposed the Fisher-Rao norm, which is a metric invariant with respect to known scaling issues. While the Fisher-Rao norm has interesting functional equivalence, in addition to invariance, it has been hard to directly connect the metric to generalization.

²More comparison with Wang et al. (2018) is available in Appendix L in the supplementary material.

We propose a novel scale-invariant sharpness metric for loss-landscapes. Our definition distinguishes itself by its direct connection to PAC-Bayesian generalization error bounds. While previous methods for normalization have been based on per-parameter normalization, ours is based on per-node normalization. This difference is critical from the viewpoint of our PAC-Bayesian analysis in terms of the tightness of the bounds.

3. Flat minima from PAC-Bayesian perspective

This section reviews a fundamental PAC-Bayesian generalization error bound and its connection to flat minima provided by prior work. Table 1 in Appendix A in the supplementary material summarizes the notations used in this paper.

3.1. PAC-Bayesian generalization error bound

The following is one of the most basic PAC-Bayesian generalization error bounds (Germain et al., 2016; Alquier et al., 2016). Let \mathcal{D} be a distribution over input-output space \mathcal{Z} , S be a set of m samples drawn *i.i.d.* from \mathcal{D} , \mathcal{H} be a set of parameters θ , $\ell : X \times \mathcal{H} \rightarrow [0, 1]$ be a loss function, \mathbb{P} be a distribution over \mathcal{H} independent of S , and \mathbb{Q} be a distribution over \mathcal{H} . For any $\delta \in (0, 1]$, any distribution \mathbb{Q} , and any nonnegative real number λ , with probability at least $1 - \delta$,

$$\begin{aligned} \mathcal{L}_{\mathcal{D}}(\mathbb{Q}(\theta | \bar{\theta})) &\leq \mathcal{L}_S(\mathbb{Q}(\theta | \bar{\theta})) + \frac{\lambda}{2m} + \frac{1}{\lambda} \ln \frac{1}{\delta} \\ &\quad + \frac{1}{\lambda} \text{KL}[\mathbb{Q}(\theta | \bar{\theta}) \| \mathbb{P}(\theta)], \end{aligned} \quad (1)$$

where

$$\mathcal{L}_{\mathcal{D}}(\mathbb{Q}(\theta | \bar{\theta})) := \mathbb{E}_{z \sim \mathcal{D}, \theta \sim \mathbb{Q}(\theta | \bar{\theta})} [\ell(z, \theta)], \quad (2)$$

$$\mathcal{L}_S(\mathbb{Q}(\theta | \bar{\theta})) := \mathbb{E}_{z \sim S, \theta \sim \mathbb{Q}(\theta | \bar{\theta})} [\ell(z, \theta)]. \quad (3)$$

We reorganize the PAC-Bayesian bound (1) for later use as

$$\begin{aligned} \mathcal{L}_{\mathcal{D}}(\mathbb{Q}(\theta | \bar{\theta})) &\leq \mathcal{L}_S(\theta) + \frac{\lambda}{2m} + \frac{1}{\lambda} \ln \frac{1}{\delta} \\ &\quad + \underbrace{\mathcal{L}_S(\mathbb{Q}(\theta | \bar{\theta})) - \mathcal{L}_S(\theta)}_{(A)} + \underbrace{\frac{1}{\lambda} \text{KL}[\mathbb{Q}(\theta | \bar{\theta}) \| \mathbb{P}(\theta)]}_{(B)}. \end{aligned} \quad (4)$$

Similar decompositions can be found in prior work (Dziugaite & Roy, 2017; Neyshabur et al., 2017; 2018). We use a different PAC-Bayes bound (1) for later analysis, but they are essentially the same.³ Flat minima, which

³To apply our analysis, we can also use other PAC-Bayesian

are the noise-stable solution with respect to parameters, naturally correspond to term (A) in Eq. (4). Following prior work (Langford & Caruana, 2002; Hochreiter & Schmidhuber, 1997; Dziugaite & Roy, 2017; Arora et al., 2018), we analyze the true error of the stochastic classifier \mathbb{Q} in the following sections. Nagarajan & Kolter (2019) presented a method to generalize the PAC-Bayes bounds to deterministic classifiers. We will leave combining our work and theirs as future work.

3.2. Effect of noise under second-order approximation

To connect PAC-Bayesian analysis with the Hessian of the loss-landscape as in prior work (Keskar et al., 2017; Dinh et al., 2017; Yao et al., 2018), we consider the second-order approximation of some surrogate loss functions. We use a Gaussian with a covariance matrix $\sigma^2 I$ as the posterior of parameters⁴. Then term (A) in the PAC-Bayesian bound (4) can be calculated as

$$\begin{aligned} \mathcal{L}_S(\mathbb{Q}(\theta | \bar{\theta})) - \mathcal{L}_S(\bar{\theta}) &= \mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \sigma I)} [\mathcal{L}_S(\bar{\theta} + \epsilon)] - \mathcal{L}_S(\bar{\theta}) \end{aligned} \quad (5)$$

$$\approx \frac{1}{2} \text{Tr}(\nabla_{\bar{\theta}}^2 \mathcal{L}_S(\bar{\theta}) \sigma^2). \quad (6)$$

Thus, we can approximate term (A) by the trace of the Hessian and use it as a generalization metric. The effect of the approximation is further discussed in Appendix H in the supplementary material. Note, connecting PAC-Bayes with Hessian appears in the literature repeatedly (Dziugaite & Roy, 2017; Wang et al., 2018). By tuning σ with an appropriate prior, we can balance terms (A) and (B) (Dziugaite & Roy, 2017; Neyshabur et al., 2017). However, appropriate methods to balance them have not been extensively studied. Moreover, while some prior work proposed scaling σ by parameter scales, the operation has not been justified from the PAC-Bayesian viewpoint. We address the issues in the next section.

4. PAC-Bayes and parameter-wise normalization

In this section, we propose a framework that connects parameter-wisely normalized flatness metrics to PAC-Bayesian generalization error analysis. As a first step, we virtually decompose the parameters:

$$\theta = \eta \mu \quad e^{[\sigma]}, \quad (7)$$

bounds such as Theorem 1.2.6 in Catoni (2007), which is known to be relatively tight in some cases and successfully provided empirically nontrivial bounds for ImageNet scale networks in Zhou et al. (2019).

⁴Remark: Eq. (1) holds with not only Gaussian posteriors but also arbitrary posteriors.

where \odot is the Hadamard product, and $e^{[\cdot]}$ is an element-wise exponential. The codomain of random variables θ, μ, σ is \mathbb{R}^n , and a hyperparameter η is a positive real number $\eta \in \mathbb{R}_{>0}$. Intuitively, this is a scale-sign decomposition of parameters with continuous relaxation. We extend the decomposition to a scale-direction decomposition in Sec. 6. We design a specific prior and posterior on μ and σ to mitigate the scale dependence. We first define our prior design.

$$\mathbb{P}(\mu) = \mathcal{N}(\mu \mid \mathbf{0}, \mathbf{I}), \quad (8)$$

$$\mathbb{P}(\sigma) = \mathcal{N}(\sigma \mid \beta, \eta' \mathbf{I}), \quad (9)$$

where $\beta \in \mathbb{R}^n$ and $\eta' \in \mathbb{R}_{>0}$ are hyperparameters. Next, we define our posterior design.

$$\mathbb{Q}(\mu \mid \bar{\mu}) = \mathcal{N}(\mu \mid \bar{\mu}, \mathbf{I}), \quad (10)$$

$$\mathbb{Q}(\sigma \mid \bar{\sigma}) = \mathcal{N}(\sigma \mid \bar{\sigma}, \eta' \mathbf{I}), \quad (11)$$

where $\bar{\mu}$ and $\bar{\sigma}$ are real vectors in \mathbb{R}^n satisfying

$$\bar{\theta} = \eta \bar{\mu} \odot e^{[\bar{\sigma}]}, \quad (12)$$

where a real vector $\bar{\theta}$ is a parameter assignment. Note that the choice of $\bar{\mu}$ and $\bar{\sigma}$ is not necessarily unique, and we can balance them arbitrarily. Furthermore, the parameter η was introduced to address a technical issue concerning PAC-Bayesian analysis.

Under the second-order approximation, terms (A) and (B) in the PAC-Bayesian bound (1) can now be approximated as follows⁵.

$$\begin{aligned} \text{(A)}: & \frac{\eta^2}{2} \sum_{i=1}^n (\nabla_{\theta}^2 \mathcal{L}_S \bar{\theta})_{i,i} e^{2\bar{\sigma}_i} \\ & + \frac{(\eta')^2}{2} \sum_{i=1}^n (\nabla_{\theta}^2 \mathcal{L}_S \bar{\theta})_{i,i} (\bar{\theta}_i)^2 + \mathcal{E}(\mathcal{L}_S, \mathbb{Q}), \end{aligned} \quad (13)$$

$$\text{(B)}: \frac{1}{(2\lambda)} \left(\frac{1}{\eta^2} \bar{\theta} \odot e^{[-2\bar{\sigma}]} \odot \bar{\theta} + \frac{1}{(\eta')^2} \|\bar{\sigma} - \beta\|_2^2 \right), \quad (14)$$

where $\mathcal{E}(\mathcal{L}_S, \mathbb{Q})$ is an error term introduced by the second-order approximation. We obtain the following bound by combining (13) and (14), optimizing σ , and taking the union bound with respect to η, η' , and β .

Proposition 4.1. *For any networks, any $\lambda \in \Lambda$ and $\lambda' \in \Lambda'$ where $\Lambda \subset \mathbb{R}_{>0}$ and $\Lambda' \subset \mathbb{R}_{>0}$ are a finite set of real numbers independent of S , any 0–1 bounded twice continuously differentiable loss function with respect to parameters, any finite set of real numbers $\mathcal{B} \subset \mathbb{R}$ independent of S , and any parameter assignment $\bar{\theta}$, there exists a stochastic classifier*

⁵See Appendix D.1 in the supplementary material for the derivation.

$\mathbb{Q}(\theta \mid \bar{\theta})$ such that with probability $1 - \delta$ over the training set and \mathbb{Q} , the expected error $\mathcal{L}_{\mathcal{D}}(\mathbb{Q})$ is bounded by

$$\begin{aligned} \mathcal{L}_S(\bar{\theta}) & + \frac{\lambda'}{2\sqrt{\lambda}} \text{FR} + \frac{1}{2\sqrt{\lambda}} \text{NS}_0(\lambda') \\ & + \frac{1}{\lambda} \Xi + \frac{\lambda}{2m} + \mathcal{E}(\mathcal{L}_S, \mathbb{Q}), \end{aligned} \quad (15)$$

where

$$\text{FR} = \sum_{i=1}^n (\nabla_{\theta}^2 \mathcal{L}_S \bar{\theta})_{i,i} \bar{\theta}_i^2, \quad (16)$$

$\text{NS}_0 =$

$$\sum_{i=1}^n \min_{\substack{\bar{\sigma}_i \in \mathbb{R} \\ \beta_i \in \mathcal{B}}} \left[(\nabla_{\theta}^2 \mathcal{L}_S \bar{\theta})_{i,i} e^{2\bar{\sigma}_i} + \frac{\bar{\theta}_i^2}{e^{2\bar{\sigma}_i}} + \frac{(\bar{\sigma}_i - \beta_i)^2}{\lambda'} \right], \quad (17)$$

and

$$\Xi = \ln \frac{1}{\delta} + n \ln |\mathcal{B}| + 2 \ln |\Lambda| + \ln |\Lambda'|. \quad (18)$$

The optimal choice of λ is $O(m^{-\frac{1}{2}})$.

The FR and NS in the theorem stands for Fisher-Rao and normalized sharpness respectively, where the latter will be elaborated later and defined in Sec. 7. The convergence rate of the whole bound is $O(m^{-1/4})$ with the optimal choice of λ . This is not worse than existing bounds such as those of Bartlett et al. (2017), Neyshabur et al. (2018), and Arora et al. (2018). Appendix I in the supplementary material further discusses the convergence rate. The second term, containing FR, coincides with the second-order approximation of the scaled expected sharpness suggested by Neyshabur et al. (2017), the information in weight (Achille & Soatto, 2018), Gauss-Newton norm (Zhang et al., 2019), and Fisher-Rao norm under appropriate regularity conditions (Liang et al., 2019). The third term, containing NS_0 , also has a connection to normalized loss-landscapes given by

$$\text{NS}_1 := \frac{1}{2} \lim_{\lambda' \rightarrow \infty} \text{NS}_0(\lambda') = \sum_{i=1}^n \sqrt{(\nabla_{\theta}^2 \mathcal{L}_S \bar{\theta})_{i,i} (\bar{\theta}_i)^2}. \quad (19)$$

While this equation only holds at the limit, we can match the sides of the equation by modifying the prior and posterior design. We also can remove the second term at the same time. First, we replace \mathcal{B} to be a uniform distribution on a finite set. The set is carefully designed so that it does not become too large but sufficiently dense so that the scale dependence can be ignored. Next, we replace both of the prior and the posterior of σ with a distribution which takes 1 on a real vector and 0 anywhere else. The replacement

corresponds to setting η' very small in their original prior and posterior. A more concrete explanation is provided in the proof of the next proposition, provided in Appendix D.3 in the supplementary material.

Proposition 4.2. *For any networks, any $\lambda \in \Lambda$ where $\Lambda \subset \mathbb{R}_{>0}$ is a finite set of real number independent of S , any positive number $\epsilon \in \mathbb{R}_{>0}$, any 0–1 bounded twice continuously differentiable loss function with respect to parameters, and any parameter assignment θ such that $\max_i \theta_i \leq b$, if the diagonal elements of the Hessian of the training loss \mathcal{L}_S is bounded by $M \in \mathbb{R}_{>0}$ and nonnegative, there exists a stochastic classifier $\mathbb{Q}(\theta | \theta)$ such that with probability $1 - \delta$ over the training set and \mathbb{Q} , the expected error $\mathcal{L}_D(\mathbb{Q})$ is bounded by⁶*

$$\mathcal{L}_S(\theta) + \frac{1 + \epsilon + \epsilon^2}{\sqrt{\lambda}} \text{NS}_1 + \frac{1}{\lambda} \Xi_1 + \frac{\lambda}{2m} + \mathcal{E}(\mathcal{L}_S, \mathbb{Q}), \quad (20)$$

where

$$\Xi_1 = O(n) \cdot C_{\epsilon, \delta, |\Lambda|, M, b, \lambda}. \quad (21)$$

The assumptions that we can bound all elements of $\bar{\theta}$ and the diagonal of the Hessian are reasonable because we typically train neural networks on computers with finite precision. Even though ReLU networks do not satisfy the smoothness assumption, which is necessary because we use the Hessian, the proposition at least addresses the scale dependence due to normalization layers. Section 6 discusses a possible method to address the smoothness and nonnegativity assumptions for the Hessian’s diagonal elements. The method is also applicable to the above propositions. The ϵ that appears in the proposition comes from an approximation error introduced by the discretization trick. Setting $\epsilon = 1$ suggests that the parameter’s existence does not worsen the order of the bound.

5. Finer-grained scale dependence

The generalization error bounds (15) and (20) have constant terms proportional to the number of parameters. This is because specifying scaling parameters per parameter is almost as difficult as directly specifying all parameters. Since deep learning methods typically use overparametrized models, the constant terms make the bounds vacuous. Prior parameter-wisely normalized sharpness metrics also suffer from the problem explicitly or inexplicitly (Wang et al., 2018; Achille & Soatto, 2018). Fortunately, it turns out that we can control parameter scales by using fairly small number of scaling parameters. For example, we can guess that parameters belonging to the same weight matrix have similar

⁶The exact form of Ξ_1 can be found in the Appendix D.3 in the supplementary material.

scales. In this section, we discuss what controls the parameter scales in deep learning models. This discussion is critical for our improved PAC-Bayesian analysis in Sec. 6. It also reveals the scale dependence in existing matrix-norm-based generalization bounds⁷.

Here, we analyze the scale dependences appearing in networks with the ReLU activation function. The similar dependences appear in networks with normalization layers, which we defer the explanation to Appendix C.1 in the supplementary material. We point out that the matrix-wise scale dependence introduced by Dinh et al. (2017) does not fully cover scale dependences⁸. To illustrate the hidden scale dependence, we consider a simple network with a single hidden layer and ReLU activation:

$$f_{\theta}(z) = W^{(2)}(\text{ReLU}(W^{(1)}(z))), \quad (22)$$

where weight matrices $W^{(1)}$ and $W^{(2)}$ are subsets of the parameters θ , and z is an input of the network f_{θ} . We can scale the i -th column of $W^{(2)}$ by $\alpha > 0$ and i -th row of $W^{(1)}$ by $1/\alpha$ without modifying the function that the network represents.⁹ Since we are using the ReLU activation function, which has positive homogeneity, the transformation does not change the represented function. By the transformation, the diagonal elements of the Hessian corresponding to the i -th row of $W^{(1)}$ are scaled by α^2 . The transformation has essentially the same effect as the one proposed by Dinh et al. (2017). The difference is that the above transformation runs node-wise instead of matrix-wise. In terms of weight matrices, the node-wise scale dependence can be translated into a row- and column-wise scale dependence.

6. Improved PAC-Bayesian analysis

In this section, we tighten our PAC-Bayesian bounds in Sec. 4 based on Sec. 5. In Sec. 5, we pointed out the row- and column-wise scale dependencies in modern network architectures. Thus, we at least need to absorb the row and column scales of weight matrices. It seems, however, that we do not need to control all parameter scales separately. Thus, we modify the decomposition per weight matrix in Sec. 4 in the following way.

$$W^{(l)} = \eta \text{Diag}\left(e^{[\gamma^{(l)}]}\right) V^{(l)} \text{Diag}\left(e^{[\gamma'^{(l)}]}\right), \eta \in \mathbb{R}, \quad (23)$$

where $\text{Diag}\left(e^{[\gamma^{(l)}]}\right)$ and $\text{Diag}\left(e^{[\gamma'^{(l)}]}\right)$ are diagonal matrices whose diagonal elements are $e^{[\gamma^{(l)}]}$ and $e^{[\gamma'^{(l)}]}$, re-

⁷See Appendix C.2 in the supplementary material for the explanation.

⁸The same scale dependences in ReLU networks were introduced by Neyshabur et al. (2015).

⁹Running examples of the transformation can be found in Appendix B.1 in the supplementary material.

spectively. The codomains of random variables $\gamma^{(l)}$, $\gamma'^{(l)}$, and $V^{(l)}$ are $\mathbb{R}^{h_1^{(l)}}$, $\mathbb{R}^{h_2^{(l)}}$, and $\mathbb{R}^{h_1^{(l)} \times h_2^{(l)}}$, respectively. Similar decomposition with (23) has been explored for the unit-invariant SVD (Uhlmann, 2018). In the method, weight matrices are normalized using the solution of *Program II* (Rothblum & Zenios, 1992). Even though the same method is available for removing the scale dependence, we consider the Hessian jointly and use different scaling control that has some optimality from the PAC-Bayesian perspective. We first define some notations for convenience. Let $\bar{U}^{(l)}$ be a matrix defined as $\bar{U}^{(l)} = \bar{W}^{(l)} \bar{W}^{(l)}$, where $\bar{W}^{(l)}$ is a subset of a parameter assignment $\bar{\theta}$. Let $\bar{H}^{(l)}$ be a matrix such that

$$\bar{H}_{i,j}^{(l)} = \left(\frac{\partial^2 \mathcal{L}_S(\theta)}{\partial W^{(l)}_{i,j} \partial W^{(l)}_{i,j}} \Big|_{\theta = \bar{\theta}} \right). \quad (24)$$

By extending Prop. 4.2 to decomposition (23), we have the following proposition.

Proposition 6.1. *For any networks with d weight matrices, any $\lambda \in \Lambda$ where $\Lambda \subset \mathbb{R}_{>0}$ is a finite set of real number independent of S , any positive number $\epsilon \in \mathbb{R}_{>0}$, any positive number $R \in \mathbb{R}_{>0}$, any 0–1 bounded twice continuously differentiable loss function with respect to parameters, and a parameter assignment $\bar{\theta}$ such that $\max_i \bar{\theta}_i \leq b$, if the diagonal elements of the Hessian of the training loss \mathcal{L}_S are bounded by $M \in \mathbb{R}_{>0}$ and nonnegative, then there exists a stochastic classifier $\mathbb{Q}(\theta | \bar{\theta})$ such that with probability $1 - \delta$ over the training set and \mathbb{Q} , the expected error $\mathcal{L}_{\mathcal{D}}(\mathbb{Q})$ is bounded by¹⁰*

$$\mathcal{L}_S(\bar{\theta}) + \frac{(1 + \epsilon)^2}{\sqrt{\lambda}} \text{NS}_2(R) + \frac{1}{\lambda} \Xi_2 + \frac{\lambda}{2m} + \mathcal{E}(\mathcal{L}_S, \mathbb{Q}), \quad (25)$$

where $\text{NS}_2(R) =$

$$\sum_{l=1}^d \min_{\bar{\gamma}^{(l)} \in [-R, R]^{h_1^{(l)}}} \sqrt{e^{[-\bar{\gamma}^{(l)}]^\top \bar{H}^{(l)} (\bar{U}^{(l)})^\top e^{[\bar{\gamma}'^{(l)}]}}, \quad (26)$$

$$\Xi_2 = \tilde{O}(hd) \cdot C_{\epsilon, \delta, M, b, R, |\Lambda|}, \quad (27)$$

and $h = \frac{1}{d} \max \left(\sum_{l=1}^d h_1^{(l)}, \sum_{l=1}^d h_2^{(l)} \right)$.

The constant term, Ξ_2 , scales by the number of nodes, instead of parameters. The reduction of the constant term from Prop. 4.2 makes the bound meaningful in practical settings. For example, ResNet50 has tens of millions of parameters, while it only has tens of thousands of nodes. In classification on ImageNet, which has millions of images, this reduction is critical.

¹⁰The exact form of Ξ_2 can be found in the Appendix D.4 in the supplementary material.

Note that the proof of Prop. 6.1 is constructive. In the construction, the posterior $\mathbb{Q}(\theta | \bar{\theta})$ is a Gaussian centered at $\bar{\theta}$ with a diagonal covariance matrix Σ such that $\|\Sigma\|_F = \mathcal{O}(\lambda^{-\frac{1}{2}})$. Since the $\mathcal{E}(\mathcal{L}_S, \mathbb{Q})$ term comes from a second-order approximation of the loss function, as λ increases, which means a training set size m increases, the second-order approximation will hold better, and the $\mathcal{E}(\mathcal{L}_S, \mathbb{Q})$ term will decrease.

For the sake of theoretical completeness, we address the $\mathcal{E}(\mathcal{L}_S, \mathbb{Q})$ term, smoothness assumption, and nonnegative assumption. In Props. 4.2 and 6.1, the $\mathcal{E}(\mathcal{L}_S, \mathbb{Q})$ term was introduced as a term satisfying the following equation.

$$\begin{aligned} \mathcal{L}_S(\mathbb{Q}(\theta | \bar{\theta})) &= \mathcal{L}_S(\bar{\theta}) + \mathcal{E}(\mathcal{L}_S, \mathbb{Q}) \\ &+ \frac{1}{2} \mathbb{E}_{\theta \sim \mathbb{Q}(\theta | \bar{\theta})} \left[(\theta - \bar{\theta})^\top (\nabla_{\theta}^2 \mathcal{L}_S \bar{\theta}) (\theta - \bar{\theta}) \right]. \end{aligned} \quad (28)$$

As long as the equation is satisfied, we can use any vectors which mimic the Hessian’s diagonal elements. For example, we can use the following as an estimation of the Hessian diagonal.

$$\mathbf{f} = \mathbb{E}_{\mathbf{s} \in \{-1, 1\}^n} \left[\text{Div} \left[\mathbf{s} \left(g_{S, \bar{\theta}, \epsilon}(r, \mathbf{s}) - g_{S, \bar{\theta}, \epsilon}(r, -\mathbf{s}) \right), 2r(\text{Abs}[\theta] + \epsilon \mathbf{1}) \right] \right], \quad (29)$$

where

$$g_{S, \bar{\theta}, \epsilon}(r, \mathbf{s}) = \nabla_{\theta} \mathcal{L}_S(\bar{\theta} + r(\text{Abs}(\bar{\theta})\mathbf{s} + \epsilon \mathbf{1})), \quad (30)$$

$\text{Div}[\cdot, \cdot]$ is an elementwise division, and $\text{Abs}[\cdot, \cdot]$ is an elementwise absolute. We added a parameter ϵ for numerical stability. When a network is smooth, \mathbf{f} is an approximation of the diagonal elements of the Hessian. Note, the estimation is scale-invariant when $\epsilon = 0$. The estimation is defined even when networks are not smooth. Furthermore, we can carry out the following modification to \mathbf{f} to enforce the nonnegative condition.

$$\hat{\mathbf{f}} = \text{Max}[\mathbf{f}, \mathbf{0}], \quad (31)$$

where $\text{Max}[\cdot, \cdot]$ is an elementwise maximum. Let $\bar{F}^{(l)}$ be a matrix such that we substitute the diagonal elements of the Hessian by $\hat{\mathbf{f}}$ in the definition of $\bar{H}^{(l)}$ (24). Using $\bar{F}^{(l)}$ instead of $\bar{H}^{(l)}$, we have the following theorem.

Theorem 6.1. *For any networks with d weight matrices, for any $\lambda \in \Lambda$ where $\Lambda \subset \mathbb{R}_{>0}$ is a finite set of real number independent of S , a positive number $\epsilon \in (0, 0.1]$, any positive number $R \in \mathbb{R}_{>0}$, any 0–1 bounded loss function, and a parameter assignment $\bar{\theta}$ such that $\max_i \bar{\theta}_i \leq b$, if all elements of $\hat{\mathbf{f}}$ are bounded by $M \in \mathbb{R}_{>0}$, then there exists a Gaussian distribution $\mathbb{Q}(\theta | \bar{\theta})$ with a mean $\bar{\theta}$ and a covariance matrix Σ such that $\|\Sigma\|_F = \mathcal{O}(\lambda^{-1/2})$ in terms of*

λ , and with probability $1 - \delta$ over the training set \mathcal{S} and \mathbb{Q} , the expected error $\mathcal{L}_{\mathcal{D}}(\mathbb{Q})$ is bounded by¹¹

$$\mathcal{L}_{\mathcal{S}}(\mathbb{Q}(\boldsymbol{\theta} | \bar{\boldsymbol{\theta}})) + \frac{1}{\sqrt{\lambda}} \text{NS}_3(R) + \frac{1}{\lambda} \Xi_3 + \frac{\lambda}{2m}. \quad (32)$$

where $\text{NS}_2(R) =$

$$\sum_{l=1}^d \min_{\bar{\gamma}^{(l)} \in [-R, R]^{h_1^{(l)}}} \sqrt{e^{[-\bar{\gamma}^{(l)}]^\top} \bar{F}^{(l)} (\bar{U}^{(l)})^\top e^{[\bar{\gamma}'^{(l)}]}} \quad (33)$$

and

$$\Xi_3 = \tilde{O}(hd) \cdot C_{\epsilon, \delta, M, b, R, |\Lambda|}. \quad (34)$$

Note that we can statistically bound the term $\mathcal{L}_{\mathcal{S}}(\mathbb{Q}(\boldsymbol{\theta} | \bar{\boldsymbol{\theta}}))$ using the Monte-Carlo algorithm and Hoeffding’s inequality. In Theorem 6.1, we use the second-order approximation of the loss to estimate the best choice of posterior variance to minimize both terms (A) and (B) in Eq. (4).

7. Normalized sharpness definition and calculation

In this section, based on the analysis in Sec. 6, we define normalized sharpness, a sharpness and generalization metric invariant with respect to node-wise rescaling. We also describe a practical calculation technique for the metric.

We define normalized sharpness as

$$\text{NS} = \sum_{l=1}^d \inf_{\bar{\gamma}^{(l)} \in \mathbb{R}^{h_1^{(l)}}} \sqrt{e^{[\bar{\gamma}^{(l)}]^\top} \bar{F}^{(l)} (\bar{U}^{(l)})^\top e^{[-\bar{\gamma}^{(l)}]}}, \quad (35)$$

where the notations $\bar{U}^{(l)}$ and $\bar{F}^{(l)}$ were introduced in Sec. 6. Normalized sharpness appeared in Prop. 6.1 and Theorem 6.1. The reason for the scale-invariance of normalized sharpness is explained in Appendix J in the supplementary material.

In convolutional layers, since the same filter has the same scale, we can tie the scaling parameter per channel. We defer a more detailed description of normalized sharpness for convolutional layers to Appendix F in the supplementary material. Fortunately, the optimization problem (35) is convex with respect to $\bar{\gamma}^{(l)}$.¹² Thus, we can estimate a near-optimal solution to the optimization problem by gradient descent. It is straightforward to see that the convexity also holds with convolutional layers. When calculating the

¹¹The exact form of Ξ_3 can be found in the Appendix D.5 in the supplementary material.

¹²See Appendix E in the supplementary material.

normalized sharpness, we need to ensure that the choice of surrogate loss function does not introduce other scale dependences. We discuss this choice in Appendix G in the supplementary material

8. Numerical evaluation

We numerically evaluated the effectiveness of normalized sharpness (35). We mainly validated the following two points.

- Unnormalized sharpness metrics can fail to predict generalization performance.
- Normalized sharpness (35) better predicts generalization than unnormalized metrics.

For these purposes, we checked the metrics’ ability to distinguish models trained on random labels (Sec. 8.1) and predict the generalization performance of models trained with different hyperparameters (Sec. 8.2). As existing current sharpness metrics, we used the trace of the Hessian without normalization (6) and the sum of the squared Frobenius norm of the weight matrices (Neyshabur et al., 2017). We also investigated the empirical dependence of normalized sharpness on the width and depth of networks (Sec. 8.3). In all experiments, we trained multilayer perceptrons with three hidden layers, LeNet (Lecun et al., 1998), and Wide ResNet (Zagoruyko & Komodakis, 2016) with 16 layers and width factor 4 on MNIST (LeCun et al., 1998) and CIFAR-10 (Krizhevsky, 2009). More detailed experimental setups are described in Appendix M in the supplementary material.

8.1. Random labels

We first investigated whether normalized sharpness (35) and current unnormalized sharpness metrics can distinguish models trained on random labels. In this experiment, sharper minima are expected to indicate larger generalization gaps.

Results: Figure 1 shows plots of the mean sharpness metrics for models trained on datasets with different random label ratios. The results show that networks trained on random labels had larger normalized sharpness to fit the random labels. Thus, we can say that normalized sharpness provides a fairly good hierarchy in the hypothesis class. Even though sharpness without normalization and the squared Frobenius norm of weight matrices could also distinguish models trained on random labels to some extent, the signal was much weaker than that of normalized sharpness. Note that normalized sharpness has an advantage in that it does not require a biplot for both the unnormalized sharpness and the squared Frobenius norm.

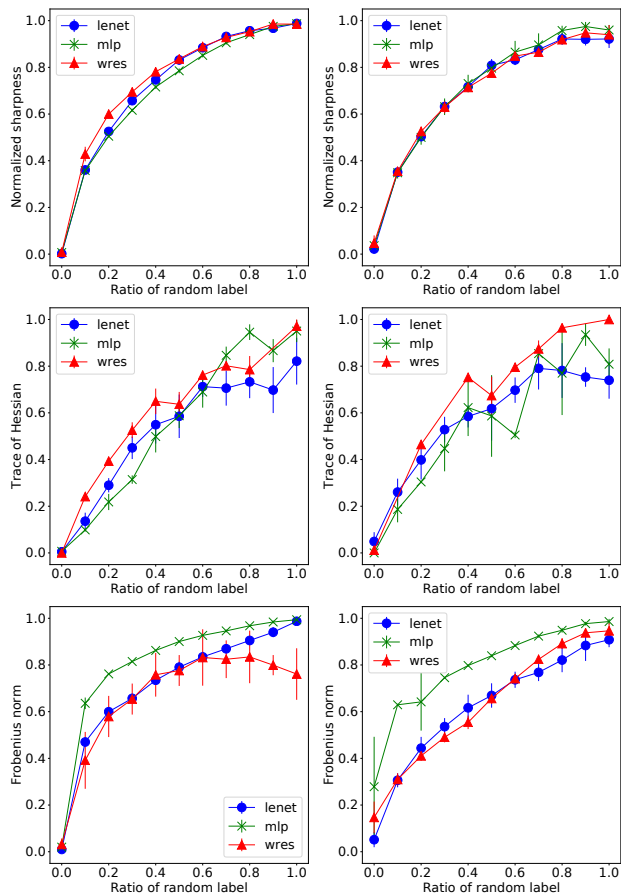


Figure 1. The figure shows the values of normalized sharpness (35), trace of the Hessian (6), and the sum of the squared Frobenius norms of weight matrices of models trained on datasets with different random label ratios. Each dot represents the mean of the generalization metric among trained networks and the error bars represent the standard deviation. The left column is the results on MNIST and the right is on CIFAR10. All generalization metrics were rescaled to $[0, 1]$ by their maximum and minimum per network architecture.

8.2. Different hyperparameters

We tested whether normalized sharpness and the existing unnormalized sharpness metrics can predict the generalization performance of models trained with different hyperparameters. We then checked the correlations between the generalization metrics and their generalization gaps, defined by $(\text{test misclassification ratio}) - (\text{train misclassification ratio})$. Models are trained with different strengths of ℓ_2 -regularization, weight decay (Loshchilov & Hutter, 2019), and dropout. This is an adversarial setting for existing unnormalized sharpness metrics because both regularizations on weights and dropout directly affects the balance between parameter scales and flatness of minima.

Table 1. Correlation coefficients between generalization gap and the generalization metrics for models trained on MNIST.

	MLP	Lenet	WRResNet
Normalized sharpness	0.73	0.73	0.59
Trace of Hessian	-0.62	-0.79	-0.59
Frobenius norm	0.58	0.58	0.49

Table 2. Correlation coefficients between generalization gap and the generalization metrics for models trained on CIFAR10.

	MLP	Lenet	WRResNet
Normalized sharpness	0.92	0.98	0.92
Trace of Hessian	-0.42	-0.51	-0.53
Frobenius norm	0.72	0.76	0.43

Results: We summarized the correlation coefficients between curvature metrics and the accuracy gaps in Table 1 and Table 2. Scatter plots can be found in Appendix N.1 in the supplementary material with more detailed results. On CIFAR10, especially for LeNet and Wide ResNet, we observed almost linear correlations between normalized sharpness and accuracy gap. Thus, we can confirm the usefulness of our generalization error bounds and normalized sharpness. On MNIST, even though there were weak correlations between normalized sharpness and the accuracy gap, the correlation was weaker than the results on CIFAR10. A possible explanation of this phenomenon is that the scale of the accuracy gap was too small on MNIST, which was at most 0.02. Since we could not create models with various accuracy gaps by merely changing the regularization parameters on MNIST, the effect of noise would have become larger. In all settings, both the trace of the Hessian and the squared Frobenius norm of the weight matrices had a weaker ability to predict the generalization. We can confirm that normalized sharpness had consistently stronger correlations with generalization. Notably, we observed a negative correlation with the trace of the Hessian and positive correlation with the squared Frobenius norm. If we use the tuples of the two as a generalization metric (Neysshabur et al., 2017), we cannot determine which models will generalize better. We explain the negative correlations as follows. In our experiments, we used three regularizers: ℓ_2 -regularization, weight-decay, and dropout. When we use ℓ_2 -regularization or weight-decay, stronger regularization decreases the generalization error and parameter scales. Since there are certain trade-offs between the sharpness and parameter scales, as shown in Eq.(4), the stronger regularization makes the Hessian larger.




Figure 2. Plot of the dependence of normalized sharpness on width and depth. The top figure plots the dependence on width of the normalized sharpness of multilayer-perceptron with different width. The bottom figure plots those on depth. The normalized sharpness showed almost linear dependence on network depth and sublinear dependence on network width.

8.3. Dependence on width and depth

To verify that the reduction of the constant term presented in Sec. 6 improves the overall bounds, we empirically investigated the dependency of normalized sharpness on network width and depth. We trained multilayer perceptron with depth- $\{1, 2, 3, 4\}$ and width- $\{32, 128, 512, 2048\}$ on MNIST three times for each and calculated normalized sharpness for every trained model. Figure 2 shows the results. It indicates that the normalized sharpness has an almost linear dependence on the network depth and sub-linear dependence on the network width, at least under this setting. This suggests that the constant term reduction presented in Sec. 6 contributes to tightening the overall generalization bounds.

9. Comparison with other generalization bounds:

We compare the tightness of our bounds and a prior PAC-Bayesian approach¹³ (Neyshabur et al., 2018). First, we remark that Neyshabur et al. (2018) relies on a margin parameter. The margin parameter controls the term (A) in Eq. (4) to some extent. Importantly, the margin loss they rely on does not decrease when the sample size increases. On the other hand the term (A) is solely controlled by the posterior choice in our bounds. As a result, the term decrease as the sample size increases. However, the difference makes the rate of convergence different and their direct comparison harder. The convergence rate is further discussed in Appendix I in the supplementary material. Nevertheless, we can still compare the effect of our constant term, which is $O(hd)$, with existing bounds. The main bound in Neyshabur et al. (2018) scales at least $O(d^2h)$. Compared to that, our constant term will not be critical to the tightness of our bound.

An $O(hd)$ generalization bound corresponds to $O(1/h)$ compression per layer. This is fairly competitive with the reported result by Arora et al. (2018). Even though our work is not directly comparable to Arora et al. (2018), this suggests that our bounds are comparably tight compared to those derived by Arora et al. (2018).

Observations in Sec. 8.3 suggest an advantage of normalized sharpness over the Fisher-Rao norm. Since the Fisher-Rao norm scales $O(d^2)$ with respect to depth, when the represented function is identical (Liang et al., 2019), normalized sharpness might be more robust against architecture changes, especially concerning depth of networks.

10. Conclusion

We have formally connected normalized loss curvatures with generalization through PAC-Bayesian analysis. The analysis bridged the known gap between theoretical understandings and empirical connections between normalized loss-landscape and generalization. The proof consists of two steps: scale-direction decompositions of parameters and discretization trick. In the analysis, we found that using a smaller number of scaling parameters is critical for meaningful generalization bounds, at least within our framework. Applying this discovery, we proposed *normalized sharpness* as a novel generalization metric. Experimental results suggest that this metric is more powerful than unnormalized loss sharpness metrics as a measure of generalization.

¹³Remark: Our primal goal is not providing the state-of-the-art tightest bound, but connecting scale-invariant flatness metric and generalization.

ACKNOWLEDGEMENTS

YT was supported by Toyota/Dwango AI scholarship. IS was supported by KAKENHI 20H04239. MS was supported by the International Research Center for Neurointelligence (WPI-IRCIN) at The University of Tokyo Institutes for Advanced Study.

References

- Achille, A. and Soatto, S. Emergence of Invariance and Disentanglement in Deep Representations. *Journal of Machine Learning Research*, 19, 2018.
- Alquier, P., Ridgway, J., and Chopin, N. On the properties of variational approximations of Gibbs posteriors. *Journal of Machine Learning Research*, 17(239):1–41, 2016.
- Arora, S., Ge, R., Neyshabur, B., and Zhang, Y. Stronger Generalization Bounds for Deep Nets via a Compression Approach. In *Proceedings of the 35th International Conference on Machine Learning*, pp. 254–263, 2018.
- Bartlett, P. L., Foster, D. J., and Telgarsky, M. J. Spectrally-normalized margin bounds for neural networks. In *Advances in Neural Information Processing Systems 30*, pp. 6240–6249. 2017.
- Catoni, O. *Pac-Bayesian Supervised Classification: The Thermodynamics of Statistical Learning*. Institute of Mathematical Statistics, 2007.
- Chaudhari, P., Choromanska, A., Soatto, S., LeCun, Y., Baldassi, C., Borgs, C., Chayes, J., Sagun, L., and Zecchina, R. Entropy-SGD: Biasing Gradient Descent Into Wide Valleys. In *International Conference on Learning Representations*, 2017.
- Devlin, J., Chang, M., Lee, K., and Toutanova, K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pp. 4171–4186, 2019.
- Dinh, L., Pascanu, R., Bengio, S., and Bengio, Y. Sharp Minima Can Generalize For Deep Nets. In *Proceedings of the 34th International Conference on Machine Learning*, pp. 1019–1028, 2017.
- Dziugaite, G. K. and Roy, D. M. Computing Nonvacuous Generalization Bounds for Deep (Stochastic) Neural Networks with Many More Parameters than Training Data. In *Proceedings of the Thirty-Third Conference on Uncertainty in Artificial Intelligence*, 2017.
- Dziugaite, G. K. and Roy, D. M. Data-dependent PAC-Bayes priors via differential privacy. In *Advances in Neural Information Processing Systems 31*, pp. 8430–8441. 2018.
- Germain, P., Bach, F., Lacoste, A., and Lacoste-Julien, S. PAC-Bayesian Theory Meets Bayesian Inference. In *Advances in Neural Information Processing Systems 29*, pp. 1884–1892. 2016.
- Guedj, B. A Primer on PAC-Bayesian Learning. *ArXiv e-prints*, abs/1901.05353, 2019.
- Hinton, G. E. and van Camp, D. Keeping the Neural Networks Simple by Minimizing the Description Length of the Weights. In *Proceedings of the Sixth Annual Conference on Computational Learning Theory, COLT '93*, pp. 5–13, 1993.
- Hochreiter, S. and Schmidhuber, J. Flat Minima. *Neural Computation*, 9(1):1–42, 1997.
- Hoffer, E., Hubara, I., and Soudry, D. Train longer, generalize better: closing the generalization gap in large batch training of neural networks. In *Advances in Neural Information Processing Systems 30*, pp. 1731–1741. 2017.
- Honkela, A. and Valpola, H. Variational Learning and Bits-back Coding: An Information-theoretic View to Bayesian Learning. *Trans. Neur. Netw.*, 15(4):800–810, 2004. ISSN 1045-9227.
- Ioffe, S. and Szegedy, C. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *Proceedings of the 32nd International Conference on Machine Learning*, pp. 448–456, 2015.
- Jiang, Y., Neyshabur, B., Krishnan, D., Mobahi, H., and Bengio, S. Fantastic Generalization Measures and Where to Find Them. In *International Conference on Learning Representations*, 2020.
- Keskar, N. S., Mudigere, D., Nocedal, J., Smelyanskiy, M., and Tang, P. T. P. On Large-Batch Training for Deep Learning: Generalization Gap and Sharp Minima. In *International Conference on Learning Representations*, 2017.
- Kingma, D. P., Salimans, T., and Welling, M. Variational Dropout and the Local Reparameterization Trick. In *Advances in Neural Information Processing Systems 28*, pp. 2575–2583. 2015.
- Krizhevsky, A. Learning Multiple Layers of Features from Tiny Images. *Master’s thesis, Department of Computer Science, University of Toronto*, 2009.

- Krizhevsky, A., Sutskever, I., and Hinton, G. E. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems 25*, pp. 1097–1105. 2012.
- Langford, J. and Caruana, R. (Not) Bounding the True Error. In *Advances in Neural Information Processing Systems 14*, pp. 809–816. 2002.
- Lecun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based Learning Applied to Document Recognition. In *Proceedings of the IEEE*, pp. 2278–2324, 1998.
- LeCun, Y., Cortes, C., and Burges, C. J. C. The MNIST Database of Handwritten Digits. 1998.
- Letarte, G., Germain, P., Guedj, B., and Laviolette, F. Dichotomize and Generalize: PAC-Bayesian Binary Activated Deep Neural Networks. In *Advances in Neural Information Processing Systems 32*, pp. 6869–6879. 2019.
- Li, H., Xu, Z., Taylor, G., and Goldstein, T. Visualizing the Loss Landscape of Neural Nets. In *Advances in Neural Information Processing Systems 31*, pp. 6391–6401. 2018.
- Liang, T., Poggio, T., Rakhlin, A., and Stokes, J. Fisher-Rao Metric, Geometry, and Complexity of Neural Networks. In *Proceedings of Machine Learning Research*, volume 89, pp. 888–896, 2019.
- Loshchilov, I. and Hutter, F. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*, 2019.
- Nagarajan, V. and Kolter, Z. Deterministic PAC-Bayesian generalization bounds for deep networks via generalizing noise-resilience. In *International Conference on Learning Representations*, 2019.
- Neklyudov, K., Molchanov, D., Ashukha, A., and Vetrov, D. P. Structured Bayesian Pruning via Log-Normal Multiplicative Noise. In *Advances in Neural Information Processing Systems 30*, pp. 6775–6784. 2017.
- Neyshabur, B., Salakhutdinov, R. R., and Srebro, N. Path-SGD: Path-Normalized Optimization in Deep Neural Networks. In *Advances in Neural Information Processing Systems 28*, pp. 2422–2430. 2015.
- Neyshabur, B., Bhojanapalli, S., McAllester, D., and Srebro, N. Exploring Generalization in Deep Learning. In *Advances in Neural Information Processing Systems 30*, pp. 5947–5956. 2017.
- Neyshabur, B., Bhojanapalli, S., and Srebro, N. A PAC-Bayesian Approach to Spectrally-Normalized Margin Bounds for Neural Networks. In *International Conference on Learning Representations*, 2018.
- Parrado-Hernández, E., Ambroladze, A., Shawe-Taylor, J., and Sun, S. PAC-Bayes Bounds with Data Dependent Priors. *Journal of Machine Learning Research*, 13:3507–3531, 2012.
- Rissanen, J. Stochastic Complexity and Modeling. *Ann. Statist.*, 14(3):1080–1100, 09 1986.
- Rothblum, U. G. and Zenios, S. A. Scalings of Matrices Satisfying Line-Product Constraints and Generalizations. *Linear Algebra and Its Applications*, pp. 159–175, 1992.
- Salimans, T. and Kingma, D. P. Weight Normalization: A Simple Reparameterization to Accelerate Training of Deep Neural Networks. In *Advances in Neural Information Processing Systems 29*, pp. 901–909. 2016.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. Intriguing properties of neural networks. In *International Conference on Learning Representations*, 2014.
- Uhlmann, J. A Generalized Matrix Inverse that is Consistent with Respect to Diagonal Transformations. *SIAM Journal on Matrix Analysis and Applications*, pp. 781–800, 2018.
- van den Oord, A., Li, Y., Babuschkin, I., Simonyan, K., Vinyals, O., Kavukcuoglu, K., van den Driessche, G., Lockhart, E., Cobo, L., Stimberg, F., Casagrande, N., Grewe, D., Noury, S., Dieleman, S., Elsen, E., Kalchbrenner, N., Zen, H., Graves, A., King, H., Walters, T., Belov, D., and Hassabis, D. Parallel WaveNet: Fast High-Fidelity Speech Synthesis. In *Proceedings of the 35th International Conference on Machine Learning*, pp. 3918–3926, 2018.
- Wang, H., Shirish Keskar, N., Xiong, C., and Socher, R. Identifying Generalization Properties in Neural Networks. *ArXiv e-prints*, abs/1809.07402, 2018.
- Yao, Z., Gholami, A., Lei, Q., Keutzer, K., and Mahoney, M. W. Hessian-based Analysis of Large Batch Training and Robustness to Adversaries. In *Advances in Neural Information Processing Systems 31*, pp. 4954–4964. 2018.
- Zagoruyko, S. and Komodakis, N. Wide Residual Networks. In *Proceedings of the British Machine Vision Conference*, pp. 87.1–87.12, 2016.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. Understanding Deep Learning Requires Rethinking Generalization. In *International Conference on Learning Representations*, 2017.
- Zhang, G., Wang, C., Xu, B., and Grosse, R. Three Mechanisms of Weight Decay Regularization. In *International Conference on Learning Representations*, 2019.

Zhou, W., Veitch, V., Austern, M., Adams, R. P., and Orbanz, P. Non-vacuous Generalization Bounds at the ImageNet Scale: a PAC-Bayesian Compression Approach. In *International Conference on Learning Representations*, 2019.