Improving Generalization of Deep Neural Networks by Leveraging Margin Distribution

Author: Shen-Huan LYU, Lu WANG, and Zhi-Hua ZHOU (Corresponding author)

A pre-printed manuscript which is accepted by Neural Networks



Learning And Mining from DatA

LAMDA Group Department of Computer Science and Technology National Key Laboratory for Novel Software Technology



Citation: Shen-Huan Lyu, Lu Wang, and Zhi-Hua Zhou. Improving generalization of neural networks by leveraging margin distribution. Neural Networks, 151:48-60, 2022. doi: https://doi.org/10.1016/j.neunet.2022.03.019.

1	1 Introduction			
2 Related work				
3	Notations	3		
4	Margin distribution rather than minimum margin	4		
	4.1 Error-resilience assumptions	4		
	4.2 Existed results	5		
	4.3 Main results	6		
5	Proofs	9		
	5.1 Proof of Lemma 4.9	10		
	5.2 Proof of Lemma 4.10	12		
	5.3 Proof of Theorem 4.11	13		
6	Optimizing margin distribution measure	14		
7	Experiments	14		
	7.1 Configuration	14		
	7.2 Ablation study	15		
	7.3 Feature visualization	15		
	7.4 Controlling capacity	16		
	7.5 Influence of the hyper-parameters	16		
8	Conclusion	17		

Improving Generalization of Deep Neural Networks by Leveraging Margin Distribution*

Shen-Huan Lyu[†] Lu Wang[‡] Zhi-Hua Zhou[§]

Abstract

Recent research has used margin theory to analyze the generalization performance for deep neural networks (DNNs). The existed results are almost based on the spectrally-normalized minimum margin. However, optimizing the minimum margin ignores a mass of information about the entire margin distribution, which is crucial to generalization performance. In this paper, we prove a generalization upper bound dominated by the statistics of the entire margin distribution. Compared with the minimum margin bounds, our bound highlights an important measure for controlling the complexity, which is the ratio of the margin standard deviation to the expected margin. We utilize a convex margin distribution loss function on the deep neural networks to validate our theoretical results by optimizing the margin ratio. Experiments and visualizations confirm the effectiveness of our approach and the correlation between generalization gap and margin ratio.

1 Introduction

Deep neural networks (DNNs) are making major advances in solving problems that have resisted the best attempts of the artificial intelligence community for many years [LBH15], especially in the field of computer vision [Gor22]. Recently, many research try to explain the practical success of DNNs via generalization, which is the ability of a classifier to perform well on unseen samples. However, some new empirical evidence has started to question this explanation. Adversarial training samples can cause the model to misclassify seriously by slight feature perturbation [GSS15; Pap+17]. On the other hand, Zhang et al. [Zha+21] find that the deep neural networks have enough complexity to fit an arbitrarily corrupted data, and a small geometric transformation may cause networks deteriorating in performance [AW19]. This complex and fragile nature of DNNs leads to a key problem, how to use the *data distribution* and *network parameters* to estimate the generalization ability of DNNs. Although several regularization techniques, such as dropout [Sri+14], batch normalization [IS15], and weight decay [KH92], do improve the generalization performance of the over-parameterized deep models, Zhang et al. [Zha+21] show that these regularizers cannot solve this problem either.

Consequently, several recent works [NTS15; BFT17; NBS18; Aro+18] have started to address this question, proving that we can control the capacity of DNNs via different upper bounds based on the minimum margin. However, the generalization bounds based on analyses of model complexity and noise stability only focus on the minimum margin, which is based on the closest distance of the training points to the decision boundary. This notion is brittle and sensitive to outliers due to a lack of the entire margin distribution information. Jiang et al. [Jia+19] propose a measure by looking at the entire distribution of distances, and conduct empirical studies on how well it can predict the generalization gap. However, how the margin distribution information affects the generalization error of the model still

^{*}This is a preprinted manuscript.

[†]Nanjing University. E-mail: lvsh@lamda.nju.edu.cn

[‡]Nanjing University. E-mail: wangl@lamda.nju.edu.cn

[§]Corresponding author. Nanjing University. E-mail: zhouzh@lamda.nju.edu.cn

needs more specific theoretical analysis, which will lead us to optimize the entire margin distribution appropriately.

The margin distribution has been shown to correspond to generalization properties in the literature on linear models and boosting algorithms, Schapire et al. [Sch+97] first introduce it to explain the phenomenon that AdaBoost seems resistant to overfitting problem. Two years later, Breiman [Bre99] indicates that the minimum margin is crucial for margin theory. Reyzin and Schapire [RS06] conjecture that the margin distribution, rather than the minimum margin, plays a key role. The debate has been finally solved by Gao and Zhou [GZ13] who theoretically proved that the AdaBoost process attempts to maximize the margin mean and minimize margin variance simultaneously; highlighting for the first time that two factors rather than a single factor are crucial for margin theory. These two factors are the first and second-order statistics describing the margin distribution, while in most cases the higherorder one is less useful. Their result successfully explains why AdaBoost seems resistant to overfitting: even when the training error reaches zero, the margin mean can be increased and/or the margin variance can be decreased further, leading to the improvement of generalization performance; it also discloses that AdaBoost will finally overfit: when the margin mean cannot be increased and margin variance cannot be decreased further. The long march of the theoretical exploration of AdaBoost is summarized in Zhou [Zho14], and Gao and Zhou [GZ13]'s result has been theoretically confirmed by Grønlund et al. [Grø+19]. Inspired by Gao and Zhou [GZ13]'s finding, powerful learning machines can be built by maximizing the margin mean and minimizing margin variance simultaneously, rather than simply maximizing the minimum margin like in traditional large-margin machines. Zhang and Zhou [ZZ19] propose the optimal margin distribution machine (ODM) for binary classification. In Zhang and Zhou [ZZ17; ZZ18a; ZZ18c], Tan et al. [Tan+20], and Zhang, Zhao, and Jin [ZZJ20], ODM is extended to many forms.

In this paper, we study a *d*-layer feed-forward network with ReLU activation functions. Our theoretical result states that the statistics of margin distribution play an important role in the generalization estimation rather than the traditional minimum margin. This result is consistent with the previous results on boosting and linear algorithms [Sch+97; GZ13; ZZ19]. It also inspires us to understand the similarities between deep learning and traditional machine learning from the perspective of margin distribution. Specially, we propose a new loss function to optimize the statistics-based measure in the theoretical results. A strong correlation between generalization and our measure is empirically shown by studying a wide range of network structures trained on the MNIST, CIFAR-10 and ImageNet datasets. The detailed contributions of this paper are as follows:

PAC guarantee Our bound shows that we can restrict the capacity of deep nets by the ratio of secondto first-order statistic of margin distribution at the last layer. Compared with the existing results based on minimum margin [BFT17; NBS18; Aro+18], our bound contains more information on the entire margin distribution to estimate the generalization error. Moreover, the empirical evaluation shows that optimizing the margin ratio can control the model capacity to alleviate the overfitting risk.

Optimization Inspired by our theoretical result, we encourage DNNs to optimize the margin ratio for better generalization performance. Therefore, we propose a new approach called <u>margin distribution</u> <u>Net</u>works (mdNet), which utilizes a convex margin distribution loss function to optimize the first- and second-order statistics of margin. Moreover, we empirically evaluate our loss function on deep neural networks across different image datasets and model structures. Specifically, empirical results demonstrate the effectiveness of the proposed method in learning tasks with limited training data.

The rest of paper is organized as follows. The related work is introduced in Section 2. Some notations are introduced in Section 3. In Section 4, we present a generalization bound leveraging margin distribution rather than minimum margin and demonstrate that the ratio of the margin standard deviation

to the expected margin is the key to control the model capacity. Section 5 list the detailed proofs for our theorems and lemmas. In Section 6, we formulate the convex loss function to optimize the margin ratio. Section 7 reports our experimental studies and empirical observations. Finally, Section 8 concludes with future work.

2 Related work

Recently, margin-based deep learning algorithms have developed rapidly. Schroff, Kalenichenko, and Philbin [SKP15] use the triplet loss to encourage a distance constraint similar to the contrastive loss. Similarly, Chan et al. [Cha+15] enhance the supervision of the learned filters by incorporating the information of class labels in the training data and learn the filters based on the idea of multi-class linear discriminant analysis (LDA) for classification task. Liu et al. [Liu+16] propose a generalized largemargin softmax loss which explicitly encourages intra-class compactness and inter-class separability in the learned representation space. It would be interesting to theoretically study feature space transformation which might be a key to understanding mysteries behind the successes of deep neural networks [Zho21]. Since Bartlett, Foster, and Telgarsky [BFT17] and Arora et al. [Aro+18] associate the generalization of deep neural networks with the minimum margin, a line of work establishes that firstorder methods can automatically maximizing the minimum margin in the settings of logistic regression [Gun+18a], deep linear networks [Sou+18; Gun+18b; JT19; LMZ18], and symmetric matrix factorization [LMZ18]. However, Wei et al. [Wei+18] point that how to extend these results to non-linear neural networks remains unclear. Recently, Wu et al. [Wu+21] propose to understand the model dynamics from the perspective of control theory. Another line of algorithm-dependent analysis of generalization [HRS16; Mou+18; CJY18] uses stability of specific optimization algorithms that satisfy certain generic properties like convexity, smoothness, etc. Specially, Keskar et al. [Kes+17], Dinh et al. [Din+17], and Zhu et al. [Zhu+19] make a connection between the sharpness of the solution obtained using the SGD algorithm and its ability to generalize well. The notion of sharpness corresponds to robustness to adversarial perturbations of parameters. Furthermore, Neyshabur et al. [Ney+17] and Neyshabur, Bhojanapalli, and Srebro [NBS18] draw a connection to the PAC-Bayesian theory for sharpness. The margin distribution measure presented in this paper is closely related to sharpness [Kes+17], because we use the statistics of the margin distribution to theoretically describe the value of the allowable perturbation. Compared with the sharpness measure which is difficult to optimize, the margin distribution measure proposed in this paper is easy to calculate, and can be directly optimized through the SGD algorithm by designing a convex loss function. Recently, Jiang et al. [Jia+19] present abundant empirical evidence to validate that the generalization in deep learning can be estimated from the margin statistics. In addition, the relevant theories of domain adaptation [MMR09; ZZY12; MS14] are also used to improve the generalization capability of deep learning [Pan+09; BCF13; KTP17; RSF19]. Domain generalization cannot see existing training source domains during training. This makes domain generalization more challenging than domain adaptation but more realistic and favorable in practical applications [Ghi+15; Dub+21; Wan+21; Mat+22].

3 Notations

Consider the multi-class task with feature domain \mathcal{X} and label domain \mathcal{Y} . Let \mathcal{D} be an unknown (underlying) distribution over $\mathcal{X} \times \mathcal{Y}$. A training set $S = \{(x_1, y_1), \ldots, (x_m, y_m)\}$ and a validation set $S' = \{(x_1, y_1), \ldots, (x_{m'}, y_{m'})\}$ are drawn identically and independently according to \mathcal{D} . We denote a labeled sample as $(x, y) \in \mathcal{D}$.

Let $f_{\boldsymbol{w}} \colon \mathcal{X}_{B,n} \to \mathcal{Y}'$ be the function represented by a *d*-layer feed-forward network with parameters

$$\boldsymbol{v} = \{ \boldsymbol{W}_1, \boldsymbol{W}_2, \dots, \boldsymbol{W}_d \}$$

U

and output domain $\mathcal{Y}' = \mathbb{R}^k$. The entire network can be formulated as

$$f_{\boldsymbol{w}}(\boldsymbol{x}) = \boldsymbol{W}_{d}\phi(\boldsymbol{W}_{d-1}\phi(\ldots\phi(\boldsymbol{W}_{1}\boldsymbol{x}))),$$

where ϕ is the ReLU activation function and let ρ be an upper bound on the number of output units in each layer.

We can define the fully connected networks (FNNs) recursively:

$$\boldsymbol{x}^1 = \boldsymbol{W}_1 \boldsymbol{x}$$
 and $\boldsymbol{x}^i = \boldsymbol{W}_i \phi(\boldsymbol{x}^{i-1}),$

where x^i denotes the output of the *i*-th layer.

The predicted label is denoted by

$$h(\boldsymbol{x}) = rg\max_{j} f_{\boldsymbol{w},j}(\boldsymbol{x}) \in \mathcal{H},$$

where $h: \mathcal{X} \to \mathcal{Y}$ is a map from the feature domain to the label domain and $f_{w,j}$ is the *j*-th element of the score vector.

In the multi-class setting [MRT18a, Chapter 9.2], the label associated to point x is the one resulting in the largest score $h(x) = \arg \max_i f_{w,i}(x)$. This naturally leads to the following definition of the margin $\gamma_h(x, y)$ of the function h at a labeled example (x, y):

$$\gamma_h(\boldsymbol{x}, y) = f_{\boldsymbol{w}, y}(\boldsymbol{x}) - \max_{j \neq y} f_{\boldsymbol{w}, j}(\boldsymbol{x}).$$
(1)

Thus, h misclassifies (\boldsymbol{x}, y) iff $\gamma_h(\boldsymbol{x}, y) < 0$.

4 Margin distribution rather than minimum margin

In Subsection 4.1, we list error-resilience assumptions that will be used. In Subsection 4.2, we introduce the existed results based on the minimum margin. In Subsection 4.3, we present our main results based on the entire margin distribution.

4.1 Error-resilience assumptions

Here we formalize the error-resilience properties for deep neural networks. Arora et al. [Aro+18] show that if we inject a scaled Gaussian noise to the input of deep nets, as it propagates up, the noise has rapidly decreasing effect on higher layers. This fact implies *compressibility* of deep nets, i.e., low rank of parameters' matrix. The empirical version of noise-sensitivity parameters are first proposed by Arora et al. [Aro+18]. It inspires us to bound the perturbation caused by Gaussian noise with the validation-based version of noise-sensitivity parameters below.

Assumption 4.1. (Layer Cushion). The layer cushion of layer *i* is defined to be largest number μ_i such that for any validation data $x \in S'$:

$$\mu_i \| \boldsymbol{W}_i \|_F \| \phi(\boldsymbol{x}^{i-1}) \|_2 \le \| \boldsymbol{x}^i \|_2.$$
(2)

Assumption 4.2. (Interlayer Cushion). For any two layers i < j, we define the interlayer cushion $\mu_{i,j}$, as the largest number such that for any validation data $x \in S'$:

$$\mu_{i,j} \|J_{\boldsymbol{x}^{i}}^{i,j}\|_{F} \|\phi(\boldsymbol{x}^{i-1})\|_{2} \le \|\boldsymbol{x}^{j}\|_{2}.$$
(3)

Furthermore, for any layer *i* we define the minimal interlayer cushion as $\mu_{i\to} = \min_{i \le j \le L} \mu_{i,j} = \min_{\{\frac{1}{\sqrt{\rho}}, \min_{i \le j \le L} \mu_{i,j}\}}$. For any two layer i < j, denote by $M^{i,j}$ the operator for composition of these layers and $J^{i,j}_{\boldsymbol{x}}$ be the Jacobian matrix (the partial derivative) of this operator at input \boldsymbol{x} . Therefore, we have $\boldsymbol{x}^j = M^{i,j}(\boldsymbol{x}^i)$. Furthermore, since the activation functions are ReLU (hence piece-wise linear), we have $M^{i,j}(\boldsymbol{x}^i) = J^{i,j}_{\boldsymbol{x}^i} \boldsymbol{x}^i$.

Assumption 4.3. (Interlayer Smoothness). For any two layers i < j, we define the interlayer smoothness ρ_{δ} as the smallest number such that with probability $1 - \delta$ over noise η for any validation data $x \in S'$:

$$\|M^{i,j}(\boldsymbol{x}^{i}+\eta) - J^{i,j}_{\boldsymbol{x}^{i}}(\boldsymbol{x}^{i}+\eta)\| \leq \frac{\|\eta\| \|\boldsymbol{x}^{j}\|}{\rho_{\delta} \|\boldsymbol{x}^{i}\|}$$
(4)

For a single layer, ρ_{δ} captures the ratio of input/weight alignment to noise/weight alignment. Arora et al. [Aro+18] show that the interlayer smoothness is indeed good: $1/\rho_{\delta}$ is a small constant.

The next two conditions qualify a common appearance: if the input in the activation and margin calculations is well-distributed and the calculations do not correlate with the magnitude of the input, then one would expect that, the effect of applying activation at any layer and margin at last layer is to decrease the norm of the vector by at most some small constant factor, i.e., c and α .

Assumption 4.4. (Activation Contraction). The activation contraction c is defined as the smallest number such that for any layer i and any validation data $x \in S'$:

$$c\|\phi(\boldsymbol{x}^{i})\|_{2} \ge \|\boldsymbol{x}^{i}\|_{2}.$$
 (5)

Assumption 4.5. (Margin Contraction). The margin contraction α is defined as the smallest number such that for any validation data $x \in S'$:

$$\alpha \|\gamma_h(\boldsymbol{x}, \boldsymbol{y})\|_2 \ge \|\boldsymbol{x}^d\|_2. \tag{6}$$

In this paper, we only use the noise-sensitivity parameters in Assumptions 4.1 - 4.5 as descriptions of error-resilience properties, from which the margin distribution term of our bound is derived. Therefore, we just need estimate these parameters based on validation data to show the magnitude of our bound rather than optimizing these parameters in the training process like Arora et al. [Aro+18] did.

4.2 Existed results

In the deep learning theory community, great efforts have been made to explain why over-parameterized deep neural networks can success, which is contrary to the classical VC dimension analysis [BMM98; HLM17]. Bartlett, Foster, and Telgarsky [BFT17] and Neyshabur, Bhojanapalli, and Srebro [NBS18] made an important stride by showing minimum margin based bounds for multi-layer neural networks. These bounds do not depend directly on the number of parameters of the network but depends on the normalized minimum margin. Theorem 4.6 provides a unified description of these bound. The only difference between them lies in the value of constants and the type of norms.

Theorem 4.6. [*BFT17*; *NBS18*] For any $d, \rho > 0$ and $||\mathbf{x}||_2 \leq B$, let $f_{\mathbf{w}} : \mathcal{X} \to \mathbb{R}^k$ be a d-layer feedforward network with ReLU activation. Then, for any $\delta > 0$, with probability $\geq 1 - \delta$ over a training set of size m, for any $\mathbf{w} = \{\mathbf{W}_1, \mathbf{W}_2, \dots, \mathbf{W}_d\}$, we have:

$$L_0(f_{\boldsymbol{w}}) \leq \widehat{L}_{\gamma}(f_{\boldsymbol{w}}) + \mathcal{O}\left(\sqrt{\frac{B^2 d^2 \rho \ln(dh) \prod_{i=1}^d \|\boldsymbol{W}_i\|_2^2 \sum_{i=1}^d \frac{\|\boldsymbol{W}_i\|_F^2}{\|\boldsymbol{W}_i\|_2^2} + \ln \frac{dm}{\delta}}{\gamma^2 m}}\right),$$
(7)

where $L_0(\cdot)$ is the 0-1 loss, $\widehat{L}_{\gamma}(\cdot) = \Pr_S[\gamma_h(\boldsymbol{x}, y) \leq \gamma]$ is the empirical estimation of γ -margin loss and $\mathcal{O}(\cdot)$ describes the limiting behavior of a function.

Based on this margin theory view, Arora et al. [Aro+18] provide an improved bound by considering the compressibility of deep nets as follows:

Theorem 4.7. [Aro+18] For any d > 0, let $f_{\boldsymbol{w}} : \mathcal{X} \to \mathbb{R}^k$ be a d-layer feed-forward network with ReLU activation. Then, for any $\delta > 0$, with probability $\geq 1 - \delta$ over a training set of size m, for any $\boldsymbol{w} = \{\boldsymbol{W}_1, \boldsymbol{W}_2, \dots, \boldsymbol{W}_d\}$, we have:

$$L_0(f_{\boldsymbol{w}}) \le \widehat{L}_{\gamma}(f_{\boldsymbol{w}}) + \mathcal{O}\left(\sqrt{\frac{c^2 d^2 \max_{\boldsymbol{x} \in S} \|f_{\boldsymbol{w}}(\boldsymbol{x})\|_2^2 \sum_{i=1}^d \frac{1}{\mu_i^2 \mu_{i \to}^2}}{\gamma^2 m}}\right)$$
(8)

where $\mu_i, \mu_{i \rightarrow}, c, \alpha$ are layer cushion, interlayer cushion, activation contraction and interlayer smoothness defined in Assumptions 4.1, 4.2, 4.4 and 4.5 respectively

These existed results follow the traditional margin theory, so they only focus on the minimum margin γ . Because they lack the description of the entire margin distribution, they can only take the minimum margin γ as the optimization target to improve the generalization performance. These methods ignore the information of the entire margin distribution. In the next subsection, we expect to prove a bound related to the entire margin distribution, so as to inspire us to directly optimize margin distribution for improving the generalization performance for DNNs.

4.3 Main results

We begin with an intuitive comparison of the minimum margin based classifier and the margin distribution based classifier. Figure 1(a) shows that maximizing the minimum margin will make the classifier easy to be misled by a small number of samples, thus ignoring the distribution information of samples, while the margin distribution based classifier considers the mean and variance of samples and generalizes better.



(a) The minimum margin based classi- (b) The (r, θ) -margin distribution loss (c) The convex margin distribution loss fier (red line) and the margin distribution function (red line). The green dotted lines function (red line). This convex funcbased classifier (black line). The blue represent the confidence area of margins, tion is used as an alternative function triangle and green square represent the which has zero loss. Lemma 4.10 shows of (r, θ) -margin distribution loss functwo classes of instances, while the dotted that the perturbation caused by u (blue ar- tion, so that the deep neural networks can ellipses represent their underlying distribution to the generalization error. optimize the margin distribution through sGD algorithm.

Figure 1: Illustration of the margin distribution analysis and loss functions.

In order to utilize the mean and variance information into the theoretical analysis, we design a new margin loss, which uses r to adjust the mean of margin and θ to adjust the variance of margin. For any parameter $r > \theta > 0$, we can define a (r, θ) -margin distribution loss function (see Figure 1(b)), which penalizes h with a cost of 1 when it predicts x with a margin smaller than $r - \theta$, but also penalizes h

when it predicts x with a margin larger than $r + \theta$. The margin distribution bound is presented in terms of this loss function, which is formally defined as follows.

Definition 4.8. (*Expected margin distribution loss function*). For any $r > \theta > 0$, the (r, θ) -margin loss is the function $L_{r,\theta}(\cdot)$ defined for all $h \in \mathcal{H}$ as:

$$L_{r,\theta}(h) = \Pr_{\mathcal{D}} \left[\gamma_h(\boldsymbol{x}, y) \le r - \theta \right] + \Pr_{\mathcal{D}} \left[\gamma_h(\boldsymbol{x}, y) > r + \theta \right].$$
(9)

Intuitively, our (r, θ) -margin distribution loss function looks for a classifier h which forces as many data points as possible into the zero-loss band $(r - \theta \le \gamma_h(x, y) < r + \theta)$. Therefore, we let $r \simeq \mathbb{E}_{\mathcal{D}}[\gamma_h(x, y)], \theta^2 \simeq \operatorname{Var}_{\mathcal{D}}[\gamma_h(x, y)]$, which implies that the expected margin is larger than the standard deviation. Actually, θ just need to be a second-order statistic, so we can re-scale $\theta^2 = a \cdot \operatorname{Var}_{\mathcal{D}}[\gamma_h(x, y)]$ to satisfy $r > \theta$. In this way, the (r, θ) -margin distribution loss is a surrogate loss function. In particular, for $r = \theta$ and $\theta \to \infty$, the zero-loss band is the positive area $(\gamma_h(x, y) > 0)$ and $L_{r,\theta}$ corresponds to the 0-1 loss L_0 . Let $\hat{L}_{r,\theta}(f_w)$ be the empirical estimate of the expected margin distribution loss. So we also denote the expected risk and the empirical risk as $L_0(f_w)$ and $\hat{L}_0(f_w)$, which are bounded between 0 and 1.

We begin with bounding the change of output caused by the noise on the classifier u with the noisesensitivity parameters and the statistics of margin distribution:

Lemma 4.9. Let $f_{\boldsymbol{w}}: \mathcal{X} \to \mathbb{R}^k$ be a *d*-layer network. For any d > 0, and $vec(\{\boldsymbol{U}_i\}_{i=1}^d) = (\boldsymbol{U}_1, \boldsymbol{U}_2, \dots, \boldsymbol{U}_d)$ is a vector of perturbation parameters with $\boldsymbol{U}_i = \beta_i \| \boldsymbol{W}_i \|_F$, and $\boldsymbol{\beta} = vec(\{\beta_i\}_{i=1}^d) = (\beta_1, \beta_2, \dots, \beta_d)$ is a vector of random vectors with $\mathbb{E}[\boldsymbol{\beta}\boldsymbol{\beta}^\top] = \sigma^2 \boldsymbol{I}$, the change of the output of the network can be bounded with a fixed probability ($\delta = 1/2$):

$$\max_{\boldsymbol{x}\in\mathcal{X}}|f_{\boldsymbol{w}+\boldsymbol{u}}(\boldsymbol{x}) - f_{\boldsymbol{w}}(\boldsymbol{x})|_{2}^{2} \leq \mathcal{O}\left(\sum_{i=1}^{d} \frac{d\alpha^{2}c^{2}\sigma^{2}(r+\theta)^{2}}{\mu_{i}^{2}\mu_{i\rightarrow}^{2}}\right).$$
(10)

The result shows that the perturbation caused by \boldsymbol{u} increases with the variance σ^2 is related to the outermost edge of margin distribution $r + \theta$, that is, the right green dotted line in Figure 1(b). The next Lemma shows that we can bound the generalization gap through a Kullback-Leibler divergence term, if we can guarantee that the perturbation caused by \boldsymbol{u} is smaller than $\frac{r-\theta}{8}$ with a constant probability. Therefore, the allowable value of σ^2 under the (r, θ) -margin distribution assumption is consistent with the intuitive understanding (see Figure 3), i.e., $\sigma^2 \propto \frac{r-\theta}{r+\theta} = \frac{1-\lambda}{1+\lambda}$, where $\lambda = \theta/r \in (0, 1)$. When the margin distribution is more compact (smaller λ), the larger noise σ^2 can be allowed, that is, it is not easy to cause misclassification. When the margin distribution is more loose (larger λ), even a small noise have misclassification risk.

Lemma 4.10. Let $f_{\boldsymbol{w}} : \mathcal{X} \to \mathbb{R}^k$ be any predictor with parameters \boldsymbol{w} , and P be any distribution on the parameters that is independent of the training data. Then, for any $r > \theta > 0$, $\delta > 0$, with probability at least $1 - \delta$ over the training set of size m, for any \boldsymbol{w} , and any random perturbation \boldsymbol{u} s.t. $\Pr_{\boldsymbol{u}}\left[\max_{\boldsymbol{x}\in\mathcal{X}}|f_{\boldsymbol{w}+\boldsymbol{u}}(\boldsymbol{x}) - f_{\boldsymbol{w}}(\boldsymbol{x})|_2 < \frac{r-\theta}{8}\right] \geq \frac{1}{2}$, we have:

$$L_0(f_{\boldsymbol{w}}) \le \widehat{L}_{r,\theta}(f_{\boldsymbol{w}}) + 4\sqrt{\frac{D_{\mathrm{KL}}(\boldsymbol{w} + \boldsymbol{u} \| P) + \ln \frac{6m}{\delta}}{m-1}}.$$
(11)

The detailed proof is presented in Section 5.2. This Lemma improves the result of Neyshabur, Bhojanapalli, and Srebro [NBS18, Lemma 1], especially using two parameters θ , r to describe the entire margin distribution instead of using one parameter γ to describe the minimum margin. Based on this result, we can derive the following generalization bound, with proof deferred to Section 5.3 showing that the Kullback-Leibler divergence term is inversely proportional to σ^2 . **Theorem 4.11.** For any $d, \rho > 0$, let $f_{w} : \mathcal{X} \to \mathbb{R}^{k}$ be a *d*-layer feed-forward network with ReLU activation. Then, for any $\delta > 0$, with probability $\geq 1 - \delta$ over a training set of size *m*, for any *w*, we have:

$$L_0(f_{\boldsymbol{w}}) \le \widehat{L}_{r,\theta}(f_{\boldsymbol{w}}) + \mathcal{O}\left(\sqrt{\frac{\left(\frac{1+\lambda}{1-\lambda}\right)^2 \left(\sum_{i=1}^d \frac{d\alpha^2 c^2}{\mu_i^2 \mu_{i\to}^2}\right) + \ln \frac{6m}{\delta}}{m}}\right).$$
(12)

where the margin ratio is defined by $\lambda = \theta/r$.

We prove an upper bound on generalization gap related to the margin ratio term, where λ is a parameter denoting the ratio of the margin standard deviation θ to the expected margin r over the underlying distribution, and the error-resilience term relies on the noise sensitivity [Aro+18] quantified by $\mu_i, \mu_{i\rightarrow}, c, \alpha$ (See Assumptions 4.1, 4.2, 4.4 and 4.5). Theorem 4.11 states that the entire margin distribution has much leverage in generalization performance rather than the minimum margin. Specifically, restricting a smaller λ (larger r and smaller θ) can effectively control the capacity of models, so as to reduce the risk of overfitting. It inspires us that optimizing margin distribution can get better generalization performance than the traditional minimum margin maximization algorithm.



Figure 2: Comparing our bound and [1] to empirical generalization error during training. All bounds are rescaled to be within the same range as the generalization error together.

Discussion. The main difference between Arora et al. [Aro+18] and our paper: Arora et al. [Aro+18] proved that the generalization performance of deep neural network is related to the *sparsity* of its parameters, focusing on how to compress the parameters of the trained model. Our paper studies the relationship between DNN generalization performance and *margin distribution* under the condition that DNN parameters is sparsity, focusing on optimizing margin distribution during training. Figure 2 shows that the improvement of our bound relative to Arora et al. [Aro+18] (the shaded part in the figure) is because the margin ratio will gradually decrease during the training process. The main difference between Jiang et al. [Jia+19] and our paper: Jiang et al. [Jia+19] conjectured that the generalization performance of DNN may be related to the interval distribution. The correlation between generalization and R^2 [GS01] is calculated experimentally, and no theoretical proof is given. Our paper proves theoretically that the generalization performance of DNNs can be bound by the margin ratio and gives the improved algorithm.



Figure 3: Illustration of the relationship between margin distribution and allowable perturbation.

5 Proofs

In this section, we provide the detailed proofs for the main theorem and lemmas. First, we present a useful lemma as follows:

Lemma 5.1. Let Q be a probability distribution over the reals. For any random variable $v, v_1, v_2, \ldots, v_{m'} \sim Q$ identically and independently (i.i.d.), we have

$$\Pr_{v \sim \mathcal{Q}} \left[v \ge \max_{v_1, \dots, v_{m'} \sim \mathcal{Q}} \{ v_1, v_2, \dots, v_{m'} \} \right] = \frac{1}{m' + 1}.$$
(13)

Proof of Lemma 5.1: Let the Cumulative Distribution Function (CDF) and Probability Density Function (PDF) of random variable v be F(x) and f(x), and we denote the maximum of a set of m' random variables by $v^{(m')} = \max_{v_1,\ldots,v_{m'}\sim Q}\{v_1, v_2, \ldots, v_{m'}\}$. Then

$$\Pr_{v_1,\dots,v_{m'}\sim\mathcal{Q}}\left[v^{(m')} \le x\right] \tag{14}$$

$$= \Pr_{v_1, \dots, v_{m'} \sim \mathcal{Q}} \left[(v_1 \le x) \land \dots \land (v_{m'} \le x) \right]$$
(15)

$$=\Pr_{\mathcal{Q}}[v_1 \le x] \times \dots \times \Pr_{\mathcal{Q}}[v_{m'} \le x] = F^{m'}(x).$$
(16)

In other word, the CDF and PDF of the minimum $v^{(m')}$ are $F^{m'}(x)$ and $m'F^{m'-1}(x)f(x)$. Then we can use the minimum value of the sample's set to bound the random variable v with a probability $\frac{m'}{m'+1}$, which converges to 1 with a rate $\mathcal{O}(1/m')$:

$$\Pr_{v,v_1,\dots,v_{m'}\sim\mathcal{Q}}\left[v\leq v^{(m')}\right] \tag{17}$$

$$=\Pr_{v,v_1,\dots,v_{m'}\sim\mathcal{Q}}\left[v-v^{(m')}\leq 0\right]$$
(18)

$$= \iint_{x \le y} f(x) F^{m'-1}(y) f(y) \,\mathrm{d}x \,\mathrm{d}y \tag{19}$$

$$= \int_{-\infty}^{+\infty} m' F^{m'-1}(y) f(y) \, \mathrm{d}y \int_{-\infty}^{y} f(x) \, \mathrm{d}x$$
(20)

$$= \int_{-\infty}^{+\infty} m' F^{m'}(y) f(y) \,\mathrm{d}y \tag{21}$$

$$=m'F^{m'+1}(y)\mid_{-\infty}^{+\infty} - \int_{-\infty}^{+\infty} m'^2 F^{m'}(y)f(y) \,\mathrm{d}y.$$
(22)

According to Eq. (21) and Eq. (22), we have

$$\Pr_{v,v_1,\dots,v_{m'}\sim\mathcal{Q}}\left[v \le v^{(m')}\right] = \frac{m'}{m'+1}.$$
(23)

5.1 Proof of Lemma 4.9

We begin with a lemma as follows:

Lemma 5.2. For any layer *i*, the point-wise compressibility of the layer-wise parameters can holds with a probability $1 - \frac{1}{m'+1}$ over $x \in D$ as follows:

$$\mu_i \| \boldsymbol{W}_i \|_F \| \phi(\boldsymbol{x}^{i-1}) \|_2 \le \| \boldsymbol{x}^i \|_2, \tag{24}$$

$$\mu_{i,j} \|J_{\boldsymbol{x}^{i}}^{i,j}\|_{F} \|\phi(\boldsymbol{x}^{i-1})\|_{2} \le \|\boldsymbol{x}^{j}\|_{2},$$
(25)

$$\|M^{i,j}(\boldsymbol{x}^{i}+\eta) - J^{i,j}_{\boldsymbol{x}^{i}}(\boldsymbol{x}^{i}+\eta)\| \leq \frac{\|\eta\| \|\boldsymbol{x}^{j}\|}{\rho_{\delta} \|\boldsymbol{x}^{i}\|}$$
(26)

$$c\|\phi(\boldsymbol{x}^{i})\|_{2} \ge \|\boldsymbol{x}^{i}\|_{2},$$
 (27)

$$\alpha \|\gamma_h(\boldsymbol{x}, \boldsymbol{y})\|_2 \ge \|\boldsymbol{x}^d\|_2, \tag{28}$$

where m' is the size of the validation set.

Proof of Lemma 5.2: According the independence between S and S', we can regard the noise-sensitivity parameters $\frac{1}{\mu_i}$, $\frac{1}{\mu_{i,j}}$, c and α as random variables over reals relying on the randomness of variable $x \in S'$. Then, the cushion parameters defined in Assumptions 4.1, 4.2, 4.3, 4.4 and 4.5 can be interpreted as choosing the maximum of multiple independent samples. We first prove Lemma 5.1 on the tail of a random variable $v \sim Q$ by choosing the maximum of multiple independent samples of the random variable. Specifically, using the following simple lemma based on the distribution of the maximum, we can guarantee the point-wise *compressibility* of the learned parameters over the underlying data distribution D with a high probability by calculating the maximum of the empirical dataset, i.e., $\frac{1}{\mu_i}$, $\frac{1}{\mu_{i,j}}$, c and α .

Proof of Lemma 4.9: First, we need to bound the perturbation of linear operator caused by injecting a scaled Gaussian noise $U = \beta ||W||_F$, $\mathbb{E}[\beta\beta^\top] = \sigma I$. For any fixed vectors a, \mathbf{v} , we have

$$\mathbb{E}_{\beta} \| \boldsymbol{a}^{\top} (\boldsymbol{W} + \boldsymbol{U}) \mathbf{v} - \boldsymbol{a}^{\top} \boldsymbol{W} \mathbf{v} \|_{2} = \mathbb{E}_{\beta} \| \mathbf{v} \|_{2} \| \boldsymbol{a}^{\top} \boldsymbol{U} \boldsymbol{U}^{\top} \boldsymbol{a} \|_{2}$$

$$= \mathbb{E}_{\beta} \| \boldsymbol{W} \|_{F} \| \mathbf{v} \|_{2} \| \boldsymbol{a}^{\top} \beta \beta^{\top} \boldsymbol{a} \|_{2}$$

$$= \sigma \| \boldsymbol{W} \|_{F} \| \boldsymbol{a} \|_{2} \| \mathbf{v} \|_{2}.$$
(29)
(30)

According the Markov inequality, we have

$$\Pr_{\beta} \left[\|\boldsymbol{a}^{\top} (\boldsymbol{W} + \boldsymbol{U}) \mathbf{v} - \boldsymbol{a}^{\top} \boldsymbol{W} \mathbf{v}\|_{2} \ge \sigma \sqrt{d} / \sqrt{\delta} \|\boldsymbol{W}\|_{F} \|\boldsymbol{a}\|_{2} \|\mathbf{v}\|_{2} \right]$$

$$\leq \frac{\sigma^{2} \|\boldsymbol{W}\|_{F}^{2} \|\boldsymbol{a}\|_{2}^{2} \|\mathbf{v}\|_{2}^{2}}{(d\sigma^{2} / \delta) \|\boldsymbol{W}\|_{F}^{2} \|\boldsymbol{a}\|_{2}^{2} \|\mathbf{v}\|_{2}^{2}} = \frac{\delta}{d}.$$
(31)

Now, we will bound the perturbation of the *d*-layer deep nets by induction. For any layer $i \ge 0$, let x^j be the output at layer *j* with original net and \hat{x}_i^j be the output at layer *j* if the weights W_1, \ldots, W_i in the first layers are replaced with $W_1 + U_1, \ldots, W_i + U_i$. The induction hypothesis is then following:

Consider any $0 < \epsilon \le 1$, the following is true with probability $1 - \frac{i\delta}{d}$ over $W_1 + U_1, \ldots, W_i + U_i$ for any $j \ge i$:

$$\|\widehat{\boldsymbol{x}}_{i}^{j} - \boldsymbol{x}^{j}\|_{2}^{2} \leq \sum_{l=1}^{i} \frac{c^{2} d\sigma^{2}}{\delta \mu_{l}^{2} \mu_{l \to}^{2}} \|\boldsymbol{x}^{j}\|_{2}^{2}.$$
(32)

For the base case i = 0, since we are not perturbing the input, the inequality is trivial. Now assuming that the induction hypothesis is true for i - 1, we consider what happens at layer i.

$$\begin{aligned} \|\widehat{\boldsymbol{x}}_{i}^{j} - \boldsymbol{x}^{j}\|_{2}^{2} &= \|(\widehat{\boldsymbol{x}}_{i}^{j} - \widehat{\boldsymbol{x}}_{i-1}^{j}) + (\widehat{\boldsymbol{x}}_{i-1}^{j} - \boldsymbol{x}^{j})\|_{2}^{2} \\ &\leq 2\|(\widehat{\boldsymbol{x}}_{i}^{j} - \widehat{\boldsymbol{x}}_{i-1}^{j})\|_{2}^{2} + 2\|\widehat{\boldsymbol{x}}_{i-1}^{j} - \boldsymbol{x}^{j}\|_{2}^{2} \end{aligned} \tag{33}$$

The second term in Eq. (33) can be bounded by $\sum_{l=1}^{i-1} \frac{c^2 \sigma^2}{\mu_l^2 \mu_{l\rightarrow}^2} \| x^j \|_2^2$ by induction hypothesis. Therefore, it is enough to show that the first term in Eq. (33) is bounded by $\frac{c^2 \sigma^2}{\mu_i^2 \mu_{i\rightarrow}^2} \| x^j \|_2^2$. We decompose the error into two error terms one of which corresponds to the error propagation through the network if activation were fixed and the other one is the error caused by change in the activations:

$$\|(\widehat{\boldsymbol{x}}_{i}^{j} - \widehat{\boldsymbol{x}}_{i-1}^{j})\| \tag{34}$$

$$= \|M^{i,j}((\boldsymbol{W}_i + \boldsymbol{U}_i)\phi(\widehat{\boldsymbol{x}}^{i-1})) - M^{i,j}(\boldsymbol{W}_i\phi(\widehat{\boldsymbol{x}}^{i-1}))\|$$

$$= \|M^{i,j}((\boldsymbol{W}_i + \boldsymbol{U}_i)\phi(\widehat{\boldsymbol{x}}^{i-1})) - M^{i,j}(\boldsymbol{W}_i\phi(\widehat{\boldsymbol{x}}^{i-1}))$$
(35)

$$+ J_{\boldsymbol{x}^{i}}^{i,j}(\boldsymbol{U}^{i}\phi(\boldsymbol{\hat{x}}^{i-1})) - J_{\boldsymbol{x}^{i}}^{i,j}(\boldsymbol{U}^{i}\phi(\boldsymbol{\hat{x}}^{i-1})) \|$$
(36)

$$\leq \|J_{\boldsymbol{x}^{i}}^{i,j}(\boldsymbol{U}^{i}\phi(\widehat{\boldsymbol{x}}^{i-1}))\| + \|M^{i,j}((\boldsymbol{W}_{i}+\boldsymbol{U}_{i})\phi(\widehat{\boldsymbol{x}}^{i-1})) - M^{i,j}(\boldsymbol{W}_{i}\phi(\widehat{\boldsymbol{x}}^{i-1})) - J_{\boldsymbol{x}^{i}}^{i,j}(\boldsymbol{U}^{i}\phi(\widehat{\boldsymbol{x}}^{i-1}))\|.$$
(37)

The first term in Eq. (37) is bounded by:

$$\|J_{\boldsymbol{x}^{i}}^{i,j}\boldsymbol{U}^{i}\phi(\hat{\boldsymbol{x}}^{i-1})\|_{2}$$
(38)

$$\leq (\sqrt{d}\sigma/\sqrt{6\delta}) \|J_{\boldsymbol{x}^{i}}^{i,j}\|_{2} \|\boldsymbol{W}_{i}\|_{F} \|\phi(\boldsymbol{\widehat{x}}^{i-1})\|_{2}$$

$$(39)$$

$$\leq (\sqrt{d}\sigma/\sqrt{6\delta}) \|J_{\boldsymbol{x}^{i}}^{i,j}\|_{2} \|\boldsymbol{W}_{i}\|_{F} \|\widehat{\boldsymbol{x}}^{i-1}\|_{2}$$

$$\tag{40}$$

$$\leq (\sqrt{d\sigma}/\sqrt{3\delta}) \|J_{\boldsymbol{x}^{i}}^{i,j}\|_{2} \|\boldsymbol{W}_{i}\|_{F} \|\boldsymbol{x}^{i-1}\|_{2}$$

$$\tag{41}$$

$$\leq (c\sqrt{d\sigma}/\sqrt{3\delta}) \|J_{\boldsymbol{x}^{i}}^{i,j}\|_{2} \|\boldsymbol{W}_{i}\|_{F} \|\phi(\boldsymbol{x}^{i-1})\|_{2}$$
(42)

$$\leq (c\sqrt{d\sigma}/\sqrt{3\delta}\mu_i) \|J_{\boldsymbol{x}^i}^{i,j}\|_2 \|\boldsymbol{W}_i\phi(\boldsymbol{x}^{i-1})\|_2$$
(43)

$$\leq (c\sqrt{d\sigma}/\sqrt{3\delta\mu_i\mu_i}) \|\boldsymbol{x}^j\|_2. \tag{44}$$

where Equation (38) is bounded by Eq. (31), Equation (39) is bounded by Lipschitzness of the activation function, Equation (40) is bounded by inductive hypothesis, Equation (41) is bounded by activation contraction, Equation (42) is bounded by layer cushion, and Equation (44) is bounded by interlayer cushion.

The second term in Eq. (37) can be bounded as:

$$\|M^{i,j}((\boldsymbol{W}_{i} + \boldsymbol{U}_{i})\phi(\hat{\boldsymbol{x}}^{i-1})) - M^{i,j}(\boldsymbol{W}_{i}\phi(\hat{\boldsymbol{x}}^{i-1})) - J^{i,j}_{\boldsymbol{x}^{i}}(\boldsymbol{U}^{i}\phi(\hat{\boldsymbol{x}}^{i-1}))\|_{2}$$
(45)

$$= \| (M^{i,j} - J^{i,j}_{\boldsymbol{x}^{i}})((\boldsymbol{W}_{i} + \boldsymbol{U}_{i})\phi(\widehat{\boldsymbol{x}}^{i-1})) - (M^{i,j} - J^{i,j}_{\boldsymbol{x}^{i}})(\boldsymbol{W}_{i}\phi(\widehat{\boldsymbol{x}}^{i-1})) \|_{2}$$
(46)

$$= \| (M^{i,j} - J^{i,j}_{\boldsymbol{x}^{i}})((\boldsymbol{W}_{i} + \boldsymbol{U}_{i})\phi(\widehat{\boldsymbol{x}}^{i-1})) \|_{2} + \| (M^{i,j} - J^{i,j}_{\boldsymbol{x}^{i}})(\boldsymbol{W}_{i}\phi(\widehat{\boldsymbol{x}}^{i-1})) \|_{2}.$$
(47)

Both terms in Eq. (47) can be bounded using Assumption 4.3. By notations we find $W^i \phi(\hat{x}^{i-1}) =$ boun terms in Eq. (47) can be bounded using Assumption 4.3. By notations we find $W^i \phi(\hat{x}^{i-1}) = \hat{x}_{i-1}^i$. By induction hypothesis, we have that $\|W^i \phi(\hat{x}^{i-1}) - x^i\|_2^2 \leq \sum_{l=1}^{i-1} \frac{c^2 d\sigma^2}{\delta \mu_l^2 \mu_{l-2}^2} \|x^i\|_2^2$. Now by inter-layer smoothness property, $\|(M^{i,j} - J_{x^i}^{i,j})(W^i \phi(\hat{x}^{i-1}))\|_2^2 \leq \frac{\sum_{l=1}^{i-1} \frac{c^2 d\sigma^2}{\delta \mu_l^2 \mu_{l-2}^2} \|x^j\|}{\rho_\delta} \leq (\sum_{l=1}^{i-1} \frac{c^2 d\sigma^2}{\delta \mu_l^2 \mu_{l-2}^2}) \|x^j\|_2^2/(3d) \simeq \frac{i-1}{3d} \frac{c^2 d\sigma^2}{\delta \mu_l^2 \mu_{l-2}^2} \|x^j\|_2^2$. Similar to this term, $\|(M^{i,j} - J_{x^i}^{i,j})((W^i + U_i)\phi(\hat{x}^{i-1}))\| \leq (\sum_{l=1}^{i} \frac{c^2 d\sigma^2}{\delta \mu_l^2 \mu_{l-2}^2}) \|x^j\|/(3d) \simeq \frac{i}{3d} \frac{c^2 d\sigma^2}{\delta \mu_l^2 \mu_{l-2}^2} \|x^j\|_2^2$. Putting everything together completes the induction with probability at least $1 - \delta$ (if i = d)

i = d).

Instead of assuming that the input domain \mathcal{X} is bounded by a constant B, we assume that the input boundary is relative to the expected value which implies the data-distribution information: $\max_x \|x^d\|_2 \leq 1$ $\mathcal{O}(\mathbb{E}_{\mathcal{D}} \| \boldsymbol{x}^d \|_2)$. According to the margin contraction property, we can use the first- and second-statistics of the margin in the last layer $\mathbb{E}_{\mathcal{D}}[\gamma_h(\boldsymbol{x}, y)] = r, \operatorname{Var}_{\mathcal{D}}[\gamma_h(\boldsymbol{x}, y)] = \theta^2$ to bound the perturbation instead of the worst situation:

$$\max_{\boldsymbol{x}} \|\boldsymbol{x}^d\|_2^2 \le \mathcal{O}(\mathbb{E}_{\mathcal{D}} \|\boldsymbol{x}^d\|_2^2) \le \mathcal{O}(\alpha^2 \mathbb{E}_{\mathcal{D}} \|\gamma_h(\boldsymbol{x}, y)\|_2^2)$$
(48)

$$\leq \mathcal{O}(\alpha^2(r^2 + \theta^2)) \leq \mathcal{O}(\alpha^2(r + \theta)^2)$$
(49)

Connecting these two inequalities we prove that the equality holds with a probability at least 1/2:

$$|f_{\boldsymbol{w}+\boldsymbol{u}}(\boldsymbol{x}) - f_{\boldsymbol{w}}(\boldsymbol{x})|_{2}^{2} \leq \mathcal{O}\left(\sum_{i=1}^{d} \frac{d\alpha^{2}c^{2}(r+\theta)^{2}\sigma^{2}}{\mu_{i}^{2}\mu_{i\rightarrow}^{2}}\right).$$
(50)

_	_

5.2 Proof of Lemma 4.10

Proof of Lemma 4.10: Let w' = w + u, Let S_w be the set of perturbations with the following property:

$$\mathcal{S}_{\boldsymbol{w}} \subseteq \left\{ \boldsymbol{w}' \middle| \max_{\boldsymbol{x} \in \mathcal{X}} |f_{\boldsymbol{w}'}(\boldsymbol{x}) - f_{\boldsymbol{w}}(\boldsymbol{x})|_2 < \frac{r-\theta}{8\sqrt{\rho}} \right\},\tag{51}$$

then we will have $\max_{\boldsymbol{x}\in\mathcal{X}}|f_{\boldsymbol{w}'}(\boldsymbol{x})-f_{\boldsymbol{w}}(\boldsymbol{x})|_{\infty}<\sqrt{\rho}\max_{\boldsymbol{x}\in\mathcal{X}}|f_{\boldsymbol{w}'}(\boldsymbol{x})-f_{\boldsymbol{w}}(\boldsymbol{x})|_2<\frac{r-\theta}{8}.$

Let q be the probability density function over the parameters w'. We construct a new distribution \hat{Q} over predictors $f_{\tilde{w}}$ where \tilde{w} is restricted to \mathcal{S}_{w} with the probability density function:

$$\tilde{q}(\tilde{\boldsymbol{w}}) = \frac{1}{Z} \begin{cases} q(\tilde{\boldsymbol{w}}) & \tilde{\boldsymbol{w}} \in \mathcal{S}_{\boldsymbol{w}} \\ 0 & \text{otherwise} \end{cases}$$
(52)

According to the lemma assumption, we have $Z = \mathbb{P}[w' \in S_w] \geq \frac{1}{2}$. Therefore, we can bound the change of the margins for any instance:

$$\max_{i,j\in[k],\boldsymbol{x}\in\mathcal{X}} |(|f_{\tilde{\boldsymbol{w}}}(\boldsymbol{x})[i] - f_{\tilde{\boldsymbol{w}}}(\boldsymbol{x})[j]|) - (|f_{\boldsymbol{w}}(\boldsymbol{x})[i] - f_{\boldsymbol{w}}(\boldsymbol{x})[j]|)| < \frac{r-\theta}{2}$$
(53)

Here we define a perturbed loss function as:

$$L_{r,\theta}'(h) = \Pr_{\mathcal{D}}\left[\gamma_h(\boldsymbol{x}, y) \le \frac{r-\theta}{2}\right] + \Pr_{\mathcal{D}}\left[\gamma_h(\boldsymbol{x}, y) > r+\theta + \frac{r-\theta}{2}\right].$$
(54)

We can get the following:

$$L_0\left(f_{\boldsymbol{w}}\right) \le L_{r,\theta}'\left(f_{\tilde{\boldsymbol{w}}}\right) \tag{55}$$

$$\widehat{L}_{r,\theta}^{\prime}\left(f_{\widetilde{\boldsymbol{w}}}\right) \le \widehat{L}_{r,\theta}\left(f_{\boldsymbol{w}}\right) \tag{56}$$

Finally, using the proof of Neyshabur, Bhojanapalli, and Srebro [NBS18, Lemma 1], with probability $1 - \delta$ over the training set we have:

$$L_{0}\left(f_{\boldsymbol{w}}\right) \leq \mathbb{E}_{\tilde{\boldsymbol{w}}}\left[L_{r,\theta}'\left(f_{\tilde{\boldsymbol{w}}}\right)\right]$$
(57)

$$\leq \mathbb{E}_{\tilde{\boldsymbol{w}}}\left[\widehat{L}_{r,\theta}'\left(f_{\tilde{\boldsymbol{w}}}\right)\right] + 2\sqrt{\frac{2\left(D_{\mathrm{KL}}(\tilde{\boldsymbol{w}}\|P) + \ln\frac{2m}{\delta}\right)}{m-1}}$$
(58)

$$\leq \widehat{L}_{r,\theta}\left(f_{\boldsymbol{w}}\right) + 2\sqrt{\frac{2\left(D_{\mathrm{KL}}(\tilde{\boldsymbol{w}}\|P) + \ln\frac{2m}{\delta}\right)}{m-1}}$$
(59)

$$\leq \widehat{L}_{r,\theta}\left(f_{\boldsymbol{w}}\right) + 4\sqrt{\frac{D_{\mathrm{KL}}\left(\boldsymbol{w}'\|P\right) + \ln\frac{6m}{\delta}}{m-1}} \tag{60}$$

5.3 Proof of Theorem 4.11

Proof of Theorem 4.11: Since Lemma 4.9 proves that the perturbation caused by random vector \boldsymbol{u} is bounded by a term relative to the variance σ , we can preset the value of σ to make the random perturbation satisfy the condition for Lemma 4.10. Bounding the Kullback-Leibler divergence term by $\|\boldsymbol{w}\|_2^2 / \|\boldsymbol{u}\|_2^2$ in PAC-Bayesian theorem, we can attain the generalization bound based on a specific margin distribution.

The proof involves chiefly two steps. In the first step we bound the maximum value of perturbation of parameters to satisfy the condition that the change of output restricted by hyper-parameters of margin r and θ , using Lemma 4.9. In the second step we prove the final margin generalization bound through Lemma 4.10 with the value of Kullback-Leibler divergence term calculated based on the bound in the first step.

$$\begin{split} |f_{\boldsymbol{w}+\boldsymbol{u}}(\boldsymbol{x}) - f_{\boldsymbol{w}}(\boldsymbol{x})|_{2}^{2} &\leq \mathcal{O}\left(\sum_{i=1}^{d} \frac{d\alpha^{2}c^{2}(r+\theta)^{2}\sigma^{2}}{\mu_{i}^{2}\mu_{i\rightarrow}^{2}}\right) \\ &= (\frac{r-\theta}{8\sqrt{\rho}})^{2} \end{split}$$

We can derive $\sigma = \frac{r-\theta}{8\alpha c d \sqrt{\rho}(r+\theta) \sqrt{\sum_{i=1}^{d} \frac{1}{\mu_i^2 \mu_{i\rightarrow}^2}}}$ from the above inequality. Naturally, we can calculate the

Kullback-Leibler divergence in Lemma 4.9 with the chosen distributions for $P \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$.

$$D_{\mathrm{KL}}(\boldsymbol{w} + \boldsymbol{u} \| P) \le \frac{|\boldsymbol{w}|^2}{2|\boldsymbol{w}|^2 |\eta\eta^\top|^2} = \frac{1}{2\sigma^2}$$
(61)

$$\leq \mathcal{O}\left(\frac{(r+\theta)^2}{(r-\theta)^2}\sum_{i=1}^d \frac{d\rho\alpha^2 c^2}{\mu_i^2 \mu_{i\to}^2}\right)$$
(62)

Put it in Lemma 4.10 and let $\lambda = \theta/r$, with probability at least $1 - \delta$ and for all w such that, we have:

$$L_0(h) \le \widehat{L}_{r,\theta}(h) + \mathcal{O}\left(\sqrt{\frac{\frac{(1+\lambda)^2}{(1-\lambda)^2}\sum_{i=1}^d \frac{d\rho\alpha^2 c^2}{\mu_i^2 \mu_{i\rightarrow}^2} + \ln \frac{dm}{\delta}}{m}}\right).$$
(63)

6 Optimizing margin distribution measure

The generalization theory shows the importance of optimizing the margin distribution ratio λ . The result inspires us to find a margin distribution band $(r - \theta \le \gamma_h(\boldsymbol{x}, y) < r + \theta)$ containing as many training samples as possible to minimize the empirical estimate loss $\hat{L}_{r,\theta}$, but also a ratio $\lambda = \theta/r$ as small as possible to minimize the generalization gap $L_0(h) - \hat{L}_{r,\theta}(h)$. This type of loss function was first proposed by Zhang and Zhou [ZZ19] to optimize the first- and second-order statistics of margin distribution. We formulate a convex margin distribution loss function for DNNs:

Definition 6.1. (Convex margin distribution loss function). For a labeled sample $(x, y) \in D$, we denote its margin by γ_h which is defined as Eq. (1). We define the margin distribution loss for networks (mdNet loss) as:

$$\ell_{r,\theta,\eta}(h(\boldsymbol{x}), y) = \begin{cases} \frac{(r-\theta-\gamma_h)^2}{(r-\theta)^2} & \gamma_h \le r-\theta\\ 0 & r-\theta < \gamma_h \le r+\theta\\ \frac{\eta(r+\theta-\gamma_h)^2}{(r+\theta)^2} & \gamma_h > r+\theta, \end{cases}$$
(64)

where r is the margin mean parameter, θ is the margin variance parameter and η is a parameter to trade off two different kinds of deviation (keeping the balance on both sides of the margin mean). Figure l(c) shows the shape of this convex loss function.

Equation (64) will produce a square loss when the margin satisfies $\gamma_h \leq r - \theta$ or $\gamma_h \geq r + \theta$. Therefore, our margin loss function will force the zero-loss band to contain as many sample points as possible. The ratio of hyper-parameters $\lambda = \theta/r$ can control the capacity measure, which implies our measure is dependent to our specific learning algorithm (loss function with specific hyper-parameters). Since our loss function aims at finding a decision boundary which is determined by the entire margin distribution, instead of the minority samples that have minimum margins, we name our method as <u>margin distribution Networks</u> (mdNet).

7 Experiments

In Subsection 7.1, we introduce the configuration of datasets and models. In Subsection 7.2, We design an ablation experiment to verify the superiority of our method. In Subsection 7.3, we show the correlation between separability of representations and margin ratio via visualization. In Subsection 7.4, we design experiments to confirm that our method can control the capacity of deep nets. In Subsection 7.5, we discuss the influence of the different hyper-parameters on the test accuracy.

7.1 Configuration

Since our method only works on the loss function part of deep models and does not change the architecture of deep neural networks, we can verify the effectiveness of mdNet on the classic CNNs (convolutional neural networks) and image classification benchmark datasets. We consider the following architectures and datasets: a LeNet architecture for MNIST dataset [LeC+98], an AlexNet architecture [KSH12] for CIFAR-10 dataset [Kri09] and a ResNet-18 architecture [He+16] for ImageNet dataset [Rus+15]. From the literature, these datasets come pre-divided into training and testing sets, therefore in our experiments, we use them in their original format. The loss functions used for comparison in the experiments are as follows: cross-entropy loss (abbr., xent), hinge loss and soft hinge loss. Hinge loss [CV95] and soft hinge loss [Liu+16] are loss functions specially proposed to optimize the minimum margin, both of them are inspired the traditional margin theory.

As for details about the architecture, we remove the weight decay [KH92], dropout [Sri+14] and batch normalization (BN) [IS15] from all the models, because the batch normalization operation and weight decay will shift the data distribution. The notable dropout technique, in which some neurons are dropped from the DNNs in each iteration, can also be viewed as an ensemble method composed of different neural networks, with different dropped neurons [BS13]. It is hard to analyze the influence of the ensemble structure on the margin distribution, so we remove this technique in these architectures in the experiments except to understand the contribution of the components to the whole models in the ablation study.

For *special hyper-parameters*, including the expected margin parameter and margin variance parameter for mdNet loss model, and margin parameter for hinge loss model, we perform hyper-parameter search. We hold out 5000 samples of the training set as a validation set, and use the remaining samples to train models with different special hyper-parameters values on all datasets. As for the common hyper-parameters, such as learning rate and momentum, we set them as the default commonly used values in PyTorch [Pas+19] for all the models. We chose batch stochastic gradient descent as the optimizer. We run all the experiments on four K80 GPU machines. As for the influence of the different hyper-parameters on the test accuracy, we discuss it empirically in Subsection 7.5.

7.2 Ablation study

Since the optimization algorithm proposed in this paper only focuses on the improvement of loss function, we design an ablation experiment to study the performance of our proposed mdnet method and the traditional benchmark loss functions under different regularizations in Table 1. The mdNet loss outperforms the others consistently across different situations, no matter whether dropout, batch normalization or the entire dataset are used or not. The experiments are evaluated on three MNIST (LeNet), CIFAR-10 (AlexNet) and ImageNet (ResNet-18) datasets. Specifically, when the amount of training samples is small, the advantage of mdNet loss is significant. Moreover, the mdNet loss function can cooperate with both batch normalization and dropout, achieving the best performance in Table 1, which is highlighted in bold red text. Unlike dropout and batch normalization which lack solid theoretical grounds, the mdNet loss function is inspired by the margin distribution bound in Theorem 4.11, which guides us to find a suitable margin ratio to restrict the model capacity and alleviate the overfitting problem efficiently.

7.3 Feature visualization

In this experiment, we use t-SNE method to visualize the learned representations on the last hidden layer. Figure 4(a), 4(b) and 4(c) plot the 2D t-SNE [MH08] embedding image on limited datasets, including MNIST (LeNet), CIFAR-10 (AlexNet) and 10-class ImageNet (ResNet-18). Consistently, we can find that the result of mdNet loss model is better than all the others, the distribution of samples which have the same label are more compact. To quantify the degree of separability of data distribution, we perform a variance decomposition on the data in the embedding space. By comparing the ratio of inter-class variance S_E to intra-class variance S_A in Figure 4(d), 4(e) and 4(f), we see that the mdNet loss always attain the most separable distribution among these four loss functions. Moreover, the visualization result is consistent with the margin distribution ratio $1/\lambda$ of these four models, which means that optimizing the margin distribution (searching an appropriate margin ratio λ) is helpful to attain a good learned representation space. This representation features space can further alleviate the overfitting problem of deep learning, we verify empirically that a network trained with mdNet loss shows stronger clustering. Specially, Figure 4(d), 4(e) and 4(f) show the relationship between the margin ratio and test error. Moreover, Figure 5 plots the test error of mdNet and the margin ratio across the different epochs. We can see that more compact margin distribution gets better prediction performance across different models and epochs. This exhibits that optimizing margin distribution can indeed improve the learning ability of deep nets.

7.4 Controlling capacity

The first two experiments have demonstrated that our mdNet can outperform other classical loss functions and our method can learn a more separable feature representation, as the corresponding margin ratio is also smaller. However, it leads to the last question:

Can smaller margin ratio reduce the capacity of models and accelerate the convergence of generalization gap?

Let's go back to the theoretical result obtained by Theorem 4.11. The generalization gap of the model is bounded by $\Lambda_{\lambda, w}/\sqrt{m}$, where the margin distribution measure $\Lambda_{\lambda, w} \propto \frac{1+\lambda}{1-\lambda}$ determines the worse case of generalization gap when the number of samples are equal:

Generalization Gap
$$\leq \mathcal{O}(\frac{\Lambda_{\lambda, \boldsymbol{w}}}{\sqrt{m}}).$$
 (65)

Therefore, we design experiments to compare the empirical value of generalization gap with the increase of training samples under different loss functions. In Figure 6, it shows that the red dotted line representing the convergence curve of our method always convergences faster than the other lines under different datasets and models. It also demonstrates that our method can effectively control the capacity of the model by optimizing the margin distribution ratio, so that the trained model has better generalization performance.

Given a fixed number of samples m, we find that the worst case of generalization gap is proportional to the model capacity. When m is large enough, the scale factor $1/\sqrt{m}$ will be close to 0, and the difference of sample complexity is not significant. The advantage of optimizing margin distribution is relatively significant when \sqrt{m} is relatively small. Therefore, in the right of Figure 6 (the ImageNet experiment), we specially truncate the most significant result of convergence rate (form 0.1% to 5% of training set), which shows that optimizing margin distribution can control the capacity of the model even on such a complex dataset.

7.5 Influence of the hyper-parameters

Figure 7 plots the 3D surface figure for the test accuracy on the MNIST, CIFAR-10 and ImageNet datasets varying with two hyper-parameter r and θ . It shows that the ratio $1/\lambda = r/\theta$ (the lower surface with rainbow colors) increases with r increasing and θ decreasing. As for the test accuracy (the upper surface with warm-cool colors), we find that its trend is consistent with the ratio $1/\lambda$. Therefore, the influence of the hyper-parameters demonstrates that our theoretical result. Within a certain range, getting a smaller ratio λ through specific optimization (the margin distribution loss function) will effectively reduce the size of the hypothesis set for deep nets (returned by the specific algorithm), so as to improve the generalization ability of the learned model. In other words, the test accuracy changes consistently with the ratio of hyper-parameters (an estimation of the ratio of margin distribution). The parameter η to trade off two different kinds of deviation (keep the balance on both sides of the margin mean) is always fixed to 0.1 in practice.

8 Conclusion

This paper proves generalization bound for deep neural networks by considering the margin distribution at the last layer instead of the minimum margin. The theoretical result inspires us to utilize a margin distribution loss function to improve the generalization performance of neural networks. Experimental results show that our method can effectively control the model capacity by optimizing the margin distribution measure, so that the trained model learns more separable representations and has better generalization performance. In future work, we will explore the effectiveness of regularization methods from a *margin theory* perspective.

Acknowledgement

This research was supported by the National Science Foundation of China (61921006) and the Collaborative Innovation Center of Novel Software Technology and Industrialization. The authors would like to thank the anonymous reviewers for constructive suggestions, as well as Yi-Xiao He and Yi-He Chen for helpful discussions.

References

[Aro+18]	Sanjeev Arora, Rong Ge, Behnam Neyshabur, and Yi Zhang. "Stronger Generalization Bounds for Deep Nets via a Compression Approach". In: <i>Proceedings of the 35th Interna-</i> <i>tional Conference on Machine Learning</i> . 2018, pp. 254–263 (Cited on pages 1–6, 8).
[AW19]	Aharon Azulay and Yair Weiss. "Why do deep convolutional networks generalize so poorly to small image transformations?" In: <i>Journal of Machine Learning Research</i> 20.184 (2019), pp. 1–25 (Cited on page 1).
[BS13]	Pierre Baldi and Peter J Sadowski. "Understanding dropout". In: <i>Advances in Neural In-</i> <i>formation Processing Systems</i> 27. 2013, pp. 2814–2822 (Cited on page 15).
[BMM98]	Peter L Bartlett, Vitaly Maiorov, and Ron Meir. "Almost linear VC-dimension bounds for piecewise polynomial networks". In: <i>Neural Computation</i> 10.8 (1998), pp. 2159–2173 (Cited on page 5).
[BM02]	Peter L Bartlett and Shahar Mendelson. "Rademacher and Gaussian complexities: Risk bounds and structural results". In: <i>Journal of Machine Learning Research</i> 3 (2002), pp. 463-482.
[BFT17]	Peter L. Bartlett, Dylan J. Foster, and Matus J. Telgarsky. "Spectrally-normalized margin bounds for neural networks". In: <i>Advances in Neural Information Processing Systems 31</i> . 2017, pp. 6241–6250 (Cited on pages 1–3, 5).
[BCF13]	Carlos J. Becker, C. Mario Christoudias, and Pascal Fua. "Non-Linear Domain Adapta- tion with Boosting". In: <i>Advances in Neural Information Processing Systems</i> 26. 2013, pp. 485–493 (Cited on page 3).
[BE02]	Olivier Bousquet and André Elisseeff. "Stability and generalization". In: <i>Journal of Machine Learning Research</i> 2.Mar (2002), pp. 499–526.
[Bre99]	Leo Breiman. "Prediction Games and Arcing Algorithms". In: <i>Neural Computation</i> 11.7 (1999), pp. 1493–1517 (Cited on page 2).
[Cha+15]	Tsung-Han Chan, Kui Jia, Shenghua Gao, Jiwen Lu, Zinan Zeng, and Yi Ma. "PCANet: A simple deep learning baseline for image classification?" In: <i>IEEE Transactions on Image</i>

Processing 24.12 (2015), pp. 5017–5032 (Cited on page 3).

- [Che+19] Hao Chen, Zhanfeng Mo, Zhouwang Yang, and Xiao Wang. "Theoretical Investigation of Generalization Bound for Residual Networks". In: *Proceedings of the 28th International Joint Conference on Artificial Intelligence*. 2019, pp. 2081–2087.
- [CJY18] Yuansi Chen, Chi Jin, and Bin Yu. "Stability and convergence trade-off of iterative optimization algorithms". In: *CoRR* abs/1804.01619 (2018) (Cited on page 3).
- [CV95] Corinna Cortes and Vladimir Vapnik. "Support-Vector Networks". In: *Machine Learning* 20.3 (1995), pp. 273–297 (Cited on page 15).
- [Din+17] Laurent Dinh, Razvan Pascanu, Samy Bengio, and Yoshua Bengio. "Sharp Minima Can Generalize For Deep Nets". In: *Proceedings of the 34th International Conference on Machine Learning*. Vol. 70. 2017, pp. 1019–1028 (Cited on page 3).
- [Dru+97] Harris Drucker, Christopher J. C. Burges, Linda Kaufman, Alex J. Smola, and Vladimir Vapnik. "Support Vector Regression Machines". In: Advances in Neural Information Processing Systems 10. 1997, pp. 155–161.
- [Dub+21] Abhimanyu Dubey, Vignesh Ramanathan, Alex Pentland, and Dhruv Mahajan. "Adaptive Methods for Real-World Domain Generalization". In: *IEEE Conference on Computer Vi*sion and Pattern Recognition. 2021, pp. 14340–14349 (Cited on page 3).
- [Els+18] Gamaleldin Elsayed, Dilip Krishnan, Hossein Mobahi, Kevin Regan, and Samy Bengio. "Large Margin Deep Networks for Classification". In: Advances in Neural Information Processing Systems 31. 2018, pp. 850–860.
- [GZ13] Wei Gao and Zhi-Hua Zhou. "On the Doubt About Margin Explanation of Boosting". In: *Artificial Intelligence* 203 (2013), pp. 1–18 (Cited on page 2).
- [Ghi+15] Muhammad Ghifary, W. Bastiaan Kleijn, Mengjie Zhang, and David Balduzzi. "Domain Generalization for Object Recognition with Multi-task Autoencoders". In: *IEEE International Conference on Computer Vision*. 2015, pp. 2551–2559 (Cited on page 3).
- [GS01] Stanton Glantz and Bryan Slinker. *Primer of Applied Regression & Analysis of Variance, ed.* McGraw-Hill, Inc., New York, 2001 (Cited on page 8).
- [GSS15] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. "Explaining and Harnessing Adversarial Examples". In: *3rd International Conference on Learning Representations*. 2015 (Cited on page 1).
- [Gor22] Marco Gori. "Ten Questions for a Theory of Vision". In: *Frontiers in Computer Science* 3 (2022), p. 701248. DOI: 10.3389/fcomp.2021.701248 (Cited on page 1).
- [Grø+19] Allan Grønlund, Lior Kamma, Kasper Green Larsen, Alexander Mathiasen, and Jelani Nelson. "Margin-Based Generalization Lower Bounds for Boosted Classifiers". In: Advances in Neural Information Processing Systems 32. 2019, pp. 11940–11949 (Cited on page 2).
- [Gun+18a] Suriya Gunasekar, Jason D. Lee, Daniel Soudry, and Nathan Srebro. "Characterizing Implicit Bias in Terms of Optimization Geometry". In: *Proceedings of the 35th International Conference on Machine Learning*. Vol. 80. 2018, pp. 1827–1836 (Cited on page 3).
- [Gun+18b] Suriya Gunasekar, Jason D. Lee, Daniel Soudry, and Nati Srebro. "Implicit Bias of Gradient Descent on Linear Convolutional Networks". In: Advances in Neural Information Processing Systems 31. 2018, pp. 9482–9491 (Cited on page 3).
- [HRS16] Moritz Hardt, Ben Recht, and Yoram Singer. "Train faster, generalize better: Stability of stochastic gradient descent". In: *Proceedings of the 33nd International Conference on Machine Learning*. Vol. 48. 2016, pp. 1225–1234 (Cited on page 3).

- [HLM17] Nick Harvey, Christopher Liaw, and Abbas Mehrabian. "Nearly-tight VC-dimension bounds for piecewise linear neural networks". In: *Proceedings of the 30th Annual Conference on Learning Theory*. 2017, pp. 1064–1068 (Cited on page 5).
- [He+16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. "Deep Residual Learning for Image Recognition". In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition. 2016, pp. 770–778 (Cited on page 15).
- [HSS12] Geoffrey Hinton, Nitish Srivastava, and Kevin Swersky. "Neural networks for machine learning lecture 6a overview of mini-batch gradient descent". In: (2012).
- [Hua+18] Furong Huang, Jordan Ash, John Langford, and Robert Schapire. "Learning Deep ResNet Blocks Sequentially using Boosting Theory". In: *Proceedings of the 35th International Conference on Machine Learning*. 2018, pp. 2058–2067.
- [IS15] Sergey Ioffe and Christian Szegedy. "Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift". In: *Proceedings of the 32nd International Conference on Machine Learning*. 2015, pp. 448–456 (Cited on pages 1, 15).
- [JT19] Ziwei Ji and Matus Telgarsky. "Gradient descent aligns the layers of deep linear networks". In: *7th International Conference on Learning Representations*. 2019 (Cited on page 3).
- [Jia+19] Yiding Jiang, Dilip Krishnan, Hossein Mobahi, and Samy Bengio. "Predicting the Generalization Gap in Deep Networks with Margin Distributions". In: *7th International Conference on Learning Representations*. 2019 (Cited on pages 1, 3, 8).
- [Kes+17] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. "On Large-Batch Training for Deep Learning: Generalization Gap and Sharp Minima". In: 5th International Conference on Learning Representations. 2017 (Cited on page 3).
- [KB15] Diederik P. Kingma and Jimmy Ba. "Adam: A Method for Stochastic Optimization". In: *3rd International Conference on Learning Representations*. 2015.
- [KTP17] Piotr Koniusz, Yusuf Tas, and Fatih Porikli. "Domain Adaptation by Mixture of Alignments of Second-or Higher-Order Scatter Tensors". In: *IEEE Conference on Computer Vision and Pattern Recognition*. 2017, pp. 7139–7148 (Cited on page 3).
- [Kri09] Alex Krizhevsky. *Learning Multiple Layers of Features from Tiny Images*. Tech. rep. 2009 (Cited on page 15).
- [KSH12] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. "ImageNet Classification with Deep Convolutional Neural Networks". In: Advances in Neural Information Processing Systems 26. 2012, pp. 1097–1105 (Cited on page 15).
- [KH92] Anders Krogh and John A Hertz. "A simple weight decay can improve generalization". In: Advances in Neural Information Processing Systems 5. 1992, pp. 950–957 (Cited on pages 1, 15).
- [LBH15] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. "Deep learning". In: *Nature* 521.7553 (2015), p. 436 (Cited on page 1).
- [LeC+98] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. "Gradient-based learning applied to document recognition". In: *Proceedings of the IEEE* 86.11 (1998), pp. 2278–2324 (Cited on page 15).
- [LMZ18] Yuanzhi Li, Tengyu Ma, and Hongyang Zhang. "Algorithmic Regularization in Overparameterized Matrix Sensing and Neural Networks with Quadratic Activations". In: *Conference On Learning Theory*. Vol. 75. 2018, pp. 2–47 (Cited on page 3).

- [Liu+16] Weiyang Liu, Yandong Wen, Zhiding Yu, and Meng Yang. "Large-Margin Softmax Loss for Convolutional Neural Networks". In: *Proceedings of the 33rd International Conference* on Machine Learning. 2016, pp. 507–516 (Cited on pages 3, 15).
- [MH08] Laurens van der Maaten and Geoffrey Hinton. "Visualizing data using t-SNE". In: *Journal* of Machine Learning Research 9.11 (2008), pp. 2579–2605 (Cited on page 15).
- [Mac+14] Lester Mackey, Michael I Jordan, Richard Y Chen, Brendan Farrell, Joel A Tropp, et al. "Matrix concentration inequalities via the method of exchangeable pairs". In: *The Annals* of *Probability* 42.3 (2014), pp. 906–945.
- [MMR09] Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. "Domain Adaptation: Learning Bounds and Algorithms". In: *The 22nd Conference on Learning Theory*. 2009 (Cited on page 3).
- [MS14] Yishay Mansour and Mariano Schain. "Robust domain adaptation". In: *Annals of Mathematics and Artificial Intelligence* 71.4 (2014), pp. 365–380 (Cited on page 3).
- [Mat+22] Alex Matskevych, Adrian Wolny, Constantin Pape, and Anna Kreshuk. "From Shallow to Deep: Exploiting Feature-Based Classifiers for Domain Adaptation in Semantic Segmentation". In: *Frontiers in Computer Science* 4 (2022), p. 805166. DOI: 10.3389/fcomp. 2022.805166 (Cited on page 3).
- [McA03] David McAllester. "Simplified PAC-Bayesian Margin Bounds". In: *Learning Theory and Kernel Machines*. 2003, pp. 203–215.
- [MRT18a] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning*. MIT press, 2018 (Cited on page 4).
- [MRT18b] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning*. MIT press, 2018.
- [Mou+18] Wenlong Mou, Liwei Wang, Xiyu Zhai, and Kai Zheng. "Generalization Bounds of SGLD for Non-convex Learning: Two Theoretical Viewpoints". In: *Conference On Learning The*ory. Vol. 75. 2018, pp. 605–638 (Cited on page 3).
- [MBS13] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. "Domain Generalization via Invariant Feature Representation". In: *Proceedings of the 30th International Conference on Machine Learning*. Vol. 28. JMLR.org, 2013, pp. 10–18.
- [Ney+17] Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nati Srebro. "Exploring Generalization in Deep Learning". In: Advances in Neural Information Processing Systems 30, 2017, pp. 5947–5956 (Cited on page 3).
- [NBS18] Behnam Neyshabur, Srinadh Bhojanapalli, and Nathan Srebro. "A PAC-Bayesian Approach to Spectrally-Normalized Margin Bounds for Neural Networks". In: *6th International Conference on Learning Representations*. 2018 (Cited on pages 1–3, 5, 7, 13).
- [NTS15] Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. "Norm-based capacity control in neural networks". In: *Proceedings of the 28th Annual Conference on Learning Theory*. 2015, pp. 1376–1401 (Cited on page 1).
- [Pan+09] Sinno Jialin Pan, Ivor W. Tsang, James T. Kwok, and Qiang Yang. "Domain Adaptation via Transfer Component Analysis". In: *Proceedings of the 21st International Joint Conference* on Artificial Intelligence. 2009, pp. 1187–1192 (Cited on page 3).
- [Pap+17] Nicolas Papernot, Patrick D. McDaniel, Ian J. Goodfellow, Somesh Jha, Z. Berkay Celik, and Ananthram Swami. "Practical Black-Box Attacks against Machine Learning". In: *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*. 2017, pp. 506–519 (Cited on page 1).

- [Pas+19] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. "PyTorch: An Imperative Style, High-Performance Deep Learning Library". In: Advances in Neural Information Processing Systems 32. 2019, pp. 8024–8035 (Cited on page 15).
- [RS06] Lev Reyzin and Robert E. Schapire. "How Boosting the Margin Can Also Boost Classifier Complexity". In: *Proceedings of the 23rd International Conference on Machine Learning*. 2006, pp. 753–760 (Cited on page 2).
- [RSF19] Artem Rozantsev, Mathieu Salzmann, and Pascal Fua. "Beyond Sharing Weights for Deep Domain Adaptation". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41.4 (2019), pp. 801–814 (Cited on page 3).
- [Rud16] Sebastian Ruder. "An overview of gradient descent optimization algorithms". In: *CoRR* abs/1609.04747 (2016).
- [Rus+15] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. "ImageNet Large Scale Visual Recognition Challenge". In: *International Journal of Computer Vision* 115.3 (2015), pp. 211–252 (Cited on page 15).
- [Sch+97] Robert E. Schapire, Yoav Freund, Peter Barlett, and Wee Sun Lee. "Boosting the margin: A new explanation for the effectiveness of voting methods". In: *Proceedings of the 14th International Conference on Machine Learning*. 1997, pp. 322–330 (Cited on page 2).
- [SKP15] Florian Schroff, Dmitry Kalenichenko, and James Philbin. "Facenet: A unified embedding for face recognition and clustering". In: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. 2015, pp. 815–823 (Cited on page 3).
- [SZ15] Karen Simonyan and Andrew Zisserman. "Very Deep Convolutional Networks for Large-Scale Image Recognition". In: 3rd International Conference on Learning Representations. 2015.
- [Sou+18] Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. "The Implicit Bias of Gradient Descent on Separable Data". In: *Journal of Machine Learning Research* 19.70 (2018), pp. 1–57 (Cited on page 3).
- [Sri+14] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. "Dropout: a simple way to prevent neural networks from overfitting". In: *Journal of Machine Learning Research* 15.1 (2014), pp. 1929–1958 (Cited on pages 1, 15).
- [Sun+14] Yi Sun, Yuheng Chen, Xiaogang Wang, and Xiaoou Tang. "Deep learning face representation by joint identification-verification". In: *Advances in Neural Information Processing Systems* 28. 2014, pp. 1988–1996.
- [Tan+20] Zhi-Hao Tan, Peng Tan, Yuan Jiang, and Zhi-Hua Zhou. "Multi-label optimal margin distribution machine". In: *Machine Learning* 109.3 (2020), pp. 623–642 (Cited on page 2).
- [Tro12] Joel A Tropp. "User-friendly tail bounds for sums of random matrices". In: *Foundations* of Computational Mathematics 12.4 (2012), pp. 389–434.
- [Vap98] Vladimir Vapnik. *Statistical Learning Theory*. Wiley press, 1998.

[Wan+21]	Jindong Wang, Cuiling Lan, Chang Liu, Yidong Ouyang, and Tao Qin. "Generalizing to Unseen Domains: A Survey on Domain Generalization". In: <i>Proceedings of the 30th International Joint Conference on Artificial Intelligence</i> . 2021, pp. 4627–4635 (Cited on page 3).
[Wan+11]	Liwei Wang, Masashi Sugiyama, Zhaoxiang Jing, Cheng Yang, Zhi-Hua Zhou, and Jufu Feng. A Refined Margin Analysis for Boosting Algorithms via Equilibrium Margin. 2011.
[Wei+18]	Colin Wei, Jason D. Lee, Qiang Liu, and Tengyu Ma. "On the Margin Theory of Feedforward Neural Networks". In: <i>CoRR</i> abs/1810.05369 (2018) (Cited on page 3).
[Wu+21]	Wei Wu, Xiaoyuan Jing, Wencai Du, and Guoliang Chen. "Learning dynamics of gradient descent optimization in deep neural networks". In: <i>Science China Information Science</i> 64.5 (2021), p. 15102. DOI: 10.1007/s11432-020-3163-0 (Cited on page 3).
[ZZY12]	Chao Zhang, Lei Zhang, and Jieping Ye. "Generalization bounds for domain adaptation". In: <i>Advances in Neural Information Processing Systems</i> 25 4 (2012), p. 3320 (Cited on page 3).
[Zha+21]	Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. "Under- standing deep learning (still) requires rethinking generalization". In: <i>Communications of</i> <i>the ACM</i> 64.3 (2021), pp. 107–115 (Cited on page 1).
[ZZJ20]	Teng Zhang, Peng Zhao, and Hai Jin. "Optimal Margin Distribution Learning in Dynamic Environments". In: <i>Proceedings of the 34th AAAI Conference on Artificial Intelligence</i> . 2020, pp. 6821–6828 (Cited on page 2).
[ZZ17]	Teng Zhang and Zhi-Hua Zhou. "Multi-Class Optimal Margin Distribution Machine". In: <i>Proceedings of the 34th International Conference on Machine Learning</i> . 2017, pp. 4063–4071 (Cited on page 2).
[ZZ18a]	Teng Zhang and Zhi-Hua Zhou. "Optimal Margin Distribution Clustering". In: <i>Proceedings of the 32nd AAAI Conference on Artificial Intelligence</i> . 2018, pp. 4474–4481 (Cited on page 2).
[ZZ18b]	Teng Zhang and Zhi-Hua Zhou. "Optimal Margin Distribution Clustering". In: <i>Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence</i> . 2018, pp. 4474–4481.
[ZZ18c]	Teng Zhang and Zhi-Hua Zhou. "Semi-Supervised Optimal Margin Distribution Machines". In: <i>Proceedings of the 27th International Joint Conference on Artificial Intelligence</i> . 2018, pp. 3104–3110 (Cited on page 2).
[ZZ19]	Teng Zhang and Zhi-Hua Zhou. "Optimal margin distribution machine". In: <i>IEEE Transac-</i> <i>tions on Knowledge and Data Engineering</i> 32.6 (2019), pp. 1143–1156 (Cited on pages 2, 14).
[Zho14]	Zhi-Hua Zhou. "Large Margin Distribution Learning". In: <i>Artificial Neural Networks in Pattern Recognition</i> . 2014, pp. 1–11 (Cited on page 2).
[Zho21]	Zhi-Hua Zhou. "Why over-parameterization of deep neural networks does not overfit?" In: <i>Science China Information Sciences</i> 64.1 (2021), pp. 1–3 (Cited on page 3).
[Zhu+19]	Zhanxing Zhu, Jingfeng Wu, Bing Yu, Lei Wu, and Jinwen Ma. "The Anisotropic Noise in Stochastic Gradient Descent: Its Behavior of Escaping from Sharp Minima and Regulariza- tion Effects". In: <i>Proceedings of the 36th International Conference on Machine Learning</i> . Vol. 97. 2019, pp. 7654–7663 (Cited on page 3).

Table 1: Test accuracy of LeNet on MNIST, AlexNet on CIFAR-10 and ResNet-18 on ImageNet datasets
with different regularization methods and different fractions of training set. The best accuracy on each
training dataset is highlighted in bold red type. The bold black text indicates the better accuracy between
the two losses with the same regularization.

		MNSIT		
	xent	hinge	soft hinge	mdNet
Accuracy(%)	Batch Normalization			
100%_Dropout	99.095 ± 0.083	98.593 ± 0.164	99.148 ± 0.039	99.161 ± 0.073
100%_Non_Dropout	98.384 ± 0.072	97.571 ± 0.178	98.475 ± 0.064	98.837 ± 0.091
5%_Dropout	97.001 ± 0.131	96.527 ± 0.219	97.112 ± 0.092	97.268 ± 0.113
5%_Non_Dropout	83.364 ± 0.452	83.292 ± 0.721	83.749 ± 0.273	84.483 ± 0.348
	Non Batch Normalization			
100%_Dropout	98.228 ± 0.079	98.029 ± 0.184	98.271 ± 0.055	98.342 ± 0.069
100%_Non_Dropout	91.728 ± 0.117	90.237 ± 0.318	91.029 ± 0.098	92.274 ± 0.121
5%_Dropout	77.842 ± 0.489	76.938 ± 0.827	77.727 ± 0.411	78.173 ± 0.619
5%_Non_Dropout	58.023 ± 0.951	57.822 ± 1.280	59.384 ± 0.827	61.379 ± 0.588
		CIFAR-10		
	xent	hinge	soft hinge	mdNet
Accuracy(%)	Batch Normalization			
100%_Dropout	85.782 ± 0.198	84.234 ± 0.748	86.744 ± 0.294	87.644 ± 0.151
100%_Non_Dropout	81.491 ± 0.143	80.938 ± 0.812	86.032 ± 0.298	86.233 ± 0.244
5%_Dropout	61.955 ± 1.945	58.363 ± 2.450	59.441 ± 1.316	67.636 ± 1.633
5%_Non_Dropout	57.753 ± 2.228	54.289 ± 3.482	56.839 ± 2.318	64.173 ± 1.982
	Non Batch Normalization			
100%_Dropout	83.517 ± 0.322	82.153 ± 1.236	81.961 ± 0.293	84.643 ± 0.255
100%_Non_Dropout	72.223 ± 1.284	69.379 ± 2.907	75.267 ± 1.027	76.793 ± 1.279
5%_Dropout	50.747 ± 3.735	42.739 ± 6.763	52.847 ± 1.823	58.739 ± 1.348
5%_Non_Dropout	36.293 ± 4.872	30.984 ± 7.736	43.265 ± 4.263	47.056 ± 3.927
		ImageNet		
A	xent	hinge	soft hinge	mdNet
Accuracy(%)		Batch Nor	malization	
100%_Dropout	70.238 ± 1.221	69.782 ± 1.933	70.284 ± 1.022	70.758 ± 1.014
100%_Non_Dropout	68.484 ± 1.265	67.918 ± 2.166	68.83 ± 1.151	69.447 ± 1.124
5%Dropout	60.176 ± 2.045	57.475 ± 2.023	61.379 ± 1.053	65.080 ± 2.373
5‰_Non_Dropout	59.574 ± 2.747	56.621 ± 2.253	60.068 ± 1.773	63.529 ± 2.012
		Non Batch N	lormalization	
100%_Dropout	66.924 ± 1.552	67.387 ± 1.764	67.462 ± 1.017	68.655 ± 1.732
100%_Non_Dropout	64.481 ± 2.183	61.820 ± 2.947	64.334 ± 2.367	65.838 ± 2.481
5‰_Dropout	54.961 ± 3.382	52.543 ± 3.722	55.757 ± 2.357	58.774 ± 3.841
5‰_Non_Dropout	47.374 ± 3.265	45.741 ± 5.349	48.798 ± 3.392	53.727 ± 4.235



Figure 4: The quality of feature representations generated by different models on the MNIST, CIFAR-10 and ImageNet datasets.



Figure 5: Test error and margin ratio across epochs on mdNet models for MNIST, CIFAR-10 and ImageNet datasets.



Figure 6: Compare the convergence rate of generalization gap with the increase of training samples under different loss functions on MNIST, CIFAR-10 and ImageNet datasets.



Figure 7: The test accuracy varying with two hyper-parameter r and θ on MNIST, CIFAR-10 and ImageNet datasets. The logarithm of ratio $\ln(1/\lambda) = \ln(r/\theta)$ is the lower surface with rainbow colors and the test accuracy is the upper surface with warm-cool colors. The test accuracy is rescaled to be within the same rage as the ratio.