

Recent advances in deep learning theory

Fengxiang He *

Dacheng Tao *

Abstract

Deep learning is usually described as an experiment-driven field under continuous criticisms of lacking theoretical foundations. This problem has been partially fixed by a large volume of literature which has so far not been well organized. This paper reviews and organizes the recent advances in deep learning theory. The literature is categorized in six groups: (1) complexity and capacity-based approaches for analyzing the generalizability of deep learning; (2) stochastic differential equations and their dynamic systems for modelling stochastic gradient descent and its variants, which characterize the optimization and generalization of deep learning, partially inspired by Bayesian inference; (3) the geometrical structures of the loss landscape that drives the trajectories of the dynamic systems; (4) the roles of over-parameterization of deep neural networks from both positive and negative perspectives; (5) theoretical foundations of several special structures in network architectures; and (6) the increasingly intensive concerns in ethics and security and their relationships with generalizability.

Keywords: neural networks, deep learning, deep learning theory.

1 Introduction

Deep learning can be broadly defined as a family of algorithms employing *artificial neural networks* to discover knowledge from *experience* for predictions or decision-making [138]. Canonical forms of the experience can be human-annotated electronic records as a dataset or the interactions between the learner or the electronic environment depending on the scenarios [169]. A normal artificial neural network in deep learning usually connects a sequence of *weight matrix* followed by a *nonlinear activation function* as a network, which is typical of a considerably large parameter size.

The term of *deep learning* was introduced to machine learning by Dechter [62] and then to brain-inspired algorithms by Aizenberg *et al.* [5], while several major concepts wherein can be dated back to early 1940s. The research of deep learning saw two rises and then falls in 1940s-1960s [162, 111, 199] and 1980s-1990s [201]. The third and current wave starts in 2006 [24, 114, 196] and lasts till now. The recent wave has substantially reshaped many real-world application domains, including computer vision [110], natural language processing [63, 184], speech processing [64], 3D point cloud processing [98], data mining [232], recommender system [247], autonomous vehicles [152, 215], medical diagnosis [135, 209], and drug discovery [43].

To date, however, the development of deep learning heavily relies on experiments but without solid theoretical foundations. Many facets of the mechanism of deep learning still remain unknown. We are continuously surprised by finding heuristic approaches can achieve excellent performance across extensive areas, though sometimes also considerably instable. Meanwhile, the suggested intuitions are usually left un-verified or

*The authors are with UBTECH Sydney AI Centre, School of Computer Science, Faculty of Engineering, the University of Sydney, Darlington NSW 2008, Australia. Email: fengxiang.he@sydney.edu.au and dacheng.tao@sydney.edu.au.

even un-checked. Such practice is tolerated and has become common in deep learning research. This black-box nature introduces unknown risks to deep learning applications. This unawareness considerably undermines our capabilities in identifying, managing, and preventing the algorithm-led disasters, and further severely hurts the confidence in applying the recent progress to many industrial sectors, especially those in security-critical areas, such as autonomous vehicles, medical diagnosis, and drug discovery. This also backlashes on the future development of innovative deep learning algorithms designing.

A major part of potential theoretical foundations is the *generalizability*, which refers to the capability of models well trained on the training data by deep learning algorithms predicting on the unseen data [224, 169]. Since the training data can not cover every future circumstance, good generalizability secures the learned model to be able to handle unseen events. This is particularly important where long-tail events constantly appear and are possible to cause fatal disasters.

Statistical learning theory has established theories for the generalizability relying on the hypothesis complexity [224, 169]. Could these tools solve the problem in deep learning theory? The answer is no in general. Conventional tools usually develops generalization bounds based on hypothesis complexities, such as VC-dimension [28, 223], Rademacher complexity [130, 129, 21], and covering number [73, 104]. These complexities, in classic results, heavily rely on the model sizes. This introduces the *Occam’s razor* principle:

Plurality should not be posited without necessity;

i.e., one needs to find a model sufficiently small to prevent overfitting as long as the model can fit the training sample. However, deep learning models usually have extremely large model sizes, which sometimes makes the generalization bound even larger than the potential maximal value of the loss function. Moreover, the Occam’s razor principle suggests a positive correlation between the generalizability and the model size, which is no longer observed in deep learning. In contrast, deeper and wider networks often have advantaged performance [38]. The irreconciliation between the excellent generalizability of deep learning and its extreme over-parameterization is like a “cloud” to conventional complexity-based learning theory.

1.1 A brief overview of the status quo

Early works have attempted to establish the theoretical foundations of deep learning [172, 90, 22, 20, 23, 158, 11] but largely stagnated with the dip of the wider deep learning research.

The recent story started from the work by Zhang *et al.* [244] in 2017. The authors conducted systematic experiments to explore the generalization ability of deep neural networks. They show that neural networks can almost perfectly fit the training data even when the training labels are random. This paper attracts the community of learning theory to the important topic that how to theoretically interpret the success of deep neural networks. Kawaguchi *et al.* [122] discuss many open problems regarding the excellent generalization ability of deep neural networks despite the large capacity, complexity, possible algorithmic instability, non-robustness, and sharp minima. The authors also provide some insights to solve the problems. Since then, the importance of deep learning theory has been widely recognized. A large volume of literature has emerged to establish theoretical foundations of deep learning. In this paper, we review the related literature and organize them into following six categories.

Complexity and capacity-based method for analyzing the generalizability of deep learning. Conventional statistical learning theory has established the a series of upper bounds on the generalization error (generalization bounds) based on the complexity of hypothesis space, such as VC-dimension [28, 223], Rademacher complexity [130, 129, 21], and covering number [73, 104]. Usually, these generalization bounds explicitly rely on the model size. They suggest that controlling the model size can help models

generalize better. However, the colossal model sizes of deep learning models also make the generalization bounds vacuous. It is thus much desired if we can develop size-independent hypothesis complexity measure and generalization bounds. A promising idea is to characterize the complexity of the “effective” hypothesis space which can be learned in deep learning. The effective hypothesis space can be significantly smaller than the whole hypothesis space. We thereby can expect to obtain a much smaller generalization guarantee.

Stochastic differential equations (SDE) for modelling stochastic gradient descent (SGD) and its variants, which dominate the optimization algorithms in deep learning. The dynamic systems of these SDEs determine the trajectory of the weights in training neural networks and their steady distributions represent the learned networks. Through the SDEs and their dynamics, many works establish guarantees for the optimization and generalization of deep learning. The “effective” hypothesis space is exactly the hypothesis space “that can be found by SGD”. Thus, studying the generalizability of deep learning via SGD would be straightforward. Moreover, this family of methods are partially inspired by Bayesian inference. This echoes the previous story of variational inference which solves Bayesian inference in an optimization manner in order to address the scaling issue. This interplay between stochastic gradient methods and Bayesian inference would help advance both areas.

The geometrical structures of the highly complex empirical risk landscape that drive the trajectories of the dynamic systems. The geometry of the loss landscape plays major roles in driving the trajectories of the SDEs: (1) the derivatives of the loss are components of the SDEs; and (2) the loss would serve as boundary conditions of the SDEs. Therefore, understanding the loss surface would be a key step in establishing theoretical foundations for deep learning.

In convention, the learnability and optimization ability are usually guaranteed in “regularized” problems.¹ The “regularization” can be characterized by a variety of terms, including convexity, Lipschitz continuity, and differentiability. However, these factors are no longer secured in deep learning, at least not obviously. Neural networks usually consist of a large number of nonlinear activations. The nonlinearity in activations makes the loss surface extremely non-smooth and non-convex. The established guarantees for convex optimization become invalid. The prohibitive complexity of loss surface checks the community touching the geometry of the loss landscape and even deep learning theory for a long time. However, the complex geometry of the loss surface exactly characterizes the behavior of deep learning. It is the “high way” to understand deep learning via its loss surface.

The roles of over-parameterization of deep neural networks. Over-parameterization is usually treated as a major barrier in developing meaningful generalization bounds for deep learning via complexity-based approaches. However, recent studies suggest that over-parameterization would make major contributions in shaping the loss surface of deep learning - making the loss surface more smooth and even “analogous” convex. Moreover, many works have proven that neural networks in the extreme over-parameterization regime are equivalent to some simpler models, such as Gaussian kernels.

Theoretical foundations for several special structures in network architectures. In the previous review, we focus on the result generally stand for all neural networks. Meanwhile, deep neural networks designing involves many special techniques. These structures also contribute significantly to deep learning’s excellent performance. We review the theoretical results on convolutional neural networks, recurrent neural networks, and networks for permutation invariant/equivariant functions.

The intensive concerns in ethics and security and their relationships with deep learning theory. Deep learning has been deployed in an increasingly wide spectrum of application domains. Some of them

¹It is worth noting that the “regularized” here or the “regularization” later refer to the approaches to make ill-posed problems solvable. It is not necessarily relating to adding “regularizers”.

involve highly private personal data, such as images and videos on mobile phones, health data, and final records. Some other scenarios may require deep learning to deliver highly sensitive decision-making, such as mortgage approval, college admission, and credit assessment. Moreover, deep learning models have been shown vulnerable to adversarial examples. How to secure a deep learning system from breaches related to privacy preservation, fairness concerning, and adversarial attacks are of significant importance.

1.2 Overview of this paper

This paper reviews the recent advances in establishing the theoretical foundations of deep learning. We acknowledge that a few papers have reviewed deep learning theory. Alom *et al.* [9] give a survey of the techniques used in deep learning. Sun [214] reviews the theory of optimization in deep learning. E *et al.* [81] summarize results and challenges related to approximation and Rademacher complexity, loss surface, and the convergence and implicit regularization of optimization in deep learning. Our survey is the most comprehensive. We organize the literature in a unique view and provide new insights for future works.

The excellent generalizability of deep learning is like a “cloud” to conventional complexity-based learning theory: the over-parameterization of deep learning makes almost all existing tools vacuous. Existing works try to tackle this problem through three major paths: (1) developing size-independent complexity measures, which can characterize the complexity of the effective hypothesis space that can be learned, instead of the whole hypothesis space. The related works are discussed in Section 2; and (2) modelling the learned hypothesis through stochastic gradient methods, the dominant optimizers in deep learning, based on the stochastic differential functions and the geometry of the associated loss functions. The related works are reviewed in Sections 3 and 4; and (3) over-parameterization surprisingly bring many good properties to the loss functions and further secures the optimization and generalization performances. The related works are given in Section 5. Parallel to these, Section 6 reviews the theoretical foundations on special structures of network architecture.

Another important facet of machine learning is the rising concern on the ethical and security issues, including privacy preservation, adversarial robustness, and fairness protection. Specifically, privacy preservation and adversarial robustness have been discovered closely related to generalizability: a good generalizability usually means a good privacy-preserving ability; and more robust algorithms might have. Efforts of understanding the interplay between these issues in the scenario of deep learning are also discussed in this paper. The related works are discussed in Section 7.

1.3 Notations

Suppose $S = \{(x_1, y_1), \dots, (x_N, y_N) | X_i \in \mathcal{X} \subset \mathbb{R}^{d_X}, Y_i \in \mathcal{Y} \subset \mathbb{R}^{d_Y}, i = 1, \dots, N\}$ is a training sample set, where d_X and d_Y are the dimensions of the feature X and the label Y , respectively. For the brevity, we define $z_i = (x_i, y_i) \in \mathcal{Z} = \mathcal{X} \times \mathcal{Y}$. All z_i are independent and identically distributed (i.i.d.) observations of the random variable $Z = (X, Y) \in \mathcal{Z}$. Machine learning algorithms are designed to learn hypothesis h from the training data; all potential hypotheses constitute a hypothesis space \mathcal{H} .

A neural network is usually defined as the following form,

$$x \mapsto W_D \sigma_{D-1}(W_{D-1} \sigma_{D-2}(\dots \sigma_1(W_1 x))),$$

where D is the depth of the neural network, W_j is the j -th weight matrix, and σ_j is the j -th nonlinear activation function. Usually, we assume the activation function σ_j is a continuous function between Euclidean spaces and $\sigma_j(0) = 0$. Popular activations in practice including softmax, sigmoid, and tanh functions.

Generalization bound measures the generalization ability of an algorithm. For any hypothesis h learned by an algorithm \mathcal{A} , the expected risk $\mathcal{R}(h)$ and empirical risk $\hat{\mathcal{R}}_S(h)$ with respect to the training sample set S are respectively defined as follows,

$$\mathcal{R}(h) = \mathbb{E}_Z l(h, Z), \quad \hat{\mathcal{R}}_S(h) = \frac{1}{N} \sum_{i=1}^N l(h, z_i).$$

Additionally, when the output hypothesis h of algorithm \mathcal{A} is stochastic, we usually calculate the expectations of the expected risk $\mathcal{R}(h)$ and empirical risk $\hat{\mathcal{R}}_S(h)$ with respect to the randomness introduced by the algorithm \mathcal{A} as follows,

$$\mathcal{R}(\mathcal{A}) = \mathbb{E}_{\mathcal{A}(S)} \mathcal{R}(\mathcal{A}(S)), \quad \hat{\mathcal{R}}_S(\mathcal{A}) = \mathbb{E}_{\mathcal{A}(S)} \hat{\mathcal{R}}_S(\mathcal{A}(S)),$$

where $\mathcal{A}(S)$ is the hypothesis learned by the algorithm \mathcal{A} on the sample S . Then, the generalization error can be defined as the difference between the expected risk and empirical risk.

2 Capacity and complexity: Advances and predicament

Conventional statistical learning theory establishes generalization guarantees based on the complexity of the hypothesis space. Following this line, some works have proposed upper bounds of the generalization error of deep learning.

2.1 Worst-case bounds based on VC-dimension

A major measure for evaluating the hypothesis complexity in conventional statistical learning theory is the Vapnik-Chervonenkis dimension (VC-dimension, [28, 223]) which is defined as follows.

Definition 1 (growth function, shattering, and VC-dimension). *For any non-negative integer m , the growth function of hypothesis space \mathcal{H} is defined as follows,*

$$\Pi_{\mathcal{H}}(m) := \max_{x_1, \dots, x_m \in \mathcal{X}} |\{(h(x_1), \dots, h(x_m)) : h \in \mathcal{H}\}|.$$

If $\Pi_{\mathcal{H}}(m) = 2^m$, we say \mathcal{H} shatters the dataset $\{x_1, \dots, x_m \in \mathcal{X}\}$. We define the VC-dimension $VCdim(\mathcal{H})$ as the largest size of shattered set.

One can obtain a uniform generalization bound via the VC-dimension as follows.

Theorem 1. *Assume hypothesis space \mathcal{H} has VC-dimension $VCdim(\mathcal{H})$. Then, for any $\delta > 0$, with probability $1 - \delta$, the following inequality holds for any $h \in \mathcal{H}$,*

$$\mathcal{R}(h) \leq \hat{\mathcal{R}}(h) + \sqrt{\frac{2VCdim(\mathcal{H}) \log \frac{em}{VCdim(\mathcal{H})}}{m}} + \sqrt{\frac{\log \frac{1}{\delta}}{2m}},$$

where m is the training sample size.

Thus, results on the VC-dimension of neural networks can characterize its generalizability. Goldberg and Jerrum [90] gives an $\mathcal{O}(W^2)$ upper bound for the VC dimension of neural networks with parameter size W and depth D . It was improved to $\mathcal{O}(W \log(WD))$ by Bartlett and Williamson [22] and then to $\mathcal{O}(WL \log W + WL^2)$ by Bartlett *et al.* [20]. Other on VC-dimension were also given by Baum and Haussler [23] and Maass [158].

The tightest upper bound for the VC-dimension so far is proven by Harvey *et al.* [102] as follows,

Theorem 2. *Consider a neural network with W parameters and U units with activation functions that are piecewise polynomials with at most p pieces and of degree at most d . Let F be the set of (real-valued) functions computed by this network. Then,*

$$VCdim(\text{sgn}(\mathcal{F})) = \mathcal{O}(WU \log((d+1)p)).$$

However, such upper bounds on the VC-dimensions all heavily rely on the model size W , which is extremely large in deep learning. In some cases, the generalization bound can be significantly larger than the largest potential value of the loss function (such as 1 for the 0-1 loss).

Harvey *et al.* [102] also give a lower bound for the VC-dimension as the following theorem.

2.2 Margin bounds: From exponential depth dependence to depth independence

Margin bounds are another family of generalization bounds [223, 207, 17, 131, 219]. Compared with the worst-case bounds based on the VC-dimension, margin bounds deliver strong guarantees for the learned models that they can achieve small empirical margin loss for large confidence margin. Similar generalization guarantees can be obtained via the luckiness [207, 131].

Margin bound. For any distribution \mathcal{D} and margin $\gamma > 0$, we define the expected margin loss for hypothesis h as follows.

$$L_\gamma(h) = \mathbb{P}_{(x,y) \sim \mathcal{D}} \left(h(x)[y] \leq \gamma + \max_{j=y} h(x)[j] \right),$$

where $h(x)[j]$ is the j -th component of the vector $h(x)$. The generalization bounds under the margin loss are called margin bounds.

One can prove a margin bound via the the *covering number* [73, 104] or *Rademacher complexity* [130, 129, 21] of the hypothesis space defined as follows.

Definition 2 (covering number). *The covering number of covering number $\mathcal{N}(\mathcal{H}, \epsilon, \|\cdot\|)$ of space \mathcal{H} is defined to be the least cardinality of any subset $V \subset \mathcal{H}$ that covers \mathcal{H} at scale ϵ under metric $\|\cdot\|$; i.e., $\sup_{A \in \mathcal{H}} \min_{B \in V} \|A - B\| \leq \epsilon$.*

Definition 3 (empirical Rademacher complexity and Rademacher complexity). *Given a real-valued function class \mathcal{H} and a dataset S , the empirical Rademacher complexity is defined as follows,*

$$\hat{\mathfrak{R}}(\mathcal{H}) = \mathbb{E}_\epsilon \left[\sup_{h \in \mathcal{H}} \frac{1}{m} \sum_{i=1}^m \epsilon_i h(x_i) \right],$$

where $\epsilon = \{\epsilon_1, \dots, \epsilon_m\}$ is a random vector distributed in the space $\{-1, +1\}^m$. Further, we can define Rademacher complexity as follows,

$$\mathfrak{R}(\mathcal{H}) = \mathbb{E}_S \hat{\mathfrak{R}}(\mathcal{H}).$$

Intuitively, covering number indicates how many balls are needed to cover the hypothesis space; and Rademacher complexity measures the maximal correlations between the output of a hypothesis and a noise vector, and thus characterizes the “goodness-of-fit” of the hypothesis space to a noise. Apparently, a larger covering number or Rademacher complexity corresponds to a larger hypothesis complexity.

One can also upper bound the Rademacher complexity via the covering number through the Dudley entropy integral [169] as follows.

Theorem 3 (Dudley entropy integral). *Let \mathcal{H} be a real-valued function class taking values in $[0, 1]$, and assume that $0 \in \mathcal{H}$. Then,*

$$\hat{\mathfrak{R}}(\mathcal{H}) \leq \inf_{\alpha > 0} \left(\frac{4\alpha}{\sqrt{m}} + \frac{12}{m} \int_{\alpha}^{\sqrt{m}} \sqrt{\log \mathcal{N}(\mathcal{H}, \epsilon, \|\cdot\|_2)} d\epsilon \right).$$

Besides, one can obtain margin bounds via PAC-Bayesian theory as below, which corporates PAC theory and Bayesian statistics [160, 161].

Theorem 4 (PAC-Bayesian theory). *Suppose the prior and the posterior of the model parameter θ is \mathcal{P} and \mathcal{Q} . Then,*

$$\mathbb{E}_{\theta} \mathcal{R}(h(\theta)) = \mathbb{E}_{\theta} \hat{\mathcal{R}}(h(\theta)) + \sqrt{\frac{KL(\mathcal{Q}||\mathcal{P}) + \log \frac{m}{\delta}}{2(m-1)}}. \quad (1)$$

When the input data $\|x\| \leq B$ and the Frobenius norms of the D weight matrices are upper bounded by $M_F(1), \dots, M_F(D)$, respectively, Neyshabur *et al.* [177] prove via Rademacher complexity that with high probability, the generalization error scales as

$$\mathcal{O} \left(\frac{B 2^D \prod_{j=1}^D M_F(j)}{\sqrt{m}} \right). \quad (2)$$

However, this bound is exponentially dependent with the depth.

Then, through covering number, Dudley entropy integral, and Rademacher complexity, Bartlett *et al.* [19] obtain the following spectrally-normalized upper bound for the generalization error,

$$\tilde{\mathcal{O}} \left(\frac{\|X\|_2 \log W}{\gamma m} R_{\mathcal{A}} + \sqrt{\frac{1/\delta}{m}} \right),$$

where

$$R_{\mathcal{A}} = \left(\prod_{i=1}^D \rho_i \|A\|_{\sigma} \right) \left(\sum_{i=1}^D \frac{\|A_i^{\top} - M_i^{\top}\|_{2,1}^{2/3}}{\|A\|_{\sigma}^{2/3}} \right)^{3/2}.$$

This bound is strictly tighter than Neyshabur *et al.* [177]. Moreover, this bound suggests that one can boost the generalizability via controlling the spectral norm of the network. This explains spectral normalization [168], an effective regularizer mainly used in generative adversarial networks, and singular value bounding [118] for boosting neural network training performance.

Concurrently, Neyshabur *et al.* [176] prove a similar generalization bound as below via the PAC-Bayesian theory,

$$\mathcal{O} \left(\sqrt{\frac{B^2 D^2 W \log(DW) \prod_{i=1}^D \frac{\|W_i\|_F^2}{\|W_i\|_2^2} + \log \frac{Dm}{\delta}}{\gamma^2 m}} \right).$$

The proofs are in two-step: (1) they prove a perturbation bound which controls the changes of the output when perturbation introduced to the weights. Thereby, we can bound the sharpness of the network; and (2)

the perturbation bound further leads to a generalization bound via the PAC-Bayesian theory. However, this bound has been shown strictly weaker than the one given by Bartlett *et al.* [19].

Through introducing some norm constraints to the weights, Golowich *et al.* [92] improve the exponential depth dependence of eq. (2) by Neyshabur *et al.* [177] to polynomial dependence as follows,

$$\mathcal{O} \left(\frac{B\sqrt{D} \prod_{j=1}^D M_F(j)}{\sqrt{m}} \right),$$

or

$$\mathcal{O} \left(\frac{B\sqrt{D + \log m} \prod_{j=1}^D M_F(j)}{\sqrt{m}} \right).$$

When assuming

$$\prod_{j=1}^D \|W_j\|^2 \leq M,$$

where $M > 0$ is a constant and W_j is the j -th weight matrix, Golowich *et al.* [92] further improves the bounds to depth-independent as follows,

$$\tilde{\mathcal{O}} \left(B \left(\prod_{j=1}^D M_F(j) \right) R_G \right),$$

where

$$R_G = \min \left\{ \sqrt{\frac{\log \left(\frac{1}{\Gamma} \prod_{j=1}^D M_F(j) \right)}{\sqrt{m}}}, \sqrt{\frac{D}{m}} \right\},$$

and Γ is the upper bound of the spectral norm of all the weight matrices.

Moreover, Golowich *et al.* [92] also prove a lower bound as follows,

$$\Omega \left(\frac{B \prod_{j=1}^D M_F(j) W^{\max\{0, \frac{1}{2} - \frac{1}{p}\}}}{\sqrt{m}} \right).$$

It suggests that a tight bound would be $\mathcal{O} \left(\frac{1}{\sqrt{m}} \right)$.

3 Stochastic gradient descent: An implicit regularization

The over-parameterization makes a colossal hypothesis space. Despite that Golowich *et al.* [92] have provided generalization bounds explicitly independent with the network size, one might observe some implicit dependence. For example, the factor $M_F(j), j = 1, \dots, D$, are the upper bounds of weight norms which might be also influenced by the model size.

Stochastic gradient methods (SGM), including stochastic gradient descent (SGD) and its variants, are the dominant optimization methods in deep learning. They only explore a small part of the hypothesis space. This performs an implicit regularization to the deep learning models that controls the effective model capacity in an algorithm-dependent and data-dependent manner. Therefore, approaching algorithm-dependent and

data-dependent bounds via SGM may overcome the conventional statistical learning theory based on VC dimension or Rademacher complexity.

3.1 Stochastic gradient methods

To optimize the expected risk, a natural tool is the gradient descent (GD). Specifically, the gradient of the empirical risk in terms of the parameter θ below,

$$g(\theta(t)) \triangleq \nabla_{\theta(t)} \mathcal{R}(\theta(t)) = \nabla_{\theta(t)} \mathbb{E}_{(X,Y)} l(F_{\theta(t)}(X), Y),$$

and the corresponding update equation are defined as follows,

$$\theta(t+1) = \theta(t) - \eta g(\theta(t)),$$

where $\theta(t)$ is the parameter at the iteration t and $\eta > 0$ is the learning rate.

Stochastic gradient descent (SGD) estimates the gradient from mini batches of the training sample set to estimate the gradient $g(\theta)$. Let S be the indices of a mini batch, in which all indices are independently and identically (i.i.d.) drawn from $\{1, 2, \dots, N\}$, where N is the training sample size. Then similar to the gradient, the iteration of SGD on the mini-batch S is defined as follows,

$$\hat{g}_S(\theta(t)) = \nabla_{\theta(t)} \hat{\mathcal{R}}(\theta(t)) = \frac{1}{|S|} \sum_{n \in S} \nabla_{\theta(t)} l(F_{\theta(t)}(X_n), Y_n), \quad (3)$$

and

$$\theta(t+1) = \theta(t) - \eta \hat{g}(\theta(t)), \quad (4)$$

where

$$\hat{\mathcal{R}}(\theta) = \frac{1}{|S|} \sum_{n \in S} l(F_{\theta}(X_n), Y_n)$$

is the empirical risk on the mini batch and $|S|$ is the cardinality of the set S . For brevity, we rewrite

$$l(F_{\theta}(X_n), Y_n) = l_n(\theta)$$

in the rest of this paper.

Also, suppose that in step i , the distribution of parameter is Q_i , the initial distribution is Q_0 , and the convergent distribution is Q . Then SGD is used to find Q from Q_0 through a series of Q_i .

3.2 Generalization bounds on convex loss surfaces

When the loss is continuous and convex, the generalizability of SGMs has been intensively studied.

Ying and Pontil [238] prove the generalization bound in the following rate, when the step size is fixed:

$$\eta_t \simeq m^{-\frac{2\zeta}{2\zeta+1}},$$

$$\mathcal{O}\left(m^{-\frac{2\zeta}{2\zeta+1}} \log m\right),$$

where ζ is a constant.

Then, Dieuleveut *et al.* [65] improves the result to the following rate for the average, when $2\zeta + \gamma > 1$,

$$\mathcal{O}\left(m^{-\frac{2\min\{\zeta, 1\}}{2\min\{\zeta, 1\} + \gamma}}\right).$$

Other related works include Lin *et al.* [149]; Lin and Rosasco [150]; Chen *et al.* [50]; Wei *et al.* [227]. However, the convex assumption for the loss surface is usually too strong in deep learning. However, the loss surface is usually highly non-convex in deep learning. This makes the aforementioned results invalid.

3.3 Generalization bounds on non-convex loss surface

Results on non-convex loss surface have also been seen in the literature. Related works include generalization bounds via algorithmic stability and PAC-Bayesian theory.

Algorithmic stability. Bousquet and Elisseeff [29] propose to employ *algorithmic stability* to measure the stability of the output hypothesis when the training sample set is disturbed. Algorithmic stability has many versions. Here, we introduce a popular one as follows.

Definition 4 (uniform stability; cf. [29], Definition 6). *A machine learning algorithm \mathcal{A} is uniformly stable, if for any neighboring sample pair S and S' which are different by only one example, we have the following inequality,*

$$|\mathbb{E}_{\mathcal{A}(S)}l(\mathcal{A}(S), Z) - \mathbb{E}_{\mathcal{A}(S')}l(\mathcal{A}(S'), Z)| \leq \beta,$$

where Z is an arbitrary example, $\mathcal{A}(S)$ and $\mathcal{A}(S')$ are the output hypotheses learned on the training sets S and S' , respectively, and β is a positive real constant. The constant β is called the uniform stability of the algorithm \mathcal{A} .

It is natural that an algorithm insensitive to the disturbance the training data have good generalization abilities. Following this intuition, some generalization bounds have been proved based on the algorithmic stability [29, 236]. For example, based on the uniform stability, Bousquet and Elisseeff [29] prove the following theorem for the following *generalization in expectation*.

Theorem 5. *Let algorithm \mathcal{A} is β -uniformly stable. Then,*

$$|\mathbb{E}_{S, \mathcal{A}}[\hat{\mathcal{R}}[\mathcal{A}(S)] - \mathcal{R}[\mathcal{A}(S)]]| \leq \beta.$$

Recent works usually model SGM via stochastic differential equations (SDEs). Mini-batch stochastic gradients introduce randomness into SGM (*i.e.*, the stochastic gradient can be decomposed into sample-estimation of the gradient on the whole training set plus a noise). So the weight update in SGM can be formulated as a stochastic process. The output hypothesis is then drawn from the steady distribution of this stochastic process. We can eventually analyze the generalization and optimization of deep learning via the steady distribution.

SGM is usually formulated by eqs. (3) and (4). The stochastic gradient introduces gradient noise into the weight updates. When the gradient noise is modeled by a Gaussian distribution, SGM reduces to stochastic gradient Langevin dynamics(SGLD) as follows [159],

$$\Delta\theta(t) = \theta(t+1) - \theta(t) = -\eta\hat{g}_S(\theta(t)) = -\eta g(\theta) + \sqrt{\frac{2\eta}{\beta}}\Delta W, \Delta W \sim N(0, I),$$

where η is the learning rate and $\beta > 0$ is the inverse temperature parameter.

Hardt *et al.* [101] prove the following upper bound for uniform stability, which exponentially depends on aggregated step sizes and smoothness parameter.

Theorem 6. *Suppose the loss function is ϵ -smooth, L -Lipschitz, and not larger than 1 for every data. We run SGM with monotonically increasing step sizes $\alpha_t \leq c/t$ for T steps. Then, SGM is uniform stable with*

$$\beta \leq \frac{1 + 1/\beta c}{m - 1} (2cL^2)^{\frac{1}{\beta c + 1}} T^{\frac{1}{\beta c + 1}}.$$

This proof of this theorem mimics the proofs for convergence. Suppose the loss function l is L -Lipschitz constant with respect to the weight for any example; *i.e.*,

$$|l(w; z) - l(w'; z)| \leq L \|w - w'\|.$$

In other words, the stabilities can be measured by the stability of the weight. Therefore, one can just analyze how the weight diverges as a function of time t . The authors then deliver generalization bounds relying on the step sizes and the iteration numbers. Further, we can easily obtain a generalization bound based on this theorem.

Under five more assumptions, Raginsky *et al.* [194] prove the following upper bound for the expected excess risk,

$$\tilde{\mathcal{O}} \left(\frac{\beta(\beta + U)^2}{\lambda_*} \delta^{1/4} \log \left(\frac{1}{\epsilon} + \epsilon \right) \right) + \tilde{\mathcal{O}} \left(\frac{(\beta + U)^2}{\lambda_* + m} + \frac{U \log(\beta + 1)}{\beta} \right),$$

provided,

$$k = \tilde{\Omega} \left(\frac{\beta(\beta + U)}{\lambda_* \epsilon^4} \log^5 \left(\frac{1}{\epsilon} \right) \right),$$

and

$$\eta \leq \left(\frac{\epsilon}{\log(1/\epsilon)} \right)^4,$$

where U is the parameter size, m is the training sample size, the inverse temperature parameter $\beta > 1 \vee 2/m$, λ_* is a spectral gap parameter that controls the exponential rate of convergence of the dynamic of SGM to its stationary distribution, d is the factor concerning the dissipativity of the hypothesis h ; *i.e.*, for any datum z and any parameter w , for some $m > 0$ and $b > 0$,

$$\langle w, \nabla h(w, z) \rangle \leq m \|w\|^2 - b,$$

and

$$\epsilon \in \left(0, \frac{m}{4M^2} \wedge e^{-\tilde{\Omega} \left(\frac{\lambda_*}{\beta(\beta + U)} \right)} \right).$$

The proofs are completed in the following three steps: (1) when the step size is small enough, the authors prove that the dynamics of SGLD converge to continuous-time Langevin dynamic processes under 2-Wasserstein distance; (2) the authors then show that the Langevin dynamic then converge to a Gibbs distribution when the training time has been sufficiently long; and (3) the authors construct a Gibbs algorithm where the output hypothesis is almost an empirical risk minimizer. Also, the Gibbs algorithm is proven to be algorithmic stable under 2-Wasserstein distance. However, this result is based on the convergence of the weight distribution to the stationary distribution, which usually exponentially relies on the model dimension.

A canonical mutual information-based generalization bound is given in Xu and Raginsky [235],

$$\sqrt{\frac{2\sigma^2}{m}I(S; \theta)}, \quad (5)$$

where S is the training sample, θ is the model parameter, and the loss is σ -sub-Gaussian under the data distribution. This result stems from the following result on the informational theoretical “stability” of a two-input function $f(X, Y)$,

$$|f(X, Y) - f(\bar{X}, \bar{Y})| \leq \sqrt{2\sigma^2 I(X; Y)},$$

if the $f(\bar{X}, \bar{Y})$ is sub-Gaussian under distribution $P_{\bar{X}, \bar{Y}} = P_X \otimes P_Y$ where P_X and P_Y are distributions of X and Y , respectively. It is worth noting that the mutual information is calculated between the learned weights and all training data; it is an “on-average” estimate of how much information has been transformed into the learned model from the data.

Based on this result, Pensia *et al.* [190] gives a generalization error bound for SGLD,

$$\sqrt{\frac{R^2}{m} \sum_{t=1}^T \frac{\eta_t^2 L^2}{\sigma_t^2}}, \quad (6)$$

where the hypothesis is assumed to R -sub-Gaussian, η_t is the learning rate in the t -th iteration, L is the upper bound of the weight update $\Delta\theta_t$, and the gradient noise is assumed to be $N(0, \sigma_t^2 I)$. This result is then extended by Bu *et al.* [33]:

- Assume the loss $l(\theta, z)$ is R -sub-Gaussian under $z \sim \mu$, where z_i is the i -th datum, θ is the model parameter, and μ is the data generating distribution. Then, the generalization error has the following upper bound,

$$\frac{1}{m} \sum_{i=1}^m \sqrt{2R^2 I(\theta; z_i)}, \quad (7)$$

where m is the sample size.

- Assume the loss $l(\tilde{\theta}, \tilde{z})$ is R -sub-Gaussian under distribution

$$P_{\tilde{\theta}, \tilde{z}} = P_{\theta} \otimes P_z.$$

Then, the generalization error has the same upper bound (eq. 7).

In contrast with Xu and Raginsky [235], Bu *et al.* [33] proposes an “individual” version metric to measure the information transformed into the learned model from the data; see eq. (7) where the whole training sample S is replaced by individual datums $I(\theta; z_i)$.

Mou *et al.* [171] models SGLD as a Langevin diffusion dynamic and then prove an $\mathcal{O}(1/m)$ generalization bound and an $\mathcal{O}(1/\sqrt{m})$ generalization bound for SGLD via the stability and the PAC-Bayesian theory, respectively:

- **Algorithmic stability:** Assume that the loss function is upper bounded by C , the hypothesis is L -Lipschitz continuous, and the step size $\eta_t \leq \frac{\log 2}{\beta L^2}$ in any iteration t , the generalization error has the

following upper bound in average,

$$\mathcal{O} \left(\frac{2LC}{m} \left(\beta \sum_{t=1}^N \eta_t \right)^2 \right). \quad (8)$$

- **PAC-Bayesian theory:** Assume the empirical risk is regularized by the l_2 norm of the weight; i.e., $\frac{\lambda}{2} \|\theta\|^2$, and the loss function is sub-Gaussian. Then, we have the following high-probability generalization bound,

$$\mathcal{O} \left(\sqrt{\frac{\beta}{m} \sum_{k=1}^N \eta_k e^{-\frac{\lambda}{3}(T_N - T_k)} \mathbb{E}[\|g_k(\theta_k)\|^2]} \right),$$

where N is the iteration number, $T_k = \sum_{i=1}^k \eta_i$, β is the temporal parameter, $g_k(\theta_k)$ is the gradient in the k iteration,

Recently, Chen *et al.* [51] proves a trade-off between convergence and stability for all iterative algorithms respectively under convex smooth setting and strong convex smooth setting, which leads to an $\mathcal{O}(1/m)$ generalization bound. Other advances include the works by He *et al.* [106], Tzen [222], Zhou [254], Wen *et al.* [230], and Li *et al.* [145].

3.4 Generalization bounds relying on data-dependent priors

In Bayesian statistics, priors are given according to the “practice experience” or domain knowledge. In algorithm designing, one usually cannot assume they have understood the data. Thus, the prior needs to be independent with the training sample. Bayesian statistics usually adopts priors containing no “information” of the data, such as uniform distributions or Gaussian distributions. Moreover, the impact of prior usually goes to vanishing when training progress. From this aspect, setting data-independent priors would not comprise the performance in an asymptotic manner. Recent works slightly relax the restrictions in convention and advance the PAC-Bayesian theory.

Distribution-dependent priors. Some works design priors relying on the data generation distribution but still not directly relying on the training data. This would be reasonable since we can assume the data distribution has been fixed before the data was collected [142]. Such *distribution-dependent* priors have shown to be able to considerably tighten the generalization bounds.

Following this line, Lever *et al.* [142] tightens the PAC-Bayesian bound (see Theorem 4) to the following bound. For any positive constant C , we have that at least probability $1 - \delta > 0$ with respect to the drawn sample, the generalization error is upper bounded as follows,

$$\mathcal{R}(Q(h)) - C^* \hat{\mathcal{R}}(Q(h)) \leq \frac{C^*}{Cm} \left(\lambda \sqrt{\frac{2}{m} \log \frac{2\xi(m)}{\delta}} + \frac{\gamma^2}{2m} + \log \frac{2}{\delta} \right),$$

where $Q(h)$ is the posterior of the learned model, its density is as below,

$$q(h) \propto e^{-\gamma \hat{\mathcal{R}}(Q(h))}, \quad \xi(m) = \mathcal{O}(\sqrt{m}), \quad C^* = \frac{C}{1 - e^{-C}},$$

and γ is a positive constant.

Li *et al.* [145] further design a *Bayes-stability* framework to deliver generalization bounds based on

distribution-dependent priors. They tighten the result of Mou *et al.* [171] (eq. 8) to the follows,

$$\mathcal{O} \left(\frac{C}{m} \sqrt{\mathbb{E}_S \left[\sum_{t=1}^T \frac{\eta_t^2}{\sigma_t^2} g_e(t) \right]} \right),$$

where the gradient noise is assumed to be $\frac{\sigma_t}{\sqrt{2}}I$, $C > 0$ is the upper bound of the loss function, T is the number of iterations, η_t is the learning rate in the t -th iteration, and

$$g_e(t) = \mathbb{E}_{W_{t-1}} \left[\frac{1}{m} \sum_{i=1}^m \|\nabla F(W_{t-1}, z_i)\|^2 \right].$$

Data-dependent priors. Negrea *et al.* [175] further push the frontier that constructs priors not independent with data. Suppose S_J is a subset of $\subset S$ with size of $n < m$. One may design a prior exploiting S_J to deliver a data-dependent forecast of the posterior Q . We denote the subset as $S_J = \{z_{j_1}, \dots, z_{j_n}\}$, where all indices constitute a set J . The index set J is randomly drawn from $\{1, \dots, m\}$. Suppose that we have the following generalization bound,

$$\mathbb{E}^{\mathcal{F}} \left[\mathcal{R}(\theta) - \hat{\mathcal{R}}_{S_J}(\theta) \right] \leq B,$$

where $\hat{\mathcal{R}}_{S_J}$ is the empirical risk on the set S_J , $B > 0$ is the bound, \mathcal{F} is a σ -field such that $\sigma(S_J) \subset \mathcal{F} \perp \sigma(S_J^c)$. Also, we denote by U the uncertainty introduced by the machine learning algorithm. Negrea *et al.* [175] prove the following theorem, which considerably improve the one given in Xu and Raginsky [235] (eq. 5),

Theorem 7. *Let $J \subset [m]$, $|J| = n$ be uniformly distributed and independent from the training sample set S and the model parameter θ . Suppose the loss is σ -sub-Gaussian under the data distribution. When the posterior $Q = \mathbb{P}^S(\theta)$ and the prior P be a $\sigma(S_J)$ -measurable data-dependent on the data space,*

$$\begin{aligned} \mathbb{E} \left[\mathcal{R}(\theta) - \hat{\mathcal{R}}(\theta) \right] &\leq \sqrt{\frac{2\sigma^2}{m-n} I(\theta; S_J^c)} \\ &\leq \sqrt{\frac{2\sigma^2}{m-n} \mathbb{E}[KL(Q||P)]}. \end{aligned}$$

Moreover, let U be independent with S and θ . When the posterior $Q = \mathbb{P}^{S,U}(\theta)$ and the prior P be a $\sigma(S_J, U)$ -measurable data-dependent on the data space,

$$\begin{aligned} \mathbb{E} \left[\mathcal{R}(\theta) - \hat{\mathcal{R}}(\theta) \right] &\leq \sqrt{\frac{2\sigma^2}{m-n} I^{S_J, U}(\theta; S_J^c)} \\ &\leq \sqrt{\frac{2\sigma^2}{m-n} \mathbb{E}^{S_J, U}[KL(Q||P)]}. \end{aligned}$$

Following this path, further works include Haghighifam *et al.* [99] cooperating with Steinke and Zakynthinou [213], Dziugaite *et al.* [80], Hellström and Durisi [112], and their applications [53].




Figure 1: Scatter plots of test accuracy to ratio of batch size to learning rate. See [117].




Figure 2: Curves of test accuracy to batch size and learning rate. The four rows are respectively for (1) ResNet-110 trained on CIFAR-10, (2) ResNet-110 trained on CIFAR-100, (3) VGG-19 trained on CIFAR-10, and (4) VGG-19 trained on CIFAR-10. Each point represents a model. Each curve is based on 20 networks. See [106].

3.5 The role of learning rate and batch size in shaping generalizability

Keskar *et al.* [124] conducted experiments which show large-batch SGM usually finds sharp local minima, while its small-batch counterpart leads to flat minima. The authors also explain that flatter local minima can make SGM generalize better. However, the experiments wherein only involve two values of batch size, which considerably limits the generality of the results. Then, Dinh *et al.* [67] argue that most current notions for sharpness or flatness are ill-defined. The flatness/sharpness is still under intensive study.

Goyal *et al.* [95] further propose a *linear scaling rule* for SGM: “when the minibatch size is multiplied by k , multiply the learning rate by k ”. A similar statement is also given by Smith and Le [210]. Jastrzebski *et al.* [117] present empirical results on the relationships between the ratio $|S|/\eta$ and the dynamics, geometry, and the generalization, but still, only two values of both the batch size and learning rate are investigated.

Similar statements on the relationship between the generalizability and the hyperparameters are also presented in an empirical work by Chaudhari and Soatto [41]. Fig. 1 gives plots of the test accuracies of a multi-layer perceptron (MLP) trained on the dataset MNIST [139] in term of the ratio of the batch size to the learning rate respectively when 25.0% and 50.0% labels are randomly settled (see the original figure in [117]). These plots express the memorization ability²

²The concept of memorization is given by [?]. It is defined as the ability of a model to fit a dataset with random labels. If the training error can be very small even when the training labels are random, we argue the model is “memorizing” data, but not “learn” the data. Therefore, it is an important index expressing the generalization ability of a model.

He *et al.* [106] proves a PAC-Bayesian generalization bound under some assumptions. The generalization bound also suggests a positive correlation between the generalizability and the ratio of the learning rate and the batch size. He *et al.* [106] then conduct experiments that show a clear correlation between the generalizability and the ratio of batch size to learning rate; see fig. 3.

For more details and advances, please see recent works by Arpit *et al.* [13], Krueger *et al.* [133], Carlini *et al.* [40], and Smith and Le [211].

3.6 Interplay of optimization and Bayesian inference

An interesting interplay is observed in the research of optimization and Bayesian inference [234]: (1) Bayesian inference (or sampling) can be treated as optimization in the probability space (*e.g.*, in stochastic gradient Markov chain Monte Carlo) and can be scaled through optimization (*e.g.*, in variational inference); and (2) stochastic gradient Markov chain Monte Carlo also performs Bayesian approximation.

Bayesian inference. For parametric machine learning models, Bayesian learning aims to obtain the posterior of model parameters and then approach the best parameter. However, the analytic expression of the posterior is unknown in most real-world cases. To solve this problem, classical methods employ Markov chain Monte Carlo (MCMC) methods to infer the posterior, such as the Metropolis-Hastings algorithms [103, 87] and hybrid Monte Carlo (HMC; [72]). Bayesian inference has been used in many areas, including topic model [136, 246], Bayesian neural network [155, 200], and generative models [127, 93, 128].

Scaling Bayesian inference via optimization. Bayesian inference would be prohibitively time-consuming on large-scale data. To address this issue,

(1) *Stochastic gradient Markov chain Monte Carlo* (SGMC; [157]) introduces stochastic grsestimates [198] into MCMC. The family of SGMCMC algorithms includes stochastic gradient Langevin dynamics (SGLD; [229]), stochastic gradient Riemannian Langevin dynamics (SGRLD; [189]), stochastic gradient Fisher scoring (SGFS; [4]), stochastic gradient Hamiltonian Monte Carlo (SGHMC; [49]), stochastic gradient Nosé-Hoover Thermostat (SGNHT; [66]), etc. This paper analyses SGLD as an example of the SGMCMC scheme.

(2) *Variational inference* [115, 27, 245] employs a two-step process to inference the posterior:

1. We define a family of distributions,

$$\mathcal{Q} = \{q_\lambda | \lambda \in \Lambda\},$$

where λ is the parameter.

2. We search for the element from the distribution family \mathcal{Q} that is the closest to the posterior $p(\theta|S)$, under KL divergence, *i.e.*,

$$\min_{\lambda} \text{KL}(q_\lambda(\theta) || p(\theta|S)).$$

Minimizing the KL divergence is equivalent to maximizing the evidence lower bound (ELBO) defined as below,

$$\text{ELBO}(\lambda, S) = \mathbb{E}_{q_\lambda} \log p(\theta, S) - \mathbb{E}_{q_\lambda} \log q_\lambda(\theta).$$

Understanding stochastic gradient-based optimization via Bayesian inference. E [228] proposed to understand the optimization in deep learning via stochastic differential equations. Mandt *et al.* [159] propose to understand SGM as approximate Bayesian inference. Following this line, a large volume of literature has been delivered on SGM as shown in the previous subsections.




Figure 3: Scatter plots of accuracy on test set to ratio of batch size to learning rate. The four colors, blue, green, orange, and purple, are for results of (1) VGG-19 on CIFAR-10, (2) ResNet-110 on CIFAR-10, (3) VGG-19 on CIFAR-100, and (4) ResNet-110 on CIFAR-100. Each point represents a model. Totally 1,600 points are plotted. See [106].

4 Taking the high way: Study on loss surface

A major barrier recognized by the whole community is that deep neural networks' loss surfaces are extremely non-convex and even non-smooth. Such non-convexity and non-smoothness make the analysis of the optimization and generalization properties prohibitively difficult. A natural idea is to bypass the geometrical properties and then approach a theoretical explanation. Some papers argue that such "intimidating" geometrical properties are exactly the major factors that shape the properties of deep neural networks, and also the key to explaining deep learning.

4.1 Does spurious local minimum exist?

If all local minima are globally optimal, then we would not need to worry about optimization for deep learning. Then, is it true?

4.1.1 Linear networks have no spurious local minima

Linear neural networks refer to the neural networks whose activations are all linear functions. For linear neural networks, the loss surfaces do not have any spurious local minima: all local minima are equally good; they are all global minima.

Kawaguchi [121] proves that linear neural networks do not have any spurious local minimum under several assumptions: (1) the loss functions are squared losses; (2) XX^T and XY^T are both full-rank, where X is the data matrix and Y is the label matrix; (3) the dimension of the input layer is larger than the dimension of the output layer; (4) all the eigenvalues of matrix $YX^T(XX^T)^{-1}XY^T$ are different from each other. Lu and Kawaguchi [156] replace the assumptions to one: the data matrix X and the label matrix Y are both full-rank, which is, however, more restrictive.

Later, Zhou and Liang [253] prove that all critical points are global minima when some conditions hold. The authors prove this based on a new result on the analytic formulation of the critical points.

4.1.2 Nonlinear activations bring infinite spurious local minima

However, this good property no longer stands when the activations are nonlinear.

Zhou and Liang [253] show that one-hidden layer ReLU neural networks have spurious local minima. However, the proof relies on that the hidden layer has only two nodes. Swirszcz *et al.* [217] relax the assumption to that most of the neurons are not activated. Safran and Shamir [202] show that two-layer ReLU networks have spurious local minima in a computer-assisted. Yun *et al.* [240] prove that one-hidden-layer ReLU neural networks have infinitely many spurious local minima when the outputs are one-dimensional. Goldblum *et al.* [91] prove that the training performance of all the local minima equals to a linear model for neural networks of any depth, and verify it by experiments. He *et al.* [108] further extend the condition to arbitrary depth, arbitrary activations, and arbitrary output dimension.

4.2 Geometric structure of loss surface.

Then, how would the loss surface look like? Knowledge of the geometrical structures can considerably advance the current understanding of deep learning.

Linear partition of the loss surface. Soudry and Hoffer [212] first highlight a *smooth and multilinear partition* exists on the loss surfaces of neural networks. The nonlinearities in the piecewise linear activations divide the loss surface into multiple smooth and multilinear open regions. Specifically, every nonlinear point in the activation functions creates a group of the non-differentiable boundaries between the cells, while the linear parts of activations correspond to the smooth and multilinear interiors.


He *et al.* [108] prove several sophisticated properties of the partition: (1) if there are local minima within an open cell, they are all equally good in the sense of empirical risk, and therefore, all local minima are global minima within this cell; (2) the local minima in any open cell form an equivalence class. Moreover, these local minima are concentrated in a valley; and (3) when all activations are linear, the partition collapses to one single cell, covering linear neural networks as a special case.

The property (2) by He *et al.* [108] explains an important property of *mode connectivity*: the solutions found by stochastic gradient descent or its variants are connected with each other by a path in the weight space, on which all points have almost the same empirical risk. This property is found in two empirical works by Garipov *et al.* [85] and Draxler *et al.* [68]. A recent theoretical work by Kuditipudi *et al.* [134] proves that the mode connectivity can be guaranteed by dropout stability and noise stability.


Correspondence between the landscapes of empirical risk and expected risk. Two works by Zhou and Feng [251] and Mei *et al.* and [163] prove a linkage between the landscapes of the empirical risk surface and expected risk. This correspondence highlights the possibility of employing the results on the geometry of loss surfaces (empirical risk surfaces) to understand of generalizability (the gap between expected risks and empirical risks). Specifically, the authors prove that all of (1) the gradient of empirical risk, (2) the stationary points of the empirical risk, and (3) the empirical risk itself can be asymptotically approximated by their population counterparts. Moreover, they prove a generalization bound for non-convex case: at least by probability $1 - \delta$,

$$\mathcal{O} \left(\tau \sqrt{\frac{9}{8} [1 + c_r (D - 1)]} \sqrt{\frac{s \log(mU/D) + \log(4/\delta)}{m}} \right),$$


where all datums x are assumed to be τ -sub-Gaussian, $c_r = \max \left(r^2/16, (r^2/16)^{l-1} \right)$, and s is the number of non-zero components in the weights. Later, [163] prove a similar correspondence between the Hessian




(a) Spectrum of train Hessian on MNIST



(b) Spectrum of train Hessian on CIFAR-10



(c) Spectrum of test Hessian on MNIST



(d) Spectrum of test Hessian on CIFAR-10

Figure 4: The spectrums of train and test Hessian matrices on MNIST and CIFAR10. See [188].

matrices of the expected risk and the empirical risk under an assumption that the sample size is larger than the parameter size.

Eigenvalues of the Hessian. Sagun *et al.* in 2016 [203] and 2018 [204] and Papyan in 2018 [188] conduct experiments to study the the eigenvalues of the Hessian of the loss surface. Sagun *et al.* [203, 204] discover that: (1) a large bulk of eigenvalues are centered close to zero; and (2) several outliers are located away from the bulk. Papyan [188] gives the full spectrums of the Hessian matrix. Fig. 4 compares among the spectrum of Hessian matrices when training and testing a VGG-11 on the datasets MNIST and CIFAR-10 (originally from a recent empirical work by [188, p. 2, figs. 1(a) and 1(c) and p. 3, figs. 2(a) and 2(c)]).

5 Reflecting the role of over-parameterization: Is it only harmful?

One might blame the over-parameterization of deep learning for its lack of theoretical foundations. However, recent works also discover that such over-parameterization would have also contributed to the success of deep learning.

Neural tangent kernel. Neal in 1995 [173] and 1996 [174], Williams [231], and Hazan and Jaakkola[105] prove that an infinite-width shallow neural network is equivalent to a Gaussian process. Damianou and Lawrence [61], Duvenaud *et al.* [74], Hensman and Lawrence [113], Lawrence and Moore [137], Bui *et al.* [34], and Lee *et al.* [140], gradually extend the equivalence to any infinite-width neural network of an arbitrary depth. Lee *et al.* [141] further prove that infinite-width neural networks of any depth behaves as linear models during training with gradient descent. Jacot *et al.* [116] prove that the learned model for the least-squares regression problem during training is characterized by a linear differential equation the setting under the infinite-width regime.

Influence of over-parameterization on the loss surface. The over-parameterization would significantly

reshape the loss surfaces of neural networks. Choromanska *et al.* [55] conducted an experiment that shows: (1) a large proportion of the local minima on the loss surface of an over-parameterized network are “equivalent” with each other (they have the same empirical risk); and (2) neural networks with small or normal sizes have spurious local minima but the proportion of the spurious local minima decreases rapidly when the network size increase. Li *et al.* [144] prove that over-parameterized fully-connected deep neural networks do not have any spurious local minimum under assumptions: (1) the activation functions are continuous; and (2) the loss functions are convex and differentiable. Nguyen *et al.* [180] extend the “no bad local minima” property to networks trained under the cross-entropy loss. Nguyen [179] further discover that the global minima are connected with each other and concentrated on a unique valley if the neural networks are sufficiently over-parameterized.

Influence of over-parameterization on the optimization. It has been shown that the over-parameterization contribute to securing the optimization performance of gradient-based optimizers. Du *et al.* [71] prove that gradient descent converges to global local minima under the quadratic loss at a linear convergence rate for two-layer neural networks as long as the network widths are sufficiently large. The authors also suggest that the over-parameterization can restrict the weights in all iterations close to the random initialization. Chizat and Bach [54] model the optimization of one-hidden-layer networks as minimizing a convex function of a measure discretized into a mixture of particles by continuous-time gradient descent. They prove that the gradient flow characterizing the optimization converges to the global minimizers.

Going beyond one hidden layer, Allen-Zhu *et al.* [8], Du *et al.* [69], and Zou *et al.* [255] concurrently prove that SGD converge to global optimal solution of over-parameterized deep neural networks in polynomial time with under slightly different assumptions on the network architectures and training data. Chen *et al.* [52] prove that the convergence of optimization is guaranteed if the network width is polylogarithmic with the sample size n and $1/\epsilon$ where ϵ is the targeted error.

Influence of over-parameterization on the generalization and learnability. Andoni *et al.* [10] study two-layer neural networks under assumptions that the node in the first layer are all linear functions while the hidden nodes in the last layer are non-linear. The authors prove that sufficiently wide neural network can learn any low-degree polynomial functions when trained by generic gradient descent algorithm and all weights are randomly initialized. Brutzkus *et al.* [32] also give generalization guarantees for two-layer neural networks under some assumptions. Li and Yuan [147] prove that over-parameterized one-hidden-layer ReLU networks have guaranteed small generalization error if they are trained by SGD for multi-class classification problems from random initialization. Also, the data is assumed to be a mixture of well-separated distributions.

Arora *et al.* [12] and Allen-Zhu *et al.* [7] prove that the sample complexity for neural networks with two or three layers is almost independent with the parameter size. Chen *et al.* [52] also show that the generalization is guaranteed if the network width is polylogarithmic with the sample size n and ϵ^{-1} where ϵ is the targeted error. Cao *et al.* [39] also prove generalization bounds for wide and deep neural networks. Wei *et al.* [226] prove that regularizers can significantly reshape the generalization and optimization.

Influence of the network depth on the generalizability. Canziani *et al.* [38] suggest that deeper neural networks could have better generalizabilities, summarizing the applications of neural networks in practice. This plays as a major part in the irreconciliation between the over-parameterization of neural networks and their excellent generalizabilities. Besides the previous results, another possible explanation is from information theory. Zhang *et al.* [249] prove a generalization bound that characterizes how the generalizability evolves when the network becomes deeper, based on some techniques developed by Xu and Raginsky [235].

6 Theoretical foundations for specific architectures

In the previous sections, we discuss the generalizabilities of neural networks in general. This section discusses the influence of some specific network structures.

Convolutional neural networks. Convolutional neural networks (CNNs) introduce convolutional layers into deep learning. It has been widely employed in computer vision [132], natural language process [239], and deep reinforcement learning [209].

Du *et al.* [70] study several statistical properties of CNNs. The authors prove that for an c -dimensional convolutional filter with linear activation needs $\tilde{O}(c/\epsilon^2)$ training examples for achieving a population prediction error ϵ on d_X -dimensional input, while a fully-connected counterpart needs $\Omega(d/\epsilon^2)$ training examples. Moreover, the authors prove that a one-hidden-layer CNN needs $\tilde{O}((c+r)/\epsilon^2)$ training examples when the output dimension is r and the ratio of stride size to filter size is fixed. Zhou and Feng [252] prove (1) an $\mathcal{O}\left(\log\left(\prod_{i=0}^D b_i\right)\right)$ generalization bound for CNN, where b_i is the magnitude of the parameter in the i -th layer; and (2) one-to-one correspondence between convergence guarantees between the stationary points on the empirical risk landscape and their population counterparts. Lin and Zhang [151] studies how the sparsity and permutation of convolutional layers would influence the spectral norm and then the generalizability.

Several works have also studied the generalizability of CNNs with residual connections [110], including He *et al.* [107] and Chen *et al.* [44].

Recurrent neural networks. Recurrent Neural Networks (RNNs) possess a recurrent structure and have demonstrated promising performance in processing sequential data analysis, and have been deployed in many real-world areas, including natural language processing [15, 216] and speech recognition [96, 97]. Chen *et al.* [48] develops a generalization bound for RNNs based on works by Neyshabur *et al.* [176] and Bartlett *et al.* [19]. Allen-Zhu and Li [6] also analyze the generalizability of RNNs.

Tu *et al.* [220] develop a gradient measure based on Fisher-Rao norm, which delivers a generalization bound as follows.

Theorem 8. Fix margin parameter α , then for any $\delta > 0$, with probability at least $1 - \delta$, the following holds for every RNN whose weight matrices $\theta = (U, W, V)$ satisfy $\|V^T\|_1 \leq \beta_V$, $\|W^T\|_1 \leq \beta_W$, $\|U^T\|_1 \leq \beta_U$ and $\|\theta\|_{fs} \leq r$:

$$\begin{aligned} & \mathbb{E}[\mathbf{1}_{\mathcal{M}_{y_L}(x,y) \leq 0}] \\ & \leq \frac{4k}{\alpha} \left(\frac{r\|X\|_F}{2n} \sqrt{\frac{1}{\lambda_{\min}(\mathbb{E}(xx^T))}} + \frac{1}{n} \beta_V \beta_U \|X^T\|_1 \Lambda \right) \\ & \quad + \frac{1}{n} \sum \mathbf{1}_{\mathcal{M}_{y_L}(x_i, y_i) \leq a} + 3\sqrt{\frac{\log \frac{2}{\delta}}{2n}}. \end{aligned} \tag{9}$$

The gradient measure considerably tightens the bound and characterizes the relationship between generalizability and trainability. It suggests adding noise to the input data to boost the generalizability. It also justifies the technique gradient clipping [165, 86, 191].

Networks for permutation invariant/equivariant functions. Zaheer *et al.* [242], Cohen and Welling [56], and Cohen *et al.* [57] adopt machine learning to tasks based on "sets", which are usually permutation invariant or equivariant. These models are then applied to a variety of scenarios, including 3D point cloud processing

[146], chemistry [82], and astronomy [182, 197]. Sannai *et al.* [206] prove a universal approximation theorem for this family of networks. Moreover, Sannai *et al.* [206] show that the *free* parameters in this network family are exponentially fewer than usual models, which further suggests a significantly better generalizability. Sannai and Imaizumi [205] prove that the deep model in Zaheer. *et al.* [242] can improve generalization bound by $\sqrt{n!}$, where n is the number of permuting coordinates of data. Brutzkus and Globerson [31] also demonstrate that larger model size can boost the generalizability from both theoretical and empirical aspects.

7 Looking beyond: Rising concerns of ethics and security

Rising concerns are seen on the ethical and security issues of deep learning, including privacy preservation, fairness protection, and adversarial robustness. Understanding the theoretical foundations of such issues are also of high importance. This section summarizes the recent advances in these areas.

7.1 Privacy preservation

Deep learning has been deployed to process massive personal data, including financial and medical records. It is of profound importance to discover the high-value population knowledge while protecting the highly sensitive privacy [77, 78, ?].

Differential privacy. A major measure in measuring the privacy-preserving ability of an algorithm is differential privacy [78]. We call a learning algorithm is (ϵ, δ) -differential private if when the training data is exposed to some small disturbance, the change in the output hypothesis is restricted as follows,

$$\log \left[\frac{\mathbb{P}_{\mathcal{A}(S)}(\mathcal{A}(S) \in B) - \delta}{\mathbb{P}_{\mathcal{A}(S')}(\mathcal{A}(S') \in B)} \right] \leq \epsilon, \quad (10)$$

where B is an arbitrary subset of the hypothesis space and (S, S') is a neighboring sample set pair in which S and S' only differ by one example. The left-hand side of eq. (10) is also called *privacy loss*.

In eq. (10), we use division operation to measure the change of the output hypothesis. Through modifying the division operation or inserting some assumptions, many variants of notions for measuring privacy-preserving ability are proposed: (1) *concentration differential privacy* assumes that the privacy loss is sub-Gaussian [79, 35]; (2) *mutual-information differential privacy* and *KL differential privacy* replace the division operation in privacy loss by mutual information and KL divergence, respectively [58, 225, 148, 42]; (3) *Rényi differential privacy* further replaces the KL divergence by Rényi divergence [167, 88]; etc.

These techniques have been also employed for privacy preservation in deep learning [1].

Generalization-privacy relationship. Some works have suggested that the privacy-preserving ability measured by differential privacy is almost equivalent to the generalizability.

Dwork *et al.* [75] prove a high-probability generalization bound for (ϵ, δ) -differentially private machine learning algorithms as follows,

$$\mathbb{P} \left[\mathcal{R}(\mathcal{A}(S)) - \hat{\mathcal{R}}_S(\mathcal{A}(S)) < 4\epsilon \right] > 1 - 8\delta^\epsilon.$$

Oneto *et al.* [183] improves the generalization bound as below

$$\mathbb{P} \left[\text{Diff } \mathcal{R} < \sqrt{6\hat{\mathcal{R}}_S(\mathcal{A}(S))\hat{\varepsilon}} + 6(\varepsilon^2 + 1/N) \right] > 1 - 3e^{-N\varepsilon^2},$$

and

$$\mathbb{P} \left[\text{Diff } \mathcal{R} < \sqrt{4\hat{V}_S(\mathcal{A}(S))\hat{\varepsilon}} + \frac{5N}{N-1}(\varepsilon^2 + 1/N) \right] > 1 - 3e^{-N\varepsilon^2},$$

where

$$\begin{aligned} \text{Diff } \mathcal{R} &= \mathcal{R}(\mathcal{A}(S)) - \hat{\mathcal{R}}_S(\mathcal{A}(S)), \\ \hat{\varepsilon} &= \varepsilon + \sqrt{1/N}, \end{aligned}$$

and $\hat{V}_S(\mathcal{A}(S))$ is the empirical variance of $l(\mathcal{A}(S), \cdot)$:

$$\hat{V}_S(\mathcal{A}(S)) = \frac{1}{2N(N-1)} \sum_{i \neq j} [\ell(\mathcal{A}(S), z_i) - \ell(\mathcal{A}(S), z_j)]^2.$$

Nissim and Stemmer [181] further prove that

$$\mathbb{P} \left[\mathcal{R}(\mathcal{A}(S)) - \hat{\mathcal{R}}_S(\mathcal{A}(S)) < 13\varepsilon \right] > 1 - \frac{2\delta}{\varepsilon} \log \left(\frac{2}{\varepsilon} \right).$$

The tightest generalization bound by far is given by He *et al.* [109] as follows,

$$\mathbb{P} \left[\left| \hat{\mathcal{R}}_S(\mathcal{A}(S)) - \mathcal{R}(\mathcal{A}(S)) \right| < 9\varepsilon \right] > 1 - \frac{e^{-\varepsilon}\delta}{\varepsilon} \ln \left(\frac{2}{\varepsilon} \right). \quad (11)$$

This generalization bound is delivered in three stages. The authors first prove an on-average generalization bound for any (ε, δ) -differentially private multi-database learning algorithm

$$\tilde{\mathcal{A}} : \vec{S} \rightarrow \mathcal{H} \times \{1, \dots, k\},$$

as follows,

$$\left| \mathbb{E}_{\vec{S} \sim \mathcal{D}^N} \left[\mathbb{E}_{\mathcal{A}(\vec{S})} \left[\hat{\mathcal{R}}_{S_{i_{\mathcal{A}(\vec{S})}}} \left(h_{\mathcal{A}(\vec{S})} \right) \right] - \mathbb{E}_{\mathcal{A}(\vec{S})} \left[\mathcal{R} \left(h_{\mathcal{A}(\vec{S})} \right) \right] \right] \right| \leq e^{-\varepsilon}k\delta + 1 - e^{-\varepsilon}, \quad (12)$$

where the loss function $\|l\|_\infty \leq 1$.

They then obtain a high-probability generalization bound for multi-database algorithms as follows,

$$\mathbb{P} \left[\hat{\mathcal{R}}_{S_{i_{\mathcal{A}(\vec{S})}}} \left(h_{\mathcal{A}(\vec{S})} \right) \leq \mathcal{R} \left(h_{\mathcal{A}(\vec{S})} \right) + ke^{-\varepsilon}\delta + 3\varepsilon \right] \geq \varepsilon. \quad (13)$$

They eventually prove eq. (11) by reduction to absurdity.

7.2 Algorithmic fairness

Deep learning has been deployed in many critical decision-making areas, such as mortgage approval [195], credit card assessment [125, 45], college admission [2, 46], employee selection [83], and recidivism prediction [30].

However, long-standing systematical discrimination has been observed in these areas against people based on their backgrounds, including races, genders, nationalities, ages, and religious. Additionally, people with minority backgrounds are institutionally under-represented in historical data. Meanwhile, it has been revealed in an increasing number of reports that the discrimination exists in the data are inherited and even intensified in the models learned by machine learning-based algorithms. The fairness concerns are particularly severe due to the black-box nature. For example, facial recognition is reportedly to be “accurate, if you are a white guy” [153]; a police system for recidivism prediction, Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) [30], tends to label minority people with significantly higher possibility to be reoffenders. Mitigating the fairness risk is timely and important.

Many approaches have been presented to address the fairness issues. They can be roughly categorized into three parts: (1) *pre-processing methods*, which aims to “rectify” data before processed by learning algorithms [36, 84, 119, 76, 37, 100, 243, 154]; (2) *post-processing methods* redress the outputs of learning systems to remove the potential discrimination [233]; and (3) *in-processing methods* adapt constraints or regularizers in the training procedures to enforce fairness [241, 3, 120, 164, 123].

7.3 Adversarial robustness

Szegedy *et al.* [218] discover that neural networks are vulnerable to adversarial examples which are slightly different from true datums but are likely to be misclassified by neural networks. Since then, many approaches have been proposed on attacking neural networks by adversarial examples or defending the adversarial attacks [26, 185, 186, 187, 94, 178, 186, 89].

Generalizability in the presence of adversarial examples. Schmidt *et al.* [208] first study the generalizability in the presence of adversarial examples. They define a new notion, *adversarial robustness generalization* for measuring the generalizability when adversarial examples exist in data. In contrast, the generalizability when there is no adversarial examples is evaluated by the *standard generalization*. They prove that to secure the same generalization, adversarial robustness generalization needs a higher sample size, when the data is drawn from some specific distributions and the adversarial examples are under the l_∞ norm.

Cullina *et al.* [59] extends the PAC-learning framework to learning with adversarial examples. The authors propose a *corrupted hypothesis classes* which is defined to be the hypothesis classes for the classification problem in the presence of adversaries. The VC dimension of a corrupted hypothesis class is then defined to be *adversarial VC-dimension*, while the *standard VC-dimension* is defined on clean datasets. They show that the adversarial VC-dimension can be either larger or smaller than the standard VC-dimension. Montasser *et al.* [170] further define an *agnostic robust PAC learnability* and *realizable robust PAC learnability*. When a learning algorithm whose rule meets the requirements of both agnostic robust PAC learnability and realizable robust PAC learnability, the learning is defined to be improper. They then show that under an improper learning rule, any hypothesis space \mathcal{H} is robustly PAC learnable, if its VC dimension is finite.

Generalizability of adversarial training. Amongst the defense approaches, adversarial training [60, 143, 16, 250] can considerably improve the adversarial robustness of deep neural networks against adversarial examples. Specifically, adversarial training can be formulated as solving the following minimax-loss problem

[221, 237, 126],

$$\min_{\theta} \frac{1}{N} \sum_{i=1}^N \max_{\|x'_i - x_i\| \leq \rho} l(h_{\theta}(x'_i), y_i),$$

where h_{θ} is the hypothesis parameterized by θ , N is the training sample size, x_i is a feature, y_i is the corresponding label, and l is the loss function. Intuitively, adversarial training optimizes neural networks according to the performance on worst-case examples, which are most likely to be adversarial examples.

Based on the minimax loss framework, several generalization bounds have been proposed [237, 221, 126]. All these results suggest that adversarial training would comprise generalizability when improving the adversarial robustness.

Yin *et al.* [237] studies how adversarial training would influence the generalizability based on the techniques developed by Bartlett *et al.* [19] and Raghunathan *et al.* [193]. They prove a tight upper bound on the Rademacher complexity which leads to a generalization bound. However, the generalization bound is only for neural networks with one hidden layer and ReLU activation functions.

To overcome this shortcoming, Khim *et al.* [126] next prove a generalization bound via function transformation. They prove an upper bound for the Rademacher complexity via techniques developed by Golowich *et al.* [92].

Tu *et al.* [221] develop another path for evaluating the generalizability: (1) they first show that one can remove the maximization operation in the objective function in the minimax framework, $\max_{\|x'_i - x_i\| \leq \rho} l(h_{\theta}(x'_i), y_i)$ via a re-parameterization mapping, while the parameter distribution would not shift much under the Wasserstein distribution; (2) they prove a local worst-case risk bound for the re-parameterized optimization problem which has been a minimization problem; and (3) they prove a generalization bound based on the local worst-case risk bound. This generalization bound is more generalization in two aspects: (1) the adversarial examples encompass all bounded adversaries under the l_q norm where l_q is arbitrary; (2) the approach can be applied to multi-class classification and almost all practical loss functions, including hinge loss and ramp loss.

On the other side, Bhagoji *et al.* [25] also prove a lower bound for the generalization error of classifiers in the presence of adversarial examples. Min *et al.* [166] discover that the performance of adversarial training does not always have a clear correlation with the training sample size. The authors categorize all learning problems in three parts: (1) in the *weak adversary regime*, an increasing training sample size corresponds to an improving generalizability; (2) in the *medium adversary regime*, one may observe a double descent curve of the generalization error, when the training sample size increases; and (3) in the *strong adversary regime*, more training data causes the generalization error to increase. Here, the three regimes can be determined by the level of adversaries. Similar phenomena to the medium adversary regime have also been discovered by Chen *et al.* [47].

Other theoretical analyses include an improved generalization bound via Rademacher complexity by Attias *et al.* [14], the trade-off between accuracy and robustness [248], a generalization bound via optimal transport and optimal couplings by Pydi and Jog [192].

8 Summary and discussion

In this paper, we review the literature on the deep learning theory from six aspects: (1) complexity and capacity-based method for analyzing the generalizability of deep learning; (2) stochastic differential equations and their dynamic systems for modelling stochastic gradient methods (SGM), which characterize the implicit

regularization introduced by the limited exploration space; (3) the geometry of the loss landscape governing the trajectories of the dynamic systems; (4) the roles of over-parameterization of deep neural networks from both positive and negative perspectives; (5) parallel to the previous parts, the theoretical foundations for some special structures in the neural network architecture; and (6) the rising concerns on ethics and security and their relationships with deep learning theory.

Many works give upper bounds for the generalization error of deep learning. The results are summarized in the following table.

Table 1: Comparisons of major generalization bounds for deep learning.

	Tool	Generalization bound
Goldberg and Jerrum [90]	VC-dimension	$\mathcal{O}\left(\sqrt{\frac{\log(mW^2)}{mW^2}}\right)$
Bartlett <i>et al.</i> [20]	VC-dimension	$\mathcal{O}\left(\sqrt{\frac{\log(mWL \log W + WL^2)}{mWL \log W + WL^2}}\right)$
Harvey <i>et al.</i> [102]	VC-dimension	$\mathcal{O}\left(\sqrt{\frac{\log(mWU \log((d+1)p))}{mWU \log((d+1)p)}}\right)$
Neyshabur <i>et al.</i> [177]	Rademacher complexity	$\mathcal{O}\left(\frac{B2^D \prod_{j=1}^D M_F(j)}{\sqrt{m}}\right)$
Bartlett <i>et al.</i> [19]	Rademacher complexity & covering number	$\tilde{\mathcal{O}}\left(\frac{\ X\ _2 \log W}{\gamma m} \left(\prod_{i=1}^D \rho_i \ A\ _\sigma\right) \left(\sum_{i=1}^D \frac{\ A_i^\top - M_i^\top\ _{2,1}^{2/3}}{\ A\ _\sigma^{2/3}}\right)^{3/2}\right)$
Neyshabur <i>et al.</i> [176]	PAC-Bayes	$\mathcal{O}\left(\sqrt{\frac{B^2 D^2 W \log(DW) \prod_{i=1}^D \frac{\ W_i\ _F^2}{\ W_i\ _2^2} + \log \frac{Dm}{\delta}}{\gamma^2 m}}\right)$
Golowich <i>et al.</i> [92]	Rademacher complexity	$\mathcal{O}\left(\frac{B\sqrt{D} \prod_{j=1}^D M_F(j)}{\sqrt{m}}\right)$ or $\mathcal{O}\left(\frac{B\sqrt{D+\log m} \prod_{j=1}^D M_F(j)}{\sqrt{m}}\right)$
Dieuleveut <i>et al.</i> [65]	SGM on convex loss	$\mathcal{O}\left(m^{-\frac{2 \min\{\zeta, 1\}}{2 \min\{\zeta, 1\} + \gamma}}\right)$
Hardt <i>et al.</i> [101]	SGM on non-convex loss & algorithmic stability	$\frac{1+1/\beta c}{m-1} (2cL^2)^{\frac{1}{\beta c+1}} T^{\frac{1}{\beta c+1}}$
Raginsky <i>et al.</i> [194]	SGM on non-convex loss & algorithmic stability	$\tilde{\mathcal{O}}\left(\frac{\beta(\beta+U)^2}{\lambda_*} \delta^{1/4} \log\left(\frac{1}{\epsilon} + \epsilon\right) + \frac{(\beta+U)^2}{\lambda_* + m} + \frac{U \log(\beta+1)}{\beta}\right)$
Pensia <i>et al.</i> [190]	SGM on non-convex loss & information theory	$\tilde{\mathcal{O}}\left(\sqrt{\frac{R^2}{m} \sum_{t=1}^T \frac{\eta_t^2 L^2}{\sigma_t^2}}\right)$
Mou <i>et al.</i> [171]	SGM on non-convex loss & algorithmic stability	$\frac{2LC}{m} \left(\beta \sum_{t=1}^N \eta_t\right)^2$
Mou <i>et al.</i> [171]	SGM on non-convex loss & PAC-Bayes	$\mathcal{O}\left(\sqrt{\frac{\beta}{m} \sum_{k=1}^N \eta_k} e^{-\frac{\lambda}{3}(T_N - T_k)} \mathbb{E}[\ g_k(\theta_k)\ ^2]\right)$
Lever <i>et al.</i> [142]	Distribution-dependent prior & PAC-Bayes	$\frac{C^*}{Cm} \left(\lambda \sqrt{\frac{2}{m} \log \frac{2\xi(m)}{\delta}} + \frac{\gamma^2}{2m} + \log \frac{2}{\delta}\right)$
Li. <i>et al.</i> [145]	Distribution-dependent prior & Bayes-stability	$\mathcal{O}\left(\frac{C}{m} \sqrt{\mathbb{E}_S \left[\sum_{t=1}^T \frac{\eta_t^2}{\sigma_t^2} g_e(t)\right]}\right)$
Negrea <i>et al.</i> [175]	Data-dependent prior & information theory	$\sqrt{\frac{2\sigma^2}{m-n} I^{S_J, U}(\theta; S_J^c)} \leq \sqrt{\frac{2\sigma^2}{m-n} \mathbb{E}^{S_J, U}[\text{KL}(Q\ P)]}$

From our view, these six directions are promising in establishing the theoretical foundations of deep learning. However, the current “states of the art” are still far away from sufficient; all of them are facing major

difficulties. In this paper, we list three main directions for future works:

- The current measures for the complexity evaluate the whole hypothesis space, but can hardly capture the complexity of the “effective” hypothesis space that the training algorithm would output. Thus, a “local” version of complexity would be helpful; see an example [18].
- The current models for SGM usually ignore some of the popular training techniques, such as momentum, adaptive learning rate, gradient clipping, batch normalization, etc. Some works have been studying on these problems, but a full “landscape” is still missing. Also, the works are usually assuming the gradient noise is drawn from some specific distributions, such as Gaussian distributions or α -stable distributions.
- The geometry of the loss surface remains premature. The extreme complexity of neural networks makes simple explanations unaccessible. Efforts on understanding some sophisticated geometrical properties are demanded, such as the sharpness/flatness of the local minima, the potential categorization of the local minima (valley) according to their performance, and the volumes of the local minima from different categories.

References

- [1] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang. Deep learning with differential privacy. In *ACM SIGSAC Conference on Computer and Communications Security*, pages 308–318, 2016.
- [2] M. Acikkar and M. F. Akay. Support vector machines for predicting the admission decision of a candidate to the school of physical education and sports at cukurova university. *Expert Systems with Applications*, 36(3):7228–7233, 2009.
- [3] A. Agarwal, A. Beygelzimer, M. Dudík, J. Langford, and H. Wallach. A reductions approach to fair classification. *arXiv preprint arXiv:1803.02453*, 2018.
- [4] S. Ahn, A. Korattikara, and M. Welling. Bayesian posterior sampling via stochastic gradient Fisher scoring. *arXiv preprint arXiv:1206.6380*, 2012.
- [5] I. Aizenberg, N. N. Aizenberg, and J. P. Vandewalle. *Multi-Valued and Universal Binary Neurons: Theory, Learning and Applications*. Springer Science & Business Media, 2000.
- [6] Z. Allen-Zhu and Y. Li. Can SGD learn recurrent neural networks with provable generalization? In *Advances in Neural Information Processing Systems*, pages 10331–10341, 2019.
- [7] Z. Allen-Zhu, Y. Li, and Y. Liang. Learning and generalization in overparameterized neural networks, going beyond two layers. In *Advances in Neural Information Processing Systems*, pages 6158–6169, 2019.
- [8] Z. Allen-Zhu, Y. Li, and Z. Song. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine learning*, 2019.
- [9] M. Z. Alom, T. M. Taha, C. Yakopcic, S. Westberg, P. Sidike, M. S. Nasrin, M. Hasan, B. C. Van Essen, A. A. Awwal, and V. K. Asari. A state-of-the-art survey on deep learning theory and architectures. *Electronics*, 8(3):292, 2019.
- [10] A. Andoni, R. Panigrahy, G. Valiant, and L. Zhang. Learning polynomials with neural networks. In *International Conference on Machine learning*, pages 1908–1916, 2014.

- [11] M. Anthony and P. L. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, 2009.
- [12] S. Arora, S. S. Du, W. Hu, Z. Li, and R. Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. *arXiv preprint arXiv:1901.08584*, 2019.
- [13] D. Arpit, S. Jastrzebski, N. Ballas, D. Krueger, E. Bengio, M. S. Kanwal, T. Maharaj, A. Fischer, A. Courville, and Y. Bengio. A closer look at memorization in deep networks. In *International Conference on Machine learning*, 2017.
- [14] I. Attias, A. Kontorovich, and Y. Mansour. Improved generalization bounds for robust learning. In *Algorithmic Learning Theory*, pages 162–183, 2019.
- [15] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations*, 2015.
- [16] S. Baluja and I. Fischer. Learning to attack: Adversarial transformation networks. In *AAAI Conference on Artificial Intelligence*, volume 1, page 3, 2018.
- [17] P. Bartlett and J. Shawe-Taylor. Generalization performance of support vector machines and other pattern classifiers. *Advances in Kernel methodsSupport Vector Learning*, pages 43–54, 1999.
- [18] P. L. Bartlett, O. Bousquet, and S. Mendelson. Local rademacher complexities. *The Annals of Statistics*, 33(4):1497–1537, 2005.
- [19] P. L. Bartlett, D. J. Foster, and M. J. Telgarsky. Spectrally-normalized margin bounds for neural networks. In *Advances in Neural Information Processing Systems*, pages 6240–6249, 2017.
- [20] P. L. Bartlett, V. Maiorov, and R. Meir. Almost linear VC dimension bounds for piecewise polynomial networks. In *Advances in Neural Information Processing Systems*, pages 190–196, 1999.
- [21] P. L. Bartlett and S. Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.
- [22] P. L. Bartlett and R. C. Williamson. The VC dimension and pseudodimension of two-layer neural networks with discrete inputs. *Neural Computation*, 8(3):625–628, 1996.
- [23] E. B. Baum and D. Haussler. What size net gives valid generalization? In *Advances in Neural Information Processing Systems*, pages 81–90, 1989.
- [24] Y. Bengio, P. Lamblin, D. Popovici, and H. Larochelle. Greedy layer-wise training of deep networks. *Advances in Neural Information Processing Systems*, 19:153–160, 2006.
- [25] A. N. Bhagoji, D. Cullina, and P. Mittal. Lower bounds on adversarial robustness from optimal transport. In *Advances in Neural Information Processing Systems*, pages 7498–7510, 2019.
- [26] B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Šrndić, P. Laskov, G. Giacinto, and F. Roli. Evasion attacks against machine learning at test time. In *European Conference on Machine Learning*, 2013.
- [27] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. Variational inference: A review for statisticians. *Journal of the American statistical Association*, 112(518):859–877, 2017.
- [28] A. Blumer, A. Ehrenfeucht, D. Haussler, and M. K. Warmuth. Learnability and the Vapnik-Chervonenkis dimension. *Journal of the ACM*, 36(4):929–965, 1989.
- [29] O. Bousquet and A. Elisseeff. Stability and generalization. *Journal of Machine Learning Research*, 2(Mar):499–526, 2002.

- [30] T. Brennan, W. Dieterich, and B. Ehret. Evaluating the predictive validity of the compas risk and needs assessment system. *Criminal Justice and Behavior*, 36(1):21–40, 2009.
- [31] A. Brutzkus and A. Globerson. Why do larger models generalize better? a theoretical perspective via the xor problem. In *International Conference on Machine learning*, pages 822–830, 2019.
- [32] A. Brutzkus, A. Globerson, E. Malach, and S. Shalev-Shwartz. SGD learns over-parameterized networks that provably generalize on linearly separable data. In *International Conference on Learning Representations*, 2018.
- [33] Y. Bu, S. Zou, and V. V. Veeravalli. Tightening mutual information based bounds on generalization error. *IEEE Journal on Selected Areas in Information Theory*, 2020.
- [34] T. Bui, D. Hernández-Lobato, J. Hernandez-Lobato, Y. Li, and R. Turner. Deep Gaussian processes for regression using approximate expectation propagation. In *International Conference on Machine learning*, pages 1472–1481, 2016.
- [35] M. Bun and T. Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*, pages 635–658, 2016.
- [36] T. Calders, F. Kamiran, and M. Pechenizkiy. Building classifiers with independency constraints. In *IEEE International Conference on Data Mining Workshops*, pages 13–18, 2009.
- [37] F. Calmon, D. Wei, B. Vinzamuri, K. Natesan Ramamurthy, and K. R. Varshney. Optimized pre-processing for discrimination prevention. *Advances in Neural Information Processing Systems*, 30:3992–4001, 2017.
- [38] A. Canziani, A. Paszke, and E. Culurciello. An analysis of deep neural network models for practical applications. *arXiv preprint arXiv:1605.07678*, 2016.
- [39] Y. Cao and Q. Gu. Generalization bounds of stochastic gradient descent for wide and deep neural networks. In *Advances in Neural Information Processing Systems*, pages 10836–10846, 2019.
- [40] N. Carlini, C. Liu, J. Kos, Ú. Erlingsson, and D. Song. The secret sharer: Measuring unintended neural network memorization & extracting secrets. *arXiv preprint arXiv:1802.08232*, 2018.
- [41] P. Chaudhari and S. Soatto. Stochastic gradient descent performs variational inference, converges to limit cycles for deep networks. In *Information Theory and Applications Workshop*, 2018.
- [42] K. Chaudhuri, J. Imola, and A. Machanavajjhala. Capacity bounded differential privacy. *arXiv preprint arXiv:1907.02159*, 2019.
- [43] H. Chen, O. Engkvist, Y. Wang, M. Olivecrona, and T. Blaschke. The rise of deep learning in drug discovery. *Drug Discovery Today*, 23(6):1241–1250, 2018.
- [44] H. Chen, Z. Mo, Z. Yang, and X. Wang. Theoretical investigation of generalization bound for residual networks. In *International Joint Conference on Artificial Intelligence*, pages 2081–2087, 2019.
- [45] J. Chen, N. Kallus, X. Mao, G. Svacha, and M. Udell. Fairness under unawareness: Assessing disparity when protected class is unobserved. In *Conference on Fairness, Accountability, and Transparency*, pages 339–348, 2019.
- [46] J.-F. Chen and Q. H. Do. Training neural networks to predict student academic performance: A comparison of cuckoo search and gravitational search algorithms. *International Journal of Computational Intelligence and Applications*, 13(01):1450005, 2014.

- [47] L. Chen, Y. Min, M. Zhang, and A. Karbasi. More data can expand the generalization gap between adversarially robust and standard models. *arXiv preprint arXiv:2002.04725*, 2020.
- [48] M. Chen, X. Li, and T. Zhao. On generalization bounds of a family of recurrent neural networks. *arXiv preprint arXiv:1910.12947*, 2019.
- [49] T. Chen, E. Fox, and C. Guestrin. Stochastic gradient Hamiltonian Monte Carlo. In *International Conference on Machine learning*, pages 1683–1691, 2014.
- [50] X. Chen, J. D. Lee, X. T. Tong, and Y. Zhang. Statistical inference for model parameters in stochastic gradient descent. *arXiv preprint arXiv:1610.08637*, 2016.
- [51] Y. Chen, C. Jin, and B. Yu. Stability and convergence trade-off of iterative optimization algorithms. *arXiv preprint arXiv:1804.01619*, 2018.
- [52] Z. Chen, Y. Cao, D. Zou, and Q. Gu. How much over-parameterization is sufficient to learn deep ReLU networks? *arXiv preprint arXiv:1911.12360*, 2019.
- [53] J. J. Cherian, A. G. Taube, R. T. McGibbon, P. Angelikopoulos, G. Blanc, M. Snarski, D. D. Richman, J. L. Klepeis, and D. E. Shaw. Efficient hyperparameter optimization by way of pac-bayes bound minimization. *arXiv preprint arXiv:2008.06431*, 2020.
- [54] L. Chizat and F. Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. *Advances in Neural Information Processing Systems*, 31:3036–3046, 2018.
- [55] A. Choromanska, M. Henaff, M. Mathieu, G. B. Arous, and Y. LeCun. The loss surfaces of multilayer networks. In *International Conference on Artificial Intelligence and Statistics*, 2015.
- [56] T. Cohen and M. Welling. Group equivariant convolutional networks. In *International Conference on Machine learning*, pages 2990–2999, 2016.
- [57] T. S. Cohen, M. Geiger, and M. Weiler. A general theory of equivariant cnns on homogeneous spaces. In *Advances in Neural Information Processing Systems*, pages 9145–9156, 2019.
- [58] P. Cuff and L. Yu. Differential privacy as a mutual information constraint. In *ACM SIGSAC Conference on Computer and Communications Security*, pages 43–54, 2016.
- [59] D. Cullina, A. N. Bhagoji, and P. Mittal. PAC-learning in the presence of adversaries. In *Advances in Neural Information Processing Systems*, pages 230–241, 2018.
- [60] H. Dai, H. Li, T. Tian, X. Huang, L. Wang, J. Zhu, and L. Song. Adversarial attack on graph structured data. *arXiv preprint arXiv:1806.02371*, 2018.
- [61] A. Damianou and N. D. Lawrence. Deep Gaussian processes. In *Artificial intelligence and statistics*, pages 207–215, 2013.
- [62] R. Dechter. Learning while searching in constraint-satisfaction problems. 1986.
- [63] L. Deng and Y. Liu. *Deep learning in natural language processing*. Springer, 2018.
- [64] L. Deng and D. Yu. Deep learning: methods and applications. *Foundations and Trends in Signal Processing*, 7(3–4):197–387, 2014.
- [65] A. Dieuleveut and F. Bach. Nonparametric stochastic approximation with large step-sizes. *The Annals of Statistics*, 44(4):1363–1399, 2016.

- [66] N. Ding, Y. Fang, R. Babbush, C. Chen, R. D. Skeel, and H. Neven. Bayesian sampling using stochastic gradient thermostats. In *Advances in Neural Information Processing Systems*, pages 3203–3211, 2014.
- [67] L. Dinh, R. Pascanu, S. Bengio, and Y. Bengio. Sharp minima can generalize for deep nets. In *International Conference on Machine learning*, 2017.
- [68] F. Draxler, K. Veschgini, M. Salmhofer, and F. Hamprecht. Essentially no barriers in neural network energy landscape. In *International Conference on Machine learning*, 2018.
- [69] S. Du, J. Lee, H. Li, L. Wang, and X. Zhai. Gradient descent finds global minima of deep neural networks. In *International Conference on Machine learning*, pages 1675–1685, 2019.
- [70] S. S. Du, Y. Wang, X. Zhai, S. Balakrishnan, R. R. Salakhutdinov, and A. Singh. How many samples are needed to estimate a convolutional neural network? In *Advances in Neural Information Processing Systems*, pages 373–383, 2018.
- [71] S. S. Du, X. Zhai, B. Póczos, and A. Singh. Gradient descent provably optimizes over-parameterized neural networks. In *International Conference on Learning Representations*, 2019.
- [72] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth. Hybrid Monte Carlo. *Physics Letters B*, 195(2):216–222, 1987.
- [73] R. M. Dudley. The sizes of compact subsets of hilbert space and continuity of Gaussian processes. *Journal of Functional Analysis*, 1(3):290–330, 1967.
- [74] D. Duvenaud, O. Rippel, R. Adams, and Z. Ghahramani. Avoiding pathologies in very deep networks. In *Artificial Intelligence and Statistics*, pages 202–210, 2014.
- [75] C. Dwork, V. Feldman, M. Hardt, T. Pitassi, O. Reingold, and A. L. Roth. Preserving statistical validity in adaptive data analysis. In *Annual ACM Symposium on Theory of Computing*, pages 117–126, 2015.
- [76] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel. Fairness through awareness. In *Innovations in Theoretical Computer Science Conference*, pages 214–226, 2012.
- [77] C. Dwork and D. K. Mulligan. It’s not privacy, and it’s not fair. *Stanford Law Review Online*, 66:35, 2013.
- [78] C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- [79] C. Dwork and G. N. Rothblum. Concentrated differential privacy. *arXiv preprint arXiv:1603.01887*, 2016.
- [80] G. K. Dziugaite, K. Hsu, W. Gharbieh, and D. M. Roy. On the role of data in pac-bayes bounds. *arXiv preprint arXiv:2006.10929*, 2020.
- [81] M. C. E. Weinan, S. Wojtowytsch, and L. Wu. Towards a mathematical understanding of neural network-based machine learning: What we know and what we don’t. *arXiv preprint arXiv:2009.10713*, 2020.
- [82] F. A. Faber, A. Lindmaa, O. A. Von Lilienfeld, and R. Armiento. Machine learning energies of 2 million elpasolite (a b c 2 d 6) crystals. *Physical Review Letters*, 117(13):135502, 2016.
- [83] E. Faliagka, K. Ramantas, A. Tsakalidis, and G. Tzimas. Application of machine learning algorithms to an online recruitment system. In *International Conference on Internet and Web Applications and Services*, 2012.

- [84] M. Feldman, S. A. Friedler, J. Moeller, C. Scheidegger, and S. Venkatasubramanian. Certifying and removing disparate impact. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 259–268, 2015.
- [85] T. Garipov, P. Izmailov, D. Podoprikin, D. P. Vetrov, and A. G. Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns. In *Advances in Neural Information Processing Systems*, 2018.
- [86] J. Gehring, M. Auli, D. Grangier, D. Yarats, and Y. N. Dauphin. Convolutional sequence to sequence learning. In *International Conference on Machine learning*, pages 1243–1252, 2017.
- [87] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (6):721–741, 1984.
- [88] J. Geumlek, S. Song, and K. Chaudhuri. Renyi differential privacy mechanisms for posterior sampling. In *Advances in Neural Information Processing Systems*, pages 5289–5298, 2017.
- [89] J. Gilmer, R. P. Adams, I. Goodfellow, D. Andersen, and G. E. Dahl. Motivating the rules of the game for adversarial example research. *arXiv preprint arXiv:1807.06732*, 2018.
- [90] P. W. Goldberg and M. R. Jerrum. Bounding the Vapnik-Chervonenkis dimension of concept classes parameterized by real numbers. *Machine Learning*, 18(2-3):131–148, 1995.
- [91] M. Goldblum, J. Geiping, A. Schwarzschild, M. Moeller, and T. Goldstein. Truth or backpropaganda? An empirical investigation of deep learning theory. In *International Conference on Learning Representations*, 2020.
- [92] N. Golowich, A. Rakhlin, and O. Shamir. Size-independent sample complexity of neural networks. In *Annual Conference on Learning Theory*, pages 297–299, 2018.
- [93] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014.
- [94] I. J. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [95] P. Goyal, P. Dollár, R. Girshick, P. Noordhuis, L. Wesolowski, A. Kyrola, A. Tulloch, Y. Jia, and K. He. Accurate, large minibatch sgd: training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017.
- [96] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *International Conference on Machine learning*, pages 369–376, 2006.
- [97] A. Graves, A.-r. Mohamed, and G. Hinton. Speech recognition with deep recurrent neural networks. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 6645–6649, 2013.
- [98] Y. Guo, H. Wang, Q. Hu, H. Liu, L. Liu, and M. Bennamoun. Deep learning for 3d point clouds: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [99] M. Haghifam, J. Negrea, A. Khisti, D. M. Roy, and G. K. Dziugaite. Sharpened generalization bounds based on conditional mutual information and an application to noisy, iterative algorithms. *arXiv preprint arXiv:2004.12983*, 2020.

- [100] M. Hardt, E. Price, and N. Srebro. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems*, pages 3315–3323, 2016.
- [101] M. Hardt, B. Recht, and Y. Singer. Train faster, generalize better: Stability of stochastic gradient descent. In *International Conference on Machine learning*, pages 1225–1234, 2016.
- [102] N. Harvey, C. Liaw, and A. Mehrabian. Nearly-tight VC-dimension bounds for piecewise linear neural networks. In *Annual Conference on Learning Theory*, pages 1064–1068, 2017.
- [103] W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. 1970.
- [104] D. Haussler. Sphere packing numbers for subsets of the boolean n -cube with bounded Vapnik-Chervonenkis dimension. *Journal of Combinatorial Theory, Series A*, 69(2):217–232, 1995.
- [105] T. Hazan and T. Jaakkola. Steps toward deep kernel methods from infinite neural networks. *arXiv preprint arXiv:1508.05133*, 2015.
- [106] F. He, T. Liu, and D. Tao. Control batch size and learning rate to generalize well: Theoretical and empirical evidence. In *Advances in Neural Information Processing Systems*, 2019.
- [107] F. He, T. Liu, and D. Tao. Why resnet works? residuals generalize. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- [108] F. He, B. Wang, and D. Tao. Piecewise linear activations substantially shape the loss surfaces of neural networks. In *International Conference on Learning Representations*, 2020.
- [109] F. He, B. Wang, and D. Tao. Tighter generalization bounds for iterative differentially private learning algorithms. *arXiv preprint arXiv:2007.09371*, 2020.
- [110] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [111] D. O. Hebb. *The Organization of Behavior: A Neuropsychological Theory*. J. Wiley; Chapman & Hall, 1949.
- [112] F. Hellström and G. Durisi. Generalization bounds via information density and conditional information density. *arXiv preprint arXiv:2005.08044*, 2020.
- [113] J. Hensman and N. D. Lawrence. Nested variational compression in deep Gaussian processes. *arXiv preprint arXiv:1412.1370*, 2014.
- [114] G. E. Hinton, S. Osindero, and Y.-W. Teh. A fast learning algorithm for deep belief nets. *Neural Computation*, 18(7):1527–1554, 2006.
- [115] M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 14(1):1303–1347, 2013.
- [116] A. Jacot, F. Gabriel, and C. Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in Neural Information Processing Systems*, pages 8571–8580, 2018.
- [117] S. Jastrzębski, Z. Kenton, D. Arpit, N. Ballas, A. Fischer, Y. Bengio, and A. Storkey. Three factors influencing minima in sgd. *arXiv preprint arXiv:1711.04623*, 2017.
- [118] K. Jia, D. Tao, S. Gao, and X. Xu. Improving training of deep neural networks via singular value bounding. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 4344–4352, 2017.

- [119] F. Kamiran and T. Calders. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems*, 33(1):1–33, 2012.
- [120] T. Kamishima, S. Akaho, H. Asoh, and J. Sakuma. Fairness-aware classifier with prejudice remover regularizer. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 35–50, 2012.
- [121] K. Kawaguchi. Deep learning without poor local minima. In *Advances in Neural Information Processing Systems*, 2016.
- [122] K. Kawaguchi, L. P. Kaelbling, and Y. Bengio. Generalization in deep learning. *arXiv:1710.05468*, 2017.
- [123] M. Kearns, S. Neel, A. Roth, and Z. S. Wu. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *International Conference on Machine learning*, pages 2564–2572, 2018.
- [124] N. S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, and P. T. P. Tang. On large-batch training for deep learning: Generalization gap and sharp minima. In *International Conference on Learning Representations*, 2017.
- [125] A. E. Khandani, A. J. Kim, and A. W. Lo. Consumer credit-risk models via machine-learning algorithms. *Journal of Banking & Finance*, 34(11):2767–2787, 2010.
- [126] J. Khim and P.-L. Loh. Adversarial risk bounds via function transformation. *arXiv preprint arXiv:1810.09519*, 2018.
- [127] D. P. Kingma and M. Welling. Auto-encoding variational Bayes. In *International Conference on Learning Representations*, 2014.
- [128] I. Kobyzev, S. Prince, and M. Brubaker. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [129] V. Koltchinskii. Rademacher penalties and structural risk minimization. *IEEE Transactions on Information Theory*, 47(5):1902–1914, 2001.
- [130] V. Koltchinskii and D. Panchenko. Rademacher processes and bounding the risk of function learning. In *High Dimensional Probability II*, pages 443–457. Springer, 2000.
- [131] V. Koltchinskii and D. Panchenko. Empirical margin distributions and bounding the generalization error of combined classifiers. *The Annals of Statistics*, 30(1):1–50, 2002.
- [132] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 1097–1105, 2012.
- [133] D. Krueger, N. Ballas, S. Jastrzebski, D. Arpit, M. S. Kanwal, T. Maharaj, E. Bengio, A. Fischer, and A. Courville. Deep nets don’t learn via memorization. 2017.
- [134] R. Kuditipudi, X. Wang, H. Lee, Y. Zhang, Z. Li, W. Hu, S. Arora, and R. Ge. Explaining landscape connectivity of low-cost solutions for multilayer nets. In *Advances in Neural Information Processing Systems*, 2019.
- [135] C. A. Kulikowski. Artificial intelligence methods and systems for medical consultation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (5):464–476, 1980.

- [136] H. Larochelle and S. Lauly. A neural autoregressive topic model. In *Advances in Neural Information Processing Systems*, pages 2708–2716, 2012.
- [137] N. D. Lawrence and A. J. Moore. Hierarchical Gaussian process latent variable models. In *International Conference on Machine learning*, pages 481–488, 2007.
- [138] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436, 2015.
- [139] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [140] J. Lee, Y. Bahri, R. Novak, S. S. Schoenholz, J. Pennington, and J. Sohl-Dickstein. Deep neural networks as Gaussian processes. In *International Conference on Learning Representations*, 2018.
- [141] J. Lee, L. Xiao, S. Schoenholz, Y. Bahri, R. Novak, J. Sohl-Dickstein, and J. Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. In *Advances in Neural Information Processing Systems*, pages 8572–8583, 2019.
- [142] G. Lever, F. Laviolette, and J. Shawe-Taylor. Tighter pac-bayes bounds through distribution-dependent priors. *Theoretical Computer Science*, 473:4–28, 2013.
- [143] B. Li, C. Chen, W. Wang, and L. Carin. Second-order adversarial attack and certifiable robustness. 2018.
- [144] D. Li, T. Ding, and R. Sun. Over-parameterized deep neural networks have no strict local minima for any continuous activations. *arXiv preprint arXiv:1812.11039*, 2018.
- [145] J. Li, X. Luo, and M. Qiao. On generalization error bounds of noisy gradient methods for non-convex learning. *arXiv preprint arXiv:1902.00621*, 2019.
- [146] Y. Li, R. Bu, M. Sun, W. Wu, X. Di, and B. Chen. Pointcnn: Convolution on x-transformed points. In *Advances in Neural Information Processing Systems*, pages 820–830, 2018.
- [147] Y. Li and Y. Yuan. Convergence analysis of two-layer neural networks with ReLU activation. In *Advances in Neural Information Processing Systems*, pages 597–607, 2017.
- [148] J. Liao, L. Sankar, V. Y. Tan, and F. du Pin Calmon. Hypothesis testing under mutual information privacy constraints in the high privacy regime. *IEEE Transactions on Information Forensics and Security*, 13(4):1058–1071, 2017.
- [149] J. Lin, R. Camoriano, and L. Rosasco. Generalization properties and implicit regularization for multiple passes SGM. In *International Conference on Machine learning*, pages 2340–2348, 2016.
- [150] J. Lin and L. Rosasco. Optimal learning for multi-pass stochastic gradient methods. In *Advances in Neural Information Processing Systems*, pages 4556–4564, 2016.
- [151] S. Lin and J. Zhang. Generalization bounds for convolutional neural networks. *arXiv preprint arXiv:1910.01487*, 2019.
- [152] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. A. Van Der Laak, B. Van Ginneken, and C. I. Sánchez. A survey on deep learning in medical image analysis. *Medical Image Analysis*, 42:60–88, 2017.
- [153] S. Lohr. Facial recognition is accurate, if youre a white guy. *New York Times*, 9, 2018.
- [154] C. Louizos, K. Swersky, Y. Li, M. Welling, and R. Zemel. The variational fair autoencoder. *arXiv preprint arXiv:1511.00830*, 2015.

- [155] C. Louizos and M. Welling. Multiplicative normalizing flows for variational Bayesian neural networks. In *International Conference on Machine learning*, pages 2218–2227, 2017.
- [156] H. Lu and K. Kawaguchi. Depth creates no bad local minima. *arXiv preprint arXiv:1702.08580*, 2017.
- [157] Y.-A. Ma, T. Chen, and E. Fox. A complete recipe for stochastic gradient mcmc. In *Advances in Neural Information Processing Systems*, pages 2917–2925, 2015.
- [158] W. Maass. Bounds for the computational power and learning complexity of analog neural nets. *SIAM Journal on Computing*, 26(3):708–732, 1997.
- [159] S. Mandt, M. D. Hoffman, and D. M. Blei. Stochastic gradient descent as approximate Bayesian inference. *Journal of Machine Learning Research*, 18(1):4873–4907, 2017.
- [160] D. A. McAllester. PAC-Bayesian model averaging. In *Annual Conference of Learning Theory*, volume 99, pages 164–170, 1999.
- [161] D. A. McAllester. Some PAC-Bayesian theorems. *Machine Learning*, 37(3):355–363, 1999.
- [162] W. S. McCulloch and W. Pitts. A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, 5(4):115–133, 1943.
- [163] S. Mei, Y. Bai, and A. Montanari. The landscape of empirical risk for nonconvex losses. *The Annals of Statistics*, 46(6A):2747–2774, 2018.
- [164] A. K. Menon and R. C. Williamson. The cost of fairness in binary classification. In *Conference on Fairness, Accountability and Transparency*, pages 107–118, 2018.
- [165] S. Merity, N. S. Keskar, and R. Socher. Regularizing and optimizing lstm language models. In *International Conference on Learning Representations*, 2018.
- [166] Y. Min, L. Chen, and A. Karbasi. The curious case of adversarially robust models: More data can help, double descend, or hurt generalization. *arXiv preprint arXiv:2002.11080*, 2020.
- [167] I. Mironov. Rényi differential privacy. In *IEEE Computer Security Foundations Symposium*, pages 263–275, 2017.
- [168] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida. Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations*, 2018.
- [169] M. Mohri, A. Rostamizadeh, and A. Talwalkar. *Foundations of machine learning*. MIT press, 2018.
- [170] O. Montasser, S. Hanneke, and N. Srebro. VC classes are adversarially robustly learnable, but only improperly. In *Conference on Learning Theory*, pages 2512–2530, 2019.
- [171] W. Mou, L. Wang, X. Zhai, and K. Zheng. Generalization bounds of sgld for non-convex learning: Two theoretical viewpoints. In *Annual Conference On Learning Theory*, 2018.
- [172] M. T. Musavi, K. H. Chan, D. M. Hummels, and K. Kalantri. On the generalization ability of neural network classifiers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(6):659–663, 1994.
- [173] R. M. Neal. *Bayesian LEARNING FOR NEURAL NETWORKS*. PhD thesis, University of Toronto, 1995.

- [174] R. M. Neal. Priors for infinite networks. In *Bayesian Learning for Neural Networks*, pages 29–53. Springer, 1996.
- [175] J. Negrea, M. Haghifam, G. K. Dziugaite, A. Khisti, and D. M. Roy. Information-theoretic generalization bounds for sgld via data-dependent estimates. In *Advances in Neural Information Processing Systems*, pages 11015–11025, 2019.
- [176] B. Neyshabur, S. Bhojanapalli, and N. Srebro. A PAC-Bayesian approach to spectrally-normalized margin bounds for neural networks. *arXiv preprint arXiv:1707.09564*, 2017.
- [177] B. Neyshabur, R. Tomioka, and N. Srebro. Norm-based capacity control in neural networks. In *Conference on Learning Theory*, pages 1376–1401, 2015.
- [178] A. Nguyen, J. Yosinski, and J. Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 427–436, 2015.
- [179] Q. Nguyen. On connected sublevel sets in deep learning. In *International Conference on Machine learning*, 2019.
- [180] Q. Nguyen, M. C. Mukkamala, and M. Hein. On the loss landscape of a class of deep neural networks with no bad local valleys. In *International Conference on Learning Representations*, 2019.
- [181] K. Nissim and U. Stemmer. On the generalization properties of differential privacy. *CoRR*, abs/1504.05800, 2015.
- [182] M. Ntampaka, H. Trac, D. J. Sutherland, S. Fromenteau, B. Póczos, and J. Schneider. Dynamical mass measurements of contaminated galaxy clusters using machine learning. *The Astrophysical Journal*, 831(2):135, 2016.
- [183] L. Oneto, S. Ridella, and D. Anguita. Differential privacy and generalization: Sharper bounds with applications. *Pattern Recognition Letters*, 89:31–38, 2017.
- [184] D. W. Otter, J. R. Medina, and J. K. Kalita. A survey of the usages of deep learning for natural language processing. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- [185] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami. Practical black-box attacks against machine learning. In *ACM on Asia Conference on Computer and Communications Security*, pages 506–519, 2017.
- [186] N. Papernot, P. McDaniel, S. Jha, M. Fredrikson, Z. B. Celik, and A. Swami. The limitations of deep learning in adversarial settings. In *IEEE European Symposium on Security and Privacy*, pages 372–387, 2016.
- [187] N. Papernot, P. McDaniel, A. Sinha, and M. Wellman. Towards the science of security and privacy in machine learning. *arXiv preprint arXiv:1611.03814*, 2016.
- [188] V. Pappas. The full spectrum of deep net Hessians at scale: Dynamics with sample size. *arXiv preprint arXiv:1811.07062*, 2018.
- [189] S. Patterson and Y. W. Teh. Stochastic gradient Riemannian Langevin dynamics on the probability simplex. In *Advances in Neural Information Processing Systems*, pages 3102–3110, 2013.
- [190] A. Pensia, V. Jog, and P.-L. Loh. Generalization error bounds for noisy, iterative algorithms. In *IEEE International Symposium on Information Theory*, 2018.

- [191] M. E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer. Deep contextualized word representations. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 2227–2237, 2018.
- [192] M. S. Pydi and V. Jog. Adversarial risk via optimal transport and optimal couplings. In *International Conference on Machine learning*, pages 7814–7823, 2020.
- [193] A. Raghunathan, J. Steinhardt, and P. Liang. Certified defenses against adversarial examples. In *International Conference on Learning Representations*, 2018.
- [194] M. Raginsky, A. Rakhlin, and M. Telgarsky. Non-convex learning via stochastic gradient Langevin dynamics: A nonasymptotic analysis. In *Conference on Learning Theory*, pages 1674–1703, 2017.
- [195] T. Ramadorai, A. Fuster, P. Goldsmith-Pinkham, and A. Walther. Predictably unequal? the effects of machine learning on credit markets. 2017.
- [196] M. Ranzato, C. Poultney, S. Chopra, and Y. Cun. Efficient learning of sparse representations with an energy-based model. *Advances in Neural Information Processing Systems*, 19:1137–1144, 2006.
- [197] S. Ravanbakhsh, J. Oliva, S. Fromenteau, L. Price, S. Ho, J. Schneider, and B. Póczos. Estimating cosmological parameters from the dark matter distribution. In *International Conference on Machine learning*, pages 2407–2416, 2016.
- [198] H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, pages 400–407, 1951.
- [199] F. Rosenblatt. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386, 1958.
- [200] W. Roth and F. Pernkopf. Bayesian neural networks with weight sharing using Dirichlet processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(1):246–252, 2018.
- [201] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986.
- [202] I. Safran and O. Shamir. Spurious local minima are common in two-layer ReLU neural networks. In *International Conference on Machine learning*, 2018.
- [203] L. Sagun, L. Bottou, and Y. LeCun. Singularity of the hessian in deep learning. *arXiv preprint arXiv:1611.07476*, 2016.
- [204] L. Sagun, U. Evci, V. U. Guney, Y. Dauphin, and L. Bottou. Empirical analysis of the hessian of over-parametrized neural networks. In *International Conference on Learning Representations Workshop*, 2018.
- [205] A. Sannai and M. Imaizumi. Improved generalization bound of permutation invariant deep neural networks. *arXiv preprint arXiv:1910.06552*, 2019.
- [206] A. Sannai, Y. Takai, and M. Cordonnier. Universal approximations of permutation invariant/equivariant functions by deep neural networks. *arXiv preprint arXiv:1903.01939*, 2019.
- [207] R. E. Schapire, Y. Freund, P. Bartlett, and W. S. Lee. Boosting the margin: A new explanation for the effectiveness of voting methods. *The Annals of Statistics*, 26(5):1651–1686, 1998.
- [208] L. Schmidt, S. Santurkar, D. Tsipras, K. Talwar, and A. Madry. Adversarially robust generalization requires more data. *Advances in Neural Information Processing Systems*, 31:5014–5026, 2018.

- [209] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, and M. Lanctot. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484, 2016.
- [210] S. L. Smith, P.-J. Kindermans, C. Ying, and Q. V. Le. Don’t decay the learning rate, increase the batch size. In *International Conference on Learning Representations*, 2018.
- [211] S. L. Smith and Q. V. Le. A Bayesian perspective on generalization and stochastic gradient descent. In *International Conference on Learning Representations*, 2018.
- [212] D. Soudry and E. Hoffer. Exponentially vanishing sub-optimal local minima in multilayer neural networks. In *International Conference on Learning Representations Workshop*, 2018.
- [213] T. Steinke and L. Zakynthinou. Reasoning about generalization via conditional mutual information. *arXiv preprint arXiv:2001.09122*, 2020.
- [214] R. Sun. Optimization for deep learning: Theory and algorithms. *arXiv preprint arXiv:1912.08957*, 2019.
- [215] Z. Sun, G. Bebis, and R. Miller. On-road vehicle detection: A review. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(5):694–711, 2006.
- [216] I. Sutskever, O. Vinyals, and Q. V. Le. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems*, pages 3104–3112, 2014.
- [217] G. Swirszcz, W. M. Czarnecki, and R. Pascanu. Local minima in training of deep networks. *arXiv preprint arXiv:1611.06310*, 2016.
- [218] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations*, 2014.
- [219] B. Taskar, C. Guestrin, and D. Koller. Max-margin Markov networks. In *Advances in Neural Information Processing Systems*, pages 25–32, 2004.
- [220] Z. Tu, F. He, and D. Tao. Understanding generalization in recurrent neural networks. In *International Conference on Learning Representations*, 2020.
- [221] Z. Tu, J. Zhang, and D. Tao. Theoretical analysis of adversarial learning: A minimax approach. In *Advances in Neural Information Processing Systems*, pages 12280–12290, 2019.
- [222] B. Tzen, T. Liang, and M. Raginsky. Local optimality and generalization guarantees for the Langevin algorithm via empirical metastability. In *Conference On Learning Theory*, pages 857–875, 2018.
- [223] V. Vapnik. *Estimation of Dependences based on Empirical Data*. Springer Science & Business Media, 2006.
- [224] V. Vapnik. *The Nature of Statistical Learning Theory*. Springer Science & Business Media, 2013.
- [225] W. Wang, L. Ying, and J. Zhang. On the relation between identifiability, differential privacy, and mutual-information privacy. *IEEE Transactions on Information Theory*, 62(9):5018–5029, 2016.
- [226] C. Wei, J. D. Lee, Q. Liu, and T. Ma. Regularization matters: Generalization and optimization of neural nets vs their induced kernel. In *Advances in Neural Information Processing Systems*, pages 9712–9724, 2019.

- [227] Y. Wei, F. Yang, and M. J. Wainwright. Early stopping for kernel boosting algorithms: A general analysis with localized complexities. In *Advances in Neural Information Processing Systems*, pages 6065–6075, 2017.
- [228] E. Weinan. A proposal on machine learning via dynamical systems. *Communications in Mathematics and Statistics*, 5(1):1–11, 2017.
- [229] M. Welling and Y. W. Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *International Conference on Machine learning*, pages 681–688, 2011.
- [230] Y. Wen, K. Luk, M. Gazeau, G. Zhang, H. Chan, and J. Ba. Interplay between optimization and generalization of stochastic gradient descent with covariance noise. *arXiv preprint arXiv:1902.08234*, 2019.
- [231] C. Williams. Computing with infinite networks. *Advances in Neural Information Processing Systems*, 9:295–301, 1996.
- [232] I. H. Witten, E. Frank, M. A. Hall, and C. J. Pal. *Data Mining: Practical machine learning tools and techniques*. Morgan Kaufmann, 2016.
- [233] B. Woodworth, S. Gunasekar, M. I. Ohannessian, and N. Srebro. Learning non-discriminatory predictors. *arXiv preprint arXiv:1702.06081*, 2017.
- [234] C. Xiang. *The Interplay between Sampling and Optimization*. PhD thesis, EECS Department, University of California, Berkeley, 2020.
- [235] A. Xu and M. Raginsky. Information-theoretic analysis of generalization capability of learning algorithms. In *Advances in Neural Information Processing Systems*, pages 2524–2533, 2017.
- [236] H. Xu, C. Caramanis, and S. Mannor. Sparse algorithms are not stable: A no-free-lunch theorem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(1):187–193, 2011.
- [237] D. Yin, R. Kannan, and P. Bartlett. Rademacher complexity for adversarially robust generalization. In *International Conference on Machine learning*, pages 7085–7094, 2019.
- [238] Y. Ying and M. Pontil. Online gradient descent learning algorithms. *Foundations of Computational Mathematics*, 8(5):561–596, 2008.
- [239] A. W. Yu, D. Dohan, M.-T. Luong, R. Zhao, K. Chen, M. Norouzi, and Q. V. Le. Qanet: Combining local convolution with global self-attention for reading comprehension. In *International Conference on Learning Representations*, 2018.
- [240] C. Yun, S. Sra, and A. Jadbabaie. Small nonlinearities in activation functions create bad local minima in neural networks. In *International Conference on Learning Representations*, 2019.
- [241] M. B. Zafar, I. Valera, M. G. Rodriguez, and K. P. Gummadi. Fairness constraints: Mechanisms for fair classification. In *Artificial Intelligence and Statistics*, pages 962–970, 2017.
- [242] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Póczos, R. Salakhutdinov, and A. J. Smola. Deep sets. In *Advances in Neural Information Processing Systems*, 2017.
- [243] R. Zemel, Y. Wu, K. Swersky, T. Pitassi, and C. Dwork. Learning fair representations. In *International Conference on Machine Learning*, pages 325–333, 2013.
- [244] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017.

- [245] C. Zhang, J. B  tepage, H. Kjellstr  m, and S. Mandt. Advances in variational inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8):2008–2026, 2018.
- [246] H. Zhang, B. Chen, Y. Cong, D. Guo, H. Liu, and M. Zhou. Deep autoencoding topic model with scalable hybrid Bayesian inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [247] H. Zhang, J. Shao, and R. Salakhutdinov. Deep neural networks with multi-branch architectures are intrinsically less non-convex. In *International Conference on Artificial Intelligence and Statistics*, 2019.
- [248] H. Zhang, Y. Yu, J. Jiao, E. P. Xing, L. E. Ghaoui, and M. I. Jordan. Theoretically principled trade-off between robustness and accuracy. *arXiv preprint arXiv:1901.08573*, 2019.
- [249] J. Zhang, T. Liu, and D. Tao. An information-theoretic view for deep learning. *arXiv preprint arXiv:1804.09060*, 2018.
- [250] T. Zheng, C. Chen, and K. Ren. Distributionally adversarial attack. In *AAAI Conference on Artificial Intelligence*, volume 33, pages 2253–2260, 2019.
- [251] P. Zhou and J. Feng. Empirical risk landscape analysis for understanding deep neural networks. In *International Conference on Learning Representations*, 2018.
- [252] P. Zhou and J. Feng. Understanding generalization and optimization performance of deep cnns. In *International Conference on Machine learning*, pages 5960–5969, 2018.
- [253] Y. Zhou and Y. Liang. Critical points of neural networks: Analytical forms and landscape properties. In *International Conference on Learning Representations*, 2018.
- [254] Y. Zhou, Y. Liang, and H. Zhang. Generalization error bounds with probabilistic guarantee for SGD in nonconvex optimization. *arXiv preprint arXiv:1802.06903*, 2018.
- [255] D. Zou, Y. Cao, D. Zhou, and Q. Gu. Gradient descent optimizes over-parameterized deep ReLU networks. *Machine Learning*, 109(3):467–492, 2020.