

Learning rates for stochastic gradient descent with nonconvex objectives

Lei, Yunwen; Tang, Ke

DOI:

[10.1109/TPAMI.2021.3068154](https://doi.org/10.1109/TPAMI.2021.3068154)

License:

Other (please specify with Rights Statement)

Document Version

Peer reviewed version

Citation for published version (Harvard):

Lei, Y & Tang, K 2021, 'Learning rates for stochastic gradient descent with nonconvex objectives', *IEEE Transactions on Pattern Analysis and Machine Intelligence*. <https://doi.org/10.1109/TPAMI.2021.3068154>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

Y. Lei and K. Tang, "Learning Rates for Stochastic Gradient Descent with Nonconvex Objectives," in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, doi: 10.1109/TPAMI.2021.3068154.

© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Learning Rates for Stochastic Gradient Descent with Nonconvex Objectives

Yunwen Lei and Ke Tang, *Senior Member, IEEE*

Abstract—Stochastic gradient descent (SGD) has become the method of choice for training highly complex and nonconvex models since it can not only recover good solutions to minimize training errors but also generalize well. Computational and statistical properties are separately studied to understand the behavior of SGD in the literature. However, there is a lacking study to jointly consider the computational and statistical properties in a nonconvex learning setting. In this paper, we develop novel learning rates of SGD for nonconvex learning by presenting *high-probability* bounds for both computational and statistical errors. We show that the complexity of SGD iterates grows in a controllable manner with respect to the iteration number, which sheds insights on how an implicit regularization can be achieved by tuning the number of passes to balance the computational and statistical errors. As a byproduct, we also slightly refine the existing studies on the uniform convergence of gradients by showing its connection to Rademacher chaos complexities.

Index Terms—Stochastic Gradient Descent, Learning Rates, Nonconvex Optimization, Early Stopping

1 INTRODUCTION

STOCHASTIC algorithms such as stochastic gradient descent (SGD) have become one of the workhorses behind many machine learning tasks [1–3]. As an iterative algorithm, SGD iteratively moves models along the reverse direction of an unbiased gradient estimate built on one or several training examples. Despite its simplicity, SGD has found great success in training highly complex and nonconvex models by the ability to identify good solutions to minimize training errors [3–6]. Surprisingly, the models trained in this way also generalize well to testing examples even in the case with significantly more parameters than training examples [7, 8], which can not be explained well by the traditional generalization analysis based only on the complexity of models [9, 10]. This is especially the case for nonconvex learning problems: while there are multiple local minima for empirical risks, SGD is prone to identify one with good generalization ability. With the recent extensive studies, it is gradually realized that the generalization performance of models depends not only on the complexity of models but also on the optimization algorithms used to train models [7]. In this spirit, there is a growing interest on the theoretical work to understand the success of SGD by considering either its computational or statistical properties.

Roughly speaking, the computational property is related to how the learning algorithm minimizes an empirical risk, while the statistical property is related to the gap between the empirical and population risks on the trained model. It is the interaction of these two factors that determines the learning performance of the trained model [5]. Therefore, it is necessary to take both factors into account for a full understanding on the behavior of the algorithm, which, however, is lacking in the nonconvex setting. Existing theoretical results on SGD in a nonconvex setting are mainly studied for computational errors [11–14]. However, our primary interests in machine learning is the generalization behavior of the trained model on testing examples, which may differ largely from the empirical behavior on training examples due to the possible noises in the data. Motivated by this observation, the generalization gap between empirical and population risks is studied by analyzing the sensitivity of SGD to a small perturbation of the training set, in the framework of algorithmic stability [15–17]. However, the step sizes there need to be very small to enjoy a good stability in the nonconvex setting, which requires an exponential number of iterations for a moderate decay of computational errors. In this way, the resulting learning rates are not quite satisfactory. Very recently, the generalization gap of nonconvex learning is studied via the uniform convergence for the gradients of empirical risks to their population counterparts [18–20]. However, these discussions ignore the interaction between the corresponding training algorithm and the training examples to produce the trained model.

In this paper, we study the learning performance of SGD for nonconvex learning problems from a joint perspective of computational and statistical properties. For the computational properties, we provide a high-probability bound on the decay rate for the gradient of empirical risks (Lemma 4). For the statistical properties, we show that the complexity of models grows in a controllable manner along the optimization process by presenting a high-probability bound on the norm of SGD iterates (Lemma 9). This joint

This work was supported in part by the Guangdong Provincial Key Laboratory under Grant 2020B121201001, the Program for Guangdong Introducing Innovative and Entrepreneurial Teams under Grant 2017ZT07X386, the National Natural Science Foundation of China under Grant 61806091, the Science and Technology Commission of Shanghai Municipality under Grant 19511120600, and the MOE University Scientific-Technological Innovation Plan Program. (Corresponding author: Ke Tang.)

Yunwen Lei was with the Department of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen 518055, China. He is now with the School of Computer Science, University of Birmingham, Birmingham B15 2TT, United Kingdom (e-mail: y.lei@bham.ac.uk).

Ke Tang is with the Guangdong Key Laboratory of Brain-Inspired Intelligent Computation, Department of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen 518055, China, and also with the Research Institute of Trustworthy Autonomous Systems, Southern University of Science and Technology, Shenzhen 518055, China (e-mail: tangk3@sustech.edu.cn).

perspective sheds intuitive insights on how the computational and statistical errors should be balanced by choosing an optimal iteration number to achieve a good learning rate. As compared to the existing stability-based approach requiring a fast-decaying step-size sequence, our statistical errors are controlled for a general polynomially-decaying step-size sequence for which the computational errors would decay significantly faster. We consider both general nonconvex and gradient-dominated objective functions, for each of which our analysis can achieve, to our best knowledge, the first learning rates that can match the associated ones in the convex setting. As a byproduct, we also derive slightly better uniform convergence rates between gradients of empirical risks and the corresponding population risks by employing the tool of Rademacher chaos complexities [21].

The paper is structured as follows. We formulate the problem in Section 2, and study uniform convergence of gradients in Section 3. We present learning rates in Section 4, and discuss related work in Section 5. We report experimental results in Section 6, and conclude the paper in Section 7. All the proofs can be found in the Supplementary Material.

2 PROBLEM FORMULATION

Let ρ be a probability measure defined over a compact sample space $\mathcal{Z} = \mathcal{X} \times \mathcal{Y} \subset \mathbb{R}^d \times \mathbb{R}$, where $d \in \mathbb{N}$ is the input dimension. We aim to learn a prediction rule from the input space \mathcal{X} to the output space \mathcal{Y} in a hypothesis space indexed by $\mathcal{W} \subseteq \mathbb{R}^d$. The error of a model at a single example z can be quantified by a loss function $f : \mathcal{W} \times \mathcal{Z} \mapsto \mathbb{R}_+$, from which we can define the population risk (generalization error) by $F : \mathcal{W} \mapsto \mathbb{R}_+$ to measure the performance of a model on testing examples. That is, $F(\mathbf{w}) = \mathbb{E}_z[f(\mathbf{w}; z)]$, where \mathbb{E}_z denotes the expectation with respect to (w.r.t.) the random variable z drawn from ρ . The population risk is not accessible since the probability measure ρ is unknown. In practice, we often draw a training dataset $S = \{z_1, \dots, z_n\}$ independently from ρ and construct an empirical approximation of F by $F_S(\mathbf{w}) = \frac{1}{n} \sum_{i=1}^n f(\mathbf{w}; z_i)$. The empirical risk F_S takes a finite-sum structure, for which an efficient method to use this structure is SGD. Let $\mathbf{w}_1 = 0$ and $\{\eta_t\}_{t \in \mathbb{N}}$ be a sequence of positive step sizes. At the t -th iteration, we first draw an index j_t from the uniform distribution over $\{1, \dots, n\}$, and update the model by

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta_t \nabla f(\mathbf{w}_t; z_{j_t}), \quad (1)$$

where $\nabla f(\mathbf{w}_t; z_{j_t})$ denotes the gradient of f w.r.t. the first argument. It should be noted that the randomness of the iterate \mathbf{w}_t comes from two sources: one from the sampling of training examples according to ρ and one from the sampling of the indices $\{j_t\}_t$ according to the uniform distribution over $\{1, \dots, n\}$. In this paper, we are interested in the generalization behavior of $\{\mathbf{w}_t\}$ measured by the gradient of population risks $\nabla F(\mathbf{w}_t)$. We introduce some notations in this paper. Denote $b = \sup_{z \in \mathcal{Z}} f(0; z)$, $\bar{b} = \sup_{z \in \mathcal{Z}} \|\nabla f(0, z)\|_2$ and $\kappa = \sup_{x \in \mathcal{X}} \|x\|_2$, where $\|\cdot\|_2$ is the Euclidean norm. For any $R > 0$, define $B_R = \{\mathbf{w} \in \mathbb{R}^d : \|\mathbf{w}\|_2 \leq R\}$. Let e be the base of the natural logarithm.

2.1 Assumptions

We need some assumptions. We say a differentiable function $g : \mathcal{W} \mapsto \mathbb{R}$ is L -smooth for a constant $L > 0$ if

$$\|\nabla g(\mathbf{w}) - \nabla g(\tilde{\mathbf{w}})\|_2 \leq L \|\mathbf{w} - \tilde{\mathbf{w}}\|_2, \quad \forall \mathbf{w}, \tilde{\mathbf{w}} \in \mathcal{W}.$$

For L -smooth g , we have two useful properties [22]

$$g(\mathbf{w}) \leq g(\tilde{\mathbf{w}}) + \langle \mathbf{w} - \tilde{\mathbf{w}}, \nabla g(\tilde{\mathbf{w}}) \rangle + 2^{-1} L \|\mathbf{w} - \tilde{\mathbf{w}}\|_2^2, \quad (2)$$

$$(2L)^{-1} \|\nabla g(\mathbf{w})\|_2^2 \leq g(\mathbf{w}) - \inf_{\tilde{\mathbf{w}}} g(\tilde{\mathbf{w}}). \quad (3)$$

Our first assumption is a standard and widely used assumption on the smoothness of loss functions [11–13, 23].

Assumption 1. We assume that for any z , the function $\mathbf{w} \mapsto f(\mathbf{w}; z)$ is L -smooth.

In our next assumption, we suppose that a weighted norm of gradient is bounded.

Assumption 2. We assume the existence of $G > 0$ such that

$$\sqrt{\eta_t} \|\nabla f(\mathbf{w}_t; z)\|_2 \leq G, \quad \forall t \in \mathbb{N}, z \in \mathcal{Z}.$$

In the literature, a bounded gradient assumption as $\|\nabla f(\mathbf{w}_t; z)\|_2 \leq G$ (i.e., the Lipschitz continuity of f) is often imposed to study statistical properties of SGD [15–17, 24] as well as the computational properties [11, 25]. As compared to this bounded gradient assumption, Assumption 2 is much milder since the step sizes should diminish to zero for the convergence of the algorithm. Indeed, typical choices of step sizes are $\eta_t = O(1/\sqrt{t})$ and $\eta_t = O(1/t)$ (depending on whether the objective functions satisfy additional conditions such as gradient-dominance condition) [13, 26], in which case, the gradients can respectively grow with the rate $O(t^{\frac{1}{4}})$ and $O(t^{\frac{1}{2}})$ without violating Assumption 2.

Our third assumption is on the boundedness of variances of stochastic gradients, which is widely adopted in the literature to study either computational errors [13, 14] or stability of SGD [16, 27] in the nonconvex setting.

Assumption 3. We assume the existence of $\sigma > 0$ such that

$$\mathbb{E}_{j_t} [\|\nabla f(\mathbf{w}_t; z_{j_t}) - \nabla F_S(\mathbf{w}_t)\|_2^2] \leq \sigma^2, \quad \forall t \in \mathbb{N},$$

where \mathbb{E}_{j_t} denotes the expectation w.r.t. j_t .

2.2 Structure of loss functions

Other than the general smooth loss functions, we also consider in particular a class of loss functions of the structure

$$f(\mathbf{w}; z) = \phi(\langle \mathbf{w}, x \rangle, y), \quad (4)$$

where $\phi : \mathbb{R}^2 \mapsto \mathbb{R}_+$. Functions of this structure can be found in *generalized linear models* and *robust regression*.

1. For generalized linear models in binary classification, we would like to learn a model of the form $\Pr\{Y = 1 | X = x\} = \ell(\langle \mathbf{w}^*, x \rangle)$ with $\mathbf{w}^* \in \mathbb{R}^d$ being a parameter vector and $\ell : \mathbb{R} \mapsto [0, 1]$ being a link function [18]. A commonly used loss is the non-linear squared loss $f(\mathbf{w}; z) = (y - \ell(\langle \mathbf{w}, x \rangle))^2$ which can be written as the form of (4) with $\phi(a, b) = (\ell(a) - b)^2$. Standard choices of ℓ include the logistic link function $\ell(s) = (1 + e^{-s})^{-1}$ and the probit link function $\ell(s) = \Phi(s)$ with Φ being the Gaussian cumulative distribution function.

2. For robust regression, we pick a potentially nonconvex function $\ell : \mathbb{R} \mapsto \mathbb{R}_+$ and assume a linear model $y = \langle \mathbf{w}^*, x \rangle + \xi$, where the noise term ξ is i.i.d. with mean zero [18, 20]. The loss function is then $f(\mathbf{w}; z) = \ell(y - \langle \mathbf{w}, x \rangle)$, which takes a form of (4) with $\phi(a, b) = \ell(b - a)$.

For loss functions of the structure (4), we will develop dimensionality-independent learning rates which would be appealing for the high-dimensional learning setting.

3 UNIFORM CONVERGENCE OF GRADIENTS

We are interested in the population gradients ∇F instead of empirical gradients ∇F_S . The gap between these two terms can be studied via the uniform convergence on the gradient of empirical risks to the population counterparts [18–20, 28]. In this section, we continue this direction by presenting a new connection between the uniform convergence of gradients and the Rademacher chaos complexities of order two, which is an extension of Rademacher complexities to U-processes [21, 29]. As we will see, this connection allows us to improve the existing uniform convergence by removing a logarithmic factor. It should be noted that the uniform convergence of gradients is not our main contribution.

Definition 1. Let $\mathcal{F} : \mathcal{Z} \times \mathcal{Z} \mapsto \mathbb{R}$ be a function class and $S = \{z_i\}_{i=1}^n \subset \mathcal{Z}$. Let $\{\epsilon_i\}_{i=1}^n$ be independent Rademacher variables with $\Pr\{\epsilon_i = 1\} = \Pr\{\epsilon_i = -1\} = 1/2$. The empirical Rademacher chaos complexity of order two for \mathcal{F} w.r.t. S is defined as $\mathcal{U}_S(\mathcal{F}) = \frac{1}{n} \mathbb{E}_\epsilon \left[\sup_{f \in \mathcal{F}} \sum_{1 \leq i < j \leq n} \epsilon_i \epsilon_j f(x_i, x_j) \right]$.

Theorem 1. Let $\delta \in (0, 1)$, $R > 0$ and $S = \{z_1, \dots, z_n\}$ be examples drawn independently from ρ . Suppose Assumption 1 holds. Then with probability at least $1 - \delta$ we have

$$\sup_{\mathbf{w} \in B_R} \|\nabla F(\mathbf{w}) - \nabla F_S(\mathbf{w})\|_2 \leq \frac{LR + \tilde{b}}{\sqrt{n}} \left(2 + \sqrt{2 \log(1/\delta)}\right) + 2\sqrt{\frac{2\mathcal{U}_S(\mathcal{F}_R)}{n}},$$

where $\mathcal{F}_R = \{(z, \tilde{z}) \mapsto \langle \nabla f(\mathbf{w}; z), \nabla f(\mathbf{w}; \tilde{z}) \rangle : \mathbf{w} \in B_R\}$.

As a corollary, we can derive the following bound by controlling the Rademacher chaos complexities in Theorem 1. The proofs of Theorem 1 and Corollary 2 are given in Section II.1 and Section II.2 (Supplementary Material).

Corollary 2. Under Assumptions of Theorem 1, the following inequality holds with probability $1 - \delta$

$$\sup_{\mathbf{w} \in B_R} \|\nabla F(\mathbf{w}) - \nabla F_S(\mathbf{w})\|_2 \leq (LR + \tilde{b})n^{-\frac{1}{2}} \times \left(2 + 2\sqrt{48e\sqrt{2}(\log 2 + d \log(3e))} + \sqrt{2 \log(1/\delta)}\right).$$

Remark 1. Based on covering numbers, it was shown [18]

$$\sup_{\mathbf{w} \in B_R} \|\nabla F(\mathbf{w}) - \nabla F_S(\mathbf{w})\|_2 = O\left(\sqrt{\frac{d \log(n) \log(1/\delta)}{n}}\right).$$

As a comparison, our argument based on Rademacher chaos complexities yields the bound $O(\sqrt{\frac{d + \log(1/\delta)}{n}})$, which is slightly tighter than the above one by removing the term $\log(n)$. Furthermore, our bound involves $d + \log(1/\delta)$ while

the above bound involves $d \log(1/\delta)$. Although the improvement by removing a logarithmic factor is marginal, the obtained bound is generally the best one can get. These improvements are achieved by establishing a new connection between uniform convergence of gradients and Rademacher chaos complexities (Theorem 1), which might be interesting on its own due to the important role of Rademacher-type complexities in learning theory.

The uniform convergence rates in Corollary 2 involves the dimensionality d . Dimensionality-free bounds are possible if we consider loss functions f of the structure (4). Proposition 3 is a direct corollary of the chain rule for vector-valued Rademacher complexities [20]. For completeness, we provide a complete proof in Section II.3 (Supplementary Material) based on Gaussian complexities.

Proposition 3. Let $\delta \in (0, 1)$, $R > 0$ and $S = \{z_1, \dots, z_n\}$ be examples drawn independently from ρ . Suppose $f : \mathcal{W} \times \mathcal{Z} \mapsto \mathbb{R}$ takes the form (4) with ϕ being L_ϕ -smooth w.r.t. the first argument. Then with probability $1 - \delta$, $\sup_{\mathbf{w} \in B_R} \|\nabla F(\mathbf{w}) - \nabla F_S(\mathbf{w})\|_2$ is upper bounded by

$$\frac{2\sqrt{2}\kappa(2L_\phi R\kappa + \tilde{b})}{\sqrt{n}} + \sqrt{\frac{2(L_\phi\kappa^2 R + \tilde{b})^2 \log(1/\delta)}{n}}.$$

4 LEARNING RATES

In this section, we present our main results on the behavior of gradients for population risks of SGD iterates. We consider general nonconvex objectives in Section 4.1, and then move to a subclass of nonconvex objectives satisfying the gradient-dominance condition in Section 4.2.

4.1 General nonconvex objective functions

Before providing learning rates, we first present a high-probability bound on the gradients of empirical risks to be proved in Section III.1 (Supplementary Material).

Lemma 4. Suppose Assumptions 1-3 hold. Let $\{\mathbf{w}_t\}_t$ be the sequence produced by (1) with $\eta_t \leq 1/(2L)$. Then for any $\delta \in (0, 1)$, the following inequality holds with probability $1 - \delta$

$$\sum_{k=1}^t \eta_k \|\nabla F_S(\mathbf{w}_k)\|_2^2 \leq C_t + C_1 \log(2/\delta), \quad (5)$$

where we introduce $C_t = 4b + 4L(\sigma^2 + G^2) \sum_{k=1}^t \eta_k^2$ and $C_1 = 8 \max\{G^2, \sigma^2/L\} + 32LG^2$.

Remark 2. Under Assumptions 1 and 3, it was shown for SGD with $\eta_t \leq 1/(2L)$ that [13] $\sum_{k=1}^t \eta_k \mathbb{E}[\|\nabla F_S(\mathbf{w}_k)\|_2^2] = O(\sum_{k=1}^t \eta_k^2 + 1)$, from which and the Markov's inequality it follows that $\Pr\{\sum_{k=1}^t \eta_k \mathbb{E}[\|\nabla F_S(\mathbf{w}_k)\|_2^2] \geq \delta\} \leq \delta^{-1} O(\sum_{k=1}^t \eta_k^2 + 1)$. As a comparison, we develop a high-probability bound where the dependency on the confidence parameter $1/\delta$ is logarithmic instead of linear.

In the following theorem, we present our main result on the decay rate of population gradients in general nonconvex cases. As we will see in the proof given in Section III.3 (Supplementary Material), the term $\|\nabla F(\mathbf{w}_t)\|_2^2$ depends on two terms: the computational error $\|\nabla F_S(\mathbf{w}_t)\|_2^2$ and the statistical error $\|\nabla F(\mathbf{w}_t) - \nabla F_S(\mathbf{w}_t)\|_2^2$. Generally, the

computational error decreases as a function of the iteration number, while the statistical error would increase along the learning process (\mathbf{w}_t traverses over a ball with an increasing radius as t increases, and therefore one has to apply Theorem 1 with an increasing R). Therefore, one should trade-off these two errors by selecting an appropriate iteration number. We denote $A \asymp B$ if there exist universal constants $\tilde{C}_1, \tilde{C}_2 > 0$ such that $\tilde{C}_1 A \leq B \leq \tilde{C}_2 A$.

Theorem 5. *Suppose Assumptions 1-3 hold. Let $\{\mathbf{w}_t\}_t$ be the sequence produced by (1) with $\eta_t = \eta_1 t^{-\theta}$, $\theta \in (0, 1)$ and $\eta_1 \leq 1/(2L)$. Then for any $\delta \in (0, 1)$, we can choose $T \asymp (nd^{-1})^{\frac{1}{2(1-\theta)}}$ to derive with probability $1 - \delta$*

$$\left(\sum_{t=1}^T \eta_t\right)^{-1} \sum_{t=1}^T \eta_t \|\nabla F(\mathbf{w}_t)\|_2^2 = \begin{cases} O\left((nd^{-1})^{\frac{\theta}{2(1-\theta)}} \log^3(1/\delta)\right), & \text{if } \theta < 1/2, \\ O\left(n^{-\frac{1}{2}} d^{\frac{1}{2}} \log(T/\delta) \log^3(1/\delta)\right), & \text{if } \theta = 1/2, \\ O\left(n^{-\frac{1}{2}} d^{\frac{1}{2}} \log^3(1/\delta)\right), & \text{if } \theta > 1/2. \end{cases} \quad (6)$$

Remark 3. According to Theorem 5, one can select an appropriate early-stopping iteration number to achieve similar learning rates for polynomially decaying step sizes with $\theta \in (1/2, 1)$. However, the corresponding computational cost measured by $T \asymp (nd^{-1})^{\frac{1}{2(1-\theta)}}$ varies for different θ , which achieves its minimum by taking θ close to $1/2$.

Remark 4. We compare our results with the existing stability bounds. For the specific step-size sequence $\eta_t = O(1/t)$, stability bounds $O(T^\alpha/n)$ with $\alpha \in (0, 1)$ were established for SGD with T iterations in nonconvex learning (Theorem 3.8 in [15]), from which one can derive $\mathbb{E}[F(\mathbf{w}_T) - F_S(\mathbf{w}_T)] = O(T^\alpha/n)$. As compared to our uniform convergence approach, this stability approach has an appealing property of implying dimensionality-independent learning rates. However, the stability-based approach requires a fast-decaying step sizes as $\eta_t = O(1/t)$, for which the associated computational errors would generally decay as

$$\min_{t=1, \dots, T} \mathbb{E}[\|\nabla F_S(\mathbf{w}_t)\|_2^2] = O\left(\left(1 + \sum_{t=1}^T \eta_t^2\right) / \sum_{t=1}^T \eta_t\right) = O(\log^{-1}(T)),$$

where the first step is due to [13] and the second step is due to the elementary inequalities $\sum_{t=1}^T t^{-2} = O(1)$, $\sum_{t=1}^T t^{-1} \asymp \log T$. To get a meaningful statistical error bound, one requires $T = o(n^{\frac{1}{\alpha}})$, for which the associated computational error bound decays as $O(\log^{-1}(n))$. As a comparison, our statistical errors are developed for the step sizes $\eta_t = O(t^{-\theta})$, $\theta \in (0, 1)$, for which the associated computational errors decay significantly faster as $O(T^{\max\{\theta-1, -\theta\}})$. The ability to balance the statistical and computational errors with a general step-size sequence allows us to develop learning rate of the order $\tilde{O}((n^{-1}d)^{\frac{1}{2}})$ in the general nonconvex learning setting, where we use the notation \tilde{O} to hide logarithmic factors. It is now clear that each of these two approaches has its own advantages over the other, depending on whether the dependency on the dimensionality or the iteration number is more interesting. In more details, the stability approach is preferable in a high-dimensional learning setting with $n = O(d)$,

while the uniform-convergence approach is better for low-dimensional learning problems with $d = o(n)$.

If the loss function takes a specific structure, we can improve the learning rates in Theorem 5 by removing the dependency on the dimensionality. Theorem 6 can be proved similarly to Theorem 5 by using Proposition 3 instead of Corollary 2, and we omit the proof for brevity.

Theorem 6. *Let assumptions in Theorem 5 hold. If $f : \mathcal{W} \times \mathcal{Z} \mapsto \mathbb{R}$ takes the form (4) with ϕ being L_ϕ -smooth w.r.t. the first argument, we can choose $T \asymp n^{\frac{1}{2(1-\theta)}}$ to derive Eq. (6) with the involved d removed.*

Remark 5. We can argue the effectiveness of learning rates in Theorem 6 as follows. If we assume ϕ is convex w.r.t. the first argument, the best known generalization bound of SGD is of the order $F(\bar{\mathbf{w}}_T) - F(\mathbf{w}^*) = \tilde{O}(n^{-\frac{1}{2}})$, where $\bar{\mathbf{w}}_T$ is an average of the first T iterates with an appropriately chosen T and $\mathbf{w}^* = \arg \min_{\mathbf{w}} F(\mathbf{w})$ [15, 30, 31]. It then follows from (3) that $\|\nabla F(\bar{\mathbf{w}}_T)\|_2^2 = \tilde{O}(n^{-\frac{1}{2}})$ with high probabilities. This matches our learning rates derived in the nonconvex learning setting in Theorem 6.

4.2 Gradient-dominated objective functions

In this section, we show that a faster learning rate is possible if we impose a gradient-dominance condition on the objective function. Roughly speaking, gradient-dominance condition means that the suboptimality of function values can be bounded by the squared magnitude of gradients. It also means that every stationary point must be a global optimum. The gradient-dominance condition is widely used in the analysis of nonconvex learning algorithms, see, e.g., [11, 17, 20, 26, 32]. Many nonconvex objective functions satisfy the gradient-dominance condition, including one hidden layer neural networks, ResNets with linear activations, phase retrieval and matrix factorization [20, 32].

Assumption 4 (Gradient-dominance condition). We assume for all $S \subset \mathcal{Z}$, there exists an $\mu_S > 0$ such that

$$F_S(\mathbf{w}) - F_S(\mathbf{w}_S) \leq (4\mu_S)^{-1} \|\nabla F_S(\mathbf{w})\|_2^2, \quad \forall \mathbf{w} \in \mathcal{W},$$

where $\mathbf{w}_S = \arg \min_{\mathbf{w}} F_S(\mathbf{w})$.

For gradient-dominated objectives, we can implement SGD with $O(n)$ iterations to achieve the learning rates $\tilde{O}(n^{-1}d)$ for excess risks instead of gradients. The proof of Theorem 7 is given in Section IV.

Theorem 7. *Suppose Assumptions 1-4 hold. Let $\{\mathbf{w}_t\}_t$ be the sequence produced by (1) with $\eta_t = 2/(\mu_S(t + t_0))$ for all $t \in \mathbb{N}$ and $t_0 \geq \max\{4L/\mu_S, 1\}$. Let $\delta \in (0, 1)$. Then we can choose $T \asymp n$ to derive the following inequality with probability $1 - \delta$*

$$F(\mathbf{w}_T) - F(\mathbf{w}^*) = O\left(n^{-1}(d \log(1/\delta)) \log^2 n \log^2(1/\delta)\right). \quad (7)$$

Remark 6. We compare our results with the existing stability analysis. For gradient-dominated objectives, computational error bounds of the order $O(1/T)$ were developed for step sizes $\eta_t = \frac{2t+1}{2\mu_S(t+1)^2}$ [26]. For this step-size sequence, uniform stability bounds were shown to decay as $O(n^{-1}T^{\frac{\alpha}{\alpha+1}})$ (Theorem 3.8 of [15]), where $\alpha = L/\mu_S$ behaves as a condition number. Therefore, the stability

approach implies $\mathbb{E}[F(\mathbf{w}_T)] - F(\mathbf{w}^*) = O(T^{-1} + n^{-1}T^{\frac{\alpha}{\alpha+1}})$, and one can choose $T \asymp n^{\frac{\alpha+1}{2\alpha+1}}$ to derive the learning rate $O(n^{-\frac{\alpha+1}{2\alpha+1}})$. As compared to Theorem 7, it is clear that the stability-based rates have an advantage of being dimensionality-free but meanwhile a worse dependency on n . Indeed, the stability approach is preferable if $n = O(d^{\frac{2\alpha+1}{\alpha}})$, while our learning rate is better if $d = O(n^{\frac{\alpha}{2\alpha+1}})$.

Analogous to the case with general nonconvex objectives, we can remove the dependency of learning rates on the dimensionality for loss functions of the structure (4). We omit the proof of Theorem 8 for brevity.

Theorem 8. *Let assumptions in Theorem 7 hold. If $f : \mathcal{W} \times \mathcal{Z} \mapsto \mathbb{R}$ takes the form (4) with ϕ being L_ϕ -smooth w.r.t. the first argument, then we can choose $T \asymp n$ to derive (7) with the involved d removed.*

Remark 7. The effectiveness of learning rates in Theorem 8 can be shown as follows. If we assume ϕ is Lipschitz continuous and strongly convex w.r.t. the first argument, the best known generalization bound of the order $F(\bar{\mathbf{w}}_T) - F(\mathbf{w}^*) = \tilde{O}(T^{-1})$ was derived in the online learning setting, where $\bar{\mathbf{w}}_T$ is an average of the first T iterates [33]. This online learning setting corresponds to the one-pass SGD since the released training examples there need to be independently drawn from ρ . Therefore the results in [33] correspond to the generalization bound $F(\bar{\mathbf{w}}_T) - F(\mathbf{w}^*) = \tilde{O}(n^{-1})$ in the stochastic learning setting with T chosen proportionally to the sample size, which matches the bounds in Theorem 8.

4.3 Sketch of the proof

We sketch here our main idea in proving our results in Section 4. Our learning rates are derived by controlling the population gradients in terms of two components as follows

$$\|\nabla F(\mathbf{w}_t)\|_2^2 \leq 2\|\nabla F_S(\mathbf{w}_t)\|_2^2 + 2\|\nabla F_S(\mathbf{w}_t) - \nabla F(\mathbf{w}_t)\|_2^2.$$

We refer to the first term as the computational error since it is related to the optimization algorithm to minimize the empirical risk F_S . The second term is called the statistical error since it is related to approximating the true gradient with its empirical counterpart based on random samples.

Computational errors have received much attention in the optimization community [11–13, 23]. However, existing results are mainly studied in expectation, while our focus here is to establish high-probability bounds (Lemma 4). Our idea to this aim is to bound $\sum_{k=1}^t \eta_k \|\nabla F_S(\mathbf{w}_k)\|_2^2$ in terms of two martingale difference sequences where the variance of martingales plays an essential role in deriving the stated high probability bounds. To apply uniform convergence of gradients in Section 3 to control statistical errors, an essential step is to control the norm of iterates in the following lemma to be proved in Section III.2 (Supplementary Material).

Lemma 9. *Suppose Assumptions 1-3 hold. Let $\{\mathbf{w}_t\}_t$ be produced by (1) with $\eta_t \leq 1/(2L)$. Then for any $\delta \in (0, 1)$, with probability $1 - \delta$ we have uniformly for all $t = 1, \dots, T$*

$$\|\mathbf{w}_{t+1}\|_2 \leq C_2 \left(\left(\sum_{k=1}^T \eta_k^2 \right)^{\frac{1}{2}} + 1 \right) \left(\left(\sum_{k=1}^t \eta_k \right)^{\frac{1}{2}} + 1 \right) \log(4/\delta), \quad (8)$$

where C_2 is independent of T and δ (explicitly given in the proof).

Our idea in proving Lemma 9 is to use (1) to get

$$\begin{aligned} \|\mathbf{w}_{t+1}\|_2 &\leq \left\| \sum_{k=1}^t \eta_k \nabla f(\mathbf{w}_k; z_{j_k}) \right\|_2 \leq \left\| \sum_{k=1}^t \eta_k \nabla F_S(\mathbf{w}_k) \right\|_2 \\ &\quad + \left\| \sum_{k=1}^t \eta_k (\nabla F_S(\mathbf{w}_k) - \nabla f(\mathbf{w}_k; z_{j_k})) \right\|_2. \end{aligned}$$

The last second term of the above inequality can be controlled by Lemma 4 together with the Schwartz's inequality. Furthermore, the term $\sum_{k=1}^t \eta_k (\nabla F_S(\mathbf{w}_k) - \nabla f(\mathbf{w}_k; z_{j_k}))$ is a summation of a martingale difference sequence, which we can control by concentration inequalities of martingales.

5 DISCUSSIONS

We discuss here related work on computational errors, uniform convergence of gradients and stability-based analysis.

Computational errors. SGD for nonconvex learning is mainly studied from the perspective of computational errors. Initially, asymptotic properties of SGD were studied [34]. The key property on the smoothness of loss functions was used to establish the first nonasymptotic convergence rates in expectation $\mathbb{E}[\|\nabla F_S(\mathbf{w}_u)\|_2^2] = O(1/\sqrt{T})$ where u is drawn from a probability distribution on $\{1, \dots, T\}$ [13]. For gradient-dominated and smooth objectives, it was further shown that $\mathbb{E}[F_S(\mathbf{w}_T)] - \inf_{\mathbf{w}} F_S(\mathbf{w}) = O(1/T)$ [26]. These discussions were then extended successfully to other variants of SGD for nonconvex optimization, including distributed SGD [25], stochastic composite optimization [6, 14] and stochastic variance reduction [11, 12]. In particular, near optimal iteration complexities were recently achieved by some interesting variance-reduced algorithms for nonconvex optimization [35–38].

Uniform convergence of gradients. In nonconvex learning, convergence rates are generally stated for the gradients of empirical risks [11–13, 39], which not necessarily means that similar convergence rates can be carried to their population counterparts. Motivated by this, the uniform convergence of gradients is recently drawing increasing attention [18–20]. Based on the tool of covering numbers, convergence rates $O(\sqrt{n^{-1} \log n})$ were established for the uniform deviation between population and empirical gradients [18, 19]. A chain rule for vector-valued Rademacher complexities was established to show the uniform convergence of gradients for nonconvex functions with the structure (4) [20]. The above mentioned uniform convergence results are established for smooth functions. Recently, an interesting graphical convergence was studied for nonconvex and nonsmooth loss functions, for which the convergence is measured by the gradient of the Moreau envelopes [40]. Similar to our results in Section 3, the graphic convergence in [40] involves a square-root dependency on the dimension d in general, and is dimensionality-independent for loss functions of the structure (4). These discussions apply to empirical risk minimization, which do not take into account the computational property of learning algorithms.

Stability-based analysis. Other than the uniform convergence approach, another popular approach to investigate the generalization behavior of a learning algorithm is to study its algorithmic stability [41–43]. An advantage of




Fig. 1. Computational errors, statistical errors and test errors versus the number of passes for some datasets. Grad_Tr, Grad_Te and Grad_diff refer to computational errors, test errors and statistical errors, respectively.




Fig. 2. Test errors versus the number of passes for step sizes of the form $\eta_t = \eta_1 t^{-\theta}$, $\theta \in \{0, 1/2, 1\}$ and some datasets.

stability-based approach is that it considers the specific model produced by the learning algorithm, and therefore can imply dimensionality-independent learning rates. As a comparison, the dependence on the dimensionality is generally inevitable for the uniform convergence approach. Initially, stability bounds are especially useful to derive optimal generalization error bounds stated *in expectation* [44]. Recently, almost optimal generalization bounds *with high probability* were developed for uniformly-stable algorithms in a series of breakthrough papers [45–47]. The uniform stability of SGD was established in a seminar paper [15], which motivates some interesting work on data-dependent stability [16] and on-average stability of SGD [17, 44]. However, stability bounds there are mainly derived for the specific step-size sequence $\eta_t = O(1/t)$ in nonconvex learning, for which one generally needs an exponential number of iterations for a moderate decay of the associated computational errors [13]. Under gradient-dominance conditions or quadratic growth conditions, black-box stability results for first-order methods were established [32]. However, the discussions there require to impose a strict assumption on the convergence of a learning algorithm to global optimum which is difficult to hold for nonconvex problems. Stability of stochastic gradient Langevin dynamics is also drawing increasing attention recently, where isotropic Gaussian noises are added to each stochastic gradient steps [24].

The primary focus of stability bounds is to control the generalization gap $F(\mathbf{w}_T) - F_S(\mathbf{w}_T)$ which ignores computational errors, while we aim to study the gradients of population risks where the interplay between computational and statistical errors plays an important role to achieve satisfactory learning rates. Unlike existing bounds

stated in expectation [15–17, 24, 32], our analysis is able to develop more challenging high-probability bounds. A key observation is to show that the complexity of SGD iterates grows in a controllable manner w.r.t. the step-size sequence (Lemma 9), which allows us to relate the gradients of empirical risks to their population counterparts via the uniform convergence of gradients. Our analysis therefore provides an alternative approach to study learning rates of SGD in nonconvex learning without considering the stability of an algorithm. As stated in Remark 5 and Remark 7, the effectiveness of our learning rates is justified by matching the existing generalization bounds in the convex learning setting. Furthermore, our statistical errors are developed for step sizes $\eta_t = O(t^{-\theta})$ with $\theta \in (0, 1)$ in the general nonconvex learning setting, for which the associated computational errors decay significantly faster than $O(\log^{-1} T)$ achieved by the step sizes $\eta_t = O(t^{-1})$. Although our learning rates are dimensionality-dependent in general, one can exploit the specific structure of loss functions to derive dimensionality-independent bounds in specific cases.

For convex learning, generalization bounds of SGD were studied via algorithmic stability [15, 16], integral operators [48–50] and tools in empirical process [30, 31].

6 SIMULATIONS

In this section, we present some simulation results to validate our theory. According to Section 2.2, we consider a generalized linear model for binary classification where the loss function takes the form $f(\mathbf{w}; z) = (\ell(\langle \mathbf{w}, x \rangle) - y)^2$ and ℓ is the logistic link function $\ell(s) = (1 + e^{-s})^{-1}$. We consider three datasets available from the LIBSVM dataset:

Segment, Fourclass and German [51]. For each dataset, we take 80 percents as the training dataset, and reserve the remaining 20 percents for testing. Let $F_S(\mathbf{w})$ and $F_{\tilde{S}}(\mathbf{w})$ be the objective function built on the training dataset S and the testing dataset \tilde{S} , i.e., $F_{\tilde{S}}(\mathbf{w}) = \frac{1}{|\tilde{S}|} \sum_{z \in \tilde{S}} f(\mathbf{w}; z)$, where $|\tilde{S}|$ denotes the cardinality of the set \tilde{S} . It is clear that $F_{\tilde{S}}$ serves as a good approximation of F . We repeat experiments 100 times and report the average of results.

Our first aim is to illustrate how the computational and statistical errors behave versus the number of passes, which is the iteration number divided by the sample size. We apply SGD with 200 passes and consider the step sizes $\eta_t = 5/\sqrt{t}$. We plot the behavior of test errors $\|\nabla F_{\tilde{S}}(\mathbf{w}_t)\|_2$, computational errors $\|\nabla F_S(\mathbf{w}_t)\|_2$ and statistical errors versus the number of passes. According to Fig. 1, along the optimization process the computational errors continue to decrease while the statistical errors continue to increase, which is well consistent with our theory. It can be also seen that the model trained by SGD also generalizes well by its resistance to overfitting, i.e., the test error $\|\nabla F_{\tilde{S}}(\mathbf{w}_t)\|_2$ does not fluctuate even if we apply SGD with 200 passes.

Our second aim is to show how test errors would behave versus different step sizes. To this aim, we consider step sizes of the form $\eta_t = \eta_1 t^{-\theta}$ with $\theta \in \{0, 1/2, 1\}$. For each θ , we validate η_1 over the set $\{1, 2, \dots, 2^8\}$ based on a validation set, and report the behavior of the associated test errors versus the number of passes in Fig. 2. In our experiments, SGD with $\theta = 1/2$ outperforms those with $\theta = 0$ and $\theta = 1$. This is consistent with Theorem 5, where $\theta = 1/2$ is shown to achieve the best learning rates (up to logarithmic factors) while achieving the minimal computation cost.

7 CONCLUSIONS

In this paper, we present learning rates of SGD for nonconvex learning problems from a joint perspective of optimization and statistics, which allows us to see how an optimal number of passes should be taken to trade-off the computational and statistical errors. We consider both general nonconvex and gradient-dominated objectives, and derive learning rates comparable to the corresponding ones in the convex case. For objective functions with a specific structure, we show that the learning rates can be dimensionality-independent. We control the statistical errors by giving high-probability bounds on the complexity of SGD iterates, which provides an alternative explanation on the generalization performance of SGD to train nonconvex models.

There remain some interesting directions to pursue. Firstly, it is interesting to extend the learning rates here to some variants of SGD for nonconvex learning, including SGD with variance reduction [52–54] and SGD with momentum [22, 55]. Secondly, the existing stability bounds for nonconvex learning are developed for the step-size sequence $\eta_t = O(1/t)$. It is very interesting to study the stability of SGD for the general polynomially-decaying step sizes.

ACKNOWLEDGMENTS

We are grateful to the editor and referees for the constructive comments, which are very helpful to improve the paper.

REFERENCES

- [1] M. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in *International Conference on Machine Learning*, 2003, pp. 928–936.
- [2] T. Zhang, "Solving large scale linear prediction problems using stochastic gradient descent algorithms," in *International Conference on Machine Learning*, 2004, pp. 919–926.
- [3] L. Bottou, F. E. Curtis, and J. Nocedal, "Optimization methods for large-scale machine learning," *SIAM Review*, vol. 60, no. 2, pp. 223–311, 2018.
- [4] S. Huang and Z. Zhou, "Fast multi-instance multi-label learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, to appear.
- [5] O. Bousquet and L. Bottou, "The tradeoffs of large scale learning," in *Advances in Neural Information Processing Systems*, 2008, pp. 161–168.
- [6] W. Zhang, L. Zhang, Z. Jin, R. Jin, D. Cai, X. Li, R. Liang, and X. He, "Sparse learning with stochastic composite optimization," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 6, pp. 1223–1236, 2017.
- [7] B. Neyshabur, S. Bhojanapalli, D. McAllester, and N. Srebro, "Exploring generalization in deep learning," in *Advances in Neural Information Processing Systems*, 2017, pp. 5947–5956.
- [8] P. L. Bartlett, D. J. Foster, and M. J. Telgarsky, "Spectrally-normalized margin bounds for neural networks," in *Advances in Neural Information Processing Systems*, 2017, pp. 6240–6249.
- [9] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, "Understanding deep learning requires rethinking generalization," in *International Conference on Learning Representations*, 2017.
- [10] I. Steinwart and A. Christmann, *Support Vector Machines*. Springer Science & Business Media, 2008.
- [11] S. Reddi, A. Hefny, S. Sra, B. Póczos, and A. Smola, "Stochastic variance reduction for nonconvex optimization," in *International Conference on Machine Learning*, 2016, pp. 314–323.
- [12] Z. Allen-Zhu and E. Hazan, "Variance reduction for faster non-convex optimization," in *International Conference on Machine Learning*, 2016, pp. 699–707.
- [13] S. Ghadimi and G. Lan, "Stochastic first-and zeroth-order methods for nonconvex stochastic programming," *SIAM Journal on Optimization*, vol. 23, no. 4, pp. 2341–2368, 2013.
- [14] S. Ghadimi, G. Lan, and H. Zhang, "Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization," *Mathematical Programming*, vol. 155, no. 1-2, pp. 267–305, 2016.
- [15] M. Hardt, B. Recht, and Y. Singer, "Train faster, generalize better: Stability of stochastic gradient descent," in *International Conference on Machine Learning*, 2016, pp. 1225–1234.
- [16] I. Kuzborskij and C. Lampert, "Data-dependent stability of stochastic gradient descent," in *International Conference on Machine Learning*, 2018, pp. 2820–2829.
- [17] Y. Zhou, Y. Liang, and H. Zhang, "Generalization error bounds with probabilistic guarantee for SGD in nonconvex optimization," *arXiv preprint arXiv:1802.06903*, 2018.
- [18] S. Mei, Y. Bai, and A. Montanari, "The landscape of empirical risk for nonconvex losses," *The Annals of Statistics*, vol. 46, no. 6A, pp. 2747–2774, 2018.
- [19] L. Zhang, T. Yang, and R. Jin, "Empirical risk minimization for stochastic convex optimization: $O(1/n)$ - and $O(1/n^2)$ -type of risk bounds," in *Conference on Learning Theory*, 2017.

- pp. 1954–1979.
- [20] D. J. Foster, A. Sekhari, and K. Sridharan, “Uniform convergence of gradients for non-convex learning and optimization,” in *Advances in Neural Information Processing Systems*, 2018, pp. 8759–8770.
- [21] V. De la Peña and E. Giné, *Decoupling: From Dependence to Independence*. New York: Springer-Verlag, 1999.
- [22] Y. Nesterov, *Introductory Lectures on Convex Optimization: A Basic Course*. Springer Science & Business Media, 2013, vol. 87.
- [23] D. P. Bertsekas and J. N. Tsitsiklis, “Gradient convergence in gradient methods with errors,” *SIAM Journal on Optimization*, vol. 10, no. 3, pp. 627–642, 2000.
- [24] W. Mou, L. Wang, X. Zhai, and K. Zheng, “Generalization bounds of sgld for non-convex learning: Two theoretical viewpoints,” in *Conference on Learning Theory*, 2018, pp. 605–638.
- [25] X. Lian, Y. Huang, Y. Li, and J. Liu, “Asynchronous parallel stochastic gradient for nonconvex optimization,” in *Advances in Neural Information Processing Systems*, 2015, pp. 2737–2745.
- [26] H. Karimi, J. Nutini, and M. Schmidt, “Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition,” in *European Conference on Machine Learning*, 2016, pp. 795–811.
- [27] Y. Zhou, Y. Liang, and H. Zhang, “Generalization error bounds with probabilistic guarantee for SGD in nonconvex optimization,” *arXiv preprint arXiv:1802.06903*, 2018.
- [28] L. Zhang and Z.-H. Zhou, “Stochastic approximation of smooth and strongly convex functions: Beyond the $\mathcal{O}(1/t)$ convergence rate,” in *Conference on Learning Theory*, 2019, pp. 3160–3179.
- [29] Y. Ying and C. Campbell, “Rademacher chaos complexities for learning the kernel problem,” *Neural computation*, vol. 22, no. 11, pp. 2858–2886, 2010.
- [30] J. Lin, R. Camoriano, and L. Rosasco, “Generalization properties and implicit regularization for multiple passes SGM,” in *International Conference on Machine Learning*, 2016, pp. 2340–2348.
- [31] Y. Lei and K. Tang, “Stochastic composite mirror descent: Optimal bounds with high probabilities,” in *Advances in Neural Information Processing Systems*, 2018, pp. 1524–1534.
- [32] Z. Charles and D. Papailiopoulos, “Stability and generalization of learning algorithms that converge to global optima,” in *International Conference on Machine Learning*, 2018, pp. 744–753.
- [33] S. M. Kakade and A. Tewari, “On the generalization ability of online strongly convex programming algorithms,” in *Advances in Neural Information Processing Systems*, 2009, pp. 801–808.
- [34] L. Bottou, “On-line learning and stochastic approximations,” in *On-line Learning in Neural Networks*, D. Saad, Ed. New York, NY, USA: Cambridge University Press, 1998, pp. 9–42.
- [35] C. Fang, C. J. Li, Z. Lin, and T. Zhang, “Spider: Near-optimal non-convex optimization via stochastic path-integrated differential estimator,” in *Advances in Neural Information Processing Systems*, 2018, pp. 689–699.
- [36] D. Zhou, P. Xu, and Q. Gu, “Stochastic nested variance reduced gradient descent for nonconvex optimization,” in *Advances in Neural Information Processing Systems*, 2018, pp. 3921–3932.
- [37] L. M. Nguyen, M. van Dijk, D. T. Phan, P. H. Nguyen, T.-W. Weng, and J. R. Kalagnanam, “Finite-sum smooth optimization with SARAH,” 2019.
- [38] Z. Wang, K. Ji, Y. Zhou, Y. Liang, and V. Tarokh, “Spiderboost and momentum: Faster variance reduction algorithms,” in *Advances in Neural Information Processing Systems*, 2019, pp. 2403–2413.
- [39] Z. Allen-Zhu, “Natasha 2: Faster non-convex optimization than sgd,” in *Advances in Neural Information Processing Systems*, 2018, pp. 2675–2686.
- [40] D. Davis and D. Drusvyatskiy, “Uniform graphical convergence of subgradients in nonconvex optimization and learning,” *arXiv preprint arXiv:1810.07590*, 2018.
- [41] A. Elisseeff, T. Evgeniou, and M. Pontil, “Stability of randomized learning algorithms,” *Journal of Machine Learning Research*, vol. 6, no. Jan, pp. 55–79, 2005.
- [42] O. Bousquet and A. Elisseeff, “Stability and generalization,” *Journal of Machine Learning Research*, vol. 2, no. Mar, pp. 499–526, 2002.
- [43] T. Liu, D. Tao, M. Song, and S. J. Maybank, “Algorithm-dependent generalization bounds for multi-task learning,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 2, pp. 227–241, 2016.
- [44] S. Shalev-Shwartz, O. Shamir, N. Srebro, and K. Sridharan, “Learnability, stability and uniform convergence,” *Journal of Machine Learning Research*, vol. 11, no. Oct, pp. 2635–2670, 2010.
- [45] V. Feldman and J. Vondrak, “Generalization bounds for uniformly stable algorithms,” in *Advances in Neural Information Processing Systems*, 2018, pp. 9747–9757.
- [46] —, “High probability generalization bounds for uniformly stable algorithms with nearly optimal rate,” in *Conference on Learning Theory*, 2019, pp. 1270–1279.
- [47] O. Bousquet, Y. Klochkov, and N. Zhivotovskiy, “Sharper bounds for uniformly stable algorithms,” in *Conference on Learning Theory*, 2020, pp. 610–626.
- [48] L. Rosasco and S. Villa, “Learning with incremental iterative regularization,” in *Advances in Neural Information Processing Systems*, 2015, pp. 1630–1638.
- [49] A. Dieuleveut and F. Bach, “Nonparametric stochastic approximation with large step-sizes,” *Annals of Statistics*, vol. 44, no. 4, pp. 1363–1399, 2016.
- [50] J. Lin and L. Rosasco, “Optimal rates for multi-pass stochastic gradient methods,” *Journal of Machine Learning Research*, vol. 18, no. 1, pp. 3375–3421, 2017.
- [51] C.-C. Chang and C.-J. Lin, “Libsvm: a library for support vector machines,” *ACM Transactions on Intelligent Systems and Technology*, vol. 2, no. 3, p. 27, 2011.
- [52] R. Johnson and T. Zhang, “Accelerating stochastic gradient descent using predictive variance reduction,” in *Advances in Neural Information Processing Systems*, 2013, pp. 315–323.
- [53] L. Zhang, M. Mahdavi, and R. Jin, “Linear convergence with condition number independent access of full gradients,” in *Advances in Neural Information Processing Systems*, 2013, pp. 980–988.
- [54] M. Schmidt, N. Le Roux, and F. Bach, “Minimizing finite sums with the stochastic average gradient,” *Mathematical Programming*, vol. 162, no. 1–2, pp. 83–112, 2017.
- [55] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *International Conference on Learning Representations*, 2015.