
Data-Dependent Stability of Stochastic Gradient Descent

Ilja Kuzborskij¹ Christoph H. Lampert²

Abstract

We establish a data-dependent notion of algorithmic stability for Stochastic Gradient Descent (SGD), and employ it to develop novel generalization bounds. This is in contrast to previous distribution-free algorithmic stability results for SGD which depend on the worst-case constants. By virtue of the data-dependent argument, our bounds provide new insights into learning with SGD on convex and non-convex problems. In the convex case, we show that the bound on the generalization error depends on the risk at the initialization point. In the non-convex case, we prove that the expected curvature of the objective function around the initialization point has crucial influence on the generalization error. In both cases, our results suggest a simple data-driven strategy to stabilize SGD by pre-screening its initialization. As a corollary, our results allow us to show optimistic generalization bounds that exhibit fast convergence rates for SGD subject to a vanishing empirical risk and low noise of stochastic gradient.

1. Introduction

Stochastic gradient descent (SGD) has become one of the workhorses of modern machine learning. In particular, it is the optimization method of choice for training highly complex and non-convex models, such as neural networks. When it was observed that these models generalize better (suffer less from overfitting) than classical machine learning theory suggests, a large theoretical interest emerged to explain this phenomenon. Given that SGD at best finds a local minimum of the non-convex objective function, it has been argued that all such minima might be equally good. However, at the same time, a large body of empirical work and tricks of trade, such as *early stopping*, suggests that in prac-

tice one might not even reach a minimum, yet nevertheless observes excellent performance.

In this work we follow an alternative route that aims to *directly* analyze the generalization ability of SGD by studying how sensitive it is to small perturbations in the training set. This is known as *algorithmic stability* approach (Bousquet & Elisseeff, 2002) and was used recently (Hardt et al., 2016) to establish generalization bounds for both convex and non-convex learning settings. To do so they employed a rather restrictive notion of stability that does not depend on the data, but captures only intrinsic characteristics of the learning algorithm and global properties of the objective function. Consequently, their analysis results in worst-case guarantees that in some cases tend to be too pessimistic. As recently pointed out in (Zhang et al., 2017), *deep learning* might indeed be such a case, as this notion of stability is insufficient to give deeper theoretical insights, and a less restrictive one is desirable.

As our main contribution in this work we establish that a data-dependent notion of algorithmic stability, very similar to the *On-Average Stability* (Shalev-Shwartz et al., 2010), holds for SGD when applied to convex as well as non-convex learning problems. As a consequence we obtain new generalization bounds that depend on the data-generating distribution and the initialization point of an algorithm. For convex loss functions, the bound on the generalization error is essentially multiplicative in the risk at the initialization point when noise of stochastic gradient is not too high. For the non-convex loss functions, besides the risk, it is also critically controlled by the expected second-order information about the objective function at the initialization point. We further corroborate our findings empirically and show that, indeed, the data-dependent generalization bound is tighter than the worst-case counterpart on non-convex objective functions. Finally, the nature of the data-dependent bounds allows us to state *optimistic* bounds that switch to the faster rate of convergence subject to the vanishing empirical risk.

In particular, our findings justify the intuition that SGD is more stable in less curved areas of the objective function and link it to the generalization ability. This also backs up numerous empirical findings in the deep learning literature that solutions with low generalization error occur in less curved regions. At the same time, in pessimistic scenarios,

¹University of Milan, Italy ²IST Austria. Correspondence to: Ilja Kuzborskij <ilja.kuzborskij@gmail.com>.

our bounds are no worse than those of (Hardt et al., 2016).

Finally, we exemplify an application of our bounds, and propose a simple yet principled *transfer learning* scheme for the convex and non-convex case, which is guaranteed to transfer from the best source of information. In addition, this approach can also be used to select a good initialization given a number of random starting positions. This is a theoretically sound alternative to the purely random commonly used in non-convex learning.

The rest of the paper is organized as follows. We revisit the connection between stability and generalization of SGD in Section 3 and introduce a data-dependent notion of stability in Section 4. We state the main results in Section 5, in particular, Theorem 3 for the convex case, and Theorem 4 for the non-convex one. Next we demonstrate empirically that the bound shown in Theorem 4 is tighter than the worst-case one in Section 5.2.1. Finally, we suggest application of these bounds by showcasing principled transfer learning approaches in Section 5.3, and we conclude in Section 7.

2. Related Work

Algorithmic stability has been a topic of interest in learning theory for a long time, however, the modern approach on the relationship between stability and generalization goes back to the milestone work of (Bousquet & Elisseeff, 2002). They analyzed several notions of stability, which fall into two categories: distribution-free and distribution-dependent ones. The first category is usually called *uniform* stability and focuses on the intrinsic stability properties of an algorithm without regard to the data-generating distribution. Uniform stability was used to analyze many algorithms, including regularized Empirical Risk Minimization (ERM) (Bousquet & Elisseeff, 2002), randomized aggregation schemes (Elisseeff et al., 2005), and recently SGD by (Hardt et al., 2016; London, 2016), and (Poggio et al., 2011). Despite the fact that uniform stability has been shown to be sufficient to guarantee learnability, it can be too pessimistic, resulting in worst-case rates.

In this work we are interested in the data-dependent behavior of SGD, thus the emphasis will fall on the distribution-dependent notion of stability, known as *on-average* stability, explored thoroughly in (Shalev-Shwartz et al., 2010). The attractive quality of this less restrictive stability type is that the resulting bounds are controlled by how stable the algorithm is under the data-generating distribution. For instance, in (Bousquet & Elisseeff, 2002) and (Devroye & Wagner, 1979), the on-average stability is related to the variance of an estimator. In (Shalev-Shwartz & Ben-David, 2014, Sec. 13), the authors show risk bounds that depend on the expected empirical risk of a solution to the regularized ERM. In turn, one can exploit this fact to state improved *optimistic* risk

bounds, for instance, ones that exhibit *fast-rate* regimes (Koren & Levy, 2015; Gonen & Shalev-Shwartz, 2017), or even to design enhanced algorithms that minimize these bounds in a data-driven way, e.g. by exploiting side information as in transfer (Kuzborskij & Orabona, 2013; Ben-David & Urner, 2013) and metric learning (Perrot & Habrard, 2015). Here, we mainly focus on the later direction in the context of SGD: how stable is SGD under the data-generating distribution given an initialization point? We also touch the former direction by taking advantage of our data-driven analysis and show optimistic bounds as a corollary.

We will study the on-average stability of SGD for both convex and non-convex loss functions. In the convex setting, we will relate stability to the risk at the initialization point, while previous data-driven stability arguments usually consider minimizers of convex ERM rather than a stochastic approximation (Shalev-Shwartz & Ben-David, 2014; Koren & Levy, 2015). Beside convex problems, our work also covers the generalization ability of SGD on non-convex problems. Here, we borrow techniques of (Hardt et al., 2016) and extend them to the distribution-dependent setting. That said, while bounds of (Hardt et al., 2016) are stated in terms of worst-case quantities, ours reveal new connections to the data-dependent second-order information. These new insights also partially justify empirical observations in deep learning about the link between the curvature and the generalization error (Hochreiter & Schmidhuber, 1997; Keskar et al., 2017; Chaudhari et al., 2017). At the same time, our work is an alternative to the theoretical studies of neural network objective functions (Choromanska et al., 2015; Kawaguchi, 2016), as we focus on the direct connection between the generalization and the curvature.

In this light, our work is also related to non-convex optimization by SGD. Literature on this subject typically studies rates of convergence to the stationary points (Ghadimi & Lan, 2013; Allen-Zhu & Hazan, 2016; Reddi et al., 2016), and ways to avoid saddles (Ge et al., 2015; Lee et al., 2016). However, unlike these works, and similarly to (Hardt et al., 2016), we are interested in the generalization ability of SGD, and thanks to the stability approach, involvement of stationary points in our analysis is not necessary.

Finally, we propose an example application of our findings in Transfer Learning (TL). For instance, by controlling the stability bound in a data-driven way, one can choose an initialization that leads to improved generalization. This is related to TL where one transfers from pre-trained models (Kuzborskij & Orabona, 2016; Tommasi et al., 2014; Pentina & Lampert, 2014; Ben-David & Urner, 2013), especially popular in deep learning due to its data-demanding nature (Galanti et al., 2016). Literature on this topic is mostly focused on the ERM setting and PAC-bounds, while our analysis of SGD yields such guarantees as a corollary.

3. Stability of Stochastic Gradient Descent

First, we introduce definitions used in the rest of the paper.

3.1. Definitions

We will denote with small and capital bold letters respectively column vectors and matrices, e.g., $\mathbf{a} = [a_1, a_2, \dots, a_d]^T \in \mathbb{R}^d$ and $\mathbf{A} \in \mathbb{R}^{d_1 \times d_2}$, $\|\mathbf{a}\|$ is understood as a Euclidean norm and $\|\mathbf{A}\|_2$ as the spectral norm. We denote enumeration by $[n] = \{1, \dots, n\}$ for $n \in \mathbb{N}$.

We indicate an example space by \mathcal{Z} and its member by $z \in \mathcal{Z}$. For instance, in a supervised setting $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$, such that \mathcal{X} is the input and \mathcal{Y} is the output space of a learning problem. We assume that training and testing examples are drawn iid from a probability distribution \mathcal{D} over \mathcal{Z} . In particular, we will denote the training set as $S = \{z_i\}_{i=1}^m \sim \mathcal{D}^m$.

For a parameter space \mathcal{H} , we define a learning algorithm as a map $A : \mathcal{Z}^m \mapsto \mathcal{H}$ and for brevity we will use the notation $A_S = A(S)$. In the following we assume that $\mathcal{H} \subseteq \mathbb{R}^d$. To measure the accuracy of a learning algorithm A , we have a *loss* function $f(\mathbf{w}, z)$, which measures the cost incurred by predicting with parameters $\mathbf{w} \in \mathcal{H}$ on an example z . The *risk* of \mathbf{w} , with respect to the distribution \mathcal{D} , and the *empirical risk* given a training set S are defined as

$$R(\mathbf{w}) := \mathbb{E}_{z \sim \mathcal{D}} [f(\mathbf{w}, z)], \text{ and } \hat{R}_S(\mathbf{w}) := \frac{1}{m} \sum_{i=1}^m f(\mathbf{w}, z_i).$$

Finally, define $R^* := \inf_{\mathbf{w} \in \mathcal{H}} R(\mathbf{w})$.

3.2. Uniform Stability and Generalization

On an intuitive level, a learning algorithm is said to be *stable* whenever a small perturbation in the training set does not affect its outcome too much. Of course, there is a number of ways to formalize the perturbation and the extent of the change in the outcome, and we will discuss some of them below. The most important consequence of a stable algorithm is that it *generalizes* from the training set to the unseen data sampled from the same distribution. In other words, the difference between the risk $R(A_S)$ and the empirical risk $\hat{R}_S(A_S)$ of the algorithm's output is controlled by the quantity that captures how stable the algorithm is. So, to observe good performance, or a decreasing true risk, we must have a stable algorithm *and* decreasing empirical risk (training error), which usually comes by design of the algorithm. In this work we focus on the stability of the Stochastic Gradient Descent (SGD) algorithm, and thus, as a consequence, we study its generalization ability.

Recently, (Hardt et al., 2016) used a stability argument to prove generalization bounds for learning with SGD. Specifically, the authors extended the notion of the *uniform stability*

originally proposed by (Bousquet & Elisseeff, 2002), to accommodate randomized algorithms.

Definition 1 (Uniform stability). *A randomized algorithm A is ϵ -uniformly stable if for all datasets $S, S^{(i)} \in \mathcal{Z}^m$ such that S and $S^{(i)}$ differ in the i -th example, we have*

$$\sup_{z \in \mathcal{Z}, i \in [m]} \left\{ \mathbb{E}_A [f(A_S, z) - f(A_{S^{(i)}}, z)] \right\} \leq \epsilon.$$

Since SGD is a randomized algorithm, we have to cope with two sources of randomness: the data-generating process and the randomization of the algorithm A itself, hence we have statements in expectation. The following theorem of (Hardt et al., 2016) shows that the uniform stability implies generalization in expectation.

Theorem 1. *Let A be ϵ -uniformly stable. Then,*

$$\left| \mathbb{E}_{S, A} \left[\hat{R}_S(A_S) - R(A_S) \right] \right| \leq \epsilon.$$

Thus it suffices to characterize the uniform stability of an algorithm to state a generalization bound. In particular, (Hardt et al., 2016) showed generalization bounds for SGD under different assumptions on the loss function f . Despite that these results hold in expectation, other forms of generalization bounds, such as high-probability ones, can be derived from the above (Shalev-Shwartz et al., 2010).

Apart from SGD, uniform stability has been used before to prove generalization bounds for many learning algorithms (Bousquet & Elisseeff, 2002). However, these bounds typically suggest worst-case generalization rates, and rather reflect intrinsic stability properties of an algorithm. In other words, uniform stability is oblivious to the data-generating process and any other side information, which might reveal scenarios where generalization occurs at a faster rate. In turn, these insights could motivate the design of improved learning algorithms. In the following we address some limitations of analysis through uniform stability by using a less restrictive notion of stability. We extend the setting of (Hardt et al., 2016) by proving data-dependent stability bounds for convex and non-convex loss functions. In addition, we also take into account the initialization point of an algorithm as a form of supplementary information, and we dedicate special attention to its interplay with the data-generating distribution. Finally, we discuss situations where one can explicitly control the stability of SGD in a data-dependent way.

4. Data-dependent Stability Bounds for SGD

In this section we describe a notion of data-dependent algorithmic stability, that allows us to state generalization bounds which depend not only on the properties of the learning algorithm, but also on the additional parameters of the

algorithm. We indicate such additional parameters by θ , and therefore we denote stability as a function $\epsilon(\theta)$. In particular, in the following we will be interested in scenarios where θ describes the data-generating distribution and the initialization point of SGD.

Definition 2 (On-Average stability). *A randomized algorithm A is $\epsilon(\theta)$ -on-average stable if it is true that*

$$\sup_{i \in [m]} \left\{ \mathbb{E}_{A, S, z} \mathbb{E} [f(A_S, z) - f(A_{S^{(i)}}, z)] \right\} \leq \epsilon(\theta),$$

where $S \stackrel{iid}{\sim} \mathcal{D}^m$ and $S^{(i)}$ is its copy with i -th example replaced by $z \stackrel{iid}{\sim} \mathcal{D}$.

Our definition of on-average stability resembles the notion introduced by (Shalev-Shwartz et al., 2010). The difference lies in the fact that we take supremum over index of replaced example. A similar notion was also used by (Bousquet & Elisseeff, 2002) and later by (Elisseeff et al., 2005) for analysis of a randomized aggregation schemes, however their definition involves absolute difference of losses. The dependence on θ also bears similarity to recent work of (London, 2016), however, there, it is used in the context of uniform stability. The following theorem shows that on-average -stable random algorithm is guaranteed to generalize in expectation.

Theorem 2. *Let an algorithm A be $\epsilon(\theta)$ -on-average stable. Then,*

$$\mathbb{E}_A \mathbb{E} \left[R(A_S) - \widehat{R}_S(A_S) \right] \leq \epsilon(\theta).$$

5. Main Results

Before presenting our main results in this section, we discuss algorithmic details and assumptions. We will study the following variant of SGD: given a training set $S = \{z_i\}_{i=1}^m \stackrel{iid}{\sim} \mathcal{D}^m$, step sizes $\{\alpha_t\}_{t=1}^T$, random indices $I = \{j_t\}_{t=1}^T$, and an initialization point \mathbf{w}_1 , perform updates

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

for $T \leq m$ steps. Moreover we will use the notation $\mathbf{w}_{S,t}$ to indicate the output of SGD ran on a training set S , at step t . We assume that the indices in I are sampled from the uniform distribution over $[m]$ without replacement, and that this is the only source of randomness for SGD. In practice this corresponds to permuting the training set before making a pass through it, as it is commonly done in practical applications. We also assume that the variance of stochastic gradients obeys

$$\mathbb{E}_{S,z} \left[\|\nabla f(\mathbf{w}_{S,t}, z) - \nabla R(\mathbf{w}_{S,t})\|^2 \right] \leq \sigma^2 \quad \forall t \in [T].$$

Next, we introduce statements about the loss functions f used in the following.

Definition 3 (Lipschitz f). *A loss function f is L -Lipschitz if $\|\nabla f(\mathbf{w}, z)\| \leq L, \forall \mathbf{w} \in \mathcal{H}$ and $\forall z \in \mathcal{Z}$. Note that this also implies that $|f(\mathbf{w}, z) - f(\mathbf{v}, z)| \leq L\|\mathbf{w} - \mathbf{v}\|$.*

Definition 4 (Smooth f). *A loss function is β -smooth if $\forall \mathbf{w}, \mathbf{v} \in \mathcal{H}$ and $\forall z \in \mathcal{Z}$, $\|\nabla f(\mathbf{w}, z) - \nabla f(\mathbf{v}, z)\| \leq \beta\|\mathbf{w} - \mathbf{v}\|$, which also implies $f(\mathbf{w}, z) - f(\mathbf{v}, z) \leq \nabla f(\mathbf{v}, z)^\top (\mathbf{w} - \mathbf{v}) + \frac{\beta}{2}\|\mathbf{w} - \mathbf{v}\|^2$.*

Definition 5 (Lipschitz Hessians). *A loss function f has a ρ -Lipschitz Hessian if $\forall \mathbf{w}, \mathbf{v} \in \mathcal{H}$ and $\forall z \in \mathcal{Z}$, $\|\nabla^2 f(\mathbf{w}, z) - \nabla^2 f(\mathbf{v}, z)\|_2 \leq \rho\|\mathbf{w} - \mathbf{v}\|$.*

The last condition is occasionally used in analysis of SGD (Ge et al., 2015) and holds whenever f has a bounded third derivative. All presented theorems assume that the loss function is non-negative, Lipschitz, and β -smooth. Examples of such commonly used loss functions are the logistic/softmax losses and neural networks with sigmoid activations. Convexity of loss functions or Lipschitzness of Hessians will only be required for some results, and we will denote it when necessary. Proofs for all the statements in this section are given in the arXiv version of the paper¹.

5.1. Convex Losses

First, we present a new and data-dependent stability result for convex losses.

Theorem 3. *Assume that f is convex, and that SGD's step sizes satisfy $\alpha_t = \frac{c}{\sqrt{t}} \leq \frac{1}{\beta}, \forall t \in [T]$. Then SGD is $\epsilon(\mathcal{D}, \mathbf{w}_1)$ -on-average stable with*

$$\epsilon(\mathcal{D}, \mathbf{w}_1) = \mathcal{O} \left(\sqrt{c(R(\mathbf{w}_1) - R^*)} \cdot \frac{\sqrt[4]{T}}{m} + c\sigma \frac{\sqrt{T}}{m} \right).$$

Under the same assumptions, taking step size of order $\mathcal{O}(1/\sqrt{t})$, Theorem 3.7 of (Hardt et al., 2016) implies a uniform stability bound $\epsilon = \mathcal{O}(\sqrt{T}/m)$. Our bound differs since it involves a multiplicative risk at the initialization point. Thus, our bound corroborates the intuition that whenever we start at a good location of the objective function, the algorithm is more stable and thus generalizes better. However, this is only the case, whenever the variance of stochastic gradient σ^2 is not too large. In the deterministic case, and of $R(\mathbf{w}_1) = R^*$, the theorem confirms that SGD, in expectation, does not need to make any updates and is therefore perfectly stable. On the other hand, when the variance σ^2 is large enough to make the second summand in Theorem 3 dominant, the bound does not offer improvement compared to (Hardt et al., 2016). Note, that a result of this type cannot be obtained through the more restrictive uniform stability, precisely because such bounds on the stability must hold even for a worst-case choice of data distribution and initialization. In contrast, the notion of stability we

¹<https://arxiv.org/abs/1703.01678>

empoly depends on the data-generating distribution, which allowed us to introduce dependency on the risk.

Furthermore, consider that we start at arbitrary location \mathbf{w}_1 : assuming that the loss function is bounded for a concrete \mathcal{H} and \mathcal{Z} , the rate of our bound up to a constant is no worse than that of (Hardt et al., 2016). Finally, one can always tighten this result by taking the minimum of two bounds.

5.2. Non-convex Losses

Now we state a new stability result for non-convex losses.

Theorem 4. *Assume that $f(\cdot, z) \in [0, 1]$ and has a ρ -Lipschitz Hessian, and that step sizes of a form $\alpha_t = \frac{c}{t}$ satisfy $c \leq \min \left\{ \frac{1}{\beta}, \frac{1}{4(2\beta \ln(T))^2} \right\}$. Then SGD is $\epsilon(\mathcal{D}, \mathbf{w}_1)$ -on-average stable with*

$$\epsilon(\mathcal{D}, \mathbf{w}_1) \leq \frac{1 + \frac{1}{c\gamma}}{m} (2cL^2)^{\frac{1}{1+c\gamma}} \left(\mathbb{E}_{S,A} [R(A_S)] \cdot T \right)^{\frac{c\gamma}{1+c\gamma}}, \quad (1)$$

where

$$\gamma := \tilde{\mathcal{O}} \left(\min \left\{ \beta, \mathbb{E}_z [\|\nabla^2 f(\mathbf{w}_1, z)\|_2] + \Delta_{1,\sigma^2}^* \right\} \right), \quad (2)$$

$$\Delta_{1,\sigma^2}^* := \rho \left(c\sigma + \sqrt{c(R(\mathbf{w}_1) - R^*)} \right).$$

In particular, γ characterizes how the curvature at the initialization point affects stability, and hence the generalization error of SGD. Since γ heavily affects the rate of convergence in (1), and in most situations smaller γ yields higher stability, we now look at a few cases of its behavior. Consider a regime such that γ is of the order $\tilde{\Theta} \left(\mathbb{E} [\|\nabla^2 f(\mathbf{w}_1, z)\|_2] + \sqrt{R(\mathbf{w}_1)} + \sigma \right)$, or in other words, that stability is controlled by the curvature, the risk of the initialization point \mathbf{w}_1 , and the variance of the stochastic gradient σ^2 . This suggests that starting from a point in a less curved region with low risk should yield higher stability, and therefore as predicted by our theory, allow for faster generalization. In addition, we observe that the considered stability regime offers a principled way to pre-screen a good initialization point in practice, by choosing the one that minimizes spectral norm of the Hessian and the risk.

Next, we focus on a more specific case. Suppose that we choose a step size $\alpha_t = \frac{c}{t}$ such that $\gamma = \tilde{\Theta} \left(\mathbb{E} [\|\nabla^2 f(\mathbf{w}_1, z)\|_2] \right)$, yet not too small, so that the empirical risk can still be decreased. Then, stability is dominated by the curvature around \mathbf{w}_1 . Indeed, lower generalization errors on non-convex problems, such as training deep neural networks, have been observed empirically when SGD is actively guided (Hochreiter & Schmidhuber, 1997; Goodfellow et al., 2016; Chaudhari et al., 2017) or converges to solutions with low curvature (Keskar et al., 2017). However, to the best of our knowledge, Theorem 4 is the first

to establish a theoretical link between the curvature of the loss function and the generalization ability of SGD in a data-dependent sense.

Theorem 4 immediately implies the following statement that further reinforces the effect of the initialization point on the generalization error, assuming that $\mathbb{E}_S [R(A_S)] \leq R(\mathbf{w}_1)$.

Corollary 1. *Under conditions of Theorem 4 we have that SGD is $\epsilon(\mathcal{D}, \mathbf{w}_1)$ -on-average stable with*

$$\epsilon(\mathcal{D}, \mathbf{w}_1) = \mathcal{O} \left(\frac{1 + \frac{1}{c\gamma}}{m} (R(\mathbf{w}_1) \cdot T)^{\frac{c\gamma}{1+c\gamma}} \right). \quad (3)$$

We take a moment to discuss the role of the risk term in $(R(\mathbf{w}_1) \cdot T)^{\frac{c\gamma}{1+c\gamma}}$. Observe that $\epsilon(\mathcal{D}, \mathbf{w}_1) \rightarrow 0$ as $R(\mathbf{w}_1) \rightarrow 0$, in other words, the generalization error approaches zero as the risk of the initialization point vanishes. This is an intuitive behavior, however, uniform stability does not capture this due to its distribution-free nature. Finally, we note that (Hardt et al., 2016, Theorem 3.8) showed a bound similar to (1), however, in place of γ their bound has a Lipschitz constant of the gradient. The crucial difference lies in term γ which is now not merely a Lipschitz constant, but rather depends on the data-generating distribution and initialization point of SGD. We compare to their bound by considering the worst case scenario, namely, that SGD is initialized in a point with high curvature, or altogether, that the objective function is highly curved everywhere. Then, at least our bound is no worse than the one of (Hardt et al., 2016), since $\gamma \leq \beta$. Finally, it should be noted that our bound can be compared to the one of (Hardt et al., 2016) only in a setting of a single pass. In a multiple-pass case, data-dependent analysis in a current form would not hold, since the output of SGD would not be independent from a newly observed example after the first pass. On the other hand, our results focus on the gains due to data-dependent initialization.


Theorem 4 also allows us to prove an optimistic generalization bound for learning with SGD on non-convex objectives.

Corollary 2. *Under conditions of Theorem 4 we have that the output of SGD obeys*

$$\mathbb{E}_{S,A} [R(A_S) - \hat{R}_S(A_S)] = \mathcal{O} \left(\frac{1 + \frac{1}{c\gamma}}{m} \cdot \max \left\{ \left(\mathbb{E}_{S,A} [\hat{R}_S(A_S)] \cdot T \right)^{\frac{c\gamma}{1+c\gamma}}, \left(\frac{T}{m} \right)^{c\gamma} \right\} \right).$$

An important consequence of Corollary 2, is that for a vanishing expected empirical risk, in particular for $\mathbb{E}_{S,A} [\hat{R}_S(A_S)] = \mathcal{O} \left(\frac{T^{c\gamma}}{m^{1+c\gamma}} \right)$, the generalization error behaves as $\mathcal{O} \left(\frac{T^{c\gamma}}{m^{1+c\gamma}} \right)$. Considering the full pass, that is $m = \mathcal{O}(T)$, we have an optimistic generalization error of order $\mathcal{O}(1/m)$ instead of $\mathcal{O}(m^{-\frac{1}{1+c\gamma}})$. We note that

Figure 1. Comparison of data-dependent and uniform generalization bounds evaluated by training a neural network.



PAC bounds with similar optimistic message (although not directly comparable), but without curvature information can also be obtained through empirical Bernstein bounds as in (Maurer & Pontil, 2009). However, a PAC bound does not suggest a way to minimize non-convex empirical risk in general, where SGD is known to work reasonably well.

5.2.1. TIGHTNESS OF NON-CONVEX BOUNDS

Next we empirically assess the tightness of our non-convex generalization bounds on real data. In the following experiment we train a neural network with three convolutional layers interlaced with max-pooling, followed by the fully connected layer with 16 units, on the MNIST dataset. This totals in a model with 18K parameters. Figure 1 compares our data-dependent bound (1) to the distribution-free one of (Hardt et al., 2016, Theorem 3.8). As a reference we also include an empirical estimate of the generalization error taken as an absolute difference of the validation and training average losses. Since our bound also depends on the initialization point, we plot (1) for multiple “warm-starts”, ie. with SGD initialized from a pre-trained position. We consider 7 such warm-starts at every 200 steps, and report data-dependent quantities used to compute (1) just beneath the graph. Our first observation is that, clearly, the data-dependent bound gives tighter estimate, by roughly one order of magnitude. Second, simulating start from a pre-trained position suggests even tighter estimates: we suspect that this is due to decreasing validation error which is used as an empirical estimate for $R(\mathbf{w}_1)$ which affects bound (1).

We compute an empirical estimate of the expected Hessian spectral norm by the power iteration method using an efficient Hessian-vector multiplication method (Pearlmuter, 1994). Since bounds depend on constants L , β , and ρ , we estimate them by tracking maximal values of the

gradient and Hessian norms throughout optimization. We compute bounds with estimates $\hat{L} = 78.72$, $\hat{\beta} = 1692.28$, $\hat{\rho} = 3823.73$, and $c = 10^{-3}$.

5.3. Application to Transfer Learning

One example application of data-dependent bounds presented before lies in *Transfer Learning (TL)*, where we are interested in achieving faster generalization on a *target* task by exploiting side information that originates from different but related *source* tasks. The literature on TL explored many ways to do so, and here we will focus on the one that is most compatible with our bounds. More formally, suppose that the *target* task at hand is characterized by a joint probability distribution \mathcal{D} , and as before we have a training set $\mathcal{S} \stackrel{\text{iid}}{\sim} \mathcal{D}^m$. Some TL approaches also assume access to the data sampled from the distributions associated with the *source* tasks. Here we follow a conservative approach – instead of the source data, we receive a set of *source* hypotheses $\{\mathbf{w}_k^{\text{src}}\}_{k=1}^K \subset \mathcal{H}$, trained on the source tasks. The goal of a learner is to come up with a target hypothesis, which in the optimistic scenario generalizes better by relying on source hypotheses. In the TL literature this is known as Hypothesis Transfer Learning (HTL) (Kuzborskij & Orabona, 2016), that is, we transfer from the source hypotheses which act as a proxy to the source tasks and the risk $R(\mathbf{w}_k^{\text{src}})$ quantifies how much source and target tasks are related. In the following we will consider SGD for HTL, where the source hypotheses act as initialization points. First, consider learning with convex losses: Theorem 3 depends on $R(\mathbf{w}_1)$, thus it immediately quantifies the relatedness of source and target tasks. So it is enough to pick the point that minimizes the stability bound to transfer from the most related source. Then, bounding $R(\mathbf{w}_k^{\text{src}})$ by $\hat{R}_S(\mathbf{w}_k^{\text{src}})$ through Hoeffding bound along with union bound gives with high probability that

$$\min_{k \in [K]} \epsilon(\mathcal{D}, \mathbf{w}_k^{\text{src}}) \leq \min_{k \in [K]} \mathcal{O} \left(\hat{R}_S(\mathbf{w}_k^{\text{src}}) + \sqrt{\frac{\log(K)}{m}} \right).$$

Hence, the most related source is the one that simply minimizes empirical risk. Similar conclusions were drawn in HTL literature, albeit in the context of ERM. Matters are slightly more complicated in the non-convex case. We take a similar approach, however, now we minimize stability bound (3), and for the sake of simplicity assume that we make a full pass over the data, so $T = m$. Minimizing the following empirical upper bound select the best source.

Proposition 1. Let $\hat{\gamma}_k^\pm = \Theta \left(\frac{1}{m} \sum_{i=1}^m \|\nabla^2 f(\mathbf{w}_k^{\text{src}}, z_i)\|_2 + \sqrt{\hat{R}_S(\mathbf{w}_k^{\text{src}}) \pm \sqrt{\log(K)/m}} \right)$. Then with high probability the generalization error of $\mathbf{w}_k^{\text{src}}$ is bounded by

$$\min_{k \in [K]} \mathcal{O} \left(\left(1 + \frac{1}{c\hat{\gamma}_k^-} \right) \hat{R}_S(\mathbf{w}_k^{\text{src}})^{\frac{c\hat{\gamma}_k^+}{1+c\hat{\gamma}_k^+}} \cdot \frac{\sqrt{\log(K)}}{m^{\frac{1}{1+c\hat{\gamma}_k^+}}} \right).$$

Note that $\hat{\gamma}_k^\pm$ involves estimation of the spectral norm of the Hessian, which is computationally cheaper to evaluate compared to the complete Hessian matrix (Pearlmutter, 1994). This is particularly relevant for deep learning, where computation of the Hessian matrix can be prohibitively expensive.

6. Proof Outline of Theorem 4

In this section we discuss the proof of Theorem 4, which bounds data-dependent stability for non-convex losses.

We say that the SGD gradient update rule is an operator $G_t : \mathcal{H} \mapsto \mathcal{H}$, such that

$$G_t(\mathbf{w}) := \mathbf{w} - \alpha_t \nabla f(\mathbf{w}, z_{i_t}),$$

and it is also a function of the training set S and a random index set I . Then, $\mathbf{w}_{t+1} = G_t(\mathbf{w}_t)$, throughout $t = 1, \dots, T$. Recall the use of notation $\mathbf{w}_{S,t}$ to indicate the output of SGD ran on a training set S , at step t , and define

$$\delta_t(S, z) := \|\mathbf{w}_{S,t} - \mathbf{w}_{S^{(i)},t}\|.$$

The following property of G_t will be central to our proof.

Definition 6 (Expansiveness). *A gradient update rule is η -expansive if for all \mathbf{w}, \mathbf{v} , $\|G_t(\mathbf{w}) - G_t(\mathbf{v})\| \leq \eta \|\mathbf{w} - \mathbf{v}\|$.*

The following lemma characterizes expansiveness for the gradient update rule under different assumptions on f .

Lemma 1 ((Hardt et al., 2016)). *Assume that f is β -smooth. Then, we have that G_t is $(1 + \alpha_t \beta)$ -expansive.*

The following lemma is similar to Lemma 3.11 of (Hardt et al., 2016), and is instrumental in bounding the stability of SGD. However, we make an adjustment and state it in expectation over the data. Note that it does not require convexity of the loss function.

Lemma 2. *Assume that the loss function $f(\cdot, z) \in [0, 1]$ is L -Lipschitz for all z . Then, for every $t_0 \in \{0, 1, 2, \dots, m\}$ we have that,*

$$\mathbb{E}_{S,z} \mathbb{E}_A [f(\mathbf{w}_{S,T}, z) - f(\mathbf{w}_{S^{(i)},T}, z)] \quad (4)$$

$$\leq L \mathbb{E}_{S,z} \left[\mathbb{E}_A [\delta_T(S, z) \mid \delta_{t_0}(S, z) = 0] \right] + \mathbb{E}_{S,A} [R(A_S)] \frac{t_0}{m}. \quad (5)$$

We spend a moment to highlight the role of conditional expectation. Observe that we could naively bound (4) by the Lipschitzness of f , but Lemma 2 follows a more careful argument. First note that t_0 is a free parameter. The expected distance in (5) between SGD outputs $\mathbf{w}_{S,t}$ and $\mathbf{w}_{S^{(i)},t}$ is conditioned on the fact that at step t_0 outputs of SGD are still the same. This means that the perturbed point is encountered after t_0 . Then, the conditional expectation should be a decreasing function of t_0 : the later the perturbation occurs, the smaller deviation between $\mathbf{w}_{S,t}$ and $\mathbf{w}_{S^{(i)},t}$ we should expect. Eventually we minimize (5) over t_0 .

6.1. Non-convex Losses

Our proof of a stability bound for non-convex loss functions, Theorem 4, follows a general outline of (Hardt et al., 2016, Theorem 3.8). Namely, the outputs of SGD run on a training set S and its perturbed version $S^{(i)}$ will not differ too much, because by the time a perturbation is encountered, the step size has already decayed enough. So, on the one hand, stabilization is enforced by the diminishing the step size, and on the other hand, by how much updates expand the distance between the gradients after the perturbation. Since (Hardt et al., 2016) work with uniform stability, they capture the expansiveness of post-perturbation update by the Lipschitzness of the gradient. In combination with a recursive argument, their bound has exponential dependency on the Lipschitz constant of the gradient. We argue that the Lipschitz continuity of the gradient can be too pessimistic in general. Instead, we rely on a local data-driven argument: considering that we initialize SGD at point \mathbf{w}_1 , how much do updates expand the gradient under the distribution of interest? The following crucial lemma characterizes such behavior in terms of the curvature at \mathbf{w}_1 .

Lemma 3. *Assume that the loss function $f(\cdot, z)$ is β -smooth and that its Hessian is ρ -Lipschitz. Then,*

$$\|G_t(\mathbf{w}_{S,t}) - G_t(\mathbf{w}_{S^{(i)},t})\| \leq (1 + \alpha_t \xi_t(S, z)) \delta_t(S, z)$$

Furthermore, for any $t \in [T]$,

$$\begin{aligned} \mathbb{E}_{S,z} [\xi_t(S, z)] &= \tilde{O} \left(\mathbb{E}_{S,z} [\|\nabla^2 f(\mathbf{w}_1, z_t)\|_2] + \Delta_{1,\sigma^2}^* \right), \\ \Delta_{1,\sigma^2}^* &:= \rho \left(c\sigma + \sqrt{c(R(\mathbf{w}_1) - R^*)} \right). \end{aligned}$$

Next, we need the following statement.

Proposition 2 (Bernstein-type inequality). *Let Z be a zero-mean real-valued r.v., such that $|Z| \leq b$ and $\mathbb{E}[Z^2] \leq \sigma^2$. Then for all $|c| \leq \frac{1}{2b}$, we have that $\mathbb{E}[e^{cZ}] \leq e^{c^2 \sigma^2}$.*

Now we are ready to prove Theorem 4.

Sketch proof of Theorem 4. We start from Lemma 2. Most of the proof is dedicated to bounding the first term in (5). We deal with this similarly as in (Hardt et al., 2016). Specifically, we state the bound on $\mathbb{E}_A [\delta_T(S, z) \mid \delta_{t_0}(S, z) = 0]$ by using a recursion. In our case, however, we also have an expectation w.r.t. the data, and to avoid complications with dependencies, we first unroll the recursion for the random quantities, and only then take the expectation. At this point the proof crucially relies on the product of exponentials arising from the recursion, and all relevant random quantities end up inside of them. We alleviate this by Proposition 2. Finally, we conclude by minimizing (5) w.r.t. t_0 . Thus we have three steps: 1) recursion, 2) bounding $\mathbb{E}[\exp(\dots)]$, and 3) tuning of t_0 .

1) Recursion. We begin by stating the bound on $\mathbb{E}_A [\delta_T(S, z) | \delta_{t_0}(S, z) = 0]$ by recursion. Thus we will first state the bound on $\mathbb{E}_A [\delta_{t+1}(S, z) | \delta_{t_0}(S, z) = 0]$ in terms of $\mathbb{E}_A [\delta_t(S, z) | \delta_{t_0}(S, z) = 0]$, and other relevant quantities and then unravel the recursion. We distinguish two cases: 1) SGD encounters the perturbed point at step t , that is $t = i$, and 2) the current point is the same in S and $S^{(i)}$, so $t \neq i$. For the first case, we will use worst-case boundedness of G_t , that is, we observe that $\|G_t(\mathbf{w}_{S,t}) - G_t(\mathbf{w}_{S^{(i)},t})\| \leq \delta_t(S, z) + 2\alpha_t L$. To handle the second case we will use Lemma 3, namely,

$$\|G_t(\mathbf{w}_{S,t}) - G_t(\mathbf{w}_{S^{(i)},t})\| \leq (1 + \alpha_t \xi_t(S, z)) \delta_t(S, z).$$

In addition, as a safety measure we will also take into account that the gradient update rule is at most $(1 + \alpha_t \beta)$ -expansive by Lemma 1. So we will work with the function $\psi_t(S, z) := \min\{\xi_t(S, z), \beta\}$ instead of $\xi_t(S, z)$. Now, introduce $\Delta_t(S, z) := \mathbb{E}_A[\delta_t(S, z) | \delta_{t_0}(S, z) = 0]$, and decompose the expectation w.r.t. A for a step t . Noting that SGD encounters the perturbed example with probability $\frac{1}{m}$,

$$\begin{aligned} \Delta_{t+1}(S, z) &\leq \left(1 - \frac{1}{m}\right) (1 + \alpha_t \psi_t(S, z)) \Delta_t(S, z) \\ &\quad + \frac{1}{m} (2\alpha_t L + \Delta_t(S, z)) \\ &= \left(1 + \left(1 - \frac{1}{m}\right) \alpha_t \psi_t(S, z)\right) \Delta_t(S, z) + \frac{2\alpha_t L}{m} \\ &\leq \exp(\alpha_t \psi_t(S, z)) \Delta_t(S, z) + \frac{2\alpha_t L}{m}, \end{aligned} \quad (6)$$

where the last inequality follows from $1 + x \leq \exp(x)$. This inequality is not overly loose for $x \in [0, 1]$, and, in our case it becomes instrumental in handling the recursion.

Now, observe that relation $x_{t+1} \leq a_t x_t + b_t$ with $x_{t_0} = 0$ unwinds from T to t_0 as $x_T \leq \sum_{t=t_0+1}^T b_t \prod_{k=t+1}^T a_k$. Consequently, having $\Delta_{t_0}(S, z) = 0$, we unwind (6) to get

$$\begin{aligned} \Delta_T(S, z) &\leq \sum_{t=t_0+1}^T \left(\prod_{k=t+1}^T \exp\left(\frac{c\psi_k(S, z)}{k}\right) \right) \frac{2cL}{mt} \\ &= \sum_{t=t_0+1}^T \exp\left(c \sum_{k=t+1}^T \frac{\psi_k(S, z)}{k}\right) \frac{2cL}{mt}. \end{aligned} \quad (7)$$

2) Bounding $\mathbb{E}[\exp(\dots)]$. We take expectation w.r.t. S and z on both sides and focus on the expectation of the exponential in (7). First, introduce $\mu_k := \mathbb{E}_{S,z}[\psi_k(S, z)]$, and observe that the zero-mean version of $\psi_k(S, z)$ is bounded as $\sum_{k=t+1}^T \frac{1}{k} |\psi_k(S, z) - \mu_k| \leq 2\beta \ln(T)$. Assuming the setting of $c \leq \frac{1}{2(2\beta \ln(T))^2}$, we apply Proposition 2 and get

$$\mathbb{E}_{S,z} \left[\exp\left(c \sum_{k=t+1}^T \frac{\psi_k(S, z)}{k}\right) \right] \leq \exp\left(c \sum_{k=t+1}^T \frac{2\mu_k}{k}\right), \quad (8)$$

where we bounded variance by μ_k thanks to the setting of c . Next, we give an upper-bound on μ_k , that is $\mu_k \leq \min\{\beta, \mathbb{E}_{S,z}[\xi_k(S, z)]\}$. Finally, we bound $\mathbb{E}_{S,z}[\xi_k(S, z)]$ using the second result of Lemma 3, which holds for any $k \in [T]$, to get that $\mu_k \leq \gamma$, with γ defined in (2).

3) Tuning of t_0 . Now we turn our attention back to (7). Considering that we took an expectation w.r.t. the data, we use (8) and the fact that $\mu_k \leq \gamma$ to get that

$$\begin{aligned} \mathbb{E}_{S,z} [\Delta_T(S, z)] &\leq \sum_{t=t_0+1}^T \exp\left(2c\gamma \sum_{k=t+1}^T \frac{1}{k}\right) \frac{2cL}{mt} \\ &\leq \sum_{t=t_0+1}^T \exp\left(2c\gamma \ln\left(\frac{T}{t}\right)\right) \frac{2cL}{mt} \leq \frac{L}{\gamma m} \left(\frac{T}{t_0}\right)^{2c\gamma}. \end{aligned}$$

Plug the above into (5) to get

$$\begin{aligned} \mathbb{E}_{S,z} \mathbb{E}_A [f(\mathbf{w}_{S,T}, z) - f(\mathbf{w}_{S^{(i)},T}, z)] \\ \leq \frac{L^2}{\gamma m} \left(\frac{T}{t_0}\right)^{2c\gamma} + \frac{t_0}{m}. \end{aligned} \quad (9)$$

Let $q = 2c\gamma$. Then, setting $t_0 = (2cL^2)^{\frac{1}{1+q}} T^{\frac{q}{1+q}}$ minimizes (9). Plugging t_0 back we get that (9) equals to $\frac{1+\frac{1}{q}}{m} (2cL^2)^{\frac{1}{1+q}} T^{\frac{q}{1+q}}$. This completes the proof. \square

7. Conclusions and Future Work

In this work we proved data-dependent stability bounds for SGD and revisited its generalization ability. We presented novel bounds for convex and non-convex smooth loss functions, partially controlled by data-dependent quantities, while previous stability bounds for SGD were derived through the worst-case analysis. In particular, for non-convex learning, we demonstrated theoretically that generalization of SGD is heavily affected by the expected curvature around the initialization point. We demonstrated empirically that our bound is indeed tighter compared to the uniform one. In addition, our data-dependent analysis also allowed us to show optimistic bounds on the generalization error of SGD, which exhibit fast rates subject to the vanishing empirical risk of the algorithm's output.

In future work we further intend to explore our theoretical findings experimentally and evaluate the feasibility of the transfer learning based on the second-order information. Another direction lies in making our bounds adaptive. So far we have presented bounds that have data-dependent components, however the step size cannot be adjusted depending on the data, e.g. as in (Zhao & Zhang, 2015). This was partially addressed by (London, 2016), albeit in the context of uniform stability, and we plan to extend this idea to the context of data-dependent stability.

Acknowledgments

This work was in parts funded by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement no 637076). This work was in parts funded by the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement no 308036.

References

- Allen-Zhu, Z. and Hazan, E. Variance reduction for faster non-convex optimization. In *International Conference on Machine Learning (ICML)*, pp. 699–707, 2016.
- Ben-David, S. and Uner, R. Domain adaptation as learning with auxiliary information. NIPS workshop on New Directions in Transfer and Multi-Task, 2013.
- Bousquet, O. and Elisseeff, A. Stability and Generalization. *Journal of Machine Learning Research*, 2:499–526, 2002.
- Chaudhari, P., Choromanska, A., Soatto, S., LeCun, Y., Baldassi, C., Borgs, C., Chayes, J., Sagun, L., and Zecchina, R. Entropy-sgd: Biasing gradient descent into wide valleys. In *International Conference on Learning Representations (ICLR)*, 2017.
- Choromanska, A., Henaff, M., Mathieu, M., Arous, G. B., and LeCun, Y. The loss surfaces of multilayer networks. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2015.
- Devroye, L. and Wagner, T. Distribution-free performance bounds for potential function rules. *IEEE Transactions on Information Theory*, 25(5):601–604, 1979.
- Elisseeff, A., Evgeniou, T., and Pontil, M. Stability of randomized learning algorithms. *Journal of Machine Learning Research*, 6(Jan):55–79, 2005.
- Galanti, T., Wolf, L., and Hazan, T. A theoretical framework for deep transfer learning. *Information and Inference*, pp. iaw008, 2016.
- Ge, R., Huang, F., Jin, C., and Yuan, Y. Escaping from saddle points-online stochastic gradient for tensor decomposition. In *Conference on Learning Theory (COLT)*, pp. 797–842, 2015.
- Ghadimi, S. and Lan, G. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- Gonen, A. and Shalev-Shwartz, S. Fast rates for empirical risk minimization of strict saddle problems. *arXiv preprint arXiv:1701.04271*, 2017.
- Goodfellow, I., Bengio, Y., and Courville, A. *Deep learning*. The MIT Press, 2016.
- Hardt, M., Recht, B., and Singer, Y. Train faster, generalize better: Stability of stochastic gradient descent. In *International Conference on Machine Learning (ICML)*, 2016.
- Hochreiter, S. and Schmidhuber, J. Flat minima. *Neural Computation*, 9(1):1–42, 1997.
- Kawaguchi, K. Deep learning without poor local minima. In *Conference on Neural Information Processing Systems (NIPS)*, pp. 586–594, 2016.
- Keskar, N. S., Mudigere, D., Nocedal, J., Smelyanskiy, M., and Tang, P. T. P. On large-batch training for deep learning: Generalization gap and sharp minima. In *International Conference on Learning Representations (ICLR)*, 2017.
- Koren, T. and Levy, K. Fast rates for exp-concave empirical risk minimization. In *Conference on Neural Information Processing Systems (NIPS)*, pp. 1477–1485, 2015.
- Kuzborskij, I. and Orabona, F. Stability and hypothesis transfer learning. In *International Conference on Machine Learning (ICML)*, pp. 942–950, 2013.
- Kuzborskij, I. and Orabona, F. Fast Rates by Transferring from Auxiliary Hypotheses. *Machine Learning*, pp. 1–25, 2016.
- Lee, J. D., Simchowitz, M., Jordan, M. I., and Recht, B. Gradient descent only converges to minimizers. In *Conference on Learning Theory (COLT)*, pp. 1246–1257, 2016.
- London, B. Generalization bounds for randomized learning with application to stochastic gradient descent. In *NIPS Workshop on Optimizing the Optimizers*, 2016.
- Maurer, A. and Pontil, M. Empirical bernstein bounds and sample variance penalization. *arXiv preprint arXiv:0907.3740*, 2009.
- Pearlmutter, B. A. Fast exact multiplication by the hessian. *Neural Computation*, 6(1):147–160, 1994.
- Pentina, A. and Lampert, C. H. A pac-bayesian bound for lifelong learning. In *International Conference on Machine Learning (ICML)*, 2014.
- Perrot, M. and Habrard, A. A theoretical analysis of metric hypothesis transfer learning. In *International Conference on Machine Learning (ICML)*, pp. 1708–1717, 2015.
- Poggio, T., Voinea, S., and Rosasco, L. Online learning, stability, and stochastic gradient descent. *arXiv preprint arXiv:1105.4701*, 2011.

- Reddi, S. J., Hefny, A., Sra, S., Póczos, B., and Smola, A. Stochastic variance reduction for nonconvex optimization. In *International Conference on Machine Learning (ICML)*, pp. 314–323, 2016.
- Shalev-Shwartz, S. and Ben-David, S. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.
- Shalev-Shwartz, S., Shamir, O., Srebro, N., and Sridharan, K. Learnability, stability and uniform convergence. *Journal of Machine Learning Research*, 11(Oct):2635–2670, 2010.
- Tommasi, T., Orabona, F., and Caputo, B. Learning categories from few examples with multi model knowledge transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(5):928–941, 2014.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations (ICLR)*, 2017.
- Zhao, P. and Zhang, T. Stochastic optimization with importance sampling for regularized loss minimization. In *International Conference on Machine Learning (ICML)*, pp. 1–9, 2015.