



Calhoun: The NPS Institutional Archive
DSpace Repository

NPS Scholarship

Publications

2019-10-31

Diametrical Risk Minimization: Theory and Computations

Norton, Matthew; Royset, Johannes O.

ArXiv

Norton, Matthew, and Johannes O. Royset. "Diametrical Risk Minimization: Theory and Computations." arXiv preprint arXiv:1910.10844 (2019).

<https://hdl.handle.net/10945/65601>

This publication is a work of the U.S. Government as defined in Title 17, United States Code, Section 101. Copyright protection is not available for this work in the United States.

Downloaded from NPS Archive: Calhoun



Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>

Diametrical Risk Minimization: Theory and Computations

Matthew Norton and Johannes O. Royset

Operations Research Department
Naval Postgraduate School
mnorton@nps.edu, jroyset@nps.edu

Abstract. The theoretical and empirical performance of Empirical Risk Minimization (ERM) often suffers when loss functions are poorly behaved with large Lipschitz moduli and spurious sharp minimizers. We propose and analyze a counterpart to ERM called Diametrical Risk Minimization (DRM), which accounts for worst-case empirical risks within neighborhoods in parameter space. DRM has generalization bounds that are independent of Lipschitz moduli for convex as well as nonconvex problems and it can be implemented using a practical algorithm based on stochastic gradient descent. Numerical results illustrate the ability of DRM to find quality solutions with low generalization error in chaotic landscapes from benchmark neural network classification problems with corrupted labels.

Keywords: empirical risk minimization, generalization error, solution stability.

Date: October 31, 2019

1 Introduction

In stochastic optimization, the minimum value of the empirical risk exhibits a downward bias and the corresponding minimizers are therefore often poor in terms of their true risk. Lipschitz continuity is often brought in as a critical component in attempts to assess the quality of such minimizers, with the Lipschitz moduli of loss functions entering in generalization bounds and other results for Empirical Risk Minimization (ERM) problems; see for example [32, 11, 7, 6, 1]. In this work, we propose a counterpart to ERM called Diametrical Risk Minimization (DRM) that possesses a generalization bound which is independent of Lipschitz moduli for convex as well as nonconvex loss functions. Preliminary simulations on benchmark Neural Network (NN) classification problems with MNIST and CIFAR-10 datasets support the hypothesis that when problems have large Lipschitz moduli (sharp empirical risk landscapes), DRM is able to locate quality solutions with low generalization error, while ERM comparatively struggles. Instead of the empirical risk, DRM considers the *diametrical risk* at a point in the parameter space, which is given by the worst-case empirical risk in a neighborhood of the point. This provides DRM

with a broader view of the local empirical risk landscape than ERM and results in improved performance when the landscape is “chaotic” with large Lipschitz moduli and sharp minimizers.

Dealing with empirical risk landscapes that have large Lipschitz moduli and sharp minimizers is a challenge that has seen renewed attention in recent years under the heading of sharp vs flat minimizers in landscapes generated by NN. It is hypothesized that the landscape of NN problems are chaotic [21, 16] and that flat minimizers have better generalization properties compared to sharp ones [13, 31, 12, 8]. The potential effects can be seen in Figure 1, which depicts the empirical risk and true risk for a loss function with a large Lipschitz modulus. The spurious dip of the empirical risk is substantial because the Lipschitz modulus is large. Since ERM seeks out such dips, the resulting minimizer is poor when assessed using the true risk. We would have preferred that ERM found the left-most valley where any of its obtained minimizers would have had a low true risk. Although conceptual, it is believed that landscapes of NN’s may exhibit similar behavior [13]. Thus, it has been a goal of many researchers to either locate the flat minimizers of such problems or to construct objective functions and/or NN architectures which do not have a large number of sharp minimizers [13, 31, 8, 10]

Figure 1b illustrates a different kind of sharpness, which can be induced by introducing label noise as explored in [36, 24]. When a portion of labels are randomly flipped, it has been shown that NN’s are capable of fitting the training data perfectly, achieving a zero-training error solution. However, this solution clearly will not generalize and lies near a sharp minimizer which is associated with the fitting of the incorrectly labeled training data. Nevertheless, it has been shown that, even in the presence of label noise, there still exist good solutions such as the left-most minimizer in Figure 1b. It is immediately clear that DRM may perform well in landscapes with spurious sharp minimizers due its awareness of the neighborhood of a candidate parameter vector and thus be guided to flat minimizers.

Lipschitz moduli (that would be large in these examples) frequently play a direct role in statistical learning theory. For example, a paradigm in learning theory is the analysis of algorithmic stability [32, 11, 7, 6]. However, a majority of these learning bounds require some notion of smoothness in terms of either a Lipschitz continuous objective function and/or Lipschitz continuous gradient. The Lipschitz moduli then enter the resulting generalization bounds and influence the (theoretical) stability of the algorithm used to perform ERM. The reliance of these and other generalization bounds [1] on the magnitude of the Lipschitz moduli, as well as the growing support for the sharp-vs-flat hypothesis, have even given rise to research centered on Lipschitz regularization [22, 10, 26] for improving the generalization of NN’s. We provide generalization bounds, however, that do not rely upon the Lipschitz moduli: DRM can be applied to sharp landscapes with resulting solutions having low generalization error. Eliminating the dependence on the Lipschitz modulus does, however, come with a cost. The provided bounds rely somewhat unfavorably on the dimension of the parameter space. However, experiments indicate that this is a limitation of the proof approach and not a fundamental limitation of DRM in general. We carry out experiments in a NN setting where the dimension of the parameter space is larger than the number of training samples. Even in this setting, we find that DRM




Figure 1

solutions generalize favorably compared to solutions found via ERM.

The downward bias associated with ERM has been known since the early days of stochastic optimization and M -estimators. Traditional remedies include a variety of regularization schemes, focused on alteration of the objective function [6, 17, 3] or the optimization procedure itself with, for example, early stopping [11, 28, 30]. Another remedy is to replace ERM by the problem of minimizing the distributionally worst-case empirical risk; see for example [35, 37, 29, 9, 4] and references therein. Typically, the worst-case is defined in terms of a ball in some metric on a space of probability measures centered on the empirical distribution generated by the available data. DRM is distinct from this approach by perturbing the parameter vector instead of the distribution governing the data.

Perturbation of a parameter vector as a means to account for “implementation error” of a decision or design specified by the vector is included as a motivation for Robust Optimization [2]; see [34] for application in the context of meta-models and [20] for fabrication problems. The latter reference as well as [14, 15] lay theoretical and computational foundations for minimizing functions of the form $w \mapsto \sup_{v \in V} f(w + v)$ that include establishing Lipschitz continuity even if f is rather general. In particular, the minimization of such sup-functions can be achieved by semidefinite programming [14, 19] when f is convex and quadratic. Examples of “robustification” by considering a worst-case parameter vector are also found in signal processing [18, 25].

The remainder of this paper is organized as follows. Section 2 introduces DRM and illustrates the difficulties faced by ERM when loss functions are poorly behaved. In Section 3, we provide a theoretical analysis of DRM that includes generalization bounds independent of Lipschitz moduli. In Section 4, we propose a practical algorithm for performing DRM based on stochastic gradient descent (SGD). We then provide an experimental study in Section 5, with a focus on supporting the idea that DRM can find good solutions to problems with sharp empirical risk landscapes. An appendix furnishes proofs and code for experiments is available online¹.

¹ https://github.com/matthew-norton/Diametrical_Learning

2 Diametrical Risk Minimization

For a loss function $\ell : \mathbb{R}^n \times \mathbb{R}^d \rightarrow \mathbb{R}$ and sample $S = \{z_1, \dots, z_m\} \subset \mathbb{R}^d$, it is well known that the ERM problem

$$\text{minimize}_{w \in \mathbb{R}^n} R_m(w) = \frac{1}{m} \sum_{i=1}^m \ell(w, z_i)$$

can be a poor surrogate for the actual problem of minimizing the population risk $R(w) = \mathbb{E}_z[\ell(w, z)]$. Specifically, a (near-)minimizer w_m^* of R_m tends to have population risk $R(w_m^*)$ significantly higher than that stipulated by $R_m(w_m^*)$. The effect worsens when the loss $\ell(w, z)$ varies dramatically under changing parameters w , which is the case when $\ell(\cdot, z)$ has a large Lipschitz modulus.

In this section, we propose an alternative that we coin *Diametrical Risk Minimization (DRM)*. In contrast to common robustification strategies based on perturbing the data set, DRM perturbs the parameters and thereby obtains stability even for poorly behaved loss functions. Instead of minimizing R_m directly as in ERM, any learned parameter vector w is “diametrically” modified before the empirical risk is minimized. Specifically, for $\gamma \in [0, \infty)$, let

$$R_m^\gamma(w) = \sup_{\|v\| \leq \gamma} R_m(w + v)$$

be the *diametrical risk* of parameter vector $w \in \mathbb{R}^n$ and is therefore the worst possible empirical risk in a neighborhood of w . Any norm $\|\cdot\|$ can be used. Trivially, $R_m^0(w) = R_m(w)$, but generally $R_m^\gamma(w) \geq R_m(w)$. For some set $W \subset \mathbb{R}^n$ of permissible parameter vectors, the *Algorithm for Diametrical Risk Minimization* amounts to

$$\text{minimize}_{w \in W} R_m^\gamma(w)$$

and results in a solution w_m^γ , which might be a global minimizer, a local minimizer, a stationary point, or some other parameter vector with “low” diametrical risk.

As we show in Theorem 3.2, under mild assumptions,

$$R(w_m^\gamma) - R_m^\gamma(w_m^\gamma) \leq \beta m^{-1/2} \tag{1}$$

with high probability for some constant β regardless of the exact nature of w_m^γ . In particular, w_m^γ generalizes even if obtained after aggressive minimization of the diametrical risk; DRM is inherently resistant to overfitting.

Two examples illustrate the challenge faced by ERM when loss functions have large Lipschitz moduli (with respect to parameters w). In both examples, we will see that the generalization

error for DRM is dramatically smaller than for ERM. For $\kappa \in (1, \infty)$ and $\gamma \in (0, 1)$, let

$$\ell(w, z) = \begin{cases} \kappa w/\gamma + \kappa & \text{if } w \in [-\gamma, 0), z = 0 \\ -\kappa w/\gamma - \kappa & \text{if } w \in [-\gamma, 0), z = 1 \\ -\kappa w/\gamma + \kappa & \text{if } w \in [0, \gamma), z = 0 \\ \kappa w/\gamma - \kappa & \text{if } w \in [0, \gamma), z = 1 \\ 0 & \text{otherwise.} \end{cases}$$

If z takes the values 0 and 1, each with probability $\frac{1}{2}$, then $R(w) = \mathbb{E}_z[\ell(w, z)] = 0$ for all $w \in \mathbb{R}$. In contrast,

$$R_m(w) = \begin{cases} \frac{1}{m}\rho_m(\kappa w/\gamma + \kappa) & \text{if } w \in [-\gamma, 0) \\ \frac{1}{m}\rho_m(-\kappa w/\gamma + \kappa) & \text{if } w \in [0, \gamma) \\ 0 & \text{otherwise.} \end{cases}$$

where ρ_m is the number of zeros minus the number of ones in the data $\{z_1, \dots, z_m\}$. Viewing the data as random, we obtain that with probability nearly $\frac{1}{2}$, $\rho_m < 0$ and thus $w_m^* = 0$ minimizes R_m for such outcomes of the data and $R_m(w_m^*) = \rho_m \kappa/m$. Also with probability near $\frac{1}{2}$, $\rho_m \geq 0$ and then $w_m^* = 1$ minimizes R_m and $R_m(w_m^*) = R(w_m^*) = 0$. Consequently, $R_m(w_m^*)$ has a downward bias. Although

$$R(w_m^*) - R_m(w_m^*) \leq \max\{0, -\rho_m\}\kappa/m \tag{2}$$

with probability one, the right-hand side includes the constant κ , which is proportional to the Lipschitz modulus κ/γ of $\ell(\cdot, z)$. This illustrates the well-known fact that generalization tends to be poor for loss functions with large Lipschitz moduli. However, considering diametrical risk, we have that $R(w_m^\gamma) - R_m^\gamma(w_m^\gamma) \leq 0$ with probability one.

The situation deteriorates further when the loss function is not Lipschitz continuous. Let

$$\ell(w, z) = \begin{cases} 1/w & \text{if } w \in (0, \infty), z = 0 \\ -1/w & \text{if } w \in (0, \infty), z = 1 \\ 0 & \text{otherwise.} \end{cases}$$

Again, with z and ρ_m as above, $R(w) = \mathbb{E}_z[\ell(w, z)] = 0$ for all $w \in \mathbb{R}$ and

$$R_m(w) = \begin{cases} \frac{1}{m}\rho_m/w & \text{if } w \in (0, \infty) \\ 0 & \text{otherwise.} \end{cases}$$

Then, $\inf_{w \in \mathbb{R}} R_m(w) = -\infty$ when $\rho_m < 0$, which takes place with probability nearly $\frac{1}{2}$. The downward bias is now unbounded. However, we have that $R(w_m^\gamma) - R_m^\gamma(w_m^\gamma) \leq \beta\gamma^{-1}m^{-1/2}$ with high probability for some constant β .

From these simple examples, it is clear that ERM can lead to arbitrarily slow learning when the loss function is poorly behaved. In the first example above, DRM has a generalization error (right-hand side of (2)) equal to zero and thus is independent of the Lipschitz modulus. In the second example, DRM reduces the unbounded generalization error encountered by ERM to a quantity proportional to $m^{-1/2}$ as we see in the following section.

3 Theoretical Analysis of DRM

We begin by formalizing the setting and recall that $f : \mathbb{R}^n \times \Omega \rightarrow \mathbb{R}$ is a Caratheodory function relative to a probability space (Ω, \mathcal{A}, P) if for all $\omega \in \Omega$, $f(\cdot, \omega)$ is continuous and for all $w \in \mathbb{R}^n$, $f(w, \cdot)$ is \mathcal{A} -measurable. In the following, we assume that the data comprises d -dimensional random vectors generated by independent sampling according to the distribution \mathbb{P} and thus consider the m -fold product probability space $(Z^m, \mathcal{Z}^m, \mathbb{P}^m)$ constructed from a probability space $(Z, \mathcal{Z}, \mathbb{P})$, with $Z \subset \mathbb{R}^d$. If $\ell : \mathbb{R}^n \times Z \rightarrow \mathbb{R}$ is a Caratheodory function relative to $(Z, \mathcal{Z}, \mathbb{P})$, then R_m , now viewed as a function on $\mathbb{R}^n \times Z^m$, is a Caratheodory function relative to $(Z^m, \mathcal{Z}^m, \mathbb{P}^m)$; see for example [27, Prop. 14.44; Ex. 14.29]. Likewise, we view R_m^γ as a function on $\mathbb{R}^n \times Z^m$. It is real-valued by virtue of being the maximum value of the continuous R_m over a compact set. For all $w \in \mathbb{R}^n$, $R_m^\gamma(w)$ is \mathcal{Z}^m -measurable when R_m is a Caratheodory function [27, Thm. 14.37]. Since R_m^γ is continuous (in w) for all $(z_1, \dots, z_m) \in \mathcal{Z}^m$ by [27, Thm. 1.17], we conclude that R_m^γ is a Caratheodory function relative to $(Z^m, \mathcal{Z}^m, \mathbb{P}^m)$. In view of [27, Thm. 14.37; Ex. 14.32], $\sup_{w \in W} R_m^\gamma(w)$ is \mathcal{Z}^m -measurable as long as $W \subset \mathbb{R}^n$ is closed. In effect, any concern about measurability in the below statements are put to rest if ℓ is a Caratheodory function and W is closed.

We denote by \mathbb{E} the expectation with respect to \mathbb{P} so that for $w \in \mathbb{R}^n$, $R(w) = \mathbb{E}[\ell(w, z)] = \int \ell(w, z) d\mathbb{P}(z)$. When $\ell : \mathbb{R}^n \times Z$ is a Caratheodory function, we say it is locally sup-integrable if for all $\bar{w} \in \mathbb{R}^n$, there exists $\rho > 0$ such that $\int \sup\{\ell(w, z) \mid \|w - \bar{w}\| \leq \rho\} d\mathbb{P}(z) < \infty$. It is locally inf-integrable if “sup” is replaced by “inf” in the above statement. The moment-generating function of a random variable X is $\tau \mapsto E[\exp(\tau X)]$. Proofs for the following results can be found in the Appendix.

3.1 Theorem (generalization error in DRM). *Suppose that $W \subset \mathbb{R}^n$ is compact, $\ell : \mathbb{R}^n \times Z \rightarrow \mathbb{R}$ is a locally sup-integrable Caratheodory function, and for all $w \in W$, the moment generating function of $\ell(w, \cdot) - R(w)$ is real-valued in a neighborhood of zero. Then, for any $\varepsilon, \gamma > 0$ and m , there exist $\eta, \beta > 0$ (independent of m) such that*

$$\mathbb{P}^m \left(\sup_{w \in W} \{R(w) - R_m^\gamma(w)\} \leq \varepsilon \right) \geq 1 - \eta e^{-\beta m}.$$

The theorem furnishes a uniform bound on R , which implies in particular that

$$R(w_m^\gamma) \leq R_m^\gamma(w_m^\gamma) + \varepsilon \text{ with high probability}$$

for any parameter vector w_m^γ produced by DRM. Thus, there is a strong justification for minimizing R_m^γ : lower values of the diametrical risk ensure better guarantees on the true risk. The goal now becomes to develop good methods for producing w_m^γ with low $R_m^\gamma(w_m^\gamma)$. The issue of overfitting is largely removed.

The assumptions in the theorem are generally mild: $\ell(\cdot, z)$ only needs to be continuous and the condition on the moment generating function is just checked pointwise. The requirement about locally sup-integrable amounts to determine an integrable random variable at every \bar{w} that upper bounds ℓ in a neighborhood of \bar{w} . The constant β depends on ε , while η is a function of γ and the diameter of W . In particular, they are independent of Lipschitz moduli of ℓ , which may not even be finite.

If the value of the moment generating function of $\ell(w, \cdot) - R(w)$ can be quantified near zero, then we can examine the effect as the error ε vanishes as seen next. We recall that a random variable X is subgaussian with variance proxy σ^2 if its mean $E[X] = 0$ and its moment generating function satisfies $E[\exp(\tau X)] \leq \exp(\frac{1}{2}\sigma^2\tau^2)$ for all $\tau \in \mathbb{R}$.

3.2 Theorem (rate of convergence in DRM). *Suppose that $W \subset \mathbb{R}^n$ is compact, $\ell : \mathbb{R}^n \times Z \rightarrow \mathbb{R}$ is a locally sup-integrable Caratheodory function, and for all $w \in W$, $\ell(w, \cdot) - R(w)$ is subgaussian. Then, for any $\alpha \in (0, 1)$, $\gamma > 0$, and m , there exists $\beta > 0$ (independent of m) such that*

$$\mathbb{P}^m \left(\sup_{w \in W} \{R(w) - R_m^\gamma(w)\} \leq \beta m^{-1/2} \right) \geq 1 - \alpha.$$

The constant β in the theorem is given in the proof and depends on the largest variance proxy, denoted by σ^2 , for $\ell(w, \cdot) - R(w)$ at a finite number of different w . It also depends on a parameter ξ that measures the size of X and this depends on the norm used. For example, if R_m^γ is defined in terms of the sup-norm, then the balls W_1, \dots, W_N in the proof can be constructed according to that norm and the number required is simply² $N = \lceil \delta/\gamma \rceil^n$, where $\delta = \sup_{w, \bar{w} \in W} \|w - \bar{w}\|_\infty$. This leads to

$$\beta = \sigma \sqrt{2n \log \lceil \delta/\gamma \rceil - 2 \log \alpha}.$$

The constant β in the theorem depends unfavorably on n . One can attempt to reduce the effect of n by enlarging γ as n increases. For example, under the sup-norm one may set $\gamma = \delta [\exp(\zeta n^{-\nu})]^{-1}$ for positive constants ζ and ν . Then, assuming that δ/γ is an integer (which can be achieved by enlarging δ when needed),

$$\beta = \sigma \sqrt{2n \log \lceil \delta/\gamma \rceil - 2 \log \alpha} = \sigma \sqrt{2\zeta n^{1-\nu} - 2 \log \alpha}.$$

For example, if $\nu = 1$, then β becomes independent of n at the expense of having to grow γ rather quickly as n increase. A compromise could be $\nu = \frac{1}{2}$, in which case β grows only as $n^{1/4}$ and γ grows somewhat slower too. Still, in the limit as $n \rightarrow \infty$, R_m^γ involves maximization over all of W , which of course leads to an upper bound.

²We note by $\lceil c \rceil$ be lowest integer at least as high as the scalar c .

For fixed n , we may also ask what is the right value of γ . Since a large value might imply additional computational burden and also lead to overly conservative upper bounds, it would be ideal to select it as small as possible without ruining the rate (significantly). One possibility could be to set γ proportional to m^{-1} because then the rate deteriorate only with a logarithmic factor from $m^{-1/2}$ to $\sqrt{m^{-1} \log m}$.

It is clear from the proof of the theorem that the assumptions about independent sampling and subgaussian random variables can be relaxed. We only needed that the error in $R_m(w)$ relative to $R(w)$ can be bounded for a finite number of w ; see [5, 23] for possible extensions.

It may be of interest to determine the error of an obtained parameter vector w_m^γ relative to the set of actually optimal parameters $\operatorname{argmin}_{w \in W} R(w)$. Theorem 3.2 yields immediately that for any $\delta \in \mathbb{R}$,

$$\{w \in W \mid R_m^\gamma(w) \leq \delta\} \subset \{w \in W \mid R(w) \leq \delta + \beta m^{-1/2}\} \text{ with probability at least } 1 - \alpha.$$

We now examine the harder question of confidence regions for “good” parameter vectors. For two sets $A, B \subset \mathbb{R}^n$, we denote the excess of A over B by

$$\operatorname{exs}(A; B) = \sup_{w \in A} \inf_{\bar{w} \in B} \|w - \bar{w}\| \text{ for nonempty } A, B;$$

with the convention that $\operatorname{exs}(A; B) = \infty$ if $A \neq \emptyset$ and $B = \emptyset$; $\operatorname{exs}(A; B) = 0$ otherwise. Below, the sets of interest are lower level-sets of R_m , possible with a random level. Arguing by means of Prop. 14.33, Thm. 14.37, and Ex. 14.32 in [27], we see that the excess involving such sets is measurable.

3.3 Theorem (confidence region in DRM). *Suppose that $W \subset \mathbb{R}^n$ is compact, $\ell : \mathbb{R}^n \times Z \rightarrow \mathbb{R}$ is a locally inf-integrable Caratheodory function, and for all $w \in W$, $\ell(w, \cdot) - R(w)$ is subgaussian. Then, for any $\alpha \in (0, 1)$, $\gamma > 0$, $\delta \in \mathbb{R}$, and m , there exists $\beta > 0$ (independent of m) such that*

$$\begin{aligned} \mathbb{P}^m \left(\operatorname{exs} \left(\{w \in W \mid R(w) \leq \delta\}; \{w \in W \mid R_m(w) \leq \delta + \beta m^{-1/2}\} \right) \leq \gamma, \right. \\ \left. \operatorname{exs} \left(\operatorname{argmin}_{w \in W} R(w); \{w \in W \mid R_m(w) \leq \inf_{\bar{w} \in W} R_m^\gamma(\bar{w}) + 2\beta m^{-1/2}\} \right) \leq \gamma \right) \geq 1 - \alpha. \end{aligned}$$

Since the constant β is nearly of the same form as in Theorem 3.2, the discussion following that theorem carries over to the present context. In particular, we note that a lower-level set of R_m , enlarged with γ , covers $\operatorname{argmin}_{w \in W} R(w)$ with high probability.

4 Algorithms for Diametrical Risk Minimization

Although there are some computational challenges associated with DRM, most of the existing optimization procedures for ERM can be adapted. Significantly, if the empirical risk function

R_m is convex, then the diametrical risk function R_m^γ is also convex. Moreover, regardless of convexity,

$$\frac{1}{m} \sum_{i=1}^m \nabla_w \ell(\bar{w} + \bar{v}, z_i)$$

is a subgradient (in the general sense, cf. [27, Ch. 8]) of R_m^γ at \bar{w} under weak assumptions, where $\bar{v} \in \operatorname{argmax}_{\|v\| \leq \gamma} R_m^\gamma(\bar{w} + v)$; see for example [27, Cor. 10.9]. This implies that standard (sub)gradient methods apply provided that \bar{v} can be computed. Since γ might very well be small, this could be within reach, at least approximately, by carrying one iteration of gradient *ascent*. However, this could become costly as computation of such subgradients need to access all data points.

We utilize a less costly approach based on the application of SGD to an outer approximation formed via sampling. In short, we approximate the inner maximization by maximizing over a finite set of random points inside the γ -neighborhood at the current solution w^t of each iteration. We find this approach to be effective, even when working with problems involving NN's where the dimension of w is in the millions. In these applications an outer approximation of R_m^γ using as little as 10-20 samples from $\{v \mid \|v\| = \gamma\}$ suffices to achieve improvement over ERM.

We observe also that DRM is related to but distinct from ERM with early termination. If from a minimizer w_m^* of R_m^γ the process of maximizing $R_m(w_m^* + v)$ subject $\|v\| \leq \gamma$ follows the trajectory along which the algorithm approached w_m^* in the first place, then DRM would be equivalent to ERM that terminates a distance γ from a minimizer. However, this equivalence will only take place when w_m^* is approached along such direction. It appears that this will occur only occasionally.

4.1 Gradient Based Algorithm

We propose two variations of an SGD-based algorithm for DRM which we denote by Simple-SGD-DRM and SGD-DRM. We start with a simple version of the main algorithm that is easier to follow and then introduce the full algorithm with minor alterations aimed toward improving efficiency. In the following, let $\operatorname{prj}_W(w)$ denote the projection of w on W and let $R_{B_t}(w) = \frac{1}{|B_t|} \sum_{z \in B_t} \ell(w, z)$ denote the empirical risk over a batch $B_t \subset S$.

Algorithm 1: Simple-SGD-DRM.

Step 0. Initialize $w^0 \in W$, $r \in \mathbb{N}$, $t = 0$. Initialize sequence of batches $B_t \subset S$ and learning rates $\lambda_t > 0$ for $t = 1, \dots, T$.

Step 1. Sample r random perturbations (directions): $U = \{u_1, \dots, u_r \mid \|u\| = \gamma\}$

Step 2. Select $u^* \in \operatorname{argmax}_{u \in U} \frac{1}{|B_t|} \sum_{z \in B_t} \ell(w^t + u, z)$

Step 3. Compute $w^{t+1} = \operatorname{prj}_W(w^t - \lambda_t \nabla_w R_{B_t}(w^t + u^*))$.

Step 4. If $t = T$, stop. Else, $t \leftarrow t + 1$ and return to Step 2.

The Simple-SGD-DRM algorithm, at each iteration t , performs an SGD update towards minimizing the approximating objective function

$$w \mapsto \max_{u \in U} \frac{1}{|B_t|} \sum_{z \in B_t} \ell(w + u, z_i).$$

The algorithm does so by first forming a set of r random directions (vectors) $U = \{u_1, \dots, u_r\}$ with norm equal to γ . Then, it determines the more critical $u \in U$, i.e., $u^* \in \operatorname{argmax}_{u \in U} \frac{1}{|B_t|} \sum_{z \in B_t} \ell(w^t + u, z)$. A subgradient of the approximating objective function is then $\nabla_w R_{B_t}(w^t + u^*)$.

This algorithm, however, does have drawbacks. First, sampling r new vectors in Step 1 at every iteration can be computationally expensive. It may be enough, as we will see in the experiments, to only perform Step 1 intermittently.³ Second, it may be beneficial to reuse one or more of the sampled vectors from Step 1 in future iterations, particularly if we decide to perform sampling only intermittently.

The following algorithm, which we simply call SGD-DRM, includes these options explicitly. One will notice that it can be made equivalent to simple-SGD-DRM with particular choices, and is thus an extension with more options to save computation by limiting sampling.

Algorithm 2: SGD-DRM.

- Step 0.** Initialize $V_{-1} = \{\}$, $w^0 \in W$, $q \in \mathbb{N}$, $r \in \mathbb{N}$, $p \in [0, 1]$, $t = 0$. Initialize sequence of batches $B_t \subset S$ and learning rates $\lambda_t > 0$ for $t = 1, \dots, T$.
- Step 1.** Sample r random perturbations (directions): $U = \{u_1, \dots, u_r \mid \|u\| = \gamma\}$
- Step 2.** Select $u^* \in \operatorname{argmax}_{u \in U} \frac{1}{|B_t|} \sum_{z \in B_t} \ell(w^t + u, z)$
- Step 3.** Let $V_t = V_{t-1} \cup \{u^*\}$. If $|V_t| > q$, remove oldest element from V_t
- Step 4.** Select $v^* \in \operatorname{argmax}_{v \in V_t} \frac{1}{|B_t|} \sum_{z \in B_t} \ell(w^t + v, z)$
- Step 5.** Compute $w^{t+1} = \operatorname{prj}_W(w^t - \lambda_t \nabla_w R_{B_t}(w^t + v^*))$.
- Step 6.** If $t = T$, stop. Else, $t \leftarrow t + 1$ and with probability p , return to Step 1; with probability $1 - p$ return to Step 4 with $V_t = V_{t-1}$.

The primary difference between this and the former algorithm is within Steps 3-4 and Step 6. Step 6 allows one to skip the expensive sampling in Step 1 at some frequency represented by p . The new set V_t is introduced in Steps 3-4 to allow the reuse of one or more vectors u from previous iterations. As the algorithm progresses, the set V_t acts like a queue with maximal size q .⁴ Every time the sampling of Step 1 is not skipped, V_t will be equal to the set V_{t-1} with its oldest element replaced by u^* . For iterations $t \leq q$, the oldest element need not be removed since the queue has not reached its maximum length of q .

³In experiments, we sample only every 5th iteration.

⁴In experiments, we set $q = 1$.

4.2 Implementation for Neural Networks

In the following experiments, we consider NN classifiers. Because of the structure of NN’s we implement the perturbations $w + v$ and $w + u$ by considering the groupings of parameters that correspond to the structure of each NN layer. For example, a two layer NN might have parameter matrices $\{W_1, W_2\}$, each belonging to the separate network layers. Because of this, we select perturbations in Step 1 such that the Frobenius norm of each layer-wise perturbation matrix is equal to γ . So, for the network with separable parameters $\{W_1, W_2\}$, a single sample from Step 1 would look like $\{U_1, U_2\}$ with $\|U_1\|_F = \gamma, \|U_2\|_F = \gamma$. Additionally, in the implementation, we first sample each component from a standard normal distribution, then normalize the resulting vector (or matrix) to have norm equal to γ .

Additionally, in the experiments we implement the coin flip (based on p) from Step 6 deterministically. We only perform Step 2 and 3 at every 5th step. Otherwise, we let $V_t = V_{t-1}$. This allows us to save computation time, particularly since the sampling of U can be expensive for NN’s with millions of parameters.

5 Experiments

We aim to illustrate that DRM is resistant to overfitting and that its solutions have different local characteristics compared to those from ERM. In particular, we hypothesize that minimizers found by DRM lie in flatter regions of the empirical risk landscape. To illustrate these aspects, we focus on the problem of classification with NN’s in the presence of label noise, which is suitable because it is prone to overfitting as ERM settles into sharp minimizers; see [24, 36]. These works show that NN’s have the unique ability to perfectly fit⁵ both a data set with label noise as well as the same data with correct labels, with large generalization error in the former but small generalization error in the latter.

In these experiments, we train NN’s on a subset of the MNIST and CIFAR-10 datasets with large amounts of label noise (50% of training labels flipped to an incorrect class). Using standard SGD for ERM (labelled SGD-ERM), the NN’s indeed settle into solutions with high generalization error. On the other hand, SGD-DRM is remarkably resistant to overfitting, finding solutions with dramatically lower generalization error. Additionally, we find empirical evidence that the SGD-DRM solution lies in a “wide” valley compared to the SGD-ERM solution which exists around a much sharper landscape.

Code associated with these experiments and a PyTorch⁶ based implementation of SGD-DRM is available online⁷.

⁵with zero training error

⁶ <https://pytorch.org/>

⁷ https://github.com/matthew-norton/Diametrical_Learning

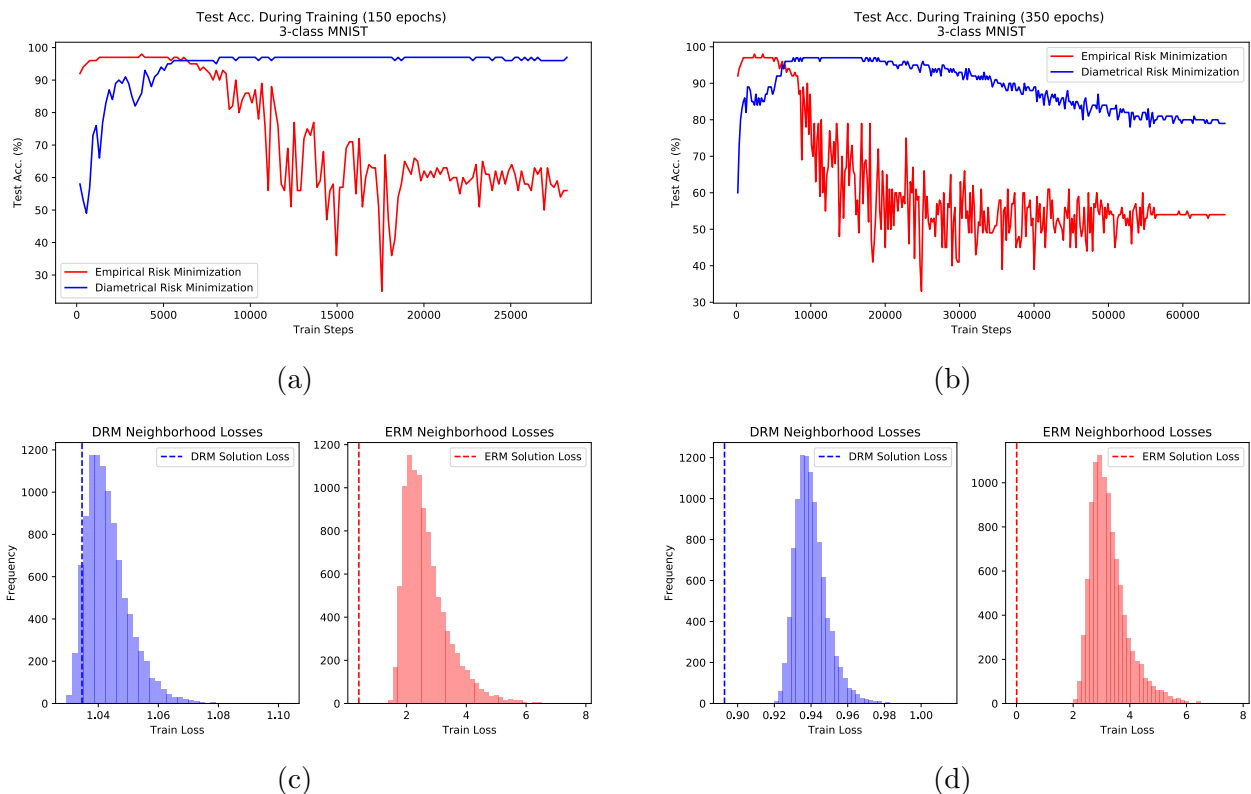


Figure 2: MNIST experiments: (a) Test accuracy when trained for 150 epochs. (b) Test accuracy when trained for 350 epochs. (c) Distribution of empirical risk for 10,000 random points in γ -neighborhood of final 150 epoch DRM and ERM solution. (d) Distribution of empirical risk for 10,000 random points in γ -neighborhood of final 350 epoch DRM and ERM solution.

5.1 MNIST

In the first set of experiments, we train a fully connected, 3 layer NN with hidden units per layer being (320, 320, 200) with ReLU nonlinearities and an additional fully connected output layer feeding into a 3-class softmax negative log-likelihood objective function. For the MNIST dataset, we use only the handwritten digits zero, one, and two so that SGD-ERM achieves nearly zero training error. We flip 50% of the training labels to an incorrect class. For both SGD-ERM and SGD-DRM, we use standard SGD updates with batch size 100 and learning rate .01 until the last 50 epochs when decreased to .001. For SGD-DRM, the Step 1 perturbations are treated on a layer-wise basis (see Section 4.2) with $\gamma = 10$. We also use $r = 20$ and $q = 1$ for the size of U and V , respectively. We implement Step 6 deterministically with Step 1 being performed every 5th iteration (batch).

We train twice. First, we train the network for 150 epochs total with learning rate .01

for the first 100 and .001 for the final 50. Test accuracy can be seen in Figure 2a. We first note the behavior of SGD-ERM. It begins by finding a good solution that generalizes, but then continues minimizing the empirical risk and settles into a solution that fits the training data with incorrect labels and thus suffers from a sharp decline in test accuracy (a sharp increase in generalization error). In contrast, SGD-DRM resists overfitting. Once it finds a good solution that generalizes, it is able to stay there, resisting the fall into a poor solution. As mentioned earlier, it is hypothesized that solutions in wide valleys generalize better than those in sharp regions. Figure 2c plots the distribution of empirical risk for 10,000 random points in the γ -neighborhood of the final 150-epoch DRM and ERM solutions and illustrates that, indeed, the DRM solution is in a much flatter region than the ERM solution. The dotted line represents the value of the empirical risk at the found solution and the rest of the plot is the distribution of empirical risk at points surrounding the solution⁸. Figure 2c also illustrates that, while the empirical risk of the SGD-ERM solution is lower (the red dotted line) than that of SGD-DRM it has much larger diametrical risk.


To allow SGD-ERM enough time to achieve nearly zero training error, we also train for 350 epochs with learning rate .01 for the first 300 and .001 for the final 50. Figure 2b depicts test accuracy. Again, we see the same behavior for SGD-ERM, as it chaotically falls into a poor solution that does not generalize. SGD-DRM remains resistant to overfitting. While it does experience some degradation in test accuracy, it still ends at a much better solution and its path is not nearly as chaotic; see the smooth vs choppy lines in Figure 2b. We also see, again, that the SGD-DRM solution lies in a flatter region than that from SGD-ERM. Figure 2d illustrates that, while the empirical risk of the SGD-ERM solution is lower (the dotted line at zero) it has (approximately) much larger diametrical risk equal to around 6. The SGD-DRM solution, on the other hand, has higher empirical risk ($\approx .89$) but much smaller diametrical risk ($\approx .98$).

5.2 CIFAR-10


We perform similar experiments on CIFAR-10 using subclasses airplane, bird, and car and two network architectures: a fully connected and ResNet20. First, we use the same fully connected architecture as before with hidden 3 layers of size (320, 320, 200) with ReLU nonlinearities and an additional fully connected output layer feeding into a 3-class softmax negative log-likelihood objective function. We train the network on a dataset with 50% of training labels flipped to an incorrect class. For both SGD-ERM and SGD-DRM, we use standard SGD updates with batch size 100 and learning rate .01 for 750 epochs and .001 for the final 50 epochs. For SGD-DRM, the Step 1 perturbations are treated on a layer-wise basis (see Section 4.2) with $\gamma = 5$. We adopt $r = 20$, $q = 1$, and a deterministic implementation of Step 6 as above.

Results are nearly identical to those for MNIST; see Figure 3a. SGD-ERM begins by finding


⁸The points are sampled in the same way as for U in Step 1 of SGD-DRM, with neighborhood points sampled on a layer-wise basis as $\{w^* + u \mid \|u\| = \gamma\}$. Additionally, we use the same set of points u for approximating the neighborhood of ERM and DRM solutions.




(a) Fully Connected



(b) ResNet20



(c) Fully Connected



(d) ResNet20

Figure 3: CIFAR-10 experiments: (a-b) Test accuracy when trained for 800 epochs. (c-d) Distribution of empirical risk for 10,000 random points in γ -neighborhood of final 800 epoch DRM and ERM solution.

a good solution that generalizes, but then continues minimizing the empirical risk and settles into a solution that fits the training data with incorrect labels and thus suffers from a sharp decline in test accuracy. SGD-DRM is again resistant to overfitting and is able to find a good solution that generalizes. Figure 3c illustrates that the DRM solution is in a much flatter region than the ERM solution and achieves much smaller diametrical risk at the expense of larger empirical risk. The figure plots the distribution of empirical risk for 10,000 random points in the γ -neighborhood of the final 800-epoch DRM and ERM solutions. The dotted line represents the value of the empirical risk at the found solution and the rest of the plot is the distribution of empirical risk at points surrounding the solution.

For the ResNet20 architecture, we utilize identical settings except that $\gamma = 1$ due to the smaller number of parameters (per layer) and we omit perturbations to the batchnorm layer parameters. Figure 3b reports test accuracy. Unlike for the fully connect architecture, SGD-DRM suffers nearly zero degradation of test accuracy as training progresses. It also achieves, and maintains, higher test accuracy than is achieved by SGD-ERM. Furthermore, Figure 3d illustrates again the flatness of the landscape surrounding the SGD-SRM solution. While the

diametrical risk is similar for both solutions, there is still a significant gap between the empirical risk of the SGD-ERM solution and points within its neighborhood. For SGD-DRM, this gap is much smaller, indicating a flatter landscape.

5.3 Discussion

While these experiments support the proposition that minimization of diametrical risk leads to good generalization and can be used to handle problems with large Lipschitz moduli, there is still much to be explored with DRM. In particular, SGD-DRM can be improved by considering different policies for choosing hyperparameters γ , λ , q , d , batch size, and p . The choice of γ , for example, could be chosen adaptively at every iteration along with the step size λ , mimicking adaptive SGD implementations such as Adam. Also, different choices of γ could be made for different groupings of parameters corresponding to the layers in a NN.

An additional algorithmic component that could be made more efficient is the estimation of the diametrical risk, which currently is a bottleneck. First, the sampling in Step 1 can be expensive, especially if done at every iteration or for large value of d . Second, if d and/or q is large, the memory requirements of storing U and V can be large. Third, the maximization in Steps 2 and 4 can also be expensive, particularly if the batch size is large and the loss is expensive to compute. Many of these issues, however, can be reduced with parallel implementation. For example, independent workers can each produce a single sample u_i and calculate the value of the objective function. Then, Step 2 can be performed by considering only the collection of d function values produced by the set of workers. We leave these tasks to future work, however, and use the presented results to encourage more work in this direction.

A Proofs

Proof of Theorem 3.1. Suppose that $w^\nu \rightarrow \bar{w}$. Since ℓ is a sup-integrable Caratheodory function, Fatou's Lemma establishes that $\liminf \mathbb{E}[-\ell(w^\nu, z)] \geq \mathbb{E}[\liminf(-\ell(w^\nu, z))] = \mathbb{E}[-\ell(\bar{w}, z)]$. Thus, R is upper-semicontinuous. Let $\{W_k \subset \mathbb{R}^n, k = 1, \dots, N\}$ be a finite cover of W consisting of closed balls with radius $\gamma/2$. Without loss of generality, suppose that $W_k \cap W \neq \emptyset$ for all $k = 1, \dots, N$. Let $w^k \in \operatorname{argmax}_{w \in W_k \cap W} R(w)$, which exists for all $k = 1, \dots, N$ because $W_k \cap W$ is nonempty and compact, and R is upper-semicontinuous.

For $k = 1, \dots, N$, let $\tau \mapsto M_k(\tau)$ be the moment generating function of $R(w^k) - \ell(w^k, \cdot)$ and $I_k(\varepsilon) = \sup_{\tau \in \mathbb{R}} \{\varepsilon\tau - \log M_k(\tau)\}$, which is positive since M_k is real-valued in a neighborhood of zero. Then, by the upper bound in Cramer's Large Deviation Theorem (see for example [33, Sec. 7.2.8])

$$\mathbb{P}^m(R(w^k) - R_m(w^k) \geq \varepsilon) \leq e^{-mI_k(\varepsilon)}.$$

Moreover, with $\beta = \min_{k=1, \dots, N} I_k(\varepsilon)$,

$$\mathbb{P}^m \left(\max_{k=1, \dots, N} \{R(w^k) - R_m(w^k)\} \geq \varepsilon \right) \leq \sum_{k=1}^N e^{-mI_k(\varepsilon)} \leq N e^{-\beta m}.$$

Consider an event for which $\max_{k=1, \dots, N} \{R(w^k) - R_m(w^k)\} \leq \varepsilon$, which takes place with probability at least $1 - N e^{-\beta m}$, and let $\bar{w} \in W$. There exists $\bar{k} \in \{1, \dots, N\}$ such that $\bar{w} \in W_{\bar{k}}$. Since $R_m^\gamma(\bar{w}) = \sup_{\|v\| \leq \gamma} R_m(\bar{w} + v) \geq R_m(w^{\bar{k}})$,

$$\begin{aligned} R(\bar{w}) - R_m^\gamma(\bar{w}) &\leq R(\bar{w}) - R_m(w^{\bar{k}}) \\ &\leq R(\bar{w}) - R(w^{\bar{k}}) + \varepsilon \\ &\leq R(w^{\bar{k}}) - R(w^{\bar{k}}) + \varepsilon = \varepsilon, \end{aligned}$$

where the last inequality follows by the fact that $R(w^{\bar{k}}) = \sup_{w \in W_{\bar{k}}} R(w)$. The conclusion then follows with $\eta = N$ because \bar{w} is arbitrary. \square

Proof of Theorem 3.2. The first paragraph in the proof of Theorem 3.1 carries over. There exists $\xi \in (0, \infty)$, which may depend on n , such that the number of closed balls N of radius $\gamma/2$ required to cover W is no greater than $(\xi/\gamma)^n$. Since $\ell(w^k, \cdot) - R(w^k)$ is subgaussian, say with variance proxy σ_k^2 , we have by Bernstein's inequality that

$$\mathbb{P}^m (R(w^k) - R_m(w^k) > \varepsilon) \leq \exp \left(-\frac{1}{2} m \varepsilon^2 / \sigma_k^2 \right) \quad \text{for all } \varepsilon \in [0, \infty).$$

Let $\sigma = \max_{k=1, \dots, N} \sigma_k$. Thus,

$$\mathbb{P}^m \left(\max_{k=1, \dots, N} \{R(w^k) - R_m(w^k)\} > \varepsilon \right) \leq N \exp \left(-\frac{1}{2} m \varepsilon^2 / \sigma^2 \right) \leq \alpha$$

provided that $\varepsilon \geq \beta m^{-1/2}$ and

$$\beta = \sigma \sqrt{2n \log(\xi/\gamma) - 2 \log \alpha}.$$

Consider the event for which $\max_{k=1, \dots, N} \{R(w^k) - R_m(w^k)\} \leq \varepsilon$, which takes place with probability at least $1 - \alpha$ for such ε . The arguments in the proof of Theorem 3.1 establishes that $\sup_{w \in W} \{R(w) - R_m^\gamma(w)\} \leq \varepsilon$ for this event and the conclusion follows. \square

Proof of Theorem 3.3. Suppose that $w^\nu \rightarrow \bar{w}$. Since ℓ is an inf-integrable Caratheodory function, Fatou's Lemma establishes that $\liminf \mathbb{E}[\ell(w^\nu, z)] \geq \mathbb{E}[\liminf(\ell(w^\nu, z))] = \mathbb{E}[\ell(\bar{w}, z)]$. Thus, R is lower-semicontinuous.

From the compactness of $W \subset \mathbb{R}^n$, we obtain $\xi \in (0, \infty)$, which may depend on n , such that the number of closed balls N with radius $\gamma/2$ required to cover W is no greater than $(\xi/\gamma)^n$. Suppose that $\{W_k \subset \mathbb{R}^n, k = 1, \dots, N\}$ is a collection of such balls with $W_k \cap W \neq \emptyset$ for all

$k = 1, \dots, N$. Let $w^k \in \operatorname{argmin}_{w \in W_k \cap W} R(w)$, which exists for all $k = 1, \dots, N$ because $W_k \cap W$ is nonempty and compact, and R is lower-semicontinuous.

Since $\ell(w^k, \cdot) - R(w^k)$ is subgaussian, say with variance proxy σ_k^2 , we have by Bernstein's inequality that

$$\mathbb{P}^m \left(|R_m(w^k) - R(w^k)| > \varepsilon \right) \leq 2 \exp(-\frac{1}{2} m \varepsilon^2 / \sigma_k^2) \quad \text{for all } \varepsilon \in [0, \infty).$$

Let $\sigma = \max_{k=1, \dots, N} \sigma_k$. Thus,

$$\mathbb{P}^m \left(\max_{k=1, \dots, N} |R_m(w^k) - R(w^k)| > \varepsilon \right) \leq 2N \exp(-\frac{1}{2} m \varepsilon^2 / \sigma^2) \leq \alpha$$

provided that $\varepsilon \geq \beta m^{-1/2}$ and

$$\beta = \sigma \sqrt{2n \log(\xi/\gamma) - 2 \log(\alpha/2)}.$$

Consider the event for which $\max_{k=1, \dots, N} |R_m(w^k) - R(w^k)| \leq \varepsilon$. Let $\bar{w} \in W$ satisfy $R(\bar{w}) \leq \delta$. Then, there exists \bar{k} such that $\bar{w} \in W_{\bar{k}}$ and because $R(w_{\bar{k}}) = \inf_{w \in W_{\bar{k}} \cap W} R(w)$ we obtain that

$$R_m(w_{\bar{k}}) \leq R(w_{\bar{k}}) + \varepsilon \leq R(\bar{w}) + \varepsilon \leq \delta + \varepsilon.$$

Since $\|\bar{w} - w_{\bar{k}}\| \leq \gamma$, we conclude that

$$\operatorname{exs} \left(\{w \in W \mid R(w) \leq \delta\}; \{w \in W \mid R_m(w) \leq \delta + \varepsilon\} \right) \leq \gamma.$$

We next turn to the result for $\operatorname{argmin}_{w \in W} R(w)$ and let w^* be a point in that set. Then, there is k^* such that $w^* \in W_{k^*}$ and

$$R_m(w^{k^*}) \leq R(w^{k^*}) + \varepsilon \leq R(w^*) + \varepsilon = \inf_{w \in W} R(w) + \varepsilon.$$

Moreover, let $\bar{w} \in \operatorname{argmin}_{w \in W} R_m^\gamma(w)$. Then, there is \bar{k} such that $\bar{w} \in W_{\bar{k}}$ and

$$\inf_{w \in W} R_m^\gamma(w) = \sup_{\|v\| \leq \gamma} R_m(\bar{w} + v) \geq R_m(w_{\bar{k}}) \geq R(w_{\bar{k}}) - \varepsilon \geq \inf_{w \in W} R(w) - \varepsilon.$$

Combining these inequalities, we obtain that

$$R_m(w^{k^*}) \leq \inf_{w \in W} R_m^\gamma(w) + 2\varepsilon.$$

Since $\|w^* - w^{k^*}\| \leq \gamma$, we have established that

$$\operatorname{exs} \left(\operatorname{argmin}_{w \in W} R(w); \{w \in W \mid R_m(w) \leq \inf_{\bar{w} \in W} R_m^\gamma(\bar{w}) + 2\varepsilon\} \right) \leq \gamma$$

and the conclusion follows. \square

Acknowledgement. This work is supported in part by AFOSR under F4FGA08272G001.

References

- [1] P. Bartlett, D. Foster, and M. Telgarsky. Spectrally-normalized margin bounds for neural networks. In *Advances in Neural Information Processing Systems*, pages 6240–6249, 2017.
- [2] A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Mathematics of Operations Research*, 23(4):769–805, 1998.
- [3] D. Bertsimas and M. Copenhaver. Characterization of the equivalence of robustification and regularization in linear and matrix regression. *European J. Operational Research*, 270(3):931–942, 2018.
- [4] D. Bertsimas, V. Gupta, and N. Kallus. Robust sample average approximation. *Mathematical Programming*, 171(1):217–282, 2018.
- [5] S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2016.
- [6] O. Bousquet and A. Elisseeff. Stability and generalization. *Journal of machine learning research*, 2(Mar):499–526, 2002.
- [7] Z. Charles and D. Papailiopoulos. Stability and generalization of learning algorithms that converge to global optima. In *International Conference on Machine Learning*, pages 744–753, 2018.
- [8] P. Chaudhari, A. Choromanska, S. Soatto, Y. LeCun, C. Baldassi, C. Borgs, J. Chayes, L. Sagun, and R. Zecchina. Entropy-sgd: Biasing gradient descent into wide valleys. In *International Conference on Learning Representations (ICLR)*, 2017.
- [9] J. Duchi, P. Glynn, and H. Namkoong. Statistics of robust optimization: A generalized empirical likelihood approach. *arXiv preprint arXiv:1610.03425*, 2018.
- [10] H. Gouk, E. Frank, B. Pfahringer, and M. Cree. Regularisation of neural networks by enforcing lipschitz continuity. *arXiv preprint arXiv:1804.04368*, 2018.
- [11] M. Hardt, B. Recht, and Y. Singer. Train faster, generalize better: Stability of stochastic gradient descent. In *International Conference on Machine Learning*, pages 1225–1234, 2016.
- [12] S. Hochreiter and J. Schmidhuber. Flat minima. *Neural Computation*, 9(1):1–42, 1997.
- [13] N. S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, and P. T. P. Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*, 2016.

- [14] A. Lewis. Robust regularization. Tech. report, School of ORIE, Cornell University, Ithaca, NY, 2002.
- [15] A. Lewis and C. Pang. Lipschitz behavior of the robust regularization. *SIAM J. Control and Optimization*, 48(5):3080–3104, 2010.
- [16] H. Li, Z. Xu, G. Taylor, C. Studer, and T. Goldstein. Visualizing the loss landscape of neural nets. In *Advances in Neural Information Processing Systems*, pages 6389–6399, 2018.
- [17] H. Liu, X. Wang, T. Yao, R. Li, and Y. Ye. Sample average approximation with sparsity-inducing penalty for high-dimensional stochastic programming. *Mathematical Programming*, 2018.
- [18] Z.-Q. Luo. Applications of convex optimization in signal processing and digital communication. *Mathematical Programming*, 97(1-2):177–207, 2003.
- [19] Z.-Q. Luo, J. Sturm, and S. Zhang. Multivariate nonnegative quadratic mappings. *SIAM J. Optimization*, 14(4):1140–1162, 2004.
- [20] H. Men, R. Freund, N. Nguyen, J. Saa-Seoane, and J. Peraire. Fabrication-adaptive optimization with an application to photonic crystal design. *Operations Research*, 62(2):418–434, 2014.
- [21] Q. Nguyen and M. Hein. The loss surface of deep and wide neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2603–2612. JMLR.org, 2017.
- [22] A. M. Oberman and J. Calder. Lipschitz regularized deep neural networks converge and generalize. *arXiv preprint arXiv:1808.09540*, 2018.
- [23] R. Oliveira and P. Thompson. Sample average approximation with heavier tails i: non-asymptotic bounds with weak assumptions and stochastic constraints. *arXiv preprint arXiv:1705.00822*, 2017.
- [24] S. Oymak, Z. Fabian, M. Li, and M. Soltanolkotabi. Generalization guarantees for neural networks via harnessing the low-rank structure of the jacobian. *arXiv preprint arXiv:1906.05392*, 2019.
- [25] M. Pinar and O. Arikan. On robust solutions to linear least squares problems affected by data uncertainty and implementation errors with application to stochastic signal modeling. *Linear Algebra and Its Applications*, 391:223–243, 2004.
- [26] H. Qian and M. N. Wegman. L2-nonexpansive neural networks. *arXiv preprint arXiv:1802.07896*, 2018.

- [27] R. T. Rockafellar and R. J-B Wets. *Variational Analysis*, volume 317 of *Grundlehren der Mathematischen Wissenschaft*. Springer, 3rd printing-2009 edition, 1998.
- [28] J. O. Royset. On sample size control in sample average approximations for solving smooth stochastic programs. *Computational Optimization and Applications*, 55(2):265–309, 2013.
- [29] J. O. Royset and R. J-B Wets. Variational theory for optimization under stochastic ambiguity. *SIAM J. Optimization*, 27(2):1118–1149, 2017.
- [30] J. O. Royset and R. Szechtman. Optimal budget allocation for sample average approximation. *Operations Research*, 61:762–776, 2013.
- [31] L. Sagun, L. Bottou, and Y. LeCun. Eigenvalues of the hessian in deep learning: Singularity and beyond. *arXiv preprint arXiv:1611.07476*, 2016.
- [32] S. Shalev-Shwartz, O. Shamir, N. Srebro, and K. Sridharan. Learnability, stability and uniform convergence. *Journal of Machine Learning Research*, 11(Oct):2635–2670, 2010.
- [33] A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on Stochastic Programming: Modeling and Theory*. SIAM, 2009.
- [34] E. Stinstra and D. den Hertog. Robust optimization using computer experiments. *European J. Operations Research*, 191:816–837, 2008.
- [35] W. Wiesemann, D. Kuhn, and M. Sim. Distributionally robust convex optimization. *Operations Research*, 62(6):1358–1376, 2014.
- [36] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.
- [37] J. Zhang, H. Xu, and L. Zhang. Quantitative stability analysis for distributionally robust optimization with moment constraints. *SIAM J. Optimization*, 26(3):1855–1882, 2016.