
The sample complexity of ERM in stochastic convex optimization

Daniel Carmon
Technion-IIT

Roi Livni
Tel Aviv University

Amir Yehudayoff
University of Copenhagen
Technion-IIT

Abstract

Stochastic convex optimization is one of the most well-studied models for learning in modern machine learning. Nevertheless, a central fundamental question in this setup remained unresolved:

how many data points must be observed so that any empirical risk minimizer (ERM) shows good performance on the true population?

This question was proposed by Feldman who proved that $\Omega(\frac{d}{\epsilon} + \frac{1}{\epsilon^2})$ data points are necessary (where d is the dimension and $\epsilon > 0$ is the accuracy parameter). Proving an $\omega(\frac{d}{\epsilon} + \frac{1}{\epsilon^2})$ lower bound was left as an open problem. In this work we show that in fact $\tilde{O}(\frac{d}{\epsilon} + \frac{1}{\epsilon^2})$ data points are also sufficient. This settles the question and yields a new separation between ERM and uniform convergence.

This sample complexity holds for the classical setup of learning bounded convex Lipschitz functions over the Euclidean unit ball. We further generalize the result and show that a similar upper bound holds for all symmetric convex bodies. The general bound is composed of two terms: (i) a term of the form $\tilde{O}(\frac{d}{\epsilon})$ with an inverse-linear dependence on the accuracy parameter, and (ii) a term that depends on the statistical complexity of the class of *linear* functions (captured by the Rademacher complexity). The proof builds a mechanism for controlling the behavior of stochastic convex optimization problems.

1 INTRODUCTION

Stochastic convex optimization (SCO) is a benchmark framework that is widely used for studying stochastic optimization algorithms such as gradient descent and its variants. This is often justified by the simplicity of the framework and the possibility of a rigorous analysis that can hint at the pros and cons of various optimization techniques in practical setups such as machine learning. It has also been studied in the optimization literature under the name sample average approximation (SAA); see for example (13; 14) and references within.

Stochastic convex optimization is particularly useful for understanding the interaction between optimization and generalization in complex scenarios. The works of (30), and subsequently (31), demonstrated how the choice of an algorithm is crucial not only for optimization reasons, but also for generalization reasons. Namely, to avoid overfitting without careful algorithmic choices, one must use dimension-dependent sample size. On the other hand, with the correct algorithm, one can avoid overfitting with far less data points. For this reason, SCO became a prototypical model for researching *over-parameterization* (23; 19; 25; 24). These works aim to understand how can algorithms avoid overfitting, even when the number of data points is significantly smaller than the number of free parameters.

In more detail, a classical paradigm for learning is to draw n i.i.d. data points z_1, \dots, z_n from an unknown distribution and to optimize the *empirical risk* defined as

$$\hat{F}(x) = \frac{1}{n} \sum_{i=1}^n f(x, z_i),$$

where $f(x, z)$ is some loss function that measures the performance of parameter x on the data point z . When n is sufficiently large, a minimizer of \hat{F} should also demonstrate good performance on the population loss $F(x) = \mathbb{E}_{z \sim D}[f(x, z)]$. This is exactly the setting of SCO, except that SCO makes the additional assumption that the loss functions are convex (more details

are provided below).

It is a well-known fact that in the Probably Approximately Correct framework (20; 2), all optimization algorithms share the same statistical rate (20; 21). Namely, the statistical rate of *all* learning algorithms is the same. And understanding the rate of *any* learning algorithm is equivalent, then, to understanding the rate of any *specific* ERM. In stochastic convex optimization, however, not all algorithms are equal. While there are learning algorithms that perform well with only $n = O\left(\frac{1}{\varepsilon^2}\right)$ data points (30; 3), others must observe at least $\Omega\left(\frac{d}{\varepsilon}\right)$ points to guarantee good performance where d is the dimension of the parameter space $x \in \mathbb{R}^d$ (31).

And yet, even though separations between the statistical performances of different ERMs are known, one of the most fundamental questions in stochastic convex optimization remained unanswered:

what is the worst-case sample complexity of ERMs in SCO?

A standard covering argument (e.g. (10; 5)) shows it to be at most $\tilde{O}\left(\frac{d}{\varepsilon^2}\right)$. (31) demonstrated it to be at least $\Omega\left(\frac{d}{\varepsilon} + \frac{1}{\varepsilon^2}\right)$. The gap between the lower bound and the upper bound remained open, and was proposed as an open problem in (31).

1.1 Our contribution

We resolve the aforementioned open problem and complete the picture by proving that the sample complexity of ERMs in SCO is actually at most

$$\tilde{O}\left(\frac{d}{\varepsilon} + \frac{1}{\varepsilon^2}\right).$$

Notice that this bound separates the statistical complexity of ERMs from uniform convergence. The statistical complexity of uniform convergence is the number of examples required so that typically for *every* model in the parameter space, the empirical loss is close to the true population loss. In SCO, (31) provided a lower bound that matches the well-known upper bound (5; 10) and demonstrated that the statistical complexity of uniform convergence is $\tilde{\Theta}\left(\frac{d}{\varepsilon^2}\right)$. Our result shows that, while different ERMs may exhibit different performances, all ERMs are better than the worst-case uniform convergence rate. Even without algorithmic assumptions, learning in SCO is easier than uniform convergence (distinctively from other classical models such as PAC learning).

A second notable aspect of our result is that, usually, the term $\frac{1}{\varepsilon^2}$ corresponds to noise or to the agnostic setting of learning. The term $\frac{d}{\varepsilon}$, on the other hand,

corresponds to a realizable setting. It is somewhat surprising that these two different types of “behaviors” are simultaneously appearing in a single model. In stochastic convex optimization, the ERM principle somehow intrinsically combines low dimensional agnostic learning and high-dimensional realizable learning.

The above result is true for learning bounded convex Lipschitz functions with respect to an ℓ_2 -bounded domain. We generalize our result and analyze the statistical complexity of ERMs with respect to general norms. For any norm and its corresponding unit ball \mathcal{K} , and for learning bounded convex Lipschitz functions, the statistical complexity of ERMs with error $O(\varepsilon)$ is at most

$$\tilde{O}\left(\frac{d}{\varepsilon} + \text{Rad}_{\mathcal{K}}^{-1}(\varepsilon)\right),$$

where $\text{Rad}_{\mathcal{K}}^{-1}(\varepsilon)$ measures the statistical complexity of learning over \mathcal{K} when we restrict the observed losses to be linear. The first term is the dimension-dependent part of the statistical complexity and is the same for all norms. The second term, which is the dimension-independent part for the ℓ_2 norm, is not the classical Rademacher complexity of the full class (“convex functions”), but is the Rademacher complexity of a smaller class (“linear functions”). The above result yields a tight bound for the ℓ_2 norm, but also yields tight bounds for general ℓ_p norms ((31) proved matching lower bounds for ℓ_p norms).

2 FORMAL SETUP AND MAIN RESULT

Let us formally state the algorithmic problem we explore. A one-dimensional illustration appears in ???. We consider a general norm $\|\cdot\|$ on \mathbb{R}^d and its unit ball \mathcal{K} . We also consider a general domain Z , which for concreteness we assume to be finite. Our result hold when Z is infinite, but then one needs to carefully deal with measurability issues. We refer to (6) for the relevant definitions and formal setup in which our result hold in full generality.

For each $z \in Z$, we consider a function $f_z(x)$ that is convex and L -Lipschitz with respect to x . That is, for every x, y ,

$$|f_z(x) - f_z(y)| \leq L \|x - y\|.$$

We are interested in the problem of optimization over the parameter domain \mathcal{K} . We note, though, that there is no loss of generality in assuming f_z is defined over the whole space \mathbb{R}^d , because every L -Lipschitz convex function over \mathcal{K} can be extended to an L -Lipschitz convex function over all of \mathbb{R}^d .

We next assume a distribution \mathcal{D} that is supported on the domain Z . We consider the convex function F defined by

$$F(x) = \mathbb{E}_{z \sim \mathcal{D}}[f_z(x)].$$

The function F is often referred to as the *population loss* or *true risk*, and is also L -Lipschitz. The algorithmic goal is to find a minimizer x^* of F . The set of minimizers of F in \mathcal{K} is

$$\operatorname{argmin}_{\mathcal{K}} F = \{x^* \in \mathcal{K} : \forall x \in \mathcal{K} \ F(x^*) \leq F(x)\}.$$

The learning algorithm does not know the distribution \mathcal{D} . Instead, one assumes that the learner can observe an i.i.d. sample $S = (z_1, \dots, z_n)$ of n examples drawn from the distribution \mathcal{D} . A natural strategy in this setting is to follow the ERM principle. Given S , we consider the *empirical risk* or *empirical loss*

$$\hat{F} = \frac{1}{n} \sum_{j \in [n]} f_{z_j}.$$

An ERM learning rule is any algorithm that outputs

$$\hat{x} \in \operatorname{argmin}_{\mathcal{K}} \hat{F}.$$

For an accuracy parameter $\epsilon > 0$, the ERM principle is considered successful if

$$F(\hat{x}) \leq F(x^*) + \epsilon.$$

We would like to have the guarantee that with as few samples as possible, any algorithm that follows the ERM principle is successful.

Our results provide such a guarantee. We first state our result in the cleanest form possible. Accordingly, the only two parameters we focus on are the dimension d and accuracy parameter ϵ , and we work over the ℓ_2 norm and its unit ball

$$\mathcal{B} = \{x \in \mathbb{R}^d : \|x\|_2 \leq 1\}.$$

General norms and other parameters are treated in Section 2.1 below.

Theorem 1. *Suppose that for every $z \in Z$, the function f_z is convex, 1-Lipschitz with respect to $\|\cdot\|_2$ and for every $x \in \mathcal{B}$ we have $|f_z(x)| \leq 1$. For every distribution \mathcal{D} on Z , and for all accuracy parameters $0 < \epsilon < 1$, if*

$$n \geq n_0 = \frac{3d \ln(\frac{40}{\epsilon})}{\epsilon} + \frac{40}{\epsilon^2}$$

then with probability at least $\frac{3}{4}$ over $S \sim \mathcal{D}^n$, for all $x \in \mathcal{B}$,

$$\hat{F}(x) \leq \hat{F}(x^*) + \epsilon \implies F(x) \leq F(x^*) + 40\epsilon.$$

The theorem shows that for $n \geq n_0$, not only that the ERM principle is successful, but it is also robust in the sense that it is safe to output any point \hat{x} that is merely *close* to being a minimizer of \hat{F} . There is a distinction between ERMs (algorithms that output a minimizer of \hat{F}) and approximate ERMs (algorithms that output an approximate minimizer of \hat{F}). This subtle distinction is important in SCO. Approximate ERMs can be more efficient than exact ERMs; see (10). For example, for $\epsilon \geq \Omega(1)$, for all problems, sample complexity of $O(1)$ can be achieved by a regularized ERM which is an approximate ERM, but for some problems, any exact ERM must have sample complexity $\Omega(\ln d)$.

2.1 General norms

We next provide a more fine-grained result that generalizes Theorem 1 for arbitrary norms and their unit balls. We start by introducing the necessary terminology. Recall that $\|\cdot\|$ is a norm on \mathbb{R}^d and $\mathcal{K} \subset \mathbb{R}^d$ is its unit ball. Denote by $\|\cdot\|_*$ the dual norm

$$\|g\|_* = \sup_{x \in \mathcal{K}} \langle g, x \rangle.$$

The Rademacher complexity of \mathcal{K} with respect to a sample $S = (g_1, \dots, g_n) \in (\mathbb{R}^d)^n$ is

$$\operatorname{Rad}(\mathcal{K}, S) = \mathbb{E}_{\sigma} \left[\left\| \frac{1}{n} \sum_{j=1}^n \sigma_j g_j \right\|_* \right],$$

where the expectation is over Rademacher random variables $\sigma_1, \dots, \sigma_n$ (i.e., i.i.d. uniform in $\{\pm 1\}$). Observe that $\operatorname{Rad}(\mathcal{K}, S)$ corresponds to the standard Rademacher complexity if we think of \mathcal{K} as a class of linear functions operating on the dual ball; indeed, for every S ,

$$\operatorname{Rad}(\mathcal{K}, S) = \mathbb{E}_{\sigma} \left[\sup_{x \in \mathcal{K}} \frac{1}{n} \sum_{j=1}^n \sigma_j \langle g_j, x \rangle \right].$$

For an integer n , define

$$\operatorname{Rad}(\mathcal{K}, n) = \sup_S \operatorname{Rad}(\mathcal{K}, S),$$

where the supremum is over $S = (g_1, \dots, g_n)$ so that $\|g_j\|_* \leq 1$ for all $j \in [n]$. For $\epsilon > 0$, define

$$\operatorname{Rad}_{\mathcal{K}}^{-1}(\epsilon) = \min \{n \in \mathbb{N} : \operatorname{Rad}(\mathcal{K}, n) < \epsilon\}.$$

It is worth noting that because

$$\begin{aligned} & \mathbb{E} \left[\left\| \frac{1}{n+1} \sum_{j=1}^{n+1} \sigma_j g_j \right\|_{\star} \right] \\ &= \mathbb{E} \left[\left\| \frac{1}{n+1} \sum_{I \subset [n+1]: |I|=n} \frac{1}{n} \sum_{i \in I} \sigma_i g_i \right\|_{\star} \right] \\ &\leq \frac{1}{n+1} \sum_{I \subset [n+1]: |I|=n} \mathbb{E} \left[\left\| \frac{1}{n} \sum_{i \in I} \sigma_i g_i \right\|_{\star} \right], \end{aligned}$$

the following monotonicity holds:

$$\text{Rad}(\mathcal{K}, n+1) \leq \text{Rad}(\mathcal{K}, n).$$

Therefore, if $n \geq \text{Rad}_{\mathcal{K}}^{-1}(\varepsilon)$ then $\text{Rad}(\mathcal{K}, n) < \varepsilon$.

The general bound on the sample complexity of ERM is described in the following theorem.

Theorem 2. *Suppose that for every $z \in Z$, the function f_z is convex, L -Lipschitz with respect to $\|\cdot\|$ and for every $x \in \mathcal{K}$ we have $|f_z(x)| \leq c$. For every distribution \mathcal{D} on Z , and for all $0 < \varepsilon, \delta < 1$, if*

$$n \geq n_0 = \frac{12cd}{\varepsilon} \ln\left(\frac{3L}{\varepsilon}\right) + \text{Rad}_{\mathcal{K}}^{-1}\left(\frac{\varepsilon}{2L}\right) + \frac{8c^2}{\varepsilon^2} \ln\left(\frac{4}{\delta}\right)$$

then with probability at least $1 - \delta$ over $S \sim \mathcal{D}^n$, for all $x \in \mathcal{K}$,

$$\hat{F}(x) \leq \hat{F}(x^*) + \varepsilon \implies F(x) \leq F(x^*) + 40\varepsilon.$$

The sample complexity n_0 consists of three different terms. The first term is of the form $\frac{cd}{\varepsilon} \ln(\frac{L}{\varepsilon})$, and also appears in the ℓ_2 setting. In a nutshell, the $d \ln(\frac{L}{\varepsilon})$ part comes from the logarithm of the cover number (see Lemma 1 below). The second term $\text{Rad}_{\mathcal{K}}^{-1}(\frac{\varepsilon}{L})$ correspond to learning the special case of linear functions. It heavily depends on the norm and its unit ball \mathcal{K} . For example, when $\frac{\varepsilon}{L} \geq \Omega(1)$, for the ℓ_2 norm it is $O(1)$, for the ℓ_1 norm it is $O(\ln d)$, and for the ℓ_∞ norm it is $O(d)$. The third term $O(\frac{c^2}{\varepsilon^2} \ln(\frac{1}{\delta}))$ is the only place the confidence parameter δ appears in.

3 ON THE PROOF

We now outline the proof of Theorem 1. The proof has two central ingredients. One ingredient is a first-order optimality condition for stochastic convex optimization problems (Theorem 3 in Section 4.1). The second ingredient connects properties of ERM to the Bregman divergence and prove a strong concentration bound for it. The mechanism underlying these steps may be useful for other convex optimization problems. The proof of the theorem is finally given in Section 5.1

The proof of Theorem 2 requires a couple of additional ideas, which are explained in Section 6.

A key property is identified in Theorem 6. This claim allows to replace the $\frac{d}{\varepsilon^2}$ term which is “expected to hold” by the term $\frac{d}{\varepsilon}$ which is “correct”. The intuition for this gain is that distinguishing between a coin with bias $\frac{1}{2} + \varepsilon$ and a coin with bias $\frac{1}{2} - \varepsilon$ requires $\frac{1}{\varepsilon^2}$ samples, but distinguishing between a coin with bias ε and a coin with bias 2ε can be done with only $\frac{1}{\varepsilon}$ samples.

Where is the coin? For the simplicity of the exposition, we assume that the f_z ’s are smooth and that the minimizer x^* is in the interior of \mathcal{B} (we also included an illustration in Figure 1). In this case, F is also smooth and first-order optimality implies that $\nabla F(x^*) = 0$. For each $x \in \mathcal{B}$, we know that $F(x) - F(x^*) \geq 0$, and our goal is to control the empirical difference $\hat{F}(x) - \hat{F}(x^*)$. This is achieved, in a nutshell, through the non-negativity of the Bregman divergence. The Bregman divergence associated with f_z between x and x^* is

$$D_{f_z}(x, x^*) := f_z(x) - f_z(x^*) - \langle \nabla f_z(x^*), x - x^* \rangle.$$

It is always non-negative, and it is also bounded from above by four. The key observation is that

$$F(x) - F(x^*) = D_F(x, x^*) = \mathbb{E}_{z \sim \mathcal{D}}[D_{f_z}(x, x^*)],$$

which is explained by

$$\mathbb{E}_{z \sim \mathcal{D}}[\nabla f_z(x^*)] = \nabla F(x^*) = 0.$$

We also see that “ x is a bad output” when $D_F(x, x^*) \geq 2\varepsilon$. Instead of directly controlling $\hat{F}(x) - \hat{F}(x^*)$, we aim to control the empirical Bregman divergence

$$D_{\hat{F}}(x, x^*) = \hat{F}(x) - \hat{F}(x^*) - \langle \nabla \hat{F}(x^*), x - x^* \rangle.$$

We can think of $D_F(x, x^*)$ as the bias of a coin taking values in $[0, 4]$. The non-negativity and boundedness of the Bregman divergence allows to prove, via Bernstein inequality, that if $D_F(x, x^*) \geq 2\varepsilon$ then

$$\Pr[D_{\hat{F}}(x, x^*) \leq \varepsilon] \leq \exp(-\Omega(\varepsilon n)).$$

This seems pretty close to our goal; if “ x is a bad output” then it is very likely that $D_{\hat{F}}(x, x^*)$ is large. Because we actually care about $\hat{F}(x) - \hat{F}(x^*)$ and not $D_{\hat{F}}(x, x^*)$, we also need to control $\nabla \hat{F}(x^*)$. This can be achieved via a simple concentration argument because

$$\mathbb{E}[\|\nabla \hat{F}(x^*)\|_2^2] \leq \frac{1}{n}.$$

By taking a standard union bound over an ε -net N for \mathcal{B} of size $|N| \approx (\frac{1}{\varepsilon})^d$, we get the final bound

$$\begin{aligned} & \Pr[\exists x \in \mathcal{B} F(x) - F(x^*) > 4\varepsilon, \hat{F}(x) - \hat{F}(x^*) < \varepsilon] \\ & \leq \frac{1}{\varepsilon^2 n} + |N| \exp(-\Omega(\varepsilon n)). \end{aligned}$$




Figure 1: An illustration of the role the Bregman divergence plays in the proof. The graphs of F and \hat{F} are from ???. The solid line is the tangent to F at $x^* = 0$, and the dashed line is the tangent to \hat{F} at x^* . The tangents are always below the graphs due to convexity. The differences between the graphs and the tangents are the Bregman divergences. By controlling the tangent to \hat{F} , we can control the behavior of ERM. Controlling the tangent to \hat{F} is easier than controlling \hat{F} because it is an affine function.

This argument actually works also when the f_z 's are *not* assumed to be smooth. For this, we need to replace “gradients” by “subgradients”; see Section 4.1 below. But when x^* is on the boundary of \mathcal{B} , even in the smooth case, we do not know that $\nabla F(x^*) = 0$. In this case, we additionally need to understand what happens when $D_F(x, x^*) < 2\epsilon$. This is achieved in Theorem 4.

4 PRELIMINARIES

We next develop several preliminary results, and provide the background needed for the proof.

4.1 First-order optimality

The subgradient of $f : \mathbb{R}^d \rightarrow \mathbb{R}$ at $x \in \mathbb{R}^d$ is

$$\partial f(x) = \{g \in \mathbb{R}^d : \forall y \in \mathbb{R}^d f(y) \geq f(x) + \langle g, y - x \rangle\}.$$

If f is convex then the set $\partial f(x)$ is never empty (33). If f is convex and L -Lipschitz with respect to $\|\cdot\|$ then

$$g \in \partial f(x) \implies \|g\|_* \leq L.$$

A deep property of stochastic functions follows from the subgradient sum property (because it is a central idea, we included a proof in the material). For a stochastic convex function, the subgradient can be

thought as the expectation of the subgradients (6); namely, if $F = \mathbb{E}_{z \sim \mathcal{D}}[f_z]$, then for every $g \in \partial F(x)$, for each $z \in Z$, there exists $g_z \in \partial f_z(x)$ so that

$$g = \mathbb{E}_{z \sim \mathcal{D}}[g_z].$$

This leads to the following stochastic first-order optimality condition.

Proposition 3 (Stochastic first-order condition). *Let $\|\cdot\|$ be a norm on \mathbb{R}^d and denote by \mathcal{K} its unit ball. Let \mathcal{D} be a distribution on the finite set Z . Suppose that for every $z \in Z$, the function f_z is convex and L -Lipschitz. Assume that*

$$x^* \in \operatorname{argmin}_{\mathcal{K}} F$$

where $F = \mathbb{E}_{z \sim \mathcal{D}}[f_z]$. Then, for each $z \in Z$, there is $g_z \in \partial f_z(x^*)$ so that $\|g_z\|_* \leq L$ and for all $x \in \mathcal{K}$,

$$\langle G, x - x^* \rangle \geq 0 \quad (1)$$

where $G = \mathbb{E}_{z \sim \mathcal{D}}[g_z] \in \partial F(x^*)$.

Proof. The first-order optimality condition (see e.g. (33; 7)) applied to F states that there is $G \in \partial F(x^*)$ so that for all $x \in \mathcal{K}$,

$$\langle G, x - x^* \rangle \geq 0.$$

By the subgradient sum property, there are $g_z \in \partial f_z(x)$ such that $G = \mathbb{E}[g_z]$. Because f_z is L -Lipschitz, we know that $\|g_z\|_* \leq L$. \square

4.2 The Bregman divergence

The Bregman divergence measures the difference between a convex function and its first-order approximation. It is often defined for smooth functions, but for our purposes, we require a definition that makes sense for general SCO problems. For a stochastic function $F = \mathbb{E}_{z \sim \mathcal{D}}[f_z]$ and a point x^* , the Bregman divergence associated with F at x^* is

$$D(x, x^*) = D_F(x, x^*) = F(x) - F(x^*) - \langle G, x - x^* \rangle,$$

for some $G \in \partial F(x^*)$. This mapping depends on the choice of G , but to avoid cumbersome notations we will suppress this dependence. Importantly, we care about the Bregman divergence at x^* which is a minimizer of F , and we assume that $G \in \partial F(x^*)$ is chosen to be a sub-gradient whose existence is guaranteed by Theorem 3.

Given a sample $S = (z_1, \dots, z_n)$, the empirical Bregman divergence is

$$\hat{D}(x, x^*) = \hat{F}(x) - \hat{F}(x^*) - \langle \hat{G}, x - x^* \rangle,$$

where $\hat{G} = \frac{1}{n} \sum_{i=1}^n g_{z_i} \in \partial \hat{F}(x^*)$. Two key observations are that by the convexity of all the f_z 's, both $D(x, x^*)$ and $\hat{D}(x, x^*)$ are non-negative, and

$$\mathbb{E}_{S \sim \mathcal{D}^n}[\hat{D}(x, x^*)] = D(x, x^*).$$

4.3 Covering numbers

For $\varepsilon > 0$, the ε -cover number $\text{Cov}_{\mathcal{K}}(\varepsilon)$ of \mathcal{K} is the minimum integer m so that there is a net $N \subset \mathcal{K}$ of cardinality $|N| = m$ so that for every $x \in \mathcal{K}$ there is $y \in N$ so that $\|x - y\| \leq \varepsilon$. The following is a standard bound on the cover numbers of unit balls of norms.

Lemma 1. *Let $\|\cdot\|$ be a norm on \mathbb{R}^d and denote by \mathcal{K} its unit ball. For every $\varepsilon > 0$,*

$$\text{Cov}_{\mathcal{K}}(\varepsilon) \leq \left(\frac{2(1+\varepsilon)}{\varepsilon} \right)^d.$$

Proof. Let Y be a subset of \mathcal{K} of maximum size so that for every $y \neq y'$ in Y we have $\|y - y'\| > \varepsilon$. It follows that

$$\begin{aligned} (1+\varepsilon)^d |\mathcal{K}| &= |(1+\varepsilon)\mathcal{K}| \geq \left| \bigcup_{y \in Y} \left(y + \frac{\varepsilon}{2}\mathcal{K} \right) \right| \\ &= |Y| \cdot \left| \frac{\varepsilon}{2}\mathcal{K} \right| = \left(\frac{\varepsilon}{2} \right)^d \cdot |Y| \cdot |\mathcal{K}|, \end{aligned}$$

where $|\cdot|$ denotes volume. Because Y is maximal, for every $x \in \mathcal{K}$ there is $y \in Y$ so that $\|x - y\| \leq 2\varepsilon$. \square

4.4 Bernstein's inequality

Finally, we recall the following concentration bound, due to Bernstein:

Lemma (Bernstein's inequality). *Let R_1, \dots, R_n be i.i.d. random variables taking values in $[-M, M]$, each with mean zero and variance σ^2 . Then, for all $t > 0$, we have*

$$\Pr \left[\sum_{j \in [n]} R_j \geq t \right] \leq \exp \left(- \frac{t^2}{2\sigma^2 n + (2Mt/3)} \right).$$

5 THE EUCLIDEAN NORM

In this section, we prove Theorem [1](#). Let x^* be an arbitrary minimizer of F in \mathcal{B} . Let $G \in \partial F(x^*)$ be the subgradient given in Theorem [3](#). For a fixed sample, we also denote:

$$\hat{G} = \frac{1}{n} \sum_{i=1}^n g_{z_i}.$$

We start the proof by providing both a lower bound on the Bregman divergence at spurious empirical risk minimizers, as well as an upper bound for the empirical Bregman divergence at empirical risk minimizers.

The following claim says that for a point to be a spurious risk minimizer, the Bregman divergence must be “large”.

Claim 4. *For all $x \in \mathcal{B}$, if*

$$(\hat{F}(x) - \hat{F}(x^*) \leq 5\varepsilon) \wedge (\|\hat{G} - G\|_2 \leq \varepsilon)$$

then

$$D(x, x^*) \geq F(x) - F(x^*) - 7\varepsilon.$$

Proof. Because

$$\begin{aligned} \langle G, x - x^* \rangle &= \langle G - \hat{G}, x - x^* \rangle + \langle \hat{G}, x - x^* \rangle \\ &\leq 2\varepsilon + \langle \hat{G}, x - x^* \rangle \\ &= 2\varepsilon + \hat{F}(x^*) + \langle \hat{G}, x - x^* \rangle - \hat{F}(x^*) \\ &= 2\varepsilon + \hat{F}(x) - D(x, x^*) - \hat{F}(x^*) \\ &\leq 2\varepsilon + \hat{F}(x) - \hat{F}(x^*) \\ &\leq 7\varepsilon, \end{aligned}$$

we can bound

$$\begin{aligned} D(x, x^*) &= F(x) - F(x^*) - \langle G, x - x^* \rangle \\ &\geq F(x) - F(x^*) - 7\varepsilon. \end{aligned}$$

\square

Next, we relate data of the form “ x is an ERM” to an upper bound on the empirical Bregman divergence.

Claim 5. *For all $x \in \mathcal{B}$, if*

$$(\hat{F}(x) - \hat{F}(x^*) \leq 5\varepsilon) \wedge (\|\hat{G} - G\|_2 \leq \varepsilon)$$

then

$$\hat{D}(x, x^*) \leq 7\varepsilon.$$

Proof.

$$\begin{aligned} \hat{D}(x, x^*) &= \hat{F}(x) - \hat{F}(x^*) - \langle \hat{G}, x - x^* \rangle \\ &\leq 5\varepsilon - \langle \hat{G}, x - x^* \rangle \\ &= 5\varepsilon - \langle \hat{G} - G, x - x^* \rangle - \langle G, x - x^* \rangle \\ &\leq 5\varepsilon - \langle \hat{G} - G, x - x^* \rangle \\ &\leq 7\varepsilon. \end{aligned}$$

Eq. [\(1\)](#)

\square

So far, we showed that for a point to be a spurious empirical risk minimizer, the Bregman divergence must be large yet the empirical Bregman divergence is small. The last claim we need is a concentration bound that relates the empirical and true Bregman divergences.

Claim 6. *Suppose that for each $z \in Z$, the function f_z is convex, 1-Lipschitz with respect to $\|\cdot\|_2$, and for all $x \in \mathcal{B}$ we have $|f_z(x)| \leq 1$. Then, for every $x \in \mathcal{B}$,*

$$\Pr \left[\hat{D}(x, x^*) \leq \frac{D(x, x^*)}{2} \right] \leq \exp \left(- \frac{D(x, x^*)n}{40} \right),$$

where the probability is over z_1, \dots, z_n i.i.d. samples from \mathcal{D} .

Proof. The random variable

$$\hat{D}(x, x^*) = \frac{1}{n} \sum_j D_{z_j}(x, x^*)$$

is the average of n i.i.d. variables R_1, \dots, R_n , each taking values in $[0, 4]$ with expectation $\lambda := D(x, x^*)$. By the Bhatia–Davis inequality, the variance of the random variables is at most

$$(4 - \lambda)\lambda \leq 4\lambda.$$

By Bernstein’s inequality,

$$\begin{aligned} \Pr \left[\hat{D}(x, x^*) \leq \frac{\lambda}{2} \right] &= \Pr \left[\sum_j (R_j - \lambda) \leq -\frac{\lambda n}{2} \right] \\ &\leq \exp \left(-\frac{\lambda^2 n^2}{4(8\lambda n + 2\lambda n)} \right) \\ &\leq \exp \left(-\frac{\lambda n}{40} \right). \end{aligned}$$

□

5.1 Putting it together

Assume that $n \geq n_0$. We start with a simple claim that controls the behavior of the linear part, and we show that:

$$\Pr[\|\hat{G} - G\|_2^2 > \varepsilon^2] \leq \frac{4}{\varepsilon^2 n}. \quad (2)$$

Indeed,

$$\begin{aligned} \mathbb{E}\|\hat{G} - G\|_2^2 &= \frac{1}{n^2} \sum_{j, j'} \mathbb{E}\langle g_{z_j} - G, g_{z_{j'}} - G \rangle \\ &= \frac{1}{n^2} \sum_j \mathbb{E}\|g_{z_j} - G\|^2 \leq \frac{4}{n}. \end{aligned}$$

Now, Eq. (2) follows from Markov’s inequality.

Next, denote by E the event that there is $x \in \mathcal{B}$ so that

$$(\hat{F}(x) - \hat{F}(x^*) \leq \varepsilon) \wedge (F(x) - F(x^*) > 40\varepsilon).$$

To prove the final result, we need to show that

$$\Pr(E) \leq \frac{1}{4}. \quad (3)$$

Towards this, denote by H the event that

$$\|\hat{G} - G\|_2 \leq \varepsilon.$$

Next, let N be an $\frac{\varepsilon}{3}$ -net for \mathcal{B} of size at most $(12/\varepsilon)^d$. Denote by \hat{N} the set of points $x \in N$ so that

$$(\hat{F}(x) - \hat{F}(x^*) \leq 5\varepsilon) \wedge (F(x) - F(x^*) > 35\varepsilon).$$

Because F and \hat{F} are Lipschitz, the net property implies that if E holds then \hat{N} is a non-empty set. By Theorem 4, conditioned on H , if $x \in \hat{N}$ then

$$D(x, x^*) \geq 16\varepsilon.$$

On the other hand, for every x , conditioned on H , if $x \in \hat{N}$ then by Theorem 5 we have

$$\hat{D}(x, x^*) \leq 8\varepsilon.$$

Let E_x be the event that $\hat{D}(x, x^*) \leq 8\varepsilon$. We obtain, therefore:

$$E \cap H \subseteq \bigcup_{x \in N: D(x, x^*) > 16\varepsilon} E_x.$$

By Theorem 6 if $D(x, x^*) > 16\varepsilon$ then

$$\Pr[E_x] \leq \exp \left(-\frac{\varepsilon n}{3} \right).$$

By the union bound, Lemma 1, Eq. (2), and because $n \geq n_0$,

$$\begin{aligned} \Pr[E] &\leq \Pr[\neg H] + |N| \exp \left(-\frac{\varepsilon n}{3} \right) \\ &\leq \frac{4}{\varepsilon^2 n} + \left(\frac{12}{\varepsilon} \right)^d \exp \left(-\frac{\varepsilon n}{3} \right) \\ &\leq \frac{1}{4}. \end{aligned}$$

This completes the proof of Theorem 1.

6 GENERAL NORMS

In this section, we prove Theorem 2. Our first task is to handle the Lipschitz constant L and the bound c on the ℓ_∞ norm. To properly handle general L and c , we need to introduce a truncated version of the Bregman divergence. Second, we address general confidence parameters. This requires a uniform concentration result for the gradient at the optimum, combined with an additional argument that bounds a contractive term that may depend on it.

For the rest of this section, we fix a norm $\|\cdot\|$ and its unit ball \mathcal{K} . We also fix an optimal point x^* , and using Theorem 3 we fix subgradients $g_z \in \partial f_z(x^*)$ such that the first-order condition Eq. (1) holds for $G = \mathbb{E}_{z \sim \mathcal{D}}[g_z]$. As before, we denote by \hat{G} the empirical subgradient

$$\hat{G} = \frac{1}{n} \sum_{j=1}^n g_{z_j}.$$

6.1 The Lipschitz and ℓ_∞ constants

Special care is required when $c \ll L$. The reason is that the Bregman divergence can be as large as L regardless of c . We, therefore, introduce a truncated

version of the Bregman divergence. The truncated divergence associated with $f : \mathcal{K} \rightarrow [-c, c]$ is

$$T_f(x, x^*) = f(x) - f(x^*) - \max\{-2c, \langle g, x - x^* \rangle\}.$$

The truncated divergence, as the Bregman divergence, is always non-negative. But, in contrast to the Bregman divergence, it is always at most $4c$ because f is bounded by c . Similarly to the Bregman divergence, we use the following notation:

$$T_z = T_{f_z}, \quad T = \mathbb{E}_{z \sim \mathcal{D}} T_z \quad \text{and} \quad \hat{T} = \frac{1}{n} \sum_{j=1}^n T_{z_j}.$$

The random variable $\hat{T}(x, x^*)$ is the average of n i.i.d. variables taking values in $[0, 4c]$ with expectation $T(x, x^*)$. Analogously to Theorem 6, we can deduce that:

Claim 7. For every $x \in \mathcal{K}$,

$$\Pr \left[\hat{T}(x, x^*) \leq \frac{T(x, x^*)}{2} \right] \leq \exp \left(- \frac{T(x, x^*)n}{40c} \right).$$

We now want to develop analogue statements to Theorem 4 and Theorem 5. For that it will be convenient to add the following notation

$$\ell_g(a) = \max\{-2c, a - \langle g, x^* \rangle\}. \quad (4)$$

The function $\ell_g : \mathbb{R} \rightarrow \mathbb{R}$ is always convex and 1-Lipschitz. Denote by \mathcal{L} the function $\mathcal{L}(x) = \mathbb{E}_{z \sim \mathcal{D}}[\ell_{g_z}(\langle g_z, x \rangle)]$ and by $\hat{\mathcal{L}}$ its empirical analogue. Both \mathcal{L} and $\hat{\mathcal{L}}$ are convex. For a sample $S = (z_1, \dots, z_n)$, the *representativeness*¹ of S is defined to be

$$\text{Rep}(S) = \sup_{x \in \mathcal{K}} (\mathcal{L}(x) - \hat{\mathcal{L}}(x)).$$

Claim 8. Suppose that $\text{Rep}(S) \leq 2\varepsilon$. Then, for every $x \in \mathcal{K}$, if

$$\hat{F}(x) - \hat{F}(x^*) \leq 6\varepsilon$$

then

$$\hat{T}(x, x^*) \leq 8\varepsilon.$$

Proof.

$$\begin{aligned} \hat{T}(x, x^*) &= \hat{F}(x) - \hat{F}(x^*) - \hat{\mathcal{L}}(x) \\ &\leq 6\varepsilon + \mathcal{L}(x) - \hat{\mathcal{L}}(x) - \mathcal{L}(x) \\ &\leq 8\varepsilon - \mathbb{E}_{z \sim \mathcal{D}} \max\{-2c, \langle g_z, x - x^* \rangle\} \\ &\leq 8\varepsilon - \max\{-2c, \mathbb{E}_{z \sim \mathcal{D}} \langle g_z, x - x^* \rangle\} \\ &\leq 8\varepsilon - \langle G, x - x^* \rangle \\ &\leq 8\varepsilon. \end{aligned} \quad \text{Eq. (1)}$$

□

¹That is the terminology from the book (10).

Claim 9. Suppose that $\text{Rep}(S) \leq 2\varepsilon$. For all $x \in \mathcal{K}$ if

$$\hat{F}(x) - \hat{F}(x^*) \leq 5\varepsilon$$

then

$$T(x, x^*) \geq F(x) - F(x^*) - 7\varepsilon.$$

Proof. Because

$$\begin{aligned} \mathcal{L}(x) &= \mathcal{L}(x) - \hat{\mathcal{L}}(x) + \hat{\mathcal{L}}(x) && \text{Rep}(S) \leq 2\varepsilon \\ &\leq 2\varepsilon + \hat{\mathcal{L}}(x) \\ &= 2\varepsilon + \hat{F}(x^*) + \hat{\mathcal{L}}(x) - \hat{F}(x^*) \\ &\leq 2\varepsilon + \hat{F}(x) - \hat{T}(x, x^*) - \hat{F}(x^*) \\ &\leq 2\varepsilon + \hat{F}(x) - \hat{F}(x^*) && \hat{T}(x, x^*) \geq 0 \\ &\leq 7\varepsilon, \end{aligned}$$

we can bound

$$\begin{aligned} T(x, x^*) &= F(x) - F(x^*) - \mathcal{L}(x) \\ &\geq F(x) - F(x^*) - 7\varepsilon. \end{aligned}$$

□

6.2 The linear part

In this section, we explain how to control $\text{Rep}(S)$. We are going to use standard results on the Rademacher complexity, and for that we use the book (10). In the book, the Rademacher complexity of a set $A \subset \mathbb{R}^n$ is defined to be (and denoted by R):

$$\text{Rad}(A) = \mathbb{E}_\sigma \left[\sup_{a \in A} \frac{1}{n} \sum_{j=1}^n \sigma_j a_j \right].$$

For a fixed sample $S = (z_1, \dots, z_n)$, we consider the set

$$A = A_S = \{(\langle g_{z_1}, x \rangle, \dots, \langle g_{z_n}, x \rangle) \in \mathbb{R}^n : x \in \mathcal{K}\}$$

so that

$$\text{Rad}(A) = \text{Rad}(\mathcal{K}, S).$$

We also use the 1-Lipschitz maps $\ell_g : \mathbb{R} \rightarrow \mathbb{R}$ that are defined in Eq. (4). We denote by $\ell = \ell_S$ the sequence $\ell = (\ell_{g_{z_1}}, \dots, \ell_{g_{z_n}})$, and consider the set

$$\ell \circ A = \{(\ell_{g_{z_1}}(\langle g_{z_1}, x \rangle), \dots, \ell_{g_{z_n}}(\langle g_{z_n}, x \rangle)) \in \mathbb{R}^n : x \in \mathcal{K}\}.$$

It follows from Rademacher calculus (10, see Lemma 26.9 (contraction lemma)) that

$$\text{Rad}(\ell \circ A, S) \leq \text{Rad}(A) = \text{Rad}(\mathcal{K}, S). \quad (5)$$

We also rely on the McDiarmid's inequality following (see e.g. (10, Lemma 26.4)).

Lemma 2 (McDiarmid’s inequality). *Let $h : Z^n \rightarrow \mathbb{R}$ be so that for some $c_0 > 0$, for all $j \in [n]$ and for all $z_1, \dots, z_n, z'_i \in Z$,*

$$|h(z_1, \dots, z_n) - h(z_1, \dots, z_{i-1}, z'_i, z_{i+1}, \dots, z_n)| \leq c_0. \quad (6)$$

Let \mathcal{D} be a distribution on Z . Then, for all $\delta > 0$,

$$\Pr_{S \sim \mathcal{D}^n} \left[|h(S) - \mu| > c_0 \sqrt{\frac{\ln(2/\delta)n}{2}} \right] \leq \delta$$

where $\mu = \mathbb{E}_{S \sim \mathcal{D}^n} [h(S)]$.

We obtain the following important corollary (see also (I10, Theorem 26.5)).

Corollary 10. *Suppose that for each $z \in Z$, the function f_z is convex and L -Lipschitz with respect to $\|\cdot\|$ and for every $x \in \mathcal{K}$, we have $|f_z(x)| \leq c$. For every distribution \mathcal{D} over Z and $\delta \geq 0$, with probability at least $1 - \delta$ over $S \sim \mathcal{D}^n$,*

$$\text{Rep}(S) \leq 2L \cdot \text{Rad}(\mathcal{K}, n) + c \sqrt{\frac{2 \ln(2/\delta)}{n}}.$$

Proof. For each $z \in Z$, we have $|\ell_{g_z}(x)| \leq 4c$ for all $x \in \mathcal{K}$. The map $S \mapsto \text{Rep}(S)$ satisfies Eq. (6) with $c_0 = \frac{2c}{n}$. By Eq. (5) and (I10, Lemma 26.2), because $\|g_z\|_* \leq L$ for all $z \in Z$,

$$\begin{aligned} \mathbb{E}_{S \sim \mathcal{D}^n} [\text{Rep}(S)] &\leq 2\mathbb{E}_{S \sim \mathcal{D}^n} [\text{Rad}(\ell_S \circ A_S)] \\ &\leq 2L \cdot \text{Rad}(\mathcal{K}, n). \end{aligned}$$

By McDiarmid’s inequality, therefore, with probability at least $1 - \delta$,

$$\text{Rep}(S) \leq 2L \cdot \text{Rad}(\mathcal{K}, n) + \frac{2c}{n} \sqrt{\frac{\ln(2/\delta)n}{2}}.$$

□

6.3 Putting it together

The proof from here is similar to the simpler setup in Theorem 1. Assume that $n \geq n_0$, as in the theorem statement. Next, denote by E the event that there is $x \in \mathcal{K}$ so that

$$(\hat{F}(x) - \hat{F}(x^*) \leq \varepsilon) \wedge (F(x) - F(x^*) > 40\varepsilon).$$

The final result then follows if we prove that

$$\Pr(E) \leq \delta. \quad (7)$$

Towards this, denote by H the event that

$$\text{Rep}(S) \leq 2\varepsilon$$

By Theorem 10,

$$\Pr[\neg H] \leq \frac{\delta}{2}.$$

Next, let N be an $\frac{\varepsilon}{3L}$ -net for \mathcal{K} of minimal size. Denote by \hat{N} the set of points $x \in N$ so that

$$(\hat{F}(x) - \hat{F}(x^*) \leq 5\varepsilon) \wedge (F(x) - F(x^*) > 35\varepsilon).$$

Because F and \hat{F} are Lipschitz, the net property implies that if E holds then \hat{N} is a non-empty set. By Theorem 9, conditioned on H , if $x \in \hat{N}$ then

$$T(x, x^*) \geq 16\varepsilon.$$

On the other hand, by Theorem 8, for every $x \in \mathcal{K}$, conditioned on H , if $x \in \hat{N}$ then

$$\hat{T}(x, x^*) \leq 8\varepsilon.$$

Let E_x be the event that $\hat{T}(x, x^*) \leq 8\varepsilon$. We obtain, then:

$$E \cap H \subseteq \bigcup_{x \in N: T(x, x^*) > 16\varepsilon} E_x.$$

By Theorem 7, if $T(x, x^*) > 16\varepsilon$ then

$$\Pr[E_x] \leq \exp\left(-\frac{\varepsilon n}{3c}\right).$$

By the union bound and Lemma 1,

$$\begin{aligned} \Pr[E] &\leq \Pr[\neg H] + |N| \exp\left(-\frac{\varepsilon n}{3c}\right) \\ &\leq \frac{\delta}{2} + \left(\frac{12L}{\varepsilon}\right)^d \exp\left(-\frac{\varepsilon n}{3c}\right). \end{aligned}$$

The proof of Theorem 2 is now complete because

$$\frac{3c^2 \ln(4/\delta)}{\varepsilon^2} \geq \frac{3c \ln(2/\delta)}{\varepsilon};$$

the last inequality is true for $c \geq \varepsilon$, which we can assume because otherwise the result trivially holds.

References

- K. Jogdeo and S. M. Samuels (1968). Monotone convergence of binomial probabilities and a generalization of Ramanujan’s equation. *The Annals of Mathematical Statistics* **39**(4):1191–1195.
- L. G. Valiant (1984). A theory of the learnable. *Communications of the ACM* **27**(11):1134–1142.
- O. Bousquet and A. Elisseeff (2002). Stability and generalization. *The Journal of Machine Learning Research* **2**:499–526.
- S. Boucheron, G. Lugosi, and P. Massart (2013). *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press.

- M. Anthony and P. L. Bartlett (1999). *Neural network learning: Theoretical foundations*, vol. 9. Cambridge: Cambridge University Press.
- D. P. Bertsekas (1973). Stochastic optimization problems with nondifferentiable cost functionals. *Journal of Optimization Theory and Applications* **12**(2):218–231.
- S. Bubeck (2015). Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning* **8**(3-4):231–357.
- F. Chung and L. Lu (2006). Concentration inequalities and martingale inequalities: a survey. *Internet Mathematics* **3**(1):79–127.
- C. McDiarmid (1998). Concentration. In *Probabilistic methods for algorithmic discrete mathematics*, 195–248. Springer.
- S. Shalev-Shwartz and S. Ben-David (2014). *Understanding machine learning: From theory to algorithms*. Cambridge University Press.
- Kleywegt, Anton J., Shapiro, Alexander, Homem-de-Mello, Tito (2002). The sample average approximation method for stochastic discrete optimization. *SIAM Journal on optimization* **12** (2):479–502.
- Nemirovskij, Arkadij Semenovič, Yudin, David Borisovich (1983). Problem complexity and method efficiency in optimization.
- Shapiro, Alexander (2003). Monte Carlo sampling methods. *Handbooks in operations research and management science* **10** ():353–425.
- A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro (2009). Robust stochastic approximation approach to stochastic programming. *SIAM Journal on optimization* **19**(4), 1574–1609.
- A. Shapiro and A. Nemirovski (2005). On complexity of stochastic programming problems. *Continuous optimization: Current trends and modern applications*, 111–146. Springer.
- V. Feldman, C. Guzman, and S. Vempala (2017). Statistical query algorithms for mean vector estimation and stochastic convex optimization. In *Symposium on Discrete Algorithms*, 1265–1277.
- D. Davis, D. Drusvyatskiy, L. Xiao, and J. Zhang (2021). From low probability to high confidence in stochastic convex optimization. *The Journal of Machine Learning Research* **22**(1), 2237–2274.
- S. Chaturapruek, J. C. Duchi, and C. Ré (2015). Asynchronous stochastic convex optimization: the noise is in the noise and SGD don’t care. In *Advances in Neural Information Processing Systems*, **28**, 28.
- I. Amir, T. Koren, and R. Livni (2021). SGD generalizes better than GD (and regularization doesn’t help). In *Conference on Learning Theory*, 63–92. PMLR.
- V. Vapnik and A. Chervonenkis (2015). On the uniform convergence of relative frequencies of events to their probabilities. In *Measures of complexity: festschrift for Alexey Chervonenkis*, 11–30. Springer.
- A. Blumer, A. Ehrenfeucht, D. Haussler, and M. K. Warmuth (1989). Learnability and the Vapnik-Chervonenkis dimension. *Journal of the ACM* **36**(4), 929–965.
- M. Yan (2012). Extension of convex function. *arXiv:1207.0944*.
- A. Dauber, M. Feder, T. Koren, and R. Livni (2020). Can implicit bias explain generalization? stochastic convex optimization as a case study. In *Advances in Neural Information Processing Systems*, **33**, 7743–7753.
- A. Sekhari, K. Sridharan, and S. Kale (2021). Sgd: The role of implicit regularization, batch-size and multiple-epochs. In *Advances In Neural Information Processing Systems*, **34**, 27422–27433.
- T. Koren, R. Livni, Y. Mansour, and U. Sherman (2022). Benign underfitting of stochastic gradient descent. In *Advances in Neural Information Processing Systems*, **35**, 19605–19617.
- R. Bassily, C. Guzmán, and A. Nandi (2021). Non-euclidean differentially private stochastic convex optimization. In *COLT*, 474–499.
- O. Shamir (2013). On the complexity of bandit and derivative-free stochastic convex optimization. In *Conference on Learning Theory*, 3–24.
- D. J. Foster, A. Sekhari, O. Shamir, N. Srebro, K. Sridharan, and B. Woodworth (2019). The complexity of making the gradient small in stochastic convex optimization. In *Conference on Learning Theory*, 1319–1345.
- I. Amir, Y. Carmon, T. Koren, and R. Livni (2021). Never go full batch (in stochastic convex optimization). In *Advances in Neural Information Processing Systems*, **34**, 25033–25043.
- S. Shalev-Shwartz, O. Shamir, N. Srebro, and K. Sridharan (2009). Stochastic Convex Optimization. In *COLT*.
- V. Feldman (2016). Generalization of erm in stochastic convex optimization: The dimension strikes back. In *Advances in Neural Information Processing Systems*, **29**.
- J.-L. Verger-Gaugry (2005). Covering a ball with smaller equal balls in \mathcal{R}^n . *Discrete and Computational Geometry* **33**, 143–155. Springer.
- R. T. Rockafellar (1997). *Convex analysis*. Princeton University Press, **11**.

- A. Nemirovski and D. Yudin (2004). *Problem complexity and method efficiency in optimization*. John Wiley & Sons.
- Y. Zheng, Y. Cheng, T. Zhang, and R. Jin (2015). A stochastic gradient method with an exponential convergence for finite training sets. In *Advances in Neural Information Processing Systems*, **28**.
- R. I. Bot, P. T. Vuong, and V. N. H. Liem (2015). Subgradient extragradient methods extended to stochastic equilibrium problems. *Vietnam Journal of Mathematics* **43**(2), 413–430.
- N. N. Vakhania, V. I. Tarieladze, and S. A. Chobanyan (1987). *Probability distributions on Banach spaces*. Springer.
- H. J. Kushner and G. G. Yin (2003). *Stochastic approximation algorithms and applications*. Springer.
- S. Ghadimi and G. Lan (2013). Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization* **23**(4), 2341–2368.
- J. C. Spall (2003). *Introduction to stochastic search and optimization*. John Wiley & Sons.
- A. Beck and M. Teboulle (2003). Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters* **31**(3), 167–175.
- Y. Nesterov (2018). *Lectures on convex optimization*. Springer.
- J. C. Duchi, E. Hazan, and Y. Singer (2011). Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research* **12**, 2121–2159.
- B. T. Polyak and Y. O. Shcherbakov (1964). On the convergence of stochastic approximation algorithms. *Soviet Mathematics Doklady* **5**, 1156–1159.
- D. P. Bertsekas (1997). *Nonlinear programming*. Athena Scientific.
- G. Goh, M. Magdon-Ismail, and C. M. Mitchell (2017). Why momentum really works. *arXiv preprint arXiv:1705.08292*.
- G. Goh, M. Magdon-Ismail, and C. M. Mitchell (2017). Momentum optimization: A simple yet effective algorithm for non-convex optimization problems. In *Advances in Neural Information Processing Systems*, **30**, 3365–3373.
- Z. Chen, P. Zhao, and D. B. Dunson (2016). Optimal Bayesian predictive synthesis. *Journal of the American Statistical Association* **111**(514), 202–213.
- H. Kim, J. A. Heller, and M. E. Taddy (2018). Bayesian space-time forecasting using high-dimensional data. *Journal of the American Statistical Association* **113**(522), 665–675.
- A. Venkatraman, T. D. Tran, and M. Dunson (2019). Bayesian nonparametric predictive synthesis. *Journal of the American Statistical Association* **114**(527), 460–472.

Appendix: The subgradient sum theorem

For completeness, we provide a proof of the subgradient sum theorem (see Rockafellar's book ?). We added a couple of figures that illustrate the main idea (Figure 1 and Figure 2). We work in the more general setting when the function is defined on a convex subset of \mathbb{R}^d . For $f : \mathcal{K} \rightarrow \mathbb{R}$ where $\mathcal{K} \subseteq \mathbb{R}^d$ is convex, define

$$\partial_{\mathcal{K}} f(x) = \{g \in \mathbb{R}^d : \forall y \in \mathcal{K} \ f(y) \geq f(x) + \langle g, y - x \rangle\}.$$

Theorem (subgradient sum). *Let $\mathcal{K} \subseteq \mathbb{R}^d$ be a convex set with non-empty interior. If $f_1 : \mathbb{R}^d \rightarrow \mathbb{R}$ and $f_2 : \mathcal{K} \rightarrow \mathbb{R}$ are convex and Lipschitz, then for all $x \in \mathcal{K}$,*

$$\partial_{\mathcal{K}}(f_1 + f_2)(x) = \partial_{\mathbb{R}^d} f_1(x) + \partial_{\mathcal{K}} f_2(x).$$

This generality helps, for example, to deduce the first-order optimality condition we rely on.

Claim (first-order optimality). *Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be Lipschitz and convex. Let \mathcal{K} be convex and closed, and let $x_0 \in \operatorname{argmin}_{\mathcal{K}} f$. Then, there is $g \in \partial_{\mathbb{R}^d} f(x_0)$ so that $\langle g, x - x_0 \rangle \geq 0$ for all $x \in \mathcal{K}$.*

Proof. The point x_0 is also a minimizer of $f + 1$ as a function on \mathcal{K} . Therefore, $0 \in \partial_{\mathcal{K}}(f + 1)(x_0)$. By the subgradient sum theorem, there is $g \in \partial_{\mathbb{R}^d} f(x_0)$ so that $-g \in \partial_{\mathcal{K}} 1(x_0)$. The latter says that for all $x \in \mathcal{K}$, we have $1 \geq 1 + \langle -g, x - x_0 \rangle$. \square

Proof of the subgradient sum theorem. Let $g \in \partial_{\mathcal{K}}(f_1 + f_2)(x)$. Define the convex Lipschitz functions h_1, h_2 :

$$h_1(y) = f_1(y) - f_1(x) - \langle g, y - x \rangle$$

and

$$h_2(y) = f_2(y) - f_2(x).$$

So, x is a minimizer of $h_1 + h_2$ in \mathcal{B} , and $h_1(x) = h_2(x) = 0$. Define the convex sets

$$Z_1 = \{(y_1, \lambda_1) : y_1 \in \mathbb{R}^d, \lambda_1 \geq h_1(y_1)\} \subset \mathbb{R}^d \times \mathbb{R}$$

and

$$Z_2 = \{(y_2, \lambda_2) : y_2 \in \mathcal{K}, \lambda_2 \leq -h_2(y_2)\} \subset \mathbb{R}^d \times \mathbb{R}.$$

The interiors of these convex sets are disjoint because the minimum of $h_1 + h_2$ is zero. It follows that there is a hyperplane separating them. This hyperplane is not "vertical" because the projection of both to \mathbb{R}^d contains \mathcal{K} ; that is, the separating hyperplane is a graph of an affine function on \mathbb{R}^d . The point $(x, 0)$ belongs to both sets, so the hyperplane must contain it. Overall, there is $v \in \mathbb{R}^d$ so that

$$\forall (y_1, \lambda_1) \in Z_1 \quad \langle v, y_1 - x \rangle \leq \lambda_1$$

and

$$\forall (y_2, \lambda_2) \in Z_2 \quad \langle v, y_2 - x \rangle \geq \lambda_2.$$

In other words, for all $y \in \mathbb{R}^d$,

$$\langle v, y - x \rangle \leq h_1(y)$$

and for all $y \in \mathcal{K}$,

$$\langle v, y - x \rangle \geq -h_2(y).$$

Going back to f_1, f_2 , we get that for all $y \in \mathbb{R}^d$,

$$\langle v, y - x \rangle \leq f_1(y) - f_1(x) - \langle g, y - x \rangle$$

and for all $y \in \mathcal{K}$

$$\langle -v, y - x \rangle \leq f_2(y) - f_2(x).$$

The first inequality says that $g + v \in \partial_{\mathbb{R}^d} f_1(x)$ and the second that $-v \in \partial_{\mathcal{K}} f_2(x)$. \square




Figure 1: Part I of the illustration of the proof of the subgradient sum theorem. The function $f_1(x) = (x - 0.1)^2 - \frac{1}{25}$ is the gray line. The function $f_2(x) = |x + 0.1|$ is the dashed line. Their sum $f_1 + f_2$ is the solid black line. The point of interest is $x = -\frac{1}{10}$, which is the minimum of $f_1 + f_2$.




Figure 2: Part II of the illustration of the proof of the subgradient sum theorem. The convex set Z_1 is the epigraph of f_1 . The convex set Z_2 is “minus” the epigraph of f_2 . The point of interest is $x = -\frac{1}{10}$, and $0 \in \partial_{\mathcal{B}}(f_1 + f_2)(x)$. The two convex sets meet at $(x, 0)$. There is a line separating Z_1 and Z_2 going through $(x, 0)$. This line shows that we can write 0 as a subgradient of f_1 plus a subgradient of f_2 .