

Information-theoretic lower bounds on the oracle complexity of stochastic convex optimization

Alekh Agarwal, Peter L. Bartlett, Pradeep Ravikumar, Martin J. Wainwright, *Senior Member, IEEE*.

Abstract—Relative to the large literature on upper bounds on complexity of convex optimization, lesser attention has been paid to the fundamental hardness of these problems. Given the extensive use of convex optimization in machine learning and statistics, gaining an understanding of these complexity-theoretic issues is important. In this paper, we study the complexity of stochastic convex optimization in an oracle model of computation. We introduce a new notion of discrepancy between functions, and use it to reduce problems of stochastic convex optimization to statistical parameter estimation, which can be lower bounded using information-theoretic methods. Using this approach, we improve upon known results and obtain tight minimax complexity estimates for various function classes.

Keywords: Convex optimization, oracle complexity, computational learning theory; Fano’s inequality; Minimax analysis.

1

I. INTRODUCTION

Convex optimization forms the backbone of many algorithms for statistical learning and estimation. Given that many statistical estimation problems are large-scale in nature—with the problem dimension and/or sample size being large—it is essential to make efficient use of computational resources. Stochastic optimization algorithms are an attractive class of methods, known to yield moderately accurate solutions in a relatively short time [1]. Given the popularity of such stochastic optimization methods, understanding the fundamental computational complexity of stochastic convex optimization is thus a key issue for large-scale learning. A large body of literature is devoted to obtaining rates of convergence of specific procedures for various classes of convex optimization problems. A typical outcome of such analysis is an upper bound on the error—for instance, gap to the optimal cost—as a function of the number of iterations. Such analyses have been performed for many standard optimization algorithms, among them gradient descent, mirror descent, interior point programming, and stochastic gradient descent, to name a few.

A. Agarwal is with the Department of EECS at the University of California, Berkeley. P. Ravikumar is with the Department of Computer Science at University of Texas, Austin. M. J. Wainwright is with the Departments of Statistics and EECS, University of California, Berkeley. P. Bartlett is with the Departments of Statistics and EECS, University of California, Berkeley, and the Department of Mathematical Sciences, QUT, Brisbane, Australia.

¹AA and PLB gratefully acknowledge partial support from NSF awards DMS-0707060 and DMS-0830410 and DARPA-HR0011-08-2-0002. AA was also supported in part by a Microsoft Research Fellowship. MJW and PR were partially supported by funding from the National Science Foundation (DMS-0605165, and DMS-0907632). In addition, MJW received funding from the Air Force Office of Scientific Research (AFOSR-09NL184). We also thank the anonymous reviewers for helpful suggestions, and corrections to our results and for pointing out the optimality of our bounds in the primal-dual norm setting. This paper was presented in part at NIPS 2009 conference.

We refer the reader to various standard texts on optimization (e.g., [2], [3], [4]) for further details on such results.

On the other hand, there has been relatively little study of the inherent complexity of convex optimization problems. To the best of our knowledge, the first formal study in this area was undertaken in the seminal work of Nemirovski and Yudin [5], hereafter referred to as NY. One obstacle to a classical complexity-theoretic analysis, as these authors observed, is that of casting convex optimization problems in a Turing Machine model. They avoided this problem by instead considering a natural oracle model of complexity, in which at every round the optimization procedure queries an oracle for certain information on the function being optimized. This information can be either noiseless or noisy, depending on whether the goal is to lower bound the oracle complexity of deterministic or stochastic optimization algorithms. Working within this framework, the authors obtained a series of lower bounds on the computational complexity of convex optimization problems, both in deterministic and stochastic settings. In addition to the original text NY [5], we refer the interested reader to the book by Nesterov [4], and the lecture notes by Nemirovski [6] for further background.

In this paper, we consider the computational complexity of stochastic convex optimization within this oracle model. In particular, we improve upon the work of NY [5] for stochastic convex optimization in two ways. First, our lower bounds have an improved dependence on the dimension of the space. In the context of statistical estimation, these bounds show how the difficulty of the estimation problem increases with the number of parameters. Second, our techniques naturally extend to give sharper results for optimization over simpler function classes. We show that the complexity of optimization for strongly convex losses is smaller than that for convex, Lipschitz losses. Third, we show that for a fixed function class, if the set of optimizers is assumed to have special structure such as sparsity, then the fundamental complexity of optimization can be significantly smaller. All of our proofs exploit a new notion of the discrepancy between two functions that appears to be natural for optimization problems. They involve a reduction from stochastic optimization to a statistical parameter estimation problem, and an application of information-theoretic lower bounds for the estimation problem. We note that special cases of the first two results in this paper appeared in the extended abstract [7], and that a related study was independently undertaken by Raginsky and Rakhlin [8].

The remainder of this paper is organized as follows. We begin in Section II with background on oracle complexity, and a precise formulation of the problems addressed in this

paper. Section III is devoted to the statement of our main results, and discussion of their consequences. In Section IV, we provide the proofs of our main results, which all exploit a common framework of four steps. More technical aspects of these proofs are deferred to the appendices.

a) Notation: For the convenience of the reader, we collect here some notation used throughout the paper. For $p \in [1, \infty]$, we use $\|x\|_p$ to denote the ℓ_p -norm of a vector $x \in \mathbb{R}^p$, and we let q denote the conjugate exponent, satisfying $\frac{1}{p} + \frac{1}{q} = 1$. For two distributions \mathbb{P} and \mathbb{Q} , we use $D(\mathbb{P} \parallel \mathbb{Q})$ to denote the Kullback-Leibler (KL) divergence between the distributions. The notation $\mathbb{I}(A)$ refers to the 0-1 valued indicator random variable of the set A . For two vectors $\alpha, \beta \in \{-1, +1\}^d$, we define the Hamming distance $\Delta_H(\alpha, \beta) := \sum_{i=1}^d \mathbb{I}[\alpha_i \neq \beta_i]$. Given a convex function $f: \mathbb{R}^d \rightarrow \mathbb{R}$, the subdifferential of f at x is the set $\partial f(x)$ given by

$$\{z \in \mathbb{R}^d \mid f(y) \geq f(x) + \langle z, y - x \rangle \text{ for all } y \in \mathbb{R}^d\}.$$

II. BACKGROUND AND PROBLEM FORMULATION

We begin by introducing background on the oracle model of convex optimization, and then turn to a precise specification of the problem to be studied.

A. Convex optimization in the oracle model

Convex optimization is the task of minimizing a convex function f over a convex set $\mathbb{S} \subseteq \mathbb{R}^d$. Assuming that the minimum is achieved, it corresponds to computing an element x_f^* that achieves the minimum—that is, an element $x_f^* \in \arg \min_{x \in \mathbb{S}} f(x)$. An *optimization method* is any procedure that solves this task, typically by repeatedly selecting values from \mathbb{S} . For a given class of optimization problems, our primary focus in this paper is to determine lower bounds on the computational cost, as measured in terms of the number of (noisy) function and subgradient evaluations, required to obtain an ϵ -optimal solution to any optimization problem within the class.

More specifically, we follow the approach of Nemirovski and Yudin [5], and measure computational cost based on the oracle model of optimization. The main components of this model are an *oracle* and an *information set*. An *oracle* is a (possibly random) function $\phi: \mathbb{S} \mapsto \mathcal{I}$ that answers any query $x \in \mathbb{S}$ by returning an element $\phi(x)$ in an information set \mathcal{I} . The information set varies depending on the oracle; for instance, for an exact oracle of m^{th} order, the answer to a query x_t consists of $f(x_t)$ and the first m derivatives of f at x_t . For the case of stochastic oracles studied in this paper, these values are corrupted with zero-mean noise with bounded variance. We then measure the computational labor of any optimization method as the number of queries it poses to the oracle.

In particular, given a positive integer T corresponding to the number of iterations, an optimization method \mathcal{M} designed to approximately minimize the convex function f over the convex set \mathbb{S} proceeds as follows. At any given iteration $t = 1, \dots, T$,

the method \mathcal{M} queries at $x_t \in \mathbb{S}$, and the oracle reveals the information $\phi(x_t, f)$. The method then uses the information $\{\phi(x_1, f), \dots, \phi(x_t, f)\}$ to decide at which point x_{t+1} the next query should be made. For a given oracle function ϕ , let \mathbb{M}_T denote the class of all optimization methods \mathcal{M} that make T queries according to the procedure outlined above. For any method $\mathcal{M} \in \mathbb{M}_T$, we define its error on function f after T steps as

$$\epsilon_T(\mathcal{M}, f, \mathbb{S}, \phi) := f(x_T) - \min_{x \in \mathbb{S}} f(x) = f(x_T) - f(x_f^*), \quad (1)$$

where x_T is the method's query at time T . Note that by definition of x_f^* as a minimizing argument, this error is a non-negative quantity.

When the oracle is stochastic, the method's query x_T at time T is itself random, since it depends on the random answers provided by the oracle. In this case, the optimization error $\epsilon_T(\mathcal{M}, f, \mathbb{S}, \phi)$ is also a random variable. Accordingly, for the case of stochastic oracles, we measure the accuracy in terms of the expected value $\mathbb{E}_\phi[\epsilon_T(\mathcal{M}, f, \mathbb{S}, \phi)]$, where the expectation is taken over the oracle randomness. Given a class of functions \mathcal{F} defined over a convex set \mathbb{S} and a class \mathbb{M}_T of all optimization methods based on T oracle queries, we define the minimax error

$$\epsilon_T^*(\mathcal{F}, \mathbb{S}; \phi) := \inf_{\mathcal{M} \in \mathbb{M}_T} \sup_{f \in \mathcal{F}} \mathbb{E}_\phi[\epsilon_T(\mathcal{M}, f, \mathbb{S}, \phi)]. \quad (2)$$

In the sequel, we provide results for particular classes of oracles. So as to ease the notation, when the oracle ϕ is clear from the context, we simply write $\epsilon_T^*(\mathcal{F}, \mathbb{S})$.

B. Stochastic first-order oracles

In this paper, we study stochastic oracles for which the information set $\mathcal{I} \subset \mathbb{R} \times \mathbb{R}^d$ consists of pairs of noisy function and subgradient evaluations. More precisely, we have:

Definition 1. For a given set \mathbb{S} and function class \mathcal{F} , the class of first-order stochastic oracles consists of random mappings $\phi: \mathbb{S} \times \mathcal{F} \rightarrow \mathcal{I}$ of the form $\phi(x, f) = (f(x), \hat{z}(x))$ such that

$$\begin{aligned} \mathbb{E}[\hat{f}(x)] &= f(x), & \mathbb{E}[\hat{z}(x)] &\in \partial f(x), & \text{and} \\ \mathbb{E}[\|\hat{z}(x)\|_p^2] &\leq \sigma^2. \end{aligned} \quad (3)$$

We use $\mathbb{O}_{p, \sigma}$ to denote the class of all stochastic first-order oracles with parameters (p, σ) . Note that the first two conditions imply that $\hat{f}(x)$ is an unbiased estimate of the function value $f(x)$, and that $\hat{z}(x)$ is an unbiased estimate of a subgradient $z \in \partial f(x)$. When f is actually differentiable, then $\hat{z}(x)$ is an unbiased estimate of the gradient $\nabla f(x)$. The third condition in equation (3) controls the “noisiness” of the subgradient estimates in terms of the ℓ_p -norm.

Stochastic gradient methods are a widely used class of algorithms that can be understood as operating based on information provided by a stochastic first-order oracle. As a particular example, consider a function of the separable form $f(x) = \frac{1}{n} \sum_{i=1}^n h_i(x)$, where each h_i is differentiable. Functions of this form arise very frequently in statistical problems, where each term i corresponds to a different sample

and the overall cost function is some type of statistical loss (e.g., maximum likelihood, support vector machines, boosting etc.) The natural stochastic gradient method for this problem is to choose an index $i \in \{1, 2, \dots, n\}$ uniformly at random, and then to return the pair $(h_i(x), \nabla h_i(x))$. Taking averages over the randomly chosen index i yields $\frac{1}{n} \sum_{i=1}^n h_i(x) = f(x)$, so that $h_i(x)$ is an unbiased estimate of $f(x)$, with an analogous unbiased property holding for the gradient of $h_i(x)$.

C. Function classes of interest

We now turn to the classes \mathcal{F} of convex functions for which we study oracle complexity. In all cases, we consider real-valued convex functions defined over some convex set \mathbb{S} . We assume without loss of generality that \mathbb{S} contains an open set around 0, and many of our lower bounds involve the maximum radius $r = r(\mathbb{S}) > 0$ such that

$$\mathbb{S} \supseteq \mathbb{B}_\infty(r) := \{x \in \mathbb{R}^d \mid \|x\|_\infty \leq r\}. \quad (4)$$

Our first class consists of *convex Lipschitz functions*:

Definition 2. For a given convex set $\mathbb{S} \subseteq \mathbb{R}^d$ and parameter $p \in [1, \infty]$, the class $\mathcal{F}_{\text{cv}}(\mathbb{S}, L, p)$ consists of all convex functions $f : \mathbb{S} \rightarrow \mathbb{R}$ such that

$$|f(x) - f(y)| \leq L \|x - y\|_q \quad \text{for all } x, y \in \mathbb{S}, \quad (5)$$

where $\frac{1}{q} = 1 - \frac{1}{p}$.

We have defined the Lipschitz condition (5) in terms of the conjugate exponent $q \in [1, \infty]$, defined by the relation $\frac{1}{q} = 1 - \frac{1}{p}$. To be clear, our motivation in doing so is to maintain consistency with our definition of the stochastic first-order oracle, in which we assumed that $\mathbb{E}[\|\hat{z}(x)\|_p^2] \leq \sigma^2$. We note that the Lipschitz condition (5) is equivalent to the condition

$$\|z\|_p \leq L \quad \forall z \in \partial f(x), \quad \text{and for all } x \in \text{int}(\mathbb{S}).$$

If we consider the case of a differentiable function f , the unbiasedness condition in Definition 1 implies that

$$\begin{aligned} \|\nabla f(x)\|_p &= \|\mathbb{E}[\hat{z}(x)]\|_p \stackrel{(a)}{\leq} \mathbb{E}\|\hat{z}(x)\|_p \\ &\stackrel{(b)}{\leq} \sqrt{\mathbb{E}\|\hat{z}(x)\|_p^2} \leq \sigma, \end{aligned}$$

where inequality (a) follows from the convexity of the ℓ_p -norm and Jensen's inequality, and inequality (b) is a result of Jensen's inequality applied to the concave function \sqrt{x} . This bound implies that f must be Lipschitz with constant at most σ with respect to the dual ℓ_q -norm. Therefore, we necessarily must have $L \leq \sigma$, in order for the function class from Definition 2 to be consistent with the stochastic first-order oracle.

A second function class consists of strongly convex functions, defined as follows:

Definition 3. For a given convex set $\mathbb{S} \subseteq \mathbb{R}^d$ and parameter $p \in [1, \infty]$, the class $\mathcal{F}_{\text{scv}}(\mathbb{S}, p; L, \gamma)$ consists of all convex functions $f : \mathbb{S} \rightarrow \mathbb{R}$ such that the Lipschitz condition (5)

holds, and such that f satisfies the ℓ_2 -strong convexity condition

$$\begin{aligned} f(\alpha x + (1 - \alpha)y) &\geq \\ \alpha f(x) + (1 - \alpha)f(y) + \alpha(1 - \alpha) \frac{\gamma^2}{2} \|x - y\|_2^2 &\quad (6) \end{aligned}$$

for all $x, y \in \mathbb{S}$.

In this paper, we restrict our attention to the case of strong convexity with respect to the ℓ_2 -norm. (Similar results on the oracle complexity for strong convexity with respect to different norms can be obtained by straightforward modifications of the arguments given here). For future reference, it should be noted that the Lipschitz constant L and strong convexity constant γ interact with one another. In particular, whenever $\mathbb{S} \subset \mathbb{R}^d$ contains the ℓ_∞ -ball of radius r , the Lipschitz L and strong convexity γ constants must satisfy the inequality

$$\frac{L}{\gamma^2} \geq \frac{r}{4} d^{1/p}. \quad (7)$$

In order to establish this inequality, we note that strong convexity condition with $\alpha = 1/2$ implies that

$$\frac{\gamma^2}{8} \leq \frac{2f\left(\frac{x+y}{2}\right) - f(x) - f(y)}{2\|x - y\|_2^2} \leq \frac{L\|x - y\|_q}{2\|x - y\|_2^2}$$

We now choose the pair $x, y \in \mathbb{S}$ such that $\|x - y\|_\infty = r$ and $\|x - y\|_2 = r\sqrt{d}$. Such a choice is possible whenever \mathbb{S} contains the ℓ_∞ ball of radius r . Since we have $\|x - y\|_q \leq d^{1/q} \|x - y\|_\infty$, this choice yields $\frac{\gamma^2}{4} \leq \frac{Ld^{\frac{1}{q}-1}}{r}$, which establishes the claim (7).

As a third example, we study the oracle complexity of optimization over the class of convex functions that have sparse minimizers. This class of functions is well-motivated, since a large body of statistical work has studied the estimation of vectors, matrices and functions under various types of sparsity constraints. A common theme in this line of work is that the ambient dimension d enters the rates only logarithmically, and so has a mild effect. Consequently, it is natural to investigate whether the complexity of optimization methods also enjoys such a mild dependence on ambient dimension under sparsity assumptions.

For a vector $x \in \mathbb{R}^d$, we use $\|x\|_0$ to denote the number of non-zero elements in x . Recalling the set $\mathcal{F}_{\text{cv}}(\mathbb{S}, L, p)$ from Definition 2, we now define a class of Lipschitz functions with sparse minimizers.

Definition 4. For a convex set $\mathbb{S} \subset \mathbb{R}^d$ and positive integer $k \leq \lfloor d/2 \rfloor$, let $\mathcal{F}_{\text{sp}}(k; \mathbb{S}, L)$ be the set of be the class of all convex functions that are L -Lipschitz in the ℓ_∞ -norm, and have at least one k -sparse optimizer, meaning that there exists some

$$x^* \in \arg \min_{x \in \mathbb{S}} f(x) \quad \text{satisfying } \|x^*\|_0 \leq k, \quad (8)$$

We frequently use the shorthand notation $\mathcal{F}_{\text{sp}}(k)$ when the set \mathbb{S} and parameter L are clear from context.

III. MAIN RESULTS AND THEIR CONSEQUENCES

With the setup of stochastic convex optimization in place, we are now in a position to state the main results of this paper, and to discuss some of their consequences. As previously mentioned, a subset of our results assume that the set \mathbb{S} contains an ℓ_∞ ball of radius $r = r(\mathbb{S})$. Our bounds scale with r , thereby reflecting the natural dependence on the size of the set \mathbb{S} . Also, we set the oracle second moment bound σ to be the same as the Lipschitz constant L in our results.

A. Oracle complexity for convex Lipschitz functions

We begin by analyzing the minimax oracle complexity of optimization for the class of bounded and convex Lipschitz functions \mathcal{F}_{cv} from Definition 2.

Theorem 1. *Let $\mathbb{S} \subset \mathbb{R}^d$ be a convex set such that $\mathbb{S} \supseteq \mathbb{B}_\infty(r)$ for some $r > 0$. Then there exists a universal constant $c_0 > 0$ such that the minimax oracle complexity over the class $\mathcal{F}_{\text{cv}}(\mathbb{S}, L, p)$ satisfies the following lower bounds:*

(a) For $1 \leq p \leq 2$,

$$\sup_{\phi \in \mathcal{O}_{p,L}} \epsilon_T^*(\mathcal{F}_{\text{cv}}, \mathbb{S}; \phi) \geq \min \left\{ c_0 L r \sqrt{\frac{d}{T}}, \frac{Lr}{144} \right\}. \quad (9)$$

(b) For $p > 2$,

$$\sup_{\phi \in \mathcal{O}_{p,L}} \epsilon_T^*(\mathcal{F}_{\text{cv}}, \mathbb{S}; \phi) \geq \min \left\{ c_0 L r \frac{d^{1-\frac{1}{p}}}{\sqrt{T}}, \frac{Ld^{1-1/p}r}{72} \right\}. \quad (10)$$

Remarks: Nemirovski and Yudin [5] proved the lower bound $\Omega(\frac{1}{\sqrt{T}})$ for the function class \mathcal{F}_{cv} , in the special case that \mathbb{S} is the unit ball of a given norm, and the functions are Lipschitz in the corresponding *dual norm*. For $p \geq 2$, they established the minimax optimality of this dimension-independent result by appealing to a matching upper bound achieved by the method of mirror descent. In contrast, here we do not require the two norms—namely, that constraining the set \mathbb{S} and that for the Lipschitz constraint—to be dual to one other; instead, we give lower bounds in terms of the largest ℓ_∞ ball contained within the constraint set \mathbb{S} . As discussed below, our bounds do include the results for the dual setting of past work as a special case, but more generally, by examining the relative geometry of an arbitrary set with respect to the ℓ_∞ ball, we obtain results for arbitrary sets. (We note that the ℓ_∞ constraint is natural in many optimization problems arising in machine learning settings, in which upper and lower bounds on variables are often imposed.) Thus, in contrast to the past work of NY on stochastic optimization, our analysis gives sharper dimension dependence under more general settings. It also highlights the role of the geometry of the set \mathbb{S} in determining the oracle complexity.

In general, our lower bounds cannot be improved, and hence specify the optimal minimax oracle complexity. We consider here some examples to illustrate their sharpness. Throughout we assume that T is large enough to ensure that the $1/\sqrt{T}$ term attains the lower bound and not the $L/144$ term. (This condition is reasonable given our goal of understanding the rate as T increases, as opposed to the

transient behavior over the first few iterations.)

(a) We start from the special case that has been primarily considered in past works. We consider the class $\mathcal{F}_{\text{cv}}(\mathbb{B}_q(1), L, p)$ with $q = 1 - 1/p$ and the stochastic first-order oracles $\mathcal{O}_{p,L}$ for this class. Then the radius r of the largest ℓ_∞ ball inscribed within the $\mathbb{B}_q(1)$ scales as $r = d^{-1/q}$. By inspection of the lower bounds (9) and (10), we see that

$$\sup_{\phi \in \mathcal{O}_{p,L}} \epsilon_T^*(\mathcal{F}_{\text{cv}}, \mathbb{B}_q(1); \phi) = \begin{cases} \Omega\left(L \frac{d^{1/2-1/q}}{\sqrt{T}}\right) & \text{for } 1 \leq p \leq 2 \\ \Omega\left(\frac{L}{\sqrt{T}}\right) & \text{for } p \geq 2. \end{cases} \quad (11)$$

As mentioned previously, the dimension-independent lower bound for the case $p \geq 2$ was demonstrated in Chapter 5 of NY, and shown to be optimal² since it is achieved using mirror descent with the prox-function $\|\cdot\|_q^2$. For the case of $1 \leq p < 2$, the lower bounds are also unimprovable, since they are again achieved (up to constant factors) by stochastic gradient descent. See Appendix C for further details on these matching upper bounds.

(b) Let us now consider how our bounds can also make sharp predictions for non-dual geometries, using the special case $\mathbb{S} = \mathbb{B}_\infty(1)$. For this choice, we have $r(\mathbb{S}) = 1$, and hence Theorem 1 implies that for all $p \in [1, 2]$, the minimax oracle complexity is lower bounded as

$$\sup_{\phi \in \mathcal{O}_{p,L}} \epsilon_T^*(\mathcal{F}_{\text{cv}}, \mathbb{B}_\infty(1); \phi) = \Omega\left(L \sqrt{\frac{d}{T}}\right).$$

Up to constant factors, this lower bound is sharp for all $p \in [1, 2]$. Indeed, for any convex set \mathbb{S} , stochastic gradient descent achieves a matching upper bound (see Section 5.2.4, p. 196 of NY [5], as well as Appendix C in this paper for further discussion).

(c) As another example, suppose that $\mathbb{S} = \mathbb{B}_2(1)$. Observe that this ℓ_2 -norm unit ball satisfies the relation $\mathbb{B}_2(1) \supseteq \frac{1}{\sqrt{d}}\mathbb{B}_\infty(1)$, so that we have $r(\mathbb{B}_2(1)) = 1/\sqrt{d}$. Consequently, for this choice, the lower bound (9) takes the form

$$\sup_{\phi \in \mathcal{O}_{p,L}} \epsilon_T^*(\mathcal{F}_{\text{cv}}, \mathbb{B}_2(1); \phi) = \Omega\left(L \frac{1}{\sqrt{T}}\right),$$

which is a dimension-independent lower bound. This lower bound for $\mathbb{B}_2(1)$ is indeed tight for $p \in [1, 2]$, and as before, this rate is achieved by stochastic gradient descent [5].

(d) Turning to the case of $p > 2$, when $\mathbb{S} = \mathbb{B}_\infty(1)$, the lower bound (10) can be achieved (up to constant factors) using mirror descent with the dual norm $\|\cdot\|_q^2$; for further discussion, we again refer the reader to Section 5.2.1, p. 190 of NY [5], as well as to Appendix C of this paper. Also, even though this lower bound requires the oracle to have only bounded variance, our proof actually uses a stochastic oracle based on Bernoulli random variables, for which all moments exist. Consequently, at least in general, our results show that there

²There is an additional logarithmic factor in the upper bounds for $p = \Omega(\log d)$.

is no hope of achieving faster rates by restricting to oracles with bounds on higher-order moments. This is an interesting contrast to the case of having *less* than two moments, in which the rates are slower. For instance, as shown in Section 5.3.1 of NY [5], suppose that the gradient estimates in a stochastic oracle satisfy the moment bound $\mathbb{E}\|\hat{z}(x)\|_p^b \leq \sigma^2$ for some $b \in [1, 2)$. In this setting, the oracle complexity is lower bounded by $\Omega(T^{-(b-1)/b})$. Since $T^{\frac{b-1}{b}} \ll T^{\frac{1}{2}}$ for all $b \in [1, 2)$, there is a significant penalty in convergence rates for having less than two bounded moments.

(e) Even though the results have been stated in a first-order stochastic oracle model, they actually hold in a stronger sense. Let $\nabla^i f(x)$ denote the i th-order derivative of f evaluated at x , when it exists. With this notation, our results apply to an oracle that responds with a random function \hat{f}_t such that

$$\mathbb{E}[\hat{f}_t(x)] = \mathbb{E}[f(x)], \quad \text{and} \quad \mathbb{E}[\nabla^i \hat{f}_t(x)] = \nabla^i f(x)$$

for all $x \in \mathbb{S}$ and i such that $\nabla^i f(x)$ exists, along with appropriately bounded second moments of all the derivatives. Consequently, higher-order gradient information cannot improve convergence rates in a worst-case setting. Indeed, the result continues to hold even for the significantly stronger oracle that responds with a random function that is a noisy realization of the true function. In this sense, our result is close in spirit to a statistical sample complexity lower bound. Our proof technique is based on constructing a ‘‘packing set’’ of functions, and thus has some similarity to techniques used in statistical minimax analysis (e.g., [9], [10], [11], [12]) and learning theory (e.g., [13], [14], [15]). A significant difference, as will be shown shortly, is that the metric of interest for optimization is very different than those typically studied in statistical minimax theory.

B. Oracle complexity for strongly convex Lipschitz functions

We now turn to the statement of lower bounds over the class of Lipschitz and strongly convex functions \mathcal{F}_{scv} from Definition 3. In all these statements, we assume that $\gamma^2 \leq \frac{4Ld^{-1/p}}{r}$, as is required for the definition of \mathcal{F}_{scv} to be sensible.

Theorem 2. *Let $\mathbb{S} = \mathbb{B}_\infty(r)$. Then there exist universal constants $c_1, c_2 > 0$ such that the minimax oracle complexity over the class $\mathcal{F}_{\text{scv}}(\mathbb{S}; p; L, \gamma)$ satisfies the following lower bounds:*

(a) *For $p = 1$, the oracle complexity $\sup_{\phi \in \mathcal{O}_{p,L}} \epsilon^*(\mathcal{F}_{\text{scv}}, \phi)$ is lower bounded by*

$$\min \left\{ c_1 \frac{L^2}{\gamma^2 T}, c_2 Lr \sqrt{\frac{d}{T}}, \frac{L^2}{1152\gamma^2 d}, \frac{Lr}{144} \right\}. \quad (12)$$

(b) *For $p > 2$, the oracle complexity $\sup_{\phi \in \mathcal{O}_{p,L}} \epsilon^*(\mathcal{F}_{\text{scv}}, \phi)$ is lower bounded by*

$$\min \left\{ c_1 \frac{L^2 d^{1-2/p}}{\gamma^2 T}, c_2 \frac{Lr d^{1-1/p}}{\sqrt{T}}, \frac{L^2 d^{1-2/p}}{1152\gamma^2}, \frac{Lr d^{1-1/p}}{144} \right\}. \quad (13)$$

As with Theorem 1, these lower bounds are sharp. In particular, for $S = \mathbb{B}_\infty(1)$, stochastic gradient descent achieves the rate (12) up to logarithmic factors [16], and closely related

algorithms proposed in very recent works [17], [18] match the lower bound exactly up to constant factors. It should be noted Theorem 2 exhibits an interesting phase transition between two regimes. On one hand, suppose that the strong convexity parameter γ^2 is large: then as long as T is sufficiently large, the first term $\Omega(1/T)$ determines the minimax rate, which corresponds to the fast rate possible under strong convexity. In contrast, if we consider a poorly conditioned objective with $\gamma \approx 0$, then the term involving $\Omega(1/\sqrt{T})$ is dominant, corresponding to the rate for a convex objective. This behavior is natural, since Theorem 2 recovers (as a special case) the convex result with $\gamma = 0$. However, it should be noted that Theorem 2 applies only to the set $\mathbb{B}_\infty(r)$, and not to arbitrary sets \mathbb{S} like Theorem 1. Consequently, the generalization of Theorem 2 to arbitrary convex, compact sets remains an interesting open question.

C. Oracle complexity for convex Lipschitz functions with sparse optima

Finally, we turn to the oracle complexity of optimization over the class \mathcal{F}_{sp} from Definition 4.

Theorem 3. *Let \mathcal{F}_{sp} be the class of all convex functions that are L -Lipschitz with respect to the $\|\cdot\|_\infty$ norm and that have a k -sparse optimizer. Let $\mathbb{S} \subset \mathbb{R}^d$ be a convex set with $\mathbb{B}_\infty(r) \subseteq \mathbb{S}$. Then there exists a universal constant $c > 0$ such that for all $k \leq \lfloor \frac{d}{2} \rfloor$, we have*

$$\sup_{\phi \in \mathcal{O}_{\infty,L}} \epsilon^*(\mathcal{F}_{\text{sp}}, \phi) \geq \min \left\{ cLr \sqrt{\frac{k^2 \log \frac{d}{k}}{T}}, \frac{Lkr}{432} \right\}. \quad (14)$$

b) Remark: If $k = \mathcal{O}(d^{1-\delta})$ for some $\delta \in (0, 1)$ (so that $\log \frac{d}{k} = \Theta(\log d)$), then this bound is sharp up to constant factors. In particular, suppose that we use mirror descent based on the $\|\cdot\|_{1+\varepsilon}$ norm with $\varepsilon = 2 \log d / (2 \log d - 1)$. As we discuss in more detail in Appendix C, it can be shown that this technique will achieve a solution accurate to $\mathcal{O}(\sqrt{\frac{k^2 \log d}{T}})$ within T iterations; this achievable result matches our lower bound (14) up to constant factors under the assumed scaling $k = \mathcal{O}(d^{1-\delta})$. To the best of our knowledge, Theorem 3 provides the first tight lower bound on the oracle complexity of sparse optimization.

IV. PROOFS OF RESULTS

We now turn to the proofs of our main results. We begin in Section IV-A by outlining the framework and establishing some basic results on which our proofs are based. Sections IV-B through IV-D are devoted to the proofs of Theorems 1 through 3 respectively.

A. Framework and basic results

We begin by establishing a basic set of results that are exploited in the proofs of the main results. At a high-level, our main idea is to show that the problem of convex optimization is at least as hard as estimating the parameters of Bernoulli variables—that is, the biases of d independent coins. In order

to perform this embedding, for a given error tolerance ϵ , we start with an appropriately chosen subset of the vertices of a d -dimensional hypercube, each of which corresponds to some values of the d Bernoulli parameters. For a given function class, we then construct a ‘‘difficult’’ subclass of functions that are indexed by these vertices of the hypercube. We then show that being able to optimize any function in this subclass to ϵ -accuracy requires identifying the hypercube vertex. This is a multiway hypothesis test based on the observations provided by T queries to the stochastic oracle, and we apply Fano’s inequality [19] or Le Cam’s bound [20], [12] to lower bound the probability of error. In the remainder of this section, we provide more detail on each of steps involved in this embedding.

1) *Constructing a difficult subclass of functions:* Our first step is to construct a subclass of functions $\mathcal{G} \subseteq \mathcal{F}$ that we use to derive lower bounds. Any such subclass is parametrized by a subset $\mathcal{V} \subseteq \{-1, +1\}^d$ of the hypercube, chosen as follows. Recalling that Δ_H denotes the Hamming metric, we let $\mathcal{V} = \{\alpha^1, \dots, \alpha^M\}$ be a subset of the vertices of the hypercube such that

$$\Delta_H(\alpha^j, \alpha^k) \geq \frac{d}{4} \quad \text{for all } j \neq k, \quad (15)$$

meaning that \mathcal{V} is a $\frac{d}{4}$ -packing in the Hamming norm. It is a classical fact (e.g., [21]) that one can construct such a set with cardinality $|\mathcal{V}| \geq (2/\sqrt{e})^{d/2}$.

Now let $\mathcal{G}_{\text{base}} = \{f_i^+, f_i^-, i = 1, \dots, d\}$ denote some base set of $2d$ functions defined on the convex set \mathbb{S} , to be chosen appropriately depending on the problem at hand. For a given tolerance $\delta \in (0, \frac{1}{4}]$, we define, for each vertex $\alpha \in \mathcal{V}$, the function $x \mapsto g_\alpha(x)$ given by

$$\frac{c}{d} \sum_{i=1}^d \left\{ (1/2 + \alpha_i \delta) f_i^+(x) + (1/2 - \alpha_i \delta) f_i^-(x) \right\}. \quad (16)$$

Depending on the result to be proven, our choice of the base functions $\{f_i^+, f_i^-\}$ and the pre-factor c will ensure that each g_α satisfies the appropriate Lipschitz and/or strong convexity properties over \mathbb{S} . Moreover, we will ensure that that all minimizers x_α of each g_α are contained within \mathbb{S} .

Based on these functions and the packing set \mathcal{V} , we define the function class

$$\mathcal{G}(\delta) := \{g_\alpha, \alpha \in \mathcal{V}\}. \quad (17)$$

Note that $\mathcal{G}(\delta)$ contains a total of $|\mathcal{V}|$ functions by construction, and as mentioned previously, our choices of the base functions etc. will ensure that $\mathcal{G}(\delta) \subseteq \mathcal{F}$. We demonstrate specific choices of the class $\mathcal{G}(\delta)$ in the proofs of Theorems 1 through 3 to follow.

2) *Optimizing well is equivalent to function identification:* We now claim that if a method can optimize over the subclass $\mathcal{G}(\delta)$ up to a certain tolerance, then it must be capable of identifying which function $g_\alpha \in \mathcal{G}(\delta)$ was chosen. We first require a measure for the *closeness* of functions in terms of their behavior near each others’ minima. Recall that we use $x_f^* \in \mathbb{R}^d$ to denote a minimizing point of the function f . Given

a convex set $S \subseteq \mathbb{R}^d$ and two functions f, g , we define

$$\rho(f, g) := \inf_{x \in \mathbb{S}} [f(x) + g(x) - f(x_f^*) - g(x_g^*)]. \quad (18)$$

This discrepancy measure is non-negative, symmetric in its

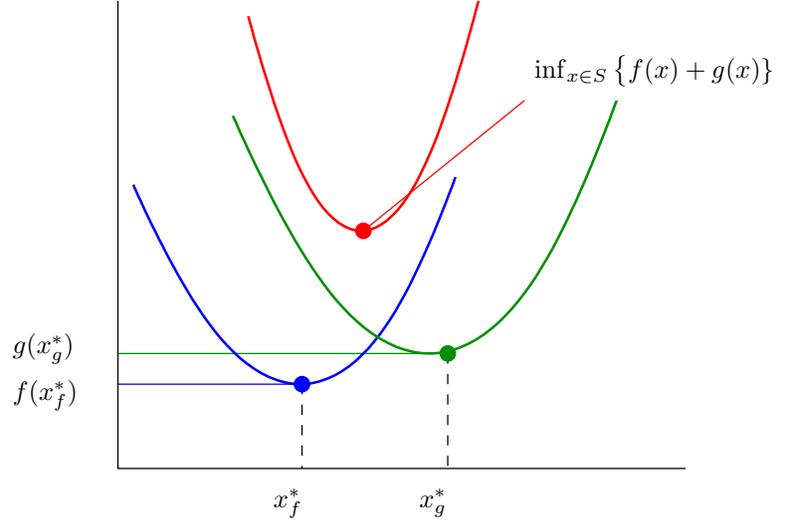


Fig. 1. Illustration of the discrepancy function $\rho(f, g)$. The functions f and g achieve their minimum values $f(x_f^*)$ and $g(x_g^*)$ at the points x_f^* and x_g^* respectively.

arguments, and satisfies $\rho(f, g) = 0$ if and only if $x_f^* = x_g^*$, so that we may refer to it as a premetric. (It does not satisfy the triangle inequality nor the condition that $\rho(f, g) = 0$ if and only if $f = g$, both of which are required for ρ to be a metric.)

Given the subclass $\mathcal{G}(\delta)$, we quantify how densely it is packed with respect to the premetric ρ using the quantity

$$\psi(\mathcal{G}(\delta)) := \min_{\alpha \neq \beta \in \mathcal{V}} \rho(g_\alpha, g_\beta). \quad (19)$$

We denote this quantity by $\psi(\delta)$ when the class \mathcal{G} is clear from the context. We now state a simple result that demonstrates the utility of maintaining a separation under ρ among functions in $\mathcal{G}(\delta)$.

Lemma 1. *For any $\tilde{x} \in \mathbb{S}$, there can be at most one function $g_\alpha \in \mathcal{G}(\delta)$ such that*

$$g_\alpha(\tilde{x}) - \inf_{x \in \mathbb{S}} g_\alpha(x) \leq \frac{\psi(\delta)}{3}. \quad (20)$$

Thus, if we have an element $\tilde{x} \in \mathbb{S}$ that approximately minimizes one function in the set $\mathcal{G}(\delta)$ up to tolerance $\psi(\delta)$, then it cannot approximately minimize any other function in the set.

Proof: For a given $\tilde{x} \in \mathbb{S}$, suppose that there exists an $\alpha \in \mathcal{V}$ such that $g_\alpha(\tilde{x}) - g_\alpha(x_\alpha^*) \leq \frac{\psi(\delta)}{3}$. From the definition of $\psi(\delta)$ in (19), for any $\beta \in \mathcal{V}$, $\beta \neq \alpha$, we have

$$\begin{aligned} \psi(\delta) &\leq g_\alpha(\tilde{x}) - \inf_{x \in \mathbb{S}} g_\alpha(x) + g_\beta(\tilde{x}) - \inf_{x \in \mathbb{S}} g_\beta(x) \\ &\leq \frac{\psi(\delta)}{3} + g_\beta(\tilde{x}) - \inf_{x \in \mathbb{S}} g_\beta(x). \end{aligned}$$

Re-arranging yields the inequality $g_\beta(\tilde{x}) - g_\beta(x_\beta^*) \geq \frac{2}{3} \psi(\delta)$, from which the claim (20) follows.

Suppose that for some fixed but unknown function $g_{\alpha^*} \in \mathcal{G}(\delta)$, some method \mathcal{M}_T is allowed to make T queries to an oracle with information function $\phi(\cdot; g_{\alpha^*})$, thereby obtaining the information sequence

$$\phi(x_1^T; g_{\alpha^*}) := \{\phi(x_t; g_{\alpha^*}), t = 1, 2, \dots, T\}.$$

Our next lemma shows that if the method \mathcal{M}_T achieves a low minimax error over the class $\mathcal{G}(\delta)$, then one can use its output to construct a hypothesis test that returns the true parameter α^* at least $2/3$ of the time. (In this statement, we recall the definition (2) of the minimax error in optimization.)

Lemma 2. *Suppose that based on the data $\phi(x_1^T; g_{\alpha^*})$, there exists a method \mathcal{M}_T that achieves a minimax error satisfying*

$$\mathbb{E}[\epsilon_T(\mathcal{M}_T, \mathcal{G}(\delta), \mathbb{S}, \phi)] \leq \frac{\psi(\delta)}{9}. \quad (21)$$

Based on such a method \mathcal{M}_T , one can construct a hypothesis test $\hat{\alpha} : \phi(x_1^T; g_{\alpha^}) \rightarrow \mathcal{V}$ such that $\max_{\alpha^* \in \mathcal{V}} \mathbb{P}_{\phi}[\hat{\alpha} \neq \alpha^*] \leq \frac{1}{3}$.*

Proof: Given a method \mathcal{M}_T that satisfies the bound (21), we construct an estimator $\hat{\alpha}(\mathcal{M}_T)$ of the true vertex α^* as follows. If there exists some $\alpha \in \mathcal{V}$ such that $g_{\alpha}(x_T) - g_{\alpha}(x_{\alpha}) \leq \frac{\psi(\delta)}{3}$ then we set $\hat{\alpha}(\mathcal{M}_T)$ equal to α . If no such α exists, then we choose $\hat{\alpha}(\mathcal{M}_T)$ uniformly at random from \mathcal{V} . From Lemma 1, there can exist only one such $\alpha \in \mathcal{V}$ that satisfies this inequality. Consequently, using Markov's inequality, we have $\mathbb{P}_{\phi}[\hat{\alpha}(\mathcal{M}_T) \neq \alpha^*] \leq \mathbb{P}_{\phi}[\epsilon_T(\mathcal{M}_T, g_{\alpha^*}, \mathbb{S}, \phi) \geq \psi(\delta)/3] \leq \frac{1}{3}$. Maximizing over α^* completes the proof. \blacksquare We have thus shown that having a low minimax optimization error over $\mathcal{G}(\delta)$ implies that the vertex $\alpha^* \in \mathcal{V}$ can be identified most of the time.

3) *Oracle answers and coin tosses:* We now describe stochastic first order oracles ϕ for which the samples $\phi(x_1^T; g_{\alpha})$ can be related to coin tosses. In particular, we associate a coin with each dimension $i \in \{1, 2, \dots, d\}$, and consider the set of coin bias vectors lying in the set

$$\Theta(\delta) = \{(1/2 + \alpha_1\delta, \dots, 1/2 + \alpha_d\delta) \mid \alpha \in \mathcal{V}\}, \quad (22)$$

Given a particular function $g_{\alpha} \in \mathcal{G}(\delta)$ —or equivalently, vertex $\alpha \in \mathcal{V}$ —we consider two different types of stochastic first-order oracles ϕ , defined as follows:

By construction, the function value and gradients returned by Oracle A are unbiased estimates of those of g_{α} . In particular, since each co-ordinate i is chosen with probability $1/d$, the expectation $\mathbb{E}[\hat{g}_{\alpha, A}(x)]$ is given by

$$\frac{c}{d} \sum_{i=1}^d [\mathbb{E}[b_i] f_i^+(x) + \mathbb{E}[1 - b_i] f_i^-(x)] = g_{\alpha}(x),$$

with a similar relation for the gradient. Furthermore, as long as the base functions f_i^+ and f_i^- have gradients bounded by 1, we have $\mathbb{E}[\|\hat{z}_{\alpha, A}(x)\|_p] \leq c$ for all $p \in [1, \infty]$.

Parts of proofs are based on an oracle which responds with function values and gradients that are d -dimensional in nature.

As with Oracle A, this oracle returns unbiased estimates of the function values and gradients. We frequently work with functions f_i^+, f_i^- that depend only on the i^{th} coordinate $x(i)$. In such cases, under the assumptions $|\frac{\partial f_i^+}{\partial x(i)}| \leq 1$ and $|\frac{\partial f_i^-}{\partial x(i)}| \leq 1$, we have

$$\begin{aligned} \|\hat{z}_{\alpha, B}(x)\|_p^2 &= \frac{c^2}{d^2} \left(\sum_{i=1}^d \left| b_i \frac{\partial f_i^+(x)}{\partial x(i)} + (1 - b_i) \frac{\partial f_i^-(x)}{\partial x(i)} \right|^p \right)^{2/p} \\ &\leq c^2 d^{2/p-2}. \end{aligned} \quad (23)$$

In our later uses of Oracles A and B, we choose the pre-factor c appropriately so as to produce the desired Lipschitz constants.

4) *Lower bounds on coin-tossing:* Finally, we use information-theoretic methods to lower bound the probability of correctly estimating the true parameter $\alpha^* \in \mathcal{V}$ in our model. At each round of either Oracle A or Oracle B, we can consider a set of d coin tosses, with an associated vector $\theta^* = (\frac{1}{2} + \alpha_1^*\delta, \dots, \frac{1}{2} + \alpha_d^*\delta)$ of parameters. At any round, the output of Oracle A can (at most) reveal the instantiation $b_i \in \{0, 1\}$ of a randomly chosen index, whereas Oracle B can at most reveal the entire vector (b_1, b_2, \dots, b_d) . Our goal is to lower bound the probability of estimating the true parameter α^* , based on a sequence of length T . As noted previously in remarks following Theorem 1, this part of our proof exploits classical techniques from statistical minimax theory, including the use of Fano's inequality (e.g., [9], [10], [11], [12]) and Le Cam's bound (e.g., [20], [12]).

Lemma 3. *Suppose that the Bernoulli parameter vector α^* is chosen uniformly at random from the packing set \mathcal{V} , and suppose that the outcome of $\ell \leq d$ coins chosen uniformly at random is revealed at each round $t = 1, \dots, T$. Then for any $\delta \in (0, 1/4]$, any hypothesis test $\hat{\alpha}$ satisfies*

$$\mathbb{P}[\hat{\alpha} \neq \alpha^*] \geq 1 - \frac{16\ell T \delta^2 + \log 2}{\frac{d}{2} \log(2/\sqrt{e})}, \quad (24)$$

where the probability is taken over both randomness in the oracle and the choice of α^* .

Note that we will apply the lower bound (24) with $\ell = 1$ in the case of Oracle A, and $\ell = d$ in the case of Oracle B.

Proof: For each time $t = 1, 2, \dots, T$, let U_t denote the randomly chosen subset of size ℓ , $X_{t,i}$ be the outcome of oracle's coin toss at time t for coordinate i and let $Y_t \in \{-1, 0, 1\}^d$ be a random vector with entries

$$Y_{t,i} = \begin{cases} X_{t,i} & \text{if } i \in U_t, \text{ and} \\ -1 & \text{if } i \notin U_t. \end{cases}$$

By Fano's inequality [19], we have the lower bound

$$\mathbb{P}[\hat{\alpha} \neq \alpha^*] \geq 1 - \frac{I(\{(U_t, Y_t)\}_{t=1}^T; \alpha^*) + \log 2}{\log |\mathcal{V}|},$$

where $I(\{(U_t, Y_t)\}_{t=1}^T; \alpha^*)$ denotes the mutual information between the sequence $\{(U_t, Y_t)\}_{t=1}^T$ and the random parameter vector α^* . As discussed earlier, we are guaranteed that $\log |\mathcal{V}| \geq \frac{d}{2} \log(2/\sqrt{e})$. Consequently, in order to prove the

Oracle A: 1-dimensional unbiased gradients

- (a) Pick an index $i \in \{1, \dots, d\}$ uniformly at random.
- (b) Draw $b_i \in \{0, 1\}$ according to a Bernoulli distribution with parameter $1/2 + \alpha_i \delta$.
- (c) For the given input $x \in \mathbb{S}$, return the value $\widehat{g}_{\alpha, A}(x)$ and a sub-gradient $\widehat{z}_{\alpha, A}(x) \in \partial \widehat{g}_{\alpha, A}(x)$ of the function

$$\widehat{g}_{\alpha, A} := c[b_i f_i^+ + (1 - b_i) f_i^-].$$

Oracle B: d -dimensional unbiased gradients

- (a) For $i = 1, \dots, d$, draw $b_i \in \{0, 1\}$ according to a Bernoulli distribution with parameter $1/2 + \alpha_i \delta$.
- (b) For the given input $x \in \mathbb{S}$, return the value $\widehat{g}_{\alpha, B}(x)$ and a sub-gradient $\widehat{z}_{\alpha, B}(x) \in \partial \widehat{g}_{\alpha, B}(x)$ of the function

$$\widehat{g}_{\alpha, B} := \frac{c}{d} \sum_{i=1}^d [b_i f_i^+ + (1 - b_i) f_i^-].$$

lower bound (24), it suffices to establish the upper bound $I(\{U_t, Y_t\}_{t=1}^T; \alpha^*) \leq 16T \ell \delta^2$.

By the independent and identically distributed nature of the sampling model, we have

$$\begin{aligned} I(\{(U_1, Y_1), \dots, (U_T, Y_T)\}; \alpha^*) &= \sum_{t=1}^T I((U_t, Y_t); \alpha^*) \\ &= T I((U_1, Y_1); \alpha^*), \end{aligned}$$

so that it suffices to upper bound the mutual information for a single round. To simplify notation, from here onwards we write (Y, U) to mean the pair (Y_1, U_1) . With this notation, the remainder of our proof is devoted to establishing that $I(Y; U) \leq 16 \ell \delta^2$,

By chain rule for mutual information [19], we have

$$I((U, Y); \alpha^*) = I(Y; \alpha^* | U) + I(\alpha^*; U). \quad (25)$$

Since the subset U is chosen independently of α^* , we have $I(\alpha^*; U) = 0$, and so it suffices to upper bound the first term. By definition of conditional mutual information [19], we have

$$I(Y; \alpha^* | U) = \mathbb{E}_U [D(\mathbb{P}_{Y|\alpha^*, U} \| \mathbb{P}_{Y|U})]$$

Since α has a uniform distribution over \mathcal{V} , we have $\mathbb{P}_{Y|U} = \frac{1}{|\mathcal{V}|} \sum_{\alpha \in \mathcal{V}} \mathbb{P}_{Y|\alpha, U}$, and convexity of the Kullback-Leibler (KL) divergence yields the upper bound

$$D(\mathbb{P}_{Y|\alpha^*, U} \| \mathbb{P}_{Y|U}) \leq \frac{1}{|\mathcal{V}|} \sum_{\alpha \in \mathcal{V}} D(\mathbb{P}_{Y|\alpha^*, U} \| \mathbb{P}_{Y|\alpha, U}). \quad (26)$$

Now for any pair $\alpha^*, \alpha \in \mathcal{V}$, the KL divergence $D(\mathbb{P}_{Y|\alpha^*, U} \| \mathbb{P}_{Y|\alpha, U})$ can be at most the KL divergence between ℓ independent pairs of Bernoulli variates with parameters $\frac{1}{2} + \delta$ and $\frac{1}{2} - \delta$. Letting $D(\delta)$ denote the Kullback-Leibler divergence between a single pair of Bernoulli variables with parameters $\frac{1}{2} + \delta$ and $\frac{1}{2} - \delta$, a little calculation yields

$$\begin{aligned} D(\delta) &= \left(\frac{1}{2} + \delta\right) \log \frac{\frac{1}{2} + \delta}{\frac{1}{2} - \delta} + \left(\frac{1}{2} - \delta\right) \log \frac{\frac{1}{2} - \delta}{\frac{1}{2} + \delta} \\ &= 2\delta \log \left(1 + \frac{4\delta}{1 - 2\delta}\right) \\ &\leq \frac{8\delta^2}{1 - 2\delta}. \end{aligned}$$

Consequently, as long as $\delta \leq 1/4$, we have $D(\delta) \leq 16\delta^2$. Returning to the bound (26), we conclude that $D(\mathbb{P}_{Y|\alpha^*, U} \| \mathbb{P}_{Y|U}) \leq 16 \ell \delta^2$. Taking averages over U , we obtain the bound $I(Y; \alpha^* | U) \leq 16 \ell \delta^2$, and applying the decomposition (25) yields the upper bound $I((U, Y); \alpha^*) \leq 16 \ell \delta^2$, thereby completing the proof. ■

The reader might have observed that Fano's inequality yields a non-trivial lower bound only when $|\mathcal{V}|$ is large enough. Since $|\mathcal{V}|$ depends on the dimension d for our construction, we can apply the Fano lower bound only for d large enough. Smaller values of d can be lower bounded by reduction to the case $d = 1$; here we state a simple lower bound for estimating the bias of a single coin, which is a straightforward application of Le Cam's bounding technique [20], [12]. In this special case, we have $\mathcal{V} = \{1/2 + \delta, 1/2 - \delta\}$, and we recall that the estimator $\widehat{\alpha}(\mathcal{M}_T)$ takes values in \mathcal{V} .

Lemma 4. *Given a sample size $T \geq 1$ and a parameter $\alpha^* \in \mathcal{V}$, let $\{X_1, \dots, X_T\}$ be T i.i.d Bernoulli variables with parameter α^* . Let $\widehat{\alpha}$ be any test function based on these samples and returning an element of \mathcal{V} . Then for any $\delta \in (0, 1/4]$, we have the lower bound*

$$\sup_{\alpha^* \in \{\frac{1}{2} + \delta, \frac{1}{2} - \delta\}} \mathbb{P}_{\alpha^*}[\widehat{\alpha} \neq \alpha^*] \geq 1 - \sqrt{8T\delta^2}.$$

Proof: We observe first that for $\widehat{\alpha} \in \mathcal{V}$, we have

$$\mathbb{E}_{\alpha^*} [|\widehat{\alpha} - \alpha^*|] = 2\delta \mathbb{P}_{\alpha^*}[\widehat{\alpha} \neq \alpha^*],$$

so that it suffices to lower bound the expected error. To ease notation, let \mathbb{Q}_1 and \mathbb{Q}_{-1} denote the probability distributions indexed by $\alpha = \frac{1}{2} + \delta$ and $\alpha = \frac{1}{2} - \delta$ respectively. By Lemma 1 of Yu [12], we have

$$\sup_{\alpha^* \in \mathcal{V}} \mathbb{E}_{\alpha^*} [|\widehat{\alpha} - \alpha^*|] \geq 2\delta \left\{1 - \|\mathbb{Q}_1 - \mathbb{Q}_{-1}\|_1 / 2\right\}.$$

where we use the fact that $|(1/2 + \delta) - (1/2 - \delta)| = 2\delta$. Thus, we need to upper bound the total variation distance $\|\mathbb{Q}_1 - \mathbb{Q}_{-1}\|_1$. From Pinsker's inequality [19], we have

$$\|\mathbb{Q}_1 - \mathbb{Q}_{-1}\|_1 \leq \sqrt{2D(\mathbb{Q}_1 \| \mathbb{Q}_{-1})} \stackrel{(i)}{\leq} \sqrt{32T\delta^2},$$

where inequality (i) follows from the calculation following equation 26 (see proof of Lemma 3), and uses our assumption that $\delta \in (0, 1/4]$. Putting together the pieces, we obtain a lower bound on the probability of error

$$\sup_{\alpha^* \in \mathcal{V}} \mathbb{P}[\hat{\alpha} \neq \alpha^*] = \sup_{\alpha^* \in \mathcal{V}} \frac{\mathbb{E}|\hat{\alpha} - \alpha^*|}{2\delta} \geq 1 - \sqrt{8T\delta^2},$$

as claimed. \blacksquare

Equipped with these tools, we are now prepared to prove our main results.

B. Proof of Theorem 1

We begin with oracle complexity for bounded Lipschitz functions, as stated in Theorem 1. We first prove the result for the set $\mathbb{S} = \mathbb{B}_\infty(\frac{1}{2})$.

a) *Proof for $p \in [1, 2]$:* Consider Oracle A that returns the quantities $(\hat{g}_{\alpha, A}(x), \hat{z}_{\alpha, A}(x))$. By definition of the oracle, each round reveals only at most one coin flip, meaning that we can apply Lemma 3 with $\ell = 1$, thereby obtaining the lower bound

$$\mathbb{P}[\hat{\alpha}(\mathcal{M}_T) \neq \alpha] \geq 1 - 2 \frac{16T\delta^2 + \log 2}{d \log(2/\sqrt{e})}. \quad (27)$$

We now seek an upper bound $\mathbb{P}[\hat{\alpha}(\mathcal{M}_T) \neq \alpha]$ using Lemma 2. In order to do so, we need to specify the base functions (f_i^+, f_i^-) involved. For $i = 1, \dots, d$, we define

$$f_i^+(x) := \left| x(i) + \frac{1}{2} \right|, \quad \text{and} \quad f_i^-(x) := \left| x(i) - \frac{1}{2} \right|. \quad (28)$$

Given that $\mathbb{S} = \mathbb{B}_\infty(\frac{1}{2})$, we see that the minimizers of g_α are contained in S . Also, both the functions are 1-Lipschitz in the ℓ_1 -norm. By the construction (16), we are guaranteed that for any subgradient of g_α , we have

$$\|\hat{z}_{\alpha, A}(x)\|_p \leq 2c \quad \text{for all } p \geq 1.$$

Therefore, in order to ensure that g_α is L -Lipschitz in the dual ℓ_q -norm, it suffices to set $c = L/2$.

Let us now lower bound the discrepancy function (18). We first observe that each function g_α is minimized over the set $\mathbb{B}_\infty(\frac{1}{2})$ at the vector $x_\alpha := -\alpha/2$, at which point it achieves its minimum value

$$\min_{x \in \mathbb{B}_\infty(\frac{1}{2})} g_\alpha(x) = \frac{c}{2} - c\delta.$$

Furthermore, we note that for any $\alpha \neq \beta$, we have

$$\begin{aligned} g_\alpha(x) + g_\beta(x) &= \frac{c}{d} \sum_{i=1}^d \left[\left(\frac{1}{2} + \alpha_i \delta + \frac{1}{2} + \beta_i \delta \right) f_i^+(x) \right. \\ &\quad \left. + \left(\frac{1}{2} - \alpha_i \delta + \frac{1}{2} - \beta_i \delta \right) f_i^-(x) \right] \\ &= \frac{c}{d} \sum_{i=1}^d \left[(1 + \alpha_i \delta + \beta_i \delta) f_i^+(x) \right. \\ &\quad \left. + (1 - \alpha_i \delta - \beta_i \delta) f_i^-(x) \right] \\ &= \frac{c}{d} \sum_{i=1}^d \left[(f_i^+(x) + f_i^-(x)) \mathbb{I}(\alpha_i \neq \beta_i) \right. \\ &\quad \left. + ((1 + 2\alpha_i \delta) f_i^+(x) + (1 - 2\alpha_i \delta) f_i^-(x)) \mathbb{I}(\alpha_i = \beta_i) \right]. \end{aligned}$$

When $\alpha_i = \beta_i$ then $x_\alpha(i) = x_\beta(i) = -\alpha_i/2$, so that this co-ordinate does not make a contribution to the discrepancy function $\rho(g_\alpha, g_\beta)$. On the other hand, when $\alpha_i \neq \beta_i$, we have

$$f_i^+(x) + f_i^-(x) = \left| x(i) + \frac{1}{2} \right| + \left| x(i) - \frac{1}{2} \right| \geq 1$$

for all $x \in \mathbb{R}$. Consequently, any such co-ordinate yields a contribution of $2c\delta/d$ to the discrepancy. Recalling our packing set (15) with $d/4$ separation in Hamming norm, we conclude that for any distinct $\alpha \neq \beta$ within our packing set,

$$\rho(g_\alpha, g_\beta) = \frac{2c\delta}{d} \Delta_H(\alpha, \beta) \geq \frac{c\delta}{2},$$

so that by definition of ψ , we have established the lower bound $\psi(\delta) \geq \frac{c\delta}{2}$.

Setting the target error $\epsilon := \frac{c\delta}{18}$, we observe that this choice ensures that $\epsilon < \frac{\psi(\delta)}{9}$. Recalling the requirement $\delta < 1/4$, we have $\epsilon < c/72$. In this regime, we may apply Lemma 2 to obtain the upper bound $\mathbb{P}_\phi[\hat{\alpha}(\mathcal{M}_T) \neq \alpha] \leq \frac{1}{3}$. Combining this upper bound with the lower bound (27) yields the inequality

$$\frac{1}{3} \geq 1 - 2 \frac{16T\delta^2 + \log 2}{d \log(2/\sqrt{e})}.$$

Recalling that $c = \frac{L}{2}$, making the substitution $\delta = \frac{18\epsilon}{c} = \frac{36\epsilon}{L}$, and performing some algebra yields

$$T = \Omega\left(\frac{L^2 d}{\epsilon^2}\right) \quad \text{for all } d \geq 11 \text{ and for all } \epsilon \leq \frac{L}{144}.$$

Combined with Theorem 5.3.1 of NY [5] (or by using the lower bound of Lemma 4 instead of Lemma 3), we conclude that this lower bound holds for all dimensions d .

b) *Proof for $p > 2$:* The preceding proof based on Oracle A is also valid for $p > 2$, but yields a relatively weak result. Here we show how the use of Oracle B yields the stronger claim stated in Theorem 1(b). When using this oracle, all d coin tosses at each round are revealed, so that Lemma 3 with $\ell = d$ yields the lower bound

$$\mathbb{P}[\hat{\alpha}(\mathcal{M}_T) \neq \alpha] \geq 1 - 2 \frac{16Td\delta^2 + \log 2}{d \log(2/\sqrt{e})}. \quad (29)$$

We now seek an upper bound on $\mathbb{P}[\hat{\alpha}(\mathcal{M}_T) \neq \alpha]$. As before, we use the set $\mathbb{S} = \mathbb{B}_\infty(\frac{1}{2})$, and the previous definitions (28) of $f_i^+(x)$ and $f_i^-(x)$. From our earlier analysis (in particular, equation (23)), the quantity $\|\hat{z}_{\alpha, B}(x)\|_p$ is at most $cd^{1/p-1}$, so that setting $c = Ld^{1-1/p}$ yields functions that are Lipschitz with parameter L .

As before, for any distinct pair $\alpha, \beta \in \mathcal{V}$, we have the lower bound

$$\rho(g_\alpha, g_\beta) = \frac{2c\delta}{d} \Delta_H(\alpha, \beta) \geq \frac{c\delta}{2},$$

so that $\psi(\delta) \geq \frac{c\delta}{2}$. Consequently, if we set the target error $\epsilon := \frac{c\delta}{18}$, then we are guaranteed that $\epsilon < \frac{\psi(\delta)}{9}$, as is required for applying Lemma 2. Application of this lemma yields the upper bound $\mathbb{P}_\phi[\hat{\alpha}(\mathcal{M}_T) \neq \alpha] \leq \frac{1}{3}$. Combined with the lower bound (29), we obtain the inequality

$$\frac{1}{3} \geq 1 - 2 \frac{16dT\delta^2 + \log 2}{d \log(2/\sqrt{e})}.$$

Substituting $\delta = 18\epsilon/c$ yields the scaling $\epsilon = \Omega(\frac{c}{\sqrt{T}})$ for all $d \geq 11$ and $\epsilon \leq c/72$. Recalling that $c = Ld^{1-1/p}$, we obtain the bound (10). Combining this bound with Theorem 5.3.1 of NY [5], or alternatively, by using the lower bound of Lemma 4 instead of Lemma 3, we conclude that the claim holds for all dimensions.

We have thus completed the proof of Theorem 1 in the special case $\mathbb{S} = \mathbb{B}_\infty(\frac{1}{2})$. In order to prove the general claims, which scale with r when $B_\infty(r) \subseteq \mathbb{S}$, we note that our preceding proof required only that $\mathbb{S} \supseteq \mathbb{B}_\infty(\frac{1}{2})$ so that the minimizing points $x_\alpha = -\alpha/2 \in \mathbb{S}$ for all α (in particular, the Lipschitz constant of g_α does not depend on \mathbb{S} for our construction). In the general case, we define our base functions to be

$$f_i^+(x) = \left| x(i) + \frac{r}{2} \right|, \quad \text{and} \quad f_i^-(x) = \left| x(i) - \frac{r}{2} \right|.$$

With this choice, the functions $g_\alpha(x)$ are minimized at $x_\alpha = -r\alpha/2$, and $\inf_{x \in \mathbb{S}} g_\alpha(x) = cd/2 - cr\delta$. Mimicking the previous steps with $r = 1/2$, we obtain the lower bound

$$\rho(g_\alpha, g_\beta) \geq \frac{cr\delta}{2} \quad \forall \alpha \neq \beta \in \mathcal{V}.$$

The rest of the proof above did not depend on \mathbb{S} , so that we again obtain the lower bound $T = \Omega(\frac{d}{\delta^2})$ or $T = \Omega(\frac{1}{\delta^2})$ depending on the oracle used. In this case, the difference in ρ computation means that $\epsilon = \frac{L\delta r}{36} \leq \frac{Lr}{144}$, from which the general claims follow.

C. Proof of Theorem 2

We now turn to the proof of lower bounds on the oracle complexity of the class of strongly convex functions from Definition 3. In this case, we work with the following family of base functions, parametrized by a scalar $\theta \in [0, 1)$:

$$\begin{aligned} f_i^+(x) &:= r\theta|x(i) + r| + \frac{(1-\theta)}{4}(x(i) + r)^2, \quad \text{and} \\ f_i^-(x) &:= r\theta|x(i) - r| + \frac{(1-\theta)}{4}(x(i) - r)^2. \end{aligned} \quad (30)$$

A key ingredient of the proof is a uniform lower bound on the discrepancy ρ between pairs of these functions:

Lemma 5. *Using an ensemble based on the base functions (30), we have*

$$\rho(g_\alpha, g_\beta) \geq \begin{cases} \frac{2c\delta^2 r^2}{(1-\theta)d} \Delta_H(\alpha, \beta) & \text{if } 1-\theta \geq \frac{4\delta}{1+2\delta} \\ \frac{c\delta r^2}{d} \Delta_H(\alpha, \beta) & \text{if } 1-\theta < \frac{4\delta}{1+2\delta}. \end{cases} \quad (31)$$

The proof of this lemma is provided in Appendix A. Let us now proceed to the proofs of the main theorem claims.

c) Proof for $p = 1$: We observe that both the functions f_i^+, f_i^- are r -Lipschitz with respect to the $\|\cdot\|_1$ norm by construction. Hence, g_α is cr -Lipschitz and furthermore, by the definition of Oracle A, we have $\mathbb{E}\|\hat{z}_{\alpha, A}(x)\|_1^2 \leq c^2 r^2$. In addition, the function g_α is $(1-\theta)c/(4d)$ -strongly convex with respect to the Euclidean norm. We now follow the same steps as the proof of Theorem 1, but this time exploiting the ensemble formed by the base functions (30), and the lower

bound on the discrepancy $\rho(g_\alpha, g_\beta)$ from Lemma 5. We split our analysis into two sub-cases.

Case 1: First suppose that $1-\theta \geq 4\delta/(1+2\delta)$, in which case Lemma 5 yields the lower bound

$$\begin{aligned} \rho(g_\alpha, g_\beta) &\geq \frac{2c\delta^2 r^2}{(1-\theta)d} \Delta_H(\alpha, \beta) \\ &\stackrel{(i)}{\geq} \frac{c\delta^2 r^2}{2(1-\theta)} \quad \forall \alpha \neq \beta \in \mathcal{V}, \end{aligned}$$

where inequality (i) uses the fact that $\Delta_H(\alpha, \beta) \geq d/4$ by definition of \mathcal{V} . Hence by definition of ψ , we have established the lower bound $\psi(\delta) \geq \frac{c\delta^2 r^2}{2(1-\theta)}$. Setting the target error $\epsilon := c\delta^2 r^2/(18(1-\theta))$, we observe that this ensures $\epsilon \leq \psi(\delta)/9$. Recalling the requirement $\delta < 1/4$, we note that $\epsilon < cr^2/(288(1-\theta))$. In this regime, we may apply Lemma 2 to obtain the upper bound $\mathbb{P}_\phi[\hat{\alpha}(\mathcal{M}_T) \neq \alpha] \leq \frac{1}{3}$. Combining this upper bound with the lower bound (24) yields the inequality

$$\begin{aligned} \frac{1}{3} &\geq 1 - 2 \frac{16T\delta^2 + \log 2}{d \log(2/\sqrt{\epsilon})} \\ &\geq 1 - 2 \frac{\frac{288T\epsilon(1-\theta)}{cr^2} + \log 2}{d \log(2/\sqrt{\epsilon})}. \end{aligned}$$

Simplifying the above expression yields that for $d \geq 11$, we have the lower bound

$$\begin{aligned} T &\geq cr^2 \left(\frac{\frac{d}{3} \log(2/\sqrt{\epsilon}) - \log 2}{288\epsilon(1-\theta)} \right) \\ &\geq cr^2 \frac{d \log(2/\sqrt{\epsilon})}{28800\epsilon(1-\theta)}. \end{aligned} \quad (32)$$

Finally, we observe that $L = cr$ and $\gamma^2 = (1-\theta)c/(4d)$ which gives $1-\theta = 4dr\gamma^2/L$. Substituting the above relations in the lower bound (32) gives the first term in the stated result for $d \geq 11$.

To obtain lower bounds for dimensions $d < 11$, we use an argument based on $d = 1$. For this special case, we consider f^+ and f^- to be the two functions of the single coordinate coming out of definition (30). The packing set \mathcal{V} consists of only two elements now, corresponding to $\alpha = 1$ and $\alpha = -1$. Specializing the result of Lemma 5 to this case, we see that the two functions are $2c\delta^2 r^2/(1-\theta)$ separated. Now we again apply Lemma 2 to get an upper bound on the error probability and Lemma 4 to get a lower bound, which gives the result for $d \leq 11$.

Case 2: On the other hand, suppose that $1-\theta \leq 4\delta/(1+2\delta)$. In this case, appealing to Lemma 5 gives us that $\rho(g_\alpha, g_\beta) \geq c\delta r^2/4$ for $\alpha \neq \beta \in \mathcal{V}$. Recalling that $L = cr$, we set the desired accuracy $\epsilon := c\delta r^2/36 = L\delta r/36$. From this point onwards, we mimic the proof of Theorem 1; doing so yields that for all $\delta \in (0, 1/4)$, we have

$$T = \Omega\left(\frac{d}{\delta^2}\right) = \Omega\left(\frac{L^2 dr^2}{\epsilon^2}\right),$$

corresponding to the second term in Theorem 1.

Finally, the third and fourth terms are obtained just like Theorem 1 by checking the condition $\delta < 1/4$ in the two cases above. Overall, this completes the proof for the case $p = 1$.

d) Proof for $p > 2$: As with the proof of Theorem 1(b), we use Oracle B that returns d -dimensional values and gradients in this case, with the base functions defined in equation 30. With this choice, we have the upper bound

$$\mathbb{E}\|\widehat{z}_{\alpha,B}(x)\|_p^2 \leq c^2 d^{2/p-2} r^2,$$

so that setting the constant $c = Ld^{1-1/p}/r$ ensures that $\mathbb{E}\|\widehat{z}_{\alpha,B}(x)\|_p^2 \leq L^2$. As before, we have the strong convexity parameter

$$\gamma^2 = \frac{c(1-\theta)}{4d} = \frac{Ld^{-1/p}(1-\theta)}{4r},$$

Also $\rho(g_\alpha, g_\beta)$ is given by Lemma 5. In particular, let us consider the case $1-\theta \geq 4\delta/(1+2\delta)$ so that $\psi(\delta) \geq \frac{c\delta^2 r^2}{2(1-\theta)}$, and we set the desired accuracy $\epsilon := \frac{c\delta^2 r^2}{18(1-\theta)}$ as before. With this setting of ϵ , we invoke Lemma 2 as before to argue that $\mathbb{P}_\phi[\widehat{\alpha}(\mathcal{M}_T) \neq \alpha] \leq \frac{1}{3}$. To lower bound the error probability, we appeal to Lemma 3 with $\ell = d$ just like Theorem 1(b) and obtain the inequality

$$\frac{1}{3} \geq 1 - 2 \frac{16 d T \delta^2 + \log 2}{d \log(2/\sqrt{\epsilon})}.$$

Rearranging terms and substituting $\epsilon = \frac{c\delta^2 r^2}{18(1-\theta)}$, we obtain for $d \geq 11$

$$T = \Omega\left(\frac{1}{\delta^2}\right) = \Omega\left(\frac{cr^2}{\epsilon(1-\theta)}\right).$$

The stated result can now be attained by recalling $c = Ld^{1-1/p}/r$ and $\gamma^2 = Ld^{-1/p}(1-\theta)/r$ for $1-\theta \geq 4\delta/(1+2\delta)$ and $d \geq 11$. For $d < 11$, the cases of $p > 2$ and $p = 1$ are identical up to constant factors in the lower bounds we state. This completes the proof for $1-\theta \geq 4\delta/(1+2\delta)$.

Finally, the case for $1-\theta < 4\delta/(1+2\delta)$ involves similar modifications as part(a) by using the different expression for $\rho(g_\alpha, g_\beta)$. Thus we have completed the proof of this theorem.

D. Proof of Theorem 3

We begin by constructing an appropriate subset of $\mathcal{F}_{\text{sp}}(k)$ over which the Fano method can be applied. Let $\mathcal{V}(k) := \{\alpha^1, \dots, \alpha^M\}$ be a set of vectors, such that each $\alpha^j \in \{-1, 0, +1\}^d$ satisfies

$$\|\alpha^j\|_0 = k \quad \forall j = 1, \dots, M \text{ and } \Delta_H(\alpha^j, \alpha^\ell) \geq \frac{k}{2}$$

for all $j \neq \ell$. It can be shown that there exists such a packing set with $|\mathcal{V}(k)| \geq \exp\left(\frac{k}{2} \log \frac{d-k}{k/2}\right)$ elements (e.g., see Lemma 5 in Raskutti et al. [22]).

For any $\alpha \in \mathcal{V}(k)$, we define the function $x \mapsto g_\alpha(x)$ via

$$c \left[\sum_{i=1}^d \left\{ \left(\frac{1}{2} + \alpha_i \delta \right) |x(i) + r| + \left(\frac{1}{2} - \alpha_i \delta \right) |x(i) - r| \right\} + \delta \sum_{i=1}^d |x(i)| \right]. \quad (33)$$

In this definition, the quantity $c > 0$ is a pre-factor to be chosen later, and $\delta \in (0, \frac{1}{4}]$ is a given error tolerance. Observe that each function $g_\alpha \in \mathcal{G}(\delta; k)$ is convex, and Lipschitz with parameter c with respect to the $\|\cdot\|_\infty$ norm.

Central to the remainder of the proof is the function class $\mathcal{G}(\delta; k) := \{g_\alpha, \alpha \in \mathcal{V}(k)\}$. In particular, we need to control the discrepancy $\psi(\delta; k) := \psi(\mathcal{G}(\delta; k))$ for this class. The following result, proven in Appendix B, provides a suitable lower bound:

Lemma 6. *We have*

$$\psi(\delta; k) = \inf_{\alpha \neq \beta \in \mathcal{V}(k)} \rho(g_\alpha, g_\beta) \geq \frac{ck\delta r}{4}. \quad (34)$$

Using Lemma 6, we may complete the proof of Theorem 3. Define the base functions

$$f_i^+(x) := d(|x(i) + r| + \delta|x(i)|), \quad \text{and} \\ f_i^-(x) := d(|x(i) - r| + \delta|x(i)|).$$

Consider Oracle B, which returns d -dimensional gradients based on the function

$$\widehat{g}_{\alpha,B}(x) = \frac{c}{d} \sum_{i=1}^d [b_i f_i^+(x) + (1-b_i) f_i^-(x)],$$

where $\{b_i\}$ are Bernoulli variables. By construction, the function $\widehat{g}_{\alpha,B}$ is at most $3c$ -Lipschitz in ℓ_∞ norm (i.e. $\|\widehat{z}_{\alpha,B}(x)\|_\infty \leq 3c$), so that setting $c = \frac{L}{3}$ yields an L -Lipschitz function.

Our next step is to use Fano's inequality [19] to lower bound the probability of error in the multiway testing problem associated with this stochastic oracle, following an argument similar to (but somewhat simpler than) the proof of Lemma 3. Fano's inequality yields the lower bound

$$\mathbb{P}[\widehat{\alpha} \neq \alpha^*] \geq 1 - \frac{\frac{1}{\binom{|\mathcal{V}|}{2}} \sum_{\alpha \neq \beta} D(\mathbb{P}_\alpha \| \mathbb{P}_\beta) + \log 2}{\log |\mathcal{V}|}. \quad (35)$$

(As in the proof of Lemma 3, we have used convexity of mutual information [19] to bound it by the average of the pairwise KL divergences.) By construction, any two parameters $\alpha, \beta \in \mathcal{V}$ differ in at most $2k$ places, and the remaining entries are all zeroes in both vectors. The proof of Lemma 3 shows that for $\delta \in [0, \frac{1}{4}]$, each of these $2k$ places makes a contribution of at most $16\delta^2$. Recalling that we have T samples, we conclude that $D(\mathbb{P}_\alpha \| \mathbb{P}_\beta) \leq 32kT\delta^2$. Substituting this upper bound into the Fano lower bound (35) and recalling that the cardinality of \mathcal{V} is at least $\exp\left(\frac{k}{2} \log \frac{d-k}{k/2}\right)$, we obtain

$$\mathbb{P}[\widehat{\alpha}(\mathcal{M}_T) \neq \alpha] \geq 1 - 2 \left(\frac{32kT\delta^2 + \log 2}{\frac{k}{2} \log \frac{d-k}{k/2}} \right) \quad (36)$$

By Lemma 6 and our choice $c = L/3$, we have

$$\psi(\delta) \geq \frac{ck\delta r}{4} = \frac{Lk\delta r}{12},$$

Therefore, if we aim for the target error $\epsilon = \frac{Lk\delta r}{108}$, then we are guaranteed that $\epsilon \leq \frac{\psi(\delta)}{9}$, as is required for the application of Lemma 2. Recalling the requirement $\delta \leq 1/4$ gives $\epsilon \leq$

$Lk\delta r/432$. Now Lemma 2 implies that $\mathbb{P}[\widehat{\alpha}(\mathcal{M}_T) \neq \alpha] \leq 1/3$, which when combined with the earlier bound (36) yields

$$\frac{1}{3} \geq 1 - 2 \left(\frac{32kT\delta^2 + \log 2}{\frac{k}{2} \log \frac{d-k}{k/2}} \right).$$

Rearranging yields the lower bound

$$T = \Omega \left(\frac{\log \frac{d-k}{k/2}}{\delta^2} \right) = \Omega \left(L^2 \gamma^2 k^2 \frac{\log \frac{d-k}{k/2}}{\epsilon^2} \right),$$

where the second step uses the relation $\delta = \frac{108\epsilon}{Lkr}$. As long as $k \leq \lfloor d/2 \rfloor$, we have $\log \frac{d-k}{k/2} = \Theta \left(\log \frac{d}{k} \right)$, which completes the proof.

V. DISCUSSION

In this paper, we have studied the complexity of convex optimization within the stochastic first-order oracle model. We derived lower bounds for various function classes, including convex functions, strongly convex functions, and convex functions with sparse optima. As we discussed, our lower bounds are sharp in general, since there are matching upper bounds achieved by known algorithms, among them stochastic gradient descent and stochastic mirror descent. Our bounds also reveal various dimension-dependent and geometric aspects of the stochastic oracle complexity of convex optimization. An interesting aspect of our proof technique is the use of tools common in statistical minimax theory. In particular, our proofs are based on constructing packing sets, defined with respect to a pre-metric that measures how the degree of separation between the optima of different functions. We then leveraged information-theoretic techniques, in particular Fano's inequality and its variants, in order to establish lower bounds.

There are various directions for future research. It would be interesting to consider the effect of memory constraints on the complexity of convex optimization, or to derive lower bounds for problems of distributed optimization. We suspect that the proof techniques developed in this paper may be useful for studying these related problems.

APPENDIX

A. Proof of Lemma 5

Let g_α and g_β be an arbitrary pair of functions in our class, and recall that the constraint set \mathbb{S} is given by the ball $\mathbb{B}_\infty(r)$. From the definition (18) of the discrepancy ρ , we need to compute the single function infimum $\inf_{x \in \mathbb{B}_\infty(r)} g_\alpha(x)$, as well as the quantity $\inf_{x \in \mathbb{B}_\infty(r)} \{g_\alpha(x) + g_\beta(x)\}$.

a) Evaluating the single function infimum: Beginning with the former quantity, first observe that for any $x \in \mathbb{B}_\infty(r)$, we have

$$|x(i) + r| = x(i) + r \quad \text{and} \quad |x(i) - r| = r - x(i).$$

Consequently, using the definition (30) of the base functions, some algebra yields the relations

$$\begin{aligned} f_i^+(x) &= \frac{1-\theta}{4}x(i)^2 + \frac{1+3\theta}{4}r^2 + \frac{(1+\theta)}{2}rx(i), \quad \text{and} \\ f_i^-(x) &= \frac{1-\theta}{4}x(i)^2 + \frac{1+3\theta}{4}r^2 - \frac{(1+\theta)}{2}rx(i). \end{aligned}$$

Using these expressions for f_i^+ and f_i^- , we obtain that the quantity $h_i(x) := \left(\frac{1}{2} + \alpha_i\delta\right) f_i^+(x) + \left(\frac{1}{2} - \alpha_i\delta\right) f_i^-(x)$ can be written as

$$\begin{aligned} h_i(x) &= \frac{1}{2}(f_i^+(x) + f_i^-(x)) + \alpha_i\delta(f_i^+(x) - f_i^-(x)) \\ &= \frac{1-\theta}{4}x(i)^2 + \frac{1+3\theta}{4}r^2 + (1+\theta)\alpha_i\delta rx(i). \end{aligned}$$

A little calculation shows that constrained minimum of the univariate function h_i over the interval $[-r, r]$ is achieved at

$$x^*(i) := \begin{cases} \frac{-2\alpha_i\delta r(1+\theta)}{1-\theta} & \text{if } \frac{1-\theta}{1+\theta} \geq 2\delta \\ -\alpha_i r & \text{if } \frac{1-\theta}{1+\theta} < 2\delta, \end{cases}$$

where we have recalled that α_i takes values in $\{-1, +1\}$. Substituting the minimizing argument $x^*(i)$, we find that the minimum value is given by

$$h_i(x^*(i)) = \begin{cases} \frac{1+3\theta}{4}r^2 - \frac{\delta^2 r^2(1+\theta)^2}{(1-\theta)} & \text{if } \frac{1-\theta}{1+\theta} \geq 2\delta \\ \frac{1+\theta}{2}r^2 - (1+\theta)\delta r^2 & \text{if } \frac{1-\theta}{1+\theta} < 2\delta. \end{cases}$$

Summing over all co-ordinates $i \in \{1, 2, \dots, d\}$ yields that $\inf_{x \in \mathbb{B}_\infty(r)} g_\alpha(x) = \frac{c}{d} \sum_{i=1}^d h_i(x^*(i))$, and hence that

$$\inf_{x \in \mathbb{B}_\infty(r)} g_\alpha(x) = \begin{cases} -\frac{\delta^2 r^2 c(1+\theta)^2}{(1-\theta)} + \frac{cr^2(1+3\theta)}{4} & \text{if } \frac{1-\theta}{1+\theta} \geq 2\delta \\ \frac{1+\theta}{2}cr^2 - (1+\theta)c\delta r^2 & \text{if } \frac{1-\theta}{1+\theta} < 2\delta. \end{cases} \quad (37)$$

b) Evaluating the joint infimum: Here we begin by observing that for any two $\alpha, \beta \in \mathcal{V}$, we have

$$\begin{aligned} g_\alpha(x) + g_\beta(x) &= \frac{c}{d} \sum_{i=1}^d \left[\frac{1-\theta}{2}x(i)^2 + \frac{1+3\theta}{2}r^2 \right. \\ &\quad \left. + 2(1+\theta)\alpha_i\delta rx(i)\mathbb{I}(\alpha_i = \beta_i) \right]. \quad (38) \end{aligned}$$

As in our previous calculation, the only coordinates that contribute to $\rho(g_\alpha, g_\beta)$ are the ones where $\alpha_i \neq \beta_i$, and for such coordinates, the function above is minimized at $x^*(i) = 0$. Furthermore, the minimum value for any such coordinate is $(1+3\theta)cr^2/(2d)$.

We split the remainder of our analysis into two cases: first, if we suppose that $\frac{1-\theta}{1+\theta} \geq 2\delta$, or equivalently that $1-\theta \geq 4\delta/(1+2\delta)$, then equation (38) yields that

$$\begin{aligned} \inf_{x \in \mathbb{B}_\infty(r)} \{g_\alpha(x) + g_\beta(x)\} &= \\ &= \frac{c}{d} \sum_{i=1}^d \left[\frac{1+3\theta}{2}r^2 - \frac{2\delta^2 r^2(1+\theta)^2}{1-\theta} \mathbb{I}(\alpha_i = \beta_i) \right]. \end{aligned}$$

Combined with our earlier expression (37) for the single function infimum, we obtain that the discrepancy is given by

$$\begin{aligned} \rho(g_\alpha, g_\beta) &= \frac{2\delta^2 r^2 c(1+\theta)^2}{d(1-\theta)} \Delta_H(\alpha, \beta) \\ &\geq \frac{2\delta^2 r^2 c}{d(1-\theta)} \Delta_H(\alpha, \beta). \end{aligned}$$

On the other hand, if we assume that $\frac{1-\theta}{1+\theta} < 2\delta$, or equivalently that $1-\theta < 4\delta/(1+2\delta)$, then we obtain

$$\begin{aligned} &\inf_{x \in \mathbb{B}_\infty(r)} \{g_\alpha(x) + g_\beta(x)\} = \\ &\frac{c}{d} \sum_{i=1}^d \left[\frac{1+3\theta}{2} r^2 - \left(2(1+\theta)r^2\delta - \frac{1-\theta}{2} r^2 \right) \mathbb{I}(\alpha_i = \beta_i) \right]. \end{aligned}$$

Combined with our earlier expression (37) for the single function infimum, we obtain

$$\begin{aligned} \rho(g_\alpha, g_\beta) &= \frac{c}{d} \left(2(1+\theta)r^2\delta - \frac{1-\theta}{2} r^2 \right) \Delta_H(\alpha, \beta) \\ &\stackrel{(i)}{\geq} \frac{c(1+\theta)r^2\delta}{d} \Delta_H(\alpha, \beta), \end{aligned}$$

where step (i) uses the bound $1-\theta < 2\delta(1+\theta)$. Noting that $\theta \geq 0$ completes the proof of the lemma.

B. Proof of Lemma 6

Recall that the constraint set \mathbb{S} in this lemma is the ball $\mathbb{B}_\infty(r)$. Thus, recalling the definition (18) of the discrepancy ρ , we need to compute the single function infimum $\inf_{x \in \mathbb{B}_\infty(r)} g_\alpha(x)$, as well as the quantity $\inf_{x \in \mathbb{B}_\infty(r)} \{g_\alpha(x) + g_\beta(x)\}$.

c) Evaluating the single function infimum: Beginning with the former quantity, first observe that for any $x \in \mathbb{B}_\infty(r)$, we have

$$\begin{aligned} &\left[\frac{1}{2} + \alpha_i \delta \right] |x(i) + r| + \left[\frac{1}{2} - \alpha_i \delta \right] |x(i) - r| \\ &= r + 2\alpha_i \delta x(i). \end{aligned} \quad (39)$$

We now consider one of the individual terms arising in the definition (16) of the function g_α . Using the relation (39), we see that

$$\begin{aligned} &\frac{1}{d} \left[\left(\frac{1}{2} + \alpha_i \delta \right) f_i^+(x) + \left(\frac{1}{2} - \alpha_i \delta \right) f_i^-(x) \right] = \\ &\left(\frac{1}{2} + \alpha_i \delta \right) |x(i) + r| + \left(\frac{1}{2} - \alpha_i \delta \right) |x(i) - r| + \delta |x(i)|, \end{aligned}$$

which is equal to

$$\begin{cases} r + (2\alpha_i + 1)\delta x(i) & \text{if } x(i) \geq 0 \\ r + (2\alpha_i - 1)\delta x(i) & \text{if } x(i) \leq 0 \end{cases}$$

From this representation, we see that whenever $\alpha_i \neq 0$, then the i^{th} term in the summation defining g_α minimized at $x(i) = -r\alpha_i$, at which point it takes on its minimum value $r(1-\delta)$. On the other hand, for any term with $\alpha_i = 0$, the function is minimized at $x(i) = 0$ with associated minimum value of r . Combining these two facts shows that the vector $-\alpha r$ is an element of the set $\arg \min_{x \in \mathbb{S}} g_\alpha(x)$, and moreover that

$$\inf_{x \in \mathbb{S}} g_\alpha(x) = cr(d - k\delta). \quad (40)$$

d) Evaluating the joint infimum: We now turn to the computation of $\inf_{x \in \mathbb{B}_\infty(r)} \{g_\alpha(x) + g_\beta(x)\}$. From the relation (39) and the definitions of g_α and g_β , some algebra yields

$$\begin{aligned} &\inf_{x \in \mathbb{S}} \{g_\alpha(x) + g_\beta(x)\} = \\ &c \inf_{x \in \mathbb{S}} \sum_{i=1}^d \{2r + 2\delta [(\alpha_i + \beta_i)x(i) + |x(i)|]\}. \end{aligned} \quad (41)$$

Let us consider the minimizer of the i^{th} term in this summation. First, suppose that $\alpha_i \neq \beta_i$, in which case there are two possibilities.

- If $\alpha_i \neq \beta_i$ and neither α_i nor β_i is zero, then we must have $\alpha_i + \beta_i = 0$, so that the minimum value of $2r$ is achieved at $x(i) = 0$.
- Otherwise, suppose that $\alpha_i \neq 0$ and $\beta_i = 0$. In this case, we see from equation (41) that it is equivalent to minimizing $\alpha_i x(i) + |x(i)|$. Setting $x(i) = -\alpha_i$ achieves the minimum value of $2r$.

In the remaining two cases, we have $\alpha_i = \beta_i$.

- If $\alpha_i = \beta_i \neq 0$, then the component is minimized at $x(i) = -\alpha_i r$ and the minimum value along the component is $2r(1-\delta)$.
- If $\alpha_i = \beta_i = 0$, then the minimum value is $2r$, achieved at $x(i) = 0$.

Consequently, accumulating all of these individual cases into a single expression, we obtain

$$\inf_{x \in \mathbb{S}} \{g_\alpha(x) + g_\beta(x)\} = 2cr \left(d - \delta \sum_{i=1}^d \mathbb{I}[\alpha_i = \beta_i \neq 0] \right). \quad (42)$$

Finally, combining equations (40) and (42) in the definition of ρ , we find that

$$\begin{aligned} \rho(g_\alpha, g_\beta) &= 2cr \left[d - \delta \sum_{i=1}^d \mathbb{I}[\alpha_i = \beta_i \neq 0] - (d - k\delta) \right] \\ &= 2c\delta r \left[k - \sum_{i=1}^d \mathbb{I}[\alpha_i = \beta_i \neq 0] \right] \\ &= cr\delta \Delta_H(\alpha, \beta), \end{aligned}$$

where the second equality follows since α and β have exactly k non-zero elements each. Finally, since \mathcal{V} is an $k/2$ -packing set in Hamming distance, we have $\Delta_H(\alpha, \beta) \geq k/2$, which completes the proof.

C. Upper bounds via mirror descent

This appendix is devoted to background on the family of mirror descent methods. We first describe the basic form of the algorithm and some known convergence results, before showing that different forms of mirror descent provide matching upper bounds for several of the lower bounds established in this paper, as discussed in the main text.

1) *Background on mirror descent*: Mirror descent is a generalization of (projected) stochastic gradient descent, first introduced by Nemirovski and Yudin [5]; here we follow a more recent presentation of it due to Beck and Teboulle [23]. For a given norm $\|\cdot\|$, let $\Phi : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ be a differentiable function that is 1-strongly convex with respect to $\|\cdot\|$, meaning that

$$\Phi(y) \geq \Phi(x) + \langle \nabla \Phi(x), y - x \rangle + \frac{1}{2} \|y - x\|^2.$$

We assume that Φ is a function of Legendre type [24], [25], which implies that the conjugate dual Φ^* is differentiable on its domain with $\nabla \Phi^* = (\nabla \Phi)^{-1}$. For a given proximal function, we let D_Φ be the Bregman divergence induced by Φ , given by

$$D_\Phi(x, y) := \Phi(x) - \Phi(y) - \langle \nabla \Phi(y), x - y \rangle. \quad (43)$$

With this set-up, we can now describe the mirror descent algorithm based on the proximal function Φ for minimizing a convex function f over a convex set \mathbb{S} contained within the domain of Φ . Starting with an arbitrary initial $x_0 \in \mathbb{S}$, it generates a sequence $\{x_t\}_{t=0}^\infty$ contained within \mathbb{S} via the updates

$$x_{t+1} = \arg \min_{x \in \mathbb{S}} \{ \eta_t \langle x, \nabla f(x_t) \rangle + D_\Phi(x, x_t) \}, \quad (44)$$

where $\eta_t > 0$ is a stepsize. In case of stochastic optimization, $\nabla f(x_t)$ is simply replaced by the noisy version $\widehat{z}(x_t)$.

A special case of this algorithm is obtained by choosing the proximal function $\Phi(x) = \frac{1}{2} \|x\|_2^2$, which is 1-strongly convex with respect to the Euclidean norm. The associated Bregman divergence $D_\Phi(x, y) = \frac{1}{2} \|x - y\|_2^2$ is simply (a scaled version of) the Euclidean norm, so that the updates (44) correspond to a standard projected gradient descent method. If one receives only an unbiased estimate of the gradient $\nabla f(x_t)$, then this algorithm corresponds to a form of projected stochastic gradient descent. Moreover, other choices of the proximal function lead to different stochastic algorithms, as discussed below.

Explicit convergence rates for this algorithm can be obtained under appropriate convexity and Lipschitz assumptions for f . Following the set-up used in our lower bound analysis, we assume that $\mathbb{E} \|\nabla \widehat{z}(x_t)\|_*^2 \leq L^2$ for all $x \in \mathbb{S}$, where $\|v\|_* := \sup_{\|x\| \leq 1} \langle x, v \rangle$ is the dual norm defined by $\|\cdot\|$. Given stochastic mirror descent based on unbiased estimates of the gradient, it can be showed that (see e.g., Chapter 5.1 of NY [5] or Beck and Teboulle [23]) with the initialization $x_0 = \arg \min_{x \in \mathbb{S}} \Phi(x)$ and stepsizes $\eta_t = 1/\sqrt{t}$, the optimization error of the sequence $\{x_t\}$ is bounded as

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E} [f(x_t) - f(x^*)] &\leq L \sqrt{\frac{D_\Phi(x^*, x_1)}{T}} \\ &\leq L \sqrt{\frac{\Phi(x^*)}{T}} \end{aligned} \quad (45)$$

Note that this averaged convergence is a little different from the convergence of x_T discussed in our lower bounds. In order to relate the two quantities, observe that by Jensen's inequality

$$\mathbb{E} \left[f \left(\frac{\sum_{t=1}^T x_t}{T} \right) \right] \leq \frac{1}{T} \mathbb{E} [f(x_t)].$$

Consequently, based on mirror descent for $T - 1$ rounds, we may set $x_T = \frac{1}{T-1} \sum_{t=1}^{T-1} x_t$ so as to obtain the same convergence bounds up to constant factors. In the following discussion, we assume this choice of x_T for comparing the mirror descent upper bounds to our lower bounds.

2) *Matching upper bounds*: Now consider the form of mirror descent obtained by choosing the proximal function

$$\Phi_a(x) := \frac{1}{2(a-1)} \|x\|_a^2 \quad \text{for } 1 < a \leq 2. \quad (46)$$

Note that this proximal function is 1-strongly convex with respect to the ℓ_a -norm for $1 < a \leq 2$, meaning that $\frac{1}{2(a-1)} \|x\|_a^2$ is lower bounded by

$$\frac{1}{2(a-1)} \|y\|_a^2 + \left(\nabla \frac{1}{2(a-1)} \|x\|_a^2 \right)^T (x - y) + \frac{1}{2} \|x - y\|_a^2.$$

a) *Upper bounds for dual setting*: Let us start from the case $1 \leq p \leq 2$. In this case we use stochastic gradient descent with \cdot , and the choice of p ensures that $\mathbb{E} \|\widehat{z}(x)\|_2^2 \leq \mathbb{E} \|\widehat{z}(x)\|_p^2 \leq L^2$ (the second inequality is true by assumption of Theorem 1). Also a straightforward calculation shows that $\|x^*\|_2 \leq \|x^*\|_q d^{1/2-1/q}$, which leads to

$$\mathbb{E} [f(x_T) - f(x^*)] = \mathcal{O} \left(\frac{L d^{1/2-1/q}}{\sqrt{T}} \right).$$

This upper bound matches the lower bound from equation (11) in this case. For $p \geq 2$, we use mirror descent with $a = q = p/(p-1)$. In this case, $\mathbb{E} \|\widehat{z}(x)\|_p^2 \leq L^2$ and $\|x^*\|_q \leq 1$ for the convex set $\mathbb{B}_q(1)$ and the function class $\mathcal{F}_{\text{cv}}(\mathbb{B}_q(1), L, p)$. Hence in this case, the upper bound from equation 45 is $\mathcal{O}(L/\sqrt{T})$ as long as $p = o(\log d)$, which again matches our lower bound from Equation 11. Finally, for $p = \Omega(\log d)$, we use mirror descent with $a = 2 \log d / (2 \log d - 1)$, which gives an upper bound of $\mathcal{O}(L \sqrt{\log d / T})$, since $1/(a-1) = \mathcal{O}(\log d)$ in this regime.

b) *Upper bounds for ℓ_∞ ball*: For this case, we use mirror descent based on the proximal function Φ_a with $a = q$. Under the condition $\|x^*\|_\infty \leq 1$, a condition which holds in our lower bounds, we obtain

$$\|x^*\|_q \leq \|x^*\|_\infty d^{1/q} = d^{1/q},$$

which implies that $\Phi_q(x^*) = \mathcal{O}(d^{2/q})$. Under the conditions of Theorem 1, we have $\mathbb{E} \|\widehat{z}(x_t)\|_p^2 \leq L^2$ where $p = q/(q-1)$ defines the dual norm. Note that the condition $1 < q \leq 2$ implies that $p \geq 2$. Substituting this in the upper bound (45) yields

$$\begin{aligned} \mathbb{E} [f(x_T) - f(x^*)] &= \mathcal{O} \left(L \sqrt{d^{2/q} / T} \right) \\ &= \mathcal{O} \left(L d^{1-1/p} \sqrt{\frac{1}{T}} \right), \end{aligned}$$

which matches the lower bound from Theorem 1(b). (Note that there is an additional log factor, as in the previous discussion, which we ignore.)

For $1 \leq p \leq 2$, we use stochastic gradient descent with $q = 2$, in which case $\|x^*\|_2 \leq \sqrt{d}$ and $\mathbb{E} \|\widehat{z}(x_t)\|_2^2 \leq \mathbb{E} \|\widehat{z}(x_t)\|_p^2 \leq L^2$ by assumption. Substituting these in the upper bound for mirror descent yields an upper bound to match the lower bound of Theorem 1(a).

c) *Upper bounds for Theorem 3:* In order to recover matching upper bounds in this case, we use the function Φ_a from equation (46) with $a = \frac{2 \log d}{2 \log d - 1}$. In this case, the resulting upper bound (45) on the convergence rate takes the form

$$\begin{aligned} \mathbb{E}[f(x_T) - f(x^*)] &= \mathcal{O}\left(L\sqrt{\frac{\|x^*\|_a^2}{2(a-1)T}}\right) \\ &= \mathcal{O}\left(L\sqrt{\frac{\|x^*\|_a^2 \log d}{T}}\right), \end{aligned} \quad (47)$$

since $\frac{1}{a-1} = 2 \log d - 1$. Based on the conditions of Theorem 3, we are guaranteed that x^* is k -sparse, with every component bounded by 1 in absolute value, so that $\|x^*\|_a^2 \leq k^{2/a} \leq k^2$, where the final inequality follows since $a > 1$. Substituting this upper bound back into equation (47) yields

$$\mathbb{E}[f(x_T) - f(x^*)] = \mathcal{O}\left(L\sqrt{\frac{k^2 \log d}{T}}\right).$$

Note that whenever $k = \mathcal{O}(d^{1-\delta})$ for some $\delta > 0$, then we have $\log d = \Theta(\log \frac{d}{k})$, in which case this upper bound matches the lower bound from Theorem 3 up to constant factors, as claimed.

REFERENCES

- [1] L. Bottou and O. Bousquet, "The tradeoffs of large scale learning," in *NIPS*, 2008.
- [2] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge, UK: Cambridge University Press, 2004.
- [3] D. Bertsekas, *Nonlinear programming*. Belmont, MA: Athena Scientific, 1995.
- [4] Y. Nesterov, *Introductory lectures on convex optimization: Basic course*. Kluwer Academic Publishers, 2004.
- [5] A. S. Nemirovski and D. B. Yudin, *Problem Complexity and Method Efficiency in Optimization*. John Wiley UK/USA, 1983.
- [6] A. S. Nemirovski, "Efficient methods in convex programming," Georgia Tech, Lecture notes, Tech. Rep., 2010. [Online]. Available: http://www2.isye.gatech.edu/~nemirovs/OPTI_LectureNotes.pdf
- [7] A. Agarwal, P. Bartlett, P. Ravikumar, and M. J. Wainwright, "Information-theoretic lower bounds on the oracle complexity of convex optimization," in *Advances in Neural Information Processing Systems* 22, 2009, pp. 1–9.
- [8] M. Raginsky and A. Rakhlin, "Information-based complexity, feedback and dynamics in convex programming," *IEEE Transactions on Information Theory*, 2011, to appear.
- [9] R. Z. Has'minskii, "A lower bound on the risks of nonparametric estimates of densities in the uniform metric," *Theory Prob. Appl.*, vol. 23, pp. 794–798, 1978.
- [10] L. Birgé, "Approximation dans les espaces metriques et theorie de l'estimation," *Z. Wahrsch. verw. Gebiete*, vol. 65, pp. 181–327, 1983.
- [11] Y. Yang and A. Barron, "Information-theoretic determination of minimax rates of convergence," *Annals of Statistics*, vol. 27, no. 5, pp. 1564–1599, 1999.
- [12] B. Yu, "Assouad, Fano and Le Cam," in *Festschrift for Lucien Le Cam*. Berlin: Springer-Verlag, 1997, pp. 423–435.
- [13] V. N. Vapnik and A. Chervonenkis, *Theory of Pattern Recognition*, 1974, in Russian.
- [14] A. Ehrenfeucht, D. Haussler, M. Kearns, and L. Valiant, "A general lower bound on the number of examples needed for learning," in *Proceedings of the first annual workshop on Computational learning theory*, ser. COLT '88, 1988, pp. 139–154.
- [15] H. U. Simon, "Bounds on the number of examples needed for learning functions," *SIAM J. Comput.*, vol. 26, pp. 751–763, June 1997.
- [16] E. Hazan, A. Agarwal, and S. Kale, "Logarithmic regret algorithms for online convex optimization," *Mach. Learn.*, vol. 69, no. 2-3, pp. 169–192, 2007.
- [17] E. Hazan and S. Kale, "Beyond the regret minimization barrier: an optimal algorithm for stochastic strongly-convex optimization," in *COLT*, 2010.
- [18] A. Juditsky and Y. Nesterov, "Primal-dual subgradient methods for minimizing uniformly convex functions," Universite Joseph Fourier, Tech. Rep., 2010, uRL <http://hal.archives-ouvertes.fr/hal-00508933/>.
- [19] T. Cover and J. Thomas, *Elements of Information Theory*. New York: John Wiley and Sons, 1991.
- [20] L. Le Cam, "Convergence of estimates under dimensionality restrictions," *Annals of Statistics*, vol. 1, no. 1, pp. 38–53, 1973.
- [21] J. Matousek, *Lectures on discrete geometry*. New York: Springer-Verlag, 2002.
- [22] G. Raskutti, M. J. Wainwright, and B. Yu, "Minimax rates of estimation for high-dimensional linear regression over ℓ_q -balls," *IEEE Trans. Information Theory*, vol. 57, no. 10, pp. 6976–6994, October 2011.
- [23] A. Beck and M. Teboulle, "Mirror descent and nonlinear projected subgradient methods for convex optimization," *Operations Research Letters*, vol. 31, no. 3, pp. 167 – 175, 2003.
- [24] G. Rockafellar, *Convex Analysis*. Princeton: Princeton University Press, 1970.
- [25] J. Hiriart-Urruty and C. Lemaréchal, *Convex Analysis and Minimization Algorithms*. New York: Springer-Verlag, 1993, vol. 1.

Alekh Agarwal Alekh Agarwal received the B. Tech degree in Computer Science and Engineering from Indian Institute of Technology Bombay, Mumbai, India in 2007, and an M. A. in Statistics from University of California Berkeley in 2009 where he is currently pursuing a Ph.D. in Computer Science. He received the Microsoft Research Fellowship in 2009 and the Google PhD Fellowship in 2011.

Peter L. Bartlett Peter Bartlett is a professor in the Division of Computer Science and Department of Statistics at the University of California at Berkeley. He is the co-author, with Martin Anthony, of the book *Learning in Neural Networks: Theoretical Foundations*, has edited three other books, and has co-authored many papers in the areas of machine learning and statistical learning theory. He has served as an associate editor of the journals *Machine Learning*, *Mathematics of Control Signals and Systems*, the *Journal of Machine Learning Research*, the *Journal of Artificial Intelligence Research*, and the *IEEE Transactions on Information Theory*, as a member of the editorial boards of *Machine Learning*, the *Journal of Artificial Intelligence Research*, and *Foundations and Trends in Machine Learning*, and as a member of the steering committees of the *Conference on Computational Learning Theory* and the *Algorithmic Learning Theory Workshop*. He has consulted to a number of organizations, including General Electric, Telstra, and SAC Capital Advisors. In 2001, he was awarded the Malcolm McIntosh Prize for Physical Scientist of the Year in Australia, for his work in statistical learning theory. He was a Miller Institute Visiting Research Professor in Statistics and Computer Science at U.C. Berkeley in Fall 2001, and a fellow, senior fellow and professor in the Research School of Information Sciences and Engineering at the Australian National University's Institute for Advanced Studies (1993-2003), and an honorary professor in the School of Information Technology and Electrical Engineering at the University of Queensland. His research interests include machine learning, statistical learning theory, and adaptive control.

Pradeep Ravikumar Pradeep Ravikumar received his BTech in Computer Science and Engineering from the Indian Institute of Technology, Bombay, and his PhD in Machine Learning from the School of Computer Science at Carnegie Mellon University. He was then a postdoctoral scholar in the Department of Statistics at the UC Berkeley. He is now an Assistant Professor in the Department of Computer Science, at the University of Texas at Austin. He is also affiliated with the Division of Statistics and Scientific Computation, and the Institute for Computational Engineering and Sciences at UT Austin. His thesis has received honorable mentions in the ACM SIGKDD Dissertation award and the CMU School of Computer Science Distinguished Dissertation award.

Martin J. Wainwright Martin Wainwright is currently a professor at University of California at Berkeley, with a joint appointment between the Department of Statistics and the Department of Electrical Engineering and Computer Sciences. He received a Bachelor's degree in Mathematics from University of Waterloo, Canada, and Ph.D. degree in Electrical Engineering and Computer Science (EECS) from Massachusetts Institute of Technology (MIT). His research interests include statistical signal processing, coding and information theory, statistical machine learning, and high-dimensional statistics. He has been awarded an Alfred P. Sloan Foundation Fellowship, an NSF CAREER Award, the George M. Sprowls Prize for his dissertation research (EECS department, MIT), a Natural Sciences and Engineering Research Council of Canada 1967 Fellowship, Best Paper Awards from the IEEE Signal Processing Society (2008) and IEEE Communications Society (2010), and several outstanding conference paper awards.