

# No Quantum Speedup over Gradient Descent for Non-Smooth Convex Optimization

**Ankit Garg**

Microsoft Research India, Bangalore, India  
garga@microsoft.com

**Robin Kothari**

Microsoft Quantum and Microsoft Research, Redmond, WA, USA  
robin.kothari@microsoft.com

**Praneeth Netrapalli**

Microsoft Research India, Bangalore, India  
praneeth@microsoft.com

**Suhail Sherif**

Microsoft Research India, Bangalore, India  
Tata Institute of Fundamental Research, Mumbai, India  
suhail.sherif@gmail.com

---

## Abstract

We study the first-order convex optimization problem, where we have black-box access to a (not necessarily smooth) function  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  and its (sub)gradient. Our goal is to find an  $\epsilon$ -approximate minimum of  $f$  starting from a point that is distance at most  $R$  from the true minimum. If  $f$  is  $G$ -Lipschitz, then the classic gradient descent algorithm solves this problem with  $O((GR/\epsilon)^2)$  queries. Importantly, the number of queries is independent of the dimension  $n$  and gradient descent is optimal in this regard: No deterministic or randomized algorithm can achieve better complexity that is still independent of the dimension  $n$ .

In this paper we reprove the randomized lower bound of  $\Omega((GR/\epsilon)^2)$  using a simpler argument than previous lower bounds. We then show that although the function family used in the lower bound is hard for randomized algorithms, it can be solved using  $O(GR/\epsilon)$  quantum queries. We then show an improved lower bound against quantum algorithms using a different set of instances and establish our main result that in general even quantum algorithms need  $\Omega((GR/\epsilon)^2)$  queries to solve the problem. Hence there is no quantum speedup over gradient descent for black-box first-order convex optimization without further assumptions on the function family.

**2012 ACM Subject Classification** Theory of computation  $\rightarrow$  Quantum computation theory; Theory of computation  $\rightarrow$  Convex optimization

**Keywords and phrases** Quantum algorithms, Gradient descent, Convex optimization

**Digital Object Identifier** 10.4230/LIPIcs.ITCS.2021.53

**Related Version** A full version of the paper is available at <https://arxiv.org/abs/2010.01801>.

**Acknowledgements** We thank Sébastien Bubeck, Ronald de Wolf, and András Gilyén for helpful conversations about this work. RK thanks Vamsi Pritham Pingali for many helpful conversations about multivariable calculus. We would also like to thank Matthias Kleinmann for asking us some questions about Lemma 17 that helped us improve its presentation.

## 1 Introduction

The classic gradient descent algorithm, first proposed by Cauchy in 1847, is a popular algorithm for minimizing functions in high-dimensional spaces. For some problems, such as the case of convex function minimization that we consider in this paper, gradient descent



© Ankit Garg, Robin Kothari, Praneeth Netrapalli, and Suhail Sherif;  
licensed under Creative Commons License CC-BY

12th Innovations in Theoretical Computer Science Conference (ITCS 2021).

Editor: James R. Lee; Article No. 53; pp. 53:1–53:20



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

provably converges to the function’s global minimum. For other problems, such as finding good parameters of a deep neural network, gradient descent does not necessarily converge to a global minimum, and yet it has remarkable performance in practice.

Given the algorithm’s popularity, it is interesting to ask if gradient descent can be sped up on a quantum computer. However, it’s not obvious how to formalize this question since it’s not clear what it means for a quantum algorithm to speed up a given classical algorithm. For example, the best known classical algorithm for integer factorization is the general number field sieve (GNFS). Does Shor’s quantum algorithm for integer factorization speed up GNFS, or is it simply a different algorithm that solves the same problem?

One way to formalize the question *Can quantum computers speed up gradient descent?* is to consider a computational problem that is provably solved by gradient descent, and for which gradient descent is optimal among all classical algorithms. We can then ask if quantum algorithms can solve this problem faster than gradient descent. The second condition, that gradient descent is optimal among classical algorithms, is required since otherwise quantum computers would trivially be able to outperform gradient descent by using the best classical algorithm.

Fortunately, there is a canonical optimization task that is solved optimally by gradient descent: convex optimization with black-box first-order oracles. A more thorough introduction to the theory of black-box convex optimization can be found in the textbooks by Nemirovsky and Yudin [22], Nesterov [23, 24], and the monograph by Bubeck [11].

## 1.1 First-order convex optimization

Let’s start with the unconstrained convex minimization problem for a convex function  $f : \mathbb{R}^n \rightarrow \mathbb{R}$ . Here we want to find an  $x \in \mathbb{R}^n$  that’s  $\epsilon$ -close to minimizing the function  $f$ . More precisely, if we let  $x^* := \operatorname{argmin}_{x \in \mathbb{R}^n} f(x)$ , then our goal is to find any  $x \in \mathbb{R}^n$  such that  $f(x) - f(x^*) \leq \epsilon$ .

To obtain algorithms that are very general, this problem is often studied in the setting of black-box optimization. Here we do not assume any particular structure of the function  $f$  (e.g., that  $f$  is a low-degree polynomial), and only assume that we have some efficient method of computing  $f$  by an algorithm or circuit. In other words, we view  $f$  as a black box.

If we only had access to a black-box computing  $f$ , this would be zeroth-order optimization. In first-order optimization, we additionally assume we can also compute the gradient of  $f$ , or more precisely, since the gradient may not exist, we assume we can compute some subgradient of  $f$  (defined in Section 2). We call this oracle the *first-order oracle* and denote it by  $\mathcal{FO}(f)$ . In this work we consider arbitrary convex functions that are not necessarily smooth,<sup>1</sup> and so we cannot assume that the gradient exists. Our goal is to solve the function minimization problem while minimizing the number of calls or queries to the black boxes for  $f$  and some subgradient of  $f$ .

One might wonder why we consider queries to  $f$  and the subgradient of  $f$  to cost the same. This assumption is justified in many practical situations because of the *cheap gradient principle* [16], which says that “the cost to evaluate the gradient  $\nabla f$  is bounded above by a small constant times the cost to evaluate the function itself.” This provably holds in many models of computation; E.g., for arithmetic circuits over  $+$  and  $\times$ , it can be proved that the complexity of computing the gradient is at most 5 times the complexity

---

<sup>1</sup> In the optimization literature, a *smooth function* is a function that is differentiable everywhere in its domain, so the gradient is well defined, and whose gradient has bounded Lipschitz constant.

of computing  $f$  [6]. The conversion of source code computing  $f$  to code computing  $\nabla f$  can often be done automatically in many programming languages, and such methods are called *automatic differentiation* or *algorithmic differentiation* [16]. These same principles essentially carry over to the computation of subgradients [18]. In the quantum setting, there is additional motivation to assume that a function and its gradient cost roughly the same since we can obtain the gradient (or a subgradient) of a function from a black-box computing the function, as shown in a sequence of papers that make increasingly weaker assumptions on the function oracle [17, 15, 13, 3].

Now that we have black-box access to  $f$  and  $\mathcal{FO}(f)$ , we also need a starting point  $x_0 \in \mathbb{R}^n$  to begin our search for a minimum. We require this to be an input, and the complexity will depend on how close this is to  $x^*$ , since otherwise the interesting portion of the function where the minimum is achieved might be hiding in some small corner of  $\mathbb{R}^n$  that we cannot efficiently locate with only black-box access. Since we can easily shift the function by a fixed vector, without loss of generality we assume  $x_0 = \vec{0}$  is the origin. Let the distance between  $x_0 = \vec{0}$  and  $x^*$  be  $R := \|x^*\|$ .<sup>3</sup> For convenience, we will assume that  $R$  is part of the input as well, although this can be relaxed by binary searching for the correct value of  $R$ .

Finally, it is also reasonable that the complexity of our algorithms depend on how quickly  $f$  can change, since the value of  $f$  at some point only constrains its values at nearby points if the function does not change too rapidly. Let  $G$  be an upper bound on the Lipschitz constant of  $f$  (defined in Section 2), and we assume this is part of the input as well.

We are now ready to formally define the first-order convex minimization problem in the black-box setting. We use  $B(x, R) := \{y : \|x - y\| \leq R\}$  to denote an  $\ell_2$ -ball of radius  $R$  around  $x$ .

► **Problem 1** (First-order convex minimization). *Let  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  have Lipschitz constant at most  $G$  on  $B(\vec{0}, R)$ , and let*

$$x^* := \operatorname{argmin}_{x \in B(\vec{0}, R)} f(x). \quad (1)$$

*Then given  $n$ ,  $G$ ,  $R$ , and  $\epsilon > 0$ , the goal is to output a solution  $x \in B(\vec{0}, R)$  such that  $f(x) - f(x^*) \leq \epsilon$  while minimizing the number of queries to  $f$  and  $\mathcal{FO}(f)$ .<sup>4</sup>*

Note that we allow algorithms to query the function and gradient oracles at any point in  $\mathbb{R}^n$  even though the domain we are minimizing over is  $B(\vec{0}, R)$ . This only makes our lower bounds stronger, and the algorithms discussed in this paper never query the oracles outside the domain.

As we discuss in Section 2, although the problem seems to involve 4 parameters, the parameters  $G$ ,  $R$ , and  $\epsilon$  are not independent since we can rescale the input and output spaces of  $f$  and assume  $G = 1$  and  $R = 1$  without loss of generality. Thus any upper or lower bound on the complexity of this problem will be a function of  $n$  and  $GR/\epsilon$ .

## 1.2 Classical algorithms for first-order convex minimization

Gradient descent, or in this case subgradient descent, is a simple algorithm that starts from a point  $x_0$  and takes a small step (governed by a step size  $\eta$ ) in the opposite direction of the subgradient returned at  $x_0$ . Intuitively this brings us closer to the minimum since we are stepping in the direction where  $f$  decreases the most.

<sup>2</sup> If  $x^*$  is not unique, we can let  $R$  be the distance between  $x_0$  and the closest  $x^*$  to it.

<sup>3</sup> Throughout this paper  $\|\cdot\|$  always denotes the standard  $\ell_2$  norm in  $\mathbb{R}^n$  defined as  $\|z\| := \sqrt{\sum_i z_i^2}$ .

<sup>4</sup> For simplicity, we assume that these oracles output real numbers to arbitrarily many bits of precision. Since the main results of this paper are lower bounds, this only makes our results stronger.

We can now describe the performance of subgradient descent for Problem 1. Since this is a constrained optimization problem, we use the projected subgradient descent algorithm, which is subgradient descent with the added step of projecting the current vector back onto the ball  $B(\vec{0}, R)$  after every step.

► **Theorem 2** (Complexity of projected subgradient descent). *The projected subgradient descent algorithm solves Problem 1 using  $(GR/\epsilon)^2$  queries to  $f$  and  $\mathcal{FO}(f)$ .*

A proof of this is included in the full version of this paper. Observe that the query complexity of this algorithm, the number of queries made by the algorithm, is independent of  $n$ .<sup>5</sup> This is quite surprising at first and partly explains why gradient descent and its variants are popular in high-dimensional applications. More generally, we call such algorithms *dimension-independent algorithms*.

There also exist dimension-dependent algorithms for Problem 1 that work well when  $n$  is small. For example, the center of gravity method [11] solves this problem with  $O(n \log(GR/\epsilon))$  queries, which is very reasonable when  $n$  is small (and the algorithm is very efficient in terms of  $\epsilon$ ). In this work we focus on dimension-independent algorithms and assume that  $n$  is polynomially larger than the other parameters in the problem.

When  $n$  is large, we cannot improve over projected subgradient descent (Theorem 2) using any deterministic or randomized algorithm. We reprove the (well known) optimality of this algorithm among deterministic and randomized algorithms. This result is presented in Section 3.

► **Theorem 3** (Randomized lower bound). *For any  $G$ ,  $R$ , and  $\epsilon$ , there exists a family of convex functions  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  with  $n = O((GR/\epsilon)^2)$ , with Lipschitz constant at most  $G$  on  $B(\vec{0}, R)$ , such that any classical (deterministic or bounded-error randomized) algorithm that solves Problem 1 on this function family must make  $\Omega((GR/\epsilon)^2)$  queries to  $f$  or  $\mathcal{FO}(f)$  in the worst case.*

This lower bound on query complexity has been shown in several prior works [22, 26, 12], but we believe our proof is simpler and the dimension  $n$  required in our proof seems to be smaller than that in prior works. Note that while several expositions of gradient descent prove the lower bound for deterministic algorithms, very few sources establish a lower bound against randomized algorithms.

Our lower bound uses the following hard family of functions: For any  $z \in \{-1, +1\}^n$ , let  $f_z(x_1, \dots, x_n) = \max_{i \in [n]} z_i x_i$ ,<sup>6</sup> where  $n = O(1/\epsilon^2)$ . These functions are convex with Lipschitz constant 1. We show that finding an  $\epsilon$ -approximate minimum within  $B(\vec{0}, 1)$  requires  $\Omega(n)$  queries to the oracles. We establish the lower bound by showing that with high probability, every query of a randomized algorithm only reveals  $O(1)$  bits of information about the string  $z$ , but an  $\epsilon$ -approximate solution to this problem allows us to reconstruct the string  $z$ , which has  $n$  bits of information.

### 1.3 Quantum algorithms for first-order convex minimization

We then turn to quantum algorithms for solving Problem 1. At first, it might seem that since gradient descent is a sequential, adaptive algorithm where each step depends on the previous one, there is little hope of quantum algorithms outperforming gradient descent.

<sup>5</sup> Of course, the time complexity of implementing this algorithm will be at least linear in  $n$  since each query to either oracle requires us to manipulate a vector of length  $n$ .

<sup>6</sup> We use  $[n]$  to denote the set of positive integers less than or equal to  $n$ , i.e.,  $[n] := \{1, \dots, n\}$ .

On the other hand, consider the hard family of functions described above that witnesses the classical randomized lower bound in Theorem 3. While this is hard for classical algorithms, we show in Section 3.1 that there is a quantum algorithm that solves the problem on this family obtaining a quadratic speedup over any classical algorithm (and in particular, over gradient descent).

► **Theorem 4** (Quantum algorithm for classically hard function family). *There is a quantum algorithm that solves Problem 1 on the class of functions that appear in the classical lower bound of Theorem 3 using  $O(GR/\epsilon)$  queries to the oracle for  $f$ .*

Notably, unlike most quadratic speedups in quantum computing, the source of this quadratic speedup is not Grover’s algorithm or amplitude amplification. Theorem 4 uses Belovs’ quantum algorithm for learning symmetric juntas, which is constructed by exhibiting a feasible solution to the dual semidefinite program of the negative-weights adversary bound [7].

Now that we have shown a quadratic quantum speedup on a family of instances known to be hard for classical algorithms, there is some hope that quantum algorithms may provide some speedup for the general first-order convex minimization problem. Alas, our next result (established in Section 4), which is our main result, shows that this is not the case, and quantum algorithms cannot in general yield a speedup over classical algorithms for first-order convex minimization.

► **Theorem 5** (Quantum lower bound). *For any  $G$ ,  $R$ , and  $\epsilon$ , there exists a family of convex functions  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  with  $n = \tilde{O}((GR/\epsilon)^4)$ , with Lipschitz constant at most  $G$  on  $B(\vec{0}, R)$ , such that any quantum algorithm that solves Problem 1 with high probability on this function family must make  $\Omega((GR/\epsilon)^2)$  queries to  $f$  or  $\mathcal{FO}(f)$  in the worst case.*

Our lower bound uses ideas from the lower bound against parallel randomized algorithms recently established by Bubeck, Jiang, Lee, Li, and Sidford [12].

At a high level, the hard family of functions used in the randomized lower bound does not work for quantum algorithms because although classical algorithms can only learn  $O(1)$  bits of information per query, quantum algorithms can make queries in superposition and learn a little information about many bits simultaneously. We remedy this by choosing a new family of functions in which with high probability, no matter what query the quantum algorithm makes, the oracle’s response is essentially the same. This allows us to control what the quantum algorithm learns per query, but now the instance is more complicated and the quantum algorithm learns  $O(n)$  bits of information per query. Since the final output of the algorithm is a vector in  $\mathbb{R}^n$ , we cannot use the argument used before that simply compared the information learned per query to the total information that needs to be learned. Instead we use the venerable hybrid argument [8] to control what the quantum algorithm learns and show that it cannot find an  $\epsilon$ -approximate solution to the minimization problem.

## 1.4 Related work

Classically, there is a long history of the study of oracle complexity (also known as black-box complexity or query complexity) for deterministic and randomized algorithms for non-smooth and smooth convex optimization. The setting considered in this paper, first-order convex optimization, where the algorithm has query access to the function value and the gradient, is very well studied. This topic is too vast to survey here, but we refer the reader to [22, 23, 24, 11] for more information about upper and lower bounds that can be shown in this setting.

There also has been work in the classical parallel setting, where in each round the algorithm is allowed to query polynomially many points and the goal is to minimize the number of rounds [21, 4, 14, 12]. Our work is most closely related to this setting and borrows many ideas from these works. Although quantum algorithms and parallel classical algorithms are incomparable in power, the constructions used to thwart parallel classical algorithms in these papers also help with showing quantum lower bounds.

In the quantum setting, there has been some work on convex optimization in the oracle model. There is also work on quantum gradient descent not in the oracle model. For example, one situation studied is where the dimension  $n$  of the optimization space is very large and the vectors are encoded in quantum states of dimension  $\log n$ . See [25, 19] and the references therein for more information. Another setting is the work on semidefinite programming, an important special case of convex optimization, but these algorithms exploit the specific structure of semidefinite programs [10, 2, 9, 1] and are not directly related to our work.

While in the classical setting, in general, a function value oracle is weaker than a gradient oracle, this is not the case in the quantum setting. Given a function value oracle, one can get a gradient oracle quite efficiently (with an  $\tilde{O}(1)$  overhead) [17, 15, 3, 13]. A similar result also holds for simulating a separation oracle given a membership oracle for convex bodies [3, 13]. As discussed before, our focus in this paper is to see if quantum algorithms can outperform classical algorithms when given a function oracle and gradient oracle since in many relevant settings, gradient computation is cheap in the classical case as well.

The most related works are the papers by Chakrabarti, Childs, Li, and Wu [13] and van Apeldoorn, Gilyén, Gribling, and de Wolf [3]. These papers establish very similar results so we cover them together. These papers study the problem of black-box convex optimization, and their results are phrased in the slightly different language of membership and separation oracles, but this is not the main difference between their work and our work. Indeed, it is possible to recast our problem in their setting (see the discussion in the introduction in [3] for how to do this). The main difference is that their algorithms are dimension-dependent and have complexities that depend on  $n$ , whereas we're working in the parameter regime where  $n$  is large and so we seek algorithms that are independent of  $n$ .

Specifically, [13] and [3] consider the problem of minimizing a linear function over a convex body given via a membership or separation oracle. A membership oracle for a convex body tells us whether a given point  $x$  is in the convex body and a separation oracle in addition when  $x$  is not in the body outputs a hyperplane that separates  $x$  from the convex body. Classically, the problem of outputting an  $\epsilon$ -approximate solution can be solved with  $O(n^2 \text{polylog}(\cdot))$  queries to a membership oracle [20], where we are suppressing polylogarithmic dependence on several parameters (including  $\epsilon$ ). These two papers show a quantum algorithm that makes only  $O(n \text{polylog}(\cdot))$  membership queries. The key technical component of this is a construction of a separation oracle from a membership oracle with only polylogarithmic overhead. To do this, they first show how to obtain an approximate subgradient oracle from a function oracle with only polylogarithmic overhead.

There are also several lower bounds shown in these papers. In [3], the authors prove that quantum algorithms do not give any advantage over classical algorithms in the setting where we are not given a point inside the convex body to start with. This setting is not directly comparable to our setting, as far as we are aware. In the setting where we do know a point inside the convex body, which is very similar to our setting, [3, 13] prove a lower bound of  $\Omega(\sqrt{n})$ , which is quadratically worse than their algorithm. While, in general, their results are incomparable to our results, one specific comparison to our results is that [13, Theorem 3.3] essentially shows a  $\tilde{\Omega}(\min\{GR/\epsilon, \sqrt{n}\})$  lower bound on the number of oracle

calls to a function value oracle for the setting in Problem 1.<sup>7</sup> Note that this is quadratically worse than our tight lower bound (Theorem 5) in the dimension-independent setting (i.e., when the dimension  $n$  is large compared to  $GR/\epsilon$ ).

## 1.5 Paper organization and summary of contributions

We first present some preliminaries on convex optimization in Section 2. In Section 3 we reprove the lower bound for randomized algorithms (Theorem 3) using a simpler argument compared to prior works. In Section 3.1, we show that quantum algorithms can solve the hard instance from Theorem 3 faster than randomized algorithms, obtaining a quadratic speedup (Theorem 4). In Section 4, we present a different hard instance and show our main result that quantum algorithms cannot obtain any speedup over gradient descent for the first-order convex optimization problem (Theorem 5). We conclude with open problems in Section 5.

## 2 Convex optimization preliminaries

As described in the introduction, we are interested in approximately minimizing a convex function  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  on some closed convex set  $\mathcal{K} \subseteq \mathbb{R}^n$ . A function  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  is convex if for all  $x, y \in \mathbb{R}^n$  and  $t \in [0, 1]$ ,

$$tf(x) + (1-t)f(y) \geq f(tx + (1-t)y). \quad (2)$$

A set  $\mathcal{K} \subseteq \mathbb{R}^n$  is convex if the line segment joining two points in  $\mathcal{K}$  is also contained in  $\mathcal{K}$ . We will consider convex sets of bounded size, and specifically let  $2R$  be the diameter of  $\mathcal{K}$ , i.e.,

$$\max_{x, y \in \mathcal{K}} \|x - y\| \leq 2R, \quad (3)$$

where  $\|z\| := \sqrt{\sum z_i^2}$  is the Euclidean norm.

It turns out that the query complexity of first-order convex optimization depends only on  $R$  no matter how complicated the set  $\mathcal{K}$  happens to be. However, to obtain an algorithm with efficient time complexity we require that the set  $\mathcal{K}$  be simple enough that we can efficiently implement a projection operator for  $\mathcal{K}$ . This means given any  $y \in \mathbb{R}^n$ , we can efficiently compute  $\mathcal{P}_{\mathcal{K}}(y) \in \mathcal{K}$ , which satisfies  $\|\mathcal{P}_{\mathcal{K}}(y) - y\| = \min_{z \in \mathcal{K}} \|z - y\|$ . Since the main result of this paper is a lower bound, our lower bound is stronger if shown for a simple convex set  $\mathcal{K}$ . So throughout this paper we work with the set  $\mathcal{K} = B(\vec{0}, R)$ , the  $\ell_2$ -ball of radius  $R$  around the origin.

In the model of first-order black-box optimization, we have access to a black-box that computes the function  $f$  on any input  $x \in \mathbb{R}^n$ . In addition to this, we also have a *first-order oracle*,  $\mathcal{FO}(f)$ , which when queried at any point  $x \in \mathbb{R}^n$  returns some vector  $g_x \in \mathbb{R}^n$  that satisfies for all  $y \in \mathbb{R}^n$ ,

$$f(y) \geq f(x) + \langle g_x, y - x \rangle. \quad (4)$$

Since  $f$  is convex, it is known that such a vector  $g_x$  exists for all  $x \in \mathbb{R}^n$  [23]. Any vector  $g_x$  satisfying (4) is called a subgradient of  $f$  at  $x$ , and the set of all subgradients at  $x$  is called the subdifferential at  $x$  and denoted by  $\partial f(x)$ . If  $f$  is differentiable at  $x$  then  $g_x$  is unique and equal to  $\nabla f(x)$ , the gradient of  $f$  at  $x$ , defined as

<sup>7</sup> This is equivalent to our setting, where we have a function value and gradient oracle, due to their results.

$$\nabla f(x) := \left( \frac{\partial f(x)}{\partial x_1}, \dots, \frac{\partial f(x)}{\partial x_n} \right). \quad (5)$$

We will say that the function  $f$  has Lipschitz constant at most  $G$  in  $\mathcal{K}$  if  $\|g_x\| \leq G$  for every  $x \in \mathcal{K}$ .<sup>8</sup>

As described in Problem 1, we are interested in algorithms that take as inputs the parameters  $n$ ,  $G$ ,  $R$ , and  $\epsilon > 0$ , and have access to  $f$  and a first-order oracle  $\mathcal{FO}(f)$ , and output  $x \in B(\vec{0}, R)$  such that  $f(x) - f(x^*) \leq \epsilon$ , where  $x^* := \operatorname{argmin}_{x \in B(\vec{0}, R)} f(x)$ .

In the quantum setting, we have quantum analogues of these oracles. There is a straightforward generalization of any oracle to the quantum setting, which makes the classical oracle reversible and then allows queries in superposition to this oracle. This quantum generalization of the oracle is justified by the fact that if we had a classical circuit or algorithm computing a function  $f$ , then it is possible in a completely black-box manner to construct the quantum oracle corresponding to  $f$ , and this oracle would then support superposition queries. We discuss quantum oracles in more detail in Section 4, but for now it is sufficient to consider them as computing the same functions as the classical oracles, except that they can additionally be queried in superposition.

Note that it is sufficient to consider the special case of the problem where  $G = R = 1$ . While this seems like a special case, given an  $f$  and  $\mathcal{K}$  with Lipschitz constant  $G$ , radius  $R$ , and optimization accuracy  $\epsilon$ , we can instead minimize  $\hat{f}(x) := \frac{1}{GR}f(Rx)$  over  $\hat{\mathcal{K}} := \mathcal{K}/R$ , which have Lipschitz constant and radius 1 up to an accuracy of  $\frac{\epsilon}{GR}$ . So we consider  $G = R = 1$  without loss of generality, or for general  $G$  and  $R$ , the complexity must be a function of  $GR/\epsilon$ .

The query complexity of an algorithm that solves Problem 1 is the maximum number of oracle calls it makes for fixed values of  $n$ ,  $G$ ,  $R$ , and  $\epsilon$ , where the maximum is taken over all convex functions  $f$  with Lipschitz constant at most  $G$ , and all first order oracles  $\mathcal{FO}(f)$  for  $f$  (i.e., the algorithm must work for any choice of first-order oracle that correctly outputs some subgradient of  $f$  at  $x$ ). As discussed, the query complexity must be a function of  $n$  and  $GR/\epsilon$ . Furthermore, since we're interested in dimension-independent algorithms, we study algorithms that only depend on  $GR/\epsilon$  and not on  $n$ .

Given a class of algorithms, such as deterministic, randomized, or quantum algorithms, the query complexity of first-order Lipschitz convex optimization for that class of algorithms is the minimum query complexity of any algorithm in that class that solves Problem 1.

As we will see in the next section, the randomized query complexity of this problem (and hence the deterministic query complexity) is at least  $\Omega((GR/\epsilon)^2)$  in the dimension-independent setting.

### 3 Randomized Lower Bound

In this section, we prove a lower bound for randomized first-order methods for non-smooth convex optimization, restated here for convenience:

<sup>8</sup> This is slightly different from the usual definition of the Lipschitz constant where we would say  $f$  is  $G$ -Lipschitz in  $\mathcal{K}$  if for all  $x, y \in \mathcal{K}$ ,  $|f(x) - f(y)| \leq G\|x - y\|$ . Our definition is the same as requiring the function  $f$  to be  $G$ -Lipschitz according to this definition in an open set that contains  $\mathcal{K}$ .



► **Theorem 3** (Randomized lower bound). *For any  $G$ ,  $R$ , and  $\epsilon$ , there exists a family of convex functions  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  with  $n = O((GR/\epsilon)^2)$ , with Lipschitz constant at most  $G$  on  $B(\vec{0}, R)$ , such that any classical (deterministic or bounded-error randomized) algorithm that solves Problem 1 on this function family must make  $\Omega((GR/\epsilon)^2)$  queries to  $f$  or  $\mathcal{FO}(f)$  in the worst case.*

This randomized lower bound is known and multiple proofs can be found in the literature [22, 26]. Our proof is elementary and we did not find it written anywhere, although it is conceptually similar to the one in [22], and so we include it here for completeness. Our proof also has the dimension  $n = \Theta(1/\epsilon^2)$ , without any log factors, which is the best possible. As far as we are aware, the previous proofs (for randomized algorithms) required larger dimension than this.<sup>9</sup> As we will see later, the family of instances used is also interesting because we can get a quantum speedup for it, because of which we have to look at other instances to prove the quantum lower bound.

We can now define the family of convex functions used in the lower bound. For any  $\epsilon > 0$ , we set  $n = \lfloor .9/\epsilon^2 \rfloor$  and look at the following class of functions.

► **Definition 6.** *Let  $z \in \{-1, +1\}^n$ . Let  $f_z : \mathbb{R}^n \rightarrow \mathbb{R}$  be defined as*

$$f_z(x_1, \dots, x_n) = \max_{i \in [n]} z_i x_i. \quad (6)$$

Each such function is convex since it is a maximum of convex functions [23, Theorem 3.1.5]. Note that if  $f_z(x) = z_i x_i$  for some  $i \in [n]$ , then  $z_i e_i$  is a subgradient of  $f_z$  at  $x$  (since  $f_z(x) + \langle z_i e_i, y - x \rangle = z_i y_i \leq f_z(y)$ ). Hence the function is 1-Lipschitz. We can also see that within the unit ball the function is minimized at the point

$$x^* = \frac{-1}{\sqrt{n}} \sum_{i \in [n]} z_i e_i, \quad (7)$$

and  $f_z(x^*) = -1/\sqrt{n}$ . Clearly given  $x^*$  we can recover  $z$  from it. We now show  $z$  can even be recovered from an  $\epsilon$ -approximate minimum of  $f_z$ .

► **Lemma 7.** *Let  $x$  be such that  $f_z(x) - f_z(x^*) \leq \epsilon$ . Then we can recover  $z \in \{-1, +1\}^n$  from  $x \in \mathbb{R}^n$ .*

**Proof.** Let  $s_x \in \{-1, +1\}^n$  be the vector with  $(s_x)_i = \text{sign}(x_i)$ , where  $\text{sign}(a) = +1$  if  $a \geq 0$  and  $\text{sign}(a) = -1$  otherwise. We claim that  $z = -s_x$ . Toward a contradiction, if  $(s_x)_i \neq -z_i$  for some  $i$ , then  $(s_x)_i = z_i$ , since these only take values in  $\{-1, +1\}$ . In this case,  $x_i$  and  $z_i$  agree in sign, and hence  $f_z(x) \geq z_i x_i \geq 0$ . Since  $\epsilon < 1/\sqrt{n}$  (because of our choice of  $n$  above) the point  $x$  cannot satisfy  $f_z(x) - f_z(x^*) \leq \epsilon$ . ◀

Since this function is not differentiable everywhere, for our lower bound we need to specify the behavior of the subgradient oracle on all inputs. The function is not differentiable only at  $x \in \mathbb{R}^n$  where the maximum is achieved at multiple indices. In this case, the subgradient oracle responds as if the maximum was achieved on the smallest such index  $i$ , i.e., it responds with  $z_i e_i$ . Note that for this function, querying the subgradient oracle allows us to simulate a call to the function oracle as well, since the response is  $z_i e_i$  for the index  $i$  that achieves the maximum, so the function evaluates to  $z_i x_i$  at that point, which we can compute since we know  $x$ . So we can assume without loss of generality that an algorithm only queries the subgradient oracle.

<sup>9</sup> The result of [26] requires dimension roughly  $1/\epsilon^8$ . We also believe the result in [22] requires a larger dimension than our claim.

Now that the problem is fully specified, we will show that any randomized optimization algorithm using the function oracle and this subgradient oracle will require  $\Omega(n)$  queries in order to solve Problem 1 with a constant probability of success.

The following will be the crux of the lower bound. Let  $I \subseteq [n]$ . We say a distribution  $\mathcal{D}$  over  $\{-1, +1\}^n$  is  $I$ -fixed if for  $z \sim \mathcal{D}$  the random variable  $z_I$  is fixed and  $z_{\bar{I}}$  is uniform over  $\{-1, +1\}^{\bar{I}}$ .

► **Lemma 8.** *Let  $z$  be distributed according to an  $I$ -fixed distribution. Let  $x$  be an arbitrary query made to the  $f_z$  oracle. After one query to the subgradient oracles, the conditional distribution on  $z$  given the answer is  $I'$ -fixed with  $I \subseteq I'$  and  $\mathbb{E}[|I'|] \leq |I| + 2$ .*

**Proof.** Let  $x$  be the algorithm's query. The index  $i$  that achieves the maximum in the definition of  $f_z(x)$  can be computed as follows. Let  $i_1, \dots, i_n$  be the ordering of the indices 1 to  $n$  in decreasing order of  $|x_i|$ , with ties broken with the natural ordering on integers. The oracle outputs  $f_z(x) = z_{i_j} x_{i_j}$  and chooses the subgradient  $z_{i_j} e_{i_j}$  where  $j$  is the smallest index for which  $x_{i_j}$  agrees in sign with  $z_{i_j}$ , and if no such index exists, then  $j = n$ .

Since  $f_z(x)$  can be computed given the subgradient  $z_{i_j} e_{i_j}$ , the only information obtained from a query is the prefix  $\{z_{i_k}\}_{k \leq j}$ . In other words, if the subgradient oracle responds with  $z_{i_j} e_{i_j}$ , then we have learned that for all indices  $k \leq j$ , we must have  $\text{sign}(x_{i_k}) = -z_{i_k}$ , but we have not learned any more since the oracle's output does not depend on the bits of  $z$  with index  $i_k$  with  $k > j$ . After this query, we know the bits  $z_{i_k}$  with  $k \leq j$ , but conditioned on these, the distribution on the remaining bits of  $z$  continues to be uniform. This is an  $I'$ -fixed distribution with  $I' = I \cup \{i_k\}_{k \leq j}$ . Intuitively,  $I'$  cannot be much larger than  $I$  since an index  $i_k$  is part of this set only if the algorithm correctly guessed the sign of  $z_{i_k}$  for this index and all indices with a smaller value of  $k$ . Since the initial distribution  $z$  was uniformly at random outside of  $I$  and  $x$  is fixed, the probability of correctly guessing the first index (according to the  $i_j$  ordering) that was not fixed is  $1/2$ , the probability of guessing the first two is  $1/4$  and so on. Thus the expected number of new entries fixed by one query is  $\sum_{k=1}^{n \setminus |I|} k \cdot \frac{1}{2^k} \leq 2$ . ◀

We can use this to show establish the final claim.

► **Lemma 9.** *Let  $z$  be sampled uniformly at random from  $\{-1, +1\}^n$ . If a randomized algorithm  $\mathcal{A}$  outputs an  $x$  with  $f_z(x) - f_z(x^*) \leq \epsilon$  with probability at least  $2/3$ , then its query complexity is at least  $n/3 - 1$ .*

**Proof.** When  $\mathcal{A}$  outputs a point  $x$ , we will require it to also query the oracle at  $x$  to see if it is indeed  $\epsilon$ -optimal. This can increase its query complexity by at most one. Let the query complexity of this modified  $\mathcal{A}$  be  $t$ . Whenever  $\mathcal{A}$  does output an  $\epsilon$ -optimal point, Lemma 7 implies that the conditional distribution on  $z$  is  $[n]$ -fixed. For each  $i \in [0, \dots, t]$ , let  $I_i$  be the random variable such that the distribution on  $z$  after  $i$  queries of  $\mathcal{A}$  is  $I_i$ -fixed (Lemma 8 implies that after any sequence of queries it will be an  $I$ -fixed distribution for some  $I$ ). Since  $z$  is sampled uniformly at random from  $\{-1, +1\}^n$ ,  $I_0 = \emptyset$ . And since we want the algorithm to succeed with probability at least  $2/3$ ,  $\mathbb{E}[|I_t|] \geq 2n/3$ .

However,  $|I_t| = \sum_{i=1}^t |I_i| - |I_{i-1}|$ , and it is a simple consequence of Lemma 8 that  $\mathbb{E}[|I_i| - |I_{i-1}|] \leq 2$  for all  $i$ . So by the linearity of expectation,  $\mathbb{E}[|I_t|] \leq 2t$  and hence  $t \geq n/3$ . ◀

This proves a lower bound of  $\Omega(1/\epsilon^2)$  on the randomized query complexity of first-order convex minimization for a function with  $G = R = 1$ . As noted earlier, this is without loss of generality and implies the more general bound in Theorem 3.

### 3.1 Quantum speedup

In this section we discuss Theorem 4, restated for convenience:

► **Theorem 4** (Quantum algorithm for classically hard function family). *There is a quantum algorithm that solves Problem 1 on the class of functions that appear in the classical lower bound of Theorem 3 using  $O(GR/\epsilon)$  queries to the oracle for  $f$ .*

The quantum speedup for the above class of functions relies on Belovs' quantum algorithm for Combinatorial Group Testing [7]. Belovs showed that given access to an oracle making OR queries to an  $n$ -bit string, the  $n$ -bit string can be learned in  $O(\sqrt{n})$  quantum queries. The proof of Theorem 4 then follows by noting that for any  $S \subseteq [n]$ , one can find out if there is an  $i \in S$  such that  $z_i = 1$  by querying  $x = \frac{1}{\sqrt{n}} \sum_{i \in S} e_i$ . Hence within  $O(\sqrt{n})$  queries, the entire string  $z \in \{-1, +1\}^n$  can be learned. This proof is provided in more detail in the full version of the paper.

A similar problem is also studied in [13], which they solve using a quantum algorithm for the search with wildcards problem, which is a special case of the combinatorial group testing problem studied by Belovs.

Note that this quantum algorithm is also essentially optimal for this problem and it is not hard to show an  $\Omega(\sqrt{n}/\log n)$  lower bound for quantum algorithms. A similar lower bound is shown in [13, Theorem 3.3], and a simpler sketch of such a lower bound is also given in the full version.

## 4 Quantum lower bound

In this section, we show that for any  $\epsilon$ , there exists a 1-Lipschitz family of functions such that any quantum algorithm that solves Problem 1 on the unit ball must make  $\frac{1}{100\epsilon^2}$  queries. In other words, there is no quantum first-order convex optimization algorithm that always outperforms the classical gradient descent algorithm described in Theorem 2. The function we will use was introduced by Nemirovsky and Yudin [22]. To show the quantum lower bound, we adapt to the quantum setting the lower bound strategy of Bubeck et al. [12] in the model of parallel algorithms.

We restate the main result proved in this section for convenience:

► **Theorem 5** (Quantum lower bound). *For any  $G$ ,  $R$ , and  $\epsilon$ , there exists a family of convex functions  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  with  $n = \tilde{O}((GR/\epsilon)^4)$ , with Lipschitz constant at most  $G$  on  $B(\vec{0}, R)$ , such that any quantum algorithm that solves Problem 1 with high probability on this function family must make  $\Omega((GR/\epsilon)^2)$  queries to  $f$  or  $\mathcal{FO}(f)$  in the worst case.*

We start by first proving a qualitatively similar, but simpler result with a larger value of  $n = \tilde{O}((GR/\epsilon)^6)$  in Section 4.4. If we only care about the optimality of gradient descent in the dimension-independent setting, this lower bound is sufficient. But if we also want to understand the trade-off between dimension-independent and dimension-dependent algorithms, then we would like to show this lower bound with as small a value of  $n$  as we can. In the full version of the paper, we improve the lower bound to achieve the value of  $n$  stated in this theorem.

### 4.1 Function family and basic properties

We start by defining the family of functions  $\mathcal{F} = \{f : \mathbb{R}^n \rightarrow \mathbb{R}\}$  that we use. The function family  $\mathcal{F}$  depends on the dimension  $n$  and two other parameters  $k$  and  $\gamma$ . Since the function family we choose depends on  $\epsilon$ , the parameters  $n$ ,  $k$ , and  $\gamma$  will be functions of  $\epsilon$ . Our choice of  $n$ ,  $k$ , and  $\gamma$  will become clear later, but for now we simply choose them as follows. Let

## 53:12 No Quantum Speedup over Gradient Descent for Non-Smooth Convex Optimization

$$k := \frac{1}{100\epsilon^2} \implies \epsilon = \frac{1}{10\sqrt{k}} \quad \text{and} \quad \gamma := \frac{1}{10k^{3/2}} = 100\epsilon^3. \quad (8)$$

We choose  $n$  such that it satisfies

$$\gamma \geq 8\sqrt{\frac{\log n}{n}} \implies n := O\left(\frac{\log(1/\epsilon)}{\epsilon^6}\right) = \tilde{O}\left(\frac{1}{\epsilon^6}\right). \quad (9)$$

The discussion before Lemma 11 explains the choice of  $k$  and the discussion after Lemma 12 explains the choice of  $\gamma$ . For the dimension  $n$ , see the discussion at the beginning of Section 4.2.

We now define the function family for these specific choices of  $n$ ,  $k$ , and  $\gamma$ .

► **Definition 10** (Hard function family). *Let  $\mathcal{V} = \{(v_1, \dots, v_k) \mid \forall i, j, \in [k], \langle v_i, v_j \rangle = \delta_{ij}\}$  be the set of all  $k$ -tuples of orthonormal vectors in  $\mathbb{R}^n$ . Let the family of functions  $\mathcal{F} = \{f_V\}_{V \in \mathcal{V}}$  be defined as*

$$f_{(v_1, v_2, \dots, v_k)}(x) := \max_{i \in [k]} \{g_V^{(i)}(x)\}, \quad \text{where} \quad g_V^{(i)}(x) := \langle v_i, x \rangle + (k - i)\gamma\|x\|. \quad (10)$$

We will show that any quantum algorithm that solves Problem 1 on the functions in this family must make  $k$  queries. As we will prove, informally what happens is each query of the quantum algorithm to the gradient oracle only reveals a single direction  $v_i$  to the algorithm. In fact, with very high probability the vectors are revealed in order, so that the algorithm first learns  $v_1$ , then  $v_2$ , and so on. As we will show in Lemma 12, any  $\epsilon$ -optimal solution must overlap significantly with all  $v_i$ , and thus any quantum algorithm must make  $k$  queries. Since we want to show an  $\Omega(1/\epsilon^2)$  bound, we choose  $k$  to be a small multiple of  $1/\epsilon^2$ , which explains our choice for  $k$  in Equation (8).

We now establish some basic properties of these functions.

► **Lemma 11** (Properties of  $f_V$ ). *For any  $V \in \mathcal{V}$ , let  $f_V$  and  $g_V^{(i)}$  be as in Definition 10. Then  $f_V$  is convex with Lipschitz constant at most  $1 + k\gamma \leq 2$  on  $B(\vec{0}, 1)$ , and*

$$\text{for } x \neq \vec{0}, \quad \nabla g_V^{(i)}(x) = v_i + (k - i)\gamma x / \|x\|, \quad \text{and} \quad (11)$$

$$\text{for } x = \vec{0}, \quad \partial g_V^{(i)}(\vec{0}) = \{v_i + (k - i)\gamma u \mid u \in B(\vec{0}, 1)\}, \quad \text{and} \quad (12)$$

$$\text{for any } x, \quad \partial f_V(x) = \text{ConvexHull}(\{u \in \partial g_V^{(i)}(x) \mid g_V^{(i)}(x) = f_V(x)\}), \quad (13)$$

where the convex hull of a set of vectors is the set of all convex combinations of vectors in the set. Lastly, for any  $\alpha > 0$ ,  $f_V(\alpha x) = \alpha f_V(x)$  and  $\partial f_V(\alpha x) = \partial f_V(x)$ .

**Proof.** For all  $V \in \mathcal{V}$ ,  $f_V : \mathbb{R}^n \rightarrow \mathbb{R}$  is convex. This follows because linear functions and norms are convex functions [23, Example 3.1.1], and the sum or maximum of convex functions is convex [23, Theorem 3.1.5].

Let us now compute the subgradients of  $g_V^{(i)}(x) = \langle v_i, x \rangle + (k - i)\gamma\|x\|$ . The linear function  $\langle v_i, x \rangle$  is differentiable and its gradient is simply  $v_i$ . The Euclidian norm  $\|x\|$  is differentiable everywhere except at  $x = \vec{0}$ . At  $x \neq \vec{0}$ , the gradient of  $\|x\|$  is  $x/\|x\|$  and at  $x = 0$ , the set of subgradients is  $B(\vec{0}, 1)$  [23, Example 3.1.5]. We also know that  $\partial(\alpha_1 f_1(x) + \alpha_2 f_2(x)) = \alpha_1 \partial f_1(x) + \alpha_2 \partial f_2(x)$  [23, Lemma 3.1.9], which gives us the expressions for the subgradients of  $g_V^{(i)}$ .

For a function that is the maximum of functions  $g_V^{(i)}$ , we know that the set of subgradients is simply the convex hull of subgradients of those  $g_V^{(i)}$  which achieve the maximum at the given point  $x$  [23, Lemma 3.1.10].

The Lipschitz constant of a function is the maximum norm of any subgradient of the function. Since any vector in  $\partial g_V^{(i)}$  has norm  $1 + k\gamma$ , and any vector in  $\partial f_V$  is the convex combination of vectors with norm at most  $1 + k\gamma$ , the Lipschitz constant of  $f_V$  is at most  $1 + k\gamma \leq 2$ .

Finally, it is easy to see from the definition of  $f_V$  that for  $\alpha > 0$ ,  $f_V(\alpha x) = \alpha f(x)$  since each term in the max gets multiplied by  $\alpha$ . For  $\partial f_V(\alpha x)$ , note that this is a convex combination of  $\partial g_V^{(i)}(\alpha x)$ , and these do not depend on  $\alpha$ . ◀

For convenience we work with this family of functions with Lipschitz constant at most 2 instead of 1, which doesn't change the asymptotic bounds since we could just divide every function  $f_V$  by 2.

The last property essentially says that querying the function or its subgradient on a scalar multiple of a vector  $x$  gives us only as much information as querying it on  $x$ . Thus we can assume that an algorithm only queries the oracles within the unit ball without loss of generality.

Now let us discuss the vector  $x^* \in B(\vec{0}, 1)$  that minimizes  $f_V(x)$  and vectors that  $\epsilon$ -approximately solve the minimization problem. First note that if  $\gamma$  were equal to 0, then the function would simply be  $\max_{i \in [k]} \langle v_i, x \rangle$ , which requires us to minimize the component of  $x$  in  $k$  different directions subject to it being a unit vector. The solution to this is simply  $\frac{-1}{\sqrt{k}} \sum_i v_i$ . Now  $-1/\sqrt{k} = -10\epsilon$ , so the overlap of  $x$  with each direction  $v_i$  is a large multiple of  $\epsilon$ . So even an  $\epsilon$ -approximate solution must have reasonable overlap with each of the vectors  $v_i$ . Specifically, each overlap must be at least  $-9\epsilon$ . Now in our function  $f_V$  the term  $\gamma$  is not 0, but that term at most perturbs the function by  $k\gamma = \epsilon$ , which again is much smaller than  $10\epsilon$ , and thus even approximate solutions must have significant overlaps with all  $v_i$ . We formalize these properties below.

► **Lemma 12** (Properties of the minimum). *For any  $V \in \mathcal{V}$ , let  $f_V : \mathbb{R}^n \rightarrow \mathbb{R}$  be the function in Definition 10 and let  $x^* := \operatorname{argmin}_{x \in B(\vec{0}, 1)} f_V(x)$ . Then  $f_V(x^*) \leq -9\epsilon$ . Furthermore, any  $x \in \mathbb{R}^n$  that satisfies  $|f_V(x) - f_V(x^*)| \leq \epsilon$  must satisfy for all  $i \in [k]$ ,  $\langle v_i, x \rangle \leq -8\epsilon$ .*

**Proof.** Consider the vector  $y = \frac{-1}{\sqrt{k}} \sum_{i \in [k]} v_i$ . This is a vector in  $B(\vec{0}, 1)$ , satisfying  $f_V(y) \leq \frac{-1}{\sqrt{k}} + (k-1)\gamma \leq \frac{-1}{\sqrt{k}} + k\gamma = -10\epsilon + \epsilon = -9\epsilon$ , because we have  $10\epsilon = \frac{1}{\sqrt{k}}$  and  $k\gamma = \frac{1}{10\sqrt{k}} = \epsilon$ . Thus  $f_V(x^*) \leq f_V(y) \leq -9\epsilon$ .

Now consider any vector  $x$  with  $|f_V(x) - f_V(x^*)| \leq \epsilon$ , which implies  $f_V(x) \leq -8\epsilon$ . If  $\langle v_i, x \rangle > -8\epsilon$  for any  $i \in [k]$ , then  $f_V(x) \geq \langle v_i, x \rangle + (k-i)\gamma \|x\| > -8\epsilon$ , which is a contradiction. ◀

This result crucially uses the relation between  $\gamma$  and  $k$  and because we want  $k\gamma$  to be a constant factor (say 10) smaller than  $\sqrt{1/k}$ , this informs our choice of  $\gamma$  in Equation (8). Our choice of  $n$  in Equation (9) will be discussed in the next section.

## 4.2 Probabilistic facts about the function family

So far all the properties we have discussed of our function family hold for any  $V \in \mathcal{V}$ , but now we want to talk about a hard distribution over such functions. Specifically we want to talk about choosing a uniformly random (according to the Haar measure)  $V$  from the infinite set  $\mathcal{V}$ . It is easy to see how to sample a random  $V$  once we can sample unit vectors from a subspace. We start by choosing  $v_1$  to be a Haar random unit vector from  $\mathbb{R}^n$ , let  $v_2$  be a Haar random unit vector from  $\operatorname{span}(v_1)^\perp$ , and so on, until  $v_k$  is a Haar random unit vector in  $\operatorname{span}(v_1, v_2, \dots, v_{k-1})^\perp$ .

We can now discuss what determines our choice of  $n$ . By construction, the family of functions  $\mathcal{F}$  has the property that if the input vector  $x$  has equal inner product with all vectors  $v_i$ , then the maximum will be achieved uniquely on the first term  $i = 1$  because the additive term  $(k - i)\gamma\|x\|$  is largest for  $i = 1$ . Now what we want to ensure is that this property holds even when  $x$  does not have equal inner product with all  $v_i$ , but  $x$  is chosen uniformly at random from  $B(\vec{0}, 1)$ . Or equivalently, we want this property to hold when  $x$  is fixed, but the set  $V$  is chosen uniformly at random.

In either case, the inner product of  $x$  with a random unit vector  $v$  will be a random variable with mean 0 due to symmetry. But the expected value of  $|\langle v, x \rangle|^2$  for a random unit vector  $v$  is  $1/n$ , and in fact it will be tightly concentrated around  $1/n$ . The following proposition follows from [5, Lemma 2.2].

► **Proposition 13.** *Let  $x \in B(\vec{0}, 1)$ . Then for a random unit vector  $v$ , and all  $c > 0$ ,*

$$\Pr_v(|\langle x, v \rangle| \geq c) \leq 2e^{-nc^2/2}. \quad (14)$$

We choose  $\gamma$  so that it is very unlikely (polynomially small in  $n$ ) that the maximum is not achieved at  $i = 1$ . From Proposition 13, we see that the probability of any  $|\langle v_i, x \rangle|^2$  being larger than a constant multiple of  $\log n/n$  is inverse polynomially small. So it is sufficient to take  $\gamma^2$  to be a large constant multiple of  $\log n/n$  as in Equation (9).

In our lower bound we will need a slightly stronger result. We can show that if the vectors  $v_1, \dots, v_{t-1}$  are fixed (and hence known to the algorithm), and the remaining vectors  $v_t, \dots, v_k$  are chosen uniformly at random such that the set of vectors  $\{v_1, \dots, v_k\}$  is orthonormal, then the maximum will be achieved in the set  $[t]$  with high probability. This generalizes the previous claim, which is the case of  $t = 1$ , where none of the vectors were fixed.

► **Lemma 14 (Most probable argmax).** *Let  $1 \leq t \leq k$  be integers and  $\{v_1, \dots, v_{t-1}\}$  be a set of orthonormal vectors. Let  $\{v_t, \dots, v_k\}$  be chosen uniformly at random so that the set  $\{v_1, \dots, v_k\}$  is orthonormal. Then*

$$\forall x \in B(\vec{0}, 1) : \Pr_{v_t, \dots, v_k} \left( \max_{i \in [k]} \langle v_i, x \rangle + (k - i)\gamma\|x\| \neq \max_{i \in [t]} \langle v_i, x \rangle + (k - i)\gamma\|x\| \right) \leq \frac{1}{n^7}. \quad (15)$$

**Proof.** Let  $E_x$  denote the event whose probability we want to upper bound. Since  $E_x$  and  $E_{\alpha x}$ , for any  $\alpha \in [0, 1]$ , are the same event, we can assume without loss of generality that  $\|x\| = 1$ . If event  $E_x$  occurs, then it must hold that

$$\max_{i \in \{t+1, \dots, k\}} \langle v_i, x \rangle + (k - i)\gamma > \max_{i \in [t]} \langle v_i, x \rangle + (k - i)\gamma \geq \langle v_t, x \rangle + (k - t)\gamma. \quad (16)$$

We want to show that this event is very unlikely. To do so, let  $F_x$  be the event that for all  $i \in \{t, \dots, k\}$ ,  $\langle v_i, x \rangle \in [-\frac{\gamma}{2}, +\frac{\gamma}{2}]$ . Note that if  $F_x$  occurs, then the terms in the max are in decreasing order, and we have

$$\langle v_t, x \rangle + (k - t)\gamma \geq \langle v_{t+1}, x \rangle + (k - t - 1)\gamma \geq \dots \geq \langle v_{k-1}, x \rangle + \gamma \geq \langle v_k, x \rangle, \quad (17)$$

which contradicts Equation (16). Thus if  $E_x$  holds then the complement of  $F_x$ ,  $\bar{F}_x$  must hold, which means  $\Pr(E_x) \leq \Pr(\bar{F}_x)$ . So let us show that  $F_x$  is very likely.

The event  $\bar{F}_x$  holds only if there exists an  $i \in \{t, \dots, k\}$  such that  $\langle v_i, x \rangle \notin [-\frac{\gamma}{2}, +\frac{\gamma}{2}]$ . We can upper bound this probability for any particular  $i \in \{t, \dots, k\}$  using Proposition 13 and the fact that  $v_i$  is chosen uniformly at random from an  $n - t + 1$ -dimension ball. This probability is at most  $2e^{-(n-t+1)\gamma^2/8} = 2e^{-(n-t+1) \cdot 8 \frac{\log n}{n}} \leq 2 \cdot 2^{-8 \log n} = 2/n^8$ , with the inequality holding because  $n > 4t$ . The probability that this happens for any  $i$  is at most  $(k - t + 1) \leq k$  times this probability, by the union bound. Using the fact that  $2k < n$ , we get that  $\Pr(E_x) \leq \Pr(\bar{F}_x) < 1/n^7$ . ◀

Finally, we show that even if we knew the vectors  $v_1, \dots, v_{k-1}$ , we cannot guess a vector  $x$  that is an  $\epsilon$ -approximate solution to our problem, because it won't have enough overlap with  $v_k$ , which is unknown. In other words, for an algorithm to output an  $\epsilon$ -optimal solution, it essentially must know the entire set  $V$ .

► **Lemma 15** (Cannot guess  $x^*$ ). *Let  $k > 0$  be an integer and  $\{v_1, \dots, v_{k-1}\}$  be a set of orthonormal vectors. Let  $v_k$  be chosen uniformly at random from  $\text{span}(v_1, \dots, v_{k-1})^\perp$  and let  $V = (v_1, \dots, v_k)$ . Then*

$$\forall x \in B(\vec{0}, 1) : \Pr_{v_k}(f_V(x) - f_V(x^*) \leq \epsilon) \leq 2e^{-\Omega(k^2)}. \quad (18)$$

**Proof.** From Lemma 12, we know that an  $\epsilon$ -optimal solution  $x$  must satisfy  $\langle v_k, x \rangle \leq -8\epsilon$ . But  $v_k$  is chosen uniformly at random from the space  $\text{span}(v_1, \dots, v_{k-1})^\perp$  and any vector  $x \in B(\vec{0}, 1)$  projected to that space also has length at most 1. So from Proposition 13 we know that for any  $x \in B(\vec{0}, 1)$ ,

$$\Pr_{v_k}(\langle v_k, x \rangle \leq -8\epsilon) \leq \Pr_{v_k}(|\langle v_k, x \rangle| \geq 8\epsilon) \leq 2e^{-32(n-k+1)\epsilon^2} \leq 2e^{-\Omega(k^2)}. \quad (19)$$

### 4.3 Quantum query model

We now formally define the quantum query model in our setting. In the usual quantum query model the set of allowed queries is finite, whereas in our setting it is natural to allow the quantum algorithm to query the oracles at any point  $x \in \mathbb{R}^n$ . Due to Lemma 11, it is sufficient to allow the algorithm to query any  $x \in B(\vec{0}, 1)$ , but this is still a continuous space of queries, and hence a query vector could be a superposition over infinitely many states. Instead of formalizing this notion of quantum algorithms, we allow the algorithm to make discrete queries only, but to arbitrarily high precision. The reader is encouraged to not get bogged down by details and to think of the registers as storing the real values that they ideally should, but in the rest of this section we define these algorithms more carefully so that all the spaces involved are finite and well defined. This formalization is not specific to the quantum setting and is done classically as well if we do not want to manipulate real numbers as atomic objects.

All the real numbers that appear will be represented using some  $b$  bits of precision, where  $b$  can be chosen by the algorithm. The reader should imagine  $b$  being arbitrarily large, say exponentially larger than all the parameters involved in the problem, so that the inaccuracy involved by using this representation is negligible. Then the algorithm represents the input  $x \in B(\vec{0}, 1)$  using  $b$  bits of precision per coordinate. The oracle's response will also use  $b$  bits of precision per real number. For a given choice of  $b$ , the quantum algorithm will have some probability of success of solving the problem at hand. We then define the success probability of quantum algorithms that make  $q$  queries by taking a supremum over all  $b$  of  $q$ -query algorithms that solve the problem.

We can now define the oracles more precisely. Classically, the function oracle for a function  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  would simply implement the map  $x \mapsto f(x)$ , where we represent each entry of  $x$  and the output  $f(x)$  using  $b$  bits, so  $x \in \{0, 1\}^{bn}$  and  $f(x) \in \{0, 1\}^b$ . Let's say we have a classical circuit that implements this map using  $G$  gates, say over the gate set of AND, OR, and NOT gates. Then it is easy to construct, in a completely black-box way, a quantum circuit using  $O(G)$  gates (say over the gate set of Hadamard, CNOT, and T) that performs the unitary  $U|x\rangle|y\rangle = |x\rangle|y \oplus f(x)\rangle$ , for every  $x \in \{0, 1\}^{bn}$ , and  $y \in \{0, 1\}^b$ . This is why it is standard to assume that the quantum oracle corresponding to the classical map  $x \mapsto f(x)$  is a unitary that performs  $U|x\rangle|y\rangle = |x\rangle|y \oplus f(x)\rangle$ . We apply the same construction for

the  $\mathcal{FO}(f)$  oracle to get the quantum analogue of the classical map  $x \mapsto g_x$ , where  $g_x$  is some subgradient of  $f$  at  $x$ . Lastly, for convenience we will combine both the function and subgradient oracle into one oracle that when queried with  $x$  returns  $f(x)$  and a subgradient at  $x$ . Since our function family is parameterized by  $V \in \mathcal{V}$ , we call this oracle  $O_V$ .

Let  $\mathcal{A}$  be a quantum query algorithm that makes  $q$  queries.  $\mathcal{A}$  is described by a sequence of unitaries  $U_q O_V U_{q-1} O_V U_{q-2} O_V \cdots U_1 O_V U_0$  applied to an initial state, say  $|0\rangle$ . We assume that the output of  $\mathcal{A}$ , which is a vector  $x$ , is determined by measuring the first  $n$  registers storing real numbers using  $b$  bits.

#### 4.4 Lower bound

We can now prove the quantum lower bound. Let  $\mathcal{A}$  be a  $k - 1$  query quantum algorithm that solves Problem 1 on all the functions  $f_V$  for  $V \in \mathcal{V}$ . Due to Lemma 11, we can assume that the algorithm only queries the oracles with vectors  $x \in B(\vec{0}, 1)$ . We also need to describe the behavior of the subgradient oracle on inputs where the subgradient is not unique. On such inputs  $x$ , the subgradient is not unique because several indices  $i \in [k]$  simultaneously achieve the maximum in  $f_V(x)$ . In this case, the subgradient will answer as if the smallest index  $i$  in this set achieved the maximum. Now let  $\mathcal{A}$  be described by the sequence of unitaries  $U_{k-1} O_V U_{k-2} O_V \cdots O_V U_1 O_V U_0$  acting on the starting state  $|0\rangle$ . Let this sequence of unitaries be called  $A$ . Then the final state of the algorithm is  $A|0\rangle$ .

Recall that we defined  $f_V(x) = \max_{i \in [k]} \{g_V^{(i)}(x)\}$ . Let us also define functions  $f_V^{(j)}$  where the maximization is only over the first  $j$  indices instead of all  $k$  indices. Specifically, let  $f_V^{(j)} := \max_{i \in [j]} \{g_V^{(i)}(x)\}$ . We previously defined the oracle  $O_V$  as corresponding to the function  $f_V$ . Let  $O_V^{(j)}$  be the oracle corresponding to the functions  $f_V^{(j)}$ .

Now we define a sequence of unitaries starting with  $A_0 = A$  as follows:

$$\begin{aligned} A_0 &:= U_{k-1} O_V U_{k-2} O_V \cdots O_V U_1 O_V U_0 \\ A_1 &:= U_{k-1} O_V U_{k-2} O_V \cdots O_V U_1 O_V^{(1)} U_0 \\ A_2 &:= U_{k-1} O_V U_{k-2} O_V \cdots O_V^{(2)} U_1 O_V^{(1)} U_0 \\ &\vdots \\ A_{k-1} &:= U_{k-1} O_V^{(k-1)} U_{k-2} O_V^{(k-2)} \cdots O_V^{(2)} U_1 O_V^{(1)} U_0 \end{aligned} \tag{20}$$

We want to show that the algorithm  $A_0$  does not solve our problem. To do so, we will employ the hybrid argument, in which we show that the output of the algorithm  $A_i$  and  $A_{i+1}$  is close, and thus the output of  $A_0$  and  $A_{k-1}$  is close. Finally, we argue that the algorithm  $A_{k-1}$  does not solve our problem because the oracles in the algorithm do not know  $v_k$ . Let us first establish these two claims.

► **Lemma 16** ( $A_{k-1}$  does not solve the problem). *Let  $\mathcal{A}$  be a  $k - 1$  query algorithm and let  $A_{k-1}$  be defined as above. Let  $p_V$  be the probability distribution over  $x \in B(\vec{0}, 1)$  obtained by measuring the output state  $A_{k-1}|0\rangle$ . Then  $\Pr_{V, x \sim p_V} (f_V(x) - f_V(x^*) \leq \epsilon) \leq 2e^{-\Omega(k^2)}$ .*

**Proof.** We want to show that the probability (over the random choice of  $V$  and the internal randomness of the algorithm) that  $A_{k-1}$  outputs an  $x$  that satisfies  $f(x) - f(x^*) \leq \epsilon$  is very small.

Let us establish the claim for any fixed choice of  $v_1, \dots, v_{k-1}$ , since if the claim holds for any fixed choice of these vectors, then it also holds for any probability distribution over them. For a fixed choice of vectors, this claim is just  $\Pr_{v_k, x \sim p_V} (f_V(x) - f_V(x^*) \leq \epsilon) \leq 2e^{-\Omega(k^2)}$ .



Now since the algorithm  $A_{k-1}$  only has oracles  $O_V^{(i)}$  for  $i < k$ , the probability distribution  $p_V$  only depends on  $v_1, \dots, v_{k-1}$ . Since these are fixed, this is just a fixed distribution  $p$ . So we can instead establish our claim for all  $x \in B(\vec{0}, 1)$ , which will also establish it for any distribution.

So what we need to establish is that for any  $x \in B(\vec{0}, 1)$ ,  $\Pr_{v_k}(f_V(x) - f_V(x^*) \leq \epsilon) \leq 2e^{-\Omega(k^2)}$ , which is exactly what we showed in Lemma 15. ◀

► **Lemma 17** ( $A_t$  and  $A_{t-1}$  have similar outputs). *Let  $\mathcal{A}$  be a  $k-1$  query algorithm and let  $A_t$  for  $t \in [k-1]$  be the unitaries defined in Equation (20). Then*

$$\mathbb{E}_V(\|A_t|0\rangle - A_{t-1}|0\rangle\|^2) \leq \frac{4}{n^7}. \quad (21)$$

**Proof.** From the definition of the unitaries in Equation (20) and the unitary invariance of the spectral norm, we see that  $\|A_t|0\rangle - A_{t-1}|0\rangle\| = \|(O_V^{(t)} - O_V)U_{t-1}O_V^{(t-1)} \dots O_V^{(1)}U_0|0\rangle\|$ . Let us again prove the claim for any fixed choice of vectors  $v_1, \dots, v_{t-1}$ , which will imply the claim for any distribution over those vectors. Once we have fixed these vectors, the state  $U_{t-1}O_V^{(t-1)} \dots O_V^{(1)}U_0|0\rangle$  is a fixed state, which we can call  $|\psi\rangle$ . Thus our problem reduces to showing for all quantum states  $|\psi\rangle$ ,

$$\mathbb{E}_{v_t, \dots, v_k}(\|(O_V^{(t)} - O_V)|\psi\rangle\|^2) \leq \frac{4}{n^7}. \quad (22)$$

Now we can write an arbitrary quantum state as  $|\psi\rangle = \sum_x \alpha_x |x\rangle|\phi_x\rangle$ , where  $x$  is the query made to the oracle, and  $\sum_x |\alpha_x|^2 = 1$ . Thus the LHS of Equation (22) is equal to

$$\mathbb{E}_{v_t, \dots, v_k} \left( \left\| \sum_x \alpha_x (O_V^{(t)} - O_V)|x\rangle|\phi_x\rangle \right\|^2 \right) \leq \sum_x |\alpha_x|^2 \mathbb{E}_{v_t, \dots, v_k} \left( \|(O_V^{(t)} - O_V)|x\rangle|\phi_x\rangle\|^2 \right). \quad (23)$$

Since  $|\alpha_x|^2$  defines a probability distribution over  $x$ , we can again upper bound the right hand side for any  $x$  instead. Since  $O_V^{(t)}$  and  $O_V$  behave identically for some inputs  $x$ , the only nonzero terms are those where the oracles respond differently, which can only happen if  $f_V^{(t)}(x) \neq f_V(x)$ . When the response is different, we can upper bound  $\|(O_V^{(t)} - O_V)|x\rangle|\phi_x\rangle\|^2$  by 4 using the triangle inequality. Thus for any  $x \in B(\vec{0}, 1)$ , we have

$$\mathbb{E}_{v_t, \dots, v_k} \left( \|(O_V^{(t)} - O_V)|x\rangle|\phi_x\rangle\|^2 \right) \leq 4 \Pr_{v_t, \dots, v_k} (f_V^{(t)}(x) \neq f_V(x)) \leq 4/n^7, \quad (24)$$

where the second inequality follows from Lemma 14. The first inequality requires more explanation. It is based on the claim that if  $f_V^{(t)}(x) = f_V(x)$  then on such inputs  $x$ , the oracles  $O_V^{(t)}$  and  $O_V$  must behave identically. In other words, for such an input  $x$ , for any  $|\phi_x\rangle$ , we have  $O_V^{(t)}|x\rangle|\phi_x\rangle = O_V|x\rangle|\phi_x\rangle$ . The oracle responds with the function value and the subgradient, and we have already assumed that the function values are equal, so we only need to show that if  $f_V^{(t)}(x) = f_V(x)$ , then the subgradient oracle's response is also identical for both functions. Since  $f_V^{(t)}(x) = f_V(x)$ , the maximum in the definition of  $f_V$  was achieved by some index in the first  $t$  indices (and possibly other indices as well). But our definition of the subgradient oracle says that if multiple indices achieve the maximum, the subgradient oracle will respond as if the smallest index achieved the maximum. Hence the response of the subgradient oracle for the two functions will also be identical. ◀

Finally we can put these two lemmas together to prove our lower bound.

► **Lemma 18** ( $\mathcal{A}$  does not solve the problem). *Let  $\mathcal{A}$  be a  $k - 1$  query algorithm. Let  $p_V$  be the probability distribution over  $x \in B(\vec{0}, 1)$  obtained by measuring the output state  $A|0\rangle$ . Then  $\Pr_{V, x \sim p_V}(f_V(x) - f_V(x^*) \leq \epsilon) \leq \frac{1}{\text{poly}(n)}$ .*

**Proof.** Let  $P_V$  be the projection operator that projects a quantum state  $|\psi\rangle$  onto the space spanned by vectors  $|x\rangle$  for  $x$  such that  $f_V(x) - f_V(x^*) \leq \epsilon$ . Then  $\|P_V A|0\rangle\|^2 = \Pr_{x \sim p_V}(f_V(x) - f_V(x^*) \leq \epsilon)$ . We know from Lemma 16 that  $\mathbb{E}_V(\|P_V A_{k-1}|0\rangle\|^2) \leq 2e^{-\Omega(k^2)}$ . We prove our upper bound on the probability by showing that it is approximately the same as  $\mathbb{E}_V(\|P_V A_{k-1}|0\rangle\|^2)$ .

Lemma 17 states that for all  $1 \leq t < k$ ,  $\mathbb{E}_V(\|A_t|0\rangle - A_{t-1}|0\rangle\|^2) \leq \frac{4}{n^7}$ . Using telescoping sums and the Cauchy-Schwarz inequality, we see that

$$\mathbb{E}_V(\|A_{k-1}|0\rangle - A|0\rangle\|^2) \leq \mathbb{E}_V\left(\left(\sum_{t \in [k-1]} \|A_t|0\rangle - A_{t-1}|0\rangle\|\right)^2\right) \quad (25)$$

$$\leq \mathbb{E}_V\left(\sum_{t \in [k-1]} \|A_t|0\rangle - A_{t-1}|0\rangle\|^2\right) \left(\sum_{t \in [k-1]} 1^2\right) \leq \frac{4k}{n^7} \cdot k. \quad (26)$$

For all  $V$ ,

$$\| \|P_V A_{k-1}|0\rangle\| - \|P_V A|0\rangle\| \| \leq \|P_V A_{k-1}|0\rangle - P_V A|0\rangle\| \leq \|A_{k-1}|0\rangle - A|0\rangle\| \quad (27)$$

and so

$$\mathbb{E}_V(\left(\|P_V A_{k-1}|0\rangle\| - \|P_V A|0\rangle\|\right)^2) \leq \frac{4k^2}{n^7}. \quad (28)$$

We want an upper bound on  $\mathbb{E}_V(\|P_V A|0\rangle\|^2 - \|P_V A_{k-1}|0\rangle\|^2)$ , which is no larger than  $2\mathbb{E}_V(\|P_V A|0\rangle\| - \|P_V A_{k-1}|0\rangle\|)$  since  $\|P_V A|0\rangle\| + \|P_V A_{k-1}|0\rangle\| \leq 2$ . We get such a bound by applying Jensen's inequality to Equation (28):  $\mathbb{E}_V(\|P_V A|0\rangle\| - \|P_V A_{k-1}|0\rangle\|) \leq 2k/\sqrt{n^7}$ , and so  $\mathbb{E}_V(\|P_V A|0\rangle\|^2 - \|P_V A_{k-1}|0\rangle\|^2) \leq 4k/\sqrt{n^7}$ .

We can now use linearity of expectation and upper bound our required probability as

$$\Pr_{V, x \sim p_V}(f_V(x) - f_V(x^*) \leq \epsilon) = \mathbb{E}_V(\|P_V A|0\rangle\|^2) \leq 2e^{-\Omega(k^2)} + 4k/\sqrt{n^7} \leq \frac{1}{\text{poly}(n)}. \quad (29)$$

Note that this establishes a statement similar to Theorem 5, except with a polynomially larger value of  $n$ . This result is sufficient to establish the optimality of gradient descent in the dimension-independent setting. To reduce the value of  $n$ , we modify the function above using a *wall function* introduced by [12]. The details of this modification and the resulting quantitatively better lower bound can be found in the full version of the paper.

## 5 Open problems

We showed that in the black-box setting, no quantum algorithm can beat gradient descent in general, in the dimension-independent regime. Here are some interesting questions left open by our work:

1. We showed in Theorem 4 that the class of functions used in the randomized lower bound can be solved faster with quantum queries. Is there a more interesting class of functions on which we can achieve a quantum speedup?

2. Can the quantum lower bound in Section 4 be made to work using the simpler class of functions  $f_V(x) = \max_i \langle v_i, x \rangle$ , which is our function with  $\gamma = 0$ ? If so, this might also decrease the dimension  $n$  required.
3. Can we establish tight quantum lower bounds in the parameter regime where dimension-dependent algorithms outperform gradient descent? When  $1/\epsilon$  is a large polynomial in  $n$ , the complexity of gradient descent is also a large polynomial in  $n$ , but a dimension-dependent algorithm such as the center of gravity method [11] yields an  $O(n \log n)$  upper bound. Can we establish an  $\tilde{\Omega}(n)$  lower bound in this regime? The function used in the randomized lower bound in Section 3 yields an  $\tilde{\Omega}(\sqrt{n})$  lower bound and this is the best bound we are aware of. This is essentially the same as the problem left open by [13, 3], but phrased in the language of membership and separation oracles.
4. What can we say about other standard settings in convex optimization beyond first-order non-smooth convex optimization? Other natural settings include assuming the function is smooth, having the ability to query a prox oracle instead of a subgradient oracle, etc. Can quantum algorithms provide a speedup in the black-box model in these settings over the respective best classical algorithms in that setting?

---

## References

- 1 Joran van Apeldoorn and András Gilyén. Improvements in Quantum SDP-Solving with Applications. In *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, volume 132 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 99:1–99:15. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2019. doi:10.4230/LIPIcs.ICALP.2019.99.
- 2 Joran van Apeldoorn, András Gilyén, Sander Gribling, and Ronald de Wolf. Quantum SDP-solvers: Better upper and lower bounds. In *58th Annual Symposium on Foundations of Computer Science (FOCS 2017)*, October 2017. doi:10.1109/focs.2017.44.
- 3 Joran van Apeldoorn, András Gilyén, Sander Gribling, and Ronald de Wolf. Convex optimization using quantum oracles. *Quantum*, 4:220, January 2020. doi:10.22331/q-2020-01-13-220.
- 4 Eric Balkanski and Yaron Singer. Parallelization does not accelerate convex optimization: Adaptivity lower bounds for non-smooth convex minimization. *arXiv preprint*, 2018. arXiv:1808.03880.
- 5 Keith Ball. An elementary introduction to modern convex geometry. In Silvio Levy, editor, *Flavors of geometry*, volume 31, pages 1–58. Cambridge University Press, 1997. URL: <http://library.msri.org/books/Book31/files/ball.pdf>.
- 6 Walter Baur and Volker Strassen. The complexity of partial derivatives. *Theoretical Computer Science*, 22(3):317–330, 1983. doi:10.1016/0304-3975(83)90110-X.
- 7 Aleksandrs Belovs. Quantum algorithms for learning symmetric juntas via adversary bound. In *Proceedings of the 2014 IEEE 29th Conference on Computational Complexity (CCC 2014)*, CCC '14, page 22–31, 2014. doi:10.1109/CCC.2014.11.
- 8 Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. *SIAM Journal on Computing*, 26(5):1510–1523, 1997. doi:10.1137/S0097539796300933.
- 9 Fernando G. S. L. Brandão, Amir Kalev, Tongyang Li, Cedric Yen-Yu Lin, Krysta M. Svore, and Xiaodi Wu. Quantum SDP Solvers: Large Speed-Ups, Optimality, and Applications to Quantum Learning. In *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, volume 132 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 27:1–27:14. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2019. doi:10.4230/LIPIcs.ICALP.2019.27.

- 10 Fernando G.S.L. Brandão and Krysta M. Svore. Quantum speed-ups for solving semidefinite programs. In *58th Annual Symposium on Foundations of Computer Science (FOCS 2017)*, October 2017. doi:10.1109/focs.2017.45.
- 11 Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Found. Trends Mach. Learn.*, 8(3–4):231–357, November 2015. doi:10.1561/22000000050.
- 12 Sébastien Bubeck, Qijia Jiang, Yin Tat Lee, Yuanzhi Li, and Aaron Sidford. Complexity of highly parallel non-smooth convex optimization. In *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*, pages 13900–13909, 2019. URL: <http://papers.nips.cc/paper/9541-complexity-of-highly-parallel-non-smooth-convex-optimization>.
- 13 Shouvanik Chakrabarti, Andrew M. Childs, Tongyang Li, and Xiaodi Wu. Quantum algorithms and lower bounds for convex optimization. *Quantum*, 4:221, January 2020. doi:10.22331/q-2020-01-13-221.
- 14 Jelena Diakonikolas and Cristóbal Guzmán. Lower bounds for parallel and randomized convex optimization. In *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99, pages 1132–1157. PMLR, 2019. URL: <http://proceedings.mlr.press/v99/diakonikolas19c.html>.
- 15 András Gilyén, Srinivasan Arunachalam, and Nathan Wiebe. Optimizing quantum optimization algorithms via faster quantum gradient computation. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '19*, page 1425–1444, 2019. doi:10.1137/1.9781611975482.87.
- 16 Andreas Griewank and Andrea Walther. *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation, Second Edition*. Other Titles in Applied Mathematics. Society for Industrial and Applied Mathematics, 2008. URL: <https://books.google.com/books?id=xoiLaRxcBEC>.
- 17 Stephen P. Jordan. Fast quantum algorithm for numerical gradient estimation. *Phys. Rev. Lett.*, 95:050501, July 2005. doi:10.1103/PhysRevLett.95.050501.
- 18 Sham M Kakade and Jason D Lee. Provably correct automatic sub-differentiation for qualified programs. In *Advances in Neural Information Processing Systems 31 (NeurIPS 2018)*, pages 7125–7135, 2018. URL: <http://papers.nips.cc/paper/7943-provably-correct-automatic-sub-differentiation-for-qualified-programs>.
- 19 Iordanis Kerenidis and Anupam Prakash. Quantum gradient descent for linear systems and least squares. *Physical Review A*, 101(2):022316, 2020. doi:10.1103/PhysRevA.101.022316.
- 20 Yin Tat Lee, Aaron Sidford, and Santosh S. Vempala. Efficient convex optimization with membership oracles. In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet, editors, *Conference On Learning Theory, COLT 2018, Stockholm, Sweden, 6-9 July 2018*, volume 75 of *Proceedings of Machine Learning Research*, pages 1292–1294. PMLR, 06–09 July 2018. URL: <http://proceedings.mlr.press/v75/lee18a.html>.
- 21 A. Nemirovski. On parallel complexity of nonsmooth convex optimization. *Journal of Complexity*, 10(4):451–463, 1994. doi:10.1006/jcom.1994.1025.
- 22 Arkadii Nemirovsky and David Borisovich Yudin. *Problem complexity and method efficiency in optimization*. Wiley, 1983.
- 23 Yurii Nesterov. *Introductory Lectures on Convex Optimization*. Springer US, 2004. doi:10.1007/978-1-4419-8853-9.
- 24 Yurii Nesterov. *Lectures on convex optimization*, volume 137. Springer, 2018. doi:10.1007/978-3-319-91578-4.
- 25 Patrick Rebentrost, Maria Schuld, Leonard Wossnig, Francesco Petruccione, and Seth Lloyd. Quantum gradient descent and Newton’s method for constrained polynomial optimization. *New Journal of Physics*, 21(7):073023, 2019. doi:10.1088/1367-2630/ab2a9e.
- 26 Blake Woodworth and Nathan Srebro. Lower bound for randomized first order convex optimization. *arXiv preprint*, 2017. arXiv:1709.03594.