



Lower bounds for finding stationary points II: first-order methods

Yair Carmon¹  · John C. Duchi² · Oliver Hinder³ · Aaron Sidford³

Received: 13 December 2017 / Accepted: 6 September 2019 / Published online: 20 September 2019
© Springer-Verlag GmbH Germany, part of Springer Nature and Mathematical Optimization Society 2019

Abstract

We establish lower bounds on the complexity of finding ϵ -stationary points of smooth, non-convex high-dimensional functions using first-order methods. We prove that deterministic first-order methods, even applied to arbitrarily smooth functions, cannot achieve convergence rates in ϵ better than $\epsilon^{-8/5}$, which is within $\epsilon^{-1/15} \log \frac{1}{\epsilon}$ of the best known rate for such methods. Moreover, for functions with Lipschitz first and second derivatives, we prove that no deterministic first-order method can achieve convergence rates better than $\epsilon^{-12/7}$, while ϵ^{-2} is a lower bound for functions with only Lipschitz gradient. For *convex* functions with Lipschitz gradient, accelerated gradient descent achieves a better rate, showing that finding stationary points is easier given convexity.

Keywords Non-convex optimization · Information-based complexity · Dimension-free rates · Gradient methods · Accelerated gradient descent

OH was supported by the PACCAR INC fellowship. YC and JCD were partially supported by the SAIL-Toyota Center for AI Research, NSF-CAREER Award 1553086, and a Sloan Foundation Fellowship in Mathematics. YC was partially supported by the Stanford Graduate Fellowship and the Numerical Technologies Fellowship. AS was supported by the National Science Foundation (CCF-1844855).

✉ Yair Carmon
yairc@stanford.edu

John C. Duchi
jduchi@stanford.edu

Oliver Hinder
ohinder@stanford.edu

Aaron Sidford
sidford@stanford.edu

¹ Department of Electrical Engineering, Stanford University, Stanford, CA 94305, USA

² Departments of Statistics and Electrical Engineering, Stanford University, Stanford, CA 94305, USA

³ Department of Management Science and Engineering, Stanford University, Stanford, CA 94305, USA

1 Introduction

We study the oracle complexity of finding approximate stationary points of a smooth function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, that is, a point x such that

$$\|\nabla f(x)\| \leq \epsilon. \quad (1)$$

Part I of this series [10] establishes the complexity of finding an ϵ -stationary point (1) for algorithms that, at a query point x , have access to all derivatives of f . In contrast, this paper focuses on *first-order methods*, which only query function values and gradients. To keep our presentation concise this paper only touches briefly on the most relevant research; see Part I [10] for a more detailed review of related work.

First-order methods are important in large-scale optimization for many reasons. Perhaps the two most salient are that each iteration is often inexpensive, and that on many problems, the number of iterations grows slowly (or not at all) with the problem dimension d . From a theoretical perspective, the latter property is captured by *dimension-free* convergence rates, where the worst case iteration count depends polynomially on the desired accuracy and measures of function regularity but has no explicit dependence on d . In non-convex optimization problems, regularity often comes by assuming bounded function value at the initial point $x^{(0)}$, i.e. $f(x^{(0)}) - \inf_x f(x) \leq \Delta$ for some $\Delta > 0$, and that ∇f is L_1 -Lipschitz continuous. Under these conditions, classical gradient descent finds an ϵ -stationary point in $2L_1\Delta\epsilon^{-2}$ iterations [23], a dimension-free guarantee, and this is unimprovable for this class [10].

Developing first-order methods for finding stationary points of non-convex functions with improved dimension-free rates of convergence is an area of active research [1, 2, 8, 9]. Under the additional assumption of Lipschitz second derivatives, we [9] and Agarwal et al. [1] propose randomized first-order methods with nearly dimension free rate $\epsilon^{-7/4} \log \frac{d}{\epsilon}$ (ignoring other problem-dependant constants). In a later paper [8], we propose a deterministic accelerated gradient-based method with complexity $\epsilon^{-7/4} \log \frac{1}{\epsilon}$, and under the further assumption that f has Lipschitz third derivatives, we show the same method attains rates of $\epsilon^{-5/3} \log \frac{1}{\epsilon}$. This raises the main question we address in this paper: how much further can we improve this ϵ dependence, and what Lipschitz continuity assumptions are necessary?

1.1 Our contributions

In Table 1 we summarize our results, along with corresponding known upper bounds. We establish lower bounds on the *worst-case oracle complexity* of finding ϵ -stationary points, where algorithms may access f only through queries to an information oracle that returns the value and some number of (or potentially all) derivatives of f at the queried point. A lower bound T_ϵ means that for every algorithm A , there exists a function f in the allowed function class (e.g. functions with $f(x^{(0)}) - \inf_x f(x) \leq \Delta$

Table 1 The number of iterations required to find ϵ -stationary points of high dimensional functions

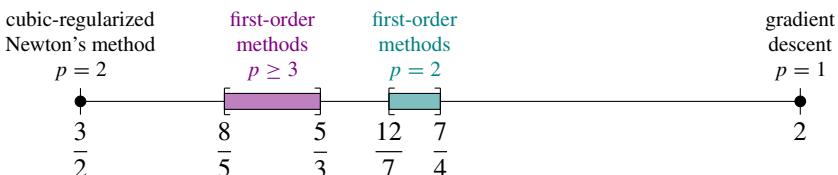
Oracle	f has Lipschitz	Upper bound	Lower bound	Gap
General	p th-order derivative	$O(\epsilon^{-(p+1)/p})$ [5]	$\Omega(\epsilon^{-(p+1)/p})$ Part I	$O(1)$
First-order	Gradient and Hessian	$\tilde{O}(\epsilon^{-7/4})$ [8]	$\Omega(\epsilon^{-12/7})$ Theorem 2	$\tilde{O}(\epsilon^{-1/28})$
First-order	q th derivative $\forall q \leq p, p \geq 3$	$\tilde{O}(\epsilon^{-5/3})$ [8]	$\Omega(\epsilon^{-8/5})$ Theorem 2	$\tilde{O}(\epsilon^{-1/15})$
First-order	Gradient + f convex	$\tilde{O}(\epsilon^{-1})$ Proposition 2	$\Omega(\epsilon^{-1})$ Theorem 1	$\tilde{O}(1)$

The bounds apply for functions f with $f(x^{(0)}) - \inf_x f(x) \leq O(1)$. The first column indexes the type of oracle access: general (all derivatives) or first-order (function value and gradient). In the first row, deterministic p th-order methods achieve the upper bounds, and the lower bounds apply to all randomized methods of arbitrary order. In the other rows, the lower bounds apply to all deterministic first-order methods, and such methods achieve the upper bounds

and L_1 -Lipschitz gradient) for which \mathcal{A} requires at least T_ϵ oracle queries before returning an ϵ -stationary point of f .

In Part I [10] of this series we prove that no algorithm, even one given all derivatives of f at each iteration, can improve on the ϵ^{-2} rate of gradient descent for the class of functions with bounded initial value and Lipschitz continuous gradient. Therefore, in distinction with the convex case, acceleration of gradient descent for non-convex optimization [8] fundamentally depends on higher-order smoothness. We further show that, for the class of functions with p th order Lipschitz derivatives, no method can improve the rate $\epsilon^{-(p+1)/p}$ achieved by a p th-order method [5]. However, this does not get at the crux of the issue we consider here—what is the best possible rate for *first-order* methods, given that higher-order derivatives are Lipschitz?

In this paper we show that the ϵ -dependencies we establish in our work [8] are almost tight. More precisely, consider the function class with L_p -Lipschitz derivatives for all $q \in \{1, \dots, p\}$, where $p \in \mathbb{N}$; for this class there does not exist a deterministic first-order algorithm with iteration complexity better than $\epsilon^{-8/5}$. If $p = 2$ this complexity lower bound strengthens to $\epsilon^{-12/7}$. In the following diagram, we compare the exponents of $1/\epsilon$ in our lower bounds and known upper bounds (smaller is better).



Thus, we establish two separations. First, *no* deterministic first-order method can achieve the rate of convergence $\epsilon^{-3/2}$ of Newton's method. Second, the rate $\epsilon^{-5/3} \log \frac{1}{\epsilon}$ we achieve [8] *requires* the assumption of Lipschitz third derivatives, as first-order methods assuming only Lipschitz Hessian must compute at least $\epsilon^{-12/7}$ function values and gradients to find an ϵ -stationary point. We also show that the optimal rate for finding ϵ -stationary points of *convex* functions with bounded initial value (i.e.

$f(x^{(0)}) - \inf_x f(x) \leq \Delta$ and L_1 -Lipschitz gradient is $\tilde{\Theta}(\sqrt{L_1 \Delta} \epsilon^{-1})$.¹ Finding stationary points is thus fundamentally easier for convex functions.

The starting point of our development is Nesterov's [23, § 2.1.2] "worst function in the world,"

$$f_{\text{Nesterov}}(x) := \frac{1}{2}(x_1 - 1)^2 + \frac{1}{2} \sum_{i=1}^{d-1} (x_i - x_{i+1})^2, \quad (2)$$

which is instrumental in proving lower bounds for convex optimization [4, 23, 30] due to its "chain-like" structure. We establish our ϵ^{-1} lower bound for finding stationary points of convex functions by making a minor modification to the construction (2). To prove our larger lower bounds for non-convex functions, we augment f_{Nesterov} with a non-convex separable function $\sum_{i=1}^d \Upsilon(x_i)$, with $\Upsilon : \mathbb{R} \rightarrow \mathbb{R}$ carefully chosen to render Nesterov's "worst function" even worse.

By using this augmented chain function, we are able to prove lower bounds for all deterministic first-order methods. This is in contrast to prior work proving lower bounds for finding stationary points [11–14, 18, 29]. The papers [11–14] consider class of algorithms that receive p th order information and minimize p th-order approximations. While this class encapsulates many of the algorithms proposed in the literature, it misses cutting-plane and grid-search algorithms [17, 29]. Vavasis's lower bounds [29] apply to 2-dimensional problems (in contrast to ours, which require higher dimensionality), but are loose at least for high-dimensional problems.

Paper organization. Throughout, we use $\text{P}1.k$ to reference an item k of Part I of this sequence [10], as we build off of many ideas there. In Sect. 2 we briefly summarize our framework (Sects. P1.2 and P1.3). Sect. 3 begins the new analysis and contains lower bounds for finding stationary points of *convex* functions. In Sect. 4 we construct our hard non-convex instance, while in Sect. 5 we use this function to establish our main result: a lower bound on the complexity of finding stationary points using deterministic first-order methods. In Sects. 6 we discuss some difficulties in sharpening or extending our lower bounds. Sect. 7 concludes by situating our work in the current literature and reflecting on its implications for future research.

Notation. Before continuing, we provide the conventions we adopt throughout the paper; our notation mirrors Part I [10], so we describe it only briefly. For a sequence of vectors, subscripts denote coordinate index, while parenthesized superscripts denote element index, i.e. $x_j^{(i)}$ is the j th coordinate of the i th entry in the sequence $\{x^{(t)}\}_{t \in \mathbb{N}}$. For any $p \geq 1$ and p times continuously differentiable $f : \mathbb{R}^d \rightarrow \mathbb{R}$, we let $\nabla^p f(x)$ denote the symmetric tensor of p th order partial derivatives of f at point x . We let $\langle \cdot, \cdot \rangle$ be the Euclidean inner product on tensors, defined for order k tensors T and M by $\langle T, M \rangle = \sum_{i_1, \dots, i_k} T_{i_1, \dots, i_k} M_{i_1, \dots, i_k}$. We use \otimes to denote the Kronecker product and $\otimes^k d$ denote $d \times \dots \times d$, k times, so that $T \in \mathbb{R}^{\otimes^k d}$ denotes an order k tensor.

For a vector $v \in \mathbb{R}^d$ we let $\|v\| := \sqrt{\langle v, v \rangle}$ denote the Euclidean (ℓ_2) norm of v . For a tensor $T \in \mathbb{R}^{\otimes^k d}$, the ℓ_2 -operator norm of T is $\|T\|_{\text{op}} := \sup_{v^{(1)}, \dots, v^{(k)}} \{\langle v^{(1)} \otimes \dots \otimes v^{(k)}, T \rangle \mid \|v^{(i)}\| \leq 1, i = 1, \dots, k\}$, where we recall [31] that if T is symmetric

¹ Given a bound $\|x^{(0)} - x^*\| \leq D$ where $x^* \in \arg \min f$, as is standard for convex optimization, the optimal rate is $\tilde{\Theta}(\sqrt{L_1 D} \epsilon^{-1/2})$ [24]. The two rates are not directly comparable.

then $\|T\|_{\text{op}} = \sup_{\|v\|=1} |\langle v^{\otimes k}, T \rangle|$ where $v^{\otimes k}$ denotes the k -th Kronecker power of v . For vectors the ℓ_2 and ℓ_2 -operator norms are identical.

For any $n \in \mathbb{N}$, we let $[n] := \{1, \dots, n\}$ denote the set of positive integers less than or equal to n . We let \mathcal{C}^∞ denote the set of infinitely differentiable functions. We denote the i th standard basis vector by $e^{(i)}$, and let $I_d \in \mathbb{R}^{d \times d}$ denote the $d \times d$ identity matrix; we drop the subscript d when it is clear from context. For any set \mathcal{S} and functions $g, h : \mathcal{S} \rightarrow [0, \infty)$ we write $g \lesssim h$ or $g = O(h)$ if there exists a numerical constant $c < \infty$ such that $g(s) \leq c \cdot h(s)$ for every $s \in \mathcal{S}$. We write $g = \tilde{O}(h)$ if $g \lesssim h \log(h + 2)$.

2 A framework for lower bounds

For ease of reference, this section provides a condensed version of Sections PI.2 and PI.3 of the first part of this series [10] that lays out the notation, concepts and strategy we use to prove lower bounds. Here, we are deliberately brief; see [10] for motivation, intuition and background for our definitions, as well as exposition of randomized and higher-order methods.

2.1 Function classes

Typically, one designs optimization algorithms for certain classes of appropriately regular functions [6, 22, 23]. We thus focus on two notions of regularity that have been important for both convex and non-convex optimization: Lipschitzian properties of derivatives and bounds on function value. A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ has L_p -Lipschitz p th order derivatives if it is p times continuously differentiable, and for every $x \in \mathbb{R}^d$ and $v \in \mathbb{R}^d$, $\|v\| = 1$, the directional projection $t \mapsto f_{x,v}(t) := f(x + t \cdot v)$ of f satisfies

$$\left| f_{x,v}^{(p)}(t) - f_{x,v}^{(p)}(t') \right| \leq L_p |t - t'| \quad \text{for } t, t' \in \mathbb{R},$$

where $f_{x,v}^{(p)}(\cdot)$ is the p th derivative of $t \mapsto f_{x,v}(t)$. We occasionally refer to a function with Lipschitz p th order derivatives as p th-order smooth.

Definition 1 Let $p \geq 1$, $\Delta > 0$ and $L_p > 0$. Then the set

$$\mathcal{F}_p(\Delta, L_p)$$

denotes the union, over $d \in \mathbb{N}$, of the collection of \mathcal{C}^∞ functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ with L_p -Lipschitz p th derivative and $f(0) - \inf_x f(x) \leq \Delta$. For positive Δ and L_1, \dots, L_p we define

$$\mathcal{F}_{1:p}(\Delta, L_1, \dots, L_p) := \bigcap_{q \leq p} \mathcal{F}_q(\Delta, L_q).$$

The function classes $\mathcal{F}_p(\Delta, L_p)$ include functions on \mathbb{R}^d for all $d \in \mathbb{N}$, following the established practice of studying “dimension free” problems [10, 22, 23].

We also require the following important invariance notion [22, Ch. 7.2].

Definition 2 (Orthogonal invariance) A class of functions \mathcal{F} is *orthogonally invariant* if for every $f \in \mathcal{F}$, $f : \mathbb{R}^d \rightarrow \mathbb{R}$, and every matrix $U \in \mathbb{R}^{d \times d}$ such that $U^\top U = I_d$, the function $f_U : \mathbb{R}^d \rightarrow \mathbb{R}$ defined by $f_U(x) = f(U^\top x)$ belongs to \mathcal{F} .

Every function class we consider is orthogonally invariant.

2.2 Algorithm classes

For any dimension $d \in \mathbb{N}$, an *algorithm* \mathbf{A} (also referred to as a *method*) maps functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ to a sequence of *iterates* in \mathbb{R}^d ; that is, \mathbf{A} is defined separately for every finite d . We let

$$\mathbf{A}[f] = \{x^{(t)}\}_{t=1}^\infty$$

denote the sequence $x^{(t)} \in \mathbb{R}^d$ of iterates that \mathbf{A} generates when operating on f . Throughout this paper, we focus on *first-order deterministic algorithms*. Such an algorithm \mathbf{A} is one that, operating on $f : \mathbb{R}^d \rightarrow \mathbb{R}$, produces iterates of the form

$$x^{(i)} = \mathbf{A}^{(i)} \left(f(x^{(1)}), \nabla f(x^{(1)}), \dots, f(x^{(i-1)}), \nabla f(x^{(i-1)}) \right) \text{ for } i \in \mathbb{N},$$

where $\mathbf{A}^{(i)} : \mathbb{R}^{d(i-1)+i} \rightarrow \mathbb{R}^d$ is measurable (the dependence on dimension d is implicit). We denote the collection of first-order deterministic algorithms by $\mathcal{A}_{\text{det}}^{(1)}$.

Key to our development are *zero-respecting algorithms* (see Sec. P1.2.2 for more information). For $v \in \mathbb{R}^d$ we let $\text{supp}\{v\} := \{i \in [d] \mid v_i \neq 0\}$ denote the support (non-zero indices) of v . Then we say that the sequence $x^{(1)}, x^{(2)}, \dots$ is *first-order zero-respecting with respect to f* if

$$\text{supp}\{x^{(t)}\} \subseteq \bigcup_{s < t} \text{supp}\{\nabla f(x^{(s)})\} \text{ for each } t \in \mathbb{N}. \quad (3)$$

The definition (3) says that $x_i^{(t)} = 0$ whenever the partial derivatives of f with respect to coordinate x_i are zero for all preceding iterations. Extending the definition (3) in the obvious way, an algorithm \mathbf{A} is *first-order zero-respecting* if for any $f : \mathbb{R}^d \rightarrow \mathbb{R}$, the iterate sequence $\mathbf{A}[f]$ is zero-respecting with respect to f . The set $\mathcal{A}_{\text{zr}}^{(1)}$ comprises all such first-order algorithms.

2.3 Complexity measures

For a sequence $\{x^{(t)}\}_{t \in \mathbb{N}}$ we define the *complexity* of the sequence $\{x^{(t)}\}_{t \in \mathbb{N}}$ on f by

$$\mathsf{T}_\epsilon(\{x^{(t)}\}_{t \in \mathbb{N}}, f) := \inf \left\{ t \in \mathbb{N} \mid \|\nabla f(x^{(t)})\| \leq \epsilon \right\},$$

the index of the first element in $\{x^{(t)}\}_{t \in \mathbb{N}}$ that is an ϵ -stationary point of f . The complexity of algorithm \mathbf{A} on f is simply the complexity of the sequence $\mathbf{A}[f]$ on f , so we define

$$\mathsf{T}_\epsilon(\mathbf{A}, f) := \mathsf{T}_\epsilon(\mathbf{A}[f], f).$$

We define the *complexity of algorithm class \mathcal{A} on function class \mathcal{F}* as

$$\mathcal{T}_\epsilon(\mathcal{A}, \mathcal{F}) := \inf_{\mathcal{A} \in \mathcal{A}} \sup_{f \in \mathcal{F}} \mathcal{T}_\epsilon(\mathcal{A}, f). \quad (4)$$

Table 1 provides upper and lower bounds on the quantity (4) for different choices of \mathcal{A} and \mathcal{F} . For example, gradient descent guarantees $\mathcal{T}_\epsilon(\mathcal{A}_{\text{det}}^{(1)} \cap \mathcal{A}_{\text{zr}}^{(1)}, \mathcal{F}_1(\Delta, L_1)) \leq 2\Delta L_1 \epsilon^{-2}$.

2.4 How to prove a lower bound

The last step in our preliminaries is to give an overview of our proof strategy; this is an abbreviated version of Section PI.3. There, we abstract classical techniques from convex optimization [22,23], presenting a generic method for proving lower bounds on deterministic methods (of any order) applied to functions in any orthogonally invariant class.

Our starting point is what we call a *zero-chain*, which distills the “chain-like” structure of Nesterov’s construction (2).

Definition 3 A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is a *first-order zero-chain* if for every $x \in \mathbb{R}^d$,

$$\text{supp}\{x\} \subseteq \{1, \dots, i-1\} \text{ implies } \text{supp}\{\nabla f(x)\} \subset \{1, \dots, i\}.$$

In Definition PI.3 [10], we extend zero-chains to higher orders (unnecessary for this paper); in our terminology Nesterov’s function (2) is a first-order zero-chain, but not a second-order zero-chain. A first-order zero-chain limits the rate that zero-respecting algorithms acquire information from derivatives, forcing them to “discover” coordinates one by one, as the following observation makes clear.

Observation 1 Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a first-order zero-chain and let $x^{(1)} = 0, x^{(2)}, \dots$ be a first-order zero-respecting sequence with respect to f . Then $x_j^{(t)} = 0$ for $j \geq t$ and all $t \leq d$.

The important insight, essentially due to Nemirovski and Yudin [22] is that by using a resisting oracle [22,23] that can adversarially rotate the function f , any lower bound for zero-respecting algorithms implies an identical bound for all deterministic algorithms:

Proposition 1 Let \mathcal{F} be an orthogonally invariant function class, $f \in \mathcal{F}$ with domain of dimension d , and $\epsilon > 0$. If $\mathcal{T}_\epsilon(\mathcal{A}_{\text{zr}}^{(1)}, \{f\}) \geq T$, then

$$\mathcal{T}_\epsilon(\mathcal{A}_{\text{det}}^{(1)}, \mathcal{F}) \geq \mathcal{T}_\epsilon(\mathcal{A}_{\text{det}}^{(1)}, \{f_U \mid U \in O(d+T, d)\}) \geq T,$$

where $f_U := f(U^\top z)$ and $O(d+T, d)$ is the set of $(d+T) \times d$ orthogonal matrices, so that $\{f_U \mid U \in O(d+T, d)\}$ contains only functions with domain of dimension $d+T$.

See Propositions PI.1 and PI.2 for more general versions of this result.

With this proposition, our strategy, inspired by Nesterov [23], becomes clear. To prove a lower bound on first-order deterministic algorithms for a function class \mathcal{F} , we find $f_\epsilon : \mathbb{R}^T \rightarrow \mathbb{R}$ such that (i) f_ϵ is a first-order zero-chain, (ii) $f_\epsilon \in \mathcal{F}$, and (iii) $\|\nabla f_\epsilon(x)\| > \epsilon$ for every x such that $x_T = 0$. Then for $\mathbf{A} \in \mathcal{A}_{\text{zr}}^{(1)}$ and $\{x^{(t)}\}_{t \in \mathbb{N}} = \mathbf{A}[f]$, Observation 1 shows that $x_T^{(t)} = 0$ for $t \leq T$, and the large gradient property (iii) guarantees the non-stationarity $\|\nabla f_\epsilon(x^{(t)})\| > \epsilon$ for all $t \leq T$. We immediately obtain the complexity lower bound

$$\mathcal{T}_\epsilon(\mathcal{A}_{\text{zr}}^{(1)}, \mathcal{F}) = \inf_{\mathbf{A} \in \mathcal{A}_{\text{zr}}^{(1)}} \sup_{f \in \mathcal{F}} \mathcal{T}_\epsilon(\mathbf{A}, f) \geq \inf_{\mathbf{A} \in \mathcal{A}_{\text{zr}}^{(1)}} \mathcal{T}_\epsilon(\mathbf{A}, f_\epsilon) > T,$$

and the same lower bound holds for $\mathcal{T}_\epsilon(\mathcal{A}_{\text{det}}^{(1)}, \mathcal{F})$ by Proposition 1.

Hard instance dimensionality. To establish the lower bound $\mathcal{T}_\epsilon(\mathcal{A}_{\text{det}}^{(1)}, \mathcal{F}) \geq T_\epsilon$ we require functions in \mathcal{F} with domain of dimension $2T_\epsilon$; the first-order zero-chain has domain of dimension $d = T_\epsilon$, and Proposition 1 doubles this dimension via orthogonal projections to $\mathbb{R}^{d+T_\epsilon} = \mathbb{R}^{2T_\epsilon}$. Thus, the dimensionality of our hard instances grows inversely $1/\epsilon$ with precisely the same rate as the lower bounds themselves. In Section PI.1.2 we explain why such high-dimensional constructions are unavoidable for tight dimension-free lower bounds.

3 Lower bounds for finding stationary points of convex functions

While for *convex* optimization guarantees of small gradients are atypical topics of study, we nonetheless begin by considering the complexity of finding stationary points of smooth convex functions. This serves two purposes. First, it is a baseline for finding stationary points in the non-convex setting; based on algorithmic upper bounds due to Nesterov [24], we see that convexity makes this task fundamentally easier. Second, our lower bound construction for convex problems underpins our construction and analysis for general smooth (non-convex) functions in the sequel, allowing us to demonstrate our techniques in a simpler setting. Of course, in convex optimization, it is typically more useful to find points x with small optimality gap, $f(x) \leq \inf_z f(z) + \epsilon$. Convexity allows efficient algorithms for guaranteeing such optimality, and typically one ignores the magnitude of the gradient in favor of small optimality or duality gaps [6]. Nonetheless, in some situations—such as certifying (near) dual feasibility or small constraint residuals in primal-dual or operator splitting algorithms [e.g. 7]—achieving small gradients is important.

We proceed as follows. In Sect. 3.1 we define the class of convex functions under consideration and a quadratic subclass. In Sect. 3.2, we construct a hard quadratic instance, and verify its key properties. Finally, in Sect. 3.3, we state, discuss and prove our lower bounds.

3.1 Convex function classes

The collections of functions we consider are the following.

Definition 4 Let $L_1 > 0$ and $\Delta > 0$. The set

$$\mathcal{K}_1(\Delta, L_1)$$

denotes the union, over $d \in \mathbb{N}$, of the collections of \mathcal{C}^∞ convex functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ with L_1 -Lipschitz gradient and $f(0) - \inf_x f(x) \leq \Delta$. Additionally,

$$\mathcal{Q}(\Delta, L_1) \subset \mathcal{K}_1(\Delta, L_1)$$

is the set of convex quadratic functions satisfying the above conditions.

Our results, following Nemirovski and Yudin [22] and Nesterov [23], demonstrate that for deterministic first-order methods, the class $\mathcal{Q}(\Delta, L_1)$ is “hard enough,” in that it provides nearly sharp lower bounds for first-order methods, which immediately apply to $\mathcal{K}_1(\Delta, L_1)$ and $\mathcal{F}_1(\Delta, L_1)$. We also have $\mathcal{Q}(\Delta, L_1) = \mathcal{F}_{1:p}(\Delta, L_1, 0, \dots, 0)$ or any $p \geq 2$.

In addition to functions restricted by initial optimality gap, we consider the following initial distance-based definition.

Definition 5 Let $D > 0$ and $L_1 > 0$. The set

$$\mathcal{K}_1^{\text{dist}}(D, L_1)$$

denotes the union, over $d \in \mathbb{N}$, of the collections of \mathcal{C}^∞ convex functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ with L_1 -Lipschitz gradient satisfying $\|x^*\| \leq D$ for all $x^* \in \arg \min_x f(x)$. Additionally,

$$\mathcal{Q}^{\text{dist}}(D, L_1) \subset \mathcal{K}_1^{\text{dist}}(D, L_1)$$

is the set of convex quadratic functions satisfying the above conditions.

Standard convergence results in (smooth) convex optimization [e.g. 23] apply to functions with bounded domain, i.e. $f \in \mathcal{K}_1^{\text{dist}}(D, L_1)$ rather than $\mathcal{K}_1(\Delta, L_1)$. This is for good reason: for any pair Δ, L_1 , any $\epsilon < \Delta$, any first-order zero-respecting or deterministic algorithm \mathcal{A} , and any $T \in \mathbb{N}$, there exists a function $f \in \mathcal{Q}(\Delta, L_1)$ with L_1 -Lipschitz gradient such for $\{x^{(t)}\}_{t \in \mathbb{N}} = \mathcal{A}[f]$ we have

$$\inf_{t \in \mathbb{N}} \left\{ t \mid f(x^{(t)}) \leq \inf_x f(x) + \epsilon \right\} > T.$$

(See Appendix A.2, Lemma 6 for a proof of this claim.) Since this holds for any $T \in \mathbb{N}$ and $\epsilon < \Delta$, making even the slightest function value improvement to functions in $\mathcal{Q}(\Delta, L_1)$ may take arbitrarily long. Thus, when we consider the function classes of Definition 4, we can only hope to give convergence guarantees in terms of stationarity—as is common in the non-convex case.

3.2 The worst function in the (convex) world

We now construct the functions that are difficult for any zero-respecting first-order method. For parameters $T \in \mathbb{N}$ and $\alpha \leq 1$ we define the (unscaled) hard function $\hat{f}_{T,\alpha} : \mathbb{R}^T \rightarrow \mathbb{R}$ by

$$\hat{f}_{T,\alpha}(x) = \frac{\alpha}{2}(x_1 - 1)^2 + \frac{1}{2} \sum_{i=1}^{T-1} (x_i - x_{i+1})^2. \quad (5)$$

For $\alpha = 1$, $\hat{f}_{T,1}$ is Nesterov's "worst function in the world" [23, § 2.1.2]. The parameter α allows us to control $f(0)$ and thus provides a degree of freedom in satisfying the constraint $f(0) - \inf_x f(x) \leq \Delta$ for our lower bounds. By inspection,

$$\hat{f}_{T,\alpha}(x) = \frac{1}{2}x^\top Lx - b^\top x + \frac{\alpha}{2}$$

where

$$L = \begin{bmatrix} 1+\alpha & -1 & & & \\ -1 & 2 & -1 & & \\ & \ddots & \ddots & \ddots & \\ & & -1 & 2 & -1 \\ & & & -1 & 1 \end{bmatrix} \in \mathbb{R}^{T \times T} \quad (6)$$

is the unnormalized graph Laplacian of the simple path on T vertices (see [16]) plus the term α in the position L_{11} , and $b = \alpha e^{(1)}$.

Let us now verify that $\hat{f}_{T,\alpha}$ meets the three requirements of our lower bound strategy.

Lemma 1 *For all $T \in \mathbb{N}$ and $\alpha \leq 1$, $\hat{f}_{T,\alpha}$ has the following properties.*

- i. **Zero-chain** $\hat{f}_{T,\alpha}$ is a first-order zero-chain.
- ii. **Membership in function class**
 - (a) $\hat{f}_{T,\alpha}$ has 4-Lipschitz continuous gradient.
 - (b) $\hat{f}_{T,\alpha}(0) - \inf_{x \in \mathbb{R}^T} \hat{f}_{T,\alpha}(x) = \alpha/2$.
 - (c) The unique minimizer of $\hat{f}_{T,\alpha}(x)$ is $x^* = \mathbf{1}$, and $\|x^*\| = \sqrt{T}$.
- iii. **Large gradient** For every $x \in \mathbb{R}^T$ such that $x_T = 0$, $\|\nabla \hat{f}_{T,\alpha}(x)\| > (T - 1 + \frac{1}{\alpha})^{-3/2}$.

Proof Part i is immediate from Definition 3, since for every $i \in [d]$, $\nabla_i \hat{f}_{T,\alpha}(x) = 0$ whenever $x_{i-1} = x_i = x_{i+1} = 0$. Part ii is also immediate, as $\hat{f}_{T,\alpha}(\mathbf{1}) = \inf_x \hat{f}_{T,\alpha}(x) = 0$, and $\|L\|_{\text{op}} \leq 4$ (apply the triangle inequality to $\|Lv\|$). To establish part iii, we calculate the minimum value of $\|\nabla \hat{f}_{T,\alpha}(x)\|^2$ obtainable by any vector $x \in \mathbb{R}^T$ with $x_T = 0$. Letting $M = L[I_{T-1} \ 0_{T-1}]^\top \in \mathbb{R}^{T \times (T-1)}$ be the matrix L of (6) with its last column removed and recalling $b = \alpha e^{(1)}$, this becomes the least

squares problem, whose solutions is the squared norm of the projection of b to the (one-dimensional) nullspace of M^\top :

$$\begin{aligned} \inf_{x \in \mathbb{R}^T, x_T = 0} \left\| \nabla \hat{f}_{T,\alpha}(x) \right\|^2 &= \inf_{v \in \mathbb{R}^{T-1}} \|Mv - b\|^2 \\ &= b^\top \left(I_T - M(M^\top M)^{-1}M^\top \right) b = \left(z^\top b \right)^2, \end{aligned} \quad (7)$$

where $z \in \mathbb{R}^T$ is the unique (up to sign) unit-norm solution to $M^\top z = 0$. A calculation shows that

$$z_j = \frac{j - 1 + \frac{1}{\alpha}}{\sqrt{\sum_{i=1}^T (i - 1 + \frac{1}{\alpha})^2}}.$$

Substituting z and b into Eq. (7), we have that $x_T = 0$ implies

$$\left\| \nabla \hat{f}_{T,\alpha}(x) \right\|^2 \geq \frac{1}{\sum_{i=1}^T (i - 1 + \frac{1}{\alpha})^2} > \frac{1}{(T - 1 + \frac{1}{\alpha})^3},$$

giving the result. \square

3.3 Scaling argument and final bound

With our hard instance in place, we provide our lower bounds for finding stationary points of convex functions. We note that the lower bound for the class $\mathcal{Q}^{\text{dist}}(D, L_1)$ also follows from the standard lower bounds on finding ϵ -suboptimal points, since for every $q \in \mathcal{Q}^{\text{dist}}(D, L_1)$ an ϵ -stationary point is also ϵD -suboptimal.

Theorem 1 *Let ϵ, Δ, D , and L_1 be positive. Then*

$$\mathcal{T}_\epsilon(\mathcal{A}_{\text{det}}^{(1)}, \mathcal{K}_1(\Delta, L_1)) \geq \mathcal{T}_\epsilon(\mathcal{A}_{\text{zr}}^{(1)}, \mathcal{Q}(\Delta, L_1)) \geq \frac{\sqrt{L_1 \Delta}}{4} \epsilon^{-1}, \quad (8a)$$

and

$$\mathcal{T}_\epsilon(\mathcal{A}_{\text{det}}^{(1)}, \mathcal{K}_1^{\text{dist}}(D, L_1)) \geq \mathcal{T}_\epsilon(\mathcal{A}_{\text{zr}}^{(1)}, \mathcal{Q}^{\text{dist}}(D, L_1)) \geq \frac{\sqrt{L_1 D}}{2} \epsilon^{-1/2}. \quad (8b)$$

Let us discuss Theorem 1 briefly. Nesterov [24] shows that for any $f \in \mathcal{K}_1^{\text{dist}}(D, L_1)$, accelerated gradient descent applied to a regularized version of f yields a point x satisfying $\|\nabla f(x)\| \leq \epsilon$ after at most $O(\sqrt{L_1 D} \epsilon^{-1/2} \log \frac{L_1 D}{\epsilon})$ iterations. For $f \in \mathcal{K}_1(\Delta, L_1)$, a similar technique to Nesterov's, which we provide for completeness in Appendix A.1, yields an upper complexity bound of $O(\sqrt{L_1 \Delta} \epsilon^{-1} \log \frac{L_1 \Delta}{\epsilon^2})$. Thus, to within logarithmic factors both bounds of Theorem 1 are sharp. It is illustrative to compare Theorem 1 to our results for *non-convex* but smooth functions, and we do so in detail in Sect. 7.1. The comparison shows that finding stationary points of

smooth convex functions with first-order methods is fundamentally easier than finding stationary points of non-convex functions, even with higher-order smoothness and using higher-order methods.

While we prove our lower bounds for the algorithm classes $\mathcal{A}_{\text{det}}^{(1)}$ and $\mathcal{A}_{\text{zr}}^{(1)}$, similar lower bounds apply to the collection $\mathcal{A}_{\text{rand}}$ of all randomized algorithms based on arbitrarily high-order derivatives when applied to worst-case functions from function class $\mathcal{K}_1(\Delta, L_1)$. While this is not our focus here, using the techniques of Woodworth and Srebro [30] and Section P1.5, it is possible to construct a distribution \mathbb{P} on $\mathcal{K}_1(\Delta, L_1)$ such that for any $\mathbf{A} \in \mathcal{A}_{\text{rand}}$, with high probability over $f \sim \mathbb{P}$ we have $\mathsf{T}_\epsilon(\mathbf{A}, f) \gtrsim \sqrt{L_1 \Delta \epsilon}^{-1}$. That is, neither randomization nor higher-order derivative information can improve performance on $\mathcal{K}_1(\Delta, L_1)$. Such an extension fails for $\mathcal{Q}(\Delta, L_1)$, as Newton's method finds the global minimizer of every $f \in \mathcal{Q}(\Delta, L_1)$ in one step.

3.4 Proof of Theorem 1

As we outline in Sect. 2.4, we establish our lower bounds by constructing a zero-chain $f : \mathbb{R}^T \rightarrow \mathbb{R}$ such that $f \in \mathcal{Q}(\Delta, L_1)$ (or $\mathcal{Q}^{\text{dist}}(D, L_1)$), and that $\|\nabla f(x)\| > \epsilon$ for any x such that $x_T = 0$. By Observation 1 we immediately have that for every $\mathbf{A} \in \mathcal{A}_{\text{zr}}^{(1)}$, the iterates $\{x^{(t)}\}_{t \in \mathbb{N}} = \mathbf{A}[f]$ produced by \mathbf{A} operating on f satisfy $x_T^{(t)} = 0$ for every $t \leq T$ and hence $\|\nabla f(x^{(t)})\| > \epsilon$. Consequently, $\inf_{\mathbf{A} \in \mathcal{A}_{\text{zr}}^{(1)}} \mathsf{T}_\epsilon(\mathbf{A}, f) \geq 1 + T$, which implies lower bounds on the required quantities by means of $\mathcal{Q}(\Delta, L_1) \subset \mathcal{K}_1(\Delta, L_1)$ and Proposition 1.

To define the difficult zero-chain f , we scale $\hat{f}_{T,\alpha}$ using two scalar parameters $\lambda, \sigma > 0$, which we determine later, defining

$$f(x) := \lambda \sigma^2 \hat{f}_{T,\alpha}(x/\sigma).$$

We use the parameter $\lambda > 0$ to control the first-order smoothness of f , as $\nabla^2 f(x) = \lambda \nabla^2 \hat{f}_{T,\alpha}(x/\sigma)$, while the parameter σ controls the lower bound on $\|\nabla f(x)\|$ for $x_T = 0$. We first show how to choose σ , depending on T, ϵ, α , and λ . By Lemma 1.iii, for every x with $x_T = 0$ we have

$$\|\nabla f(x)\| = \lambda \sigma \left\| \nabla \hat{f}_{T,\alpha}(x) \right\| > \frac{\lambda \sigma}{(T - 1 + \frac{1}{\alpha})^{3/2}}.$$

Setting

$$\sigma = \frac{1}{\lambda} \left(T - 1 + \frac{1}{\alpha} \right)^{3/2} \epsilon,$$

guarantees $\|\nabla f(x)\| > \epsilon$ for any x such that $x_T = 0$ and hence $\|\nabla f(x^{(t)})\| > \epsilon$ for all $t \leq T$.

All that remains is to choose λ , T , and α to guarantee that f belongs to the appropriate quadratic class. By Lemma 1.ii, f has 4λ Lipschitz gradient, so we take

$$\lambda = L_1/4$$

and guarantee that f has L_1 -Lipschitz gradient. To guarantee that $f \in \mathcal{Q}(\Delta, L_1)$, Lemma 1.ii yields

$$f(0) - \inf_x f(x) = \lambda \sigma^2 \alpha / 2 = \frac{2\alpha}{L_1} \left(T - 1 + \frac{1}{\alpha} \right)^3 \epsilon^2,$$

where we have substituted our choice of σ and λ in the final equality. Defining

$$\alpha = 1/T \leq 1 \text{ we obtain } f(0) - \inf_x f(x) \leq 16T^2 \epsilon^2 / L_1,$$

so to guarantee $f(0) - \inf_x f(x) \leq \Delta$, it suffices to choose

$$T = \left\lfloor \frac{\sqrt{L_1 \Delta}}{4} \epsilon^{-1} \right\rfloor.$$

This gives the first part (8a) of the theorem. For inequality (8b), we must have $f \in \mathcal{Q}^{\text{dist}}(D, L_1)$. Let $x^* = \sigma \mathbf{1}$ denote the minimizer of f , so that

$$\|x^*\| = \sigma \sqrt{T} = \frac{4}{L_1} \left(T - 1 + \frac{1}{\alpha} \right)^{3/2} \epsilon \sqrt{T},$$

where again we have substituted our choices of σ and λ in the final equality. Consequently, to guarantee $\|x^*\| \leq D$ it suffices to take

$$\alpha = 1 \text{ and } T = \left\lfloor \frac{\sqrt{L_1 D}}{2} \epsilon^{-1/2} \right\rfloor,$$

giving the bound (8b).

4 Constructing the non-convex hard instance

We now relax the assumption of convexity, and design a first-order zero-chain that provides bounds stronger than those of Theorem PI.2, when we restrict the algorithm class to first-order methods. The basis of our construction is the convex zero-chain (5), which we augment with non-convexity to strengthen the gradient lower bound in Lemma 1.iii, while ensuring that all derivatives remain Lipschitz continuous. With this in mind, for each $T \in \mathbb{N}$, we define the unscaled hard instance $\bar{f}_{T,\mu,r} : \mathbb{R}^{T+1} \rightarrow \mathbb{R}$ as

$$\bar{f}_{T,\mu,r}(x) = \frac{\sqrt{\mu}}{2} (x_1 - 1)^2 + \frac{1}{2} \sum_{i=1}^T (x_{i+1} - x_i)^2 + \mu \sum_{i=1}^T \Upsilon_r(x_i). \quad (9)$$

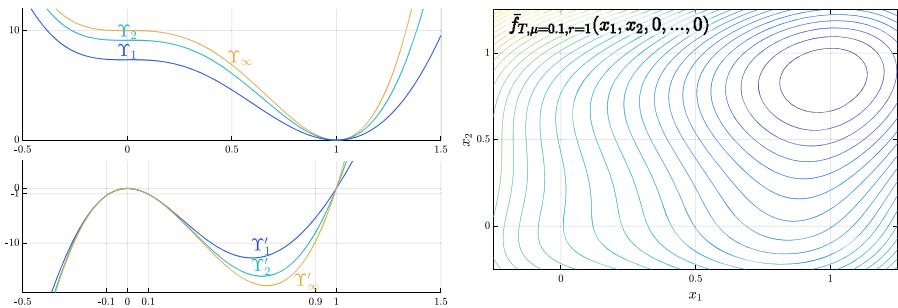


Fig. 1 Hard instance for first-order methods. Left: the non-convexity Υ_r (top) and its derivative (bottom), for different values of r . Right: contour plot of a two-dimensional cross-section of the hard instance $\bar{f}_{T,\mu,r}$

where the non-convex function $\Upsilon_r : \mathbb{R} \rightarrow \mathbb{R}$, parameterized by $r \geq 1$, is

$$\Upsilon_r(x) = 120 \int_1^x \frac{t^2(t-1)}{1+(t/r)^2} dt. \quad (10)$$

We illustrate the construction $\bar{f}_{T,\mu,r}$ in Fig. 1; it is the sum of the convex hard instance (5) (with $\alpha = \sqrt{\mu}$) and a separable non-convex function. In the following lemma, which we prove in Appendix B.1, we list the important properties of Υ_r .

Lemma 2 *The function Υ_r satisfies the following.*

- i. *We have $\Upsilon'_r(0) = \Upsilon'_r(1) = 0$.*
- ii. *For all $x \leq 1$, $\Upsilon'_r(x) \leq 0$, and for all $x \geq 1$, $\Upsilon'_r(x) \geq 0$.*
- iii. *For all $x \in \mathbb{R}$ we have $\Upsilon_r(x) \geq \Upsilon_r(1) = 0$, and for all r , $\Upsilon_r(0) \leq 10$.*
- iv. *For every $r \geq 1$, $\Upsilon'_r(x) < -1$ for every $x \in (-\infty, -0.1] \cup [0.1, 0.9]$.*
- v. *For every $r \geq 1$ and every $p \geq 1$, the p -th order derivatives of Υ_r are $r^{3-p}\ell_p$ -Lipschitz continuous, where $\ell_p \leq \exp(\frac{3}{2}p \log p + cp)$ for a numerical constant $c < \infty$.*

Before formally stating the properties of $\bar{f}_{T,\mu,r}$, we provide a high-level explanation of the choice of Υ_r . First, a necessary and sufficient condition for $\bar{f}_{T,\mu,r}$ to be a first-order zero-chain is that $\Upsilon'_r(0) = 0$. Second, examining the proof Lemma 1.iii we see that the gradient of the quadratic chain is smallest for vectors x with entries x_1, x_2, \dots, x_T that slowly decrease from 1 to 0. We design Υ_r to “punish” such slowly varying vectors, by demanding that $\Upsilon'_r(x)$ be large for any x far from both 0 and 1 (Lemma 2.iv); this is the key to improving Lemma 1.iii and the most important property of Υ_r . Third, for every finite r all the derivatives of Υ_r are Lipschitz, and as r increases $\bar{f}_{T,\mu,r}$ converges to a quartic polynomial; in the limit $r = \infty$ we have $\Upsilon_\infty(x) = 30x^4 - 40x^3 + 10$. This allows us to establish that Lipschitz continuity of derivatives beyond the third does not alter the ϵ dependence of our bounds. However, we cannot simply use Υ_∞ , as its first three derivatives are unbounded. Lastly, we place the minimum of $\Upsilon_r(x)$ at $x = 1$, so that the all-ones vector is the global minimizer of $\bar{f}_{T,\mu,r}$, and $\bar{f}_{T,\mu,r}(\mathbf{1}) = 0$; this is simply convenient for our analysis.

With our considerations explained, we verify the three components of our general strategy: f is a first-order zero-chain, belongs to the relevant function classes, and has large gradient whenever $x_T = 0$. We begin with the zero-chain property, which follows trivially from Lemma 2.i.

Observation 2 For any $T \in \mathbb{N}$, and positive μ and r , $\bar{f}_{T,\mu,r}$ is a first-order zero-chain.

Crucially, $\bar{f}_{T,\mu,r}$ is *only* a first-order zero-chain (see Definition PI.3); were it a second-order zero-chain, the resulting lower bounds would apply to second-order algorithms as well, where Newton's method achieves the rate $\epsilon^{-3/2}$ [25], which is strictly better than all of our lower bounds. We next show that any point x for which $x_T = x_{T+1} = 0$ has large gradient. This is the core technical result of our analysis.

Lemma 3 Let $r \geq 1$ and $\mu \leq 1$. For any $x \in \mathbb{R}^{T+1}$ such that $x_T = x_{T+1} = 0$,

$$\|\nabla \bar{f}_{T,\mu,r}(x)\| > \mu^{3/4}/4.$$

We defer the full proof of this lemma to Appendix B.2 and sketch its main idea here. We may view any vector meeting the conditions of the lemma as a sequence going from $x_0 := 1$ to $x_T = 0$. Every such sequence must have a “transition region”, which we define roughly as the subsequence starting after the last i such that $x_i > \frac{9}{10}$ and ending at the first (subsequent) j such that $x_j < \frac{1}{10}$ (see Fig. 2). Letting $m \in \{1, \dots, T\}$ denote the length of this subsequence and ignoring constant factors, we establish that

$$\|\nabla \bar{f}_{T,\mu,r}(x)\| \geq \max \left\{ (m + 1/\sqrt{\mu})^{-3/2}, \mu\sqrt{m} \right\}.$$

The $(m + 1/\sqrt{\mu})^{-3/2}$ bound comes from the quadratic chain in $\bar{f}_{T,\mu,r}$, which has large gradient for any sequence x with sharp transitions; this is essentially Lemma 1.iii with $T = m$ and $\alpha = \sqrt{\mu}$. The $\mu\sqrt{m}$ bound is due to the non-convex Υ_r terms in $\bar{f}_{T,\mu,r}$, which by Lemma 2.iv contribute a term of magnitude μ to every entry of $\nabla \bar{f}_{T,\mu,r}$ in the transition region. These two bounds intersect at $m \approx 1/\sqrt{\mu}$, so the gradient has norm at least $\mu^{3/4}$ for every value of m .

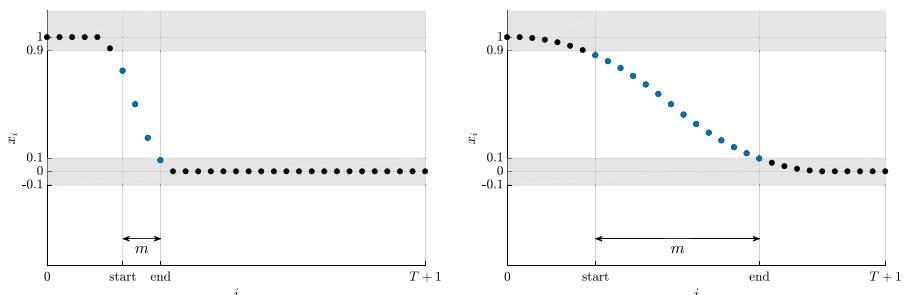


Fig. 2 Illustration of the “transition region” concept used to prove Lemma 3. Each plot shows the entries of a vector $x \in \mathbb{R}^{T+1}$ that satisfies $x_T = x_{T+1} = 0$, with entries of x belonging to the transition region marked in blue. Short transitions (left) incur large gradients due to the convex quadratic term in $\bar{f}_{T,\mu,r}$, while long transitions (right) incur large gradients due to the non-convex Υ_r terms and Lemma 2.iv

Finally, we list the boundedness properties of our construction.

Lemma 4 *The function $\bar{f}_{T,\mu,r}$ satisfies the following.*

- i. $\bar{f}_{T,\mu,r}(0) - \inf_x \bar{f}_{T,\mu,r}(x) \leq \frac{\sqrt{\mu}}{2} + 10\mu T$
- ii. For $\mu \leq 1, r \geq 1$ and every $p \geq 1$, the p -th order derivatives of $\bar{f}_{T,\mu,r}$ are $(1_{(p=1)} + r^{3-p}\mu)\ell_p$ -Lipschitz continuous, where $\ell_p \leq e^{\frac{3p}{2}\log p + cp}$ for a numerical constant $c < \infty$.

Proof The first part of the lemma follows from Lemma 2, which shows that $\inf_x \bar{f}_{T,\mu,r}(x) = \bar{f}_{T,\mu,r}(\mathbf{1}) = 0$, while $\bar{f}_{T,\mu,r}(0) = \sqrt{\mu}/2 + T\mu\Upsilon_r(0) \leq \sqrt{\mu}/2 + 10\mu T$. The second part of the lemma follows directly from Lemma 2.v and that the quadratic chain $f(x) = \frac{\sqrt{\mu}}{2}(x_1 - 1)^2 + \frac{1}{2} \sum_i (x_i - x_{i+1})^2$ has 4-Lipschitz gradient and 0-Lipschitz higher order derivatives. \square

5 Lower bounds for first-order methods

We now give our main result: lower bounds for the complexity of finding ϵ -stationary points using the class $\mathcal{A}_{\text{det}}^{(1)} \cup \mathcal{A}_{\text{zf}}^{(1)}$ of first-order deterministic and/or zero-respecting algorithms, applied to functions in the class

$$\mathcal{F}_{1:p}(\Delta, L_1, \dots, L_p) := \bigcap_{q \leq p} \mathcal{F}_q(\Delta, L_q)$$

containing all functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$, $d \in \mathbb{N}$, such that $f(0) - \inf_x f(x) \leq \Delta$ and $\nabla^q f$ is L_q -Lipschitz continuous for $1 \leq q \leq p$.

Theorem 2 *There exist numerical constants $c, C \in \mathbb{R}_+$ and $\ell_q \leq e^{\frac{3q}{2}\log q + Cq}$ for every $q \in \mathbb{N}$ such that the following lower bound holds. Let $p \in \mathbb{N}$, and let $\Delta, L_1, L_2, \dots, L_p, \epsilon$ be positive. Assume additionally that $\epsilon \leq (L_1^q/L_q)^{1/(q-1)}$ for each $q \in \{2, \dots, p\}$. Then for $p \geq 3$,*

$$\begin{aligned} T_\epsilon(\mathcal{A}_{\text{det}}^{(1)} \cup \mathcal{A}_{\text{zf}}^{(1)}, \mathcal{F}_{1:p}(\Delta, L_1, \dots, L_p)) &\geq c \cdot \Delta \\ &\cdot \min_{q \in \{2, \dots, p\}} \left\{ \left(\frac{L_1}{\ell_1} \right)^{\frac{3}{5} - \frac{2}{5(q-1)}} \left(\frac{L_q}{\ell_q} \right)^{\frac{2}{5(q-1)}} \right\} \epsilon^{-8/5}. \end{aligned}$$

Moreover, for $p = 2$,

$$T_\epsilon(\mathcal{A}_{\text{det}}^{(1)} \cup \mathcal{A}_{\text{zf}}^{(1)}, \mathcal{F}_{1:2}(\Delta, L_1, L_2)) \geq c \cdot \Delta \left(\frac{L_1}{\ell_1} \right)^{\frac{3}{7}} \left(\frac{L_2}{\ell_2} \right)^{\frac{2}{7}} \epsilon^{-12/7}.$$

We prove Theorem 2 in Sect. 5.2 to come, providing a brief overview of the argument here, and then providing some discussion. In the proof, we construct the hard instance $f : \mathbb{R}^{T+1} \rightarrow \mathbb{R}$ as $f(x) = \lambda\sigma^2 \bar{f}_{T,\mu,r}(x/\sigma)$, where we must choose the parameters

$\lambda, \sigma > 0$ as well as μ, r , and T to guarantee that f is (a) hard to optimize, i.e. $T_\epsilon(\mathbf{A}, f) > T$ for every $\mathbf{A} \in \mathcal{A}_{\text{zr}}^{(1)}$, and (b) meets the smoothness and boundedness requirements of the function class.

We begin by sketching the argument for $p = 2$. In this case, we may take $r = 1$, $\lambda \propto L_1$ and $\mu \propto L_2\sigma/\lambda$. This choice guarantees that f has L_1 -Lipschitz gradient and L_2 -Lipschitz Hessian when $\mu \leq 1$, which we later verify using the assumption $\epsilon \leq L_1^2/L_2$. We then use Observation 2, Lemma 3 and Observation 1 to show that $T_\epsilon(\mathbf{A}, f) \geq T + 1$ for every $\mathbf{A} \in \mathcal{A}_{\text{zr}}^{(1)}$ whenever $\lambda\sigma\mu^{3/4}/4 \geq \epsilon$, and conclude that σ may scale as $\lambda^{1/7}\epsilon^{4/7}$ (since $\mu \propto L_2\sigma/\lambda$). By Lemma 4.i we have

$$f(0) - \inf_x f(x) \leq \lambda\sqrt{\mu\sigma^2}/2 + 10\lambda\mu\sigma^2 T,$$

so we can take $T \propto \Delta/(\lambda\mu\sigma^2) \propto \Delta/(L_2\sigma^3)$ to guarantee $f(0) - \inf_x f(x) \leq \Delta$, where we assume without loss of generality that $\lambda\sqrt{\mu\sigma^2} \leq \Delta$ (otherwise Theorem 1 dominates our bound). Substituting the expressions for σ , μ and λ into the expression for T gives the result for $p = 2$.

For $p \geq 3$ we require a more careful argument, as we must simultaneously handle all orders of smoothness. To do so, we let $\mu = \bar{\mu}\sigma^2/\lambda$ and $r = \bar{r}/\sigma$, and show how to take \bar{r} and $\bar{\mu}$ independently of ϵ (depending only on L_1, \dots, L_p). This allows us to obtain identical ϵ -dependence for all $p \geq 3$.

To better understand the theorem, we give a few additional remarks.

Near-achievability of the lower bounds. In the paper [8], we propose the method ‘‘convex until proven guilty,’’ which augments Nesterov’s accelerated gradient method with implicit negative curvature descent. For the function classes $\mathcal{F}_{1:2}(\Delta, L_1, L_2)$ and $\mathcal{F}_{1:3}(\Delta, L_1, L_2, L_3)$, it achieves rates of convergence $\tilde{O}(\Delta L_1^{1/2} L_2^{1/4} \epsilon^{-7/4})$ and $\tilde{O}(\Delta L_1^{1/2} L_3^{1/6} \epsilon^{-5/3})$, respectively. These results nearly match our lower bounds in Theorem 2; in the case of $p = 2$, the gap (in terms of ϵ) is of order $\epsilon^{-\frac{1}{28}} \log \frac{1}{\epsilon}$, while for $p \geq 3$, the gap is of order $\epsilon^{-\frac{1}{15}} \log \frac{1}{\epsilon}$. See further discussion in Sect. 7.1.

Choice of function class. The focus on the more restricted function classes $\mathcal{F}_{1:p}(\Delta, L_1, \dots, L_p)$ —rather than the classes $\mathcal{F}_p(\Delta, L_p)$ we study in Part I [10]—makes our lower bounds stronger. Moreover, it is necessary for non-trivial results, since for any $p \geq 2$ and $\Delta, L_p > 0$, the class $\mathcal{F}_p(\Delta, L_p)$ contains functions impossible for first-order methods. Indeed, the class $\mathcal{Q}(\Delta, L_1)$ of Δ -bounded L_1 -smooth convex quadratics is a subset of $\mathcal{F}_p(\Delta, L_p)$ for any $L_1 < \infty$ and $L_p > 0$. Therefore, by Theorem 1,

$$\begin{aligned} T_\epsilon(\mathcal{A}_{\text{det}}^{(1)} \cup \mathcal{A}_{\text{zr}}^{(1)}, \mathcal{F}_p(\Delta, L_p)) &\geq \sup_{L_1 < \infty} T_\epsilon(\mathcal{A}_{\text{det}}^{(1)} \cup \mathcal{A}_{\text{zr}}^{(1)}, \mathcal{Q}(\Delta, L_1)) \\ &\geq \sup_{L_1 < \infty} \frac{\sqrt{L_1 \Delta}}{4} \epsilon^{-1} = \infty. \end{aligned}$$

We thus limit our scope to functions with smooth lower order derivatives.

Conditions on the accuracy ϵ . In Theorem 2 we require that $\epsilon^{q-1} \leq L_1^q / L_q$ for all $q \in \{2, \dots, p\}$. For each q , we may rewrite this as $L_q^{1/q} \Delta \epsilon^{-(1+q)/q} \leq L_1 \Delta \epsilon^{-2}$. In other words, these conditions ensure that q th order regularization-based methods have stronger convergence guarantees than gradient descent [5, 10].

The case $p = 1$. We state our bounds in Theorem 2 for $p \geq 2$. It is possible to use the construction (9) to prove a lower bound of $O(\Delta L_1 \epsilon^{-2})$ on the time necessary for a deterministic first-order algorithm to find an ϵ -stationary point for the class $\mathcal{F}_1(\Delta, L_1)$. As Theorem PI.2 shows this lower bound holds for all randomized high-order algorithms, we do not pursue this.

The case $p = 3$. We can slightly strengthen our lower bound in the case $p = 3$, making it independent of L_2 for sufficiently small ϵ . To achieve this we set $r = 1$ in the definition of Υ_r , take $\lambda \mu \propto L_3 \sigma^2$, and argue that the resulting construction has $O(\sigma)$ -Lipschitz continuous Hessian, and σ tends to zero as $\epsilon \rightarrow 0$. For sufficiently small ϵ , we can then replace the minimum over $q \in \{2, 3\}$ in the first claim of Theorem 2 with $L_1^{2/5} L_3^{1/5} \epsilon^{-8/5}$.

Hard instance dimensionality. Let $T_{p,\epsilon}$ be such that $\mathcal{T}_\epsilon(\mathcal{A}_{\text{det}}^{(1)} \cup \mathcal{A}_{\text{zr}}^{(1)}, \mathcal{F}_{1:p}(\Delta, L_1, \dots, L_p)) \geq T_{p,\epsilon}$ according to Theorem 2. We construct functions $f \in \mathcal{F}_{1:p}(\Delta, L_1, \dots, L_p)$ that witness this lower bound and have domain of dimension $2T_{p,\epsilon}$; see Sect. 2.4 and Section PI.1.2 for more details.

The commentary on Theorem PI.1 in Section PI.4.2 is relevant also to Theorem 2. In particular, there we discuss the polynomial scaling of $\ell_q^{1/q}$ in q .

5.1 Lower bounds based on distance to optimality

For convex optimization problems, typical convergence guarantees depend on the distance of the initial point to the globally optimal set $\arg \min_x f(x)$; the dependence on this distance may be polynomial for general convex optimization problems [22, 23], while for smooth strongly convex problems, the convergence guarantees depend only logarithmically on it. In the non-convex case, we can provide lower bounds that depend on the distance rather than the gap $\Delta := f(x^{(0)}) - \inf_x f(x)$. To that end, we consider the class

$$\mathcal{F}_{1:p}^{\text{dist}}(D, L_1, \dots, L_p)$$

functions with L_q -Lipschitz q th derivatives (for each $q \in [p]$) and all global minima x^* satisfying $\|x^*\| \leq D$. We obtain the bound, analogously to our results in Section PI.6, by “hiding” a sharp global minimum near the origin.

To state the theorem, we require an additional piece of notation. Let $B_\epsilon(\Delta, L_1, \dots, L_p)$ be the lower bound Theorem 2 provides on $\mathcal{T}_\epsilon(\mathcal{A}_{\text{det}} \cup \mathcal{A}_{\text{zr}}^{(1)}, \mathcal{F}_{1:p}(\Delta, L_1, \dots, L_p))$, so

$$\mathcal{T}_\epsilon(\mathcal{A}_{\text{det}} \cup \mathcal{A}_{\text{zr}}^{(1)}, \mathcal{F}_{1:p}(\Delta, L_1, \dots, L_p)) \geq B_\epsilon(\Delta, L_1, \dots, L_p),$$

where we take $B_\epsilon = 1$ if $\epsilon > 0$ is larger than the settings Theorem 2 requires. Then by a reduction from our lower bounds on the complexity of $\mathcal{F}_{1:p}(\Delta, L_1, \dots, L_p)$, we obtain the following result.

Theorem 3 *There exists a numerical constant $c < \infty$ such that the following lower bound holds. Let $p \geq 2$, $p \in \mathbb{N}$, and let D, L_1, L_2, \dots, L_p , and ϵ be positive. Then*

$$\mathcal{T}_\epsilon(\mathcal{A}_{\text{det}}^{(1)} \cup \mathcal{A}_{\text{zr}}^{(1)}, \mathcal{F}_{1:p}^{\text{dist}}(D, L_1, \dots, L_p)) \geq B_\epsilon \left(\min_{q \in [p]} \left\{ \frac{L_q}{2\tilde{\ell}_q} D^{q+1} \right\}, \frac{L_1}{2}, \frac{L_2}{2}, \dots, \frac{L_p}{2} \right),$$

where $\tilde{\ell}_q \leq \exp(cq \log q + c)$.

We prove Theorem 3 in Appendix B.3. Theorem 3 shows that the lower bounds of Theorem 2 apply almost identically (to constant factors), except that we replace the function gap Δ in the lower bound with the quantity $\min_{q \in [p]} L_q D^{q+1}$. As the dependence of the lower bound on ϵ does not change, distance-based assumptions seem unlikely to help in the design of efficient optimization algorithms for non-convex functions.

5.2 Proof of Theorem 2

We have five parameters with which to scale our hard function; the function $\bar{f}_{T,\mu,r}$ requires definition of the dimension $T \in \mathbb{N}$, multiplier $\mu \leq 1$ on the Υ_r terms, and scalar $r \geq 1$ that trades between higher order ($r = \infty$) smoothness and lower order ($r = 1$) smoothness of Υ_r . We additionally scale the function with $\lambda > 0$ and a perspective term $\sigma > 0$, defining

$$f(x) := \lambda\sigma^2 \bar{f}_{T,\mu,r}(x/\sigma). \quad (11)$$

We must choose these parameters to guarantee the membership

$$f \in \mathcal{F}_{1:p}(\Delta, L_1, \dots, L_p).$$

This containment requires both bounded function values and derivatives, for which we can provide sufficient conditions. Recall the definition $\ell_p \leq e^{\frac{3}{2}p \log p + cp}$ from Lemma 4 of the smoothness constant of $\bar{f}_{T,\mu,r}$. Then by Lemma 4.ii, to guarantee that f has L_q -Lipschitz q th order derivatives for every $q \in [p]$ it suffices to choose λ, r, σ , and μ such that

$$(1_{(q=1)} + r^{3-q}\sigma^{1-q}\mu)\ell_q\lambda \leq L_q \text{ for every } q \in [p]. \quad (12)$$

For the bounded values constraint $f(0) - \inf_x f(x) \leq \Delta$, by Lemma 4.i it suffices to take

$$T = \left\lfloor \frac{\Delta - \lambda\sqrt{\mu}\sigma^2/2}{10\lambda\mu\sigma^2} \right\rfloor \quad (13)$$

Thus, so long as we choose the constants μ, σ, λ, r to satisfy inequality (12), the preceding choice of T guarantees $f \in \mathcal{F}_{1:p}(\Delta, L_1, \dots, L_p)$.

With this membership guaranteed, we consider the choices for λ , μ , and σ such that after T iterations of a zero respecting first-order method, we have $\|\nabla f(x)\| \geq \epsilon$. Indeed, Observations 1 and 2 imply that if $x^{(1)} = 0, x^{(2)}, \dots$ are the sequence of iterates produced by applying any zero-respecting (first-order) method to f , then $x_T^{(t)} = x_{T+1}^{(t)} = 0$ for all $t \leq T$. Lemma 3 implies that $\|\nabla \bar{f}_{T,\mu,r}(x^{(t)}/\sigma)\| > \mu^{3/4}/4$ for any such iterate. Therefore, if we choose $\lambda > 0$, $\mu \leq 1$, and $\sigma > 0$ such that

$$\lambda \mu^{3/4} \sigma \geq 4\epsilon, \quad (14)$$

then $\|\nabla f(x^{(t)})\| = \lambda \sigma \|\nabla \bar{f}_T(x^{(t)}/\sigma)\| > \lambda \mu^{3/4} \sigma / 4 \geq \epsilon$ for all $t \leq T$. We thus obtain the guarantee

$$\mathsf{T}_\epsilon(\mathcal{A}_{\text{zf}}^{(1)}, \mathcal{F}_{1:p}(\Delta, L_1, \dots, L_p)) \geq \inf_{A \in \mathcal{A}_{\text{zf}}^{(1)}} \mathsf{T}_\epsilon(A, f) \geq T+1 \geq \frac{\Delta - \lambda \sqrt{\mu \sigma^2}/2}{10\lambda \mu \sigma^2}. \quad (15)$$

The same bound for the class $\mathcal{A}_{\text{det}}^{(1)}$ then follows from Proposition 1. Our strategy is now the obvious one: we select $\lambda > 0$, $0 < \mu \leq 1$, $r \geq 1$, and $\sigma > 0$ to satisfy the function membership constraints (12) and the large gradient guarantee (14). Substituting our choices into the bound (15) will then yield the lower bound in the theorem. We begin with the general case $p \geq 2$ and later provide a tighter construction for $p = 2$.

General smoothness orders. To simplify the derivation, we define, for any $q \in [p]$

$$\bar{L}_q := L_q / \bar{\ell}_q \quad \text{where} \quad \bar{\ell}_q := \begin{cases} 2\ell_1 & q = 1 \\ \max \{ \ell_q, 4^{q-1} (2\ell_1)^q \} & q > 1, \end{cases} \quad (16)$$

where we note that $\bar{\ell}_q \leq e^{\frac{3}{2}q \log q + cq}$ for some numerical constant $c < \infty$, as $\ell_q \leq e^{\frac{3}{2}q \log q + cq}$ for a (possibly different) numerical constant. In order to further simplify our calculations, we then define

$$\lambda = \bar{L}_1, \quad \bar{\mu} := \frac{\lambda \mu}{\sigma^2} \quad \text{and} \quad \bar{r} := \sigma r. \quad (17)$$

Substituting these definitions into the constraints (12), we see that our choice of $\bar{\ell}_1 = 2\ell_1$ implies that the constraint (12) holds whenever

$$\bar{r}^{3-q} \bar{\mu} \leq \bar{L}_q \quad \text{for all } q \in [p]. \quad (18)$$

We choose \bar{r} and $\bar{\mu}$ to guarantee that f is appropriately smooth; in the sequel, we will choose σ and λ so that the gradient bound condition (14) holds. In this sense, we may choose \bar{r} and $\bar{\mu}$ without consideration of ϵ . Taking $\bar{r} = (\bar{L}_1/\bar{\mu})^{1/2}$ guarantees the inequality (18) holds for $q = 1$. Substituting this choice into the identical inequality for $q \in \{2, \dots, p\}$ shows that we must have $\bar{\mu}^{(q-1)/2} \leq \bar{L}_q \bar{L}_1^{(q-3)/2}$ for each such q . Thus, the choice

$$\bar{\mu} = \bar{L}_1 \min_{q \in \{2, \dots, p\}} (\bar{L}_q / \bar{L}_1)^{2/(q-1)}$$

satisfies inequality (18), and consequently, the smoothness condition (12) as well. We may therefore write $\bar{\mu}$ and \bar{r} as

$$\bar{\mu} = \bar{L}_1 \left(\bar{L}_{q^*} / \bar{L}_1 \right)^{\frac{2}{q^*-1}} \text{ and } \bar{r} = \left(\bar{L}_1 / \bar{L}_{q^*} \right)^{\frac{1}{q^*-1}}, \text{ where } q^* := \arg \min_{q \in \{2, \dots, p\}} \left(\bar{L}_q / \bar{L}_1 \right)^{\frac{1}{q-1}}.$$

It remains to choose λ, σ , depending on ϵ , to guarantee our gradient lower bound condition (14) holds, i.e. $4\epsilon \leq \lambda \mu^{3/4} \sigma = \lambda^{1/4} \bar{\mu}^{3/4} \sigma^{5/2}$. We thus set

$$\sigma = \left[4\bar{L}_1^{-1} \left(\bar{L}_{q^*} / \bar{L}_1 \right)^{-\frac{3}{2(q^*-1)}} \epsilon \right]^{2/5}.$$

We can now substitute back into our definitions $r = \bar{r}/\sigma$ and $\mu = \bar{\mu}\sigma^2/\lambda$ in Eq. (17) and verify that $r \geq 1$ and $\mu \leq 1$. For r , we have

$$r = \frac{\bar{r}}{\sigma} = \left(\frac{\bar{L}_1^{q^*/(q^*-1)}}{4\bar{L}_{q^*}^{1/(q^*-1)} \epsilon} \right)^{2/5} = \left(\frac{\bar{\ell}_{q^*}}{4^{q^*-1} \bar{\ell}_1^{q^*}} \right)^{\frac{2}{5(q^*-1)}} \left(\frac{\bar{L}_1^{q^*/(q^*-1)}}{L_{q^*}^{1/(q^*-1)} \epsilon} \right)^{2/5} \geq 1,$$

where the last transition uses $\bar{\ell}_{q^*} \geq 4^{q^*-1} \bar{\ell}_1^{q^*}$ by the definition (16), and we used the assumption in the theorem statement that $\epsilon \leq L_1^{q^*/(q^*-1)} / L_{q^*}^{1/(q^*-1)}$. Similarly, our choice of μ satisfies

$$\mu = \frac{\bar{\mu}\sigma^2}{\lambda} = \left(\frac{4^{q^*-1} \bar{\ell}_1^{q^*}}{\bar{\ell}_{q^*}} \right)^{\frac{4}{5(q^*-1)}} \left(\frac{L_{q^*}^{1/(q^*-1)} \epsilon}{L_1^{q^*/(q^*-1)}} \right)^{4/5} \leq 1.$$

We now consider two cases; $\lambda\sqrt{\mu\sigma^2} \leq \Delta$ and $\lambda\sqrt{\mu\sigma^2} > \Delta$. In the first case (which holds for sufficiently small ϵ), we substitute our choices of σ, λ and μ into the time lower bound (15),

$$\begin{aligned} T + 1 &\geq \frac{\Delta - \lambda\sqrt{\mu\sigma^2}/2}{10\lambda\mu\sigma^2} \stackrel{(i)}{\geq} \frac{\Delta}{20\lambda\mu\sigma^2} = \frac{\Delta}{20\bar{\mu}\sigma^4} \\ &= \frac{\Delta \bar{L}_1^{\frac{3}{5}} \left[\left(\bar{L}_{q^*} / \bar{L}_1 \right)^{\frac{1}{q^*-1}} \right]^{\frac{2}{5}}}{20 \cdot 4^{8/5} \cdot \epsilon^{8/5}} \\ &= \frac{\Delta}{20 \cdot 4^{8/5}} \\ &\quad \cdot \min_{q \in \{2, \dots, p\}} \left\{ \left(L_1 / \bar{\ell}_1 \right)^{\frac{3}{5} - \frac{2}{5(q-1)}} \left(L_q / \bar{\ell}_q \right)^{\frac{2}{5(q-1)}} \right\} \epsilon^{-8/5}, \end{aligned}$$

which is the desired bound, where in step (i) we made use of $\lambda\sqrt{\mu\sigma^2} \leq \Delta$. When $\lambda\sqrt{\mu\sigma^2} > \Delta$, we show that the above bound is in fact smaller than the convex lower

bound in Theorem 1. Indeed, substituting in our choices of λ, μ and σ , we see that $\lambda\sqrt{\mu}\sigma^2 > \Delta$ implies

$$\begin{aligned}\Delta &< \bar{L}_1^{3/2} \left(\bar{L}_{q^*}/\bar{L}_1 \right)^{\frac{1}{q^*-1}} \left[4\bar{L}_1^{-1} \left(\bar{L}_{q^*}/\bar{L}_1 \right)^{-\frac{3}{2(q^*-1)}} \epsilon \right]^{6/5} \\ &= 4^{\frac{6}{5}} \bar{L}_1^{-\frac{1}{5}} \left(\bar{L}_{q^*}/\bar{L}_1 \right)^{-\frac{4}{5(q^*-1)}} \epsilon^{\frac{6}{5}}.\end{aligned}$$

Taking a square root and substituting to our lower bound gives

$$\frac{\Delta \bar{L}_1^{\frac{3}{5}} \left[\left(\bar{L}_{q^*}/\bar{L}_1 \right)^{\frac{1}{q^*-1}} \right]^{\frac{2}{5}}}{20 \cdot 4^{8/5} \cdot \epsilon^{8/5}} < (80 \cdot \bar{\ell}_1^{1/2})^{-1} \frac{\sqrt{\Delta L_1}}{\epsilon},$$

and therefore (recalling that $\mathcal{Q}(\Delta, L_1) \subset \mathcal{F}_{1:p}(\Delta, L_1, \dots, L_p)$ and that $\bar{\ell}_1 \geq 8$), by Theorem 1 we have

$$\begin{aligned}\mathcal{T}_\epsilon(\mathcal{A}_{\text{zr}}^{(1)}, \mathcal{F}_{1:p}(\Delta, L_1, \dots, L_p)) &\geq \mathcal{T}_\epsilon(\mathcal{A}_{\text{zr}}^{(1)}, \mathcal{Q}(\Delta, L_1)) \\ &\geq \frac{\sqrt{\Delta L_1}}{4\epsilon} \\ &\geq \frac{\Delta \bar{L}_1^{\frac{3}{5}} \left[\left(\bar{L}_{q^*}/\bar{L}_1 \right)^{\frac{1}{q^*-1}} \right]^{\frac{2}{5}}}{20 \cdot 4^{8/5} \cdot \epsilon^{8/5}},\end{aligned}$$

completing the proof in the general case.

Functions with Lipschitz Hessian. For $p = 2$, we keep the definitions (16) but replace the particular rescaling choices (17) with

$$\lambda = \bar{L}_1, \quad \mu = \frac{\bar{L}_2\sigma}{\lambda} \quad \text{and} \quad r = 1.$$

Using $\mu \leq 1$, the above parameter setting satisfies inequality (12); f has L_q -Lipschitz q th-order derivatives for $q = 1, 2$. To satisfy the gradient lower bound (14), i.e. $4\epsilon \leq \lambda\mu^{3/4}\sigma = \bar{L}_1^{1/4}\bar{L}_2^{3/4}\sigma^{7/4}$, we set

$$\sigma = \left[4\bar{L}_1^{-1/4}\bar{L}_2^{-3/4}\epsilon \right]^{4/7}.$$

We can substitute into the definition $\mu = \frac{\bar{L}_2\sigma}{\lambda}$ to verify that $\mu \leq 1$:

$$\mu = \frac{\bar{L}_2\sigma}{\lambda} = \left(\frac{4^{2-1}(\bar{\ell}_1)^2}{\ell_2} \right)^{4/7} \left(\frac{L_2\epsilon}{L_1^2} \right)^{4/7} \leq 1$$

by the definition (16) of $\bar{\ell}_p$ and the assumption $\epsilon \leq L_1^2/L_2$. As in the general case, we first assume $\lambda\sqrt{\mu\sigma^2} \leq \Delta$, where substituting into (15) yields the desired lower bound

$$T + 1 \geq \frac{\Delta - \lambda\sqrt{\mu\sigma^2}/2}{10\lambda\mu\sigma^2} \geq \frac{\Delta}{20\lambda\mu\sigma^2} = \frac{\Delta}{20\bar{L}_2\sigma^3} = \frac{\Delta\bar{L}_1^{\frac{3}{7}}\bar{L}_2^{\frac{2}{7}}}{20 \cdot 4^{12/7} \cdot \epsilon^{12/7}}.$$

If $\lambda\sqrt{\mu\sigma^2} \leq \Delta$ does not hold, we have

$$\Delta < \bar{L}_1^{1/2}\bar{L}_2^{1/2} \left[4\bar{L}_1^{-1/4}\bar{L}_2^{-3/4}\epsilon \right]^{\frac{4}{7} \cdot \frac{5}{2}} < 4^{\frac{10}{7}}\bar{L}_1^{-\frac{1}{7}}\bar{L}_2^{-\frac{2}{7}}\epsilon^{\frac{10}{7}}.$$

Taking a square root and substituting to our lower bound gives

$$\begin{aligned} \frac{\Delta\bar{L}_1^{\frac{3}{7}}\bar{L}_2^{\frac{2}{7}}}{20 \cdot 4^{12/7} \cdot \epsilon^{12/7}} &< (20 \cdot \bar{\ell}_1^{1/2})^{-1} \frac{\sqrt{\Delta L_1}}{\epsilon} < \mathcal{T}_\epsilon(\mathcal{A}_{\text{zr}}^{(1)}, \mathcal{Q}(\Delta, L_1)) \\ &\leq \mathcal{T}_\epsilon(\mathcal{A}_{\text{zr}}^{(1)}, \mathcal{F}_{1:2}(\Delta, L_1, L_2)), \end{aligned}$$

due to Theorem 1, establishing the case $p = 2$.

6 The challenge of strengthening Theorem 2

The lower bounds in Theorem 2 leave two avenues for improvement. The first is tightening our $\epsilon^{-12/7}$ and $\epsilon^{-8/5}$ lower bounds to match the known upper bounds of $\epsilon^{-7/4}$ and $\epsilon^{-5/3}$, for $p = 2$ and $p = 3$, respectively. The second improvement is to extend our lower bounds to randomized algorithms, as we did for the case of full derivative information in Section PI.5. We discuss each of these in turn.

6.1 Tightness of lower bound construction

The core of our first-order lower bounds is Lemma 3, which establishes a lower bound of the form $\|\bar{f}_{T,\mu,r}(x)\| > \mu^{3/4}/4$ for vectors x such that $x_T = x_{T+1} = 0$ (i.e. any point that a first-order zero-respecting method can produce after T iterations), where $\bar{f}_{T,\mu,r}$ is our unscaled hard instance (see Definition (9)). Here we consider a slightly more general form,

$$\widetilde{f}_{T,\alpha,\mu}(x) := \alpha \cdot \Lambda(x_1 - 1) + \sum_{i=1}^T \Lambda(x_{i+1} - x_i) + \mu \sum_{i=1}^T \widetilde{\Upsilon}(x_i), \quad (19)$$

where $\Lambda : \mathbb{R} \rightarrow \mathbb{R}$ and $\widetilde{\Upsilon} : \mathbb{R} \rightarrow \mathbb{R}$ are \mathcal{C}^∞ , and we assume $T \in \mathbb{N}$, $\alpha > 0$, and $0 < \mu \leq 1$. The chain $\bar{f}_{T,\mu,r}$ corresponds to the special case $\alpha = \sqrt{\mu}$, $\Lambda(x) = x^2/2$ and $\widetilde{\Upsilon}(x) = \Upsilon_r(x)$ (defined in Eq. (10)).

We claim that if we can show that the norm of $\nabla \tilde{f}_{T,\alpha,\mu}(x)$ is *not too large* for some $x \in \mathbb{R}^{T+1}$ with $x_T = x_{T+1} = 0$, then our lower bound cannot be improved. More concretely, suppose that for every T , $\mu \leq 1$ and $\alpha \leq 1$ we could find $x \in \mathbb{R}^{T+1}$ such that $x_T = x_{T+1} = 0$, and $\|\nabla \tilde{f}_{T,\alpha,\mu}(x)\| \leq C\mu^{3/4}$ for some constant C independent of T, r and μ , matching Lemma 3 to a constant. We can then trace the scaling arguments in the proof of Theorem 2 “in reverse,” showing that any choice of T, λ, μ, σ and r for which the function $f(x) = \lambda\sigma^2 \tilde{f}_{T,\alpha,\mu}(x/\sigma)$ satisfies both (i) $f \in \mathcal{F}_{1:p}(\Delta, L_1, \dots, L_p)$ and (ii) $\|\nabla f(x)\| > \epsilon$ for all x such that $x_T = x_{T+1} = 0$, we have $T \leq c \cdot \epsilon^{-8/5}$ for $p \geq 3$ and $T \leq c \cdot \epsilon^{-12/7}$ for $p = 2$, where c is some problem-dependent constant independent of ϵ .

The next lemma, whose proof we provide in Appendix B.4, shows such gradient norm upper bound for constructions of the form (19).

Lemma 5 *Let $T \in \mathbb{N}$, $0 < \alpha \leq 1$, $\mu \in [T^{-2}, 1]$ and $\tilde{f}_{T,\alpha,\mu}$ be defined as in (19), with Λ and $\tilde{\Upsilon}$ satisfying*

$$\Lambda'(0) = \tilde{\Upsilon}'(0) = 0 \text{ and } \Lambda' \text{ is 1-Lipschitz continuous and } \max_{z \in [0,1]} |\tilde{\Upsilon}'(z)| \leq G,$$

for $G > 0$ independent of T, α and μ . Then there exists $x \in \mathbb{R}^{T+1}$ such that $x_T = x_{T+1} = 0$ and

$$\|\nabla \tilde{f}_{T,\alpha,\mu}(x)\| < C\mu^{3/4},$$

where $C \leq 27 + \sqrt{3}G$.

Let us discuss the lemma. The condition that $\Lambda'(0) = \tilde{\Upsilon}'(0) = 0$ is essential for any zero-chain-based proof, as otherwise $\tilde{f}_{T,\alpha,\mu}$ is not a first-order zero-chain (if $\alpha = 1$ then we may have $\Lambda'(0) \neq 0$; Lemma 5 holds in this case as well). The requirement that the multiplier $\mu \geq 1/T^2$ on $\tilde{\Upsilon}$ is also benign, as our proofs require $\mu \gtrsim 1/\sqrt{T} \gg 1/T^2$ (further decreasing μ weakens the lower bound as it makes $\tilde{f}_{T,\alpha,\mu}$ too smooth; inspection of the scaling argument in the proof of Theorem 2 shows this rigorously). The function Λ must have Lipschitz derivatives with parameter independent of μ, T , as otherwise $\tilde{f}_{T,\alpha,\mu}$ cannot be scaled to meet the smoothness requirements. Finally, the requirement $\max_{z \in [0,1]} |\tilde{\Upsilon}'(z)| < \infty$ holds for every \mathcal{C}^∞ function. Moreover, a calculation shows it holds with $G = \sqrt{10\ell_3}$ independent of r for every Υ_r that satisfies Lemma 2.

Summarizing, tightening our lower bounds seems to require a construction that is not of the form (19). This does not eliminate more general (non-convex) interactions, e.g. of the form $\Lambda(x_i, x_{i+1})$ rather than $\Lambda(x_{i+1} - x_i)$. The proof technique of Lemma 3 should provide useful “sanity checks” when considering alternative constructions.

6.2 A bound for randomized algorithms

In Section P1.5 we extend our lower bound for arbitrary deterministic local algorithms to the broader class of randomized algorithms $\mathcal{A}_{\text{rand}}$ with access to all derivatives at query point x . We do this by making our hard function insensitive: the individual “linking” terms $\Psi(x_i)\Phi(x_{i+1})$ (analogues of the terms $(x_i - x_{i+1})^2$ in constructions (2)

and (9)) are identically zero for x_i near 0. A natural question is whether the same methodology (originally proposed in [30]) can extend Theorem 2 to the class of randomized first-order algorithms, $\mathcal{A}_{\text{rand}}^{(1)}$. Direct application of that technique cannot work in our case, for a simple reason: it applies to randomized algorithms of *any* order. In other words, if we modify our hard instance construction (9) to be a robust zero-chain (Definition PI.4), any lower bounds it implies hold for all algorithms in $\mathcal{A}_{\text{rand}}$, where $\epsilon^{-(p+1)/p}$ rates are achievable, so we could not provide sharper lower bounds than Theorem PI.2.

Nevertheless, the ideas introduced in Section PI.5 might still be of use. Specifically, consider a modification of the construction (9) where $\Upsilon_r(x)$ is identically zero for sufficiently small x , say $|x| < 0.05$, while still satisfying Lemma 2, thus making the non-convex component of $\bar{f}_{T,\mu,r}$ insensitive. As explained above, also making the convex quadratic component of $\bar{f}_{T,\mu,r}$ insensitive (as Woodworth and Srebro [30] do) is unworkable in our setting, as it results in a robust zero-chain equally hard for all high-order algorithms. Instead, we may keep the quadratic component unchanged—and hence sensitive—and try to carry out the proof of Lemma PI.4. Doing so, we see that the inductive argument allows us to ignore the insensitive non-convex component of $\bar{f}_{T,\mu,r}$, leaving us to contend only with the (randomly rotated) quadratic chain. Thus, the difficulty here appears closely related to proving a lower bound for minimizing convex quadratic with randomized first-order methods.

In recent work, Simchowitz [27] proves such a lower bound. However, his proof technique is quite different, relying on reductions to statistical estimation of a planted vector in a deformed Wigner model rather than a random rotation of the quadratic chain. Since the analysis of the non-convexity in our construction relies closely on the chain argument's fine-grained control over information flow, it is unclear how to extend the reduction in [27] to finding stationary points of non-convex functions. We believe that proving a lower bound for convex quadratic minimization using the rotated-chain approach could lead to randomized lower bounds in our setting as well.

7 Concluding remarks

Here we discuss implications of our two-paper series and provide a few possible extensions.

7.1 Commentary on our results

In conjunction with known upper bounds, our lower bounds characterize the optimal rates for finding stationary points. Our lower bounds are sharp to within constant factors for algorithms with full derivative information [10], and (perhaps) slightly loose for first-order algorithms. These characterizations yield a few insights.

First-order methods vs. high-order methods. For the class $\mathcal{F}_1(\Delta, L_1)$ of L_1 -smooth functions, first-order methods—specifically gradient descent—attain the optimal rate $L_1 \Delta \epsilon^{-2}$; no higher-order randomized method can attain improved performance over the entire function class. The intuition here is that $\mathcal{F}_1(\Delta, L_1)$ contains functions whose

Hessian and higher order derivatives may vary arbitrarily sharply, providing no useful information for optimization.

When higher-order derivatives are also Lipschitz continuous the picture changes fundamentally: there is a strict separation between (deterministic) second-order and first-order methods. In particular, cubic regularization of Newton's method [25] achieves ϵ dependence $\epsilon^{-3/2}$ for functions with Lipschitz Hessian, while *no* deterministic first-order method can have better time complexity than $\epsilon^{-8/5}$, regardless of how many derivatives are Lipschitz. Note that when the Hessian is Lipschitz, our definition of first-order algorithms allows for algorithms that rely on Hessian-vector products, as they can be estimated to arbitrary accuracy in two gradient evaluations.

The effect of high-order smoothness on first-order methods. For $\mathcal{F}_{1:2}(\Delta, L_1, L_2)$, the class of functions with Lipschitz gradient and Hessian, our lower bound scales as $\epsilon^{-12/7}$, while for the class $\mathcal{F}_{1:3}(\Delta, L_1, L_2, L_3)$ of functions with Lipschitz third order derivative our “convex until proven guilty” method [8] achieves the rate $\epsilon^{-5/3} \log \frac{1}{\epsilon}$. As $\frac{5}{3} < \frac{12}{7}$, this proves a separation between the optimal rate for first-order methods with second- and third-order smoothness.

In contrast, orders of smoothness beyond the third offer limited room for improvement in ϵ dependence; the lower bound $\epsilon^{-8/5}$ holds for all function classes $\mathcal{F}_{1:p}(\Delta, L_1, \dots, L_p)$ with $p \geq 3$, while the method [8] does not enjoy improved guarantees with Lipschitz fourth-order derivatives. The “robustness” of the lower bound to higher-order smoothness stems from the fact that our hard instance $\bar{f}_{T,\mu,r}$ becomes a quartic polynomial in the limit $r \rightarrow \infty$, and we choose r inversely proportional to ϵ . As we discuss in [8, Lemma 4], our guarantee $\epsilon^{-5/3} \log \frac{1}{\epsilon}$ cannot improve using fourth-order smoothness because of symmetries in the fourth-order Taylor expansion. Quartic polynomials thus appear to play a central role in the complexity of first-order methods for smooth optimization. Due to the gap between the upper and lower bounds, it remains an open question whether smoothness beyond the third order has any effect the complexity of finding stationary point with first-order functions.

Convex vs. non-convex functions. Convexity makes finding stationary points fundamentally—and significantly—easier. For first-order methods and functions with bounded initial sub-optimality, the rate $\epsilon^{-1} \log \frac{1}{\epsilon}$ is achievable for first-order smooth convex functions, while the lower bound $\epsilon^{-8/5}$ holds for non-convex functions with arbitrarily high-order smoothness. For methods using higher-order derivatives, our lower bounds [10] for finding stationary points of non-convex functions are $\epsilon^{-(p+1)/p} \rightarrow \epsilon^{-1}$ as the order p of smoothness grows. However, similar to Appendix A.1, the papers [4, 21] show that for convex functions with Lipschitz Hessian, a second-order method achieves the strictly better rate $\epsilon^{-6/7} \log \frac{1}{\epsilon}$.

Another striking difference between convex and non-convex functions is the effect of replacing the bound on the initial function value (i.e. $f(x^{(0)}) - \inf_x f(x) \leq \Delta$) with a bound on the initial distance to the global minimizer x^* (i.e. $\|x^{(0)} - x^*\| \leq D$). For non-convex function classes, we show lower bounds with the same ϵ dependence regardless of the type of bound. In contrast, for convex function the optimal rates

scale as $\sqrt{\Delta}\epsilon^{-1}$ and $\sqrt{D}\epsilon^{-1/2}$, again a gap in ϵ dependence. The rates are not directly comparable; one can construct families of functions where D grows with the dimension while Δ remains constant.

Returning to Δ -value-bounded function classes, we see one more large difference between the convex and non-convex case; convex rates scale as $\sqrt{\Delta}$ while all the non-convex rates scale linearly with Δ . This arises from fundamental differences in the convergence ‘‘mechanism’’ for convex and non-convex optimization. The analysis of non-convex optimization schemes typically [8,11,23–25] revolves around a *progress argument*, where one shows that, as long as $\|\nabla f(x^{(t)})\| > \epsilon$, the guarantee $f(x^{(t+1)}) \leq f(x^{(t)}) - p_\epsilon$ holds for some quantity p_ϵ (e.g. for gradient descent $p_\epsilon = \epsilon^2/(2L_1)$). The number of iterations to find an ϵ -stationary point x_s is therefore at most $[f(x^{(0)}) - f(x_s)]/p_\epsilon \leq \Delta/p_\epsilon$, which scales linearly in Δ . By our lower bounds, such progress arguments are, in a sense, optimal. Conversely, in convex optimization we may control either the gap $f(x^{(t)}) - f(x^*)$ or the distance $\|x^{(t)} - x^*\|$, and this interplay (see Appendix A.1) allows stronger arguments than those based purely on function progress.

7.2 Further research

Closing the gap in first-order bounds. There exists a gap in polynomial ϵ dependence between our lower bounds (Theorem 2) and the best known upper bounds [8] for first-order methods with higher-order smoothness. We do not believe the upper bounds of [8] are improvable by different analysis or by any algorithmic change that maintains the general structure of alternating between accelerated gradient descent and negative curvature exploitation. In conjunction with our arguments in Sect. 6.1 about the structure of our lower bounds, resolution of the optimal rate will likely provide either a method with a substantially different approach to accelerating gradient descent in the smooth non-convex setting or a new lower bound construction.

Finite sum and stochastic problems. Smooth, non-convex, finite-sum and stochastic optimization problems are important, arising (for example) in the training of neural networks. This motivates the design and analysis of efficient methods for finding stationary points in such problems, and researchers have successfully developed variance reduction and acceleration techniques for these settings [2,3,20,26]. However, no corresponding lower bounds are available. Woodworth and Srebro [30], show how to establish lower bounds for *convex* finite sum problems. Combined with the developments in our paper, we believe their techniques should extend to finding stationary points of non-convex problems. An important conclusion of [30] is that randomized selection of the component function is crucial to efficient convergence: in contrast to our results, they show a separation between deterministic and randomized finite sum complexity.

Second-order stationary points. Approximate stationary points are not always close to local minima, and so it is interesting to consider stronger convergence guarantees. Second-order stationarity (also known as the second-order necessary condition for local optimality) is the most popular example; for a function f , a point x is (ϵ_1, ϵ_2) -

second-order stationary if $\|\nabla f(x)\| \leq \epsilon_1$ and $\nabla^2 f(x) \succeq -\epsilon_2 I$. Efficient first-order methods for finding second-order stationary points exist [1, 9, 19]. Moreover, it is possible to generically transform methods for finding ϵ -stationary points into methods that find $(\epsilon, O(\epsilon^s))$ -second-order stationary points, for some $0 < s < 1$, without changing the ϵ dependence of the complexity [8, Appendix C], but such modifications introduce dependence logarithmic in the problem dimension d .

Clearly, lower bounds for finding ϵ_1 -stationary points also apply to finding (ϵ_1, ϵ_2) -second-order stationary points. However, attaining second-order stationarity with first-order methods is fundamentally more difficult than attaining only stationarity. There are no dimension-free guarantees: the results of Simchowitz et al. [28] imply $\Omega(\log d)$ dimension dependence for all randomized first-order algorithms that escape saddle points. Moreover, for deterministic first-order algorithms it is easy to construct a resisting oracle that forces $\Omega(d)$ dimension dependence (consider $f(U^\top x)$ with $f(x) = -x_1^2$ and adversarially chosen rotation U), implying strong separation between deterministic and randomized first-order methods for finding second-order stationary points. It will be interesting to investigate such issues further.

A Additional results for convex functions

A.1 An upper bound for finding stationary points of value-bounded functions

Here we give a first-order method that finds ϵ -stationary points of a function $f \in \mathcal{K}_1(\Delta, L_1)$ in $O(\sqrt{L_1 \Delta} \epsilon^{-1} \log \frac{L_1 \Delta}{\epsilon^2})$ iterations. The method consists of Nesterov's accelerated gradient descent (AGD) applied on the sum of f and a standard quadratic regularizer.

Our starting point is AGD for *strongly convex* functions; a function f is σ -strongly convex if

$$f(y) \geq f(x) + \langle \nabla f(x), y - x \rangle + \frac{\sigma}{2} \|y - x\|^2,$$

for every x, y in the domain of f . Let $\text{AGD}_{\sigma, L_1} \in \mathcal{A}_{\text{zr}}^{(1)} \cap \mathcal{A}_{\text{det}}^{(1)}$ be the accelerated gradient scheme developed in [23, §2.2.1] for σ -strongly convex functions with L_1 -Lipschitz gradient, initialized at $x^{(1)} = 0$ (the exact step size scheme is not important). For any L_1 -smooth f with global minimizer x_f^* , $\epsilon^2/(2L_1)$ -suboptimality guarantees ϵ -stationarity, since $\|\nabla f(x)\|^2 \leq 2L_1(f(x) - f(x_f^*))$ [6, Eq. (9.14)]. Therefore, adapting [23, Thm. 2.2.2] to our notation gives

$$T_\epsilon(\text{AGD}_{\sigma, L_1}, f) \leq 1 + 2\sqrt{\frac{L_1}{\sigma}} \log_+ \left(\frac{L_1 \|x_f^*\|}{\epsilon} \right), \quad (20)$$

with $\log_+(x) := \max\{0, \log x\}$.

Now suppose that f is convex with L_1 -Lipschitz gradient but not necessarily strongly-convex. We can add strong convexity to f by means of a proximal term; for any $\sigma > 0$, the function

$$f_\sigma(x) := f(x) + \frac{\sigma}{2} \|x\|^2$$

is σ -strongly-convex with $(L_1 + \sigma)$ -Lipschitz gradient. With this in mind, we define a proximal version of AGD as follows,

$$\text{PAGD}_{\sigma, L_1}[f] := \text{AGD}_{\sigma, L_1 + \sigma}[f_\sigma] = \text{AGD}_{\sigma, L_1 + \sigma} \left[f(\cdot) + \frac{\sigma}{2} \|\cdot\|^2 \right].$$

Proposition 2 Let Δ, L_1 and ϵ be positive, and let $\sigma = \frac{\epsilon^2}{3\Delta}$. Then, algorithm $\text{PAGD}_{\sigma, L_1} \in \mathcal{A}_{\text{det}}^{(1)}$ satisfies

$$\begin{aligned} \mathcal{T}_\epsilon(\mathcal{A}_{\text{det}}^{(1)}, \mathcal{K}_1(\Delta, L_1)) &\leq \sup_{f \in \mathcal{K}_1(\Delta, L_1)} T_\epsilon(\text{PAGD}_{\sigma, L_1}, f) \\ &\leq 1 + 5 \frac{\sqrt{L_1 \Delta}}{\epsilon} \log_+ \left(\frac{25L_1 \Delta}{\epsilon^2} \right). \end{aligned}$$

Proof For any $f \in \mathcal{K}_1(\Delta, L_1)$, recall that $f_\sigma(x) := f(x) + \frac{\sigma}{2} \|x\|^2$ and let $\{x^{(t)}\}_{t \in \mathbb{N}} = \text{PAGD}_{\sigma, L_1}[f] = \text{AGD}_{\sigma, L_1 + \sigma}[f_\sigma]$ be the sequence of iterates $\text{PAGD}_{\sigma, L_1}$ produces on f . Then by guarantee (20), we have

$$\|\nabla f_\sigma(x^{(T)})\| \leq \epsilon/6 \quad (21)$$

for some T such that

$$T \leq 1 + 2 \sqrt{1 + \frac{L_1}{\sigma}} \log_+ \left(\frac{6(L_1 + \sigma) \|x_{f_\sigma}^*\|}{\epsilon} \right). \quad (22)$$

For any point y such that $f_\sigma(y) = f(y) + \frac{\sigma}{2} \|y\|^2 \leq f_\sigma(0) = f(0)$, we have

$$\|y\|^2 \leq \frac{2(f(0) - f(y))}{\sigma} \leq \frac{2(f(0) - \inf_x f(x))}{\sigma} \leq \frac{2\Delta}{\sigma}.$$

Clearly, $f_\sigma(x_{f_\sigma}^*) \leq f_\sigma(0)$ and [23, Thm. 2.2.2] also guarantees that $f_\sigma(x^{(T)}) \leq f_\sigma(0)$. Consequently,

$$\max \left\{ \|x^{(T)}\|, \|x_{f_\sigma}^*\| \right\} \leq \sqrt{\frac{2\Delta}{\sigma}}, \quad (23)$$

and so

$$\begin{aligned} \|\nabla f(x^{(T)})\| &= \|\nabla f_\sigma(x^{(T)}) - \sigma \cdot x^{(T)}\| \leq \|\nabla f_\sigma(x^{(T)})\| + \sigma \|x^{(T)}\| \stackrel{(i)}{\leq} \frac{\epsilon}{6} \\ &\quad + \sqrt{2\sigma \Delta} \stackrel{(ii)}{\leq} \epsilon. \end{aligned}$$

In inequality (i) we substituted bounds (21) and (23), and in (ii) we used $\sigma = \epsilon^2/(3\Delta)$. We conclude that $T_\epsilon(\text{PAGD}_{\sigma, L_1}, f) \leq T$, and substituting (23) and the definition of σ into (22) we have

$$T \leq 1 + 2\sqrt{1 + \frac{3L_1\Delta}{\epsilon^2}} \log_+ \left(6\sqrt{\frac{2}{3}} + \frac{6\sqrt{6}L_1\Delta}{\epsilon^2} \right).$$

Without loss of generality, we may assume $\frac{2L_1\Delta}{\epsilon^2} \geq 1$, as otherwise $T_\epsilon(\text{PAGD}_{\sigma, L_1}, f) = 1$. We thus simplify the expression slightly to obtain the proposition. \square

A.2 The impossibility of approximate optimality without a bounded domain

Lemma 6 *Let $L_1, \Delta > 0$ and $\epsilon < \Delta$. For any first-order algorithm $A \in \mathcal{A}_{\text{det}}^{(1)} \cup \mathcal{A}_{\text{zr}}^{(1)}$ and any $T \in \mathbb{N}$, there exists a function $f \in \mathcal{Q}(\Delta, L_1)$ such that the iterates $\{x^{(t)}\}_{t \in \mathbb{N}} = A[f]$ satisfy*

$$\inf_{t \in \mathbb{N}} \left\{ t \mid f(x^{(t)}) \leq \inf_x f(x) + \epsilon \right\} > T.$$

Proof By Proposition 1 it suffices to consider $A \in \mathcal{A}_{\text{zr}}^{(1)}$ (see additional discussion of the generality of Proposition 1 in Section PI.3.3). Consider the function $f : \mathbb{R}^T \rightarrow \mathbb{R}$,

$$f(x) = \lambda \left[(\sigma - \beta x_1)^2 + \sum_{i=1}^{T-1} (x_i - \beta x_{i+1})^2 \right], \quad (24)$$

where $0 < \beta < 1$, and we take

$$\lambda := \frac{L_1}{2(1+2\beta+\beta^2)} \text{ and } \sigma := \sqrt{\frac{\Delta}{\lambda}}.$$

Since $f(x)$ is of the form $\lambda \|Ax - b\|^2$ where $\|A\|_{\text{op}} \leq 1 + \beta$, we have $\|\nabla^2 f(x)\|_{\text{op}} \leq 2\lambda \|A\|_{\text{op}}^2$ for every $x \in \mathbb{R}^T$ and therefore f has $2\lambda(1+2\beta+\beta^2)$ -Lipschitz gradient. Additionally, f satisfies $\inf_x f(x) = 0$ and $f(0) = \lambda\sigma^2$, and so the above choices of λ and σ guarantee that $f \in \mathcal{Q}(\Delta, L_1)$. Moreover, f is a first-order zero-chain (Definition 3), and thus for any $A \in \mathcal{A}_{\text{zr}}^{(1)}$ and $\{x^{(t)}\}_{t \in \mathbb{N}} = A[f]$, we have $x_T^{(t)} = 0$ for $t \leq T$ (Observation 1). Therefore, it suffices to show that $f(x) > \inf_y f(y) + \epsilon$ whenever $x_T = 0$.

We make the following inductive claim: if $f(x) \leq \inf_y f(y) + \epsilon = \epsilon$, then

$$\left| x_i - \sigma\beta^{-i} \right| \leq \sum_{j=1}^i \beta^{-j} \sqrt{\frac{\epsilon}{\lambda}} < \frac{\beta^{-i}}{1-\beta} \sqrt{\frac{\epsilon}{\lambda}} \quad (25)$$

for all $i \leq T$. Indeed, each term in the sum (24) defining f is non-negative, so for the base case of the induction $i = 1$, we have $\lambda(\sigma - \beta x_1)^2 \leq \epsilon$, or $|x_1 - \sigma\beta^{-1}| \leq$

$\beta^{-1}\sqrt{\epsilon/\lambda}$. For $i < T$, assuming that x_i satisfies the bound (25), we have that $\lambda(x_i - \beta x_{i+1})^2 \leq \epsilon$, which implies

$$\begin{aligned} |x_{i+1} - \sigma\beta^{-(i+1)}| &\leq |x_{i+1} - \beta^{-1}x_i| + \beta^{-1}|x_i - \sigma\beta^{-i}| \\ &\leq \beta^{-1}\sqrt{\frac{\epsilon}{\lambda}} + \sum_{j=1}^i \beta^{-(j+1)}\sqrt{\frac{\epsilon}{\lambda}} \\ &= \sum_{j=1}^{i+1} \beta^{-j}\sqrt{\frac{\epsilon}{\lambda}}, \end{aligned}$$

which is the desired claim (25) for x_{i+1} .

The bound (25) implies $x_i \neq 0$ for all $i \leq T$ whenever $\sigma \geq (1 - \beta)^{-1}\sqrt{\epsilon/\lambda}$. Therefore, we choose β to satisfy $\sigma = (1 - \beta)^{-1}\sqrt{\epsilon/\lambda}$, that is

$$\beta := 1 - \sqrt{\frac{\epsilon}{\lambda\sigma^2}} = 1 - \sqrt{\frac{\epsilon}{\Delta}},$$

for which $0 < \beta < 1$ since we assume $\epsilon < \Delta$. Thus, we guarantee that when $x_T = 0$ we must have $f(x) > \inf_y f(y) + \epsilon$, giving the result. \square

B Technical results

B.1 Proof of Lemma 2

Lemma 2 *The function Υ_r satisfies the following.*

- i. *We have $\Upsilon'_r(0) = \Upsilon'_r(1) = 0$.*
- ii. *For all $x \leq 1$, $\Upsilon'_r(x) \leq 0$, and for all $x \geq 1$, $\Upsilon'_r(x) \geq 0$.*
- iii. *For all $x \in \mathbb{R}$ we have $\Upsilon_r(x) \geq \Upsilon_r(1) = 0$, and for all r , $\Upsilon_r(0) \leq 10$.*
- iv. *For every $r \geq 1$, $\Upsilon'_r(x) < -1$ for every $x \in (-\infty, -0.1] \cup [0.1, 0.9]$.*
- v. *For every $r \geq 1$ and every $p \geq 1$, the p -th order derivatives of Υ_r are $r^{3-p}\ell_p$ -Lipschitz continuous, where $\ell_p \leq \exp(\frac{3}{2}p \log p + cp)$ for a numerical constant $c < \infty$.*

Proof Parts i and ii are evident from inspection, as

$$\Upsilon'_r(x) = 120 \frac{x^2(x-1)}{1 + (x/r)^2}.$$

To see the part iii, note that Υ_r is non-increasing for every $x < 1$ and non-decreasing for every $x > 1$ and therefore $x = 1$ is its global minimum. That $\Upsilon_r(1) = 0$ is immediate from its definition, and, for every r , $\Upsilon_r(0) = 120 \int_0^1 \frac{t^2(1-t)}{1+(t/r)^2} dt \leq 120 \int_0^1 t^2(1-t) dt = 10$. To see part iv, note that $|\Upsilon'_r(x)| \geq |\Upsilon'_1(x)|$ for every $r \geq 1$, and a calculation shows $|\Upsilon'_1(x)| > 1$ for $x \in (-\infty, -0.1] \cup [0.1, 0.9]$ (see Fig. 1).

To see the fifth part of the claim, note that

$$\begin{aligned}\Upsilon'_r(x) &= 120r^2(x-1)\left(1 - \frac{1}{1+(x/r)^2}\right) \\ &= 120\left[r^2(x-1) - r^3\varphi_1(x/r) + r^2\varphi_2(x/r)\right],\end{aligned}$$

where the functions φ_1 and φ_2 are $\varphi_1(\xi) = \xi/(1+\xi^2)$ and $\varphi_2(\xi) = 1/(1+\xi^2)$. We thus bound the derivatives of φ_1 and φ_2 . We begin with φ_2 , which we can write as the composition $\varphi_2(x) = (h \circ g)(x)$ where $h(x) = \frac{1}{x}$ and $g(x) = 1+x^2$. Let $\mathcal{P}_{k,2}$ denote the collection of all partitions of $\{1, \dots, k\}$ where each element of the partition has at most 2 indices. That is, if $P \in \mathcal{P}_{k,2}$, then $P = (S_1, \dots, S_l)$ for some $l \leq k$, the S_i are disjoint, $1 \leq |S_i| \leq 2$, and $\cup_i S_i = [k]$. The cardinality $|\mathcal{P}_{k,2}|$ is the number of matchings in the complete graph on k vertices, or the k th telephone number, which has bound [15, Lemma 2]

$$|\mathcal{P}_{k,2}| \leq \exp\left(\frac{k}{2} \log k + k \log 2\right).$$

We may then apply Faà di Bruno's formula for the chain rule to obtain

$$\begin{aligned}\varphi_2^{(k)}(x) &= \sum_{P \in \mathcal{P}_k} h^{(|P|)}(g(x)) \prod_{S \in P} g^{(|S|)}(x) \\ &= \sum_{P \in \mathcal{P}_{k,2}} (-1)^{|P|} \frac{(|P|-1)!}{(1+x^2)^{|P|}} (2x)^{C_1(P)} 2^{C_2(P)},\end{aligned}$$

where $C_i(P)$ denotes the number of sets in P with precisely i elements. Of course, we have $|x|^{C_1(P)} / (1+x^2)^{|P|} \leq 1$, and thus

$$|\varphi_2^{(k)}(x)| \leq \sum_{P \in \mathcal{P}_{k,2}} (|P|-1)! 2^{|P|} \leq |\mathcal{P}_{k,2}| \cdot (k-1)! \cdot 2^k \leq e^{\frac{3k}{2} \log k + 2k \log 2}.$$

The proof of the upper bound on $\varphi_1^{(k)}(x)$ is similar ($2\varphi_1(x) = \frac{d}{dx}[(\hat{h} \circ g)(x)]$ with $\hat{h}(x) = \log x$ and g as defined above), so for every $r \geq 1$ and $p \geq 1$, the $p+1$ -th derivative of Υ_r has the bound

$$\begin{aligned}|\Upsilon_r^{(p+1)}(x)| &\leq 120\left[r^2 1_{(p=1)} + r^{3-p} |\varphi_1^{(p)}(x/r)| + r^{2-p} |\varphi_2^{(p)}(x/r)|\right] \\ &\leq 120r^{3-p} e^{\frac{3}{2} \log p + cp},\end{aligned}$$

where $c < \infty$ is a numerical constant. \square

B.2 Proof of Lemma 3

Lemma 3 *Let $r \geq 1$ and $\mu \leq 1$. For any $x \in \mathbb{R}^{T+1}$ such that $x_T = x_{T+1} = 0$,*

$$\|\nabla \tilde{f}_{T,\mu,r}(x)\| > \mu^{3/4}/4.$$

Proof Throughout the proof, we fix $x \in \mathbb{R}^{T+1}$ such that $x_T = x_{T+1} = 0$; for convenience in notation, we define $x_0 := 1$. Our strategy is to carefully pick two indices $i_1 \in \{0, \dots, T-1\}$ and $i_2 \in \{i_1+1, \dots, T\}$, such that $\|\nabla \tilde{f}_{T,\mu,r}(x)\|^2 \geq \sum_{i=i_1+1}^{i_2} |\nabla_i \tilde{f}_{T,\mu,r}(x)|^2 > (\mu^{3/4}/4)^2$. We call the set of indices from i_1+1 to i_2 the *transition region*, and construct it as follows.

Let $i_1 \geq 0$ be the largest i such that $x_i > 0.9$,

so that $x_j \leq 0.9$ for every $j > i_1$. Note that $i_1 = 0$ when $x_i \leq 0.9$ for every $i \in [T+1]$. This is a somewhat special case due to the coefficient $\sqrt{\mu} \leq 1$ of the first “link” in the quadratic chain term in (9). To handle it cleanly we define

$$\alpha := \begin{cases} 1 & i_1 > 0 \\ \sqrt{\mu} & i_1 = 0. \end{cases}$$

Continuing with construction of the transition region, we make the following definition.

Let $i'_2 \leq T$ be the smallest j such that $j > i_1$ and $x_j < 0.1$,

and let $m' = i'_2 - i_1$, so $m' \geq 1$. Roughly, our transition region consists of the m' indices i_1+1, \dots, i'_2 , but for technical reasons we attach to it the following decreasing ‘tail’.

Let i_2 be the smallest k such that $k \geq i'_2$ and $x_{k+1} \geq x_k - \frac{0.2}{m' - 1 + 1/\alpha} 1_{(x_k > -0.1)}$.

With these definitions, i_2 is well-defined and $0 \leq i_1 < i_2 \leq T$, since $x_{T+1} - x_T = 0$. We denote the transition region and associated length by

$$\mathcal{I}_{\text{trans}} := \{i_1+1, \dots, i_2\} \quad \text{and} \quad m := i_2 - i_1 \geq 1. \quad (26)$$

We illustrate our definition of the transition region in Fig. 3.

Let us describe the transition region. In the “head” of the region, we have $0.1 \leq x_i \leq 0.9$ for every $i \in \{i_1+1, \dots, i'_2-1\}$; a total of $m'-1$ indices. The “tail” of the transition region is strictly decreasing, $x_{i_2} < x_{i_2-1} < \dots < x_{i'_2}$. Moreover, for any $j \in \{i'_2+1, \dots, i_2-1\}$ such that $x_j > -0.1$, the decrease is rapid; $x_j < x_{j-1} - 0.2/(m' - 1 + 1/\alpha)$. This descriptions leads us to the following technical properties.

Lemma 7 *Let the transition region $\mathcal{I}_{\text{trans}}$ be defined as above (26). Then*

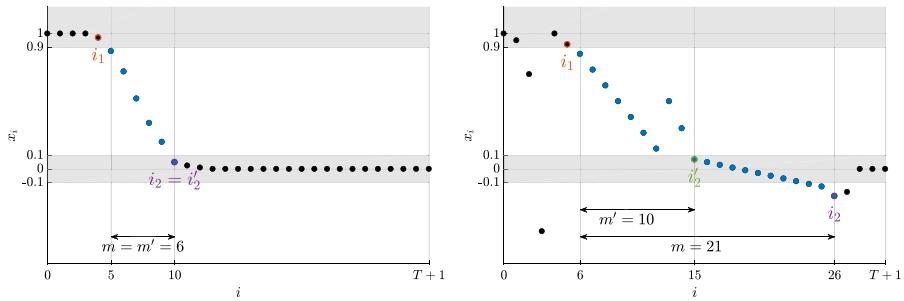


Fig. 3 The transition region (26) in the proof of Lemma 3. Each plot shows the entries of a vector $x \in \mathbb{R}^{T+1}$ that satisfies $x_T = x_{T+1} = 0$. The entries of x belonging to the transition region $\mathcal{I}_{\text{trans}}$ are blue (color figure online)

- $x_{i_1} > 0.9 > 0.1 > x_{i_2}$ and $-x_{i_2} + (m - 1 + \alpha^{-1})(x_{i_2+1} - x_{i_2}) > -0.3$.
- $\Upsilon'_r(x_i) \leq 0$ for every $i \in \mathcal{I}_{\text{trans}}$, and $\Upsilon'_r(x_i) < -1$ for at least $(m - \alpha^{-1})/2$ indices in $\mathcal{I}_{\text{trans}}$.

We defer the proof of the lemma to the end of this section, continuing the proof assuming it.

We now lower bound $\|\nabla \bar{f}_{T,\mu,r}(x)\|$. For notational convenience, define $g_i = \mu \Upsilon'_r(x_i)$, and recalling that $x_T = x_{T+1} = 0$, we see that the norm of the gradient of $\bar{f}_{T,\mu,r}$ is

$$\begin{aligned} \|\nabla \bar{f}_{T,\mu,r}(x)\|^2 &= ((1 + \sqrt{\mu})x_1 - \sqrt{\mu} - x_2 + g_1)^2 + \sum_{i=1}^T (2x_i - x_{i-1} - x_{i+1} + g_i)^2 \\ &\geq ((1 + \alpha)x_{i_1+1} - \alpha x_{i_1} - x_{i_1+2} + g_{i_1+1})^2 \\ &\quad + \sum_{i=i_1+2}^{i_2} (2x_i - x_{i-1} - x_{i+1} + g_i)^2, \end{aligned} \quad (27)$$

where we made use of the notation $\alpha := 1$ if $i_1 > 0$ and $\alpha := \sqrt{\mu}$ if $i_1 = 0$. We obtain a lower bound for the final sum of m squares (27) by fixing x_{i_1}, x_{i_2} , and $g_{i_1+1}, \dots, g_{i_2}$, then minimizing the quadratic form explicitly over the $m - 1$ variables $x_{i_1+1}, \dots, x_{i_2-1}$. We obtain

$$\begin{aligned} \|\nabla \bar{f}_{T,\mu,r}(x)\|^2 &\geq \inf_{v \in \mathbb{R}^{m-1}} \left\{ ((1 + \alpha)v_1 - \alpha x_{i_1} - v_2 + g_{i_1+1})^2 \right. \\ &\quad + \sum_{j=2}^{m-2} (2v_j - v_{j-1} - v_{j+1} + g_{i_1+j})^2 \\ &\quad \left. + (2v_{m-1} - v_{m-2} - x_{i_2} + g_{i_2-1})^2 + (2x_{i_2} - v_m - x_{i_2+1} + g_{i_2})^2 \right\} \\ &= \inf_{v \in \mathbb{R}^{m-1}} \|Av - b\|^2 = b^\top \left(I - A \left(A^\top A \right)^{-1} A^\top \right) b = (z^\top b)^2, \end{aligned}$$

where the matrix A and vector b have definitions

$$A = \begin{bmatrix} 1 + \alpha & -1 & & & \\ -1 & 2 & -1 & & \\ & \ddots & \ddots & \ddots & \\ & & -1 & 2 & -1 \\ & & & -1 & 2 \\ & & & & -1 \end{bmatrix} \in \mathbb{R}^{m \times (m-1)} \text{ and } b = \begin{bmatrix} \alpha x_{i_1} - g_{i_1+1} \\ -g_{i_1+2} \\ \vdots \\ -g_{i_2-2} \\ x_{i_2} - g_{i_2-1} \\ -2x_{i_2} + x_{i_2+1} - g_{i_2} \end{bmatrix} \in \mathbb{R}^m,$$

and $z \in \mathbb{R}^m$ is a unit-norm solution to $A^\top z = 0$. The vector $z \in \mathbb{R}^m$ with

$$z_j = \frac{j - 1 + \frac{1}{\alpha}}{\sqrt{\sum_{i=1}^m (i - 1 + \frac{1}{\alpha})^2}}$$

is such a solution. Thus

$$\begin{aligned} & \|\nabla \tilde{f}_{T,\mu,r}(x)\|^2 \\ & \geq \frac{\left(x_{i_1} - \sum_{j=1}^m (j - 1 + \frac{1}{\alpha}) \cdot g_{i_1+j} + (m - 2 + \frac{1}{\alpha}) \cdot x_{i_2} + (m - 1 + \frac{1}{\alpha}) (-2x_{i_2} + x_{i_2+1})\right)^2}{\sum_{i=1}^m (i - 1 + \frac{1}{\alpha})^2} \\ & = \frac{1}{\sum_{i=1}^m (i - 1 + \frac{1}{\alpha})^2} \left(x_{i_1} - x_{i_2} + \left(m - 1 + \frac{1}{\alpha}\right) (x_{i_2+1} - x_{i_2}) \right. \\ & \quad \left. - \sum_{j=1}^m \left(j - 1 + \frac{1}{\alpha}\right) \cdot g_{i_1+j} \right)^2. \end{aligned} \tag{28}$$

We now bring to bear the properties of the transition region Lemma 7 supplies. By Lemma 7.i,

$$x_{i_1} - x_{i_2} + (m - 1 + \alpha^{-1}) (x_{i_2+1} - x_{i_2}) \geq 0.9 - 0.3 = \frac{3}{5}, \tag{29}$$

and by Lemma 7.ii, using $1 \leq \alpha^{-1} \leq 1/\sqrt{\mu}$,

$$\begin{aligned} & - \sum_{j=1}^m (j - 1 + \alpha^{-1}) g_{i_1+j} \geq \mu \sum_{j=1}^{(m-\alpha^{-1})/2} (j - 1 + \alpha^{-1}) \\ & \geq \frac{\mu}{8} \left[m^2 - \frac{1}{\alpha^2} \right]_+ \\ & \geq \frac{1}{8} \left[\mu m^2 - 1 \right]_+. \end{aligned} \tag{30}$$

Substituting $\sum_{i=1}^m (i - 1 + \alpha^{-1})^2 \leq \frac{1}{2}m(m + 1/\sqrt{\mu})(m + 2/\sqrt{\mu})$ and the bounds (29) and (30) into the gradient lower bound (28), we have that

$$\begin{aligned} \|\nabla \bar{f}_{T,\mu,r}(x)\| &\geq \mu^{3/4} \cdot \zeta(m\sqrt{\mu}) \text{ where } \zeta(t) \\ &:= \sqrt{\frac{2}{t(t+1)(t+2)}} \left(\frac{3}{5} + \frac{1}{8} \left[t^2 - 1 \right]_+ \right). \end{aligned}$$

A quick computation reveals that $\inf_{t>0} \zeta(t) \approx 0.28 > 1/4$, which gives the result. \square

Proof of lemma 7 We have by definition that $x_{i_1} > 0.9$ and $x_{i_2} \leq x_{i'_2} < 0.1$. To see that

$$-x_{i_2} + (m - 1 + \alpha^{-1})(x_{i_2+1} - x_{i_2}) \geq -0.3$$

holds, consider the two cases that $x_{i_2} \leq -0.1$ or $x_{i_2} > -0.1$. In the first case that $x_{i_2} \leq -0.1$, by definition $x_{i_2+1} \geq x_{i_2}$ so $-x_{i_2} + (m - 1 + \alpha^{-1})(x_{i_2+1} - x_{i_2}) > 0.1 > -0.3$. The second case that $x_{i_2} > -0.1$ is a bit more subtle. By definition of the sequence $x_{i_2}, \dots, x_{i'_2}$, we have

$$\begin{aligned} -0.1 < x_{i_2} < x_{i_2-1} - \frac{0.2}{m' - 1 + \frac{1}{\alpha}} < \dots \leq x_{i'_2} - \frac{0.2}{m' - 1 + \frac{1}{\alpha}}(i_2 - i'_2) \\ &< 0.1 - 0.2 \frac{m - m'}{m' - 1 + \frac{1}{\alpha}}. \end{aligned} \tag{31}$$

Combining this bound on x_{i_2} and the inequality $x_{i_2+1} \geq x_{i_2} - \frac{0.2}{m' - 1 + 1/\alpha}$ due to the construction of i_2 , we obtain

$$\begin{aligned} -x_{i_2} + (m - 1 + \alpha^{-1})(x_{i_2+1} - x_{i_2}) &> -0.1 + 0.2 \frac{m - m'}{m' - 1 + \frac{1}{\alpha}} \\ -0.2 \frac{m - 1 + \frac{1}{\alpha}}{m' - 1 + \frac{1}{\alpha}} &= -0.3. \end{aligned}$$

We note for the proof of property ii that the chain of inequalities (31) is possible only for $m \leq 2m' - 1 + 1/\alpha$, which implies there are at most $m' - 1 + 1/\alpha$ indices $i \in \mathcal{I}_{\text{trans}}$ such that $|x_i| < 0.1$.

The first part of property ii follows from Lemma 2.ii, since $x_i \leq 0.9 \leq 1$ for every $i \in \mathcal{I}_{\text{trans}}$. To see that the second part of the property holds, let N be the number of indices in $i \in \mathcal{I}_{\text{trans}}$ for which $\Upsilon'_r(x_i) < -1$. By Lemma 2.iv and the fact that $0.1 \leq x_i \leq 0.9$ for every $i \in \{i_1 + 1, \dots, i'_2 - 1\}$, $N \geq m' - 1$. Moreover, since there can be at most $m' - 1 + 1/\alpha$ indices $i \in \mathcal{I}_{\text{trans}}$ for which $|x_i| < 0.1$, $N \geq m - (m' - 1 + 1/\alpha)$. Averaging the two lower bounds gives $N \geq (m - 1/\alpha)/2$. \square

B.3 Proof of Theorem 3

Theorem 3 *There exists a numerical constant $c < \infty$ such that the following lower bound holds. Let $p \geq 2$, $p \in \mathbb{N}$, and let D, L_1, L_2, \dots, L_p , and ϵ be positive. Then*

$$\mathcal{T}_\epsilon(\mathcal{A}_{\text{det}}^{(1)} \cup \mathcal{A}_{\text{zr}}^{(1)}, \mathcal{F}_{1:p}^{\text{dist}}(D, L_1, \dots, L_p)) \geq B_\epsilon \left(\min_{q \in [p]} \left\{ \frac{L_q}{2\tilde{\ell}_q} D^{q+1} \right\}, \frac{L_1}{2}, \frac{L_2}{2}, \dots, \frac{L_p}{2} \right),$$

where $\tilde{\ell}_q \leq \exp(cq \log q + c)$.

Proof The proof builds off of those of Theorems 2 and PI.3. We begin by recalling the following bump function construction

$$\bar{h}_T(x) := \Psi \left(1 - \frac{25}{2} \left\| x - \frac{4}{5} e^{(T)} \right\|^2 \right) \text{ where } \Psi(t) := e \cdot \exp \left(-\frac{1}{[2t - 1]_+^2} \right). \quad (32)$$

Adding a scaled version of $-\bar{h}_T$ to our hard instance construction allows us to “plant” a global minimum that is both close to the origin and essentially invisible to zero-respecting method. For convenience, we restate Lemma PI.10,

Lemma 8 *The function \bar{h}_T satisfies the following.*

- i. *For all $x \in \mathbb{R}^T$ we have $\bar{h}_T(x) \in [0, 1]$, and $\bar{h}_T(0.8e^{(T)}) = 1$.*
- ii. *On the set $\{x \in \mathbb{R}^d \mid x_T \leq \frac{3}{5}\} \cup \{x \mid \|x\| \geq 1\}$, we have $\bar{h}_T(x) = 0$.*
- iii. *For every $p \geq 1$, the p th order derivative of \bar{h}_T is $\tilde{\ell}_p$ -Lipschitz continuous, where $\tilde{\ell}_p \leq e^{cp \log p + c}$ for a numerical constant $c < \infty$.*

With this lemma in place, we follow the broad outline of the proof of Theorem 2, with modifications to make sure the norm of the minimizers of f is small. Indeed, letting $\lambda, \sigma > 0$, we define our scaled hard instance $f : \mathbb{R}^{T+2} \rightarrow \mathbb{R}$ by

$$f(x) = \lambda \sigma^2 \bar{f}_{T,\mu,r}(x_1/\sigma, \dots, x_{T+1}/\sigma) - \tilde{\lambda} \bar{h}_{T+2}(x/D), \quad (33)$$

that is, the hard instance we construct in Theorem 2 minus a scaled bump function (32). For every $p \in \mathbb{N}$, we set the parameters λ, σ, μ and r as in the proof of Theorem 2, so that we satisfy inequality (12) except we replace L_q with $L_q/2$ for every $q \in [p]$ (including in the definitions of λ, σ, μ). Thus, as in inequality (12), for each $q \in \mathbb{N}$ the function $f_0(x) := \lambda \sigma^2 \bar{f}_{T,\mu,r}(x/\sigma)$ has $L_q/2$ -Lipschitz q th order derivative and satisfies $\|\nabla f_0(x)\| > \epsilon$ for all $x \in \mathbb{R}^{T+1}$ with $x_T = x_{T+1} = 0$. By Lemma 8.iii, setting

$$\tilde{\lambda} = \min_{q \in [p]} \frac{1}{2\tilde{\ell}_q} L_q D^{q+1} \quad (34)$$

guarantees that the function $x \mapsto -\tilde{\lambda} \cdot \bar{h}_{T+2}(x/D)$ also has $L_q/2$ -Lipschitz q th order derivatives, so that overall, for each $q \in [p]$ the function f defined in Eq. (33) has L_q -Lipschitz q th order derivative.

We note that by Lemma 8.ii, $\bar{h}_{T+2}(x)$ is identically 0 at a neighborhood of any x with $x_{T+2} = 0$, which immediate implies that \bar{h}_{T+2} and f are zero-chains. Therefore for any $\mathbf{A} \in \mathcal{A}_{\text{zr}}^{(1)}$ producing iterates $x^{(1)} = 0, x^{(2)}, x^{(3)}, \dots$ when operating on f , we have $x_T^{(t)} = x_{T+1}^{(t)} = x_{T+2}^{(t)} = 0$ for any $t \leq T$. Thus, by our choices of λ, σ, μ and r , $\|\nabla f(x^{(t)})\| = \|\nabla f_0(x^{(t)})\| > \epsilon$ for every $t \leq T$, and so

$$\mathsf{T}_\epsilon(\mathcal{A}_{\text{zr}}^{(1)}, \mathcal{F}_{1:p}^{\text{dist}}(D, L_1, \dots, L_p)) \geq \inf_{\mathbf{A} \in \mathcal{A}_{\text{zr}}^{(1)}} \mathsf{T}_\epsilon(\mathbf{A}, f) \geq T + 1.$$

To establish that $f \in \mathcal{F}_{1:p}^{\text{dist}}(D, L_1, \dots, L_p)$, it remains to show that every global minimizer of f has norm at most D . Let x^* denote a global minimizer of f , and temporarily assume that

$$f(0.8D \cdot e^{(T+2)}) < 0, \quad (35)$$

Therefore, $f(x^*) < f(0.8D \cdot e^{(T+2)}) < 0$ and $\bar{h}_{T+2}(x^*/D) \neq 0$, as otherwise we have the contradiction $f(x^*) = \lambda\sigma^2 \bar{f}_{T,\mu,r}(x^*/\sigma) \geq 0$. By the definition (32), $\bar{h}_{T+2}(x^*/D) \neq 0$ implies that $1 - \frac{25}{2} \|x^*/D - 0.8e^{(T+2)}\|^2 \geq 0.5$, and therefore $\|x^*\| \leq D$. To verify the assumed inequality (35), we use Lemma 8.i to obtain

$$\begin{aligned} f(0.8D \cdot e^{(T+2)}) &= \lambda\sigma^2 \cdot \bar{f}_{T,\mu,r}(0) - \tilde{\lambda} \cdot \bar{h}_{T+2}(0.8 \cdot e^{(T+2)}) \\ &= \frac{\lambda\sqrt{\mu}\sigma^2}{2} + 10\lambda\sigma^2\mu T - \tilde{\lambda}. \end{aligned}$$

Therefore, if we set

$$T = \left\lfloor \frac{\tilde{\lambda} - \lambda\sqrt{\mu}\sigma^2/2}{10\lambda\mu\sigma^2} \right\rfloor \quad (36)$$

then inequality (35) holds and $\|x^*\| \leq D$, and so $f \in \mathcal{F}_{1:p}^{\text{dist}}(D, L_1, \dots, L_p)$. Comparing the setting (36) of T above to the setting (13) of T in the proof of Theorem 2, we see they are identical except that we replace the term Δ in (13) with $\tilde{\lambda} := \min_{q \in [p]} (2\tilde{\ell}_q)^{-1} L_q D^{q+1}$. Thus, mimicking the proof of Theorem 2 after the step (13), *mutatis mutandis*, yields the result. \square

B.4 Proof of Lemma 5

Lemma 5 *Let $T \in \mathbb{N}$, $0 < \alpha \leq 1$, $\mu \in [T^{-2}, 1]$ and $\tilde{f}_{T,\alpha,\mu}$ be defined as in (19), with Λ and $\tilde{\Upsilon}$ satisfying*

$$\Lambda'(0) = \tilde{\Upsilon}'(0) = 0 \text{ and } \Lambda' \text{ is 1-Lipschitz continuous and } \max_{z \in [0,1]} |\tilde{\Upsilon}'(z)| \leq G,$$

for $G > 0$ independent of T, α and μ . Then there exists $x \in \mathbb{R}^{T+1}$ such that $x_T = x_{T+1} = 0$ and

$$\|\nabla \tilde{f}_{T,\alpha,\mu}(x)\| < C\mu^{3/4},$$

where $C \leq 27 + \sqrt{3}G$.

Proof We construct x as follows. We let $x_1 = 1$, and for $n > 1$ let (with $x_0 := 1$),

$$x_n = x_{n-1} - (x_{n-2} - x_{n-1}) - \delta_{n-1} = 1 - \sum_{i=1}^{n-1} \sum_{j=1}^i \delta_i,$$

where we take

$$\delta_n = \frac{1}{m(m+1)} \begin{cases} 1 & n \leq m \\ 0 & n = m+1 \text{ or } n > 2m+1 \\ -1 & m+1 < n \leq 2m+1 \end{cases}$$

for some $m \in \mathbb{N}$ which we will later determine. The elements of $\nabla \tilde{f}_{T,\alpha,\mu}$ are given by

$$\nabla_n \tilde{f}_{T,\alpha,\mu}(x) = \Lambda'(x_n - x_{n-1}) - \Lambda'(x_{n+1} - x_n) + \mu \tilde{\Upsilon}'(x_n),$$

where for $n = 1$ we used $x_1 = 1$ and $\Lambda'(0) = 0$ to write $\alpha \cdot \Lambda'(x_1 - 1) = 0 = \Lambda'(x_1 - 1)$. Since Λ' is 1-Lipschitz, we have

$$|\Lambda'(x_n - x_{n-1}) - \Lambda'(x_{n+1} - x_n)| \leq |(x_n - x_{n-1}) - (x_{n+1} - x_n)| = |\delta_n|.$$

Moreover, one can readily verify that $\tilde{x}_n \in [0, 1]$ for every n and that $x_n = 0$ for every $n > 2m+1$. Therefore, using $\tilde{\Upsilon}'(0) = 0$ and $\max_{z \in [0, 1]} |\tilde{\Upsilon}'(z)| \leq G$ we have that $|\tilde{\Upsilon}'(x_n)| \leq G \cdot 1_{(n \leq 2m+1)}$, which gives the overall bound

$$|\nabla_n \tilde{f}_{T,\alpha,\mu}(x)| \leq |\delta_n| + \mu |\tilde{\Upsilon}(x_n)| \leq \left(\frac{1}{m^2} + G\mu \right) 1_{(n \leq 2m+1)},$$

and thus,

$$\|\nabla \tilde{f}_{T,\alpha,\mu}(x)\| \leq \sqrt{2m+1} \left(m^{-2} + G\mu \right) \leq \sqrt{3} \left(m^{-3/2} + \sqrt{m}G\mu \right).$$

Taking $m = \lceil \frac{1}{3\sqrt{\mu}} \rceil$, we have

$$\|\nabla \tilde{f}_{T,\alpha,\mu}(x)\| \leq \sqrt{3} \left(\left\lceil \frac{1}{3\sqrt{\mu}} \right\rceil^{-3/2} + G \left\lceil \frac{1}{3\sqrt{\mu}} \right\rceil^{1/2} \mu \right) \leq (27 + \sqrt{3}G) \mu^{3/4},$$

where we have used $\lceil 1/(3\sqrt{\mu}) \rceil \leq 1/\sqrt{\mu}$ since $\mu \leq 1$. Thus, $\|\nabla \tilde{f}_{T,\alpha,\mu}(x)\| \leq C\mu^{3/4}$ holds for $C = 27 + \sqrt{3}G$. For $T \geq 8$, since $\mu \geq T^{-2}$, we have $2m+1 \leq 2\lceil T/3 \rceil + 1 < T$ and therefore $x_T = x_{T+1} = 0$ holds as required (since $x_n = 0$ for every $n > 2m+1$). In the edge case $T \leq 8$ we have $\mu \geq T^{-2} \geq 1/64$ and therefore $x = 0$ yields $\|\nabla \tilde{f}_{T,\alpha,\mu}(x)\| = \alpha \leq 1 \leq 27 \cdot (1/64)^{3/4} \leq C\mu^{3/4}$. \square

References

1. Agarwal, N., Allen-Zhu, Z., Bullins, B., Hazan, E., Ma, T.: Finding approximate local minima faster than gradient descent. In: Proceedings of the Forty-Ninth Annual ACM Symposium on the Theory of Computing (2017)
2. Allen-Zhu, Z.: Natasha 2: Faster non-convex optimization than SGD (2017). [arXiv:1708.08694](https://arxiv.org/abs/1708.08694) [math.OC]
3. Allen-Zhu, Z., Hazan, E.: Variance reduction for faster non-convex optimization. In: Proceedings of the 33rd International Conference on Machine Learning (2016)
4. Arjevani, Y., Shamir, O., Shiff, R.: Oracle complexity of second-order methods for smooth convex optimization (2017). [arXiv:1705.07260](https://arxiv.org/abs/1705.07260) [math.OC]
5. Birgin, E.G., Gardenghi, J.L., Martínez, J.M., Santos, S.A., Toint, P.L.: Worst-case evaluation complexity for unconstrained nonlinear optimization using high-order regularized models. *Math. Program.* **163**(1–2), 359–368 (2017)
6. Boyd, S., Vandenberghe, L.: *Convex Optimization*. Cambridge University Press, Cambridge (2004)
7. Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J.: Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.* **3**(1), 1–122 (2011)
8. Carmon, Y., Duchi, J.C., Hinder, O., Sidford, A.: Convex until proven guilty: dimension-free acceleration of gradient descent on non-convex functions. In: Proceedings of the 34th International Conference on Machine Learning (2017)
9. Carmon, Y., Duchi, J.C., Hinder, O., Sidford, A.: Accelerated methods for non-convex optimization. *SIAM J. Optim.* **28**(2), 1751–1772 (2018)
10. Carmon, Y., Duchi, J.C., Hinder, O., Sidford, A.: Lower bounds for finding stationary points I. *Math. Program.* (to Appear) (2019). <https://doi.org/10.1007/s10107-019-01406-y>
11. Cartis, C., Gould, N.I., Toint, P.L.: On the complexity of steepest descent, Newton's and regularized Newton's methods for nonconvex unconstrained optimization problems. *SIAM J. Optim.* **20**(6), 2833–2852 (2010)
12. Cartis, C., Gould, N.I., Toint, P.L.: Complexity bounds for second-order optimality in unconstrained optimization. *J. Complex.* **28**(1), 93–108 (2012)
13. Cartis, C., Gould, N.I.M., Toint, P.L.: How much patience do you have? A worst-case perspective on smooth nonconvex optimization. *Optima* **88** (2012)
14. Cartis, C., Gould, N.I.M., Toint, P.L.: Worst-case evaluation complexity and optimality of second-order methods for nonconvex smooth optimization (2017). [arXiv:1709.07180](https://arxiv.org/abs/1709.07180) [math.OC]
15. Chowla, S., Herstein, I.N., Moore, W.K.: On recursions connected with symmetric groups I. *Can. J. Math.* **3**, 328–334 (1951)
16. Chung, F.R.K.: *Spectral Graph Theory*. AMS (1998)
17. Hinder, O.: Cutting plane methods can be extended into nonconvex optimization. In: Proceedings of the Thirty First Annual Conference on Computational Learning Theory (2018)
18. Jarre, F.: On Nesterov's smooth Chebyshev–Rosenbrock function. *Optim. Methods Softw.* **28**(3), 478–500 (2013)
19. Jin, C., Ge, R., Netrapalli, P., Kakade, S.M., Jordan, M. I.: How to escape saddle points efficiently. In: Proceedings of the 34th International Conference on Machine Learning (2017)
20. Lei, L., Ju, C., Chen, J., Jordan, M. I.: Nonconvex finite-sum optimization via SCSG methods. In: *Advances in Neural Information Processing Systems*, vol. 31 (2017)
21. Monteiro, R.D., Svaiter, B.F.: An accelerated hybrid proximal extragradient method for convex optimization and its implications to second-order methods. *SIAM J. Optim.* **23**(2), 1092–1125 (2013)
22. Nemirovski, A., Yudin, D.: *Problem Complexity and Method Efficiency in Optimization*. Wiley, Hoboken (1983)
23. Nesterov, Y.: *Introductory Lectures on Convex Optimization*. Kluwer Academic Publishers, Dordrecht (2004)
24. Nesterov, Y.: How to make the gradients small. *Optima* **88** (2012)
25. Nesterov, Y., Polyak, B.: Cubic regularization of Newton method and its global performance. *Math. Program. Ser. A* **108**, 177–205 (2006)
26. Reddi, S.J., Hefny, A., Sra, S., Poczos, B., Smola, A.: Stochastic variance reduction for nonconvex optimization. In: Proceedings of the 33rd International Conference on Machine Learning (2016)
27. Simchowitz, M.: On the randomized complexity of minimizing a convex quadratic function (2018). [arXiv:1807.09386](https://arxiv.org/abs/1807.09386) [cs.LG]

28. Simchowitz, M., Aloui, A.E., Recht, B.: Tight query complexity lower bounds for PCA via finite sample deformed Wigner law. In: Proceedings of the Fiftieth Annual ACM Symposium on the Theory of Computing (2018)
29. Vavasis, S.A.: Black-box complexity of local minimization. *SIAM J. Optim.* **3**(1), 60–80 (1993)
30. Woodworth, B.E., Srebro, N.: Tight complexity bounds for optimizing composite objectives. *Adv. Neural Inf. Process. Syst.* **30**, 3639–3647 (2016)
31. Zhang, X., Ling, C., Qi, L.: The best rank-1 approximation of a symmetric tensor and related spherical optimization problems. *SIAM J. Matrix Anal. Appl.* **33**(3), 806–821 (2012)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.