# Thinking Inside the Ball: Near-Optimal Minimization of the Maximal Loss 

## Yair Carmon

Tel Aviv University

Arun Jambulapati<br>Stanford University<br>Yujia Jin<br>Stanford University

Aaron Sidford
Stanford University

Editors: Mikhail Belkin and Samory Kpotufe


#### Abstract

We characterize the complexity of minimizing $\max _{i \in[N]} f_{i}(x)$ for convex, Lipschitz functions $f_{1}, \ldots, f_{N}$. For non-smooth functions, existing methods require $O\left(N \epsilon^{-2}\right)$ queries to a first-order oracle to compute an $\epsilon$-suboptimal point and $\widetilde{O}\left(N \epsilon^{-1}\right)$ queries if the $f_{i}$ are $O(1 / \epsilon)$-smooth. We develop methods with improved complexity bounds of $\widetilde{O}\left(N \epsilon^{-2 / 3}+\epsilon^{-8 / 3}\right)$ in the non-smooth case and $\widetilde{O}\left(N \epsilon^{-2 / 3}+\sqrt{N} \epsilon^{-1}\right)$ in the $O(1 / \epsilon)$-smooth case. Our methods consist of a recently proposed ball optimization oracle acceleration algorithm (which we refine) and a careful implementation of said oracle for the softmax function. We also prove an oracle complexity lower bound scaling as $\Omega\left(N \epsilon^{-2 / 3}\right)$, showing that our dependence on $N$ is optimal up to polylogarithmic factors.


Keywords: Convex optimization, Min-max problems, Monteiro-Svaiter acceleration, Ball optimization oracle, Stochastic first-order methods.

## 1. Introduction

Consider the problem of approximately minimizing the maximum of $N$ convex functions: given $f_{1}, \ldots, f_{N}$ such that for every $i \in[N]$ the function $f_{i}: \mathbb{R}^{d} \rightarrow \mathbb{R}$ is convex, Lipschitz and possibly smooth, and a target accuracy $\epsilon$,

$$
\begin{equation*}
\text { find a point } x \text { such that } F_{\max }(x)-\inf _{x_{\star} \in \mathbb{R}^{d}} F_{\max }\left(x_{\star}\right) \leq \epsilon \text { where } F_{\max }(x):=\max _{i \in[N]} f_{i}(x) \tag{1}
\end{equation*}
$$

Problems of this form play significant roles in optimization and machine learning. The maximum of $N$ functions is a canonical example of structured non-smoothness and several works develop methods for exploiting it $[31,30,36,9,12]$. The special case where the $f_{i}$ 's are linear functions is particularly important for machine learning, since it is equivalent to hard-margin SVM training (with $f_{i}$ representing the negative margin on the $i$ th example) [38, 13, 21]. Going beyond the linear case, Shalev-Shwartz and Wexler [36] argue that minimizing the maximum classification loss can have advantageous effects on training speed and generalization in the presence of rare informative examples. Moreover, minimizing the worst-case objective is the basic paradigm of robust optimization [4, 27]. In particular, since $F_{\max }(x)=\max _{p \in \Delta^{N}} \sum_{i \in[N]} p_{i} f_{i}(x)$ the problem corresponds to an extreme case of distributionally robust optimization [5] with an uncertainty set that encompasses the entire probability simplex $\Delta^{N}$.

The goal of this paper is to characterize the complexity of this fundamental problem. We are particularly interested in the regime where the number of data points $N$ and the problem dimension $d$ are large compared to the desired level of accuracy $1 / \epsilon$, as is common in modern machine learning. Consequently, we focus on dimension-independent first-order methods (i.e., methods which only rely on access to $f_{i}(x)$ and a (sub)gradient $\nabla f_{i}(x)$ as opposed to higher-order derivatives), and report complexity in terms of the number of function/gradient evaluations required to solve the problem.

### 1.1. Related work

To put our new complexity bounds in context, we first review the prior art in solving the problem (1) with first-order methods. For simplicity of presentation, throughout the introduction we assume each $f_{i}$ is 1-Lipschitz and that $F_{\max }$ has a global minimizer $x_{\star}$ with (Euclidean) norm at most 1.

The simplest approach to solving the problem (1) is the subgradient method [33]. This method finds an $\epsilon$-accurate solution in $O\left(\epsilon^{-2}\right)$ iterations, with each step computing a subgradient of $F_{\max }$, which in turn requires evaluation of all $N$ function values and a single gradient. Consequently, the complexity of this method is $O\left(N \epsilon^{-2}\right)$. We are unaware of prior work obtaining improved complexity without further assumptions. ${ }^{1}$

However, even a weak bound on smoothness helps: if each $f_{i}$ has $O(1 / \epsilon)$-Lipschitz gradient, then it is possible to minimize $F_{\max }$ to accuracy $\epsilon$ with complexity $\widetilde{O}\left(N \epsilon^{-1}\right)$ [31]. ${ }^{2}$ This result relies on the so-called "softmax" approximation of the maximum,

$$
\begin{equation*}
F_{\text {smax }, \epsilon}(x):=\epsilon^{\prime} \log \left(\sum_{i \in[N]} e^{f_{i}(x) / \epsilon^{\prime}}\right), \quad \text { where } \epsilon^{\prime}=\frac{\epsilon}{2 \log N} . \tag{2}
\end{equation*}
$$

It is straightforward to show that $\left|F_{\text {smax }, \epsilon}(x)-F_{\max }(x)\right| \leq \frac{\epsilon}{2}$ for all $x \in \mathbb{R}^{d}$, and that $\nabla F_{\text {smax }, \epsilon}$ is $\widetilde{O}(1 / \epsilon)$-Lipschitz if $\nabla f_{i}$ is $O(1 / \epsilon)$-Lipschitz for every $i$. Therefore, Nesterov's accelerated gradient descent [31] finds a minimizer of $F_{\text {smax }, \epsilon}$ to accuracy $\frac{\epsilon}{2}$ in $\widetilde{O}(\sqrt{1 / \epsilon} / \sqrt{\epsilon})$ iterations, with each iteration requiring $N$ evaluations of $f_{i}$ and $\nabla f_{i}$ to compute $\nabla F_{\text {smax }, \epsilon}$, yielding the claimed bound. The assumption that $\nabla f_{i}$ is $O(1 / \epsilon)$-Lipschitz is fairly weak; see Appendix A. 1 for additional discussion.

Given more smoothness, further improvement is possible. Nesterov [33, Section 2.3.1] shows that it suffices to solve $O\left(\sqrt{L_{g} / \epsilon}\right)$ linearized subproblems of the form $\min _{x \in \mathbb{R}^{d}} \max _{i \in[N]}\left\{f_{i}\left(y_{t}\right)+\right.$ $\left.\left(\nabla f_{i}\left(y_{t}\right)\right)^{\top}\left(x-y_{t}\right)+\frac{L_{g}}{2}\left\|x-y_{t}\right\|^{2}\right\}$. This yields a query complexity upper bound of $O\left(N \sqrt{L_{g} / \epsilon}\right)$, Though the complexity of solving each subproblem is not immediately clear, in Appendix A. 3 we explain how a first-order method [10] solves the subproblem to sufficient precision. Additional schemes for solving (1) in the special case of linear functions (i.e., $L_{g}=0$ ) are discussed in Appendix A. 2 .

A powerful technique for solving optimization problems with a large number $N$ of component functions is sampling components in order to compute cheap unbiased gradient estimates. However, both $F_{\max }$ and $F_{\mathrm{smax}, \epsilon}$ are not given as linear combinations of the $f_{i}$ 's. Consequently, it is not clear how to efficiently compute unbiased estimators for their gradients. Several works address this by

1. The center of gravity method $[24,35]$ yields a query complexity $O(N d \log (1 / \epsilon))$ which is an improvement only for sufficiently small problem dimension $d$.
2. Throughout the paper, the $\widetilde{O}(\cdot)$ and $\widetilde{\Omega}(\cdot)$ hide polylogarithmic factors.
\(\left.\begin{array}{llll}\hline Smoothness \& Method \& Upper bound \& Lower bound <br>
\hline None\left(L_{g}=\infty\right) \& Subgradient method \& N \epsilon^{-2} \& N \epsilon^{-2 / 3}+\epsilon^{-2} <br>

\hline Weak\left(L_{g} \approx 1 / \epsilon\right) \& Ours \& OGD on softmax \& N \epsilon^{-2 / 3}+\epsilon^{-8 / 3}\end{array}\right]\)|  | Ours |
| :--- | :--- |
| Strong $\left(L_{g} \ll 1 / \epsilon\right)$ | AGD on linearization* |

Table 1: The complexity of solving the problem (1) in terms of number of $(i, x)$ queries for computing and $f_{i}(x)$ and $\nabla f_{i}(x)$. The tables assume each $f_{i}$ is convex, 1 -Lipschitz and (optionally) has $L_{g}$-Lipschitz gradient, and that $F_{\text {max }}$ has a minimizer with norm at most 1 . The stated rates omit constant and (in the upper bounds) polylogarithmic factors. *For this algorithm only, the computational complexity is not simply $d$ times the query complexity; see Appendix A.3.
considering the saddle point problem

$$
\min _{x \in \mathbb{R}^{d}} \max _{p \in \Delta^{N}} F_{\mathrm{pd}}(x ; p):=\sum_{i \in[N]} p_{i} f_{i}(x)
$$

which is equivalent to minimizing to $F_{\text {max }}$. One can obtain unbiased estimators for $\nabla F_{\mathrm{pd}}(x ; p)$, and apply stochastic mirror descent to find its saddle-point [30, 36, 27]. However, all known estimators for $\nabla_{p} F_{\mathrm{pd}}$ have complexity-variance product $\Omega(N)$. Consequently, the best general guarantees known for such methods are $\widetilde{O}\left(N \epsilon^{-2}\right)$ iterations and total complexity. ${ }^{3}$ Shalev-Shwartz and Wexler [36] analyze a stochastic primal-dual method from an online learning perspective. They show that if the online method producing the primal updates admits a mistake bound (as is the case for learning halfspaces), then the complexity of the approach improves to $\widetilde{O}\left(N \epsilon^{-1}\right)$. We show that adopting a primal-only perspective and iteratively restricting $x$ to a small ball (i.e., "thinking inside the ball") allows us to make better use of the scalability of stochastic gradient methods.

### 1.2. Our contributions

To motivate our developments, note that the general complexity guarantees described above all scale linearly with the number of functions $N$. On the one hand, this is to be expected, as even evaluating the maximum of $N$ numbers requires querying all of them. On the other hand, a linear scaling in $N$ stands in sharp contrast to guarantees for minimizing the average of $N$ functions, which are typically sublinear in $N$. Since good scaling with dataset size is crucial in machine learning, we wish to precisely characterize the number of dataset passes (that is, the coefficient of $N$ ) in the complexity of minimizing $F_{\text {max }}$.

Towards that end, we prove an oracle complexity lower bound. The bound shows that any algorithm that operates by repeatedly querying $i, x$ and observing $f_{i}(x), \nabla f_{i}(x)$, must make $\Omega\left(N \epsilon^{-2 / 3}\right)$ queries in order to solve problem (1) for some convex, 1 -Lipschitz problem instance $f_{1}, \ldots, f_{N}$ with domain in the unit ball. The same bound continues to hold even when constraining the $f_{i}$ to have $O(1 / \epsilon)$-Lipschitz gradient, and when using high-order derivative oracles. This result further sharpens the contrast to average risk minimization, as it implies $\Omega\left(\epsilon^{-2 / 3}\right)$ dataset passes are required in

[^0]the worst case. However, it also suggests the potential for significant improvement over existing algorithms and their complexity bounds.

We realize this potential with new algorithms whose leading complexity term in $N$ matches our lower bound up to polylogarithmic factors. In the non-smooth case, our approach solves (1) with complexity $\widetilde{O}\left(N \epsilon^{-2 / 3}+\epsilon^{-8 / 3}\right)$, dominating prior guarantees for $N=\widetilde{\Omega}\left(\epsilon^{-2 / 3}\right)$. For $O(1 / \epsilon)-$ Lipschitz gradient functions, we obtain the stronger rate $\widetilde{O}\left(N \epsilon^{-2 / 3}+\sqrt{N} \epsilon^{-1}\right)$, which dominates prior guarantees for $N=\widetilde{\Omega}(1)$. At the core of these algorithms is a technique for accelerated optimization given a ball optimization oracle [12]; we make several improvements to this technique, which may be of independent interest.

Table 1 summarizes our results and their comparison to prior art. In addition to the results described above, the table also contains lower bounds on sublinear terms in $N$ (that follow from standard arguments), as well as a lower bound for the smooth regime where $L_{g}=o(1 / \epsilon)$. In this regime there exists a gap between the linear terms in the upper and lower bounds.

### 1.3. Overview of techniques

Our algorithms rely on a new technique introduced by Carmon et al. [12] for acceleration with a ball optimization oracle (BOO). For any $r>0$ and $F: \mathbb{R}^{d} \rightarrow \mathbb{R}$, a BOO of radius $r$ takes in a query point $\bar{x} \in \mathbb{R}^{d}$ and returns an (approximate) minimizer of $F$ in a ball of radius $r$ around $\bar{x}$. The technique, which is a variant of Monteiro-Svaiter acceleration [26, 17, 7, 9], minimizes $F$ to $\epsilon$ accuracy using $\widetilde{O}\left((1 / r)^{2 / 3}\right)$ oracle calls (with poly $(\log (1 / \epsilon))$ factors hidden). Carmon et al. [11] apply their technique to the special case of (1) with linear losses (see also Appendix A.2), showing that the log-sum-exp function is quasi-self-concordant and implementing a BOO of radius $r=\widetilde{\Theta}(\epsilon)$ using $\widetilde{O}(1)$ linear system solves. However, this approach does not extend to general $f_{i}$ because quasi-self-concordance no longer holds for $F_{\text {smax }, \epsilon}$, which might not even be differentiable.

The main technical insight of our paper is that it is possible to efficiently implement a BOO of radius $r_{\epsilon}=\widetilde{\Theta}(\epsilon)$ for $F_{\text {smax }, \epsilon}$ using stochastic first-order methods. More precisely, for any $\bar{x} \in \mathbb{R}^{d}$ we can minimize $F_{\text {smax }, \epsilon}$ in a ball of radius $r_{\epsilon}$ around $\bar{x}$ to any poly $(\epsilon)$ accuracy with precisely $N$ function evaluations and poly ( $1 / \epsilon$ ) (sub-)gradient evaluations. Using BOO acceleration, this immediately implies an $\widetilde{O}\left(N \epsilon^{-2 / 3}+\right.$ poly $\left.(1 / \epsilon)\right)$ complexity bound exhibiting optimal $N$ dependence.

To implement the BOO for $F_{\text {smax }, \epsilon}$, we consider instead the "exponentiated softmax" function
$\Gamma_{\epsilon}(x)=\epsilon^{\prime} \cdot \exp \left(\frac{F_{\text {smax }, \epsilon}(x)-F_{\text {smax }, \epsilon}(\bar{x})}{\epsilon^{\prime}}\right)=\sum_{i \in[N]} p_{i} \epsilon^{\prime} \cdot e^{\frac{f_{i}(x)-f_{i}(\bar{x})}{\epsilon^{\prime}}}$ where $p_{i}=\frac{e^{f_{i}(\bar{x}) / \epsilon^{\prime}}}{\sum_{j \in[N]} e^{f_{i}(\bar{x}) / \epsilon^{\prime}}}$,
and $\epsilon^{\prime}=\epsilon /(2 \log N)$ as in eq. (2). Note that $\Gamma_{\epsilon}$ is a monotonically increasing transformation of $F_{\text {smax }, \epsilon}$, and is therefore convex with the same minimizer as $F_{\text {smax }, \epsilon}$. Moreover, it is a (weighted) finite sum, and consequently amenable to stochastic gradient methods. It remains to verify that the functions $\xi_{i}(x)=\epsilon^{\prime} \cdot e^{\left(f_{i}(x)-f_{i}(\bar{x})\right) / \epsilon^{\prime}}$ are well-behaved, which might look difficult since exponentials are notoriously unstable. However, our choice of $r$ and Lipschitz continuity of $f_{i}$ implies that $e^{\left(f_{i}(x)-f_{i}(\bar{x})\right) / \epsilon}=\Theta(1)$ inside the ball, and consequently $\xi_{i}$ is indeed well-behaved, with Lipschitz constant $O(1)$. We thus minimize $\Gamma_{\epsilon}$ (and hence $F_{\text {smax }, \epsilon}$ ) with stochastic gradient descent [20], sampling $i$ from $p$. Moreover, if $\nabla f_{i}$ are Lipschitz, then $\nabla \xi_{i}$ are also Lipschitz, and we apply an accelerated variance reduction method [1] for better efficiency.

To complete the analysis of our methods it remains to determine how accurately we need to solve each ball subproblem. Unfortunately, the analysis of [12] makes fairly stringent accuracy
requirements, and also requires $\nabla F_{\text {smax }, \epsilon}$ to have a finite Lipschitz constant. To obtain tighter guarantees, we significantly rework the analysis in [12], modifying the algorithm to make it applicable without any differentiablility requirements. Our improved analysis takes into account the fact that the acceleration scheme only requires ball minimization with strong $\ell_{2}$ regularization, which further improves the oracle implementation complexity.

Our lower bound follows from a variation on the classical "chain constructions" in optimization lower bounds [ $29,39,18,14$ ], where in order to make a unit of progress on our constructed function, any algorithm must (with constant probability) make $\Omega(N)$ queries in order to discover a single new link in the chain. We build a chain of length $\Omega\left(\epsilon^{-2 / 3}\right)$ for which querying any $\epsilon$ minimizer of $F_{\max }$ requires discovering the entire chain, giving the $\Omega\left(N \epsilon^{-2 / 3}\right)$ complexity lower bound. To prove this result for arbitrary randomized algorithms, we randomize both the order of the functions and the rotation of the domain.

Paper outline. Section 2 provides some additional preliminaries and notation. Section 3 gives our improved derivation of the BOO acceleration method of [12], and Section 4 develops a BOO for $F_{\text {smax }, \epsilon}$, culminating in our upper complexity bounds for the problem (1), stated in Theorem 6. Section 5 gives our lower bounds with the main result stated in Theorem 10. Section 6 includes some comments on our results and potential future work.

## 2. Preliminaries

General notation. Throughout, $\|\cdot\|$ denotes the Euclidean norm. We write $\mathbb{B}_{r}(z)$ for the Euclidean ball of radius $r$ centered at $z$, and $\mathbb{B}_{r}^{d}(z)$ when emphasizing that the ball is $d$-dimensional. We use $L_{f}$ to denote a function Lipschitz constant and $L_{g}$ to denote a gradient Lipschitz constant; we say that $f$ is $L_{g}$-smooth if it has $L_{g}$-Lipschitz gradient. To disambiguate between sequence and coordinate indices, in Section 5 we denote the former with normal subscript and the latter with bracketed subscript, i.e., $x_{[i]}$ is the $i$ th coordinate of $x$ and $x_{k}$ is the $k$ th element in the sequence $x_{1}, x_{2}, \ldots$. We also write $v_{[\leq i]}$ to denote a copy of $v$ with coordinate $i+1, i+2, \ldots$ set to zero. We use $a \wedge b:=\min \{a, b\}$ to abbreviate binary minimization. We write the binary indicator of event $A$ as $\mathbb{I}\{A\}$.

Complexity model. We mainly measure complexity through the number individual function and gradient evaluations required to solve the problem (1). We write $\mathcal{T}_{f}$ for the cost of evaluating $f_{i}(x)$ for a single $i$ and $x$, and similarly write $\mathcal{T}_{g}$ for the cost of evaluating $\nabla f_{i}(x)$. Assuming $\mathcal{T}_{f}, \mathcal{T}_{g}=\Omega(d)$, our evaluation complexity upper bounds translate directly to runtime upper bounds.

Proximal operators. For any function $f$ and regularization parameter $\lambda \geq 0$, we define the standard proximal mapping $\operatorname{prox}_{\lambda}^{f}(\bar{x}):=\arg \min _{x \in \mathbb{R}^{d}}\left\{f(x)+\frac{\lambda}{2}\|x-\bar{x}\|^{2}\right\}$. We also define the ball constrained proximal mapping $\operatorname{bprox}_{\lambda, r}^{f}(\bar{x}):=\arg \min _{x \in \mathbb{B}_{r}(\bar{x})}\left\{f(x)+\frac{\lambda}{2}\|x-\bar{x}\|^{2}\right\}$. Finally, we define the notion of an approximate oracle for bprox ${ }_{\lambda}^{f}, r$, which plays a key role in our analysis.

Definition 1 (BROO) We say that a mapping $\mathcal{O}_{\lambda, \delta}(\cdot)$ is a Ball Regularized Optimization Oracle of radius $r$ ( $r$-BROO) for $f$, if for every query point $\bar{x}$, regularization parameter $\lambda$ and desired accuracy $\delta$, it return $\tilde{x}=\mathcal{O}_{\lambda, \delta}(\bar{x})$ satisfying

$$
\begin{equation*}
f(\tilde{x})+\frac{\lambda}{2}\|\tilde{x}-\bar{x}\|^{2} \leq \min _{x \in \mathbb{B}_{r}(\bar{x})}\left\{f(x)+\frac{\lambda}{2}\|x-\bar{x}\|^{2}\right\}+\frac{\lambda}{2} \delta^{2} . \tag{3}
\end{equation*}
$$

Note that when $f$ is convex, the strong convexity of $f(x)+\frac{\lambda}{2}\|x-\bar{x}\|^{2}$ and the approximation requirement (3) guarantee that $\left\|\mathcal{O}_{\lambda, \delta}(\bar{x})-\operatorname{bprox}_{\lambda, r}^{f}(\bar{x})\right\| \leq \delta$.

## 3. BROO acceleration

In this section, we describe a variant of the ball optimization acceleration scheme of Carmon et al. [12], given as Algorithm 1. Both methods follow the template of Monteiro-Svaiter acceleration [26], but our algorithm improves on [12] in two ways. First, it accesses the objective strictly through the ball oracle, while [12] also uses gradient computations. Second, our algorithm requires an oracle that solves regularized ball optimization problems, which are easier to implement. ${ }^{4}$

As a consequence of these differences, our accelerated algorithm's guarantee does not require any smoothness of the objective function. Moreover, our setup allows for far less accurate solutions to the ball optimization subproblems: Carmon et al. [12] require $\delta=O\left(\frac{\epsilon}{L_{g} R}\right)$ while we only require $\delta=O\left(\frac{\epsilon}{\lambda R}\right)$. While our requirement becomes stricter as the regularizer $\lambda$ grows, it also becomes easier to fulfill since the ball optimization problem becomes more strongly convex and hence easier to solve. Our relaxed accuracy requirement ultimately translates to an improved $\epsilon^{-1}$ dependence in the sublinear-in- $N$ term in our upper bound.

With the key innovations of Algorithm 1 explained, we now formally state its convergence guarantee; we defer the proof to Appendix B.

Theorem 2 Let $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ be convex and $L_{f}$-Lipschitz, and let $z \in \mathbb{R}^{d}$. For any domain bound $R>0$, ball radius $r \in(0, R]$, accuracy level $\epsilon>0$, and initial point $x_{0} \in \mathbb{R}^{d}$, Algorithm 1 returns a point $x \in \mathbb{R}^{d}$ satisfying $f(x)-\min _{z \in \mathbb{B}_{R}\left(x_{0}\right)} f(z) \leq \epsilon$ using at most

$$
T=O\left(\left(\frac{R}{r}\right)^{2 / 3} \log \left(\frac{\left[f\left(x_{0}\right)-\min _{z \in \mathbb{B}_{R}\left(x_{0}\right)} f(z)\right] R}{\epsilon r}\right) \log \left(\frac{L_{f} R^{2}}{\epsilon r}\right)\right)
$$

queries to an $r$-BROO. Moreover, the BROO query parameters $\left(\lambda_{(1)}, \delta_{(1)}\right), \ldots,\left(\lambda_{(T)}, \delta_{(T)}\right)$ satisfy

1. $\Omega\left(\frac{\epsilon}{r R}\right) \leq \lambda_{(i)} \leq O\left(\frac{L_{f}}{r}\right)$ and $\delta_{(i)} \geq \Omega\left(\frac{\epsilon}{\lambda_{(i)} R}\right)$ for all $i \in[T]$.
2. $\sum_{i \in[T]} \frac{1}{\sqrt{\lambda_{(i)}}} \leq O\left(\frac{R}{\sqrt{\epsilon}} \log \frac{L_{f} R^{2}}{\epsilon r}\right)$.

We remark that Theorem 2 requires a bound on the Lipschitz constant of $f$ solely to bound the complexity of the bisection procedure for finding $\left\{\lambda_{t}\right\}$.

## 4. BROO implementation

In this section, we develop efficient BROO implementations for $F_{\text {smax }, \epsilon}$, the softmax approximation of $F_{\max }(2)$. In Section 4.1 we develop our main analytical tool in the form of an "exponentiated softmax" function approximating $F_{\text {smax }, \epsilon}$ and facilitating efficient stochastic gradient estimation. We then minimize the exponentiated softmax with standard tools from stochastic convex optimization. In Section 4.2 we give a BROO implementation for the non-smooth case using restarted

[^1]```
    Algorithm 1: BROO acceleration
Input: Initial \(x_{0} \in \mathbb{R}^{d}\), Lipschitz and distance bounds \(L_{f}, R, r\), accuracy \(\epsilon\), BROO \(\mathcal{O}_{\lambda, \delta}(\cdot)\)
Output: \(x_{\mathrm{ret}}\) such that \(f\left(x_{\mathrm{ret}}\right)-\arg \min _{z \in \mathbb{B}_{R}\left(x_{0}\right)} f(z) \leq \epsilon\)
Let \(v_{0}=x_{0}, A_{0}=0\)
for \(t=0,1,2, \ldots\) do
    \(\lambda_{t+1}=\lambda\)-BISECTION \(\left(x_{t}, v_{t}, A_{t}, \lambda_{\max }=\frac{2 L_{f}}{r}, \lambda_{\min }=\frac{\epsilon}{6 r R}\right)\)
            \(\triangleright\) Finds \(\lambda_{t+1}\) such that \(x_{t+1} \approx \operatorname{prox}_{\lambda_{t+1}}^{f}\left(y_{t}\right)\) and either \(\left\|x_{t+1}-y_{t}\right\| \approx r\) or \(x_{t+1}\) is \(\epsilon\)-optimal
    \(a_{t+1}=\frac{1}{2 \lambda_{t+1}}\left(1+\sqrt{1+4 \lambda_{t+1} A_{t}}\right)\) and \(A_{t+1}=A_{t}+a_{t+1} \quad \triangleright A_{t+1}=a_{t+1}^{2} \lambda_{t+1}\)
    \(y_{t}=\frac{A_{t}}{A_{t+1}} x_{t}+\frac{a_{t+1}}{A_{t+1}} v_{t}\)
    \(x_{t+1}=\mathcal{O}_{\lambda_{t+1}, \delta_{t+1}}\left(y_{t}\right)\), where \(\delta_{t+1}=\frac{\epsilon}{12 \lambda_{t+1} R}\)
    \(v_{t+1}=\arg \min _{v \in \mathbb{B}_{R}\left(x_{0}\right)}\left\{a_{t+1} \lambda_{t+1}\left\langle y_{t}-x_{t+1}, v\right\rangle+\frac{1}{2}\left\|v-v_{t}\right\|^{2}\right\}\)
    if \(A_{t+1} \geq \frac{R^{2}}{\epsilon}, \lambda_{t+1} \leq \frac{\epsilon}{3 r R},\left\|x_{t+1}-v_{t+1}\right\|>2 R\), or \(A_{t+1}<\exp \left(\frac{r^{2 / 3}}{R^{2 / 3}}(t-1)\right) A_{1}\) then
        return \(x_{\text {ret }} \in \arg \min _{x \in\left\{x_{0}, x_{1}, \ldots, x_{t+1}\right\}} f(x)\)
function \(\lambda\)-Bisection \(\left(x, v, A, \lambda_{\text {max }}, \lambda_{\text {min }}\right)\)
    For all \(\lambda^{\prime}\), let \(y_{\lambda^{\prime}}:=\alpha_{2 A \lambda^{\prime}} \cdot x+\left(1-\alpha_{2 A \lambda^{\prime}}\right) \cdot v\), where \(\alpha_{\tau}:=\frac{\tau}{1+\tau+\sqrt{1+2 \tau}}\)
    Define \(\Delta(\lambda):=\left\|\mathcal{O}_{\lambda, \frac{r}{17}}\left(y_{\lambda}\right)-y_{\lambda}\right\| \quad \triangleright\) approximation of \(\widehat{\Delta}(\lambda):=\left\|\operatorname{bprox}_{\lambda, r}^{f}\left(y_{\lambda}\right)-y_{\lambda}\right\|\)
    Let \(\lambda=\lambda_{\max }\)
    while \(\lambda \geq \lambda_{\min }\) and \(\Delta(\lambda) \leq \frac{13 r}{16}\) do \(\lambda \leftarrow \lambda / 2 \quad \triangleright\) terminates in \(O\left(\log \frac{\lambda_{\max }}{\lambda_{\min }}\right)\) steps
    if \(\lambda \leq \lambda_{\min }\) then return \(2 \lambda \quad \triangleright\) happens only if \(\operatorname{bprox}_{2 \lambda, r}^{f}\left(y_{2 \lambda}\right)\) is \(O(\epsilon)\)-optimal for small \(\lambda_{\text {min }}\)
    Let \(\lambda_{u}=2 \lambda, \lambda_{\ell}=\lambda\) and \(\lambda_{m}=\sqrt{\lambda_{u} \lambda_{\ell}}\)
    if \(\Delta\left(\lambda_{\ell}\right) \leq \frac{15 r}{16}\) then return \(\lambda_{\ell} \quad \triangleright\) happens only if \(\Delta\left(\lambda_{\ell}\right) \in\left[\frac{13 r}{16}, \frac{15 r}{16}\right]\)
    while \(\Delta\left(\lambda_{m}\right) \notin\left[\frac{13 r}{16}, \frac{15 r}{16}\right]\) and \(\log _{2} \frac{\lambda_{u}}{\lambda_{\ell}} \geq \frac{r}{8\left(R+L_{f} / \lambda_{\ell}\right)}\) do
        if \(\Delta\left(\lambda_{m}\right)<\frac{13 r}{16}\) then \(\lambda_{u}=\lambda_{m}\) else \(\lambda_{\ell}=\lambda_{m}\)
        \(\lambda_{m}=\sqrt{\lambda_{u} \lambda_{\ell}}\)
    return \(\lambda_{m} \quad \triangleright\) the while loop terminates in \(O\left(\log \left(\frac{R}{r}+\frac{L_{f}}{\lambda_{\min r} r}\right)\right)\) steps
```

SGD [20]. In Section 4.3 we instead apply an accelerated variance reduction method (Katyusha [1]) that offers improved performance when the $f_{i}$ are even slightly smooth. Finally, in Section 4.4 we combine our BROO implementations with Algorithm 1 and its guarantees to obtain our main results: new convergence guarantees for minimizing $F_{\max }$. We defer proofs to Appendix C.

### 4.1. Exponentiating a softmax

Recall that $\epsilon^{\prime}=\epsilon /(2 \log N)$ and that (for nominal accuracy $\epsilon$ ) the softmax function $F_{\operatorname{smax}, \epsilon}(x)=$ $\epsilon^{\prime} \log \left(\sum_{i \in[N]} e^{f_{i}(x) / \epsilon^{\prime}}\right)$ approximates $F_{\max }$ to within $\epsilon / 2$ additive error. The key challenge in designing an efficient stochastic method for minimizing $F_{\mathrm{smax}, \epsilon}$ is a lack of cheap unbiased gradient
estimators. Specifically, we have $\nabla F_{\text {smax }, \epsilon}(x)=\sum_{i \in[N]} p_{i}(x) \nabla f_{i}(x)$, where

$$
\begin{equation*}
p_{i}(x)=\frac{e^{f_{i}(x) / \epsilon^{\prime}}}{\sum_{j \in[N]} e^{f_{j}(x) / \epsilon^{\prime}}} \tag{4}
\end{equation*}
$$

Given access to $p(x)$, we could easily obtain an unbiased estimator for $\nabla F_{\text {smax }, \epsilon}(x)$ by sampling $i \sim$ $p(x)$ and outputting $\nabla f_{i}(x)$. However, computing $p(x)$ itself requires evaluating all $N$ functions, making it basically as costly as computing $\nabla F_{\text {smax }, \epsilon}$ exactly.

This difficulty, however, is greatly relieved when we operate in a small ball of radius $r_{\epsilon}=\epsilon^{\prime} / L_{f}$ centered at some point $\bar{x}$. To see why, note that for every $i$ and every $x \in \mathbb{B}_{r_{\epsilon}}(\bar{x})$, Lipschitz continuity of $f_{i}$ implies $\left|f_{i}(x) / \epsilon^{\prime}-f_{i}(\bar{x}) / \epsilon^{\prime}\right| \leq L_{f} r_{\epsilon} / \epsilon^{\prime}=1$. Consequently, $p(\bar{x})$ is a multiplicative approximation for $p(x)$ throughout the ball, satisfying $e^{-2} p_{i}(\bar{x}) \leq p_{i}(x) \leq e^{2} p(\bar{x})$ for all $x \in$ $\mathbb{B}_{r_{\epsilon}}(\bar{x})$. Our high-level strategy is thus: perform a full data pass once to compute $p(\bar{x})$, and then rely on the stability of $p(x)$ within $\mathbb{B}_{r_{\epsilon}}(\bar{x})$ to efficiently estimate gradients by sampling from $p(\bar{x})$. However, simply sampling $i \sim p(\bar{x})$ and returning $\nabla f_{i}(x)$ is not enough, because it leads to a biased estimator of $\nabla F_{\text {smax }, \epsilon}(x)$. Instead, we define below a surrogate function "exponentiating the softmax" that closely approximates $F_{\text {smax }, \epsilon}$ and for which $e^{\left(f_{i}(x)-f_{i}(\bar{x})\right) / \epsilon^{\prime}} \nabla f_{i}(x)$ is an unbiased gradient estimator when $i \sim p(\bar{x}) .{ }^{5}$

To precisely define the surrogate "exponentiated softmax" function, we require some additional notation. Fixing a ball center $\bar{x}$ and regularization parameter $\lambda$, let

$$
f_{i}^{\lambda}(x):=f_{i}(x)+\frac{\lambda}{2}\|x-\bar{x}\|^{2} \text { and } F_{\mathrm{smax}, \epsilon}^{\lambda}(x):=F_{\mathrm{smax}, \epsilon}(x)+\frac{\lambda}{2}\|x-\bar{x}\|^{2}=\epsilon^{\prime} \log \left(\sum_{i \in[N]} e^{f_{i}^{\lambda}(x) / \epsilon^{\prime}}\right)
$$

be the regularized counterparts of $f_{i}$ and $F_{\text {smax }, \epsilon}$, respectively. Then, we define the exponentiated softmax as

$$
\begin{equation*}
\Gamma_{\epsilon, \lambda}(x)=\epsilon^{\prime} \cdot \exp \left(\frac{F_{\mathrm{smax}, \epsilon}^{\lambda}(x)-F_{\mathrm{smax}, \epsilon}^{\lambda}(\bar{x})}{\epsilon^{\prime}}\right)=\sum_{i \in[N]} p_{i}(\bar{x}) \gamma_{i}(x) \text { where } \gamma_{i}(x):=\epsilon^{\prime} e^{\frac{f_{i}^{\lambda}(x)-f_{i}^{\lambda}(\bar{x})}{\epsilon^{\prime}}} \tag{5}
\end{equation*}
$$

Clearly, $\Gamma_{\epsilon, \lambda}$ is a finite sum objective (weighted by $p(\bar{x})$ ), making stochastic first-order methods applicable. Moreover, as the following lemma shows, when the ball radius $r$ and $\lambda$ are not too large, $\Gamma_{\epsilon, \lambda}$ closely approximates $F_{\mathrm{smax}, \epsilon}^{\lambda}$ and is as regular as $F_{\mathrm{smax}, \epsilon}^{\lambda}$ up to a constant.
Lemma 3 Let $f_{1}, \cdots, f_{N}$ each be $L_{f}$-Lipschitz and $L_{g}$-smooth gradients. For any $c>0, r \leq$ $c \epsilon^{\prime} / L_{f}$, and $\lambda \leq c L_{f} / r$ let $C=\left(1+c+c^{2}\right) e^{c+c^{2} / 2}$. The exponentiated softmax $\Gamma_{\epsilon, \lambda}$ satisfies the following properties for any $\bar{x} \in \mathbb{R}^{d}$.

1. $F_{\mathrm{smax}, \epsilon}^{\lambda}(x)$ and $\Gamma_{\epsilon, \lambda}$ have the same minimizer $x_{\star}$ in $\mathbb{B}_{r}(\bar{x})$. Moreover, for every $x \in \mathbb{B}_{r}(\bar{x})$,

$$
F_{\mathrm{smax}, \epsilon}^{\lambda}(x)-F_{\mathrm{smax}, \epsilon}^{\lambda}\left(x_{\star}\right) \leq C\left(\Gamma_{\epsilon, \lambda}(x)-\Gamma_{\epsilon, \lambda}\left(x_{\star}\right)\right)
$$

2. Restricted to $\mathbb{B}_{r}(\bar{x})$, each function $\gamma_{i}$ defined in (5) is $C L_{f}$-Lipschitz, $C^{-1} \lambda$ strongly convex, and $C\left(L_{g}+\lambda+L_{f}^{2} / \epsilon^{\prime}\right)$-smooth.
The proof of Lemma 3 follows from a straightforward calculation, and we defer it to Appendix C.1.
[^2]
### 4.2. The non-smooth case: SGD implementation

To take advantage of the strong convexity of of $\Gamma_{\epsilon, \lambda}$ we use the restarted SGD variant of Hazan and Kale [20], which finds an $\varepsilon$-suboptimal point of a $G$-Lipschitz and $\mu$-strongly convex function with $\widetilde{O}\left(G^{2} /(\mu \varepsilon)\right)$ iterations (with high probability). To estimate the stochastic gradients, we sample $i \sim p(\bar{x})$ and output $\nabla \gamma_{i}(x)$; this takes $O\left(\mathcal{T}_{g}+\mathcal{T}_{f}\right)$ time per stochastic gradient, plus $O\left(N \mathcal{T}_{f}\right)$ preprocessing time to compute $p(\bar{x})$. We provide pseudocode for the algorithm in Appendix C.2, where we also prove the following complexity bound.

Corollary 4 Let $f_{1}, f_{2}, \cdots, f_{N}$ be $L_{f}$ Lipschitz, let $\sigma \in(0,1), \epsilon, \delta>0$ and $r_{\epsilon}=\epsilon /\left(2 \log N \cdot L_{f}\right)$. For any $\bar{x} \in \mathbb{R}^{d}$ and $\lambda \leq O\left(L_{f} / r_{\epsilon}\right)$, with probability at least $1-\sigma$, Algorithm 2 outputs a valid $r_{\epsilon}$-BROO response for $F_{\text {smax }, \epsilon}$ to query $\bar{x}$ with regularization $\lambda$ and accuracy $\delta$, and has cost

$$
\begin{equation*}
O\left(\mathcal{T}_{f} N+\left(\mathcal{T}_{g}+\mathcal{T}_{f}\right) \frac{L_{f}^{2}}{\lambda^{2} \delta^{2}} \log \left(\frac{\log \left(L_{f} / \lambda \delta\right)}{\sigma}\right)\right) \tag{6}
\end{equation*}
$$

### 4.3. The (slightly) smooth case: accelerated variance reduction implementation

If we further assume smoothness of $f_{1}, \ldots, f_{N}$, we can use stochastic variance reduction to obtain an improved runtime. With these methods, we estimate the gradient of $\Gamma_{\epsilon, \lambda}$ as $\nabla \Gamma_{\epsilon, \lambda}\left(x^{\prime}\right)+\nabla \gamma_{i}(x)-$ $\nabla \gamma_{i}\left(x^{\prime}\right)$, where $i \sim p(\bar{x})$ and $x^{\prime}$ is a reference point which we recompute $\widetilde{O}(1)$ times. Here, the $O\left(N \mathcal{T}_{f}\right)$ cost of computing $p(\bar{x})$ is essentially free compared to the $\operatorname{cost} \widetilde{O}\left(N \mathcal{T}_{g}\right)$ of computing the exact gradients of $\Gamma_{\epsilon, \lambda}$ at the reference point. We again take advantage of the regularization-induced $\lambda$-strong-convexity a variant of the Katyusha method of Allen-Zhu [1]. This results in the following complexity guarantee; see Appendix C. 3 for a proof.

Corollary 5 Let $f_{1}, \cdots, f_{N}$ be $L_{f}$-Lipschitz and $L_{g}$-smooth, let $\sigma \in(0,1), \epsilon, \delta>0, \epsilon^{\prime}=$ $\epsilon /(2 \log N)$ and $r_{\epsilon}=\epsilon^{\prime} / L_{f}$. For any $\bar{x} \in \mathbb{R}^{d}$ and $\lambda \leq O\left(L_{f} / r_{\epsilon}\right)$, with probability at least $1-\sigma$, Katyushal [1] outputs a valid $r_{\epsilon}$-BROO response to query $\bar{x}$ with regularization $\lambda$ and accuracy $\delta$, and has computational cost

$$
\begin{equation*}
O\left(\left(\mathcal{T}_{f}+\mathcal{T}_{g}\right)\left(N+\frac{\sqrt{N}\left(L_{f}+\sqrt{\epsilon^{\prime} L_{g}}\right)}{\sqrt{\lambda \epsilon^{\prime}}}\right) \log \left(\frac{L_{f} r_{\epsilon}}{\lambda \delta^{2} \sigma}\right)\right) . \tag{7}
\end{equation*}
$$

### 4.4. Main result

With our oracle implementations in hand, we are ready to state our main result.
Theorem 6 Let $f_{1}, f_{2}, \ldots, f_{N}$ be $L_{f}$-Lipschitz, let $x_{\star}$ be a minimizer of $F_{\max }(x)=\max _{i \in[N]} f_{i}(x)$ and assume $\left\|x_{0}-x_{\star}\right\| \leq R$ for a given initial point $x_{0}$ and some $R>0$. For any $\epsilon>0$, Algorithm 1 with the BROO implementation for $F_{\mathrm{smax}, \epsilon}$ in Algorithm 2 solves the problem (1) with probability at least $\frac{99}{100}$ and has computational cost

$$
\begin{equation*}
O\left(\left(\frac{L_{f} R \log N}{\epsilon}\right)^{2 / 3}\left(\mathcal{T}_{f} N+\left(\frac{L_{f} R}{\epsilon}\right)^{2} \cdot\left(\mathcal{T}_{f}+\mathcal{T}_{g}\right) \log K\right) \log ^{2} K\right) \tag{8}
\end{equation*}
$$

where $K:=L_{f} R \epsilon^{-1} \log N$. If moreover $f_{1}, f_{2}, \ldots, f_{N}$ are each $L_{g}$-smooth, then Algorithm 1 with a BROO implementation for $F_{\text {smax }, \epsilon}$ using Kayushal solves (1) with probability $\geq \frac{99}{100}$ and has cost

$$
O\left(\left(\mathcal{T}_{f}+\mathcal{T}_{g}\right)\left(\left(\frac{L_{f} R \log N}{\epsilon}\right)^{2 / 3} N+\left(\frac{L_{f} R \sqrt{\log N}}{\epsilon}+\sqrt{\frac{L_{g} R^{2}}{\epsilon}}\right) \sqrt{N}\right) \log ^{3} K\right) .
$$

The proof of Theorem 6, which we provide in Appendix C.4, follows straightforwardly from Theorem 2 and Corollaries 4 and 5 . When applying Corollary 4 with $\delta=\Omega\left(\frac{\epsilon}{\lambda R}\right)$ the dependence of the complexity on $\lambda$ cancels, and we get that each oracle call costs $\widetilde{O}\left(N \mathcal{T}_{f}+L_{f}^{2} R^{2} \epsilon^{-2}\left(\mathcal{T}_{f}+\mathcal{T}_{g}\right)\right)$. The complexity bound then follows from multiplying the per-call cost with the bound $\widetilde{O}\left(\left(R / r_{\epsilon}\right)^{-2 / 3}\right)$ that Theorem 2 provides on the total number of oracle calls. When applying Corollary 5 we obtain an oracle implementation cost of $\widetilde{O}\left(N\left(\mathcal{T}_{f}+\mathcal{T}_{g}\right)+\lambda^{-1 / 2} \sqrt{N} \sqrt{L_{f}^{2} \epsilon^{-1}+L_{g}}\left(\mathcal{T}_{f}+\mathcal{T}_{g}\right)\right)$. The complexity bound again follows by multiplying the per-call cost again with the total number of calls, except that to bound the contribution of $\sqrt{N}$ term we invoke the the guarantee $\sum_{i} \lambda_{(i)}^{-1 / 2} \leq \widetilde{O}\left(R \epsilon^{-1 / 2}\right)$ in Theorem 2 to a tighter bound.

## 5. Lower bounds

In this section, we prove oracle complexity lower bounds showing that the results of the previous section are order optimal for sufficiently large $N$ and $L_{g}$. While our algorithms are first-order methods, our lower bounds remain valid even for other algorithms that use high order derivatives, as is typical for our proof technique.

We begin by providing a formal definition of the oracle-based optimization model we consider (Section 5.1). In Section 5.2, we define an $N$-element variant for the zero-chain concept, and prove that it allows us to control the progress of any (possibly randomized) algorithm. Then, in Section 5.3 we construct a particular $N$-element zero-chain for which slow progress implies a large optimality gap. Finally, Section 5.4 ties these results together, giving our lower bound and providing some discussion.

### 5.1. Optimization protocol

Consider problem instances of the form $\left(f_{i}\right)_{i \in[N]}$, where $f_{i}: D \rightarrow \mathbb{R}$ for some common domain $D$ and all $i \in[N]$. We say that an algorithm operating on $\left(f_{i}\right)_{i \in[N]}$ is an $N$-element algorithm if it uses the following iterative protocol. At iteration $t$, the algorithm produces a query $i_{t}$, $x_{t}$, with $i_{t} \in[N]$ and $x_{t} \in D$. It then observes the output of a local oracle for $f_{i_{t}}$ at the point $x_{t}$, which we denote by $\mathcal{O}_{f_{i_{t}}}^{\text {loc }}\left(x_{t}\right)$.

Formally, $\mathcal{O}^{\text {loc }}$ can be any mapping that satisfies $\mathcal{O}_{f}^{\text {loc }}(x)=\mathcal{O}_{\tilde{f}}^{\text {loc }}(x)$ whenever $f(y)=\tilde{f}(y)$ for all $y$ in some open set containing $x$ (subsequently referred to as a "neighborhood" of $x$ ). In particular, the first-order oracle used for our upper bounds corresponds to $\mathcal{O}_{f}^{\text {loc }}(x)=(f(x), \nabla f(x))$ and is valid local oracle. The $p$ th order derivative oracle $\mathcal{O}_{f}^{\text {loc }}(x)=\left(f(x), \nabla f(x), \ldots, \nabla^{p} f\right)$ is also a valid local oracle. The notion of local oracles is classical in the literature on information-based complexity [29, 18].

The algorithms we consider may be randomized, and we use $\zeta$ to denote the algorithm's randomness. Beyond $\zeta$, the query of the algorithm at iteration $t$ may only depend on the information it
observes from the oracle. That is, for any $t \geq 1$, we have

$$
\begin{equation*}
i_{t}, x_{t}=Q_{t}\left(\zeta, \mathcal{O}_{f_{i_{1}}}^{\text {loc }}\left(x_{1}\right), \ldots, \mathcal{O}_{f_{i_{t-1}}}^{\text {loc }}\left(x_{t-1}\right)\right) \tag{9}
\end{equation*}
$$

for some measurable function $Q_{t}$.

### 5.2. Progress control argument

Following well-established methodology [33, 18, 11], instead of directly bounding the sub-optimality of the queries $x_{1}, \ldots, x_{t}$ we first bound a surrogate quantity we call progress. Informally, the progress is the highest coordinate index that the algorithm managed to "discover" using the oracle responses. Formally, we define the progress of a point $x$ as

$$
\begin{equation*}
\operatorname{prog}_{\alpha}(x):=\max \left\{i \geq 1| | x_{[i]} \mid>\alpha\right\} \quad(\text { where } \max \emptyset:=0) . \tag{10}
\end{equation*}
$$

The parameter $\alpha$ is a significance threshold for declaring a coordinate "discovered;" it allows us to prevent algorithms from trivially discovering coordinates by querying directions at random.

We next define a structural property that facilitates controlling the rate with which $\operatorname{prog}_{\alpha}\left(x_{t}\right)$ increases. For this definition, we recall that $v_{[\leq l]}$ denotes the vector whose first $l$ coordinates are identical to those of $v$ and the remainder are zero. Recall also that $\mathbb{B}_{1}^{T}(0)$ is the unit ball in $\mathbb{R}^{T}$.

Definition 7 A sequence $f_{1}, \ldots, f_{N}$ of functions $f_{i}: \mathbb{B}_{1}^{T}(0) \rightarrow \mathbb{R}$ is called an $\alpha$-robust $N$-element zero-chain if for all $x \in \mathbb{B}_{1}^{T}(0)$, all $y$ in a neighborhood of $x$, and all $i \in[N]$, we have

$$
\operatorname{prog}_{\alpha}(x) \leq p \Longrightarrow f_{i}(y)= \begin{cases}f_{i}\left(y_{[\leq p]}\right) & i<p+1  \tag{11}\\ f_{i}\left(y_{[\leq p+1]}\right) & i=p+1 \\ f_{N}\left(y_{[\leq p]}\right) & i>p+1\end{cases}
$$

To unpack this definition, consider any first-order algorithm with the following two simplifying properties: (1) the queries $i_{1}, i_{2}, \ldots$ are drawn i.i.d. from Uniform $([N])$ and (2) every query $x_{t}$ lies in the span of previously observed gradients $\nabla f_{i_{1}}\left(x_{1}\right), \ldots, \nabla f_{i_{t-1}}\left(x_{t-1}\right)$ [cf. 33]. The first query of the algorithm must be $x_{1}=0$, and consequently $\operatorname{prog}_{\alpha}\left(x_{1}\right)=0$. Definition 7 then implies that $f_{2}, \ldots, f_{N}$ are all constant in a neighborhood of $x_{1}$, while $f_{1}$ depends only on the first coordinate. Therefore, the span of the gradients (and the next query's progress) can only increase to 1 after the algorithm queries $i=1$ for the first time. With uniformly random index queries, that takes $\Omega(N)$ queries with constant probability. Repeating this argument, we see that every increase of the gradient span (and hence query progress) takes $\Omega(N)$ queries with constant probability, and therefore reaching progress $T$ takes $\Omega(N T)$ queries with high probability.

To extend this conclusion to general algorithms of the form (9), we perform two types of randomization. First, to handle arbitrary strategies for choosing $i_{t}$ (as opposed to uniform sampling), we apply a random permutation to $f_{1}, \ldots, f_{N}$. Second, to handle arbitrary queries $x_{t}$ (as opposed to queries in the span of observed gradients), we randomly rotate the coordinate system. This randomization scheme guarantees that no algorithm can materially improve on uniform sampling and span-preserving, as we formally state in the following.

Proposition 8 Let $\delta, \alpha \in(0,1)$ and let $N, T \in \mathbb{N}$ with $T \leq N / 2$. Let $\left(f_{i}\right)_{i \in[N]}$ be an $\alpha$-robust $N$-element zero-chain with domain $\mathbb{B}_{1}^{T}(0)$. For $d \geq T+\frac{2}{\alpha^{2}} \log \frac{4 N T^{2}}{\delta}$, draw $U$ uniformly from the
set of $d \times T$ orthogonal matrices, and draw $\Pi$ uniformly from the set of permutations of $[N]$. Let $\tilde{f}_{i}(x):=f_{\Pi^{-1}(i)}\left(U^{\top} x\right)$. Let $\left\{\left(i_{t}, x_{t}\right)\right\}_{t \geq 1}$ be the queries of any $N$-element algorithm operating on $\tilde{f}_{1}, \ldots, \tilde{f}_{N}$. Then with probability at least $1-\delta$ we have

$$
\operatorname{prog}_{\alpha}\left(U^{\top} x_{t}\right)<T \text { for all } t \leq \frac{1}{16} N\left(T-\log \frac{2}{\delta}\right) .
$$

See Appendix D. 1 for a proof. Our definition of $N$-element zero-chains and our proof of their progress control property builds on the notion of (single element) zero-chain functions [11]. It is also closely related to probability- $p$ zero-chains [3]; Proposition 8 essentially shows that $N$-element algorithms interacting with an $N$-element zero-chain make progress about as slowly as stochastic algorithms interacting with with a probability- $N^{-1}$ zero-chain.

### 5.3. Hard instance construction

With the progress-control machinery in hand, we proceed to constructing a specific $N$-element zerochain that also guarantees a large optimality gap for points with progress smaller than $T$. Toward that end, we first define the "link function" $\psi_{\alpha, \ell}: \mathbb{R} \rightarrow \mathbb{R}_{+}$as

$$
\psi_{\alpha, \ell}(t):= \begin{cases}0 & |t| \leq \alpha \\ \frac{\ell}{2}(t-\alpha)^{2} & \alpha \leq|t| \leq \ell^{-1}+\alpha \\ |t|-\alpha-\frac{1}{2 \ell} & \text { otherwise }\end{cases}
$$

Clearly, $\psi_{\alpha, \ell}$ is 1-Lipschitz, $\ell$-smooth, and is identically zero for all $|t| \leq \alpha$. We note that $\psi_{\alpha, \ell}$ is the composition of the Huber function [22] with $\max \{0,|t|-\alpha\}$.

Chain constructions of the form $\sum_{i \in[N]} \psi_{\alpha_{T}, \ell}\left(x_{[i]}-x_{[i-1]}\right)$ are common in lower bounds for convex optimization [cf. 33, 39]. For our construction, we instead spread the link components across the different elements. Formally, for $i \in[N]$, we define the $i$ th function in the our hard instance as

$$
\hat{f}_{i}^{\{T, N, \ell\}}(x):=\left\{\begin{array}{ll}
\psi_{\alpha_{T}, \ell}\left(\frac{x_{[i]}-x_{[i-1]}}{2}\right) & i \leq T  \tag{12}\\
0 & \text { otherwise }
\end{array} \text { where } \alpha_{T}:=\frac{1}{4 T^{3 / 2}} \text { and } x_{[0]}:=\frac{1}{\sqrt{T}} .\right.
$$

The following lemma summarizes the properties of our construction. The proof of the lemma is straightforward and we provide it in Appendix D. 2

Lemma 9 For every $T, N \in \mathbb{N}$ and $\ell \geq 0$, such that $T \leq N$, we have that

1. The hard instance $\left(\hat{f}_{i}^{\{T, N, \ell\}}\right)_{i \in N}$ is an $\alpha_{T}$-robust $N$-element zero-chain.
2. The function $\hat{f}_{i}^{\{T, N, \ell\}}$ is 1-Lipschitz and $\ell$-smooth for every $i \in[N]$.
3. For $x \in \mathbb{R}^{d}$ with $\operatorname{prog}_{\alpha_{T}}(x)<T$, the objective $\hat{F}_{\max }^{\{T, N, \ell\}}(x)=\max _{i \in[N]} \hat{f}_{i}^{\{T, N, \ell\}}(x)$ satisfies

$$
\hat{F}_{\max }^{\{T, N, \ell\}}(x)-\min _{x_{\star} \in \mathbb{B}_{1}(0)} \hat{F}_{\max }^{\{T, N, \ell\}}\left(x_{\star}\right) \geq \psi_{\alpha_{T}, \ell}\left(\frac{3}{8 T^{3 / 2}}\right) \geq \min \left\{\frac{1}{8 T^{3 / 2}}, \frac{\ell}{32 T^{3}}\right\} .
$$

### 5.4. Lower bound statement

Finally, we combine the results of the previous sections to state our lower bound. In the statement, we use $a \wedge b:=\min \{a, b\}$ to abbreviate binary minimization.

Theorem 10 Let $L_{f}, L_{g}, R>0, \epsilon<L_{f} R \wedge L_{g} R^{2}, N \in \mathbb{N}$ and $\delta \in(0,1)$. Then, for any (possibly randomized) algorithm there exists an $L_{f}$-Lipschitz and $L_{g}$-smooth functions $\left(f_{i}\right)_{i \in[n]}$ with domain $\mathbb{B}_{R}^{d}(0)$ for $d=O\left(\left[\left(\frac{L_{f} R}{\epsilon}\right)^{2} \wedge\left(\frac{L_{g} R^{2}}{\epsilon}\right)\right] \log \frac{N\left(L_{f} R \wedge L_{g} R^{2}\right)}{\epsilon}\right)$ such that with probability at least $\frac{1}{2}$ over the randomness of the algorithm, the first

$$
\begin{equation*}
\Omega\left(N\left[\left(\frac{L_{f} R}{\epsilon}\right)^{2 / 3} \wedge\left(\frac{L_{g} R^{2}}{\epsilon}\right)^{1 / 3}\right]+\left[\left(\frac{L_{f} R}{\epsilon}\right)^{2} \wedge\left(\frac{N L_{g} R^{2}}{\epsilon}\right)^{1 / 2}\right]\right) \tag{13}
\end{equation*}
$$

queries of the algorithm are all $\epsilon$-suboptimal for $F_{\max }(x)=\max _{i \in[N]} f_{i}(x)$.
See Appendix D. 3 for a proof of this result. The first (linear-in- $N$ ) term in the lower bound follows from Proposition 8 and Lemma 9 via a re-scaling argument. The second (sublinear-in- $N$ ) lower bound term is a direct consequence of existing lower bounds [14, 39, 15].

We remark that our lower bound is stated for optimization constrained to a ball of radius $R$, while our upper bounds assume unconstrained optimization given a minimizer of norm at most $R$. These two settings are essentially equivalent; in Appendix D. 4 we sketch a general technique for transferring lower bounds to the unconstrained setting.

In Table 1 we specify our lower bound in the special cases $L_{g}=\infty$ and $L_{g}=\Theta\left(L_{f}^{2} / \epsilon\right)$, showing that they match our upper bounds (up to polylogarithmic factors) for $N=\Omega\left(\left(L_{f} R / \epsilon\right)^{2}\right)$ in the former case and for any $N$ in the latter. More broadly, when $L_{g}=\Theta\left(L_{f}^{2+q} R^{q} / \epsilon^{1+q}\right)$ our lower and upper bounds match for any $N$ and $q \in[0,2 / 3]$. For $L_{g}=o\left(L_{f}^{2} / \epsilon\right)$ and $L_{g}=\omega\left(L_{f}^{8 / 3} R^{2 / 3} / \epsilon^{5 / 3}\right)$, however, there remain gaps between our upper and lower bounds. We discuss these gaps in the following section.

## 6. Discussion

To conclude the paper, we provide some commentary on our results and the possibilities of improving them. For simplicity, in this section we revert to the setting $L_{f}=R=1$ used in the introduction. We also use $a \ll b$ as a shorthand for $a=O(b)$, and ignore constant and logarithmic factors throughout.

### 6.1. Gaps between the upper and lower bounds

Regimes where a gap exists. Comparing our upper bound in Theorem 6 to our lower bound in Theorem 10, we identify two regimes where our upper and lower bounds disagree by more than polylogarithmic factors. The first is the smooth regime $L_{g} \ll \epsilon^{-1}$, the lower bound is $\Omega\left(N L_{g}^{1 / 3} \epsilon^{-1 / 3}+\right.$ $\left.\sqrt{N L_{g} \epsilon^{-1}}\right)$ while our upper bound is $\widetilde{O}\left(N \epsilon^{-2 / 3}+\sqrt{N} \epsilon^{-1}\right)$, and a different algorithm gives a better oracle complexity $O\left(N \sqrt{L_{g} \epsilon^{-1}}\right)$ (see Appendix A.3) which still falls short of the lower bound.

The second regime is the non-smooth regime $L_{g} \gg \epsilon^{-1}$, where both the upper and lower bounds share the term $N \epsilon^{-2 / 3}$. Comparing the lower bound to the variance reduced upper bound (6), we see that they disagree if and only if $N \epsilon^{-2 / 3}+\epsilon^{-2} \ll \sqrt{N L_{g} \epsilon^{-1}}$ which is equivalent to $N \ll L_{g} \epsilon^{1 / 3}$
and $N \gg \epsilon^{-3} / L_{g}$. Clearly, this is only possible only when $L_{g} \gg \epsilon^{-5 / 3}$, and so we conclude that the rate (6) is in fact optimal whenever $\epsilon^{-1} \ll L_{g} \ll \epsilon^{-5 / 3}$. Moreover, the upper bound (8) matches the lower bound whenever $N \gg \epsilon^{-2}$ for any $L_{g} \gg \epsilon^{-1}$. We conclude that gaps in the non-smooth regime exist only for $L_{g} \gg \epsilon^{-5 / 3}$ and $\epsilon^{-3} / L_{g} \ll N \ll \min \left\{\epsilon^{-2}, \epsilon^{1 / 3} L_{g}\right\}$.
Closing the gap in the non-smooth regime. Improving the bound (8) from $\widetilde{O}\left(N \epsilon^{-2 / 3}+\epsilon^{-8 / 3}\right)$ to $\widetilde{O}\left(N \epsilon^{-2 / 3}+\epsilon^{-2}\right)$ would imply that (13) gives the optimal rate for any $L_{g} \gg \epsilon^{-1}$. The main barrier for obtaining such improvement is our accuracy requirements $\delta_{t}=O\left(\epsilon / \lambda_{t}\right)$ in Algorithm 1. Meeting this requirement with SGD means that each oracle implementation costs $\widetilde{O}\left(N+\epsilon^{-2}\right)$ function/gradient evaluations, and multiplying this cost by the number of rounds $\widetilde{O}\left(\epsilon^{-2 / 3}\right)$ yields the exponent $8 / 3$. A variant of Algorithm 1 which can handle less accurate BROO outputs could close this gap by allowing a more efficient SGD-based implementation.

Closing the gap in the smooth regime. The gap between our upper and lower bounds when $L_{g} \ll \epsilon^{-1}$ is more fundamental than the one arising for $L_{g} \gg \epsilon^{-5 / 3}$, because it affects the term linear in $N$. The barrier for improving the linear term in our algorithm is the ball radius. Any $r_{\epsilon^{-}}$ BROO implementation with $\Omega(N)$ cost will have overall complexity $\Omega\left(N r_{\epsilon}^{-2 / 3}\right)$. The techniques we develop in Section 4 only allow us to support $r_{\epsilon}=\widetilde{O}(\epsilon)$, because this is the largest radius where the exponentiated softmax is stable (see Lemma 3).

Conjectures and future work. We conjecture that our lower bound is in fact optimal in both smoothness regimes. In future work we will attempt to close the remaining complexity gaps described above.

### 6.2. Some necessary algorithmic structures

Several aspects of our method, namely functions value access, individual function queries and randomization are necessary for any method that achieves (or improves on) our complexity bounds. See Appendix A. 4 for detailed discussion.

### 6.3. Practical considerations

The main purpose of the algorithms we develop in this paper is to clarify the complexity of the fundamental optimization (1). Nevertheless, since this problem formulation is relevant for a number of machine learning tasks [13, 21, 36], it is interesting to try and develop a more practical variant of algorithms. Two aspects of our method which we believe will be particularly useful in practice are the gradient estimation scheme we use in Algorithm 2 and the momentum scheme in Algorithm 1.

However, a number of aspects of our method seem rather impractical. First, the theory instructs us to constrain subproblem solutions to a very small ball of radius $r_{\epsilon}$ of roughly $\epsilon / L_{f}$. Since usually neither $\epsilon$ or $L_{f}$ are known in advance, the parameter $r_{\epsilon}$ must be tuned. Moreover, choosing $r_{\epsilon}$ to be small in keeping with the theory would likely mean very slow progress in the early stages of the algorithm. A second impractical aspect is the bisection stage in Algorithm 1: while in theory the bisection only increases complexity by a logarithmic factor, in practice it entails solving a considerable number of sub-problems without making progress. This bisection overhead is an issue with Monteiro-Svaiter acceleration more broadly and a topic of active research [37, 34].

## Acknowledgments

YC was supported in part by Len Blavatnik and the Blavatnik Family foundation, and the Yandex Machine Learning Initiative for Machine Learning. YJ was supported by Stanford Graduate Fellowship. AS was supported in part by a Microsoft Research Faculty Fellowship, NSF CAREER Award CCF-1844855, NSF Grant CCF-1955039, a PayPal research award, and a Sloan Research Fellowship.

## References

[1] Zeyuan Allen-Zhu. Katyusha: The first direct acceleration of stochastic gradient methods. arXiv preprint arXiv:1603.05953, 2016.
[2] Zeyuan Allen-Zhu and Lorenzo Orecchia. Linear coupling: An ultimate unification of gradient and mirror descent. In Innovations in Theoretical Computer Science, 2017.
[3] Yossi Arjevani, Yair Carmon, John C Duchi, Dylan J Foster, Nathan Srebro, and Blake Woodworth. Lower bounds for non-convex stochastic optimization. arXiv preprint arXiv:1912.02365, 2019.
[4] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. Robust optimization. Princeton University Press, 2009.
[5] Aharon Ben-Tal, Dick Den Hertog, Anja De Waegenaere, Bertrand Melenberg, and Gijs Rennen. Robust solutions of optimization problems affected by uncertain probabilities. Management Science, 59(2):341-357, 2013.
[6] Jan van den Brand, Yin Tat Lee, Yang P Liu, Thatchaphol Saranurak, Aaron Sidford, Zhao Song, and Di Wang. Minimum cost flows, MDPs, and $\ell_{1}$-regression in nearly linear time for dense instances. arXiv preprint arXiv:2101.05719, 2021.
[7] Sebastien Bubeck, Qijia Jiang, Yin Tat Lee, Yuanzhi Li, and Aaron Sidford. Complexity of highly parallel non-smooth convex optimization. In Advances in Neural Information Processing Systems, 2019.
[8] Sébastien Bubeck et al. Convex optimization: Algorithms and complexity. Foundations and Trends in Machine Learning, 2015.
[9] Brian Bullins. Highly smooth minimization of non-smooth problems. In Conference on Learning Theory, pages 988-1030, 2020.
[10] Yair Carmon, Yujia Jin, Aaron Sidford, and Kevin Tian. Variance reduction for matrix games. In Advances in Neural Information Processing Systems, 2019.
[11] Yair Carmon, John C. Duchi, Oliver Hinder, and Aaron Sidford. Lower bounds for finding stationary points I. Mathematical Programming, 184(1):71-120, 2020.
[12] Yair Carmon, Arun Jambulapati, Qijia Jiang, Yujia Jin, Yin Tat Lee, Aaron Sidford, and Kevin Tian. Acceleration with a ball optimization oracle. In Advances in Neural Information Processing Systems, 2020.
[13] Kenneth L Clarkson, Elad Hazan, and David P Woodruff. Sublinear optimization for machine learning. Journal of the ACM (JACM), 59(5):1-49, 2012.
[14] Jelena Diakonikolas and Cristóbal Guzmán. Lower bounds for parallel and randomized convex optimization. Journal of Machine Learning Research, 21(5):1-31, 2020.
[15] Cong Fang, Chris Junchi Li, Zhouchen Lin, and Tong Zhang. Near-optimal non-convex optimization via stochastic path integrated differential estimator. Advances in Neural Information Processing Systems, 31:689, 2018.
[16] Roy Frostig, Rong Ge, Sham Kakade, and Aaron Sidford. Un-regularizing: approximate proximal point and faster stochastic algorithms for empirical risk minimization. In International Conference on Machine Learning, 2015.
[17] Alexander Gasnikov, Pavel Dvurechensky, Eduard Gorbunov, Evgeniya Vorontsova, Daniil Selikhanovych, César A. Uribe, Bo Jiang, Haoyue Wang, Shuzhong Zhang, Sébastien Bubeck, Qijia Jiang, Yin Tat Lee, Yuanzhi Li, and Aaron Sidford. Near optimal methods for minimizing convex functions with Lipschitz p-th derivatives. In Conference on Learning Theory, 2019.
[18] Cristóbal Guzmán and Arkadi Nemirovski. On lower complexity bounds for large-scale smooth convex optimization. Journal of Complexity, 31(1):1-14, 2015.
[19] Osman Güler. New proximal point algorithms for convex minimization. SIAM Journal on Optimization, 2(4):649-664, 1992.
[20] Elad Hazan and Satyen Kale. Beyond the regret minimization barrier: optimal algorithms for stochastic strongly-convex optimization. The Journal of Machine Learning Research, 15(1): 2489-2512, 2014.
[21] Elad Hazan, Tomer Koren, and Nati Srebro. Beating SGD: Learning SVMs in sublinear time. In Advances in Neural Information Processing Systems, pages 1233-1241, 2011.
[22] Peter J Huber. Robust estimation of a location parameter. In Breakthroughs in statistics, pages 492-518. Springer, 1992.
[23] Yin Tat Lee and Aaron Sidford. Efficient inverse maintenance and faster algorithms for linear programming. In 2015 IEEE 56th Annual Symposium on Foundations of Computer Science, pages 230-249. IEEE, 2015.
[24] Anatoly Yur'evich Levin. An algorithm for minimizing convex functions. Doklady Akademii Nauk SSSR, 160(6):1244-1247, 1965.
[25] Hongzhou Lin, Julien Mairal, and Zaid Harchaoui. A universal catalyst for first-order optimization. In Advances in Neural Information Processing Systems, 2015.
[26] Renato DC Monteiro and Benar Fux Svaiter. An accelerated hybrid proximal extragradient method for convex optimization and its implications to second-order methods. SIAM Journal on Optimization, 23(2):1092-1125, 2013.
[27] Hongseok Namkoong and John C Duchi. Stochastic gradient methods for distributionally robust optimization with f-divergences. In Advances in Neural Information Processing Systems, volume 29, pages 2208-2216, 2016.
[28] Arkadi Nemirovski. Prox-method with rate of convergence $O(1 / t)$ for variational inequalities with Lipschitz continuous monotone operators and smooth convex-concave saddle point problems. SIAM Journal on Optimization, 15(1):229-251, 2004.
[29] Arkadi Nemirovski and David Yudin. Problem complexity and method efficiency in optimization. Wiley-Interscience, 1983.
[30] Arkadi Nemirovski, Anatoli Juditsky, Guanghui Lan, and Alexander Shapiro. Robust stochastic approximation approach to stochastic programming. SIAM Journal on Optimization, 19(4): 1574-1609, 2009.
[31] Yurii Nesterov. Smooth minimization of non-smooth functions. Mathematical programming, 103(1):127-152, 2005.
[32] Yurii Nesterov. Dual extrapolation and its applications to solving variational inequalities and related problems. Mathematical Programming, 109(2):319-344, 2007.
[33] Yurii Nesterov. Lectures on convex optimization, volume 137. Springer, 2018.
[34] Yurii Nesterov. Implementable tensor methods in unconstrained convex optimization. Mathematical Programming, pages 1-27, 2019.
[35] D. J. Newman. Location of the maximum on unimodal surfaces. J. ACM, 12(3):395-398, 1965.
[36] Shai Shalev-Shwartz and Yonatan Wexler. Minimizing the maximal loss: How and why. In International Conference on Machine Learning, pages 793-801, 2016.
[37] Chaobing Song, Yong Jiang, and Yi Ma. Unified acceleration of high-order algorithms under Hölder continuity and uniform convexity. arXiv preprint arXiv:1906.00582, 2019.
[38] Vladimir N Vapnik. An overview of statistical learning theory. IEEE transactions on neural networks, 10(5):988-999, 1999.
[39] Blake Woodworth and Nathan Srebro. Tight complexity bounds for optimizing composite objectives. In Advances in Neural Information Processing Systems, pages 3646-3654, 2016.

## Appendix A. Additional discussion

Here we discuss three important points that exceeded the scope of our introduction. First, Appendix A. 1 explains what makes $L_{f}^{2} / \epsilon$ a natural smoothness scale and when we can expect the $f_{i}$ to be at least that smooth. Then, Appendix A. 2 discusses the maximally smooth case of linear loss functions, and compare our results to guarantees of methods specialized to this setting. Finally, Appendix A. 3 considers the computational complexity of implementing the steps in the accelerated iterative linearization scheme of [33, Section 2.3.1].

## A.1. The generality of the smoothness assumption $L_{g}=O\left(L_{f}^{2} / \epsilon\right)$

Let $f$ be a convex, $L_{f}$ Lipschitz function. When $f$ is not continuously differentiable, it is still possible to uniformly approximate it with a continuously differentiable function, that moreover has a Lipschitz gradient. More concretely, for every $\varepsilon>0$, we may consider the infimal convolution of $f$ with a quadratic regularizer (also known as its Moreau envelope):

$$
\begin{equation*}
\tilde{f}(x)=\min _{y \in \mathbb{R}^{d}}\left\{f(y)+\frac{L_{f}^{2}}{2 \varepsilon}\|x-y\|^{2}\right\} . \tag{14}
\end{equation*}
$$

It holds that $0 \leq f(x)-\tilde{f}(x) \leq \varepsilon / 2$ for all $x \in \mathbb{R}^{d}$ and moreover that $\tilde{f}$ is $L_{f}^{2} / \varepsilon$ smooth (i.e., with $L_{f}^{2} / \varepsilon$ Lipschitz gradient) [see, e.g., 18]. Therefore, if we wish to minimize $f$ to accuracy $\epsilon$, we may choose $\varepsilon=\epsilon$ and replace $f$ with the $O\left(L_{f}^{2} / \epsilon\right)$-smooth function $\tilde{f}$.

The computational cost of such replacement depends on the application. In certain cases, smoothed versions of the $f_{i}$ 's have closed-form expression, and we may simply define the problem with them instead. Moreover, in several machine learning applications each function $f_{i}$ is "simple," and querying index $i$ really corresponds to obtaining full access to this function, in which case directly computing (14) might be feasible.

However, when we are truly restricted to accessing $f_{i}$ through a gradient and value black box, there will be some instances where computing (14) (and indeed any other smoothing) is substantially more expensive than a single oracle query. To see this, note that the worst-case oracle complexity of minimizing a single non-smooth convex function scales as $\Omega\left(L_{f}^{2} R^{2} \epsilon^{-2}\right)$, while the complexity of minimizing an $L_{g}$-smooth function scales as $O\left(\sqrt{L_{g} R^{2} \epsilon^{-1}}\right)$. If $T_{\epsilon}$ is the worst number of oracle calls required to compute an $O(\epsilon)$-accurate $O\left(L_{f}^{2} / \epsilon\right)$-smooth approximation for any $L_{f}$-Lipschitz, we immediately have a general complexity upper bound of $O\left(T_{\epsilon} L_{f} R \epsilon^{-1}\right)$. Comparing it to the lower bound $\Omega\left(L_{f}^{2} R^{2} \epsilon^{-2}\right)$ immediately yields that $T_{\epsilon}=\Omega\left(L_{f} R \epsilon^{-1}\right)$ in the worst case. Therefore, the fact that our algorithm maintains the same leading order dependence on $N$ even for $L_{g}=\infty$ is nontrivial and potentially useful.

## A.2. The special case of linear loss functions ( $L_{g}=0$ )

When the losses are linear (i.e., $f_{i}(x)=a_{i}^{\top} x+b_{i}$ ) exactly $N$ function value and gradient evaluations suffice to completely identify the problem instance, so the optimal oracle complexity should never be more than $N$. In this setting, then, it is more relevant to discuss the computational (runtime) complexity of solving the problem (1). To simplify the following discussion, we return to the setting in the introduction where each $f_{i}$ is 1-Lipschitz (i.e., $\nabla f_{i}$ has norm at most 1 ) and we assume the existence of a minimizer of $F_{\max }$ with Euclidean norm at most 1 .

In the linear setting, an equivalent form of the problem (1) is

$$
\underset{x \in \mathbb{R}^{d}}{\operatorname{minimize}} \max _{p \in \Delta^{N}} p^{\top}(A x-b)
$$

where $A$ is a matrix whose $i$ th row contains $\nabla f_{i}$. Stochastic primal-dual methods are able to take advantage of this matrix structure to obtain cheap unbiased estimates for the gradients of $p^{T} A x$ with respect to both $x$ and $p$, sampling rows and columns of $A$ respectively. Assuming that reading a row and a column of $A$ takes $O(N+d)$ time, the stochastic primal-dual method has runtime complexity $\widetilde{O}\left((N+d) \epsilon^{-2}\right)$ which, for sufficiently large $\epsilon$, is sublinear in the problem size $N d$ [13]. For lower values of $\epsilon$, a variance reduction technique [10] has preferable runtime complexity $\widetilde{O}(N d+$ $\left.\sqrt{N d(N+d)} \epsilon^{-1}\right)$.

In comparison, in the linear case our method has runtime complexity $\widetilde{O}\left(N d \epsilon^{-2 / 3}+\sqrt{N} d \epsilon^{-1}\right)$ assuming that sampling a row takes $O(d)$ time. This improves on the variance reduction method in the somewhat narrow parameter regime $N=\widetilde{\Omega}(d)$ and $d=\widetilde{O}\left(\epsilon^{-2 / 3}\right)$. However, we note that our method operates under a strictly weaker access assumption, since it only samples rows and not columns. In scenarios where accessing a column takes $\omega(N)$ time, the relative merit of the methods changes. In the most extreme case where reading a column of $A$ is as expensive as reading the entire matrix (e.g., because the rows are scattered across many devices), the stochastic primal-dual methods become less efficient than exact-gradient counterparts with runtime $\widetilde{O}\left(N d \epsilon^{-1}\right)[31,28,32]$, while the runtime of our method is unchanged and always superior. To the best of our knowledge, this is the first guarantee for a stochastic gradient method that improves on exact gradient methods in a first-order oracle model that can only provides rows of $A$.

The literature also considers high-order methods for solving the linear case of problem (1) to better accuracy but with potentially worse dependence on problem dimension. Bullins [9] proposes a fourth-order accelerated regularization method that requires $\widetilde{O}\left(\epsilon^{-4 / 5}\right)$ solutions of linear systems of the form $A^{\top} D A=b$ for a positive diagonal matrix $D$. Carmon et al. [12] use ball oracle acceleration to obtain an improved method requiring only $\widetilde{O}\left(\epsilon^{-2 / 3}\right)$ linear system solutions. They also propose to solve these systems using an efficient first-order method, resulting in a runtime guarantee of $\widetilde{O}\left(N d \epsilon^{-2 / 3}+d^{3 / 2} \epsilon^{-5 / 3}\right)$. We believe more careful reasoning about the conditioning of each linear system to be solved (as we do in Section 4.3) would improve this guarantee to $\widetilde{O}\left(N d \epsilon^{-2 / 3}+d^{3 / 2} \epsilon^{-1}\right)$ under the assumption that individual entries of $A$ take $O(1)$ time to read. In the linear case, this improves on our result when $d<N$, albeit with stronger matrix access assumptions.

For even higher accuracy, it is possible to express the linear case of problem (1) as a linear program and solve it using interior point methods. The best existing theoretical runtimes for these methods are $\widetilde{O}\left(\left(N d+(N \wedge d)^{2}\right) \sqrt{N \wedge d}\right)$ [23] or $\widetilde{O}\left(N d+d^{2.5}\right)$ [6], both depending logarithmically on the desired accuracy $1 / \epsilon$. When the problem dimensions $N$ and $d$ are sufficiently large compared to $1 / \epsilon$, first-order methods are preferable.

## A.3. The computational complexity of accelerated iterative linearization

This subsection uses our full notation defined in Section 2. We suggest considering this section after reading Sections 3 and 4 as well.

In [33, Section 2.3.1], Nesterov shows how solving $O\left(\sqrt{L_{g} R^{2} \epsilon^{-1}}\right)$ subproblems of the form

$$
\min _{x \in \mathbb{R}^{d}} \max _{i \in[N]}\left\{f_{i}\left(y_{t}\right)+\left(\nabla f_{i}\left(y_{t}\right)\right)^{\top}\left(x-y_{t}\right)+\frac{L_{g}}{2}\left\|x-y_{t}\right\|^{2}\right\}
$$

allows solving the problem (1) when the functions are $L_{g}$-smooth. We note that each subproblem is equivalent to

$$
\min _{x \in \mathbb{R}^{d}} \max _{p \in \Delta^{N}}\left\{p^{\top}(A x-b)+\frac{L_{g}}{2}\|x\|^{2}\right\}
$$

for a matrix $A \in \mathbb{R}^{N \times d}$ whose rows have norm at most $L_{f}$. Consequently we may apply a variancereduced bilinear saddle-point method to solve the subproblem to additive error $\nu$ in time

$$
\widetilde{O}\left(N d+\sqrt{N d(N+d)} \frac{L_{f}}{\sqrt{L_{g} \nu}}\right)
$$

see Proposition 6 and the subsequent discussion in the arXiv version of [10].
Applying the same arguments used to prove Theorem 2-but with $\lambda_{t}=L_{g}$ for all $t$-we have that the required subproblem solution accuracy is $O\left(\frac{\epsilon^{2}}{L_{g} R^{2}}\right)$, and consequently we can solve each problem in time

$$
\widetilde{O}\left(N d+\sqrt{N d(N+d)} \frac{L_{f} R}{\epsilon}\right) .
$$

Assuming $\mathcal{T}_{f}+\mathcal{T}_{g}=\Omega(d)$, the overall cost of the method is

$$
\widetilde{O}\left(N\left(\mathcal{T}_{f}+\mathcal{T}_{g}\right) \sqrt{\frac{L_{g} R^{2}}{\epsilon}}+\sqrt{N d(N+d)} \frac{L_{f} \sqrt{L_{g}} R^{2}}{\epsilon^{3 / 2}}\right) .
$$

## A.4. Some necessary algorithmic structures

We now argue that several aspects of our method, namely functions value access, individual function queries and randomization are necessary in any method that achieves (or improves on) our complexity bounds.

Function value access. It is possible to minimize a convex function $f$ by iterative (sub)gradient evaluations, without access to the value of $f$ itself. In contrast, all algorithms for minimizing $F_{\max }=\max _{i \in[N]} f_{i}(x)$ must query the values of the $f_{i}$ 's in addition to their gradients. To see why this is so, consider the case where $f_{i}(x)=\Pi(i)-x_{i}$, where $\Pi$ is a random permutation of $[N]$ and the domain is the unit Euclidean ball. The global minimum of $\max _{i \in[N]} f_{i}(x)$ is the $\Pi^{-1}(N)$-th standard basis vector. However, gradients provide no information about $\Pi^{-1}(N)$, since $\nabla f_{i}(x)=-e_{i}$ for all $x$, independent of $\Pi$.

Individual-function access. The algorithms from prior work in Table 1 (namely the subgradient method, AGD on softmax and AGD on linearization) are full-batch methods: they proceed by querying all $N$ functions $f_{1}, \ldots, f_{N}$ at the same point $x_{t}$ and using the result to generate the next query point $x_{t+1}$. In contrast, our BROO implementations proceed by sampling an index $i_{t}$, computing $\nabla f_{i}$ at $x_{t}$ (and potentially another point), and generating the next query $x_{t+1}$. Full-batch methods are more amenable to parallelization, but for our problem have demonstrably worse oracle complexity. To see this, consider the case where all the $f_{i}$ 's are identical and equal to standard hard instance for convex optimization. For such input, any full-batch methods will have oracle complexity $\Omega\left(N \min \left\{\epsilon^{-2}, \sqrt{L_{g} \epsilon^{-1}}\right\}\right.$ ) [14], which is worse than our upper bounds for any $L_{g} \gg \epsilon^{-1 / 3}$ and sufficiently large $N$.

Randomization. Another contrast between the prior algorithms in Table 1 and our algorithm is that the former are deterministic while ours is randomized. Woodworth and Srebro [39] prove a lower bound of $\Omega\left(N \min \left\{\epsilon^{-2}, \sqrt{L_{g} \epsilon^{-1}}\right\}\right)$ gradient queries for any deterministic method for minimizing the average of $N$ functions. Observing that the maximum of the $N$ functions in their construction has the same minimum value as their average (and that the maximum upper bounds the average in any other points), we conclude that this lower bound is also valid for any deterministic method for solving the problem (1). Therefore, randomization is necessary for obtaining our improved rates of convergence.

## Appendix B. Proof of Theorem 2

In this section we give the analysis of our accelerated algorithm. Our analysis builds off of [12] and proceeds in several parts. We first prove several standard technical results Appendix B.1. Then, in Appendix B. 2 we give the proof of Theorem 2 assuming the correctness of the $\lambda$-Bisection subroutine. Finally in Appendix B. 3 we prove this correctness.

## B.1. Preliminary technical results

First, we observe that $\mathcal{O}_{\lambda, \delta}(y)$ returns a point which is close to $\operatorname{bprox}_{\lambda, r}^{f}(y)$, the true minimizer of the proximal objective.
Lemma 11 Let $f$ be a convex function and let $x=\operatorname{bprox}_{\lambda, r}^{f}(y)$. Then if $\mathcal{O}_{\lambda, \delta}(\cdot)$ is an $r$-BROO for $f$, the point $\widetilde{x}=\mathcal{O}_{\lambda, \delta}(y)$ satisfies $\|\widetilde{x}-x\| \leq \delta$.
Proof By Definition 1 of the BROO, we have

$$
\nu=\left[f(\widetilde{x})+\frac{\lambda}{2}\|\widetilde{x}-y\|^{2}\right]-\left[f(x)+\frac{\lambda}{2}\|x-y\|^{2}\right] \leq \frac{\lambda \delta^{2}}{2} .
$$

Since the function $f(x)+\frac{\lambda}{2}\|x-y\|^{2}$ is $\lambda$-strongly convex, we obtain $\nu \geq \frac{\lambda}{2}\|x-\widetilde{x}\|_{2}^{2}$, and substituting $\nu \leq \lambda \delta^{2} / 2$ gives the result.

Second, we provide standard facts regarding proximal mappings. Though this follows from standard facts regarding subgradients we provide a self-contained proof for completeness.

Lemma 12 Let $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ be a convex function, $S \subseteq \mathbb{R}^{d}$ be a closed convex set, $\lambda \geq 0$, and $x_{\lambda}, x_{0} \in S$ satisfy

$$
x_{\lambda}=\underset{x \in S}{\arg \min }\left\{f(x)+\frac{\lambda}{2}\left\|x-x_{0}\right\|^{2}\right\}
$$

Then $g_{\lambda}:=\lambda\left(x_{0}-x_{\lambda}\right)$ is a subgradient of $f$, i.e.,

$$
\begin{equation*}
f(y) \geq f\left(x_{\lambda}\right)+\left\langle g_{\lambda}, y-x_{\lambda}\right\rangle \text { for all } y \in S . \tag{15}
\end{equation*}
$$

Further, we have

$$
\begin{equation*}
g_{\lambda}^{\top}\left(x_{\lambda}-z\right)=\frac{\lambda}{2}\left\|z-x_{0}\right\|^{2}-\frac{\lambda}{2}\left\|z-x_{\lambda}\right\|^{2}-\frac{\lambda}{2}\left\|x_{\lambda}-x_{0}\right\|^{2} \text { for all } z \in \mathbb{R}^{d} \tag{16}
\end{equation*}
$$

and

$$
\begin{equation*}
f\left(x_{\lambda}\right) \leq f(y)+\frac{\lambda}{2}\left\|y-x_{0}\right\|^{2}-\frac{\lambda}{2}\left\|y-x_{\lambda}\right\|^{2}-\frac{\lambda}{2}\left\|x_{\lambda}-x_{0}\right\|^{2} \text { for all } y \in S . \tag{17}
\end{equation*}
$$

Proof Let $x^{\alpha}:=\alpha \cdot y+(1-\alpha) \cdot x_{\lambda}$ for all $\alpha \in(0,1]$. Note that $x^{\alpha} \in S$ for all $\alpha \in(0,1]$ since $S$ is convex and $x_{\lambda}, y \in S$. Consequently, (15) and convexity of $f$ imply that for all $\alpha \in[0,1]$

$$
\begin{aligned}
f\left(x_{\lambda}\right)+\frac{\lambda}{2}\left\|x_{\lambda}-x_{0}\right\|^{2} & \leq f\left(x^{\alpha}\right)+\frac{\lambda}{2}\left\|x_{0}-x^{\alpha}\right\|^{2} \\
& \leq \alpha f(y)+(1-\alpha) f\left(x_{\lambda}\right)+\frac{\lambda}{2}\left\|\alpha\left(x_{0}-y\right)+(1-\alpha)\left(x_{0}-x_{\lambda}\right)\right\|^{2} .
\end{aligned}
$$

Rearranging yields that for all $\alpha>0$ implies

$$
f\left(x_{\lambda}\right)-f(y) \leq \frac{\lambda}{2 \alpha}\left[\alpha^{2}\left\|x_{0}-y\right\|+2 \alpha(1-\alpha)\left\langle x_{0}-y, x_{0}-x_{\lambda}\right\rangle+(1-\alpha)^{2}\left\|x_{\lambda}-x_{0}\right\|^{2}-\left\|x_{\lambda}-x_{0}\right\|^{2}\right]
$$

Taking the limit as $\alpha \rightarrow 0$ yields that

$$
f\left(x_{\lambda}\right)-f(y) \leq \lambda\left(x_{0}-y\right)^{\top}\left(x_{0}-x_{\lambda}\right)-2 \lambda\left\|x_{\lambda}-x_{0}\right\|^{2}=\lambda\left(x_{\lambda}-x_{0}\right)^{\top}\left(x_{0}-x_{\lambda}\right)
$$

The remaining claims (16) and (17) follow from direct algebraic manipulation of this inequality and the definition $g_{\lambda}=\lambda\left(x_{0}-x_{\lambda}\right)$.

Third, we bound the function error induced by proximal mapping.
Lemma 13 Let $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ be a convex function, $\lambda \geq 0$, and $x_{\lambda}, x_{0} \in \mathbb{R}^{d}$ satisfy $x_{\lambda}=\operatorname{prox}_{\lambda}^{f}\left(x_{0}\right)$. If $y \in \mathbb{R}^{d}$ satisfies $\left\|y-x_{0}\right\| \leq R$ and $\left\|x_{\lambda}-x_{0}\right\| \leq \Theta$ then $f\left(x_{\lambda}\right)-f(y) \leq \lambda \Theta R$.

Proof Bound (15) in Lemma 12 yields

$$
\begin{aligned}
f\left(x_{\lambda}\right)-f(y) & \leq-\lambda\left\langle x_{0}-x_{\lambda}, y-x_{\lambda}\right\rangle=-\lambda\left\langle x_{0}-x_{\lambda}, y-x_{0}\right\rangle-\lambda\left\|x_{0}-x_{\lambda}\right\|^{2} \\
& \leq \lambda\left\langle x_{\lambda}-x_{0}, y-x_{0}\right\rangle \leq \lambda\left\|x_{0}-x_{\lambda}\right\| \cdot\left\|y-x_{0}\right\| \leq \lambda \Theta R .
\end{aligned}
$$

Fourth, we prove that for any constrained minimizer of $f$ (denoted $x_{\star}$ ), BROO calls either decrease the distance to $x_{\star}$ or have objective value not much worse than $x_{\star}$.

Lemma 14 Let $f$ be a convex function, $\epsilon, R \geq 0, x_{0} \in \mathbb{R}^{d}, x_{\star} \in \mathbb{B}_{R}\left(x_{0}\right)$, $y \in \mathbb{B}_{R}\left(x_{\star}\right)$, and $x^{\prime}=\mathcal{O}_{\lambda, \delta}(y)$ for $\delta \leq \frac{\epsilon}{4 \lambda R}$ and $\lambda \geq \frac{\epsilon}{3 R^{2}}$. If $f\left(x^{\prime}\right)-f\left(x_{\star}\right)>\frac{\epsilon}{2}$ and $\operatorname{prox}_{\lambda}^{f}(y)=\operatorname{bprox}_{\lambda, r}^{f}(y)$ then $\left\|x^{\prime}-x_{\star}\right\|<\left\|y-x_{\star}\right\|$.

Proof Let $x_{\lambda}=\operatorname{prox}_{\lambda}^{f}(y)$. Since $\operatorname{prox}_{\lambda}^{f}(y)=\operatorname{bprox}_{\lambda, r}^{f}(y)$, Definition 1 of $\mathcal{O}_{\lambda, \delta}(\cdot)$ implies

$$
f\left(x^{\prime}\right) \leq f\left(x^{\prime}\right)+\frac{\lambda}{2}\left\|x^{\prime}-y\right\|^{2} \leq f\left(x_{\lambda}\right)+\frac{\lambda}{2}\left\|x_{\lambda}-y\right\|^{2}+\frac{\lambda \delta^{2}}{2} .
$$

Further, Equation (17) implies that

$$
f\left(x_{\lambda}\right) \leq f\left(x_{\star}\right)+\frac{\lambda}{2}\left\|x_{\star}-y\right\|^{2}-\frac{\lambda}{2}\left\|x_{\star}-x_{\lambda}\right\|^{2}-\frac{\lambda}{2}\left\|x_{\lambda}-y\right\|^{2} .
$$

Combining these inequalities, rearranging, and using $f\left(x^{\prime}\right)-f\left(x_{\star}\right)>\frac{\epsilon}{2}$ yields

$$
\begin{equation*}
\frac{\lambda}{2}\left\|x_{\lambda}-x_{\star}\right\|^{2} \leq \frac{\lambda}{2}\left\|y-x_{\star}\right\|^{2}-\frac{\epsilon}{2}+\frac{\lambda \delta^{2}}{2} \tag{18}
\end{equation*}
$$

Since $\lambda \delta^{2} \leq \frac{\epsilon^{2}}{12 \lambda R^{2}} \leq \frac{\epsilon}{6}<\frac{\epsilon}{2}$, this implies $\left\|x_{\lambda}-x_{\star}\right\| \leq\left\|y-x_{\star}\right\|$.
By the triangle inequality, we have $\left\|x^{\prime}-x_{\star}\right\| \leq\left\|x_{\lambda}-x_{\star}\right\|+\left\|x^{\prime}-x_{\lambda}\right\|$. Thus,

$$
\begin{aligned}
\left\|x^{\prime}-x_{\star}\right\|^{2} & \leq\left\|x_{\lambda}-x_{\star}\right\|^{2}+2\left\|x_{\lambda}-x_{\star}\right\| \cdot\left\|x^{\prime}-x_{\lambda}\right\|+\left\|x^{\prime}-x_{\lambda}\right\|^{2} \\
& \leq\left\|x_{\lambda}-x_{\star}\right\|^{2}+2 \delta\left\|y-x_{\star}\right\|+\delta^{2} \leq\left\|x_{\lambda}-x_{\star}\right\|^{2}+2 \delta R+\delta^{2}
\end{aligned}
$$

where we have used $\left\|x_{\lambda}-x_{\star}\right\| \leq\left\|y-x_{\star}\right\|$ (argued above) along with $\left\|x^{\prime}-\hat{x}\right\| \leq \delta$ (Lemma 11) and the assumption that $y \in \mathbb{B}_{R}\left(x_{\star}\right)$. Substituting into (18), we have

$$
\lambda\left\|x^{\prime}-x_{\star}\right\|^{2} \leq \lambda\left\|y-x_{\star}\right\|^{2}-\epsilon+2 \lambda \delta R+2 \lambda \delta^{2}
$$

To conclude the proof we note that $-\epsilon+2 \lambda \delta R+2 \lambda \delta^{2}<0$ since $\lambda \delta \leq \epsilon /(4 R)$ and $\lambda \delta^{2} \leq \epsilon / 6$.

Fifth, we bound the movement of the proximal operator by the Lipschitz continuity of the objective.

Lemma 15 Let $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ be convex and $L_{f}$-Lipschitz. Then for all $\lambda>0$ and $y \in \mathbb{R}^{d}$,

$$
\left\|\operatorname{prox}_{\lambda}^{f}(y)-y\right\| \leq \frac{L_{f}}{\lambda}
$$

Proof Let $x_{\lambda}:=\operatorname{prox}_{\lambda}^{f}(y)-y$. Equation (15) of Lemma 12 implies that

$$
f(y)-f\left(x_{\lambda}\right) \geq \lambda\left\langle x_{0}-x_{\lambda}, x_{0}-x_{\lambda}\right\rangle=\lambda\left\|x_{0}-x_{\lambda}\right\|^{2}
$$

Further, $L_{f}$-Lipschtiz continuity of $f$ implies $f(y) \leq f\left(x_{\lambda}\right)+L_{f}\left\|x_{0}-x_{\lambda}\right\|^{2}$. Combining and noting that the claim is trivial when $\left\|x_{0}-x_{\lambda}\right\|=0$ yields the claim.

Finally, we mention a standard lemma about the relation between the sequences $\left\{A_{t}\right\}$ and $\left\{\lambda_{t}\right\}$ in accelerated proximal methods.

Lemma 16 ([cf. 12, Lemma 23]) For any iteration $t$ of Algorithm 1, we have $A_{t}=a_{t}^{2} \lambda_{t}$ and

$$
\sqrt{A_{t}} \geq \frac{1}{2} \sum_{i \in[t]} \frac{1}{\sqrt{\lambda_{i}}}
$$

## B.2. Main algorithm analysis

In this section we give the analysis of our accelerated algorithm. Before going into the technical details, let us provide a brief overview of the algorithm and its analysis. At its core, our algorithm is an accelerated proximal point method [19, 25, 16]. These methods iteratively compute proximal points of the form $x_{t+1} \approx \operatorname{prox}_{\lambda_{t+1}}^{f}\left(y_{t}\right)$ and then use a momentum-like extrapolation scheme to compute $y_{t+1}$. Accelerated proximal points methods differ in the methods they employ to (approximately) compute that proximal points and the choices of $\left\{\lambda_{t}\right\}$. Our method approximates
$x_{t+1} \approx \operatorname{prox}_{\lambda_{t+1}}^{f}\left(y_{t}\right)$ using calls to a BROO, and employs a bisection procedure that finds values of $\lambda_{t}$ for which such approximation is valid because the ball constraint is inactive, i.e., when $\left\|x_{t+1}-y_{t}\right\|<r$.

The crux of the analysis of our method is showing that the optimization error decreases roughly as $\exp \left(-\Omega(1) \sum_{i \in[t]}\left(\left\|x_{t+1}-y_{t}\right\| / R\right)^{2 / 3}\right)$. By requiring our bisection procedure to find values for which $\left\|x_{t+1}-y_{t}\right\| \in[r / 2, r)$, we obtain the claimed $\widetilde{O}\left((R / r)^{2 / 3}\right)$ complexity bound. Our analysis of our method closely follows the previous ball-oracle acceleration proof of [12], which itself draws from prior analyses of Monteiro-Svaiter-type algorithms [17, 7].

The differences between our algorithm and proof and those in [12] center around handling non-smoothness and less accurate ball oracle outputs. In particular, in Line 7 of Algorithm 1 we estimate $\nabla f\left(\operatorname{prox}_{\lambda_{t+1}}^{f}\left(y_{t}\right)\right)$ as $\lambda_{t+1}\left(y_{t}-x_{t+1}\right)$, which allows us to avoid smoothness assumptions but requires a somewhat different proof of the main potential bound. We also remove the assumption that $x_{0}$ is within distance $R$ from a global minimizer of $f$, and instead compare the function value of our output to the minimizer of $f$ in a ball of radius $R$. Our bisection subroutine and its analysis also differ from its counterpart in [12]; we explain these differences in the next subsection.

Our analysis of the Algorithm 1 outer loop relies on the following guarantee for our bisection subroutine, which we prove in Appendix B.3.

Proposition 17 (Bisection) Let $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ be $L_{f}$-Lipschitz and convex, and let $x, v \in \mathbb{R}^{d}$, $\epsilon, r, R \in \mathbb{R}_{>0}$ satisfy $\epsilon \leq L_{f} R, r \leq R$ and $\|x-v\| \leq 2 R$. Given $\lambda_{\max } \geq \frac{2 L_{f}}{r}$ and $\lambda_{\min } \in$ $\left(0, \lambda_{\max }\right), \lambda$-Bisection $(x, v, A)$ outputs $\lambda \in\left[\lambda_{\min }, \lambda_{\max }\right]$ such that

$$
\operatorname{prox}_{\lambda}^{f}\left(y_{\lambda}\right)=\operatorname{bprox}_{\lambda, r}^{f}\left(y_{\lambda}\right)
$$

(i.e., $\left.\left\|\operatorname{prox}_{\lambda}^{f}\left(y_{\lambda}\right)-y_{\lambda}\right\| \leq r\right)$. The subroutine uses $O\left(\log \left(\frac{\lambda_{\max }}{\lambda_{\min }}\right)+\log \left(\frac{R+L_{f} / \lambda_{\min }}{r}\right)\right)$ calls to $\mathcal{O}_{\lambda^{\prime}, \frac{r}{17}}(\cdot)$ with $\lambda^{\prime} \in\left[\frac{1}{2} \lambda, \lambda_{\max }\right]$. Moreover, for $\alpha_{2 \lambda A}=\frac{2 \lambda A}{1+2 \lambda A+\sqrt{1+4 \lambda A}}$ and $y_{\lambda}:=\alpha_{2 \lambda A} x+\left(1-\alpha_{2 \lambda A}\right) v$ one of the following outcomes must occur:
(a) $\lambda \in\left[2 \lambda_{\min }, \lambda_{\max }\right]$ and $\left\|\operatorname{prox}_{\lambda}^{f}\left(y_{\lambda}\right)-y_{\lambda}\right\|>\frac{3 r}{4}$, or
(b) $\lambda<2 \lambda_{\text {min }}$.

When taking $\lambda_{\max }=\frac{2 L_{f}}{\epsilon}, \lambda_{\min }=\frac{\epsilon}{6 r R}$, the number of calls to $\mathcal{O}_{\lambda^{\prime}, \frac{r}{17}}(\cdot)$ is bounded by $O\left(\log \frac{L_{f} R^{2}}{r \epsilon}\right)$.
For the remainder of this section, we fix a parameter $R$ and let

$$
x_{\star} \in \underset{x \in \mathbb{B}_{R}\left(x_{0}\right)}{\arg \min } f(x)
$$

denote a minimizer of $f$ in a ball of radius $R$ around the initial point $x_{0}$. (If it is not unique, we choose one arbitrarily). Based on the iterates $\left\{x_{t}, v_{t}, A_{t}\right\}$ generated by Algorithm 1, we define the following quantities:

$$
\begin{equation*}
E_{t}:=f\left(x_{t}\right)-f\left(x_{\star}\right), \hat{E}_{t}=E_{t}-\frac{\epsilon}{4}, D_{t}=\frac{1}{2}\left\|v_{t}-x_{\star}\right\|^{2}, \text { and } P_{t}=A_{t} \hat{E}_{t}+D_{t} \tag{19}
\end{equation*}
$$

In the following lemma we prove our main potential decrease bound, under the conditions that the iterates $x_{t}, v_{t}$ are within distance $2 R$ of each other and (implicitly) that a suitably good solution has yet to be found. We establish these conditions inductively in subsequent lemmas and leverage this to lower bound the growth of $A_{t}$ and prove Theorem 2.

Lemma 18 Let $f$ be a convex function, $x_{0} \in \mathbb{R}^{d}$ and $\epsilon, R>0$. If at iteration $t$ of Algorithm 1 the following conditions hold,
(a) $\left\|x_{t}-v_{t}\right\| \leq 2 R$
(b) $\lambda_{t+1} \geq \frac{\epsilon}{3 r R}$
we have

$$
P_{t+1}-P_{t} \leq-\frac{A_{t+1} \lambda_{t+1} r^{2}}{12}
$$

Proof Let

$$
\hat{x}_{t+1}:=\operatorname{bprox}_{r, \lambda_{t+1}}^{f}\left(y_{t}\right)=\underset{x \in \mathbb{B}_{r}\left(y_{t}\right)}{\arg \min }\left\{f(x)+\frac{\lambda_{t+1}}{2}\left\|x-y_{t}\right\|^{2}\right\} .
$$

Condition (a) and Proposition 17 guarantee that the ball constraint is inactive, i.e.,

$$
\hat{x}_{t+1}=\operatorname{bprox}_{r, \lambda_{t+1}}^{f}\left(y_{t}\right)=\operatorname{prox}_{\lambda_{t+1}}^{f}\left(y_{t}\right)=\underset{x \in \mathbb{R}^{d}}{\arg \min }\left\{f(x)+\frac{\lambda_{t+1}}{2}\left\|x-y_{t}\right\|^{2}\right\} .
$$

(Note that $y_{t}$ is precisely $y_{\lambda}$ defined in Proposition 17). Consequently, by Lemma 12 we have that $g_{t+1}:=\lambda_{t+1}\left(y_{t}-\hat{x}_{t+1}\right) \in \partial f\left(\hat{x}_{t+1}\right)$, i.e.,

$$
\begin{equation*}
f(u) \geq f\left(\hat{x}_{t+1}\right)+g_{t+1}^{\top}\left(u-\hat{x}_{t+1}\right) \text { for all } u \in \mathbb{R}^{d} . \tag{20}
\end{equation*}
$$

Further, since $v_{t+1}=\arg \min _{v \in \mathbb{B}_{R}\left(x_{0}\right)}\left\{a_{t+1}\left\langle y_{t}-x_{t+1}, v\right\rangle+\frac{1}{2}\left\|v-v_{t}\right\|^{2}\right\}$ and $x_{\star} \in \mathbb{B}_{R}\left(x_{0}\right)$ applying Lemma 12 again (with $u=x_{\star} \in \mathbb{B}_{R}\left(x_{0}\right)$ ) yields that

$$
\begin{align*}
a_{t+1} \lambda_{t+1}\left\langle y_{t}-x_{t+1}, v_{t+1}-x_{\star}\right\rangle & \leq \frac{1}{2}\left\|v_{t}-x_{\star}\right\|^{2}-\frac{1}{2}\left\|v_{t+1}-x_{\star}\right\|^{2}-\frac{1}{2}\left\|v_{t+1}-v_{t}\right\|^{2} \\
& =D_{t}-D_{t+1}-\frac{1}{2}\left\|v_{t+1}-v_{t}\right\|^{2} . \tag{21}
\end{align*}
$$

Our proof strategy is to upper and lower bound the inner product $\left\langle g_{t+1}, v_{t+1}-x_{\star}\right\rangle$. In particular, we will lower bound this inner product using (20) and upper bound it using (21). Towards this end, we define the point

$$
\tilde{y}_{t}:=\frac{A_{t}}{A_{t+1}} x_{t}+\frac{a_{t+1}}{A_{t+1}} v_{t+1} .
$$

We remark that the use of $\tilde{y}_{t}$ is inspired from the acceleration analysis of Allen-Zhu and Orecchia [2]. From the definition of $y_{t}$, we obtain

$$
\begin{equation*}
v_{t}=\frac{1}{a_{t+1}}\left(A_{t+1} y_{t}-A_{t} x_{t}\right) \quad \text { and } \quad v_{t}-v_{t+1}=\frac{A_{t+1}}{a_{t+1}}\left(y_{t}-\tilde{y}_{t}\right) . \tag{22}
\end{equation*}
$$

Recalling that $A_{t+1}=A_{t}+a_{t+1}$, we have

$$
\begin{equation*}
v_{t+1}=\frac{1}{a_{t+1}}\left(A_{t+1} \tilde{y}_{t}-A_{t} x_{t}\right)=\hat{x}_{t+1}+\frac{A_{t}}{a_{t+1}}\left(\hat{x}_{t+1}-x_{t}\right)-\frac{A_{t+1}}{a_{t+1}}\left(\hat{x}_{t+1}-\tilde{y}_{t}\right) . \tag{23}
\end{equation*}
$$

To begin our inner product lower bound, we note that

$$
\begin{align*}
\left\langle g_{t+1}, v_{t+1}-x_{\star}\right\rangle & =\left\langle g_{t+1}, \hat{x}_{t+1}-x_{\star}\right\rangle+\frac{A_{t}}{a_{t+1}}\left\langle g_{t+1}, \hat{x}_{t+1}-x_{t}\right\rangle-\frac{A_{t+1}}{a_{t+1}}\left\langle g_{t+1}, \hat{x}_{t+1}-\tilde{y}_{t}\right\rangle \\
& \geq f\left(\hat{x}_{t+1}\right)-f\left(x_{\star}\right)+\frac{A_{t}}{a_{t+1}}\left[f\left(\hat{x}_{t+1}\right)-f\left(x_{t}\right)\right]-\frac{A_{t+1}}{a_{t+1}}\left\langle g_{t+1}, \hat{x}_{t+1}-\tilde{y}_{t}\right\rangle \\
& =\frac{A_{t+1}}{a_{t+1}}\left[f\left(\hat{x}_{t+1}\right)-f\left(x_{\star}\right)\right]-\frac{A_{t}}{a_{t+1}}\left[f\left(x_{t}\right)-f\left(x_{\star}\right)\right]-\frac{A_{t+1}}{a_{t+1}}\left\langle g_{t+1}, \hat{x}_{t+1}-\tilde{y}_{t}\right\rangle \tag{24}
\end{align*}
$$

where the inequality follows from (20). To relate $f\left(\hat{x}_{t+1}\right)$ to $f\left(x_{t+1}\right)$ in (24), we use the approximation guarantee (3) defining $\mathcal{O}_{\lambda, \delta}(\cdot)$ to obtain

$$
\begin{aligned}
f\left(\hat{x}_{t+1}\right) & \geq f\left(x_{t+1}\right)+\frac{\lambda_{t+1}}{2}\left\|x_{t+1}-y_{t}\right\|^{2}-\frac{\lambda_{t+1}}{2}\left\|\hat{x}_{t+1}-y_{t}\right\|^{2}-\frac{\lambda_{t+1} \delta_{t+1}^{2}}{2} \\
& \geq f\left(x_{t+1}\right)+\frac{\lambda_{t+1}}{2}\left\|x_{t+1}-y_{t}\right\|^{2}+\left\langle g_{t+1}, \hat{x}_{t+1}-\tilde{y}_{t}\right\rangle-\frac{\lambda_{t+1}}{2}\left\|y_{t}-\tilde{y}_{t}\right\|^{2}-\frac{\lambda_{t+1} \delta_{t+1}^{2}}{2}
\end{aligned}
$$

where we used $\left\langle g_{t+1}, \hat{x}_{t+1}-\tilde{y}_{t}\right\rangle=\frac{\lambda_{t+1}}{2}\left\|y_{t}-\tilde{y}_{t}\right\|^{2}-\frac{\lambda_{t+1}}{2}\left\|\hat{x}_{t+1}-\tilde{y}_{t}\right\|^{2}-\frac{\lambda_{t+1}}{2}\left\|\hat{x}_{t+1}-y_{t}\right\|^{2}$ (as in Lemma 12) to obtain the second inequality. Substituting into (24) and recalling that $\mathbb{E}_{t}=f\left(x_{t}\right)-$ $f\left(x_{\star}\right)$ yields

$$
\left\langle g_{t+1}, v_{t+1}-x_{\star}\right\rangle \geq \frac{A_{t+1}}{a_{t+1}} E_{t+1}-\frac{A_{t}}{a_{t+1}} E_{t}+\frac{A_{t+1} \lambda_{t+1}}{2 a_{t+1}}\left[\left\|x_{t+1}-y_{t}\right\|^{2}-\delta_{t+1}^{2}-\left\|y_{t}-\tilde{y}_{t}\right\|^{2}\right]
$$

The lower bound $\lambda_{t+1} \geq \frac{\epsilon}{3 r R}$ implies that $\delta_{t+1}=\frac{\epsilon}{12 \lambda_{t+1} R} \leq \frac{r}{4}$. Moreover, by condition (b) $\left(\lambda_{t+1} \geq \frac{\epsilon}{3 r R}\right.$ ) and Proposition 17 we have that $\left\|\hat{x}_{t+1}-y_{t}\right\| \geq 3 r / 4$. Applying Lemma 11 we conclude that

$$
\left\|x_{t+1}-y_{t}\right\|^{2}-\delta_{t+1}^{2} \geq\left(\left\|\hat{x}_{t+1}-y_{t}\right\|-\delta_{t+1}\right)^{2}-\delta_{t+1}^{2} \geq r^{2}\left[\left(\frac{3}{4}-\frac{1}{4}\right)^{2}-\left(\frac{1}{4}\right)^{2}\right] \geq \frac{r^{2}}{6}
$$

Substituting back, we have

$$
\begin{equation*}
\left\langle g_{t+1}, v_{t+1}-x_{\star}\right\rangle \geq \frac{A_{t+1}}{a_{t+1}} E_{t+1}+\frac{A_{t}}{a_{t+1}} E_{t}-\frac{A_{t+1} \lambda_{t+1} r^{2}}{12 a_{t+1}}-\frac{A_{t+1} \lambda_{t+1}}{2 a_{t+1}}\left\|y_{t}-\tilde{y}_{t}\right\|^{2} \tag{25}
\end{equation*}
$$

We now proceed to upper bound $\left\langle g_{t+1}, v_{t+1}-x_{\star}\right\rangle$. Recalling $g_{t+1}=\lambda_{t+1}\left(y_{t}-\hat{x}_{t+1}\right)$ we obtain

$$
\begin{align*}
\left\langle g_{t+1}, v_{t+1}-x_{\star}\right\rangle & =\lambda_{t+1}\left\langle y_{t}-x_{t+1}, v_{t+1}-x_{\star}\right\rangle+\lambda_{t+1}\left\langle x_{t+1}-\hat{x}_{t+1}, v_{t+1}-x_{\star}\right\rangle \\
& \leq \lambda_{t+1}\left\langle y_{t}-x_{t+1}, v_{t+1}-x_{\star}\right\rangle+\lambda_{t+1}\left\|x_{t+1}-\hat{x}_{t+1}\right\|\left\|v_{t+1}-x_{\star}\right\| \tag{26}
\end{align*}
$$

To bound the term $\lambda_{t+1}\left\|x_{t+1}-\hat{x}_{t+1}\right\|\left\|v_{t+1}-x_{\star}\right\|$, note that $\left\|x_{t+1}-\hat{x}_{t+1}\right\| \leq \delta_{t+1}$ by Lemma 11, $\left\|v_{t+1}-x_{\star}\right\| \leq 2 R$ since $v_{t+1}$ and $x_{\star}$ are both in $\mathbb{B}_{R}\left(x_{0}\right)$ by assumption, and $\lambda_{t+1} \delta_{t+1} \cdot 2 R=\frac{\epsilon}{6} \leq \frac{\epsilon}{4}$. To bound the term $\lambda_{t+1}\left\langle y_{t}-x_{t+1}, v_{t+1}-x_{\star}\right\rangle$ we apply (21). Applying these bounds to (26) yields

$$
\begin{align*}
\left\langle g_{t+1}, v_{t+1}-x_{\star}\right\rangle & \leq \frac{1}{a_{t+1}}\left[D_{t}-D_{t+1}-\frac{1}{2}\left\|v_{t+1}-v_{t}\right\|^{2}+\frac{a_{t+1} \epsilon}{4}\right] \\
& =\frac{1}{a_{t+1}}\left[D_{t}-D_{t+1}-\frac{A_{t+1}^{2}}{2 a_{t+1}^{2}}\left\|y_{t}-\tilde{y}_{t}\right\|^{2}+\left(A_{t+1}-A_{t}\right) \frac{\epsilon}{4}\right] \tag{27}
\end{align*}
$$

where the equality is due to (22) and $a_{t+1}=A_{t+1}-A_{t}$.
Combining the lower and upper bound (25) and (27) and rearranging, we have
$A_{t+1}\left(E_{t+1}-\frac{\epsilon}{4}\right)+D_{t+1}-A_{t}\left(E_{t+1}-\frac{\epsilon}{4}\right)-D_{t} \leq-\frac{A_{t+1} \lambda_{t+1} r^{2}}{12}+\left(A_{t+1} \lambda_{t+1}-\frac{A_{t+1}^{2}}{a_{t+1}^{2}}\right) \frac{1}{2}\left\|y_{t}-\tilde{y}_{t}\right\|^{2}$.
The proof is complete upon noticing that $A_{t+1} \lambda_{t+1}-\frac{A_{t+1}^{2}}{a_{t+1}^{2}}=0$ (since $A_{t+1}=a_{t+1}^{2} \lambda_{t+1}$ ) and that $P_{t}=A_{t} \hat{E}_{t}+D_{t}=A_{t}\left(E_{t}-\frac{\epsilon}{4}\right)+D_{t}$.

Lemma 18 shows that the potential $P_{t}$ decreases significantly whenever $\lambda_{t+1}$ is not too small and $\left\|x_{t}-v_{t}\right\| \leq 2 R$ holds; we now the latter condition inductively.

Lemma 19 Fix $t \geq 1$. In Algorithm 1 if $f\left(x_{i}\right)-f\left(x_{\star}\right)>\epsilon$ and $\lambda_{i} \geq \frac{\epsilon}{3 r R}$ for all $0 \leq i \leq t$ then

$$
\left\|x_{t}-x_{\star}\right\| \leq R, \quad\left\|v_{t}-x_{\star}\right\| \leq R, \text { and } P_{t}-P_{0} \leq-\sum_{i \in[t-1]} \frac{A_{i+1} \lambda_{i+1} r^{2}}{12}
$$

Proof We proceed by induction on $t$. For the base case of $t=0$, we note that $\left\|v_{0}-x_{\star}\right\|=$ $\left\|x_{0}-x_{\star}\right\| \leq R$ by assumption, and that $P_{0}-P_{0} \leq 0$ trivially. Therefore we assume the inductive hypothesis that the lemma statement holds for iteration $t$, and show that it also holds for $t+1$. First, we note that $\left\|x_{t}-x_{\star}\right\| \leq R,\left\|v_{t}-x_{\star}\right\| \leq R$ and $\lambda_{t+1} \geq \frac{\epsilon}{3 r R}$ satisfy the conditions of Lemma 18 and consequently $P_{t+1} \leq P_{t}-\frac{1}{12} A_{t+1} \lambda_{t+1} r^{2}$; together with the inductive hypothesis this establishes

$$
P_{t+1}-P_{0} \leq-\sum_{i=0}^{t} \frac{A_{i+1} r^{2} \lambda_{i+1}}{12} .
$$

Next, we note that $f\left(x_{t+1}\right)-f\left(x_{\star}\right)>\epsilon$ implies that $\hat{E}_{t+1}=f\left(x_{t+1}\right)-f\left(x_{\star}\right)-\frac{\epsilon}{4}>0$. Recalling the definition (19) and $A_{0}=0$, this implies

$$
\frac{1}{2}\left\|v_{t}-x_{\star}\right\|^{2}=D_{t} \leq A_{t} \hat{E}_{t}+D_{t}=P_{t} \leq P_{0}=D_{0}=\frac{1}{2}\left\|x_{0}-x_{\star}\right\|^{2} \leq \frac{1}{2} R^{2}
$$

and consequently

$$
\left\|v_{t+1}-x_{\star}\right\| \leq R
$$

To complete the induction step we need to argue that $\left\|x_{t+1}-x_{\star}\right\| \leq R$. To that end, we invoke Lemma 14 (recalling that $f\left(x_{t+1}\right)-f\left(x_{\star}\right)>\epsilon$ by assumption) which shows that $\left\|x_{t+1}-x_{\star}\right\| \leq$ $\left\|y_{t}-x_{\star}\right\|$ or $f\left(x_{t+1}\right)-f\left(x_{\star}\right) \leq \frac{\epsilon}{2}$. The definition $y_{t}=\frac{A_{t}}{A_{t+1}} x_{t}+\frac{a_{t+1}}{A_{t+1}} v_{t}$ gives

$$
\left\|x_{t+1}-x_{\star}\right\| \leq\left\|y_{t}-x_{\star}\right\| \leq \frac{A_{t}}{A_{t+1}}\left\|x_{t}-x_{\star}\right\|+\frac{a_{t+1}}{A_{t+1}}\left\|v_{t}-x_{\star}\right\| \leq R
$$

where the final bounds holds since $A_{t+1}=A_{t}+a_{t+1}$ and $\left\|x_{t}-x_{\star}\right\|,\left\|v_{t}-x_{\star}\right\| \leq R$ by the inductive assumption.

Lemma 20 Fix $t \geq$ 1. In Algorithm 1 if $f\left(x_{i}\right)-f\left(x_{\star}\right)>\epsilon$ and $\lambda_{i} \geq \frac{\epsilon}{3 r R}$ for all $0 \leq i \leq t$ then

$$
A_{t} \geq \exp \left(\frac{r^{2 / 3}}{R^{2 / 3}}(t-1)\right) A_{1} .
$$

Proof Lemma 19 implies that

$$
P_{t}-P_{0} \leq-\sum_{i=0}^{t} \frac{A_{i+1} r^{2} \lambda_{i+1}}{12} .
$$

Further, that $f\left(x_{i}\right)-f\left(x_{\star}\right) \geq \epsilon$ for all $i \leq t$ implies $\hat{E}_{t} \geq 0$ and therefore $P_{t} \geq 0$. Combining these facts with the facts that $A_{0}=0$ and the $A_{i}$ increase monotonically yields

$$
\begin{equation*}
\sum_{i=0}^{t-1} A_{i+1} \lambda_{i+1} \leq \frac{12}{r^{2}} P_{0}=\frac{12}{r^{2}} D_{0} \leq \frac{6 R^{2}}{r^{2}} \tag{28}
\end{equation*}
$$

Next, note that the reverse Hölder inequality with $p=2 / 3$ states that for any $u, v \in \mathbb{R}_{>0}^{d}$

$$
\langle u, v\rangle \geq\left(\sum_{i \in[d]} u_{i}^{2 / 3}\right)^{3 / 2} \cdot\left(\sum_{i \in[d]} v_{i}^{-2}\right)^{-1 / 2} .
$$

We therefore have

$$
\begin{aligned}
& \sqrt{A_{t}} \stackrel{(i)}{\geq} \frac{1}{2} \sum_{i \in[t]} \frac{1}{\sqrt{\lambda_{i}}} \stackrel{(i i)}{\geq} \frac{1}{2}\left(\sum_{i \in[t]}\left(\sqrt{A_{i}}\right)^{2 / 3}\right)^{3 / 2} \cdot\left(\sum_{i \in[t]}\left(\frac{1}{\sqrt{A_{i} \lambda_{i}}}\right)^{-2}\right)^{-1 / 2} \\
& \stackrel{(i i i)}{\geq} \frac{1}{2}\left(\sum_{i \in[t]}\left(\sqrt{A_{i}}\right)^{2 / 3}\right)^{3 / 2} \cdot \frac{r}{\sqrt{6} R}
\end{aligned}
$$

where we used $(i)$ Lemma 16 , $(i i)$ the reverse Hölder inequality with $u_{i}=\sqrt{A_{i}}$ and $v_{i}=1 / \sqrt{A_{i} \lambda_{i}}$, and (iii) the bound (28). Rearranging, we have

$$
\begin{equation*}
A_{t}^{1 / 3} \geq \frac{r^{2 / 3}}{3 R^{2 / 3}}\left(\sum_{i \in[t]} A_{i}^{1 / 3}\right) \tag{29}
\end{equation*}
$$

Lemma 28 of [12] shows that for any sequence of $A_{t}$ that satisfy (29) for all $t \in[T]$ also satisfies

$$
A_{T}^{1 / 3} \geq \exp \left(\frac{r^{2 / 3}}{3 R^{2 / 3}}(T-1)\right) A_{1}^{1 / 3}
$$

and the result follows.
We are now ready to prove our main theorem.
Proof [Proof of Theorem 2] This proof proceeds in parts. First, we show that whenever the algorithm terminates, i.e. one of the conditions of Line 8 is met, then Algorithm 1 outputs a point $x_{\text {ret }}$ with $f\left(x_{\mathrm{ret}}\right)-f\left(x_{\star}\right) \leq \epsilon$ on Line 9 . Next, we bound $T_{o}$, the value of $t$ at termination on Line 9 , and show that $T_{o}=O\left(\left(\frac{R}{r}\right)^{2 / 3} \log \left(\frac{\left[f\left(x_{0}-f\left(x_{\star}\right)\right)\right] \cdot R^{2}}{\epsilon \cdot r^{2}}\right)\right)$. Leveraging these facts we complete the proof.

Termination due to small $\lambda_{T_{o}+1}$. Consider the case where the algorithm terminates because $\lambda_{T_{o}+1}<\frac{\epsilon}{3 r R}$ on Line 8. Further, suppose that $f\left(x_{i}\right)-f\left(x_{\star}\right)>\epsilon$ for all $0 \leq i \leq T_{o}$ as otherwise $f\left(x_{\text {ret }}\right)-f\left(x_{\star}\right) \leq \epsilon$ as desired. By definition of $T_{o}$, we have $\lambda_{t} \geq \frac{\epsilon}{3 r R}$ for all $t \leq T_{o}$ (otherwise we would have terminated earlier). Applying Lemma 19, we conclude it must be that $\left\|x_{T_{o}}-x_{\star}\right\| \leq R$ and $\left\|v_{T_{o}}-x_{\star}\right\| \leq R$ and therefore $\left\|y_{T_{o}}-x_{\star}\right\| \leq R$. By Proposition 17 we have that $\hat{x}_{T_{o}+1}=\operatorname{prox}_{\lambda}^{f}\left(y_{T_{o}}\right)=\operatorname{bprox} \lambda_{\lambda, r}\left(y_{T_{o}}\right)$ and thus $\left\|\hat{x}_{T_{o}+1}-y_{T_{o}}\right\|<r$. Lemma 13 with $x_{0}=y_{T_{o}}$, $x_{\lambda}=\hat{x}_{T_{o}+1}, y=x_{\star}$ and $\Theta=r$ then yields

$$
f\left(\hat{x}_{T_{o}+1}\right)-f\left(x_{\star}\right) \leq \lambda_{t+1} r R<\frac{\epsilon}{3} .
$$

Moreover, the BROO guarantee (3) gives
$f\left(x_{T_{o}+1}\right) \leq f\left(\hat{x}_{T_{o}+1}\right)+\frac{\lambda_{T_{o}+1}}{2}\left[\left\|\hat{x}_{T_{o}+1}-y_{t}\right\|^{2}-\left\|x_{T_{o}+1}-y_{t}\right\|^{2}+\delta_{T_{o}+1}^{2}\right] \leq \frac{\lambda_{T_{o}+1}}{2}\left(r^{2}+\delta_{T_{o}+1}^{2}\right)$.
Noting that $\lambda_{T_{o}+1} \in\left[\frac{\epsilon}{6 r R}, \frac{\epsilon}{3 r R}\right]$ and that $\delta_{T_{o}+1}=\frac{\epsilon}{12 \lambda_{T_{o}+1} R}$, we have have $\lambda_{T_{o}+1} r^{2} \leq \epsilon \cdot \frac{r}{3 R} \leq \epsilon / 3$ and $\lambda_{T_{o}+1} \delta_{T_{o}+1}^{2} \leq \epsilon \cdot \frac{\epsilon}{12 R} \cdot \frac{r}{2} \leq \epsilon / 24$. Substituting back, we have

$$
f\left(x_{T_{o}+1}\right)-f\left(x_{\star}\right) \leq f\left(\hat{x}_{T_{o}+1}\right)-f\left(x_{\star}\right)+\frac{9 \epsilon}{49} \leq \frac{\epsilon}{3}+\frac{9 \epsilon}{49} \leq \epsilon .
$$

Termination due to large $A_{T_{o}+1}$. Next consider the case where $\lambda_{t+1} \geq \frac{\epsilon}{3 r R}$ for all $t \leq T_{o}$ but $A_{T_{o}+1} \geq \frac{R^{2}}{\epsilon}$. In this case, Lemma 19 implies that unless $f\left(x_{\mathrm{ret}}\right)-f\left(x_{\star}\right) \leq \epsilon$

$$
A_{T_{o}+1}\left(f\left(x_{T_{o}+1}\right)-f\left(x_{\star}\right)-\frac{\epsilon}{4}\right)=A_{T_{o}+1} \hat{E}_{T_{o}+1} \leq P_{T_{o}+1} \leq P_{0}=D_{0}=\frac{1}{2}\left\|x_{0}-x_{\star}\right\|^{2} \leq \frac{R^{2}}{2} .
$$

Dividing by $A_{T_{o}+1} \geq \frac{R^{2}}{\epsilon}$, we obtain

$$
f\left(x_{T_{o}+1}\right)-f\left(x_{\star}\right) \leq \frac{\epsilon}{4}+\frac{\epsilon}{2} \leq \epsilon
$$

Termination due to distant $x_{T_{o}+1}, v_{T_{o}+1}$. Next consider the case where $\lambda_{t+1} \geq \frac{\epsilon}{3 r R}$ for all $t \leq T_{o}$ but $\left\|x_{T_{o}+1}-v_{T_{o}+1}\right\|>2 R$. This implies that either $\left\|x_{T_{o}+1}-x_{\star}\right\|>R$ or $\left\|v_{T_{o}+1}-x_{\star}\right\|>R$. Consequently, Lemma 19 implies that $f\left(x_{i}\right)-f\left(x_{\star}\right) \leq \epsilon$ for some $0 \leq i \leq T_{o}+1$.

Termination due to slow $A_{t}$ growth. Next consider the case where $\lambda_{t+1} \geq \frac{\epsilon}{3 r R}$ for all $t \leq T_{o}$ but $A_{T_{o}+1}<\exp \left(\frac{r^{2 / 3}}{R^{2 / 3}}\left(T_{o}-1\right)\right) A_{1}$. In this case Lemma 20 implies that $f\left(x_{i}\right)-f\left(x_{\star}\right) \leq \epsilon$ for some $0 \leq i \leq T_{o}+1$ and therefore $f\left(x_{\mathrm{ret}}\right)-f\left(x_{\star}\right) \leq \epsilon$ as desired.

We conclude that in each of the four possible causes for the algorithm to terminate, i.e. one of the conditions in Line 8 to be true, the algorithm $x_{\text {ret }}$ such that $f\left(x_{\mathrm{ret}}\right)-f\left(x_{\star}\right) \leq \epsilon$ as desired.
Iteration Bound. We now show that $T_{o}=O\left(\frac{R^{2 / 3}}{r^{2 / 3}} \log \left(\frac{\left[f\left(x_{0}\right)-f\left(x_{\star}\right)\right] R}{\epsilon r}\right)\right)$. Note that by definition of $T_{o}$ as the iterate index $t$ for which termination on Line 9 occurs, we have

$$
\exp \left(\frac{r^{2 / 3}}{R^{2 / 3}} \cdot\left(T_{o}-1\right)\right) A_{1} \leq A_{T_{o}} \leq \frac{R^{2}}{\epsilon}
$$

since otherwise we would have terminated on at $t=T_{o}-1$. Consequently we have that

$$
\begin{equation*}
T_{o}=O\left(\frac{R^{2 / 3}}{r^{2 / 3}} \log \left(\frac{R^{2}}{A_{1} \epsilon}\right)\right) \tag{30}
\end{equation*}
$$

and it remains to lower bound $A_{1}=a_{1}=\frac{1}{2 \lambda_{1}}$ ( since $A_{0}=0$ ).
Note that since $x_{0}=v_{0}$ it is the case that $y_{0}=x_{0}$. Consequently, if $T_{o}>0$, Line 8 implies that $\lambda_{1} \geq \frac{\epsilon}{3 r R}$, i.e., item (b) of Proposition 17 does not hold). Consequently, Proposition 17 implies that $\left\|\operatorname{prox}_{\lambda_{1}}^{f}\left(x_{0}\right)-x_{0}\right\| \in\left[\frac{3 r}{4}, r\right]$ and therefore

$$
f\left(x_{0}\right) \geq f\left(\operatorname{prox}_{\lambda_{1}}^{f}\left(x_{0}\right)\right)+\frac{\lambda}{2}\left\|\operatorname{prox}_{\lambda_{1}}^{f}\left(x_{0}\right)-x_{0}\right\|^{2} \geq f\left(\operatorname{prox}_{\lambda_{1}}^{f}\left(x_{0}\right)\right)+\frac{9 \lambda r^{2}}{32} .
$$

Further, since $\left\|\operatorname{prox}_{\lambda_{1}}^{f}\left(x_{0}\right)-x_{0}\right\| \leq r \leq R$ we know that $f\left(\operatorname{prox}_{\lambda_{1}}^{f}\left(x_{0}\right)\right) \geq f\left(x_{\star}\right)$. Consequently,

$$
\lambda_{1} \leq \frac{32\left[f\left(x_{0}\right)-f\left(x_{\star}\right)\right]}{9 r^{2}} \text { and } A_{1}=\frac{1}{\lambda_{1}} \geq \frac{9 r^{2}}{32\left[f\left(x_{0}\right)-f\left(x_{\star}\right)\right]}
$$

Substituting back to (30) yields the claimed bound on $T_{o}$.
Remaining guarantees. To complete the proof, we observe that in each of the algorithm's

$$
O\left(\left(\frac{R}{r}\right)^{2 / 3} \log \left(\frac{L_{f} R^{2}}{r \epsilon}\right)\right)
$$

iterations, Proposition 17 guarantees that we perform at most $O\left(\log \frac{L_{f} R^{2}}{r \epsilon}\right)$ queries to $\mathcal{O}_{\lambda, \delta}(\cdot)$. It also guarantees that the queried $\lambda$ satisfies $\lambda \in\left[\frac{\epsilon}{24 r R}, \frac{4 L_{f}}{r}\right]$, yielding item 1 in Theorem 6. Finally, for the sequence $\lambda_{(1)}, \ldots, \lambda_{(T)}$ of $\lambda$ values queried during the execution of Algorithm 1, we have

$$
\begin{aligned}
\sum_{i \in[T]} \frac{1}{\sqrt{\lambda_{(i)}}}=\sum_{t \in\left[T_{o}+1\right]} \sum_{\substack{\text { seen in line search } \\
\text { for iteration } t}} \frac{1}{\sqrt{\lambda_{(i)}}} & \stackrel{(i)}{\leq} \sum_{i \in\left[T_{o}+1\right]} \frac{O(1)}{\sqrt{\lambda_{t}}} \log \left(\frac{L_{f} R^{2}}{r \epsilon}\right) \\
& \stackrel{(i i)}{\leq} O\left(\sqrt{A_{T_{o}}}+\frac{1}{\sqrt{\lambda_{\min }}}\right) \log \left(\frac{L_{f} R^{2}}{r \epsilon}\right),
\end{aligned}
$$

due to (i) Proposition 17 and (ii) Lemma 16 and $\lambda_{T_{o}+1} \geq \lambda_{\text {min }}=\frac{\epsilon}{6 r R}$ for all $t$. Finally, Line 8 guarantees $A_{T_{o}}<\frac{R^{2}}{\epsilon}$, and (since $\left.r \leq R\right) \sqrt{A_{T_{o}}}+\frac{1}{\sqrt{\lambda_{\text {min }}}}=O\left(\frac{R}{\sqrt{\epsilon}}\right)$, giving item 2 in Theorem 6.

## B.3. Bisection analysis

In this section we prove Proposition 17, a guarantee on our bisection subroutine, restated below.
Proposition 17 (Bisection) Let $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ be $L_{f}$-Lipschitz and convex, and let $x, v \in \mathbb{R}^{d}$, $\epsilon, r, R \in \mathbb{R}_{>0}$ satisfy $\epsilon \leq L_{f} R, r \leq R$ and $\|x-v\| \leq 2 R$. Given $\lambda_{\max } \geq \frac{2 L_{f}}{r}$ and $\lambda_{\min } \in$ $\left(0, \lambda_{\max }\right), \lambda$-Bisection $(x, v, A)$ outputs $\lambda \in\left[\lambda_{\min }, \lambda_{\max }\right]$ such that

$$
\operatorname{prox}_{\lambda}^{f}\left(y_{\lambda}\right)=\operatorname{bprox}_{\lambda, r}^{f}\left(y_{\lambda}\right)
$$

(i.e., $\left.\left\|\operatorname{prox}_{\lambda}^{f}\left(y_{\lambda}\right)-y_{\lambda}\right\| \leq r\right)$. The subroutine uses $O\left(\log \left(\frac{\lambda_{\max }}{\lambda_{\min }}\right)+\log \left(\frac{R+L_{f} / \lambda_{\min }}{r}\right)\right)$ calls to $\mathcal{O}_{\lambda^{\prime}, \frac{r}{17}}(\cdot)$ with $\lambda^{\prime} \in\left[\frac{1}{2} \lambda, \lambda_{\max }\right]$. Moreover, for $\alpha_{2 \lambda A}=\frac{2 \lambda A}{1+2 \lambda A+\sqrt{1+4 \lambda A}}$ and $y_{\lambda}:=\alpha_{2 \lambda A} x+\left(1-\alpha_{2 \lambda A}\right) v$ one of the following outcomes must occur:
(a) $\lambda \in\left[2 \lambda_{\min }, \lambda_{\max }\right]$ and $\left\|\operatorname{prox}_{\lambda}^{f}\left(y_{\lambda}\right)-y_{\lambda}\right\|>\frac{3 r}{4}$, or
(b) $\lambda<2 \lambda_{\text {min }}$.

When taking $\lambda_{\max }=\frac{2 L_{f}}{\epsilon}, \lambda_{\min }=\frac{\epsilon}{6 r R}$, the number of calls to $\mathcal{O}_{\lambda^{\prime}, \frac{r}{17}}(\cdot)$ is bounded by $O\left(\log \frac{L_{f} R^{2}}{r \epsilon}\right)$.

## B.3.1. PRELIMINARIES AND DISCUSSION

Notation. To analyze the bisection procedure, we use the following functions of $\lambda \geq 0$. Fixing $x, v \in \mathbb{R}^{d}$ and $A \geq 0$, we define

$$
y_{\lambda}:=\alpha_{2 A \lambda^{\prime}} \cdot x+\left(1-\alpha_{2 A \lambda^{\prime}}\right) \cdot, \text { where } \alpha_{\tau}:=\frac{\tau}{1+\tau+\sqrt{1+2 \tau}}
$$

We also define

$$
\hat{x}_{\lambda}:=\operatorname{prox}_{\lambda}^{f}\left(y_{\lambda}\right)=\underset{x \in \mathbb{R}^{d}}{\arg \min }\left\{f(x)+\frac{\lambda}{2}\left\|x-y_{\lambda}\right\|^{2}\right\}
$$

and

$$
\begin{equation*}
\hat{\Delta}(\lambda):=\left\|\hat{x}_{\lambda}-y_{\lambda}\right\| \tag{31}
\end{equation*}
$$

We approximate $\hat{x}_{\lambda}$ using the BROO output $\mathcal{O}_{\lambda, \delta}\left(y_{\lambda}\right)$ with $\delta=r / 17$, and write

$$
\Delta(\lambda):=\left\|\mathcal{O}_{\lambda, \delta}\left(y_{\lambda}\right)-y_{\lambda}\right\|
$$

for our approximation to $\hat{\Delta}(\lambda)$. Note that, depending on the BROO implementation, $\Delta(\lambda)$ need not be deterministic. Our bisection procedure implicitly assumes that the BROO is called only once per input $x, v, A$ and distinct value of $\lambda$, and that subsequent references to $\Delta(\lambda)$ use cached values of $\mathcal{O}_{\lambda, \delta}\left(y_{\lambda}\right)$.

Algorithm overview. The goal of $\lambda$-Bisection is to find a value of $\lambda$ where $\hat{\Delta}(\lambda)<r$ so that $\hat{x}_{\lambda}=\operatorname{prox}_{\lambda}^{f}\left(y_{\lambda}\right)=\operatorname{bprox}_{\lambda, r}^{f}\left(y_{\lambda}\right)$ is well-approximated by the BROO output $\mathcal{O}_{\lambda, \delta}\left(y_{\lambda}\right)$. In addition, the bisection has to guarantee that either

- $\hat{\Delta}(\lambda) \geq 3 r / 4$, which implies that the outer loop of Algorithm 1 makes sufficient progress (see Lemma 18), or
- $\lambda<2 \lambda_{\text {min }}$, which implies that $\hat{x}_{\lambda}$ is near-optimal (as we argue in the proof of Theorem 6).

Our procedure (given in Algorithm 1) starts with $\lambda=\lambda_{\text {max }}$ sufficiently large to guarantee $\hat{\Delta}(\lambda)<r$. It then iteratively halves $\lambda$ until finding $\lambda_{0}$ such that $\Delta\left(\lambda_{0}\right)>\frac{13 r}{16}$ or $\lambda_{0}<\lambda_{\min }$. In the latter case it returns $2 \lambda_{0}<2 \lambda_{\min }$ and we note that $\Delta\left(2 \lambda_{0}\right) \leq \frac{13 r}{16}$ implies $\hat{\Delta}\left(2 \lambda_{0}\right)<r$. In the former case we perform a binary search in the interval $\left[\lambda_{0}, 2 \lambda_{0}\right]$ until we find $\lambda_{m}$ such that $\Delta\left(\lambda_{m}\right) \in\left[\frac{13 r}{16}, \frac{15 r}{16}\right]$, which implies $\hat{\Delta}\left(\lambda_{m}\right) \in[3 r / 4, r)$.

Comparison to the bisection in [12]. Our bisection procedure essentially attempts to find $\lambda$ such that $\bar{\Delta}(\lambda)$ is close to (but smaller than) $r$. In contrast, Carmon et al. [12] use the implicit relation $y_{\lambda}-\hat{x}_{\lambda}=\frac{1}{\lambda} \nabla f\left(\hat{x}_{\lambda}\right)$ and attempt to find $\lambda$ such that $\left\|\nabla f\left(\hat{x}_{\lambda}\right)\right\| / \lambda$ is close to $r$. Consequently, the iteration count bounds on the bisection of [12] depend on the continuity $\nabla f$, whereas our bisection succeeds even when $f$ is non-smooth and hence $\nabla f$ is discontinuous. The key to this improvement is a careful analysis of the continuity of $\hat{\Delta}(\lambda)$ (see the following subsection). Another novel aspect of our procedure is the two-stage structure where we first iteratively halve $\lambda$ and only then perform a binary search. This structure allows us to guarantee that we never query the BROO with $\lambda^{\prime}$ that is much smaller than the $\lambda$ we eventually output. This guarantee is necessary for proving statement 2 in Theorem 6, which in turn is necessary for establishing an optimal complexity bound in the weakly-smooth regime $L_{g}=\Theta\left(L_{f}^{2} / \epsilon\right)$.

## B.3.2. Continuity analysis of $\hat{\Delta}(\cdot)$

We begin by proving a bound on the Jacobian of $\delta_{\lambda}=\hat{x}_{\lambda}-y_{\lambda}$ with respect to $\lambda$ (Lemma 21) under the assumption that $f$ is twice differentiable. We subsequently remove this assumption via a smoothing argument (Corollary 22).

Lemma 21 Let $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ be convex and twice differentiable, let $y_{\lambda} \in \mathbb{R}^{d}$ be a differentiable function of $\lambda>0, \hat{x}_{\lambda}=\operatorname{prox}_{\lambda}^{f}\left(y_{\lambda}\right)$, and $\delta_{\lambda}=\hat{x}_{\lambda}-y_{\lambda}$. Then for all $\lambda>0$ we have

$$
\left\|\frac{d}{d \lambda} \delta_{\lambda}\right\| \leq\left\|\frac{d}{d \lambda} y_{\lambda}\right\|+\frac{1}{\lambda}\left\|\delta_{\lambda}\right\| .
$$

Proof Note that $\nabla f\left(\hat{x}_{\lambda}\right)+\lambda\left(\hat{x}_{\lambda}-y_{\lambda}\right)=0$. Differentiating yields

$$
\nabla^{2} f\left(\hat{x}_{\lambda}\right) \cdot \frac{d}{d \lambda} \hat{x}_{\lambda}+\left(\hat{x}_{\lambda}-y_{\lambda}\right)+\lambda\left(\frac{d}{d \lambda} \hat{x}_{\lambda}-\frac{d}{d \lambda} y_{\lambda}\right)=0 .
$$

Rearranging the terms, we obtain

$$
\left(\nabla^{2} f\left(\hat{x}_{\lambda}\right)+\lambda I\right) \cdot\left(\frac{d}{d \lambda} \hat{x}_{\lambda}-\frac{d}{d \lambda} y_{\lambda}\right)=y_{\lambda}-\hat{x}_{\lambda}-\nabla^{2} f\left(\hat{x}_{\lambda}\right) \frac{d}{d \lambda} y_{\lambda} .
$$

Since $f$ is convex, $\nabla^{2} f\left(\hat{x}_{\lambda}\right)$ is PSD and therefore

$$
\left(\frac{d}{d \lambda} \hat{x}_{\lambda}-\frac{d}{d \lambda} y_{\lambda}\right)=\left(\nabla^{2} f\left(\hat{x}_{\lambda}\right)+\lambda I\right)^{-1} \cdot\left(y_{\lambda}-\hat{x}_{\lambda}-\nabla^{2} f\left(\hat{x}_{\lambda}\right) \cdot \frac{d}{d \lambda} y_{\lambda}\right) .
$$

Note that $\left(\nabla^{2} f\left(\hat{x}_{\lambda}\right)+\lambda I\right)^{-1} \nabla^{2} f\left(\hat{x}_{\lambda}\right)$ is a symmetric PSD matrix with all eigenvalues $\in[0,1]$ and $\left(\nabla^{2} f\left(\hat{x}_{\lambda}\right)+\lambda I\right)^{-1}$ is a symmetric PSD matrix with all eigenvalues at most $1 / \lambda$. Consequently, $\left\|\left(\nabla^{2} f\left(\hat{x}_{\lambda}\right)+\lambda I\right)^{-1} \nabla^{2} f(x)\right\| \leq 1$ and $\left\|\left(\nabla^{2} f\left(\hat{x}_{\lambda}\right)+\lambda I\right)^{-1}\right\| \leq \frac{1}{\lambda}$ yielding the claim.

We now relax the assumption of twice differentiability.
Corollary 22 Let $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ be $L_{f}$-Lipschitz and convex, let $y_{\lambda} \in \mathbb{R}^{d}$ be a differentiable function of $\lambda>0, \hat{x}_{\lambda}=\operatorname{prox}_{\lambda}^{f}\left(y_{\lambda}\right)$, and $\delta_{\lambda}=\hat{x}_{\lambda}-y_{\lambda}$. Then for all $\lambda_{1}, \lambda_{2}>0$ we have

$$
\left\|\delta_{\lambda_{2}}-\delta_{\lambda_{1}}\right\| \leq \int_{\lambda=\lambda_{1}}^{\lambda_{2}}\left\|\frac{d}{d \lambda} y_{\lambda}\right\| d \lambda+\int_{\lambda=\lambda_{1}}^{\lambda_{2}} \frac{1}{\lambda}\left\|\delta_{\lambda}\right\| d \lambda .
$$

Proof For $\sigma>0$, let $\nu=\mathcal{N}\left(0 ; \sigma^{2} I_{d \times d}\right)$ and let $f_{\sigma}(x)=\mathbb{E} f(x+\nu)$. Then $f_{\sigma}$ is convex, infinitely differentiable and satisfies $0 \leq f_{\sigma}(x)-f(x) \leq L_{f} \mathbb{E}\|\nu\| \leq L_{f} \sqrt{d} \sigma$ for all $x \in \mathbb{R}^{d}$. Thus, applying Lemma 21 for $f_{\sigma}$ and $\delta_{\lambda}^{\sigma}=\operatorname{prox}_{\lambda}^{f_{\sigma}}\left(y_{\lambda}\right)-y_{\lambda}$, we have

$$
\left\|\delta_{\lambda_{2}}^{\sigma}-\delta_{\lambda_{1}}^{\sigma}\right\| \leq \int_{\lambda=\lambda_{1}}^{\lambda_{2}}\left\|\frac{d}{d \lambda} y_{\lambda}\right\| d \lambda+\int_{\lambda=\lambda_{1}}^{\lambda_{2}} \frac{1}{\lambda}\left\|\delta_{\lambda}^{\sigma}\right\| d \lambda .
$$

Noting that $\operatorname{prox}_{\lambda}^{f \sigma}\left(y_{\lambda}\right)$ is at most $L_{f} \sqrt{d} \sigma$-suboptimal for $f(x)+\frac{\lambda}{2}\left\|x-y_{\lambda}\right\|^{2}$, we have

$$
\left\|\delta_{\lambda}-\delta_{\lambda}^{\sigma}\right\|=\left\|\operatorname{prox}_{\lambda}^{f}\left(y_{\lambda}\right)-\operatorname{prox}_{\lambda}^{f_{\sigma}}\left(y_{\lambda}\right)\right\| \leq \sqrt{\frac{2 L_{f} \sqrt{d} \sigma}{\lambda}}
$$

by $\lambda$-strong-convexity of $x \mapsto f(x)+\frac{\lambda}{2}\left\|x-y_{\lambda}\right\|^{2}$. Substituting back, we find that

$$
\left\|\delta_{\lambda_{2}}-\delta_{\lambda_{1}}\right\| \leq \int_{\lambda=\lambda_{1}}^{\lambda_{2}}\left\|\frac{d}{d \lambda} y_{\lambda}\right\| d \lambda+\int_{\lambda=\lambda_{1}}^{\lambda_{2}} \frac{1}{\lambda}\left\|\delta_{\lambda}\right\| d \lambda+\sqrt{\sigma} \cdot \frac{5 \lambda_{2} L_{f}^{1 / 2} d^{1 / 4}}{\lambda_{1}^{3 / 2}} .
$$

For all $\sigma>0$. Taking the limit $\sigma \rightarrow 0$ concludes the proof.
With this, we prove our desired bound on the continuity of $\hat{\Delta}(\lambda)$.
Lemma 23 Let $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ be convex and $L_{f}$-Lipschitz, and for $R>0$ let $x, v \in \mathbb{R}^{d}$ such that $\|x-v\| \leq 2 R$. For any $0<\lambda_{1} \leq \lambda_{2}$, the function $\hat{\Delta}(\lambda)$ defined in eq. (31) satisfies

$$
\left|\hat{\Delta}\left(\lambda_{1}\right)-\hat{\Delta}\left(\lambda_{2}\right)\right| \leq\left(R+\frac{L_{f}}{\lambda_{1}}\right) \log \frac{\lambda_{2}}{\lambda_{1}}
$$

Proof Note that for all $t>0$

$$
\begin{aligned}
\frac{d}{d t} \alpha_{t} & =\frac{1}{1+t+\sqrt{1+2 t}}-\frac{t+t(1+2 t)^{-1 / 2}}{(1+t+\sqrt{1+2 t})^{2}}=\frac{(1+t) \sqrt{1+2 t}+1+2 t-t \sqrt{1+2 t}-t}{(1+t+\sqrt{1+2 t})^{2} \sqrt{1+2 t}} \\
& =\frac{1}{(1+t+\sqrt{1+2 t}) \sqrt{1+2 t}} \in\left[0, \frac{1}{2 t}\right]
\end{aligned}
$$

Consequently,

$$
\left|\hat{\Delta}\left(\lambda_{1}\right)-\hat{\Delta}\left(\lambda_{2}\right)\right| \leq\left\|\frac{d}{d \lambda} y_{\lambda}\right\|=\left\|(v-x) \frac{d}{d \lambda} \alpha_{2 A \lambda}\right\| \leq \frac{2 A}{4 A \lambda}\|v-x\|=\frac{1}{2 \lambda}\|v-x\| \leq \frac{R}{\lambda} .
$$

By Lemma 21 we have

$$
\left\|\delta_{\lambda_{2}}-\delta_{\lambda_{1}}\right\| \leq \int_{\lambda=\lambda_{1}}^{\lambda_{2}}\left\|\frac{d}{d \lambda} y_{\lambda}\right\| d \lambda+\int_{\lambda=\lambda_{1}}^{\lambda_{2}} \frac{1}{\lambda}\left\|\delta_{\lambda}\right\| d \lambda .
$$

We also have $\left\|\delta_{\lambda}\right\| \leq \frac{L_{f}}{\lambda}$ from Lemma 15. Substituting back and using $\lambda \geq \lambda_{1}$, we obtain

$$
\left\|\delta_{\lambda_{2}}-\delta_{\lambda_{1}}\right\| \leq\left(R+\frac{L_{f}}{\lambda}\right) \int_{\lambda_{1}}^{\lambda_{2}} \frac{d \lambda}{\lambda} \leq\left(R+\frac{L_{f}}{\lambda_{1}}\right) \log \frac{\lambda_{2}}{\lambda_{1}}
$$

as claimed.
Proof [Proof of Proposition 17] We first prove the correctness of our procedure, recalling the notation $\hat{x}_{\lambda}=\operatorname{prox}_{\lambda}^{f}\left(y_{\lambda}\right), \hat{\Delta}(\lambda)=\left\|\hat{x}_{\lambda}-y_{\lambda}\right\|$ and $\Delta(\lambda)=\left\|\mathcal{O}_{\lambda, \delta}\left(y_{\lambda}\right)-y_{\lambda}\right\|$, where $\delta=\frac{r}{17}$ throughout. Our analysis will use the following fact: whenever $\Delta(\lambda)<r-\delta$ then (by Lemma 11) we have $\left\|\operatorname{bprox}_{\lambda, r}^{f}\left(y_{\lambda}\right)-y_{\lambda}\right\|<r$ and consequently $\operatorname{bprox}_{\lambda, r}^{f}\left(y_{\lambda}\right)=\operatorname{prox}_{\lambda}^{f}\left(y_{\lambda}\right)=\hat{x}_{\lambda}$ and $|\hat{\Delta}(\lambda)-\Delta(\lambda)| \leq \delta$.

By $\lambda_{\max } \geq \frac{2 L_{f}}{r}$ and Lemma 15, we have $\hat{\Delta}\left(\lambda_{\max }\right) \leq r / 2$ and $\operatorname{bprox}_{\lambda_{\max }, r}^{f}\left(y_{\lambda}\right)=\operatorname{prox}_{\lambda_{\max }}^{f}\left(y_{\lambda}\right)$. By Lemma 11 we have $\left\|\mathcal{O}_{\lambda_{\max }, \delta}\left(y_{\lambda_{\max }}\right)-\hat{x}_{\lambda_{\max }}\right\| \leq \delta=r / 17$, and consequently $\Delta\left(\lambda_{\max }\right) \leq$ $r / 2+r / 16 \leq 13 r / 16$. Therefore, the first while loop (Line 14) executes at least once.

Suppose the procedure terminates on Line 15 . Denoting the return value as $\lambda$, by the condition on Line 14 we have $\Delta(\lambda) \leq \frac{7 r}{8}$, which by Lemma 11 implies $\hat{\Delta}(\lambda)<\frac{15 r}{16}<r$. Consequently, outcome (b) in Proposition 17 occurs.

Next, consider the case where the procedure terminates due to the condition $\Delta(\lambda) \in\left[\frac{13 r}{16}, \frac{15 r}{16}\right]$ (either in line 17 or in line 18). Applying Lemma 11, we conclude that $\hat{\Delta}(\lambda) \in\left[\frac{3 r}{4}, r\right)$. Consequently, outcome (a) in Proposition 17 occurs.

It remains to check the case where the procedure terminates due to the condition $\log \frac{\lambda_{u}}{\lambda_{\ell}}<$ $\frac{r}{8\left(R+L_{f} / \lambda_{\ell}\right)}$ in line 18. Note that the binary search maintains the invariant $\Delta\left(\lambda_{\ell}\right)>\frac{15 r}{16}$ and $\Delta\left(\lambda_{u}\right)<\frac{13 r}{16}$. By Lemma 11, this implies

$$
\hat{\Delta}\left(\lambda_{\ell}\right)>\frac{7 r}{8}>\hat{\Delta}\left(\lambda_{u}\right) .
$$

By continuity of $\hat{\Delta}(\lambda)$, we therefore always have $\hat{\Delta}\left(\lambda^{\prime}\right)=\frac{7 r}{8}$ for some $\lambda^{\prime} \in\left(\lambda_{\ell}, \lambda_{u}\right)$. Therefore, if $\|x-v\| \leq 2 R$, Lemma 23 guarantees that

$$
\left|\hat{\Delta}\left(\lambda_{m}\right)-\frac{7}{8} r\right|=\left|\hat{\Delta}\left(\lambda_{m}\right)-\hat{\Delta}\left(\lambda^{\prime}\right)\right| \leq\left(R+\frac{L_{f}}{\min \left\{\lambda^{\prime}, \lambda_{m}\right\}}\right) \log \left|\frac{\lambda_{m}}{\lambda^{\prime}}\right| \leq \frac{1}{2}\left(R+\frac{L_{f}}{\lambda_{\ell}}\right) \log \frac{\lambda_{u}}{\lambda_{\ell}}
$$

Consequently, when $\log \frac{\lambda_{u}}{\lambda_{\ell}}<\frac{r}{8\left(R+L_{f} / \lambda_{\ell}\right)}$ and $\|x-v\| \leq 2 R$, we are guaranteed that $\hat{\Delta}\left(\lambda_{m}\right) \in$ (3r/4, r), and outcome (a) in Proposition 17 occurs, concluding the proof of correctness.

Next, we briefly justify the bounds on the $\lambda^{\prime}$ values with which we query the BROO in the $\lambda$-Bisection procedure. By construction, we have $\lambda^{\prime} \in\left[\lambda_{\min }, \lambda_{\max }\right]$. Let $\lambda_{o}$ and $\lambda_{s}$ be the procedure's output and smallest queried $\lambda$ values, respectively. When terminating on lines 15 or 17, we clearly have $\lambda_{s}=\lambda_{o}$. Furthermore, when terminating on line $18 \lambda_{s}$ equals $\lambda_{\ell}$ when entering line 18 for the first time, and consequently $\lambda_{o} \leq 2 \lambda_{s}$.

Finally, we bound the total number of BROO queries. The first while loop requires at most $\log _{2} \frac{\lambda_{\text {max }}}{\lambda_{\text {min }}}$ queries. In the second while loop, initially we have $\log _{2} \frac{\lambda_{u}}{\lambda_{\ell}}=1$ and each query decreases $\log _{2} \frac{\lambda_{u}}{\lambda_{\ell}}$ by a factor of 2 . Therefore, using $\lambda_{\ell} \geq \lambda_{\min }$, the stopping condition $\log \frac{\lambda_{u}}{\lambda_{\ell}}<\frac{r}{8\left(R+L_{f} / \lambda_{\ell}\right)}$ must hold after $O\left(\log \left(\frac{R+L_{f} / \lambda_{\min }}{r}\right)\right)$ queries. Given $\lambda_{\max }=\frac{2 L_{f}}{r}, \lambda_{\min }=\frac{\epsilon}{6 r R}$ the total number of BROO queries is bounded as

$$
\log \left(\frac{\lambda_{\max }}{\lambda_{\min }}\right)+\log \left(\frac{R+L_{f} / \lambda_{\min }}{r}\right)=O\left(\log \frac{L_{f} R^{2}}{r \epsilon}+\log \frac{L_{f} R^{2}}{r \epsilon}\right)=O\left(\log \frac{L_{f} R^{2}}{r \epsilon}\right)
$$

queries, giving the claimed complexity bound.

## Appendix C. BROO implementation

## C.1. Proof of Lemma 3

Recall that the "eponentiated softmax" function is defined as follows

$$
\Gamma_{\epsilon, \lambda}(x)=\epsilon^{\prime} \cdot \exp \left(\frac{F_{\operatorname{smax}, \epsilon}^{\lambda}(x)-F_{\operatorname{smax}, \epsilon}^{\lambda}(\bar{x})}{\epsilon^{\prime}}\right)=\sum_{i \in[N]} p_{i}(\bar{x}) \gamma_{i}(x) \text { where } \gamma_{i}(x):=\epsilon^{\prime} e^{\frac{f_{i}^{\lambda}(x)-f_{i}^{\lambda}(\bar{x})}{\epsilon_{i}}},
$$

and note that $\Gamma_{\epsilon, \lambda}(\bar{x})=\epsilon^{\prime}$.
Lemma 3 Let $f_{1}, \cdots, f_{N}$ each be $L_{f}$-Lipschitz and $L_{g}$-smooth gradients. For any $c>0, r \leq$ $c \epsilon^{\prime} / L_{f}$, and $\lambda \leq c L_{f} / r$ let $C=\left(1+c+c^{2}\right) e^{c+c^{2} / 2}$. The exponentiated softmax $\Gamma_{\epsilon, \lambda}$ satisfies the following properties for any $\bar{x} \in \mathbb{R}^{d}$.

1. $F_{\mathrm{smax}, \epsilon}^{\lambda}(x)$ and $\Gamma_{\epsilon, \lambda}$ have the same minimizer $x_{\star}$ in $\mathbb{B}_{r}(\bar{x})$. Moreover, for every $x \in \mathbb{B}_{r}(\bar{x})$,

$$
F_{\text {smax }, \epsilon}^{\lambda}(x)-F_{\operatorname{smax}, \epsilon}^{\lambda}\left(x_{\star}\right) \leq C\left(\Gamma_{\epsilon, \lambda}(x)-\Gamma_{\epsilon, \lambda}\left(x_{\star}\right)\right) .
$$

2. Restricted to $\mathbb{B}_{r}(\bar{x})$, each function $\gamma_{i}$ defined in (5) is $C L_{f}$-Lipschitz, $C^{-1} \lambda$ strongly convex, and $C\left(L_{g}+\lambda+L_{f}^{2} / \epsilon^{\prime}\right)$-smooth.
Proof For the first statement, we note that $\Gamma_{\epsilon, \lambda}$ is a monotonic increasing transformation of $F_{\mathrm{smax}, \epsilon}^{\lambda}$ and consequently they have the same minimizer $x_{\star}$ in $\mathbb{B}_{r}(\bar{x})$. Further

$$
\begin{aligned}
F_{\mathrm{smax}, \epsilon}^{\lambda}(x)-F_{\mathrm{smax}, \epsilon}^{\lambda}\left(x_{\star}\right) & =\epsilon^{\prime} \log \left(\frac{\Gamma_{\epsilon, \lambda}(x)}{\Gamma_{\epsilon, \lambda}\left(x_{\star}\right)}\right)=\Gamma_{\epsilon, \lambda}(\bar{x}) \log \left(1+\frac{\Gamma_{\epsilon, \lambda}(x)-\Gamma_{\epsilon, \lambda}\left(x_{\star}\right)}{\Gamma_{\epsilon, \lambda}\left(x_{\star}\right)}\right) \\
& \leq \frac{\Gamma_{\epsilon, \lambda}(\bar{x})}{\Gamma_{\epsilon, \lambda}\left(x_{\star}\right)}\left(\Gamma_{\epsilon, \lambda}(x)-\Gamma_{\epsilon, \lambda}\left(x_{\star}\right)\right)
\end{aligned}
$$

where the final inequality uses $\log (1+x) \leq x$. Next, note that Lipschitz continuity of each $f_{i}$ implies

$$
f_{i}^{\lambda}\left(x_{\star}\right) \geq f_{i}\left(x_{\star}\right) \geq f_{i}(\bar{x})-L_{f}\left\|x_{\star}-\bar{x}\right\|=f_{i}^{\lambda}(\bar{x})-L_{f}\left\|x_{\star}-\bar{x}\right\| .
$$

Substituting $\left\|x_{\star}-\bar{x}\right\| \leq r \leq c \epsilon^{\prime} / L_{f}$ and $C \geq e^{c}$, we have that

$$
e^{f_{i}^{\lambda}\left(x_{\star}\right) / \epsilon^{\prime}} \geq e^{f_{i}^{\lambda}(\bar{x}) / \epsilon^{\prime}-c} \geq \frac{1}{C} e^{f_{i}^{\lambda}(\bar{x})} .
$$

Consequently, we have

$$
\frac{\Gamma_{\epsilon, \lambda}(\bar{x})}{\Gamma_{\epsilon, \lambda}\left(x_{\star}\right)}=\frac{\sum_{i \in[N]} e^{f_{i}^{\lambda}(\bar{x}) / \epsilon^{\prime}}}{\sum_{i \in[N]} e^{f_{i}^{\lambda}\left(x_{\star}\right) / \epsilon^{\prime}}} \leq C
$$

which proves the first statement.
For the second statement, we first compute the gradient and Hessian of the function $\gamma_{i}(x)$ as

$$
\begin{aligned}
\nabla \gamma_{i}(x) & =\exp \left(\frac{\lambda}{2 \epsilon^{\prime}}\|x-\bar{x}\|^{2}\right) \exp \left(\frac{f_{i}(x)-f_{i}(\bar{x})}{\epsilon^{\prime}}\right)\left[\nabla f_{i}(x)+\lambda(x-\bar{x})\right], \text { and } \\
\nabla^{2} \gamma_{i}(x) & =\exp \left(\frac{f_{i}(x)-f_{i}(\bar{x})}{\epsilon^{\prime}}\right) H_{i}, \text { where } \\
H_{i} & =\nabla^{2} f_{i}(x)+\lambda I+\frac{1}{\epsilon^{\prime}} \nabla f_{i}(x) \nabla f_{i}(x)^{\top}+\frac{\lambda^{2}}{\epsilon^{\prime}}(x-\bar{x})(x-\bar{x})^{\top} .
\end{aligned}
$$

(Note that the expression for $\nabla^{2} \gamma_{i}$ assumes that $f_{i}$ is twice differentiable almost everywhere; we only rely on it for Section 4.3 where this holds.)

Now we note that for any $i \in[N]$

$$
\left\|\nabla \gamma_{i}(x)\right\| \leq \exp \left(\lambda r^{2} / 2 \epsilon^{\prime}+L_{f} r / \epsilon^{\prime}\right)\left(\lambda r+L_{f}\right)
$$

and

$$
\lambda e^{-\lambda r^{2} / 2 \epsilon^{\prime}-L_{f} r / \epsilon^{\prime}} I \preceq \nabla^{2} \gamma_{i}(x) \preceq e^{\lambda r^{2} / 2 \epsilon^{\prime}+L_{f} r / \epsilon^{\prime}}\left(L_{g}+\lambda+L_{f}^{2} / \epsilon^{\prime}+\lambda^{2} r^{2} / \epsilon^{\prime}\right) I,
$$

where for the inequality we use the fact that $f_{i}$ is $L_{f}$-Lipschitz, $L_{g}$-smooth, and that $x \in \mathbb{B}_{r}(\bar{x})$.
Now by plugging in the assumption that $r \leq c \epsilon^{\prime} / L_{f}$ and $\lambda \leq c L_{f} / r$ we have $\lambda r \leq c L_{f}$, $\lambda^{2} r^{2} / \epsilon^{\prime} \leq c^{2} \lambda$, and $\exp \left(\lambda r^{2} / 2 \epsilon^{\prime}+L_{f} r / \epsilon^{\prime}\right) \leq \exp \left(c+c^{2} / 2\right)$. Thus we can further simplify the bounds on the gradient and Hessian of $\gamma_{i}(x)$ by definition of $C$ as

$$
\left\|\nabla \gamma_{i}(x)\right\| \leq C L_{f} \quad \text { and } \quad C^{-1} \lambda I \preceq \nabla^{2} \gamma_{i}(x) \preceq C\left(L_{g}+\lambda+L_{f}^{2} / \epsilon^{\prime}\right) I,
$$

which completes the proof.

## C.2. SGD implementation

We first cite the following stochastic gradient method with restarts that obtains an $\widetilde{O}(1 / \mu T)$ bound for $\mu$-strongly convex function.

Lemma 24 (Theorem 11 in Hazan and Kale [20]) Given a $\mu$-strongly-convex objective function $f: \mathcal{X} \rightarrow \mathbb{R}$ with minimizer $x_{\star}$ with an unbiased stochastic estimator with norm at most $G$ in the convex compact set $\mathcal{X}$, Epoch-SGD-Proj algrotihm finds an approximate minimizer $\tilde{x}$ satisfying with probability $1-\sigma$

$$
f(\tilde{x})-f\left(x_{\star}\right) \leq O\left(\frac{G^{2} \log (\log (T) / \sigma)}{\mu T}\right)
$$

using $T$ stochastic gradient queries.
Applying the lemma immediately gives the following guarantee for minimizing $\Gamma_{\epsilon, \lambda}$ inside a ball of radius $r_{\epsilon}$ and hence implementing an $r_{\epsilon}$-BROO for $F_{\text {smax }, \epsilon}$.

Corollary 4 Let $f_{1}, f_{2}, \cdots, f_{N}$ be $L_{f}$ Lipschitz, let $\sigma \in(0,1), \epsilon, \delta>0$ and $r_{\epsilon}=\epsilon /\left(2 \log N \cdot L_{f}\right)$. For any $\bar{x} \in \mathbb{R}^{d}$ and $\lambda \leq O\left(L_{f} / r_{\epsilon}\right)$, with probability at least $1-\sigma$, Algorithm 2 outputs a valid $r_{\epsilon}-B R O O$ response for $F_{\text {smax, } \epsilon}$ to query $\bar{x}$ with regularization $\lambda$ and accuracy $\delta$, and has cost

$$
\begin{equation*}
O\left(\mathcal{T}_{f} N+\left(\mathcal{T}_{g}+\mathcal{T}_{f}\right) \frac{L_{f}^{2}}{\lambda^{2} \delta^{2}} \log \left(\frac{\log \left(L_{f} / \lambda \delta\right)}{\sigma}\right)\right) . \tag{6}
\end{equation*}
$$

Proof We first note that by choice of $r_{\epsilon}$ and bounds on $\lambda, \Gamma_{\epsilon, \lambda}$ is $\Omega(\lambda)$-strongly convex according to Lemma 3. At each iteration, we sample $i \in[N]$ with probability $p_{i}$ and compute stochastic gradient

$$
\nabla \gamma_{i}(x)=\gamma_{i}(x)\left(\frac{\lambda}{\epsilon^{\prime}}(x-\bar{x})+\nabla f_{i}(x) / \epsilon^{\prime}\right) \text { bounded by } G=O\left(L_{f}\right)
$$

```
    Algorithm 2: Epoch-SGD-Proj on the exponentiated softmax
Input: Functions \(f_{1}, \ldots, f_{N}\), ball center \(\bar{x}\), ball radius \(r_{\epsilon}\), regularization strength \(\lambda\), smoothing
        parameter \(\epsilon^{\prime}\), failure probability \(\sigma\)
Parameters: Step size \(\eta_{1}=1 /(3 \lambda)\), domain size \(D_{1}=\Theta(G \sqrt{\log (\log (T) / \delta)} / \lambda), T_{1}=450\)
        and total iteration budget \(T\)
Output: Approximate minimizer of \(\Gamma_{\epsilon, \lambda}\) (and hence \(F_{\text {smax }, \epsilon}^{\lambda}\) ) in \(\mathbb{B}_{r_{\epsilon}}(\bar{x})\) )
Precompute sampling probabilities \(p_{i}=e^{f_{i}(\bar{x}) / \epsilon^{\prime}} / \sum_{i \in[N]} e^{f_{i}(\bar{x}) / \epsilon^{\prime}}\) for all \(i \in[N]\)
Initialize \(x_{1}^{1} \in \mathbb{B}_{r_{\epsilon}}(\bar{x})\) arbitrarily, set \(k=1\)
while \(\sum_{i \in[k]} T_{i} \leq T\) do
    for \(t=1, \ldots, T_{k}\) do
        Sample \(i \in[N]\) with probability \(p_{i}\)
        Query stochastic gradient \(\hat{g}_{t}=e^{\left(f_{i}^{\lambda}(x)-f_{i}^{\lambda}(\bar{x})\right) / \epsilon^{\prime}} \nabla f_{i}^{\lambda}(x)\)
        Update \(x_{t+1}^{k} \leftarrow \Pi_{\mathbb{B}_{r_{\epsilon}}(\bar{x}) \cap \mathbb{B}_{D_{k}}\left(x_{1}^{k}\right)}\left(x_{t}^{k}-\eta_{k} \hat{g}_{t}\right)\)
        Let \(x_{1}^{k+1} \leftarrow \frac{1}{T_{k}} \sum_{t \in\left[T_{k}\right]} x_{t}^{k}\)
        Update parameters \(T_{k+1} \leftarrow 2 T_{k}, \eta_{k+1} \leftarrow \eta_{k} / 2, D_{k+1} \leftarrow D_{k} / \sqrt{2}, k \leftarrow k+1\)
return \(x_{1}^{k}\)
```

following from the second statement of Lemma 3. Thus by directly applying Lemma 24 with $T=$ $\Theta\left(L_{f}^{2} \lambda^{-2} \delta^{-2} \log \left(\log \left(L_{f} / \lambda \delta\right) / \sigma\right)\right)$ the algorithm outputs an approximate minimizer $\tilde{x}$ satisfying

$$
\Gamma_{\epsilon, \lambda}(\tilde{x})-\min _{x \in \mathbb{B}_{r_{\epsilon}}(\tilde{x})} \Gamma_{\epsilon, \lambda}(x) \leq O\left(\frac{L_{f}^{2}}{\lambda T} \log (\log (T) / \sigma)\right) \leq \frac{\lambda \delta^{2}}{6 e^{2}} .
$$

By the first property of Lemma 3, bound above implies that for $x_{\star}=\operatorname{bprox}_{\lambda, r}{ }_{\lambda}{ }_{\text {smax }, \epsilon}$

$$
F_{\mathrm{smax}, \epsilon}^{\lambda}(x)-F_{\mathrm{smax}, \epsilon}^{\lambda}\left(x_{\star}\right) \leq 3 e^{2}\left(\Gamma_{\epsilon, \lambda}(x)-\Gamma_{\epsilon, \lambda}\left(x_{\star}\right)\right) \leq \frac{\lambda \delta^{2}}{2},
$$

i.e., algorithm outputs a valid $r_{\epsilon}$-BROO response for $F_{\mathrm{smax}, \epsilon}$.

The bound on the total computational cost follows from noticing that the initialization cost is dominated by $N$ function value queries, and that the cost of each step in the stochastic gradient descent is dominated by a function and a gradient query at the current iteration.

## C.3. Accelerated variance reduction implementation

We first cite the following accelerated variance reduction guarantee.
Lemma 25 (Theorem 5.4 in Allen-Zhu [1]) Let $f_{1}, \ldots, f_{N}$ be L-smooth and $\mu$-strongly-convex, let $F(x)=\sum_{i \in[n]} w_{i} f_{i}(x)$ with $w_{i} \geq 0$ and $\sum_{i \in[N]} w_{i}=1$, and let $x_{\star} \in \arg \min F(x)$. For any $s \in \mathbb{N}$, Katyushal with batch size $b=1$ and initial point $\bar{x}$ finds an approximate solution $\tilde{x}_{s}$ satisfying

$$
\mathbb{E}\left[F\left(\tilde{x}_{s}\right)-F\left(x_{\star}\right)\right] \leq \frac{1}{2^{s}}\left[F(\bar{x})-F\left(x_{\star}\right)\right]
$$

using

$$
O\left(s \cdot\left(N+\sqrt{\frac{N \cdot L}{\mu}}\right)\right) \text { evaluations of } \nabla f_{i}(x) .
$$

An immediate corollary of Lemma 25 provides a high-probability guarantee.
Corollary 26 Under the same assumptions as Lemma 25, for any $\epsilon>0$ and $\sigma \in(0,1)$ Katyushal with batch size $b=1$ and initial point $\bar{x}$ finds an approximate solution $\tilde{x}$ satisfying

$$
F(\tilde{x})-F\left(x_{\star}\right) \leq \epsilon \text { with probability at least } 1-\sigma
$$

using

$$
O\left(\left(N+\sqrt{\frac{N \cdot L}{\mu}}\right) \log \left(\frac{F(\bar{x})-F\left(x_{\star}\right)}{\varepsilon \sigma}\right)\right) \text { evaluations of } \nabla f_{i}(x) .
$$

Proof Since $F\left(\tilde{x}_{s}\right)-F\left(x_{\star}\right) \geq 0$, Markov's inequality and Lemma 25 imply that for any $\epsilon>0$,

$$
\mathbb{P}\left(F\left(\tilde{x}_{s}\right)-F\left(x_{\star}\right) \geq \epsilon\right) \leq \frac{\mathbb{E}\left[F\left(\tilde{x}_{s}\right)-F\left(x_{\star}\right)\right]}{\epsilon} \leq \frac{(1 / 2)^{s}\left(f(\bar{x})-f\left(x_{\star}\right)\right)}{\epsilon}
$$

and thus the high-probability bound follows .
Specializing Corollary 26 for $\Gamma_{\epsilon, \lambda}(x)=\sum_{i \in[N]} p_{i}(\bar{x}) \gamma_{i}(x)$, we obtain the following guarantee.
Corollary 5 Let $f_{1}, \cdots, f_{N}$ be $L_{f}$-Lipschitz and $L_{g}$-smooth, let $\sigma \in(0,1), \epsilon, \delta>0, \epsilon^{\prime}=$ $\epsilon /(2 \log N)$ and $r_{\epsilon}=\epsilon^{\prime} / L_{f}$. For any $\bar{x} \in \mathbb{R}^{d}$ and $\lambda \leq O\left(L_{f} / r_{\epsilon}\right)$, with probability at least $1-\sigma$, Katyushal [1] outputs a valid $r_{\epsilon}$-BROO response to query $\bar{x}$ with regularization $\lambda$ and accuracy $\delta$, and has computational cost

$$
\begin{equation*}
O\left(\left(\mathcal{T}_{f}+\mathcal{T}_{g}\right)\left(N+\frac{\sqrt{N}\left(L_{f}+\sqrt{\epsilon^{\prime} L_{g}}\right)}{\sqrt{\lambda \epsilon^{\prime}}}\right) \log \left(\frac{L_{f} r_{\epsilon}}{\lambda \delta^{2} \sigma}\right)\right) . \tag{7}
\end{equation*}
$$

## Proof

By the choice of $r_{\epsilon}$ and the bound $\lambda=O\left(L_{f}^{2} / \epsilon^{\prime}\right)$, Lemma 3 guarantees that the $\gamma_{i}$ have strong convexity $\mu=\Omega(\lambda)$, and smoothness $L=O\left(L_{g}+\lambda+L_{f}^{2} / \epsilon^{\prime}+\lambda^{2} r_{\epsilon}^{2} / \epsilon^{\prime}\right)=O\left(L_{g}+L_{f}^{2} / \epsilon^{\prime}\right)$. In addition, for $x_{\star}:=\arg \min _{x \in \mathbb{B}_{r_{\epsilon}}(\bar{x})} \Gamma_{\epsilon, \lambda}(x)$ one has

$$
\begin{aligned}
\Gamma_{\epsilon, \lambda}(\bar{x})-\Gamma_{\epsilon, \lambda}\left(x_{\star}\right) & \leq\left\langle\nabla_{x} \Gamma_{\epsilon, \lambda}\left(x_{\star}\right), \bar{x}-x_{\star}\right\rangle \\
& \leq\left\|\nabla_{x} \Gamma_{\epsilon, \lambda}\left(x_{\star}\right)\right\|\left\|\bar{x}-x_{\star}\right\|=O\left(L_{f} r_{\epsilon}\right),
\end{aligned}
$$

where we use the second property for bounding $\nabla_{x} \Gamma_{\epsilon, \lambda}(\bar{x})$ from Lemma 3. Plugging these into the complexity of Lemma 25 and noticing that each evaluation of $\nabla \gamma_{i}$ requires evaluation of $f_{i}$ and $\nabla f_{i}$ (assuming $f_{i}(\bar{x})$ is pre-stored) gives the stated complexity bound.

## C.4. Proof of Theorem 6

With Corollaries 4 and 5 established, we prove Theorem 6.
Theorem 6 Let $f_{1}, f_{2}, \ldots, f_{N}$ be $L_{f}$-Lipschitz, let $x_{\star}$ be a minimizer of $F_{\max }(x)=\max _{i \in[N]} f_{i}(x)$ and assume $\left\|x_{0}-x_{\star}\right\| \leq R$ for a given initial point $x_{0}$ and some $R>0$. For any $\epsilon>0$, Algorithm 1 with the BROO implementation for $F_{\mathrm{smax}, \epsilon}$ in Algorithm 2 solves the problem (1) with probability at least $\frac{99}{100}$ and has computational cost

$$
\begin{equation*}
O\left(\left(\frac{L_{f} R \log N}{\epsilon}\right)^{2 / 3}\left(\mathcal{T}_{f} N+\left(\frac{L_{f} R}{\epsilon}\right)^{2} \cdot\left(\mathcal{T}_{f}+\mathcal{T}_{g}\right) \log K\right) \log ^{2} K\right) \tag{8}
\end{equation*}
$$

where $K:=L_{f} R \epsilon^{-1} \log N$. If moreover $f_{1}, f_{2}, \ldots, f_{N}$ are each $L_{g}$-smooth, then Algorithm 1 with a BROO implementation for $F_{\text {smax }, \epsilon}$ using Kayushal solves (1) with probability $\geq \frac{99}{100}$ and has cost

$$
O\left(\left(\mathcal{T}_{f}+\mathcal{T}_{g}\right)\left(\left(\frac{L_{f} R \log N}{\epsilon}\right)^{2 / 3} N+\left(\frac{L_{f} R \sqrt{\log N}}{\epsilon}+\sqrt{\frac{L_{g} R^{2}}{\epsilon}}\right) \sqrt{N}\right) \log ^{3} K\right) .
$$

Proof We use guarantees of Theorem 2 and Corollaries 4 and 5 on the problem $\min _{x} F_{\text {smax }, \epsilon}(x), \| x-$ $x_{0} \| \leq O(R)$ to prove the correctness and bound the complexity of the algorithm.
Correctness. We first note that Theorem 2 guarantees that we only make BROO calls with $\lambda \leq$ $O\left(L_{f} / r_{\epsilon}\right)$, making Corollaries 4 and 5 applicable. Let

$$
\begin{equation*}
T=O\left(\left(\frac{R}{r_{\epsilon}}\right)^{2 / 3} \log ^{2}\left(\frac{L_{f} R^{2}}{r_{\epsilon} \epsilon}\right)\right) \tag{32}
\end{equation*}
$$

be the upper bound on the total number of oracle calls guaranteed in Theorem 2. Taking $\sigma=$ $\frac{1}{100 T}$ in Corollaries 4 and 5 and applying a union bound, we see that with probability at least $99 / 100$, the outputs of the corresponding BROO implementations are valid throughout the execution of Algorithm 1. Consequently by Theorem 2 we have that Algorithm 1 with accuracy $\epsilon / 2$ outputs $x_{o}$ such that $F_{\text {smax }, \epsilon}\left(x_{o}\right)-\min _{x:\left\|x-x_{0}\right\| \leq R} F_{\text {smax }, \epsilon}(x) \leq \epsilon / 2$. Using the fact that $0 \leq$ $F_{\text {smax }, \epsilon}(x)-F_{\max }(x) \leq \epsilon / 2$ for all $x \in \mathbb{R}^{d}$ [see, e.g., 12 , Lemma 45], we conclude that $F_{\max }\left(x_{o}\right)-$ $\min _{x} F_{\max }(x)=F_{\max }\left(x_{o}\right)-\min _{\left\{x:\left\|x-x_{0}\right\| \leq R\right\}} F_{\max }(x) \leq F_{\text {smax }, \epsilon}\left(x_{o}\right)-\min _{x:\left\|x-x_{0}\right\| \leq R} F_{\text {smax }, \epsilon}(x)+$ $\epsilon / 2 \leq \epsilon$, establishing correctness.
Complexity. Substituting $r_{\epsilon}=\epsilon /\left(2 L_{f} \log N\right)$ into (32), we see that the total number of oracle calls is

$$
T=O\left(\left(\frac{L_{f} R \log N}{r_{\epsilon}}\right)^{2 / 3} \log ^{2}\left(\frac{L_{f} R \log N}{\epsilon}\right)\right)
$$

To bound to complexity of the SGD implementation, we simply multiply $T$ by the per-call complexity bound (6) where we substitute $\delta=\Omega(\epsilon /(\lambda R))$ as guaranteed by Theorem 2 and $\sigma=1 /(100 T)$.

To bound the complexity of the accelerated variance reduction implementation, we similarly substitute $\delta=\Omega(\epsilon /(\lambda R)), \lambda=O\left(L_{f} / r_{\epsilon}\right)$ and $\sigma=1 /(100 T)$ into (7). Summing the result over all oracle calls yields the complexity bound

$$
O\left(\left(\mathcal{T}_{f}+\mathcal{T}_{g}\right) \log \left(\frac{L_{f} R \log N}{\epsilon}\right)\left(N T+\sqrt{N}\left(\sqrt{L_{g}}+L_{f} \sqrt{\frac{\log N}{\epsilon}}\right) \sum_{i \in[T]} \frac{1}{\sqrt{\lambda_{(i)}}}\right)\right)
$$

where $\left\{\lambda_{(i)}\right\}$ is the sequence of $\lambda$ values with which Algorithm 1 calls the BROO implementation. Theorem 2 guarantees that $\sum_{i \in[T]} \frac{1}{\sqrt{\lambda_{(i)}}} \leq O\left(\frac{R}{\sqrt{\epsilon}} \log \frac{L_{f} R^{2}}{r \epsilon}\right)=O\left(\frac{R}{\sqrt{\epsilon}} \log \frac{L_{f} R \log N}{\epsilon}\right)$, completing the proof.

## Appendix D. Lower bound proofs

## D.1. Proof of Proposition 8

In this section, we make frequent use of the indication notation $\mathbb{I}\{\cdot\}$, where $\mathbb{I}\{A\}=1$ if event $A$ holds and $\mathbb{I}\{A\}=0$ otherwise.

Proposition 8 Let $\delta, \alpha \in(0,1)$ and let $N, T \in \mathbb{N}$ with $T \leq N / 2$. Let $\left(f_{i}\right)_{i \in[N]}$ be an $\alpha$-robust $N$-element zero-chain with domain $\mathbb{B}_{1}^{T}(0)$. For $d \geq T+\frac{2}{\alpha^{2}} \log \frac{4 N T^{2}}{\delta}$, draw $U$ uniformly from the set of $d \times T$ orthogonal matrices, and draw $\Pi$ uniformly from the set of permutations of $[N]$. Let $\tilde{f}_{i}(x):=f_{\Pi^{-1}(i)}\left(U^{\top} x\right)$. Let $\left\{\left(i_{t}, x_{t}\right)\right\}_{t \geq 1}$ be the queries of any $N$-element algorithm operating on $\tilde{f}_{1}, \ldots, \tilde{f}_{N}$. Then with probability at least $1-\delta$ we have

$$
\operatorname{prog}_{\alpha}\left(U^{\top} x_{t}\right)<T \text { for all } t \leq \frac{1}{16} N\left(T-\log \frac{2}{\delta}\right) .
$$

Proof Let us define several quantities that are important for our proof. First, we track the maximum progress attained by the algorithm queries.

$$
p_{t}:=\operatorname{prog}_{\alpha}\left(U^{\top} x_{t}\right) \text { and } \bar{p}_{t}:=\max _{s \leq t} p_{s} .
$$

Next, we recursively define a sequence that tracks the algorithm's progress in "unlocking" the relevant elements of the finite sum,

$$
B_{t}:=\mathbb{I}\left\{\Pi^{-1}\left(i_{t}\right)=C_{t}+1\right\} \text { where } C_{t}:=\min \left\{\sum_{s<t} B_{s}, T\right\} .
$$

To understand these definitions, note that $C_{t}$ is the largest number $k$ such that $1,2, \ldots, k$ is a subsequence of $\Pi^{-1}\left(i_{1}\right), \Pi^{-1}\left(i_{2}\right), \ldots, \Pi^{-1}\left(i_{t}\right)$. Therefore, a "zero-respecting" algorithm (satisfying $\left.p_{t} \leq \max _{s<t}\left\{\operatorname{prog}_{\alpha}\left(\nabla f_{\Pi^{-1}(i)}\left(U^{\top} x_{s}\right)\right)\right\}\right)$ can only query points with progress at most $C_{t}$. We define the event that the general algorithm under consideration behaves as though it was zerorespecting for the first $t$ iterations as

$$
\mathfrak{Z} \Re_{t}:=\left\{\bar{p}_{s} \leq C_{s} \text { for all } s \leq t\right\} .
$$

We also define the stopping time

$$
\Theta_{k}:=\min \left\{t \mid C_{t}=k\right\}
$$

and the difference sequence

$$
\Delta_{k}:=\Theta_{k}-\Theta_{k-1} .
$$

Let

$$
\tau:=\frac{1}{16} N\left(T-\log \frac{2}{\delta}\right)
$$

so that our goal is to prove that $\mathbb{P}\left(\bar{p}_{\lfloor\tau\rfloor}<T\right)>1-\delta$. Note that the intersection of the events $C_{\lfloor\tau\rfloor}<T$ and $\mathfrak{Z} \mathfrak{R}_{N T} \subset \mathfrak{Z} \mathfrak{R}_{\lfloor\tau\rfloor}$ imply the desired event $\bar{p}_{\lfloor\tau\rfloor}<T$. Moreover, $C_{\lfloor\tau\rfloor}<T$ is equivalent to $\Theta_{T}>\tau$. Consequently, we can upper bound $\mathbb{P}\left(\bar{p}_{\lfloor\tau]} \geq T\right)$ by the probability that $\mathfrak{Z} \mathfrak{R}_{N T}$ does not occur plus the probability that both $\Theta_{T} \leq \tau$ and $\mathfrak{Z} \mathfrak{R}_{N T}$ occur, i.e.,

$$
\mathbb{P}\left(\bar{p}_{\lfloor\tau\rfloor} \geq T\right) \leq \mathbb{P}\left(\mathfrak{Z} \mathfrak{R}_{N T}^{c}\right)+\mathbb{P}\left(\Theta_{T} \leq \tau, \mathfrak{Z} \mathfrak{R}_{N T}\right)
$$

Our strategy is to bound each of $\mathbb{P}\left(\mathfrak{Z} \Re_{N T}^{c}\right)$ and $\mathbb{P}\left(\Theta_{T} \leq \tau, \mathfrak{Z} \Re_{N T}\right)$ by $\delta / 2$.
As one final piece of of notation, we write $U_{\leq k}$ for the first $k$ columns of $U$, and $\left.\Pi\right|_{k}$ as shorthand for $\Pi(1), \ldots, \Pi(k)$.

Lemma 27 For any $t \geq 1$ and $k \leq T$, if the events $\mathfrak{Z} \Re_{t-1}$ and $C_{t} \leq k$ hold, the oracle responses to queries $\left(i_{s}, x_{s}\right)_{s<t}$, as well as the queries $\left(i_{s}, x_{s}\right)_{s \leq t}$, are deterministic (measurable) functions of $\zeta,\left.\Pi\right|_{k}$ and $U_{\leq k}$. Moreover,
(a) The random variable $\Theta_{k} \mathbb{I}\left\{\mathfrak{Z} \mathfrak{R}_{\Theta_{k}}\right\}$ is measurable w.r. $\zeta,\left.\Pi\right|_{k}$ and $U$.
(b) When $\mathfrak{Z}_{t-1}$ holds, $\mathbb{I}\left\{C_{t}=k\right\}$ is measurable w.r.t. $\zeta, \Pi$ and $U_{\leq k}$.

Proof Consider the query $i_{s}, x_{s}$ for $s<t$; we first show that under $\mathfrak{Z} \mathfrak{R}_{t-1}$ and $C_{t} \leq k$ we can compute the oracle response to this query using only $\left.\Pi\right|_{k}$ and $U_{\leq k}$. Since $\mathcal{Z} \Re_{t-1}$ holds, we have that $\operatorname{prog}_{\alpha}\left(U^{\top} x_{s}\right)=p_{s} \leq C_{s}$. Invoking Definition 7 of the $N$-element zero chain, there exists a neighborhood of $x_{s}$ such that for all $y$ in that neighborhood we have

$$
\tilde{f}_{i_{s}}(y)=f_{\Pi^{-1}\left(i_{s}\right)}\left(U^{\top} y\right)= \begin{cases}f_{\Pi^{-1}\left(i_{s}\right)}\left(U_{\leq C_{s}}^{\top} y\right) & \Pi^{-1}\left(i_{s}\right)<C_{s}+1  \tag{33}\\ f_{\Pi^{-1}\left(i_{s}\right)}\left(U_{\leq C_{s}+1}^{\top} y\right) & \Pi^{-1}\left(i_{s}\right)=C_{s}+1 \\ f_{N}\left(U_{\leq C_{s}}^{\top} y\right) & \Pi^{-1}\left(i_{s}\right)>C_{s}+1\end{cases}
$$

Recall that $B_{s}=\mathbb{I}\left\{\Pi^{-1}\left(i_{s}\right)=C_{s}+1\right\}$ and that $C_{s+1}=\min \left\{C_{s}+B_{s}, T\right\} \leq C_{t} \leq k$ and let

$$
\left.\Pi^{-1}\right|_{k}(j)= \begin{cases}\Pi^{-1}(j) & j \in \Pi([k]) \\ N & \text { otherwise }\end{cases}
$$

The relationship (33) implies that

$$
\tilde{f}_{i_{s}}(y)=f_{\left.\Pi^{-1}\right|_{k}\left(i_{s}\right)}\left(U_{\leq k}^{\top} y\right)
$$

for all $y$ in the neighborhood of $x_{s}$.
The above discussion shows that (under $\mathfrak{Z} \Re_{t-1}$ ) the oracle response to ( $i_{s}, x_{s}$ ) is measurable with respect to $i_{s}, x_{s},\left.\Pi\right|_{k}$ and $U_{\leq k}$, for every $s<t$. Moreover, $\left(i_{1}, x_{1}\right)$ is measurable with respect to $\zeta$, and consequently the first oracle response and the second query $\left(i_{2}, x_{2}\right)$ are measurable with respect to $\zeta,\left.\Pi\right|_{k}$ and $U_{\leq k}$. Repeating this argument inductively shows that (under $\mathfrak{Z} \mathfrak{R}_{t-1}$ and $C_{t} \leq$ $k$ ) we can generate the entire query sequence $\left(i_{s}, x_{s}\right)_{s \leq t}$, as well as the oracle responses to all but the last query, from $\zeta,\left.\Pi\right|_{k}$ and $U_{\leq k}$, as claimed.

To show part (a) of the lemma, note that with access to the entire matrix $U$, after generating queries $x_{1}, \ldots, x_{s}$ we can test whether $\mathfrak{Z} \Re_{s}$ holds. Moreover, the knowledge of $\left.\Pi\right|_{k}$ suffices to test (for any query sequence $i_{1}, i_{2}, \ldots$ ) whether $C_{t} \geq k$ for every $t$. Therefore, using $\zeta,\left.\Pi\right|_{k}$ and $U_{\leq k}$ we can iteratively compute $\left(i_{s}, x_{s}\right)_{s \leq t}$ until arriving at an iterate $t$ where either (a) $C_{t}<k$ but $\mathfrak{Z} \mathfrak{R}_{t}$
does not hold (which implies that $\mathfrak{Z R}_{\Theta_{k}}$ fails as well), or (b) $C_{t}=k$ (in which case we have found $\Theta_{k}$ ). In either case, we know the value of $\Theta_{k} \mathbb{I}\left\{\mathfrak{Z} \mathfrak{R}_{\Theta_{k}}\right\}$.

Finally, to show part (b) of the lemma, note that with full knowledge of $\Pi$ we can track the values of $C_{1}, \ldots, C_{t}$ for any sequence of queries $i_{1}, \ldots, i_{t-1}$. Therefore, given $\mathfrak{Z} \Re_{t-1}$, we may use $\zeta,\left.\Pi\right|_{k}$ and $U_{\leq k}$ to iteratively generate queries until either (a) we arrive at an iteration $s<t$ where $C_{s}>k$, in which case $\mathbb{I}\left\{C_{t}=k\right\}=0$, or (b) we successfully generate iteration $t-1$ in which case we can compute $C_{t}$. In either case, we know the value of $\mathbb{I}\left\{C_{t}=k\right\}$.

Lemma 28 For every $k \leq T$, and $j \geq 0$, we have

$$
\mathbb{P}\left(\Delta_{k} \leq j, \mathfrak{Z} \mathfrak{R}_{\Theta_{k}}|\zeta, \Pi|_{k-1}, U\right) \leq \frac{j}{N-k+1} .
$$

Proof Suppose the algorithm could access the random permutation $\Pi$ via an alternative oracle that, when queried at index $i \in[N]$ returns the number $\Pi^{-1}(i)$. We say that the algorithm succeeds if one of the first $j$ queries is $\Pi(k)$ (so that the oracle returns $k$ ). Given $\left.\Pi\right|_{k-1}$, the random variable $\Pi(k)$ is uniformly distributed over $\mathcal{I}=[N] \backslash \Pi([k-1])$. Therefore, for any set $\mathcal{J}$ of $j$ queries, we have

$$
\mathbb{P}\left(\Pi(k) \in \mathcal{J}|\Pi|_{k-1}\right)=\frac{\mathbb{E}|\mathcal{J} \cap \mathcal{I}|}{|\mathcal{I}|} \leq \frac{\mathbb{E}|\mathcal{J}|}{|\mathcal{I}|}=\frac{j}{N-k+1}
$$

and so for every algorithm the probability of success is at most $j /(N-k+1)$.
Now return to the original problem and the original oracle, and note that for $\Delta_{k} \leq j$ to hold, we must have $\Pi^{-1}\left(i_{t}\right)=k$ for some $t$ in $\left\{\Theta_{k-1}, \ldots, \Theta_{k-1}+j-1\right\}$, corresponding to "success" in the problem described above. Moreover, when the event $\mathfrak{Z} \mathfrak{R}_{\Theta_{k}}$ holds, Lemma 27 guarantees that the oracle responses to the queries at iteration $1, \ldots, \Theta_{k-1}-1$ are deterministic functions of $\zeta, U$ and $\left.\Pi\right|_{k-1}$ (since $C_{\Theta_{k-1}}=k-1$ by definition). Consequently, when $\mathfrak{Z} \mathfrak{R}_{\Theta_{k-1}}$ holds, the true optimization algorithm operates with no more information that the alternative oracle setting described above, giving the claimed probability bound.

Lemma 29 We have

$$
\mathbb{P}\left(\Theta_{T} \leq \frac{1}{16} N\left(T-\log \frac{2}{\delta}\right), \mathfrak{Z} \mathfrak{R}_{N T}\right) \leq \frac{\delta}{2}
$$

Proof Let $T_{ \pm \delta}:=\frac{1}{2}\left(T \pm \log \frac{2}{\delta}\right)$. We begin with a sequence of straightforward inequalities:

$$
\begin{align*}
\mathbb{P}\left(\Theta_{T}\right. & \left.\leq \frac{1}{8} N T_{-\delta}, \mathfrak{Z} \Re_{N T}\right) \leq \mathbb{P}\left(\Theta_{T} \leq \frac{1}{8} N T_{-\delta}, \mathfrak{Z} \Re_{\Theta_{T}}\right) \\
& \leq \mathbb{P}\left(\sum_{k \in[T]} \mathbb{I}\left\{\Delta_{k}>\frac{1}{8} N\right\} \leq T_{-\delta}, \mathfrak{Z} \Re_{\Theta_{T}}\right) \\
& =\mathbb{P}\left(\sum_{k \in[T]} \mathbb{I}\left\{\Delta_{k} \leq \frac{1}{8} N\right\}>T_{+\delta}, \mathfrak{Z} \Re_{\Theta_{T}}\right) \\
& \leq \mathbb{P}\left(\sum_{k \in[T]} \mathbb{I}\left\{\Delta_{k} \leq \frac{1}{8} N, \mathfrak{Z} \Re_{\Theta_{k}}\right\}>T_{+\delta}\right) \\
& \leq e^{-2 T_{+\delta} \mathbb{E}}\left[e^{2 \sum_{k \in[T]} \mathbb{I}\left\{\Delta_{k} \leq \frac{1}{8} N, \mathfrak{Z} \Re_{\Theta_{k}}\right\}}\right], \tag{34}
\end{align*}
$$

where the last transition is an application of the Chernoff bound (with parameter $\lambda=2$ ).
By Lemma 27, the random variable $\mathbb{I}\left\{\Delta_{k} \leq \frac{1}{8} N, \vec{Z} \Re_{\Theta_{k}}\right\}$ is measurable with respect to $\zeta,\left.\Pi\right|_{k}$ and $U$ since it is a function of $\Theta_{k} \mathbb{I}\left\{\mathfrak{Z}_{\Theta_{k}}\right\}$ and $\Theta_{k-1} \mathbb{I}\left\{\mathfrak{Z} \mathfrak{R}_{\Theta_{k-1}}\right\}$. Therefore,

$$
\begin{aligned}
& \mathbb{E}\left[e^{2 \sum_{k \in[T]} \mathbb{I}\left\{\Delta_{k} \leq \frac{1}{8} N, 3 \Re_{\Theta_{k}}\right\}}\right] \\
& \quad=\mathbb{E}\left[\left(1+\left(e^{2}-1\right) \mathbb{P}\left(\Delta_{T} \leq \frac{1}{8}, \mathfrak{Z} \Re_{\Theta_{T}}|\zeta, \Pi|_{T-1}, U\right)\right) e^{2 \sum_{k \in[T-1]} \mathbb{I}\left\{\Delta_{k} \leq \frac{1}{8} N, 3 \Re_{\Theta_{k}}\right\}}\right] \\
& \quad \leq\left(1+\frac{1}{4}\left(e^{2}-1\right)\right) \mathbb{E}\left[e^{2 \sum_{k \in[T-1]} \mathbb{I}\left\{\Delta_{k} \leq \frac{1}{8} N, 3 \Re_{\Theta_{k}}\right\}}\right],
\end{aligned}
$$

where the inequality follows from Lemma 28 with $j=N / 4$ and $k=T \leq N / 2$. Noting that $1+\frac{1}{4}\left(e^{2}-1\right) \leq e$ and iterating this argument, we conclude that

$$
\mathbb{E}\left[e^{2 \sum_{k \in[T]} \mathbb{I}\left\{\Delta_{k} \leq \frac{1}{8} N, 3 \Re_{\Theta_{k}}\right\}}\right] \leq e^{T} .
$$

Substituting this bound into eq. (34) and recalling that $e^{-2 T_{+\delta}}=\frac{\delta}{2} e^{-T}$ concludes the proof.
Lemma 30 We have

$$
\mathbb{P}\left(\mathfrak{Z}_{N T}^{c}\right) \leq \frac{\delta}{2} .
$$

Proof By definition, we have

$$
\begin{equation*}
\mathbb{P}\left(\mathfrak{Z} \mathfrak{R}_{N T}^{c}\right)=\sum_{t \in[N T]} \mathbb{P}\left(p_{t}>C_{t}, \mathfrak{Z} \Re_{t-1}\right) . \tag{35}
\end{equation*}
$$

We fix $t \leq N T$ and argue that $\mathbb{P}\left(p_{t}>C_{t}, \mathfrak{Z}_{t-1}\right) \leq \frac{\delta}{2 N T}$. Let $k \leq T$; by Lemma 27 events $\mathfrak{Z} \Re_{t-1}$ and $C_{t}=k$ hold, we have $x_{t}=h\left(\zeta, \Pi, U_{\leq k}\right)$ for some measurable functions $h$. Therefore, we have

$$
\begin{aligned}
& \mathbb{P}\left(p_{t}>C_{t}, \mathfrak{Z} \Re_{t-1} \mid C_{t}=k, \zeta, \Pi, U_{\leq k}\right) \stackrel{(i)}{=} \mathbb{P}\left(p_{t}>k, \mathfrak{Z} \Re_{t-1} \mid \zeta, \Pi, U_{\leq k}\right) \\
= & \mathbb{P}\left(\operatorname{prog}_{\alpha}\left(U^{\top} x_{t}\right)>k, \mathfrak{Z} \Re_{t-1} \mid \zeta, \Pi, U_{\leq k}\right) \\
= & \mathbb{P}\left(\operatorname{prog}_{\alpha}\left(U^{\top} h\left(\zeta, \Pi, U_{\leq k}\right)\right)>k, \mathfrak{Z} \Re_{t-1} \mid \zeta, \Pi, U_{\leq k}\right) \\
\leq & \mathbb{P}\left(\operatorname{prog}_{\alpha}\left(U^{\top} h\left(\zeta, \Pi, U_{\leq k}\right)\right)>k \mid \zeta, \Pi, U_{\leq k}\right) \stackrel{(i i)}{\leq} \sum_{j=k+1}^{T} \mathbb{P}\left(\left|u_{j}^{\top} h\left(\zeta, \Pi, U_{\leq k}\right)\right|>\alpha \mid \zeta, \Pi, U_{\leq k}\right),
\end{aligned}
$$

where $(i)$ follows from part (b) of Lemma 27, and (ii) follows from the definition of $\operatorname{prog}_{\alpha}$ and a union bound, where we write $u_{j}$ for the $j$ th column of $U$.

Conditionally on $\zeta, \Pi, U_{\leq k}, u_{j}$ is uniform over the unit ball in the subspace of $\mathbb{R}^{d}$ orthogonal to $U_{\leq k}$. Moreover, the magnitude of the projection of $h\left(\zeta, \Pi, U_{\leq k}\right)$ to that subspace can be at most 1 , since $x_{t}$ is a unit vector. Therefore, the probability that $\left|u_{j}^{\top} h\left(\zeta, \Pi, U_{\leq k}\right)\right|>\alpha$ holds is at most the probability that a coordinate of a uniform unit vector in $\mathbb{R}^{d-k}$ has magnitude greater than $\alpha$. By standard concentration of measure arguments, this probability is at most $2 \exp \left(-\frac{d-k+1}{2 \alpha^{2}}\right)$. By our choice of $d$ (and recalling $k \leq T$ ), we have that

$$
\mathbb{P}\left(p_{t}>C_{t}, \mathfrak{Z} \Re_{t-1} \mid C_{t}=k, \zeta, \Pi, U_{\leq k}\right) \leq(T-k) \cdot \frac{\delta}{2 N T^{2}} \leq \frac{\delta}{2 N T}
$$

and substituting back into eq. (35) concludes the proof.

## D.2. Proof of Lemma 9

Lemma 9 For every $T, N \in \mathbb{N}$ and $\ell \geq 0$, such that $T \leq N$, we have that

1. The hard instance $\left(\hat{f}_{i}^{\{T, N, \ell\}}\right)_{i \in N}$ is an $\alpha_{T}$-robust $N$-element zero-chain.
2. The function $\hat{f}_{i}^{\{T, N, \ell\}}$ is 1-Lipschitz and $\ell$-smooth for every $i \in[N]$.
3. For $x \in \mathbb{R}^{d}$ with $\operatorname{prog}_{\alpha_{T}}(x)<T$, the objective $\hat{F}_{\max }^{\{T, N, \ell\}}(x)=\max _{i \in[N]} \hat{f}_{i}^{\{T, N, \ell\}}(x)$ satisfies

$$
\hat{F}_{\max }^{\{T, N, \ell\}}(x)-\min _{x_{\star} \in \mathbb{B}_{1}(0)} \hat{F}_{\max }^{\{T, N, \ell\}}\left(x_{\star}\right) \geq \psi_{\alpha_{T}, \ell}\left(\frac{3}{8 T^{3 / 2}}\right) \geq \min \left\{\frac{1}{8 T^{3 / 2}}, \frac{\ell}{32 T^{3}}\right\}
$$

Proof To see why part 1 holds, fix $x \in \mathbb{R}^{T}$ and let $p=\operatorname{prog}_{\alpha_{T}}(x)$. First, we have $\hat{f}_{i}^{\{T, N, \ell\}}(x)=$ $\hat{f}_{i}^{\{T, N, \ell\}}\left(x_{[\leq i]}\right)$ for every $i$ and $x$, which immediately gives the first two cases of eq. (11). Second, when $i>p+1$, we have $\left|x_{[i]}-x_{[i-1]}\right|<2 \alpha_{T}$ and therefore $\hat{f}_{i}^{\{T, N, \ell\}}(x)=\psi_{\alpha_{T}, \ell}(t) \mathbb{I}\{i \leq N\}$ for some $|t|<\alpha$. Consequently $\hat{f}_{i}^{\{T, N, \ell\}}$ is identically zero in a neighborhood of $x$, giving the third and final case in eq. (11).

Part 2 is immediate because $\hat{f}_{i}^{\{T, N, \ell\}}$ is a composition of a 1 -Lipschitz and $\ell$-smooth function with the linear transformation $\left(x_{[i]}-x_{[i-1]}\right) / 2$ which has operator norm smaller than 1 .

Finally, to see part 3, first note that the global minimum of $\hat{F}_{\max }^{\{T, N, \ell\}}$ satisfies $\hat{F}_{\max }^{\{T, N, \ell\}}\left(x_{\star}\right)=0$ and $\left(x_{\star}\right)_{[i]}=\frac{1}{\sqrt{T}}$ for every $i \leq T$ (and therefore has unit norm). Consider any $x$ for which $\operatorname{prog}_{\alpha_{T}}(x)<T$, so that $x_{[T]} \leq \alpha_{T} \leq \frac{1}{\sqrt{T}}$. We have

$$
\max _{i \leq T}\left|x_{[i-1]}-x_{[i]}\right| \geq \frac{1}{T} \sum_{i \in[T]}\left|x_{[i-1]}-x_{[i]}\right| \geq \frac{1}{T}\left|x_{[0]}-x_{[T]}\right| \geq \frac{1}{T}\left(\frac{1}{\sqrt{T}}-\alpha_{T}\right) \geq \frac{3}{4 T^{3 / 2}}
$$

Since $\psi_{\alpha_{T}, \ell}(t)$ is non-decreasing in $|t|$, we have $\hat{F}_{\max }^{\{T, N, \ell\}}(x)=\psi_{\alpha_{T}, \ell}\left(\frac{1}{2} \max _{i \leq T}\left|x_{[i-1]}-x_{[i]}\right|\right)$, and consequently $\hat{F}_{\max }^{\{T, N, \ell\}}(x) \geq \psi_{\alpha_{T}, \ell}\left(\frac{3}{8 T^{3 / 2}}\right)$. To obtain the final bound, we observe that $\psi_{\alpha, \ell}(t) \geq$ $\min \left\{\frac{1}{2}(t-\alpha), \frac{\ell}{2}(t-\alpha)^{2}\right\}$ for any $t \geq \alpha$.

## D.3. Proof of Theorem 10

Theorem 10 Let $L_{f}, L_{g}, R>0, \epsilon<L_{f} R \wedge L_{g} R^{2}, N \in \mathbb{N}$ and $\delta \in(0,1)$. Then, for any (possibly randomized) algorithm there exists an $L_{f}$-Lipschitz and $L_{g}$-smooth functions $\left(f_{i}\right)_{i \in[n]}$ with domain $\mathbb{B}_{R}^{d}(0)$ for $d=O\left(\left[\left(\frac{L_{f} R}{\epsilon}\right)^{2} \wedge\left(\frac{L_{g} R^{2}}{\epsilon}\right)\right] \log \frac{N\left(L_{f} R \wedge L_{g} R^{2}\right)}{\epsilon}\right)$ such that with probability at least $\frac{1}{2}$ over the randomness of the algorithm, the first

$$
\begin{equation*}
\Omega\left(N\left[\left(\frac{L_{f} R}{\epsilon}\right)^{2 / 3} \wedge\left(\frac{L_{g} R^{2}}{\epsilon}\right)^{1 / 3}\right]+\left[\left(\frac{L_{f} R}{\epsilon}\right)^{2} \wedge\left(\frac{N L_{g} R^{2}}{\epsilon}\right)^{1 / 2}\right]\right) \tag{13}
\end{equation*}
$$

queries of the algorithm are all $\epsilon$-suboptimal for $F_{\max }(x)=\max _{i \in[N]} f_{i}(x)$.

Proof We first show that an $\Omega\left(N\left[\left(\frac{L_{f} R}{\epsilon}\right)^{2 / 3} \wedge\left(\frac{L_{g} R^{2}}{\epsilon}\right)^{1 / 3}\right]\right)$ lower bound follows from our construction in the previous section. Fix any $T>1$ and $\ell \geq 0$ and let $d=\left\lceil T+4 \alpha_{T}^{-2} \log 4 N T\right\rceil$ with $\alpha_{T}=1 /\left(4 T^{3 / 2}\right)$. Let $\Pi$ be random permutation of $[N]$ and let $U$ be drawn uniformly from the set of $d \times T$ orthogonal matrices. For $i \in[N]$, let

$$
\tilde{f}_{i}(x)=\hat{f}_{\Pi^{-1}(i)}^{\{T, N, \ell\}}\left(U^{\top} x\right)
$$

Then, by Proposition 8 and Lemma 9.1, any optimization algorithm interacting with $\left(\tilde{f}_{i}\right)_{i \leq N}$ satisfies with probability as least $1 / 2$ that $\operatorname{prog}_{\alpha_{T}}\left(U^{\top} x_{i}\right)<T$ for every $i \leq \frac{1}{64} N T$.

Set

$$
T=\left\lfloor\frac{1}{5}\left[\left(\frac{L_{f} R}{\epsilon}\right)^{2 / 3} \wedge\left(\frac{L_{g} R^{2}}{\epsilon}\right)^{1 / 3}\right]\right\rfloor \text { and } \ell=\frac{L_{g} R}{L_{f}}
$$

We may assume $T \leq N / 2$ without loss of generality, since otherwise the second term in the lower bound dominates. Let

$$
f_{i}(x)=L_{f} R \tilde{f}_{i}(x / R)
$$

With these settings, we have that $f_{i}$ is both $L_{f}$-Lipschitz and $L_{g}$-smooth for every $i$ due to Lemma 9.2, the choice of $\ell$, and the fact that $U$ is orthogonal. Moreover, if $x_{1}, x_{2}, \ldots$ are the iterates of an algorithm interacting with a finite sum oracle for $\left(f_{i}\right)_{i \in[N]}$ then $x_{1} / R, x_{2} / R, \ldots$ are the iterates of an algorithm interacting with $\left(\tilde{f}_{i}\right)_{i \leq[N]}$. Therefore, by the above discussion, with probability at least $1 / 2$ the firstp

$$
\frac{1}{64} N T=\Omega\left(N\left[\left(\frac{L_{f} R}{\epsilon}\right)^{2 / 3} \wedge\left(\frac{L_{g} R^{2}}{\epsilon}\right)^{1 / 3}\right]\right)
$$

iterates of the algorithm satisfy $\operatorname{prog}_{\alpha_{T}}\left(U^{\top} x\right)<T$, and consequently, by Lemma 9.3, they are

$$
L_{f} R \min \left\{\frac{1}{8 T^{3 / 2}}, \frac{\ell}{32 T^{3}}\right\}=\min \left\{\frac{L_{f} R}{8 T^{3 / 2}}, \frac{L_{g} R^{2}}{32 T^{3}}\right\}>\epsilon
$$

suboptimal for $\max _{i \in[N]} f_{i}$. To conclude the linear in $N$ lower bound, we note that since the suboptimality bound holds with probability at least $1 / 2$ over $\Pi, U$ and the algorithm randomness, for every algorithm there must exist fixed $\Pi$ and $U$ for which the bound holds with probability $1 / 2$ over the randomness of the algorithm alone.

To show the remaining $\Omega\left(\left[\left(\frac{L_{f} R}{\epsilon}\right)^{2} \wedge\left(\frac{N L_{g} R^{2}}{\epsilon}\right)^{1 / 2}\right]\right)$ term in the lower bound, we recall the following classical result. For every $R^{\prime}>0$ there exists a distribution over functions $F: \mathbb{R}_{d} \rightarrow$ $\mathbb{R}$ with $d^{\prime}=O\left(\left[\left(\frac{L_{f} R^{\prime}}{\epsilon}\right)^{2} \wedge\left(\frac{L_{g} R^{\prime 2}}{\epsilon}\right)^{1 / 2}\right] \log \frac{N L_{f} R^{\prime}}{\epsilon}\right)$ that are $L_{g}$-Lipschitz, $L_{f}$-smooth and has a global minimizer with norm at most $R$, such that for any algorithm interacting with $F$, with probability at least $1-\frac{1}{2 N}$, the first

$$
T_{R^{\prime}}=\Omega\left(\left[\left(\frac{L_{f} R^{\prime}}{\epsilon}\right)^{2} \wedge\left(\frac{L g R^{\prime 2}}{\epsilon}\right)^{1 / 2}\right]\right)
$$

iterations are $\epsilon$ suboptimal for $F$. This result follows from smoothing Nemirovski's function; see Diakonikolas and Guzmán [14, Theorem 14 with $p=2$ and $\kappa=0,1]$. To strengthen the smooth term in the lower bound, we consider $N$ copies of this construction with $R^{\prime}=R / \sqrt{N}$ that operate on distinct coordinates, i.e., we let $f_{i}(x)=F\left(x_{\left[1+(i-1) d^{\prime}\right]}, \ldots, x_{\left[i d^{\prime}\right]}\right)$, so that the global minimizer of
$\max _{i} f_{i}(x)$, which consists of $N$ copies of the global minimizer of $F$, has norm at most $R^{\prime} \sqrt{N}=R$, and consequently we may constrain the domain to $\mathbb{B}_{N}^{d^{\prime} N}$ without decreasing the optimality gap of any point in the ball. For any algorithm interacting with $\left(f_{i}\right)_{i \in[N]}$, the first

$$
N T_{R / \sqrt{N}}=\Omega\left(\left[\left(\frac{L_{f} R}{\epsilon}\right)^{2} \wedge\left(\frac{N L g R^{2}}{\epsilon}\right)^{1 / 2}\right]\right)
$$

queries of the algorithm must select one of the $N$ components at most $T_{R / \sqrt{N}}$ times and therefore (with probability at least $\frac{1}{2}$ ) be $\epsilon$ suboptimal for at least one component, and hence for their maximum. This gives the sublinear in $N$ term of the lower bound. We remark that the "hard instance duplication" argument is at the core of existing lower bounds for finite sum optimization [39, 15]; our argument for proving the linear in $N$ lower bound is inherently different.

## D.4. Extension to unconstrained setting

While we state and prove our lower bound for optimization problems whose domain is a ball of radius $R$, it also extends to the case of unconstrained setting of our upper bounds. That is, when the domain is $\mathbb{R}^{d}$ and we are guaranteed a local minimizer exists in a radius of $R$ from the initial point. One way to show this extension is the technique of [14] where we replace $\hat{f}_{i}^{\{T, N, \ell\}}$ with

$$
\begin{equation*}
\bar{f}_{i}^{\{N, T, \ell\}}(x):=\min _{y \in \mathbb{R}^{d}}\left\{\max \left\{\hat{f}_{i}^{\{T, N, \ell\}}(y),\|y\|-1-\ell^{-1}\right\}+\frac{\ell}{2}\|y-x\|^{2}\right\} . \tag{36}
\end{equation*}
$$

The definition above consists of two modification: pairwise maximum with $\|\cdot\|-1$ and infimal convolution with $\frac{\ell}{2}\|\cdot\|^{2}$. The pairwise maximum guarantees that large norm queries cannot break the zero-chain progress control mechanism: responses to query points with norm $\Omega(1)$ will not depend on the random transformation $U$ at all, while for responses with norm $O(1)$ we can control the progress using the zero-chain structure of $\hat{f}\{T, N, \ell\}$ as before. The infimal convolution guarantees the function remains $\ell$ smooth (and also does not increase the Lipschitz constant).

Finally, it remains to check that the new construction still satisfies property 3 of Lemma 9 up to a constant. To see that it does, first note that the global minimizer of $\bar{F}_{\max }^{\{N, T, \ell\}}(x):=$ $\max _{i \in[N]} \bar{f}_{i}^{\{N, T, \ell\}}(x)$ still satisfies $x_{\star[i]}=1 / \sqrt{T}$ for all $i \leq T$, and that $\bar{F}_{\max }^{\{N, T, \ell\}}\left(x_{\star}\right)=0$. Moreover, we clearly have

$$
\bar{f}_{i}^{\{N, T, \ell\}}(x) \geq \min _{y \in \mathbb{R}^{d}}\left\{\psi_{\alpha_{T}, \ell}(y / 2)+\frac{\ell}{2}\left\|y-\frac{x_{[i]}-x_{[i-1]}}{2}\right\|^{2}\right\}:=\tilde{\psi}\left(\frac{x_{[i]}-x_{[i-1]}}{2}\right)
$$

and it is not hard to verify that $\tilde{\psi}(t) \geq \psi_{\alpha_{T}, \ell}(c t)$ for some constant $c>0\left(\right.$ in fact, $\left.\tilde{\psi}(t)=\psi_{\alpha_{T}, c \ell}(t)\right)$.


[^0]:    3. Exact-gradient primal-dual methods such as mirror-prox [28] and dual-extrapolation [32] have complexity guarantees scaling as $\widetilde{O}\left(N \epsilon^{-1}\right)$ under the stronger smoothness assumption $L_{g}=O(1)$ [cf. 8, Section 5.2.4].
[^1]:    4. We note that $\lambda$ in our notation corresponds to $1 / \lambda$ in the notation of [12].
[^2]:    5. We remark that $g_{i}(x)=e^{\left(f_{i}(x)-f_{i}(\bar{x})\right) / \epsilon^{\prime}} \nabla f_{i}(x)$ is also nearly unbiased for $F_{\text {smax }, \epsilon}$ in the sense that $\mathbb{E} g_{i}(x)=$ $Z(x) \nabla F_{\text {smax }, \epsilon}(x)$ for some $Z(x)$ that is close to 1 when inside $\mathbb{B}_{r_{\epsilon}}(\bar{x})$. Estimators of this form suffice for SGD, but are less amenable to variance reduction.
