# A Gradient Sampling Method With Complexity Guarantees for Lipschitz Functions in High and Low Dimensions* 

Damek Davis<br>Cornell University<br>dsd95@cornell.edu

Dmitriy Drusvyatskiy<br>University of Washington, Seattle<br>ddrusv@uw.edu

Yin Tat Lee<br>University of Washington, Seattle<br>yintat@uw.edu

Swati Padmanabhan<br>University of Washington, Seattle<br>pswati@uw.edu

Guanghao Ye<br>Massachussetts Institute of Technology<br>ghye@mit.edu


#### Abstract

Zhang, Lin, Jegelka, Sra, and Jadbabaie [1] introduce a method for minimizing Lipschitz functions with an efficiency guarantee of $O\left(\varepsilon^{-4}\right)$. Their method is a novel modification of Goldstein's classical subgradient method. Their work, however, makes use of a nonstandard subgradient oracle model and requires the function to be directionally differentiable. Our first contribution in this paper is to show that both of these assumptions can be dropped by simply adding a small random perturbation in each step of their algorithm. The resulting method works on any Lipschitz function whose value and gradient can be evaluated at points of differentiability. Our second contribution is a new cutting plane algorithm that achieves better efficiency in low dimensions: $O\left(d \varepsilon^{-3}\right)$ for Lipschitz functions and $O\left(d \varepsilon^{-2}\right)$ for those that are weakly convex.


## 1 Introduction

The subgradient method [2] is a classical procedure for minimizing a nonsmooth Lipschitz function $f$ on $\mathbb{R}^{d}$. Starting from an initial iterate $x_{0}$, the method computes

$$
\begin{equation*}
x_{t+1}=x_{t}-\alpha_{t} v_{t} \text { where } v_{t} \in \partial f\left(x_{t}\right) \tag{1}
\end{equation*}
$$

Here, the positive sequence $\left\{\alpha_{t}\right\}_{t \geq 0}$ is user-specified, and the set $\partial f$ is the Clarke subdifferential [3, 4],

$$
\partial f(x)=\operatorname{conv}\left\{\lim _{i \rightarrow \infty} \nabla f\left(x_{i}\right): x_{i} \rightarrow x, x_{i} \in \operatorname{dom}(\nabla f)\right\}
$$

In classical circumstances, the subdifferential reduces to familiar objects: for example, when $f$ is $C^{1}$-smooth at $x$, the subdifferential $\partial f(x)$ comprises of only the gradient $\nabla f(x)$, while for convex functions, it reduces to the subdifferential in the sense of convex analysis.
For functions $f$ that are weakly convex - a broad class of functions first introduced in English in [5] - the limit points $\bar{x}$ of the subgradient method are known to be first-order critical, meaning

[^0]$0 \in \partial f(\bar{x})$. Recall that a function $f$ is called $\rho$-weakly convex if the quadratically perturbed function $x \mapsto f(x)+\frac{\rho}{2}\|x\|^{2}$ is convex. In particular, convex and smooth functions are weakly convex [6]. Going beyond asymptotic guarantees, finite-time complexity estimates are known for smooth, convex, or weakly convex problems [7-14].
Modern machine learning, however, has witnessed the emergence of problems far beyond the weakly convex problem class. Indeed, tremendous empirical success has been recently powered by industry-backed solvers, such as Google's TensorFlow and Facebook's PyTorch, which routinely train nonsmooth nonconvex deep networks via (stochastic) subgradient methods. Despite a vast body of work on the asymptotic convergence of subgradient methods for nonsmooth nonconvex problems [15-19], no finite-time nonasymptotic convergence rates were known outside the weakly convex setting until recently, with Zhang, Lin, Jegelka, Sra, and Jadbabaie [1] making a big leap forward towards this goal.

In particular, restricting themselves to the class of Lipschitz and directionally differentiable functions, [1] developed an efficient algorithm motivated by Goldstein's conceptual subgradient method [20]. Moreover, this was recently complemented by [21] with lower bounds for finding near-approximatestationary points for nonconvex nonsmooth functions.
While a significant breakthrough in both result and technique, one crucial limitation of [1] is that their complexity guarantees and algorithm use a nonstandard first-order oracle whose validity is unclear in examples. To elaborate, their algorithm requires the following oracle access: given $x, u \in \mathbb{R}^{d}$ solve the auxiliary convex feasibility problem:

$$
\begin{equation*}
\text { find } g \in \partial f(x) \text { subject to }\langle g, u\rangle=f^{\prime}(x, u) \tag{2}
\end{equation*}
$$

The first issue with this oracle is that no general recipe exists for representing the full subdifferential $\partial f(x)$ analytically, and evaluating even an arbitrary element of the subdifferential can be highly non-trivial [22, 23]. Moreover, $\partial f(x)$ could be a very complicated set, e.g., for a deep ReLU neural network, the subdifferential is a polyhedron with a potentially huge number of facets, making the complexity of (2) unclear.
Further, [1] claim that for a composition of directionally differentiable functions with a closed-form directional derivative for each function, we can find the desired $g$ by the chain rule. While the chain rule does compute the directional derivative $f^{\prime}(x, u)$, to the best of our knowledge, this does not translate to solving (2). This is owing to the crucial fact that the chain rule (and sum rule) can easily fail for the computation of the subdifferential ${ }^{2}$ (although these are indeed valid for the directional derivative of a composition of directionally differentiable functions). We believe that this could potentially render the oracle of [1] computationally intractable.
Finally, we are unaware of other optimization algorithms imposing this oracle model. Therefore, at face value, the convergence guarantees of [1] are not comparable to those of others.

### 1.1 Our results

Weakly convex optimization via a standard oracle. Our first contribution is to recover the complexity result of [1] under a much weaker assumption: specifically, we replace the non-standard assumption in (2) with a standard first-order oracle model. We show (Section 2) that a simple, yet critical, modification of the algorithm of [1], wherein one simply adds a small random perturbation in each iteration, works for any Lipschitz function assuming only an oracle that can compute gradients and function values at almost every point of $\mathbb{R}^{d}$ in the sense of Lebesgue measure. In particular, such oracles arise from automatic differentiation schemes routinely used in deep learning [24, 19]. Our end result is a randomized algorithm for minimizing any $L$-Lipschitz function that outputs a $(\delta, \epsilon)$-stationary point (Definition 1) after using at most $\widetilde{\mathcal{O}}\left(\frac{\Delta L^{2}}{\epsilon^{3} \delta} \log (1 / \gamma)\right)^{3}$ gradient and function evaluations. Here $\Delta$ is the initial function gap and $\gamma$ is the failure probability.

[^1]In light of the above modifications, our algorithm is implementable in the many important settings like deep neural networks that [1] is not. Along the way, we also simplify their proof techniques by providing a geometric viewpoint of the algorithm.

We would like to highlight the concurrent work of Tian, Zhou, and Man-Cho So [25], which obtains this part of our result with a very similar technique.

Improved complexity in low dimensions. Having obtained the result of [1] within the standard first-order oracle model, we then proceed to investigate the following question.

## Can we improve the efficiency of the algorithm in low dimensions?

In addition to being natural from the viewpoint of complexity theory, this question is well-grounded in applications. For instance, numerous problems in control theory involve minimization of highly irregular functions of a small number of variables. We refer the reader to the survey [26, Section 6] for an extensive list of examples, including Chebyshev approximation by exponential sums, spectral and pseudospectral abscissa minimization, maximization of the "distance to instability", and fixed-order controller design by static output feedback. We note that for many of these problems, the gradient sampling method of [26] is often used. Despite its ubiquity in applications, the gradient sampling method does not have finite-time efficiency guarantees. The algorithms we present here offer an alternative approach with a complete complexity theory.

The second contribution of our paper is an affirmative answer to the highlighted question. We present a novel algorithm that uses $\widetilde{\mathcal{O}}\left(\frac{\Delta L d}{\epsilon^{2} \delta} \log (1 / \gamma)\right)$ calls to our (weaker) oracle. Thus we are able to trade off the factor $L \epsilon^{-1}$ with $d$. Further, if the function is $\rho$-weakly convex, the complexity improves to $\widetilde{\mathcal{O}}\left(\frac{\Delta d}{\epsilon \delta} \log (\rho)\right)$, which matches the complexity in $\delta=\epsilon$ of gradient descent for smooth minimization. Strikingly, the dependence on the weak convexity constant $\rho$ is only logarithmic.
To put this contribution in perspective, assume for now $\delta=\epsilon$ : then, our algorithm's dependence on $\epsilon$ in the case of Lipschitz, weakly convex functions is likely optimal in low dimensions, following a conjecture by Bubeck and Mikulincer [27] on the optimality of gradient descent for smooth optimization in dimension $d=\log \left(\frac{1}{\epsilon}\right)$ (thus matching the lower bound by Carmon, Duchi, Hinder, and Sidford [28]). Aside from possible optimality, the logarithmic dependence on smoothness/weak convexity exhibited by our iteration complexity is a significant improvement over the prior result of either $O\left(1 / \epsilon^{4}\right)$ by [1] or Nemirovski and Yudin's rate of $O\left(1 / \epsilon^{2}\right)$ with a polynomial dependence on smoothness. In the process, we also show that the minimal-norm element of the Goldsteinsubdifferential in low dimensions can be found in time $O(\log (1 / \epsilon))$, thus settling a question open since the 70 s .

Techniques. The main idea underlying our improved dependence on $\epsilon$ in low dimensions is outlined next. The algorithm of [1] comprises of an outer loop with $\mathcal{O}\left(\frac{\Delta}{\epsilon \delta}\right)$ iterations, each performing either a decrease in the function value or an ingenious random sampling step to update the descent direction. Our observation, central to improving the $\varepsilon$ dependence, is that the violation of the descent condition can be transformed into a gradient oracle for the problem of finding a minimal norm element of the Goldstein subdifferential. This gradient oracle may then be used within a cutting plane method, which achieves better $\varepsilon$ dependence at the price of a dimension factor (Section 3).

Limitations. One limitation of our work is that our second contribution does not immediately extend to the stochastic setting. We consider this to be an interesting open problem to resolve.

Notation. Throughout, we let $\mathbb{R}^{d}$ denote a $d$-dimensional Euclidean space equipped with a dot product $\langle\cdot, \cdot\rangle$ and the Euclidean norm $\|x\|_{2}=\sqrt{\langle x, x\rangle}$. The symbol $\mathbb{B}_{r}(x)$ denotes an open Euclidean ball of radius $r>0$ around a point $x$. Throughout, we fix a function $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ that is $L$-Lipschitz, and let $\operatorname{dom}(\nabla f)$ denote the set of points where $f$ is differentiable-a full Lebesgue measure set by Rademacher's theorem. The symbol $f^{\prime}(x, u) \stackrel{\text { def }}{=} \lim _{\tau \downarrow 0} \tau^{-1}(f(x+\tau u)-f(x))$ denotes the directional derivative of $f$ at $x$ in direction $u$, whenever the limit exists.

## 2 Interpolated normalized gradient descent

In this section, we describe the results in [1] and our modified subgradient method that achieves finite-time guarantees in obtaining $(\delta, \epsilon)$-stationarity for an $L$-Lipschitz function $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$. The main construction we use is the Goldstein subdifferential [20].
Definition 1 (Goldstein subdifferential). Consider a locally Lipschitz function $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$, a point $x \in \mathbb{R}^{d}$, and a parameter $\delta>0$. The Goldstein subdifferential of $f$ at $x$ is the set

$$
\partial_{\delta} f(x) \stackrel{\text { def }}{=} \operatorname{conv}\left(\bigcup_{y \in \mathbb{B}_{\delta}(x)} \partial f(y)\right)
$$

A point $x$ is called $(\delta, \epsilon)$-stationary if $\operatorname{dist}\left(0, \partial_{\delta} f(x)\right) \leq \epsilon$.
Thus, the Goldstein subdifferential of $f$ at $x$ is the convex hull of all Clarke subgradients at points in a $\delta$-ball around $x$. Famously, [20] showed that one can significantly decrease the value of $f$ by taking a step in the direction of the minimal norm element of $\partial_{\delta} f(x)$. Throughout the rest of the section, we fix $\delta \in(0,1)$ and use the notation

$$
\begin{equation*}
\hat{g} \stackrel{\text { def }}{=} g /\|g\|_{2} \text { for any nonzero vector } g \in \mathbb{R}^{d} \tag{3}
\end{equation*}
$$

Theorem 1 ([20]). Fix a point $x$, and let $g$ be a minimal norm element of $\partial_{\delta} f(x)$. Then as long as $g \neq 0$, we have $f(x-\delta \hat{g}) \leq f(x)-\delta\|g\|_{2}$.

Theorem 1 immediately motivates the following conceptual descent algorithm:

$$
\begin{equation*}
x_{t+1}=x_{t}-\delta \hat{g}_{t}, \text { where } g_{t} \in \underset{g \in \partial_{\delta} f(x)}{\operatorname{argmin}}\|g\|_{2} \tag{4}
\end{equation*}
$$

In particular, Theorem 1 guarantees that, defining $\Delta \stackrel{\text { def }}{=} f\left(x_{0}\right)-\min f$, the approximate stationarity condition

$$
\min _{t=1, \ldots, T}\left\|g_{t}\right\|_{2} \leq \epsilon \text { holds after } T=\mathcal{O}\left(\frac{\Delta}{\delta \epsilon}\right) \text { iterations of (4). }
$$

Evaluating the minimal norm element of $\partial_{\delta} f(x)$ is impossible in general, and therefore the descent method described in (4) cannot be applied directly. Nonetheless it serves as a guiding principle for implementable algorithms. Notably, the gradient sampling algorithm [29] in each iteration forms polyhedral approximations $K_{t}$ of $\partial_{\delta} f\left(x_{t}\right)$ by sampling gradients in the ball $\mathbb{B}_{\delta}(x)$ and computes search directions $g_{t} \in \operatorname{argmin}_{g \in K_{t}}\|g\|_{2}$. These gradient sampling algorithms, however, have only asymptotic convergence guarantees [26].
The recent paper [1] remarkably shows that for any $x \in \mathbb{R}^{d}$ one can find an approximate minimal norm element of $\partial_{\delta} f(x)$ using a number of subgradient computations that is independent of the dimension. The idea of their procedure is as follows. Suppose that we have a trial vector $g \in \partial_{\delta} f(x)$ (not necessarily a minimal norm element) satisfying

$$
\begin{equation*}
f(x-\delta \hat{g}) \geq f(x)-\frac{\delta}{2}\|g\|_{2} . \tag{5}
\end{equation*}
$$

That is, the decrease in function value is not as large as guaranteed by Theorem 1 for the true minimal norm subgradient. One would like to now find a vector $u \in \partial_{\delta} f(x)$ so that the norm of some convex combination $(1-\lambda) g+\lambda u$ is smaller than that of $g$. A short computation shows that this is sure to be the case for all small $\lambda>0$ as long as $\langle u, g\rangle \leq\|g\|_{2}^{2}$. The task therefore reduces to:

$$
\text { find some } u \in \partial_{\delta} f(x) \quad \text { satisfying } \quad\langle u, g\rangle \leq\|g\|_{2}^{2}
$$

The ingenious idea of [1] is a randomized procedure for establishing exactly that in expectation. Namely, suppose for the moment that $f$ happens to be differentiable along the segment $[x, x-\delta \hat{g}]$; we will revisit this assumption shortly. Then the fundamental theorem of calculus, in conjunction with (5), yields

$$
\begin{equation*}
\frac{1}{2}\|g\|_{2} \geq \frac{f(x)-f(x-\delta \hat{g})}{\delta}=\frac{1}{\delta} \int_{0}^{\delta}\langle\nabla f(x-\tau \hat{g}), \hat{g}\rangle d \tau \tag{6}
\end{equation*}
$$

Consequently, a point $y$ chosen uniformly at random in the segment $[x, x-\delta \hat{g}]$ satisfies

$$
\begin{equation*}
\mathbb{E}\langle\nabla f(y), g\rangle \leq \frac{1}{2}\|g\|_{2}^{2} \tag{7}
\end{equation*}
$$

Therefore the vector $u=\nabla f(y)$ can act as the subgradient we seek. Indeed, the following lemma shows that, in expectation, the minimal norm element of $[g, u]$ is significantly shorter than $g$. The proof is extracted from that of [1, Theorem 8].
Lemma 2 ([1]). Fix a vector $g \in \mathbb{R}^{d}$, and let $u \in \mathbb{R}^{d}$ be a random vector satisfying $\mathbb{E}\langle u, g\rangle<\frac{1}{2}\|g\|_{2}^{2}$. Suppose moreover that the inequality $\|g\|_{2},\|u\|_{2} \leq L$ holds for some $L<\infty$. Then the minimalnorm vector $z$ in the segment $[g, u]$ satisfies:

$$
\mathbb{E}\|z\|_{2}^{2} \leq\|g\|_{2}^{2}-\frac{\|g\|_{2}^{4}}{16 L^{2}}
$$

Proof. Applying $\mathbb{E}\langle u, g\rangle \leq \frac{1}{2}\|g\|_{2}^{2}$ and $\|g\|_{2},\|u\|_{2} \leq L$, we have, for any $\lambda \in(0,1)$,

$$
\begin{aligned}
\mathbb{E}\|z\|_{2}^{2} \leq \mathbb{E}\|g+\lambda(u-g)\|_{2}^{2} & =\|g\|_{2}^{2}+2 \lambda \mathbb{E}\langle g, u-g\rangle+\lambda^{2} \mathbb{E}\|u-g\|_{2}^{2} \\
& \leq\|g\|_{2}^{2}-\lambda\|g\|_{2}^{2}+4 \lambda^{2} L^{2}
\end{aligned}
$$

Plugging in the value $\lambda=\frac{\|g\|_{2}^{2}}{8 L^{2}} \in(0,1)$ minimizes the right hand side and completes the proof.
The last technical difficulty to overcome is the requirement that $f$ be differentiable along the line segment $[g, u]$. This assumption is crucially used to obtain (6) and (7). To cope with this problem, [1] introduce extra assumptions on the function $f$ to be minimized and assume a nonstandard oracle access to subgradients.

We show, using Lemma 3, that no extra assumptions are needed if one slightly perturbs $g$.
Lemma 3. Let $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ be a Lipschitz function, and fix a point $x \in \mathbb{R}^{d}$. Then there exists a set $\mathcal{D} \subset \mathbb{R}^{d}$ of full Lebesgue measure such that for every $y \in \mathcal{D}$, the line spanned by $x$ and $y$ intersects $\operatorname{dom}(\nabla f)$ in a full Lebesgue measure set in $\mathbb{R}$. Then, for every $y \in \mathcal{D}$ and all $\tau \in \mathbb{R}$, we have

$$
f(x+\tau(y-x))-f(x)=\int_{0}^{\tau}\langle\nabla f(x+s(y-x)), y-x\rangle d s
$$

Proof. Without loss of generality, we may assume $x=0$ and $f(x)=0$. Rademacher's theorem guarantees that $\operatorname{dom}(\nabla f)$ has full Lebesgue measure in $\mathbb{R}^{d}$. Fubini's theorem then directly implies that there exists a set $\mathcal{Q} \subset \mathbb{S}^{d-1}$ of full Lebesgue measure within the sphere $\mathbb{S}^{d-1}$ such that for every $y \in \mathcal{Q}$, the intersection $\mathbb{R}_{+}\{y\} \cap(\operatorname{dom}(\nabla f))^{c}$ is Lebesgue null in $\mathbb{R}$. It follows immediately that the set $\mathcal{D}=\{\tau y: \tau>0, y \in Q\}$ has full Lebesgue measure in $\mathbb{R}^{d}$. Fix now a point $y \in \mathcal{D}$ and any $\tau \in \mathbb{R}_{+}$. Since $f$ is Lipschitz, it is absolutely continuous on any line segment and therefore
$f(x+\tau(y-x))-f(x)=\int_{0}^{\tau} f^{\prime}(x+s(y-x), y-x) d s=\int_{0}^{\tau}\langle\nabla f(x+s(y-x)), y-x\rangle d s$.
The proof is complete.
We now have all the ingredients to present a modification of the algorithm from [1], which, under a standard first-order oracle model, either significantly decreases the objective value or finds an approximate minimal norm element of $\partial_{\delta} f$.
The following theorem establishes the efficiency of Algorithm 1, and its proof is similar to that of [1, Lemma 13]. For completeness, we include the full proof in Appendix A.
Theorem 4. Let $\left\{g_{k}\right\}$ be generated by $\operatorname{MinNorm}(x, \delta, \epsilon)$. Fix an index $k \geq 0$, and define the stopping time $\tau \stackrel{\text { def }}{=} \inf \left\{k: f\left(x-\delta \hat{g}_{k}\right)<f(x)-\delta\left\|g_{k}\right\|_{2} / 4\right.$ or $\left.\left\|g_{k}\right\|_{2} \leq \epsilon\right\}$. Then, we have

$$
\mathbb{E}\left[\left\|g_{k}\right\|_{2}^{2} 1_{\tau>k}\right] \leq \frac{16 L^{2}}{16+k}
$$

An immediate consequence of Theorem 4 is that $\operatorname{MinNorm}(x, \delta, \epsilon)$ terminates with high-probability.

```
Algorithm 1 MinNorm \((x, \delta, \epsilon)\)
    Input. \(x, \delta>0\), and \(\epsilon>0\).
    Let \(k=0, g_{0}=\nabla f\left(\zeta_{0}\right)\) where \(\zeta_{0} \sim \mathbb{B}_{\delta}(x)\).
    while \(\left\|g_{k}\right\|_{2}>\epsilon\) and \(\frac{\delta}{4}\left\|g_{k}\right\|_{2} \geq f(x)-f\left(x-\delta \hat{g}_{k}\right)\) do
        Choose any \(r\) satisfying \(0<r<\left\|g_{k}\right\|_{2} \cdot \sqrt{1-\left(1-\frac{\left\|g_{k}\right\|_{2}^{2}}{128 L^{2}}\right)^{2}}\).
        Sample \(\zeta_{k}\) uniformly from \(\mathbb{B}_{r}\left(g_{k}\right)\).
        Choose \(y_{k}\) uniformly at random from the segment \(\left[x, x-\delta \widehat{\zeta}_{k}\right]\).
        \(g_{k+1}=\operatorname{argmin}_{z \in\left[g_{k}, \nabla f\left(y_{k}\right)\right]}\|z\|_{2}\).
        \(k=k+1\).
    end while
    Return \(g_{k}\).
```

Corollary 5. MinNorm $(x, \delta, \epsilon)$ terminates in at most $\left\lceil\frac{64 L^{2}}{\epsilon^{2}}\right\rceil \cdot\lceil 2 \log (1 / \gamma)\rceil$ iterations with probability at least $1-\gamma$, where we define the stopping time $\tau \stackrel{\text { def }}{=}$ $\inf \left\{k: f\left(x-\delta \hat{g}_{k}\right)<f(x)-\delta\left\|g_{k}\right\|_{2} / 4\right.$ or $\left.\left\|g_{k}\right\|_{2} \leq \epsilon\right\}$.

Combining Algorithm 1 with (4) yields Algorithm 2, with convergence guarantees summarized in Theorem 6, whose proof is identical to that of [1, Theorem 8].

```
Algorithm 2 Interpolated Normalized Gradient Descent \(\left(\operatorname{INGD}\left(x_{0}, T\right)\right)\)
    Initial \(x_{0}\), counter \(T\)
    for \(t=0, \ldots, T-1\) do
        \(g=\operatorname{MinNorm}\left(x_{t}\right)\)
        \(\triangleright\) Computational complexity \(\widetilde{\mathcal{O}}\left(L^{2} / \epsilon^{2}\right)\)
        Set \(x_{t+1}=x_{t}-\delta \hat{g}\)
                            \(\triangleright\) We define \(\hat{g}\) in (3)
    end for
    Return \(x_{T}\)
```

Theorem 6. Fix an initial point $x_{0} \in \mathbb{R}^{d}$, and define $\Delta=f\left(x_{0}\right)-\inf _{x} f(x)$. Set the number of iterations $T=\frac{4 \Delta}{\delta \epsilon}$. Then, with probability $1-\gamma$, the point $x_{T}=\operatorname{INGD}\left(x_{0}, T\right)$ satisfies $\operatorname{dist}\left(0, \partial_{\delta} f\left(x_{T}\right)\right) \leq \epsilon$ in a total of at most

$$
\left\lceil\frac{4 \Delta}{\delta \epsilon}\right\rceil \cdot\left\lceil\frac{64 L^{2}}{\epsilon^{2}}\right\rceil \cdot\left\lceil 2 \log \left(\frac{4 \Delta}{\gamma \delta \epsilon}\right)\right\rceil \quad \text { function-value and gradient evaluations. }
$$

In summary, the complexity of finding a point $x$ satisfying $\operatorname{dist}\left(0, \partial_{\delta} f(x)\right) \leq \epsilon$ is at most $\mathcal{O}\left(\frac{\Delta L^{2}}{\delta \epsilon^{3}} \log \left(\frac{4 \Delta}{\gamma \delta \epsilon}\right)\right)$ with probability $1-\gamma$. Using the identity $\partial f(x)=\lim \sup _{\delta \rightarrow 0} \partial_{\delta} f(x)$, this result also provides a strategy for finding a Clarke stationary point, albeit with no complexity guarantee. It is thus natural to ask whether one may efficiently find some point $x$ for which there exists $y \in \mathbb{B}_{\delta}(x)$ satisfying $\operatorname{dist}(0, \partial f(y)) \leq \epsilon$. This is exactly the guarantee of subgradient methods on weakly convex functions in [6]. [30] shows that for general Lipschitz functions, the number of subgradient computations required to achieve this goal by any algorithm scales with the dimension of the ambient space. Finally, we mention that the perturbation technique similarly applies to the stochastic algorithm of [1, Algorithm 2], yielding a method that matches their complexity estimate.

## 3 Faster INGD in low dimensions

In this section, we describe our modification of Algorithm 1 ("INGD") for obtaining improved runtimes in the low-dimensional setting. Our modified algorithm hinges on computations similar to (5), (6), and (7) except for the constants involved, and hence we explicitly state this setup. Given a vector $g \in \partial_{\delta} f(x)$, we say it satisfies the descent condition at $x$ if

$$
\begin{equation*}
f(x-\delta \hat{g}) \leq f(x)-\frac{\delta \epsilon}{3} \tag{8}
\end{equation*}
$$

Recall that Lemma 3 shows that for almost all $g$, we have

$$
f(x)-f(x-\delta \hat{g})=\int_{0}^{1}\langle\nabla f(x-t \delta \hat{g}), \hat{g}) d t=\delta \cdot \mathbb{E}_{z \sim \operatorname{Unif}[x-\delta \hat{g}, x]}\langle\nabla f(z), \hat{g}\rangle
$$

Hence, when $g$ does not satisfy the descent condition (8), we can output a random vector $u \in \partial_{\delta} f(x)$ such that

$$
\begin{equation*}
\mathbb{E}\langle u, g\rangle \leq \frac{\epsilon}{3}\|g\|_{2} . \tag{9}
\end{equation*}
$$

Then, an arbitrary vector $g$ either satisfies (8) or can be used to output a random vector $u$ satisfying (9). As described in Corollary 5, Algorithm 1 achieves this goal in $\widetilde{\mathcal{O}}\left(L^{2} / \epsilon^{2}\right)$ iterations.

In this section, we improve upon this oracle complexity by applying cutting plane methods (which we review shortly) to design Algorithm 3, which finds a better descent direction in $\widetilde{\mathcal{O}}(L d / \epsilon)$ oracle calls for $L$-Lipschitz functions and $\mathcal{O}(d \log (L / \epsilon) \log (\delta \rho / \epsilon))$ oracle calls for $\rho$-weakly convex functions.
Brief overview of cutting-plane methods. We first provide a brief relevant overview of cutting-plane methods here and refer the reader to standard textbooks in optimization for a more in-depth exposition. Given a convex function $f$ with its set $\mathcal{S}$ of minimizers, a cutting-plane method (CPM) minimizes $f$ by maintaining a convex search set $\mathcal{E}^{(k)} \supseteq \mathcal{S}$ in the $k^{\text {th }}$ iteration and iteratively shrinking $\mathcal{E}^{(k)}$ guided by the subgradients of $f$ that act as "separation oracles" for the set $\mathcal{S}$. Specifically, this is achieved by noting that for any $x^{(k)}$ chosen from $\mathcal{E}^{(k)}$, if the gradient oracle indicates $\nabla f\left(x^{(k)}\right) \neq 0$, (i.e. $x^{(k)} \notin \mathcal{S}$ ), then the convexity of $f$ guarantees $\mathcal{S} \subseteq \mathcal{H}^{(k)}:\left\{y:\left\langle\nabla f\left(x^{(k)}\right), y-x^{(k)}\right\rangle \leq 0\right\}$, and hence $\mathcal{S} \subseteq \mathcal{H}^{(k)} \cap \mathcal{E}^{(k)}$. The algorithm continues by choosing $\mathcal{E}^{(k+1)} \supseteq \mathcal{E}^{(k)} \cap \mathcal{H}^{(k)}$, and different choices of $x^{(k)}$ and $\mathcal{E}^{(k)}$ yield different rates of shrinkage of $\mathcal{E}^{(k)}$ until a point in $\mathcal{S}$ is found.
In light of this description, the minimization of a convex function over a constrained convex set via this cutting-plane method requires, at each iteration, merely a subgradient of the function. Our novel insight is that a lack of function decrease implies we have roughly such a subgradient, which we may then use in a cutting-plane method for computing the minimum norm element of the subdifferential faster in low dimensions (with improved complexity for weakly convex functions).

Setting the stage for our algorithm. In Appendix B, we demonstrate how to remove the expectation in (9) and turn the inequality into a high probability statement. For now, we assume the existence of an oracle $\mathscr{O}$ as in Definition 2.
Definition 2 (Inner Product Oracle). Given a vector $g \in \partial_{\delta} f(x)$ that does not satisfy the descent condition (8), the inner product oracle $\mathscr{O}(g)$ outputs a vector $u \in \partial_{\delta} f(x)$ such that

$$
\langle u, g\rangle \leq \frac{\epsilon}{2}\|g\|_{2} .
$$

We defer the proof of the lemma below to Appendix B.
Lemma 7. Fix $x \in \mathbb{R}^{d}$ and a unit vector $\hat{g} \in \mathbb{R}^{d}$ such that $f$ is differentiable almost everywhere on the line segment $[x, y]$, where $y \stackrel{\text { def }}{=} x-\delta \hat{g}$. Suppose that $z \in \mathbb{R}^{d}$ sampled uniformly from $[x, y]$ satisfies $\mathbb{E}_{z}\langle\nabla f(z), \hat{g}\rangle \leq \frac{\epsilon}{3}$. Then we can find $\bar{z} \in \mathbb{R}^{d}$ using at most $O\left(\frac{L}{\epsilon} \log (1 / \gamma)\right)$ gradient evaluations of $f$, such that with probability at least $1-\gamma$ the estimate $\langle\nabla f(\bar{z}), \hat{g}\rangle \leq \frac{\epsilon}{2}$ holds. Moreover, if $f$ is $\rho$-weakly convex, we can find $\bar{z} \in \mathbb{R}^{d}$ such that $\langle\nabla f(\bar{z}), \hat{g}\rangle \leq \frac{\epsilon}{2}$ using only $O(\log (\delta \rho / \epsilon))$ function evaluations of $f$.

Our key insight is that this oracle is almost identical to the gradient oracle of the minimal norm element problem

$$
\min _{g \in \partial_{\delta} f(x)}\|g\|_{2}
$$

Therefore, we can use it in the cutting plane method to find an approximate minimal norm element of $\partial_{\delta} f$. When there is no element of $\partial_{\delta} f$ with norm less than $\epsilon$, our algorithm will instead find a vector that satisfies the descent condition. The main result of this section is the following theorem.
Theorem 8. Let $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ be an L-Lipschitz function. Fix an initial point $x_{0} \in \mathbb{R}^{d}$, and let $\Delta \stackrel{\text { def }}{=} f\left(x_{0}\right)-\inf _{x} f(x)$. Then, there exists an algorithm that outputs a point $x \in \mathbb{R}^{d}$ satisfying $\operatorname{dist}\left(0, \partial_{\delta} f(x)\right) \leq \epsilon$ and, with probability at least $1-\gamma$, uses at most

$$
\mathcal{O}\left(\frac{\Delta L d}{\delta \epsilon^{2}} \cdot \log (L / \epsilon) \cdot \log (1 / \gamma)\right) \quad \text { function value/gradient evaluations. }
$$

If $f$ is $\rho$-weakly convex, the analogous statement holds with probability one and with the improved efficiency estimate $\mathcal{O}\left(\frac{\Delta d}{\delta \epsilon} \log (L / \epsilon) \cdot \log (\delta \rho / \epsilon)\right)$ of function value/gradient evaluations.

### 3.1 Finding a minimal norm element

In this section, we show, via Algorithm 3, how to find an approximate minimal norm element of $\partial_{\delta} f(x)$. Instead of directly working with the minimal norm problem, we note that, by CauchySchwarz inequality and the Minimax Theorem, for any closed convex set $Q$, we have

$$
\begin{equation*}
\min _{g \in Q}\|g\|_{2}=\min _{g \in Q}\left[\max _{\|v\|_{2} \leq 1}\langle g, v\rangle\right]=\max _{\|v\|_{2} \leq 1}\left[\min _{g \in Q}\langle g, v\rangle\right]=\max _{\|v\|_{2} \leq 1} \phi_{Q}(v), \tag{10}
\end{equation*}
$$

where $\phi_{Q}(v) \stackrel{\text { def }}{=} \min _{g \in Q}\langle g, v\rangle$, and Lemma 9 formally connects the problem of finding the minimal norm element with that of maximizing $\phi_{Q}$. The key observation in this section (Lemma 10) is that the inner product oracle $\mathscr{O}$ is a separation oracle for the (dual) problem $\max _{\|v\|_{2} \leq 1} \phi_{Q}(v)$ with $Q=\partial_{\delta} f(x)$ and hence can be used in cutting plane methods.
Lemma 9. Let $Q \subset \mathbb{R}^{d}$ be a closed convex set that does not contain the origin. Let $g_{Q}^{*}$ be a minimizer of $\min _{g \in Q}\|g\|_{2}$. Then, the vector $v_{Q}^{*}=g_{Q}^{*} /\left\|g_{Q}^{*}\right\|_{2}$ satisfies

$$
\left\langle v_{Q}^{*}, g\right\rangle \geq\left\|g_{Q}^{*}\right\|_{2} \quad \text { for all } g \in Q
$$

and $v_{Q}^{*}=\arg \max _{\|v\|_{2} \leq 1} \phi_{Q}(v)$.
Proof. We omit the subscript $Q$ to simplify notation. Since, by definition, $g^{*}$ minimizes $\|g\|_{2}$ over all $g \in Q$, we have

$$
\left\langle g^{*}, g\right\rangle \geq\left\|g^{*}\right\|_{2}^{2} \text { for all } g \in Q
$$

and the inequality is tight for $g=g^{*}$. Using this fact and $\phi\left(v^{*}\right)=\min _{g \in Q}\left\langle g, \frac{g^{*}}{\left\|g^{*}\right\|_{2}}\right\rangle$ gives

$$
\phi\left(v^{*}\right)=\left\|g^{*}\right\|_{2}=\min _{g \in Q}\|g\|_{2}=\min _{g \in Q} \max _{v:\|v\|_{2} \leq 1}\langle g, v\rangle=\max _{\|v\|_{2} \leq 1} \min _{g \in Q}\langle g, v\rangle=\max _{v:\|v\|_{2} \leq 1} \phi(v),
$$

where we used Sion's minimax theorem in the second to last step. This completes the proof.

Using this lemma, we can show that $\mathscr{O}$ is a separation oracle.
Lemma 10. Consider a vector $g \in \partial f_{\delta}(x)$ that does not satisfy the descent condition (8), and let the output of querying the oracle at $g$ be $u \in \mathscr{O}(g)$. Suppose that $\operatorname{dist}\left(0, \partial_{\delta} f(x)\right) \geq \frac{\epsilon}{2}$. Let $g^{*}$ be the minimal-norm element of $\partial_{\delta} f(x)$. Then the normalized vector $v^{*} \stackrel{\text { def }}{=} g^{*} /\left\|g^{*}\right\|_{2}$ satisfies the inclusion:

$$
v^{*} \in\left\{w \in \mathbb{R}^{d}:\langle u, \hat{g}-w\rangle \leq 0\right\}
$$

Proof. Set $Q=\partial_{\delta} f(x)$. By using $\langle u, \hat{g}\rangle \leq \frac{\epsilon}{2}$ (the guarantee of $\mathscr{O}$ per Definition 2) and $\left\langle u, v^{*}\right\rangle \geq$ $\left\|g^{*}\right\|_{2}$ (from Lemma 9), we have $\left\langle u, \hat{g}-v^{*}\right\rangle=\langle u, \hat{g}\rangle-\left\langle u, v^{*}\right\rangle \leq \frac{\epsilon}{2}-\left\|g^{*}\right\|_{2} \leq 0$.

Thus Lemma 10 states that if $x$ is not a $\left(\delta, \frac{\epsilon}{2}\right)$-stationary point of $f$, then the oracle $\mathscr{O}$ produces a halfspace $\mathcal{H}_{v}$ that separates $\hat{g}$ from $v^{*}$. Since $\mathscr{O}$ is a separation oracle, we can combine it with any cutting plane method to find $v^{*}$. For concreteness, we use the center of gravity method and display our algorithm in Algorithm 3. We note that $\Omega_{k}$ is an intersection of a ball and some half spaces, hence we can compute its center of gravity in polynomial time by taking an average of the empirical samples from this convex set. While we use a simple cutting-plane method, any algorithm in this class may be used; our focus is merely on minimizing the oracle complexity. Further note that in our algorithm, we use a point $\zeta_{k}$ close to the true center of gravity of $\Omega_{k}$, and therefore, we invoke a result about the perturbed center of gravity method.
Theorem 11 (Theorem 3 of [31]; see also [32]). Let $K$ be a convex set with center of gravity $\mu$ and covariance matrix $A$. For any halfspace $H$ that contains some point $x$ with $\|x-\mu\|_{A^{-1}} \leq t$, we have

$$
\operatorname{vol}(K \cap H) \leq(1-1 / e+t) \operatorname{vol}(K)
$$

```
Algorithm 3 MinNormCG \((x)\)
    Initialize center point \(x\).
    Set \(k=0\), the search region \(\Omega_{0}=\mathbb{B}_{2}(0)\), the set of gradients \(Q_{0}=\{\nabla f(x)\}\), and \(r\) satisfying
    \(0<r<\epsilon /(32 d L)\)
    while \(\min _{g \in Q_{k}}\|g\|_{2}>\epsilon\) do
        Let \(v_{k}\) be the center of gravity of \(\Omega_{k}\).
        if \(v_{k}\) satisfies the descent condition (8) at \(x\) then
            Return \(v_{k}\)
        end if
        Sample \(\zeta_{k}\) uniformly from \(\mathbb{B}_{r}\left(v_{k}\right)\)
        \(u_{k} \leftarrow \mathscr{O}\left(\zeta_{k}\right)\)
        \(\Omega_{k+1}=\Omega_{k} \cap\left\{w:\left\langle u_{k}, \zeta_{k}-w\right\rangle \leq 0\right\}\).
        \(Q_{k+1}=\operatorname{conv}\left(Q_{k} \cup\left\{u_{k}\right\}\right)\)
        \(k=k+1\)
    end while
    Return \(\arg \min _{g \in Q_{k}}\|g\|_{2}\).
```

Theorem 12 (Theorem 4.1 of [33]). Let $K$ be a convex set in $\mathbb{R}^{d}$ with center of gravity $\mu$ and covariance matrix $A$. Then,

$$
K \subset\left\{x:\|x-\mu\|_{A^{-1}} \leq \sqrt{d(d+2)}\right\}
$$

We now have all the tools to show correctness and iteration complexity of Algorithm 3.
Theorem 13. Let $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ be an L-Lipschitz function. Then Algorithm 3 returns a vector $v \in \partial_{\delta} f(x)$ that either satisfies the descent condition (8) at $x$ or satisfies $\|v\|_{2} \leq \epsilon$ in

$$
\lceil 8 d \log (8 L / \epsilon))\rceil \text { calls to } \mathscr{O} .
$$

Proof. By the description of Algorithm 3, either it returns a vector $v$ satisfying the descent condition or returns $g \in \partial_{\delta} f(x)$ with $\|g\|_{2} \leq \epsilon$. We now obtain the algorithm's claimed iteration complexity.
Consider an iteration $k$ such that $\Omega_{k}$ does contain a ball of radius $\frac{\epsilon}{4 L}$. Let $A_{k}$ be the covariance matrix of convex set $\Omega_{k}$. By Theorem 12, we have

$$
A_{k} \succeq\left(\frac{\epsilon}{8 d L}\right)^{2} I
$$

Applying this result to the observation that in Algorithm $3 \zeta_{k}$ is sampled uniformly from $\mathbb{B}_{r}\left(v_{k}\right)$ gives

$$
\begin{equation*}
\left\|v_{k}-\zeta_{k}\right\|_{A_{k}^{-1}} \leq r \cdot \frac{8 d L}{\epsilon} \leq \frac{1}{4} \tag{11}
\end{equation*}
$$

Recall from Algorithm 3 and the preceding notation that $\Omega_{k}$ has center of gravity $v_{k}$ and covariance matrix $A_{k}$. Further, the halfspace $\left\{w:\left\langle u_{k}, \zeta_{k}-w\right\rangle \leq 0\right\}$ in Algorithm 3 contains the point $\zeta_{k}$ satisfying (11). Given these statements, since Algorithm 3 sets $\Omega_{k+1}=\Omega_{k} \cap\left\{w:\left\langle u_{k}, \zeta_{k}-w\right\rangle\right\}$, we may invoke Theorem 11 to obtain

$$
\begin{equation*}
\operatorname{vol}\left(\Omega_{k}\right) \leq(1-1 / e+1 / 4)^{k} \operatorname{vol}\left(\mathbb{B}_{2}(0)\right) \leq(1-1 / 10)^{k} \operatorname{vol}\left(\mathbb{B}_{2}(0)\right) \tag{12}
\end{equation*}
$$

We claim that Algorithm 3 takes at most $T+1$ steps where $T=d \log _{\left(1-\frac{1}{10}\right)}(\epsilon /(8 L))$. For the sake of contradiction, suppose that this statement is false. Then, applying (12) with $k=T+1$ gives

$$
\begin{equation*}
\operatorname{vol}\left(\Omega_{T+1}\right) \leq\left(\frac{\epsilon}{4 L}\right)^{d} \operatorname{vol}\left(\mathbb{B}_{1}(0)\right) \tag{13}
\end{equation*}
$$

On the other hand, Algorithm 3 generates points $u_{i}=\mathscr{O}\left(\zeta_{i}\right)$ in the $i$-th call to $\mathscr{O}$ and the set $Q_{i}=\operatorname{conv}\left\{u_{1}, u_{2}, \cdots, u_{i}\right\}$. Since we assume that the algorithm takes more than $T+1$ steps, we have $\min _{g \in Q_{T+1}}\|g\|_{2} \geq \epsilon$. Using this and $u_{i} \in Q_{T+1}$, Lemma 10 lets us conclude that $v_{Q_{T+1}}^{*} \in\left\{w \in \mathbb{R}^{d}:\left\langle u_{i}, \zeta_{i}-w\right\rangle \leq 0\right\}$ for all $i \in[T+1]$. Since $\Omega_{T+1}$ is the intersection of the unit ball and these halfspaces, we have

$$
v_{Q_{T+1}}^{*} \in \Omega_{T+1}
$$

Per (13), $\Omega_{T+1}$ does not contain a ball of radius $\frac{\epsilon}{4 L}$, and therefore we may conclude that

$$
\text { there exists a point } \widetilde{v} \in \mathbb{B}_{\frac{\epsilon}{2 L}}\left(v_{Q_{T+1}}^{*}\right) \text { such that } \widetilde{v} \notin \Omega_{T+1} .
$$

Since $\widetilde{v} \in \mathbb{B}_{2}(0)$, the fact $\widetilde{v} \notin \Omega_{T+1}$ must be true due to one of the halfspaces generated in Algorithm 3. In other words, there must exist some $i \in[T+1]$ with

$$
\left\langle u_{i}, \zeta_{i}-\widetilde{v}\right\rangle>0
$$

By the guarantee of $\mathscr{O}$, we have $\left\langle u_{i}, \zeta_{i}\right\rangle \leq \frac{\epsilon}{2}$, and hence

$$
\begin{equation*}
\left\langle u_{i}, \widetilde{v}\right\rangle=\left\langle u_{i}, v_{i}\right\rangle-\left\langle u, v_{i}-\widetilde{v}\right\rangle<\frac{\epsilon}{2} . \tag{14}
\end{equation*}
$$

By applying $\widetilde{v} \in \mathbb{B}_{\frac{\epsilon}{2 L}}\left(v_{Q_{T+1}}^{*}\right), u_{i} \in \partial_{\delta} f(x), L$-Lipschitzness of $f$, and Lemma 9, we have

$$
\begin{equation*}
\left\langle u_{i}, \widetilde{v}\right\rangle \geq\left\langle u_{i}, v_{Q_{T+1}}^{*}\right\rangle-\frac{\epsilon}{2 L}\left\|u_{i}\right\|_{2} \geq\left\langle u_{i}, v_{Q_{T+1}}^{*}\right\rangle-\frac{\epsilon}{2} \geq\left\|g_{Q_{T+1}}^{*}\right\|_{2}-\frac{\epsilon}{2} \tag{15}
\end{equation*}
$$

Combining (14) and (15) yields that $\min _{g \in Q_{T+1}}\|g\|_{2}=\left\|g_{Q_{T+1}}^{*}\right\|_{2}<\epsilon$. This contradicts the assumption that the algorithm takes more than $T+1$ steps and concludes the proof.

Now, we are ready to prove the main theorem.
Proof of Theorem 8. We note that the outer loop in Algorithm 2 runs at most $\mathcal{O}\left(\frac{\Delta}{\delta \epsilon}\right)$ times because we decrease the objective by $\Omega(\delta \epsilon)$ every step. Combining this with Theorem 13 and Lemma 7, we have that with probability $1-\gamma$, the oracle complexity for $L$-Lipschitz function is

$$
\left.\left\lceil\frac{4 \Delta}{\delta \epsilon}\right\rceil \cdot\lceil 8 d \log (8 L / \epsilon))\right\rceil \cdot\left\lceil\frac{36 L}{\epsilon}\right\rceil \cdot\left\lceil 2 \log \left(\frac{4 \Delta}{\gamma \delta \epsilon}\right)\right\rceil=\mathcal{O}\left(\frac{\Delta L d}{\delta \epsilon^{2}} \cdot \log (L / \epsilon) \cdot \log (1 / \gamma)\right)
$$

and for $L$-Lipschitz and $\rho$-weakly convex function is $\mathcal{O}\left(\frac{\Delta d}{\delta \epsilon} \log (L / \epsilon) \cdot \log (\delta \rho / \epsilon)\right)$.

## Acknowledgements

Research of D. Davis was supported by an Alfred P. Sloan research fellowship and NSF DMS award 2047637. Research of D. Drusvyatskiy was supported by NSF DMS-1651851 and CCF2023166 awards. Research of YL was supported by NSF awards CCF-1749609, DMS-1839116, DMS-2023166, CCF-2105772, a Microsoft Research Faculty Fellowship, Sloan Research Fellowship, and Packard Fellowship. Research of GY was supported by an MIT Presidential Fellowship. Part of this work was done while GY was a student at the University of Washington.

## References

[1] Jingzhao Zhang, Hongzhou Lin, Stefanie Jegelka, Suvrit Sra, and Ali Jadbabaie. Complexity of finding stationary points of nonconvex nonsmooth functions. In Hal Daumé III and Aarti Singh, editors, Proceedings of the 37th International Conference on Machine Learning, volume 119 of Proceedings of Machine Learning Research, pages 11173-11182, Virtual, 13-18 Jul 2020.
[2] Naum Z. Shor, Krzysztof C Kiwiel, and Andrzej Ruszcayński. Minimization methods for non-differentiable functions, 1985.
[3] Frank H Clarke. Optimization and nonsmooth analysis. SIAM, 1990.
[4] R Tyrrell Rockafellar and Roger J-B Wets. Variational analysis, volume 317. Springer Science \& Business Media, 2009.
[5] E. A. Nurminskii. The quasigradient method for the solving of the nonlinear programming problems. Cybernetics, 9(1):145-150, Jan 1973. ISSN 1573-8337. doi: 10.1007/BF01068677. URL https://doi.org/10.1007/BF01068677.
[6] Damek Davis and Dmitriy Drusvyatskiy. Stochastic model-based minimization of weakly convex functions. SIAM Journal on Optimization, 29(1):207-239, 2019.
[7] Saeed Ghadimi and Guanghui Lan. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. SIAM Journal on Optimization, 23(4):2341-2368, 2013.
[8] Sashank J Reddi, Ahmed Hefny, Suvrit Sra, Barnabas Poczos, and Alex Smola. Stochastic variance reduction for nonconvex optimization. In International conference on machine learning, pages 314-323. PMLR, 2016.
[9] Chi Jin, Rong Ge, Praneeth Netrapalli, Sham M Kakade, and Michael I Jordan. How to escape saddle points efficiently. In International Conference on Machine Learning, pages 1724-1732. PMLR, 2017.
[10] Zeyuan Allen-Zhu. How to make the gradients small stochastically: Even faster convex and nonconvex sgd. Advances in Neural Information Processing Systems, 31, 2018.
[11] Yair Carmon, John C Duchi, Oliver Hinder, and Aaron Sidford. Accelerated methods for nonconvex optimization. SIAM Journal on Optimization, 28(2):1751-1772, 2018.
[12] Damek Davis, Dmitriy Drusvyatskiy, Kellie J MacPhee, and Courtney Paquette. Subgradient methods for sharp weakly convex functions. Journal of Optimization Theory and Applications, 179(3):962-982, 2018.
[13] Cong Fang, Chris Junchi Li, Zhouchen Lin, and Tong Zhang. Spider: Near-optimal nonconvex optimization via stochastic path-integrated differential estimator. Advances in Neural Information Processing Systems, 31, 2018.
[14] Dongruo Zhou, Pan Xu, and Quanquan Gu. Stochastic nested variance reduction for nonconvex optimization. Advances in Neural Information Processing Systems, 31, 2018.
[15] Michel Benaïm, Josef Hofbauer, and Sylvain Sorin. Stochastic approximations and differential inclusions. SIAM Journal on Control and Optimization, 44(1):328-348, 2005.
[16] Krzysztof C Kiwiel. Convergence of the gradient sampling algorithm for nonsmooth nonconvex optimization. SIAM Journal on Optimization, 18(2):379-388, 2007.
[17] Szymon Majewski, Błażej Miasojedow, and Eric Moulines. Analysis of nonsmooth stochastic approximation: the differential inclusion approach. arXiv preprint arXiv:1805.01916, 2018.
[18] Damek Davis, Dmitriy Drusvyatskiy, Sham Kakade, and Jason D Lee. Stochastic subgradient method converges on tame functions. Foundations of computational mathematics, 20(1): 119-154, 2020.
[19] Jérôme Bolte and Edouard Pauwels. Conservative set valued fields, automatic differentiation, stochastic gradient methods and deep learning. Mathematical Programming, 188(1):19-51, 2021.
[20] AA Goldstein. Optimization of lipschitz continuous functions. Mathematical Programming, 13 (1):14-22, 1977.
[21] Guy Kornowski and Ohad Shamir. Oracle complexity in nonsmooth nonconvex optimization. Advances in Neural Information Processing Systems, 34, 2021.
[22] J.V. Burke, A.S. Lewis, and M.L. Overton. Approximating subdifferentials by random sampling of gradients. Mathematics of Operations Research, 27(3):567-584, 2002. doi: 10.1287/moor. 27.3.567.317.
[23] Kamil A Khan and Paul I Barton. Evaluating an element of the clarke generalized jacobian of a composite piecewise differentiable function. ACM Transactions on Mathematical Software (TOMS), 39(4):1-28, 2013.
[24] Jerome Bolte and Edouard Pauwels. A mathematical model for automatic differentiation in machine learning. arXiv preprint arXiv:2006.02080, 2020.
[25] Lai Tian, Kaiwen Zhou, and Anthony Man-Cho So. On the finite-time complexity and practical computation of approximate stationarity concepts of lipschitz functions. In International Conference on Machine Learning, pages 21360-21379. PMLR, 2022.
[26] James V Burke, Frank E Curtis, Adrian S Lewis, Michael L Overton, and Lucas EA Simões. Gradient sampling methods for nonsmooth optimization. In Numerical Nonsmooth Optimization, pages 201-225. Springer, 2020.
[27] Sébastien Bubeck and Dan Mikulincer. How to trap a gradient flow. In Conference on Learning Theory, pages 940-960. PMLR, 2020.
[28] Yair Carmon, John C Duchi, Oliver Hinder, and Aaron Sidford. Lower bounds for finding stationary points i. Mathematical Programming, 184(1):71-120, 2020.
[29] James V Burke, Adrian S Lewis, and Michael L Overton. A robust gradient sampling algorithm for nonsmooth, nonconvex optimization. SIAM Journal on Optimization, 15(3):751-779, 2005.
[30] Ohad Shamir. Can we find near-approximately-stationary points of nonsmooth nonconvex functions? arXiv preprint arXiv:2002.11962, 2020.
[31] Dimitris Bertsimas and Santosh S. Vempala. Solving convex programs by random walks. J. ACM, 51(4):540-556, 2004.
[32] Branko Grünbaum. Partitions of mass-distributions and of convex bodies by hyperplanes. Pacific Journal of Mathematics, 10(4):1257-1261, 1960.
[33] R. Kannan, L. Lovász, and M. Simonovits. Isoperimetric problems for convex bodies and a localization lemma. Discrete Comput. Geom., 13(3-4):541-559, Dec 1995.

## Checklist

1. For all authors...
(a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
(b) Did you describe the limitations of your work? [Yes] See Section 1.1
(c) Did you discuss any potential negative societal impacts of your work? [N/A] Due to the theoretical nature of our work, we do not foresee any immediate potential negative societal impacts of it.
(d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
(a) Did you state the full set of assumptions of all theoretical results? [Yes]
(b) Did you include complete proofs of all theoretical results? [Yes] All our complete proofs are in the appendix.
3. If you ran experiments...
(a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
(a) If your work uses existing assets, did you cite the creators? [N/A]
(b) Did you mention the license of the assets? [N/A]
(c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

## A Missing proofs from Section 2

Theorem 4. Let $\left\{g_{k}\right\}$ be generated by $\operatorname{MinNorm}(x, \delta, \epsilon)$. Fix an index $k \geq 0$, and define the stopping time $\tau \stackrel{\text { def }}{=} \inf \left\{k: f\left(x-\delta \hat{g}_{k}\right)<f(x)-\delta\left\|g_{k}\right\|_{2} / 4\right.$ or $\left.\left\|g_{k}\right\|_{2} \leq \epsilon\right\}$. Then, we have

$$
\mathbb{E}\left[\left\|g_{k}\right\|_{2}^{2} 1_{\tau>k}\right] \leq \frac{16 L^{2}}{16+k}
$$

Proof. Fix an index $k$, and let $\mathbb{E}_{k}[\cdot]$ denote the conditional expectation on $g_{k}$. Suppose we are in the event $\{\tau>k\}$. Taking into account the Lipschitz continuity of $f$ and Lemma 3, we deduce that almost surely, conditioned on $g_{k}$, the following estimate holds:

$$
\begin{aligned}
\frac{1}{4}\left\|g_{k}\right\|_{2} \geq \frac{f(x)-f\left(x-\delta \hat{g}_{k}\right)}{\delta} & \geq \frac{f(x)-f\left(x-\delta \cdot \hat{\zeta}_{k}\right)}{\delta}-L\left\|\hat{g}_{k}-\hat{\zeta}_{k}\right\|_{2} \\
& =\frac{1}{\delta} \int_{0}^{\delta}\left\langle\nabla f\left(x-s \hat{\zeta}_{k}\right), \hat{\zeta}_{k}\right\rangle d s-L\left\|\hat{g}_{k}-\hat{\zeta}_{k}\right\|_{2} \\
& \geq \frac{1}{\delta} \int_{0}^{\delta}\left\langle\nabla f\left(x-s \hat{\zeta}_{k}\right), \hat{g}_{k}\right\rangle d s-2 L\left\|\hat{g}_{k}-\hat{\zeta}_{k}\right\|_{2} \\
& =\mathbb{E}_{k}\left\langle\nabla f\left(y_{k}\right), \hat{g}_{k}\right\rangle-2 L\left\|\hat{g}_{k}-\hat{\zeta}_{k}\right\|_{2}
\end{aligned}
$$

Rearranging yields $\mathbb{E}_{k}\left\langle\nabla f\left(y_{k}\right), \hat{g}_{k}\right\rangle \leq \frac{1}{4}\left\|g_{k}\right\|_{2}+2 L\left\|\hat{g}_{k}-\hat{\zeta}_{k}\right\|$. Simple algebra shows $\left\|\hat{g}_{k}-\hat{\zeta}_{k}\right\|_{2}^{2} \leq$ $2\left(1-\sqrt{1-r^{2} /\left\|g_{k}\right\|_{2}^{2}}\right) \leq \frac{\left\|g_{k}\right\|_{2}^{2}}{64 L^{2}}$. Therefore, we infer that $\mathbb{E}_{k}\left\langle\nabla f\left(y_{k}\right), \hat{g}_{k}\right\rangle<\frac{1}{2}\left\|g_{k}\right\|_{2}$. Lemma 2 then guarantees that

$$
\mathbb{E}_{k}\left[\left\|g_{k+1}\right\|_{2}^{2} 1_{\tau>k}\right] \leq\left(\left\|g_{k}\right\|_{2}^{2}-\frac{\left\|g_{k}\right\|_{2}^{4}}{16 L^{2}}\right) 1_{\tau>k}
$$

Define $b_{k}:=\left\|g_{k}\right\|_{2}^{2} 1_{\tau>k}$ for all $k \geq 0$. Then the tower rule for expectations yields

$$
\mathbb{E} b_{k+1} \leq \mathbb{E}\left[\left\|g_{k+1}\right\|_{2}^{2} 1_{\tau>k}\right] \leq \mathbb{E}\left[\left(1-\frac{b_{k}}{16 L^{2}}\right) b_{k}\right] \leq\left(1-\frac{\mathbb{E} b_{k}}{16 L^{2}}\right) \mathbb{E} b_{k}
$$

by Jensen's inequality applied to the concave function $t \mapsto\left(1-t / 16 L^{2}\right) t$. Setting $a_{k}=\mathbb{E} b_{k} / L^{2}$, this inequality becomes $a_{k+1} \leq a_{k}-a_{k}^{2} / 16$, which, upon rearranging, yields $\frac{1}{a_{k+1}} \geq \frac{1}{a_{k}\left(1-a_{k} / 16\right)} \geq$ $\frac{1}{a_{k}}+\frac{1}{16}$. Iterating the recursion and taking into account $a_{0} \leq 1$ completes the proof.

Corollary 14. MinNorm $(x, \delta, \epsilon)$ terminates in at most $\left\lceil\frac{64 L^{2}}{\epsilon^{2}}\right\rceil \cdot\lceil 2 \log (1 / \gamma)\rceil$ iterations with probability at least $1-\gamma$, where we define the stopping time $\tau \stackrel{\text { def }}{=}$ $\inf \left\{k: f\left(x-\delta \hat{g}_{k}\right)<f(x)-\delta\left\|g_{k}\right\|_{2} / 4\right.$ or $\left.\left\|g_{k}\right\|_{2} \leq \epsilon\right\}$.

Proof. Notice that when $k \geq \frac{64 L^{2}}{\varepsilon^{2}}$, we have, by Theorem 4, that

$$
\operatorname{Pr}(\tau>k) \leq \operatorname{Pr}\left(\left\|g_{k}\right\|_{2} 1_{\tau>k} \geq \epsilon\right) \leq \frac{16 L^{2}}{(16+k) \varepsilon^{2}} \leq \frac{1}{4}
$$

Similarly, for all $i \in \mathbb{N}$, we have $\operatorname{Pr}(\tau>i k \mid \tau>(i-1) k) \leq 1 / 4$. Therefore,

$$
\operatorname{Pr}(\tau>i k)=\operatorname{Pr}(\tau>i k \mid \tau>(i-1) k) \operatorname{Pr}(\tau>(i-1) k) \leq \frac{1}{4} \operatorname{Pr}(\tau>(i-1) k) \leq \frac{1}{4^{i}}
$$

Consequently, we have $\operatorname{Pr}(\tau>i k) \leq \frac{1}{4^{i}} \leq \gamma$ whenever $i \geq \log (1 / \gamma) / \log (4)$, as desired.

## B Implementation of the oracles: proof of Lemma 7

In this section, we show how to convert (9) into a deterministic guarantee.
Lemma 15. Fix a unit vector $\hat{g} \in \mathbb{R}^{d}$ and let $z \in \mathbb{R}^{d}$ be a random vector satisfying $\mathbb{E}\langle\nabla f(z), \hat{g}\rangle \leq \frac{\epsilon}{3}$. Let $z_{1}, \ldots, z_{k}$ be i.i.d realizations of $z$ with $k=\left\lceil\frac{36 L}{\epsilon}\right\rceil \cdot\left\lceil\frac{\log (1 / \gamma)}{\log (4)}\right\rceil$. Then with probability at least $1-\gamma$, one of the samples $z_{i}$ satisfies $\left\langle\nabla f\left(z_{i}\right), \hat{g}\right\rangle \leq \frac{\epsilon}{2}$.

Proof. Define the random variable $Y \stackrel{\text { def }}{=}\langle\nabla f(z), \hat{g}\rangle$, and use $p \stackrel{\text { def }}{=} \operatorname{Pr}\left[Y \leq \frac{\epsilon}{2}\right]$. We note that

$$
\mathbb{E}[Y]=p \cdot \mathbb{E}\left[Y \left\lvert\, Y \leq \frac{\epsilon}{2}\right.\right]+(1-p) \cdot \mathbb{E}\left[Y \left\lvert\, Y>\frac{\epsilon}{2}\right.\right]
$$

Rearranging the terms and using $\mathbb{E}[Y] \leq \epsilon / 3$ gives

$$
p \cdot\left(\mathbb{E}\left[Y \left\lvert\, Y>\frac{\epsilon}{2}\right.\right]-\mathbb{E}\left[Y \left\lvert\, Y \leq \frac{\epsilon}{2}\right.\right]\right) \geq \frac{\epsilon}{6}
$$

Finally, taking into account that $f$ is $L$-Lipschitz, we deduce $|Y| \leq L$, which further implies $p \geq \frac{\epsilon}{12 L}$. The results follows immediately.

Lemma 16. Let $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ be an L-Lipschitz continuous and $\rho$-weakly convex function. Fix a point $x$ and a unit vector $\hat{g} \in \mathbb{R}^{d}$ such that $f$ is differentiable almost everywhere on the line segment $[x, y]$, where $y \stackrel{\text { def }}{=} x-\delta \hat{g}$. Suppose that a random vector $z$ sampled uniformly from $[x, y]$ satisfies $\mathbb{E}_{z}\langle\nabla f(z), \hat{g}\rangle \leq \frac{\epsilon}{3}$. Then, Algorithm 4 finds $\bar{z} \in \mathbb{R}^{d}$ such that $\langle\nabla f(\bar{z}), \hat{g}\rangle \leq \frac{\epsilon}{2}$ using $3 \log (12 \delta \rho / \epsilon)$ function evaluations of $f$.

Proof. Define the new function $h:[0,1] \rightarrow \mathbb{R}$ by $h(t)=\langle\nabla f(x+t(y-x)), \hat{g}\rangle$. Clearly, we have

$$
\frac{\epsilon}{3} \geq \mathbb{E}[h(t)]=\frac{1}{2} \underbrace{\mathbb{E}[h(t) \mid t \leq 0.5]}_{P_{\leq}}+\frac{1}{2} \underbrace{\mathbb{E}[h(t) \mid t>0.5]}_{P_{>}}
$$

Therefore $P_{\leq}$or $P_{>}$is at most $\epsilon / 3$. The fundamental theorem of calculus directly implies $P_{\leq}=$ $\frac{f(x)-f\left(x-\frac{\delta}{2} \hat{g}\right)}{2 \delta}$ and $P_{>}=\frac{f\left(x-\frac{\delta}{2} \hat{g}\right)-f(y)}{2 \delta}$. Therefore with three function evaluations we may determine

```
Algorithm 4 Binary Search for \(\bar{z}\)
    Input. Line Segment \([x, y=x-\delta \hat{g}]\)
    Let \([a, b]=[0,1]\)
    while \(b-a>\frac{\epsilon}{6 \delta \rho}\) do
        if \(f(x-a \delta \hat{g})-f\left(x-\frac{a+b}{2} \delta \hat{g}\right) \leq f\left(x-\frac{a+b}{2} \delta \hat{g}\right)-f(x-b \delta \hat{g})\) then
            Let \([a, b] \leftarrow\left[a, \frac{a+b}{2}\right]\)
        else
            Let \([a, b] \leftarrow\left[\frac{a+b}{2}, b\right]\)
        end if
    end while
    Return \(x-a \delta \hat{g}\)
```

one of the two alternatives. Repeating this procedure $\log (12 \delta \rho / \epsilon)$ times, each times shrinking the interval by half, we can find an interval $[a, b] \subset[0,1]$ such that $b-a \leq \frac{\epsilon}{6 \delta \rho}$ and $\mathbb{E}_{t \in[a, b]} h(t) \leq \frac{\epsilon}{3}$. Note that for any $\bar{t} \in[a, b]$, we have $h(\bar{t})=\mathbb{E} h(t)+(h(\bar{t})-\mathbb{E} h(t))$, while weak convexity implies

$$
\begin{aligned}
h(\bar{t})-\mathbb{E} h(t) & =\frac{1}{\delta} \mathbb{E}_{t \in[a, b]}\langle\nabla f(x+\bar{t}(y-x))-\nabla f(x+t(y-x)), x-y\rangle \\
& \leq \mathbb{E}_{t \in[a, b]} \frac{\bar{t}-t}{\delta} \rho\|y-x\|^{2} \leq \frac{\epsilon}{6}
\end{aligned}
$$

We thus conclude $h(\bar{t}) \leq \frac{\epsilon}{3}+\frac{\epsilon}{6}=\frac{\epsilon}{2}$ as claimed.


[^0]:    *Authors ordered alphabetically.

[^1]:    ${ }^{2}$ We provide a simple example to demonstrate this claim: Consider the function $f(x, y)=f_{1}(x, y)+f_{2}(x, y)$ with $f_{1}(x, y)=|x|$ and $f_{2}(x, y)=-|x|$. Choose the direction $u=(0,1)$, and let $z=(0,0)$. Then, $f_{1}^{\prime}(z, u)=f_{2}^{\prime}(z, u)=0$, and $\partial f_{1}(z)=\partial f_{2}(z)=[-1,1] \times 0$. Therefore, to satisfy the oracle (2), for $f_{1}$, we may choose the subgradient $v_{1}=(-1,0)$, and for $f_{2}$, we may choose the subgradient $v_{2}=(-1,0)$ since $\left\langle v_{1}, u\right\rangle=0=\left\langle v_{2}, u\right\rangle$. However, $v_{1}+v_{2}=(-2,0)$, which is not a subgradient of $f=0$ at $z$.
    ${ }^{3}$ Throughout the paper, we use $\widetilde{\mathcal{O}}(\cdot)$ to hide poly-logarithmic factors in $L, \delta, \Delta$, and $\epsilon$.

