

The cost of nonconvexity in deterministic nonsmooth optimization

Siyu Kong ^{*} A.S. Lewis [†]

October 7, 2022

Abstract

We study the impact of nonconvexity on the complexity of nonsmooth optimization, emphasizing objectives such as piecewise linear functions, which may not be weakly convex. We focus on a dimension-independent analysis, slightly modifying a black-box algorithm of Zhang et al. [14] that approximates an ϵ -stationary point of any directionally differentiable Lipschitz objective using $O(\epsilon^{-4})$ calls to a specialized subgradient oracle and a randomized line search. Our simple black-box deterministic version, achieves $O(\epsilon^{-5})$ for any difference-of-convex objective, and $O(\epsilon^{-4})$ for the weakly convex case. Our complexity bound depends on a natural nonconvexity modulus, related, intriguingly, to the negative part of directional second derivatives of the objective, understood in the distributional sense.

Key words: nonsmooth optimization, nonconvex, Goldstein subgradient, complexity, distributional derivative

AMS Subject Classification: 90C56, 49J52, 65Y20

1 Introduction

Minimizing nonsmooth nonconvex Lipschitz functions $f: \mathbf{R}^n \rightarrow \mathbf{R}$ is in general intractable. A proxy approach, dating back to the 1970's [6], can be viewed as seeking a point in \mathbf{R}^n around which f is differentiable on some cluster of nearby points at which the gradients have a small convex combination — a small “Goldstein subgradient”, in the language of [6]. One algorithm accomplishing this goal, but without complexity analysis, appeared in [10].

^{*}ORIE, Cornell University, Ithaca, NY. sk3333@cornell.edu

[†]ORIE, Cornell University, Ithaca, NY. people.orie.cornell.edu/aslewis
Research supported in part by National Science Foundation Grant DMS-2006990.

A 2020 breakthrough [14] proposed an algorithm for this Goldstein subgradient problem with a complexity analysis depending on the radius δ of the cluster and the size ϵ of the subgradient, but *independent of the dimension n* . The algorithm assumes directional differentiability of the objective f , and, relying on an associated directional subgradient oracle, uses an innovative randomized line search to achieve an efficiency guarantee of essentially $O(\epsilon^{-3}\delta^{-1})$. Two subsequent papers [2,13] point out that small random perturbations instead allow a standard subgradient oracle.

Two recent developments [9,12]¹ raise the question of *deterministic* algorithms for this problem. While both prove positive results in the smooth case, and [9] thereby develops a “white-box” deterministic smoothing approach to the nonsmooth problem, both manuscripts also prove negative results for the general dimension-independent question.

Our contribution is to show, notwithstanding the negative results of [9,12], that a simple deterministic black-box version of the algorithm of [14] achieves, up to a *nonconvexity modulus* for the objective, a dimension-independent complexity of $O(\epsilon^{-4}\delta^{-1})$. Our analysis covers interesting objectives such as piecewise linear functions, which are not even weakly convex. We relate the nonconvexity modulus of the objective with its distributional second derivative, hinting at an intriguing relationship between such derivatives and algorithmic complexity in general optimization.

2 The optimization problem and oracle

Primarily to emphasize the elementary nature of our development, we adopt a rudimentary setting for our optimization problem. On a real inner product space \mathbf{X} with corresponding norm $|\cdot|$, we consider the problem of minimizing a function $f: \mathbf{X} \rightarrow \mathbf{R}$. The objective f may be neither smooth nor convex, and the space X may be neither finite-dimensional, nor even complete.

The method we develop, following [14], relies on an underlying *directional subgradient map* $G: \mathbf{X}^2 \rightarrow \mathbf{X}$ associated with the objective f : for all points $x \in \mathbf{X}$ and directions $e \in \mathbf{X}$, the Gâteaux directional derivative

$$f'(x; e) = \lim_{t \downarrow 0} \frac{1}{t} (f(x + te) - f(x))$$

exists and satisfies

$$\langle G(x, e), e \rangle = f'(x; e).$$

We say that G is *L-bounded* for some constant $L > 0$ if its norm $|G(x, e)|$ is never larger than L . In applications, the objective function f is L -Lipschitz and the vector $G(x, e)$ is a subgradient of some kind for f at the point x associated with

¹We became aware of these concurrent independent works after completing the initial draft of this manuscript.

the direction e , so we loosely refer to $G(x, e)$ as a “subgradient”. Nonetheless, we choose this rudimentary setting to emphasize again the elementary nature of our development, which makes no recourse to variational or Lipschitz analysis.

Example 2.1 (differentiable functions) If the function f is L -Lipschitz and has a Gâteaux derivative $\nabla f(x)$ at every point $x \in \mathbf{X}$, then the equation

$$G(x, e) = \nabla f(x)$$

defines an L -bounded directional subgradient map.

Example 2.2 (convex functions) For an L -Lipschitz convex function f with convex subdifferential ∂f , any map G satisfying

$$G(x, e) \in \operatorname{argmax}\{\langle g, e \rangle : g \in \partial f(x)\} \quad \text{for all } x, e \in \mathbf{X}$$

is an L -bounded directional subgradient map.

More generally, we have the following example, which covers many interesting objectives, including the weakly convex case. For a Lipschitz function $f: \mathbf{R}^n \rightarrow \mathbf{R}$, the *Clarke subdifferential* $\partial^c f(x)$ is the convex hull of the set of all limits of the form $\lim \nabla f(x^r)$ for points $x^r \rightarrow x$ in \mathbf{R}^n . The function f is *subdifferentially regular* when its Gâteaux directional derivative satisfies

$$f'(x; e) = \max\{\langle g, e \rangle : g \in \partial^c f(x)\} \quad \text{for all } x, e \in \mathbf{R}^n.$$

Example 2.3 (subdifferentially regular functions) For an L -Lipschitz subdifferentially regular function $f: \mathbf{R}^n \rightarrow \mathbf{R}$, any map G satisfying

$$G(x, e) \in \operatorname{argmax}\{\langle g, e \rangle : g \in \partial^c f(x)\} \quad \text{for all } x, e \in \mathbf{X}$$

is an L -bounded directional subgradient map.

Notwithstanding the generality of this example, we emphasize that our framework is not restricted to objectives that are subdifferentially regular, as the following result shows.

Proposition 2.4 (directional Clarke subgradient maps) *Any locally Lipschitz function $f: \mathbf{R}^n \rightarrow \mathbf{R}$ that is directionally differentiable has a directional subgradient map $G: \mathbf{R}^n \times \mathbf{R}^n \rightarrow \mathbf{R}^n$ satisfying $G(x, e) \in \partial^c f(x)$ for all $x, e \in \mathbf{R}^n$.*

Proof We just need to prove that if f has a Gâteaux directional derivative at the point $x \in \mathbf{R}^n$ in the direction $e \in \mathbf{R}^n$, then there exists a Clarke subgradient $g \in \partial^c f(x)$ satisfying $\langle g, e \rangle = f'(x; e)$. For $r = 1, 2, 3, \dots$, by the nonsmooth mean

value theorem, there exists a point $x_r \in [x, x + \frac{1}{r}e]$ and a subgradient $g_r \in \partial^c f(x_r)$ satisfying

$$f\left(x + \frac{1}{r}e\right) - f(x) = \langle g_r, \frac{1}{r}e \rangle.$$

Since the subdifferential $\partial^c f$ mapping is closed and locally bounded, any limit point g of the sequence $\{g_r\}$ has the desired property. \square

Minimizing nonconvex functions is in general intractable. Instead we seek a point $x \in \mathbf{X}$ that is, in some sense, approximately critical. To this end, we fix a constant $\delta > 0$, and consider what we call *Goldstein subgradients* at x : vectors of the form

$$\sum_{i=1}^k \lambda_i G(x_i, e_i)$$

for a positive integer k , positive weights λ_i summing to 1, points x_i in $B_\delta(x)$ (the closed ball of radius δ and center x) and directions $e_i \in \mathbf{X}$ for $i = 1, 2, \dots, k$. We denote the set of all Goldstein subgradients by $\partial_\delta f(x)$: loosely, it consists of all convex combinations of subgradients at nearby points. Strictly speaking, our notion is potentially smaller than the standard definition for Lipschitz $f: \mathbf{R}^n \rightarrow \mathbf{R}$, namely the closed convex hull of the set $\partial^c f(B_\delta(x))$.

We can now state our goal, which relies on a second constant $\epsilon > 0$.

Aim: Find a point $x \in \mathbf{X}$ and a Goldstein subgradient $g \in \partial_\delta f(x)$ such that $|g| \leq \epsilon$.

The development of [14] accomplishes this goal, explicitly in the setting of Proposition 2.4), and assuming that f is directionally differentiable in the (stronger) Hadamard sense. It relies on the following oracle.

Oracle 2.5 (directional subgradient)

Input:

- a point $x \in \mathbf{X}$
- a direction $e \in \mathbf{X}$.

Output:

- the objective value $f(x)$
- the directional derivative $f'(x; e)$
- a subgradient-like vector $G(x, e)$.

In this work we rely on the same directional subgradient oracle, which we should emphasize is stronger than the standard subgradient oracle typically assumed. In generic practice, for a Lipschitz objective f , we may expect that the algorithm never encounters points x where f is nondifferentiable, in which case any subgradient oracle simply returns $g = \nabla f(x)$. More formally, however, we must consider nonsmooth points. Undeterred, the authors of [14] argue that the stronger oracle is often a reasonable assumption. However, our aim here is a fully deterministic algorithm. Accordingly, we use this same stronger oracle, but instead use a deterministic line search, much as in [12]. To ensure termination, we make a mild assumption about the directional behavior of the objective function f , similar in spirit to the idea of *semismoothness* [11] common in nonsmooth computation, but simpler and weaker.

We call the objective f *directionally semismooth* if all points $x \in \mathbf{X}$ and directions $e \in \mathbf{X}$ satisfy

$$\lim_{t \downarrow 0} f'(x + te; e) = f'(x; e).$$

For comparison, following the definition in [5], a Lipschitz function $f: \mathbf{R}^n \rightarrow \mathbf{R}$ is *semismooth* if it is directionally differentiable and for all points $x \in \mathbf{R}^n$ the directional derivative satisfies the stronger property

$$f'(x + d; d) - f'(x; d) = o(d) \quad \text{as } d \rightarrow 0.$$

Most Lipschitz functions in practice are semismooth, including in particular difference-of-convex functions and semi-algebraic functions [1]. As we shall see, directional semismoothness suffices to guarantee termination of our algorithm, but before describing it, we focus first on the line search.

3 A simple line search

We pose the line search problem as a self-contained question. Consider points $p < q$ in \mathbf{R} and a function $h: [p, q] \rightarrow \mathbf{R}$ satisfying $h(p) > h(q)$. Suppose that h is right-differentiable on the interval $[p, q)$, and an oracle returns, for any input $t \in [p, q)$, the value $h(t)$ and the right derivative

$$h'_+(t) = \lim_{s \downarrow t} \frac{h(s) - h(t)}{s - t}$$

(possibly extended-valued). How difficult is it to find a point t satisfying $h'_+(t) < 0$?

When h is Lipschitz, the most basic randomized strategy — uniformly sampling random points t in the interval — solves this problem with high probability. Denoting the Lipschitz constant by L , the right derivative h'_+ always lies in the interval

$[-L, L]$. Denote the measure of the set where $h'_+ \geq 0$ by λ . Then, providing the average slope satisfies

$$\frac{h(q) - h(p)}{q - p} = -\sigma < 0,$$

the fundamental theorem of calculus implies

$$-\sigma = \frac{1}{q - p} \int_p^q h'_+ \geq L \left(\frac{\lambda}{q - p} - 1 \right).$$

Hence the probability that a uniformly distributed random point $t \in [p, q]$ fails to satisfy $h'(t) < 0$ is no larger than $1 - \frac{\sigma}{L}$. Thus for small σ , using at least $\frac{L}{\sigma}$ independent samples, the probability of success is at least $\frac{1}{2}$.

However, we seek a deterministic algorithm, so we instead consider the following simple method, similar in spirit to one described in [2]. We repeatedly bisect the interval $[p, q]$, each time discarding the subinterval over which the function h decreases the least. The algorithm checks whether the right derivative at the midpoint of the current interval is negative, in which case it terminates.

Algorithm 3.1 (bisection)

```

if  $h'_+(p) < 0$  then
  return  $p$ 
end if
 $l = p$ 
 $r = q$ 
while not done do
   $m = \frac{1}{2}(l + r)$ 
  if  $h'_+(m) < 0$  then
    return  $m$ 
  else if  $2h(m) < f(l) + f(r)$  then
     $r = m$ 
  else
     $l = m$ 
  end if
end while

```

Notice that the algorithm initially calls the oracle at the left endpoint p of the given interval, calculates the function value at the right endpoint q , and then calls the oracle once during each bisection.

In general, this algorithm may fail to terminate. It is easy to construct a Lipschitz function h satisfying $h(p) > h(q)$ and yet the derivative of h at the initial endpoint p and at every midpoint m is positive. To rule out such oscillatory examples, we

can rely on directional semismoothness of h , which in this univariate setting simply means that the right derivative exists and is right continuous,

$$\lim_{t \downarrow \bar{t}} h'_+(t) = h'_+(\bar{t}),$$

and that the left derivative also exists and is left continuous:

$$h'_-(t) = \lim_{s \uparrow t} \frac{h(s) - h(t)}{s - t} \quad \text{satisfies} \quad \lim_{t \uparrow \bar{t}} h'_-(t) = h'_-(\bar{t}).$$

For Lipschitz functions h , these two properties amount exactly to *semismoothness*. Most Lipschitz functions in practice are semismooth, including in particular, convex and concave functions, and piecewise smooth functions. Furthermore, any linear combination of semismooth functions is semismooth. When the function h is semismooth, the Clarke subdifferential is given by

$$\partial^c h(t) = \text{conv}\{h'_-(t), h'_+(t)\},$$

and the following property also holds [11, Lemma 2] and [8, Lemma 2.2]:

$$\lim_{t \downarrow \bar{t}} h'_-(t) = h'_+(\bar{t}) \quad \text{and} \quad \lim_{t \uparrow \bar{t}} h'_+(t) = h'_-(\bar{t}).$$

Semismoothness is more than enough to prove termination of the line search. The simple argument below also applies to nonlipschitz functions.

Proposition 3.2 (Termination of the line search) *Suppose that the function $h: [p, q] \rightarrow \mathbf{R}$ satisfies $h(p) > h(q)$, and that its left and right derivatives satisfy the semismoothness conditions*

$$\begin{aligned} \lim_{t \downarrow \bar{t}} h'_+(t) &= h'_+(\bar{t}) & \text{for } \bar{t} \in [p, q) \\ \lim_{t \uparrow \bar{t}} h'_+(t) &= h'_-(\bar{t}) & \text{for } \bar{t} \in (p, q]. \end{aligned}$$

Then Algorithm 3.1 terminates.

Proof If the algorithm does not terminate, then it generates monotonic sequences $l_k \uparrow$ and $r_k \downarrow$, satisfying $r_k - l_k \rightarrow 0$,

$$h'_+(l_k) \geq 0, \quad \text{and} \quad \frac{h(r_k) - h(l_k)}{r_k - l_k} \leq -\sigma < 0$$

for each iteration $k = 0, 1, 2, \dots$. Denote the two sequences' mutual limit by \bar{m} . Semismoothness ensures $h'_+(l_k) \rightarrow h'_-(\bar{m})$, so $h'_-(\bar{m}) \geq 0$.

If $r_k = \bar{m}$ for all large k , then

$$\frac{h(r_k) - h(l_k)}{r_k - l_k} = \frac{h(\bar{m}) - h(l_k)}{\bar{m} - l_k} \rightarrow h'_-(\bar{m}),$$

which is a contradiction. Hence for all large k we have $q > r_k > \bar{m}$, so $h'_+(r_k) \geq 0$, and semismoothness ensures $h'_+(r_k) \rightarrow h'_+(\bar{m})$, so $h'_+(\bar{m}) \geq 0$. We deduce

$$h(\bar{m}) - h(l_k) \geq -\frac{\sigma}{2}(\bar{m} - l_k) \quad \text{and} \quad h(r_k) - h(\bar{m}) > -\frac{\sigma}{2}(r_k - \bar{m}).$$

Adding now gives a contradiction. \square

4 The optimization algorithm

To minimize a locally Lipschitz function $f: \mathbf{X} \rightarrow \mathbf{R}$ using the directional subgradient oracle described above, we study the following algorithm. The method we describe is essentially that of [14], but with the deterministic line search described in the preceding section. Given a current point $x \in \mathbf{X}$ and a current Goldstein subgradient, the method proceeds as follows.

Algorithm 4.1 (nonsmooth minimization)

```

while not done do
  if  $|g| \leq \epsilon$  then
    return  $x$  {Small subgradient so stop.}
  end if
   $\hat{g} = \frac{g}{|g|}$  {Normalize subgradient.}
   $x' = x - \delta \hat{g}$  {Trial step of fixed length.}
  if  $f(x) - f(x') \geq \frac{\delta \epsilon}{3}$  then
     $x = x'$  {Sufficient decrease so update point.}
     $g = G(x, 0)$  {Re-initialize subgradient.}
  else
{Insufficient decrease so update subgradient.}

```

Define h on $[0, \delta]$ by

$$\begin{aligned} x(t) &= x + (t - \delta)\hat{g} \\ h(t) &= f(x(t)) - \frac{\epsilon t}{2}. \end{aligned}$$

Apply Algorithm 3.1 (bisection), using the formula

$$h'_+(t) = \langle G(x(t), \hat{g}), \hat{g} \rangle - \frac{\epsilon}{2}$$

to find $t \in [0, \delta]$ satisfying $h'_+(t) < 0$.

$g =$ shortest vector in $[g, G(x(t), \hat{g})]$

end if
end while

When the objective f is directionally semismooth, Algorithm 3.1 terminates by Proposition 3.2, which in turn guarantees termination of Algorithm 4.1, as we shall now prove. We use the following simple tool, following [14].

Lemma 4.2 *If two vectors $g, g' \in B_L(0)$ satisfy $\langle g', g \rangle \leq \frac{1}{2}|g|^2$, then the shortest vector g'' in the line segment $[g, g']$ satisfies*

$$|g''|^2 \leq |g|^2 \left(1 - \frac{|g|^2}{16L^2}\right).$$

Proof For all $t \in [0, 1]$ we have

$$|g''|^2 \leq |g + t(g' - g)|^2 = |g|^2 + t^2|g' - g|^2 + 2t\langle g, g' - g \rangle \leq (1 - t)|g|^2 + 4L^2t^2.$$

Setting $t = \frac{|g|^2}{8L^2}$ proves the result. \square

We can now prove the validity of the algorithm, again imitating parts of the argument in [14], which we reproduce for ease of reading.

Theorem 4.3 (finite termination) *Suppose that we apply Algorithm 4.1 to a directionally semismooth function $f: \mathbf{X} \rightarrow \mathbf{R}$ that is bounded below. Suppose that the directional subgradient map in Oracle 2.5 is L -bounded. Then the algorithm terminates with a point $x \in \mathbf{X}$ and a subgradient $g \in \partial_\delta f(x)$ satisfying $|g| \leq \epsilon$.*

Proof Suppose that the current subgradient g is *inadequate*, in the sense that, in the terminology of the algorithm description, it is neither small, nor generates sufficient decrease. We then apply the bisection method, Algorithm 3.1, to the given function h . The initial interval is $[p, q] = [0, \delta]$, and the average slope satisfies

$$\frac{h(q) - h(p)}{q - p} = \frac{1}{\delta} \left(f(x) - \frac{\delta\epsilon}{2} - f(x') \right) < -\frac{\epsilon}{6}.$$

We thus arrive at a subgradient $g' \in \partial_\delta f(x)$ satisfying $\langle g', \hat{g} \rangle < \frac{\epsilon}{2}$. We deduce $\langle g', g \rangle < \frac{|g|^2}{2}$. The algorithm replaces the current subgradient g by the shortest vector g'' in the line segment $[g, g']$. Since $g'' \in \partial_\delta f(x)$, we can repeat this shortening process, providing that g'' is also inadequate. Suppose that the subgradient g remains inadequate after completing k such shortening steps. Let ρ_i denote the quantity $\frac{|g|^2}{16L^2}$ after $i = 0, 1, 2, \dots, k$ steps. Then $\rho_0 \leq \frac{1}{16}$ and for each i we have

$$0 < \rho_{i+1} \leq \rho_i(1 - \rho_i),$$

so

$$\frac{1}{\rho_{i+1}} \geq \frac{1}{\rho_i} + \frac{1}{1 - \rho_i} > \frac{1}{\rho_i} + 1.$$

Consequently,

$$\frac{1}{\rho_k} \geq 16 + k,$$

so we deduce

$$\frac{\epsilon^2}{16L^2} < \frac{|g|^2}{16L^2} \leq \frac{1}{16 + k}.$$

Hence after no more than

$$16\left(\frac{L^2}{\epsilon^2} - 1\right)$$

shortening steps, each requiring one line search, we arrive at an adequate subgradient $g \in \partial_\delta f(x)$. To summarize, starting at any point with an inadequate subgradient, we require no more than $\frac{16L^2}{\epsilon^2}$ line searches before finding an adequate subgradient g .

There are now two possibilities. Either the subgradient g satisfies $|g| \leq \epsilon$, in which case we stop, or we perform a reduction step, replacing the current point x by $x - \delta \frac{g}{|g|}$, thereby decreasing the objective value by at least the quantity $\frac{\delta\epsilon}{3}$. Since the objective is bounded below, beginning from the initial point x_0 , this procedure terminates after no more than $\lceil \frac{3}{\delta\epsilon}(f(x_0) - \min f) \rceil$ reduction steps and hence

$$(4.4) \quad \left\lceil \frac{3(f(x_0) - \min f)}{\delta\epsilon} \right\rceil \cdot \frac{16L^2}{\epsilon^2}$$

line searches. Proposition 3.2 ensures that each line search requires only finitely many oracle calls, completing the proof. \square

5 Complexity of the line search

To complete our complexity analysis for the minimization algorithm, we simply need to bound the number of oracle calls needed by each line search, and multiply by our bound (4.4) on the number of line searches. Consider, therefore, the bisection method. When the function $h: [p, q] \rightarrow \mathbf{R}$ is convex, the problem is trivial: since

$$h(p) > h(q) \quad \Rightarrow \quad h'_+(p) < 0,$$

the algorithm terminates at the first oracle call. More generally, we proceed by correcting any lack of convexity in h by adding a convex perturbation $s: [p, q] \rightarrow \mathbf{R}$.

Recall that, for any interval J , a function $h: J \rightarrow \mathbf{R}$ is *difference-of-convex* when there exists a convex function $s: J \rightarrow \mathbf{R}$ such that $h + s$ is also convex. For such functions, we have the following tool.

Lemma 5.1 Consider a function $h: [p, q] \rightarrow \mathbf{R}$, and a convex function $s: [p, q] \rightarrow \mathbf{R}$ such that $h+s$ is also convex. For any points $x < y$ in the interval $[p, q]$, if $h'_+(x) \geq 0$, then

$$\frac{h(y) - h(x)}{y - x} \geq s'_+(x) - s'_-(y).$$

Proof Denote the convex function $h+s$ by r . The convex functions r and s satisfy

$$r'_+(x) \in \partial r(x) \quad \text{and} \quad s'_-(y) \in \partial s(y).$$

Hence

$$\begin{aligned} s'_-(y)(x - y) &\leq s(x) - s(y) = r(x) - h(x) - r(y) + h(y) \\ &\leq h(y) - h(x) + r'_+(x)(x - y) \leq h(y) - h(x) + s'_+(x)(x - y), \end{aligned}$$

and the result follows. \square

We can then use the change in derivative of the necessary perturbation s to bound the number of iterations in the line search.

Theorem 5.2 Consider a function $h: [p, q] \rightarrow \mathbf{R}$, and a convex function $s: [p, q] \rightarrow \mathbf{R}$ such that $h+s$ is also convex. If the bisection method, Algorithm 3.1 evaluates h'_+ , the right-derivative, $k \geq 1$ times without terminating, then

$$\frac{h(q) - h(p)}{q - p} \geq \frac{s'_+(p) - s'_-(q)}{k}.$$

Proof We proceed by induction on the number of evaluations $k = 1, 2, 3, \dots$. The case $k = 1$ follows immediately from Lemma 5.1, by setting $x = p$ and $y = q$.

Suppose that the result holds for any points $p < q$ and for any number of evaluations no larger than k . Now consider an instance of the algorithm that completes $k+1$ evaluations. After the first evaluation, consider the midpoint $m = \frac{1}{2}(p+q)$. There are two possible cases, depending on whether or not

$$(5.3) \quad 2h(m) < h(p) + h(q).$$

We consider them in turn.

Suppose first that inequality (5.3) holds. After the first evaluation, the algorithm makes k further evaluations, beginning with the initial interval $[p, m]$. Hence, by the induction hypothesis,

$$\frac{h(m) - h(p)}{m - p} \geq \frac{s'_+(p) - s'_-(m)}{k}.$$

Since the algorithm did not terminate during the first two evaluations, we know $h'_+(m) \geq 0$. Applying Lemma 5.1 with $x = m$ and $y = q$ shows

$$\frac{h(q) - h(m)}{q - m} \geq s'_+(m) - s'_-(q).$$

Hence

$$\begin{aligned} s'_+(p) - s'_-(q) &\leq (s'_+(p) - s'_-(m)) + (s'_+(m) - s'_-(q)) \\ &\leq k \frac{h(m) - h(p)}{m - p} + \frac{h(q) - h(m)}{q - m} \\ &= (k - 1) \frac{h(m) - h(p)}{m - p} + \left(\frac{h(m) - h(p)}{m - p} + \frac{h(q) - h(m)}{q - m} \right) \\ &= (k - 1) \frac{h(m) - h(p)}{m - p} + 2 \frac{h(q) - h(p)}{q - p} \\ &< (k - 1) \frac{h(q) - h(p)}{q - p} + 2 \frac{h(q) - h(p)}{q - p} = (k + 1) \frac{h(q) - h(p)}{q - p}, \end{aligned}$$

as required.

The case where inequality (5.3) fails is similar. After the first bisection, the algorithm makes k further bisections, beginning with the initial interval $[m, q]$. Hence, by the induction hypothesis,

$$\frac{h(q) - h(m)}{q - m} \geq \frac{s'_+(m) - s'_-(q)}{k}.$$

Since the algorithm did not terminate during the first bisection, we know $h'_+(p) \geq 0$. Applying Lemma 5.1 with $x = p$ and $y = m$ shows

$$\frac{h(m) - h(p)}{m - p} \geq s'_+(p) - s'_-(m).$$

Hence

$$\begin{aligned} s'_+(p) - s'_-(q) &\leq (s'_+(p) - s'_-(m)) + (s'_+(m) - s'_-(q)) \\ &\leq \frac{h(m) - h(p)}{m - p} + k \frac{h(q) - h(m)}{q - m} \\ &= \left(\frac{h(m) - h(p)}{m - p} + \frac{h(q) - h(m)}{q - m} \right) + (k - 1) \frac{h(q) - h(m)}{q - m} \\ &= 2 \frac{h(q) - h(p)}{q - p} + (k - 1) \frac{h(q) - h(m)}{q - m} \\ &\leq 2 \frac{h(q) - h(p)}{q - p} + (k - 1) \frac{h(q) - h(p)}{q - p} = (k + 1) \frac{h(q) - h(p)}{q - p}, \end{aligned}$$

as required. \square

Definition 5.4 Given any interval $J \subset \mathbf{R}$, the *concave deviation* of a function $h: J \rightarrow \mathbf{R}$ is the infimum of the Lipschitz constants of convex functions $s: J \rightarrow \mathbf{R}$ such that the sum $h + s$ is also convex.

Consider, for example, a ρ -weakly convex function h , for some constant $\rho \geq 0$, meaning that the function $t \mapsto h(t) + \frac{\rho}{2}t^2$ is convex.

Proposition 5.5 Any ρ -weakly convex function $h: [p, q] \rightarrow \mathbf{R}$, for $\rho \geq 0$, has concave deviation at most $\frac{\rho}{2}(q - p)$.

Proof The function $s(t) = \frac{\rho}{2}(t - \frac{p+q}{2})^2$ is convex, with Lipschitz constant $\frac{\rho}{2}(q - p)$, and $h + s$ is also convex. \square

The concave deviation for functions that are not weakly convex may shrink more slowly than the length of the interval. For example, on the interval $[-\delta, \delta]$, the function $h(t) = -|t|^{\frac{3}{2}}$ has concave deviation $\frac{3}{2}\sqrt{\delta}$, and the piecewise linear function $-|\cdot|$ has concave deviation 1.

The concave deviation of a function $h: [p, q] \rightarrow \mathbf{R}$ that is difference-of-convex may not be finite. An example is the concave function $\sqrt{\cdot}$ on the interval $[0, 1]$. However, if h extends to a difference-of-convex function on an open interval containing the interval $[p, q]$, then its nonconvexity bound on $[p, q]$ must be finite, since we can write h as a difference of convex functions, each of which must be Lipschitz on $[p, q]$.

Consider any continuous semi-algebraic function $h: \mathbf{R} \rightarrow \mathbf{R}$. We can partition \mathbf{R} into finitely-many closed intervals J such that each restriction $h|_J$ is either convex or concave. In general, such a partition may not guarantee that h is difference-of-convex: an example is the function $x^{\frac{1}{3}}$. However, if h is also Lipschitz, then each convex or concave ingredient $h|_J$ extends to a corresponding convex or concave Lipschitz function on \mathbf{R} , and from these we can easily decompose h into a difference of convex Lipschitz functions. Thus all semi-algebraic functions on \mathbf{R} are difference-of-convex, with finite concave deviation on any bounded interval.

Corollary 5.6 (Line search complexity) *If a function $h: [p, q] \rightarrow \mathbf{R}$ has finite concave deviation M and average rate of decrease*

$$\sigma = -\frac{h(q) - h(p)}{q - p} > 0,$$

then the number of evaluations of h'_+ , the right-derivative, required before the bisection method, Algorithm 3.1, terminates is no more than $1 + \lfloor \frac{2M}{\sigma} \rfloor$.

Proof Suppose that the bisection method evaluates the left-derivative $k \geq 1$ times without terminating. Fix any value $M' > M$. By assumption, there exists a convex

function s with Lipschitz constant less than M' such that the sum $h + s$ is also convex. From Theorem 5.2, we deduce the inequalities

$$-\sigma \geq \frac{s'_+(p) - s'_-(q)}{k} > \frac{-2M'}{k}$$

so $k < \frac{2M'}{\sigma}$. Since M' was arbitrary, we deduce $k \leq \frac{2M}{\sigma}$, and hence $k \leq \lfloor \frac{2M}{\sigma} \rfloor$. The result follows. \square

Given an open interval I , consider a difference-of-convex function $h: I \rightarrow \mathbf{R}$. As observed in [7], such functions are characterized by having left and right derivatives everywhere, which furthermore are of bounded variation on every compact interval in I . Any such function h also has a second derivative D^2h in the distributional sense: in general it is a signed Radon measure on the interval I (see [3]).

We review briefly the underlying construction. Consider any convex function $s: I \rightarrow \mathbf{R}$. Its right derivative s'_+ is nondecreasing and right-continuous, and hence defines a nonnegative Radon measure D^2s on the interval I via the property

$$(D^2s)(p, q] = s'_+(q) - s'_+(p) \quad \text{for all } p < q \text{ in } I.$$

More generally, for any difference-of-convex function $h: I \rightarrow \mathbf{R}$, consider any convex function $s: I \rightarrow \mathbf{R}$ such that $h + s$ is also convex. The second derivative D^2h is just the signed measure $D^2(h + s) - D^2s$, which is independent of the choice of s . Convexity of h is characterized by the property $D^2h \geq 0$. More generally, the Hahn decomposition decomposes D^2h uniquely into a difference of nonnegative Radon measures supported on disjoint sets. These positive and negative parts satisfy

$$D^2h = (D^2h)^+ - (D^2h)^-.$$

Theorem 5.7 *If the function $h: [p, q] \rightarrow \mathbf{R}$ is difference-of-convex, then its concave deviation is*

$$\frac{1}{2}(D^2h)^-(p, q).$$

Proof For any convex function $s: [p, q] \rightarrow \mathbf{R}$ such that $h + s$ is also convex, the second derivatives satisfy

$$0 \leq D^2s = D^2(h + s) - D^2h \geq -D^2h,$$

so $D^2s \geq (D^2h)^-$. If s is M -Lipschitz on $[p, q]$, then

$$M \geq \max\{-s'_+(p), s'_-(q)\} \geq \frac{s'_-(q) - s'_+(p)}{2} = \frac{1}{2}(D^2s)(p, q) \geq \frac{1}{2}(D^2h)^-(p, q).$$

If the right-hand side is infinite, this completes the proof, so suppose that it is finite. Define a function $g: (p, q] \rightarrow \mathbf{R}$ by

$$g(t) = (D^2h)^-(p, t).$$

Then g is a nonnegative nondecreasing left-continuous function that is bounded above and $g(t) \downarrow 0$ as $t \downarrow p$. Now define a convex function $s: [p, q] \rightarrow \mathbf{R}$ by

$$s(t) = \int_p^t g(\tau) d\tau.$$

Then $s'_-(t) = g(t)$ for all $t \in (p, q]$ and $s'_+(p) = 0$. Furthermore we have

$$D^2(h + s) = D^2h + D^2s = D^2h + (D^2h)^- = (D^2h)^+ \geq 0,$$

so the sum $h + s$ is also convex. The function

$$t \mapsto \tilde{s}(t) = s(t) - \frac{1}{2}s'_-(q)t$$

is also convex, as is $h + \tilde{s}$, and the function \tilde{s} has Lipschitz constant

$$-\tilde{s}'_+(p) = \tilde{s}'_-(q) = \frac{1}{2}s'_-(q) = \frac{1}{2}(D^2h)^-(p, q).$$

This completes the proof. □

As an illustration, we have the following result.

Corollary 5.8 (Piecewise linear functions) *Consider any continuous piecewise linear function $h: [p, q] \rightarrow \mathbf{R}$, with m derivative discontinuities $t_1 < t_2 < \dots < t_m$ in the interval (p, q) . Define $t_0 = p$ and $t_{m+1} = q$, and let g_i be the value of the derivative on the interval (t_i, t_{i+1}) for $0 \leq i \leq m$. Then h has concave deviation*

$$\frac{1}{2} \sum_{i=1}^m (g_{i-1} - g_i)^+.$$

If h is L -Lipschitz, then this bound is no larger than $\lceil \frac{m}{2} \rceil L$.

Proof Denoting a unit point mass at the point t by δ_t , we have

$$D^2h = \sum_i (g_i - g_{i-1})\delta_{t_i},$$

and hence

$$(D^2h)^- = \sum_i (g_{i-1} - g_i)^+\delta_{t_i},$$

from which the claimed equation follows. The inequality is an easy consequence, using $|g_i| \leq L$ for each i . □

As an illustration, we present an example that underlines why the line search complexity estimate in Corollary 5.6 is the best that we can expect in general. In outline, while the functions $h: [p, q] \rightarrow \mathbf{R}$ that we consider satisfy $h(p) > h(q)$, their derivatives may often be positive.

Example 5.9 (Optimality of the line search) Consider any constant $M > 0$, and a set $T \subset [0, 1)$ of cardinality strictly less than $2M$. Then there exists a function $h: [0, 1] \rightarrow \mathbf{R}$ with concave deviation less than M , that satisfies $h(0) = 0$ and $h(1) = -1$, and has strictly positive derivative throughout T .

To see this, suppose first $T \subset (0, 1)$. (The case when T contains zero is an easy modification.) Enumerate the points in increasing order:

$$t_1 < t_2 < \cdots < t_k,$$

where $k < 2M$. Define $h(0) = 0$ and $h(1) = -1$. Fix any small $\gamma > 0$, and define

$$h(t_i - \gamma) = -t_i - \gamma^2 \quad \text{and} \quad h(t_i + \gamma) = -t_i + \gamma^2 \quad \text{for } i = 1, 2, \dots, k.$$

At intermediate points in $[0, 1]$, define h by linear interpolation. A quick calculation, using Corollary 5.8, shows that h has concave deviation

$$\frac{1}{2} \sum_{i=1}^{k-1} \left(\frac{t_{i+1} - t_i + 2\gamma^2}{t_{i+1} - t_i - 2\gamma} + \gamma \right) + \frac{1}{2} \left(\frac{1 - t_k + \gamma^2}{1 - t_k - \gamma} + \gamma \right) = \frac{k}{2} + O(\gamma) < M$$

providing that γ is sufficiently small.

Now consider any line search method applicable to functions $h: [0, 1] \rightarrow \mathbf{R}$ satisfying $h(0) = 0$ and $h(1) = -1$, relying on evaluations of the value h and the right derivative h'_+ at points chosen one-by-one, and terminating once a derivative is negative. Suppose that the method is guaranteed to terminate after at most k queries providing that the underlying function h has concave deviation strictly less than some given value $M > 0$. The example above proves $k \geq 2M$.

6 Multivariate functions

To understand the complexity of Algorithm 4.1 (nonsmooth minimization), we apply our analysis in the previous section to restrictions of multivariate objectives f to line segments. For any convex set $C \subset \mathbf{X}$, consider a function $f: C \rightarrow \mathbf{R}$. Given any length $\delta > 0$, let $\Lambda(\delta)$ denote the supremum over all points $x, y \in C$ with $|x - y| \leq \delta$ of the concave deviation for the function $h: [0, \delta] \rightarrow \mathbf{R}$ defined by

$$(6.1) \quad h(t) = f\left(x + \frac{t}{\delta}(y - x)\right).$$

We call the function $\Lambda: \mathbf{R}_{++} \rightarrow [0, +\infty]$ the *nonconvexity modulus* for f . The following illustration follows immediately from Proposition 5.5.

Proposition 6.2 *The nonconvexity modulus of any ρ -weakly convex function (for $\rho \geq 0$) satisfies*

$$\Lambda(\delta) \leq \frac{\rho\delta}{2}.$$

More generally, the function f is *difference-of-convex* when there exists a convex function $q: C \rightarrow \mathbf{R}$ such that $f + q$ is also convex.

Proposition 6.3 *Consider a convex set $C \subset \mathbf{X}$ and functions $f, q: C \rightarrow \mathbf{R}$ with both q and $f + q$ convex. If q is M -Lipschitz, then the nonconvexity modulus of f satisfies $\Lambda(\delta) \leq M$ for all $\delta > 0$.*

Proof Consider any points $x, y \in C$ with $|x - y| \leq \delta$, and the function h defined by equation (6.1). The function $s: [0, \delta] \rightarrow \mathbf{R}$ defined by

$$s(t) = q\left(x + \frac{t}{\delta}(y - x)\right)$$

is convex, and M -lipschitz, and $h + s$ is convex, so the concave deviation of h is no larger than M . The result follows. \square

As a consequence, we deduce the following result.

Corollary 6.4 *Consider any convex sets $C \subset C' \subset \mathbf{R}^n$, where C is nonempty and compact and C' is open, and any difference-of-convex function $f: C' \rightarrow \mathbf{R}$. Then the nonconvexity modulus of the restriction f_C is uniformly bounded: there exists a finite constant M such that $\Lambda(\delta) \leq M$ for all $\delta > 0$.*

Corollary 6.5 *If a function $f: \mathbf{X} \rightarrow \mathbf{R}$ is the difference $p - q$ between a convex function $p: \mathbf{X} \rightarrow \mathbf{R}$ and a polyhedral convex function $q: \mathbf{X} \rightarrow \mathbf{R}$, then the nonconvexity modulus of f is no larger than any Lipschitz constant for q .*

More generally, consider a continuous function $f: \mathbf{X} \rightarrow \mathbf{R}$ that is *semi-linear*, in the sense that \mathbf{X} is a finite union of polyhedra, on each of which the function f is affine. Any such function has a Lipschitz constant L , and furthermore a uniform upper bound m on the number of possible gradient discontinuities in any function of the form (6.1). By Corollary 5.8, we deduce that the nonconvexity modulus $\Lambda(\delta)$ is no larger than $\lceil \frac{m}{2} \rceil L$.

Returning to our analysis of Algorithm 4.1, we are ready for our main result.

Theorem 6.6 (complexity of minimization) *Given two constants $\delta, \epsilon > 0$, consider a convex set $C \subset \mathbf{X}$, a function $f: C \rightarrow \mathbf{R}$ that is bounded below, an associated L -bounded directional subgradient map $G: \mathbf{X}^2 \rightarrow \mathbf{X}$, and an initial point $x_0 \in C$ such that*

$$f(x) \leq f(x_0) \quad \text{and} \quad |y| \leq \delta \quad \Rightarrow \quad x + y \in C.$$

Suppose that f has finite nonconvexity modulus $\Lambda(\delta)$. Then Algorithm 4.1 (nonsmooth minimization) requires at most

$$\left\lceil \frac{3(f(x_0) - \min f)}{\delta \epsilon} \right\rceil \cdot \frac{16L^2}{\epsilon^2} \cdot \left(1 + \left\lfloor \frac{12\Lambda(\delta)}{\epsilon} \right\rfloor\right)$$

calls to Oracle 2.5 (directional subgradient) to find a point $x \in \mathbf{X}$ and a Goldstein subgradient $g \in \partial_\delta f(x)$ satisfying $|g| \leq \epsilon$.

Proof Corollary 5.6 with $\sigma = \frac{\epsilon}{6}$ shows that the bisection method requires at most

$$1 + \left\lfloor \frac{12\Lambda(\delta)}{\epsilon} \right\rfloor$$

oracle calls to terminate. Then, one further call returns the desired subgradient $g' \in \partial_\delta f(x)$ satisfying $\langle g', g \rangle < \frac{|g|^2}{2}$. Multiplying by the bound (4.4) on the number of line searches completes the proof. \square

7 Appendix: distributional second derivatives

We saw previously that the concave deviation of a univariate function h is determined by the negative part of its second distributional derivative D^2h . In our application, we consider functions h that are restrictions of the underlying objective f to line segments of fixed length δ . We would therefore expect the nonconvexity modulus of f to be related to its own distributional second derivative. Here, we explore that relationship informally.

Consider a locally Lipschitz function $f: \mathbf{R}^n \rightarrow \mathbf{R}$. The distributional derivative of f is an n -vector Df , entries of which are distributions — linear functionals on the space of smooth, compactly supported functions $g: \mathbf{R}^n \rightarrow \mathbf{R}$ (test functions), that are continuous with respect to uniform convergence on compact sets. We can define Df through the relationship

$$\langle u^T(Df), g \rangle = - \int f(u^T \nabla g)$$

for all vectors $u \in \mathbf{R}^n$ and test functions $g: \mathbf{R}^n \rightarrow \mathbf{R}$. However, by a suitable version of Rademacher’s Theorem [4, Section 6.2, Theorem 1], the gradient ∇f exists almost everywhere and is essentially bounded, and satisfies

$$\langle u^T(Df), g \rangle = \int g(u^T \nabla f).$$

In standard terminology [4], we can identify the classical gradient ∇f with both the distributional derivative Df and the “weak” derivative of f .

The second distributional derivative of f is an n -by- n matrix D^2f , entries of which are distributions. We can define D^2f through the relationship

$$\langle u^T(D^2f)v, g \rangle = - \int (u^T \nabla f)(v^T \nabla g)$$

for all vectors $u, v \in \mathbf{R}^n$ and test functions g . If f is smooth, then D^2f is just the matrix-valued measure with density $\nabla^2 f$. More generally we must consider D^2f as a distribution, but at least for *convex* functions we can be more specific: it is a positive-semidefinite-valued Radon measure ([3] and [4, Section 6.3]).

Example 7.1 (A piecewise linear function) Consider the convex function $f: \mathbf{R}^2 \rightarrow \mathbf{R}$ defined by $f(x) = x_1^+$. For any vectors $u, v \in \mathbf{R}^2$ and smooth, compactly supported function $g: \mathbf{R}^2 \rightarrow \mathbf{R}$, we have

$$\begin{aligned} \langle u^T(D^2f)v, g \rangle &= - \int_{x_1 > 0} u_1 v^T \nabla g(x) dx = - u_1 \int_{x_1 > 0} \operatorname{div}(g(x)v) dx \\ &= -u_1 \int_{\mathbf{R}} (-e_1)^T (g(\begin{bmatrix} 0 \\ y \end{bmatrix})v) dy = u_1 v_1 \int_{\mathbf{R}} g(\begin{bmatrix} 0 \\ y \end{bmatrix}) dy, \end{aligned}$$

by the Gauss-Green formula. Thus D^2f is the matrix $\begin{bmatrix} \mu & 0 \\ 0 & 0 \end{bmatrix}$, where the measure μ is related to Lebesgue measure λ via

$$(7.2) \quad \mu(S) = \lambda\{s \in \mathbf{R} : \begin{bmatrix} 0 \\ s \end{bmatrix} \in S\},$$

for all measurable subsets of $S \subset \mathbf{R}^2$.

For a more general understanding, we begin with the univariate case.

Example 7.3 (Univariate convex functions) Consider a convex function $f: \mathbf{R} \rightarrow \mathbf{R}$. For $0 < \gamma < 1$, we can construct a smooth approximation $h_\gamma: \mathbf{R} \rightarrow [0, 1]$ of the standard step function, with the following properties:

$$h_\gamma(t) = \begin{cases} 0 & (t \leq 0) \\ \gamma^2 & (t = \gamma^2) \\ 1 - \gamma^2 & (t = \gamma - \gamma^2) \\ 1 & (t \geq \gamma), \end{cases}$$

h is convex on $[0, \gamma^2]$, linear on $[\gamma^2, \gamma - \gamma^2]$, and concave on $[\gamma - \gamma^2, \gamma]$. For any interval $(p, q] \subset \mathbf{R}$, the test function $r_\gamma: \mathbf{R} \rightarrow [0, 1]$ defined by

$$r_\gamma(t) = \begin{cases} h_\gamma(t - p) & (t \leq p + \gamma^2) \\ 1 & (p + \gamma^2 \leq t \leq q) \\ 1 - h_\gamma(t - q) & (t \geq q) \end{cases}$$

converges pointwise to the characteristic function $\chi_{(p,q]}$ pointwise as $\gamma \downarrow 0$. By Dominated Convergence we deduce

$$\int r_\gamma d(D^2f) \rightarrow \int \chi_{(p,q]} d(D^2f) = (D^2f)(p, q].$$

But the left-hand side is

$$\begin{aligned}
-\int r'_\gamma f' &= -\int_p^{p+\gamma} \left(\frac{1}{\gamma} + O(1)\right) f' - \int_q^{q+\gamma} \left(-\frac{1}{\gamma} + O(1)\right) f' \\
&= \frac{f(q+\gamma) - f(q)}{\gamma} - \frac{f(p+\gamma) - f(p)}{\gamma} + O(\gamma) \\
&= f'_+(q) - f'_+(p) + O(\gamma)
\end{aligned}$$

as $\gamma \downarrow 0$. We thus reproduce our earlier definition:

$$(D^2 f)(p, q] = f'_+(q) - f'_+(p).$$

Clearly, this fact also holds for any difference-of-convex function $f: \mathbf{R} \rightarrow \mathbf{R}$.

The modulus of nonconvexity for a function $f: \mathbf{R}^n \rightarrow \mathbf{R}$, which we denoted $\Lambda(\delta)$ (for $\delta > 0$) is the supremum of the concave deviation of the restriction of f to line segments of the form

$$S = \{z + tw : 0 \leq t \leq \delta\}$$

for some point $z \in \mathbf{R}^n$ and unit direction $w \in \mathbf{R}^n$. That concave deviation is the measure of S under the negative part of the measure $D^2(f|_S)$. We would therefore like to compare the distributional second derivative of this restriction with the distributional second derivative $D^2 f$. As we see in the next result, we should focus specifically on the *directional distributional second derivative* $w^T(D^2 f)w$.

A simple approach is furnished by mollification. We fix a *mollifier* $\phi: \mathbf{R}^n \rightarrow \mathbf{R}$: a test function satisfying $\int \phi = 1$ and with the property that, as $\gamma \downarrow 0$, the function $\phi_\gamma(x) = \gamma^{-n} \phi(\frac{1}{\gamma}x)$ converges as a distribution to the Dirac delta function. Given a Radon measure μ on \mathbf{R}^n and any test function $g: \mathbf{R}^n \rightarrow \mathbf{R}$, we can define the convolution $g \star \mu: \mathbf{R}^n \rightarrow \mathbf{R}$ by

$$(g \star \mu)(y) = \int g(y - x) d\mu(x).$$

Theorem 7.4 (Chain rule via mollification) *Consider a locally Lipschitz function $f: \mathbf{R}^n \rightarrow \mathbf{R}$, a mollifier $\phi: \mathbf{R}^n \rightarrow \mathbf{R}$, and a direction $w \in \mathbf{R}^n$. Then for almost all points $z \in \mathbf{R}^n$, the function f is differentiable almost everywhere on the line $z + \mathbf{R}w$, and the distributional second derivative $D^2 h$ of the function defined by*

$$h(t) = f(z + tw) \quad (t \in \mathbf{R})$$

is the distributional limit, as $\gamma \downarrow 0$, of the convolution

$$t \mapsto \left(\phi_\gamma \star (w^T(D^2 f)w) \right)(z + tw).$$

Proof By Rademacher's Theorem and standard properties of convolutions [4, Section 4.2, Theorem 1(iv)], there exists a full measure set $\Omega \subset \mathbf{R}^n$ on which f is differentiable and the convolution $\phi_\gamma \star (w^T \nabla f)$ converges pointwise to the essentially bounded function $w^T \nabla f$. By Fubini's Theorem, for almost all points $z \in \mathbf{R}^n$, we have $z + tw \in \Omega$ for almost all $t \in \mathbf{R}$. Restricting attention to such z , consider any test function $g: \mathbf{R} \rightarrow \mathbf{R}$. Fubini's Theorem implies

$$\begin{aligned} \int_{t \in \mathbf{R}} g(t) \int_{x \in \mathbf{R}^n} \phi_\gamma(z + tw - x) d(w^T(D^2 f)w)(x) dt \\ = \int_x \left(\int_t g(t) \phi_\gamma(z + tw - x) dt \right) d(w^T(D^2 f)w)(x) \\ = \int_x w^T \nabla f(x) \left(w^T \int_t g(t) \nabla \phi_\gamma(z + tw - x) dt \right) dx. \end{aligned}$$

(We can interchange the order of differentiation and integration since the test functions g and ϕ_γ are well behaved.) Rewriting, integrating by parts, and using Fubini's Theorem and Dominated Convergence again, we obtain

$$\begin{aligned} \int_x w^T \nabla f(x) \int_t g(t) w^T \nabla \phi_\gamma(z + tw - x) dt dx \\ = \int_x w^T \nabla f(x) \int_t g(t) \frac{d}{dt} \phi_\gamma(z + tw - x) dt dx \\ = - \int_x w^T \nabla f(x) \int_t g'(t) \phi_\gamma(z + tw - x) dt dx \\ = - \int_t g'(t) \int_x \phi_\gamma(z + tw - x) w^T \nabla f(x) dx dt \\ = - \int_t g'(t) (\phi_\gamma \star (w^T \nabla f))(z + tw) dt \\ \rightarrow - \int_t g'(t) w^T \nabla f(z + tw) dt = - \int g' h' = \int g d(D^2 h), \end{aligned}$$

as desired. \square

In this result, the effect of the convolution is to focus attention on the line through the point z in the direction w . In informal language, we deduce that *the modulus of nonconvexity* $\Lambda(\delta)$ *is determined by the concentration of the negative parts of the measures* $w^T(D^2 f)w$ *around line segments of length* δ *in unit directions* w .

For more intuition on directional distributional second derivatives of the form $w^T(D^2 f)w$, let us consider a convex function $f: \mathbf{R}^2 \rightarrow \mathbf{R}$. After a suitable choice of basis we can suppose that w is the first unit vector e^1 and therefore consider the Radon measure $(Df)_{11}$. To understand this measure, consider the integral

$$\int_{\mathbf{R}^2} r_\gamma(x_1) g(x_2) d(D^2 f)_{11}(x)$$

for the function r_γ of Example 7.3 and any test function $g: \mathbf{R} \rightarrow \mathbf{R}$. As $\gamma \downarrow 0$, we observe

$$\begin{aligned} \int_{\mathbf{R}^2} r_\gamma(x_1)g(x_2) d(e_1^T(D^2f)e_1)(x) &= - \int_{\mathbf{R}^2} (e_1^T \nabla f(x)) \cdot (e_1^T \nabla (r_\gamma(x_1)g(x_2))) dx \\ &= - \int_{\mathbf{R}^2} \frac{\partial f}{\partial x_1} r'_\gamma(x_1)g(x_2) dx_1 dx_2 \\ &\rightarrow \int (f'([\begin{smallmatrix} q \\ t \end{smallmatrix}]; [\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}])) - f'([\begin{smallmatrix} p \\ t \end{smallmatrix}]; [\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}])) g(t) dt. \end{aligned}$$

More generally, this argument suggests, loosely, that $w^T(D^2f)w$ measures the variation of the directional derivative $f'(\cdot; w)$ along the direction w .

References

- [1] J. Bolte, A. Daniilidis, and A.S. Lewis. Tame functions are semismooth. *Mathematical Programming*, 117:5–19, 2009.
- [2] D. Davis, D. Drusvyatskiy, Yin Tat Lee, S. Padmanabhan, and Guanghao Ye. A gradient sampling method with complexity guarantees for lipschitz functions in high and low dimensions. [arXiv:2112.06969v2](https://arxiv.org/abs/2112.06969v2), 2022.
- [3] R.M. Dudley. On second derivatives of convex functions. *Math. Scand*, 41:159–174, 1977.
- [4] L.C. Evans and R.F. Gariepy. *Measure Theory and Fine Properties of Functions*. CRC Press, Boca Raton, FL, 1992.
- [5] F. Facchinei and J.-S. Pang. *Finite-dimensional variational inequalities and complementarity problems. Vol. II*. Springer Series in Operations Research. Springer-Verlag, New York, 2003.
- [6] A.A. Goldstein. Optimization of Lipschitz continuous functions. *Math. Programming*, 13(1):14–22, 1977.
- [7] P. Hartman. On functions representable as a difference of convex functions. *Pacific J. Math.*, 9:707–713, 1959.
- [8] R. Henrion and J. Outrata. A subdifferential condition for calmness of multifunctions. *J. Math. Anal. Appl.*, 258:110–130, 2001.
- [9] M.I. Jordan, Tianyi Lin, and M. Zampetakis. On the complexity of deterministic nonsmooth and nonconvex optimization. [arXiv:2209.12463](https://arxiv.org/abs/2209.12463), 2022.
- [10] N. Mahdavi-Amiri and R. Yousefpour. An effective nonsmooth optimization algorithm for locally Lipschitz functions. *J. Optim. Theory Appl.*, 155:180–195, 2012.

- [11] R. Mifflin. An algorithm for constrained optimization with semismooth functions. *Math. Oper. Res.*, 2:191–207, 1977.
- [12] G. Kornowski O. Shamir. On the complexity of finding small subgradients in nonsmooth optimization. arXiv:2209.10346, 2022.
- [13] Lai Tian and Anthony Man-Cho So. Computing Goldstein (ϵ, δ) -stationary points of Lipschitz functions in $\tilde{O}(\epsilon^{-3}\delta^{-1})$ iterations via random conic perturbation. arxiv.org/abs/2112.09002.
- [14] Jingzhao Zhang, Hongzhou Lin, S. Jegelka, S. Sra, and A. Jadbabaie. Complexity of finding stationary points of nonconvex nonsmooth functions. In *ICML Proceedings*, 2020.