

On Nesterov's Nonsmooth Chebyshev-Rosenbrock Functions

Mert Gurbuzbalaban and Michael L. Overton *

January 24, 2011

Abstract

We discuss two nonsmooth functions on \mathbf{R}^n introduced by Nesterov. We show that the first variant is partly smooth in the sense of Lewis [Lew03] and that its only stationary point is the global minimizer. In contrast, we show that the second variant has 2^{n-1} Clarke stationary points, none of them local minimizers except the global minimizer. Furthermore, its only stationary point in the sense of Mordukhovich [RW98] is the global minimizer. Nonsmooth optimization algorithms from multiple starting points generate iterates that approximate all 2^{n-1} Clarke stationary points, not only the global minimizer, but it remains an open question as to whether the other Clarke stationary points are actually points of attraction for optimization algorithms.

Keywords. Nonsmooth optimization, optimization algorithms, variational analysis.

1 Introduction

In 2008, Nesterov [Nes08] introduced the following *smooth* (differentiable, in fact polynomial) function on \mathbf{R}^n :

$$\tilde{f}(x) = \frac{1}{4}(x_1 - 1)^2 + \sum_{i=1}^{n-1} (x_{i+1} - 2x_i^2 + 1)^2.$$

*Courant Institute of Mathematical Sciences, New York University, 251 Mercer Street, New York, NY 10012, USA. (mert@cims.nyu.edu and overton@cs.nyu.edu). The work of both authors was supported in part by the National Science Foundation under Grant DMS-1016325.

The only stationary point of f is the global minimizer $x^* = [1, 1, \dots, 1]^T$. Consider the point $\hat{x} = [-1, 1, 1, \dots, 1]^T$ and the manifold

$$\mathcal{M} = \{x : x_{i+1} - 2x_i^2 + 1 = 0, \quad i = 1, \dots, n-1\}$$

which contains both x^* and \hat{x} . For $x \in \mathcal{M}$,

$$x_{i+1} = 2x_i^2 - 1 = T_2(x_i) = T_{2^i}(x_1), \quad i = 1, \dots, n-1,$$

where $T_k(x)$ denotes the k th Chebyshev polynomial of the first kind [Sze39, Section 2.4].

The function \tilde{f} is the sum of a quadratic term and a nonnegative sum whose zero set is the manifold \mathcal{M} . Minimizing \tilde{f} is equivalent to minimizing the first quadratic term on \mathcal{M} . Standard optimization methods, such as Newton's method and the BFGS quasi-Newton method, when applied to minimize \tilde{f} and initiated at \hat{x} , generate iterates that, as in the well known Rosenbrock example [GMW81] and its extensions [NW06], approximately "track" \mathcal{M} to approach the minimizer. The iterates do not track \mathcal{M} exactly, but because they typically follow this highly oscillatory manifold fairly closely, the tracking process requires many iterations. To move from \hat{x} to x^* along \mathcal{M} *exactly* would require x_n to trace the graph of the 2^{n-1} th Chebyshev polynomial, which has $2^{n-1} - 1$ extrema in $(-1, 1)$, as x_1 increases from -1 to 1 . Hence, \tilde{f} is a challenging test problem for optimization methods.

Nesterov also introduced two *nonsmooth* variants of \tilde{f} , the first being

$$\hat{f}(x) = \frac{1}{4}(x_1 - 1)^2 + \sum_{i=1}^{n-1} |x_{i+1} - 2x_i^2 + 1|. \quad (1)$$

A contour plot of this function when $n = 2$ is shown on the left side of Figure 1. Again, the unique global minimizer is x^* . Like \tilde{f} , the function \hat{f} is the sum of a quadratic term and a nonnegative sum whose zero set is the manifold \mathcal{M} , so, as previously, minimizing \hat{f} is equivalent to minimizing the first quadratic term on \mathcal{M} , but unlike \tilde{f} , the function \hat{f} is not differentiable at points in \mathcal{M} . However, as we show in the next section, \hat{f} is *partly smooth* with respect to \mathcal{M} , in the sense of [Lew03], at points in \mathcal{M} . It follows that, like \tilde{f} , the function \hat{f} has only one stationary point — the global minimizer x^* — where by stationary point we mean both in the sense of Clarke and of Mordukhovich. It follows that the global minimizer x^* is the only local minimizer of \hat{f} .

The second nonsmooth variant is

$$f(x) = \frac{1}{4}|x_1 - 1| + \sum_{i=1}^{n-1} |x_{i+1} - 2|x_i| + 1|. \quad (2)$$

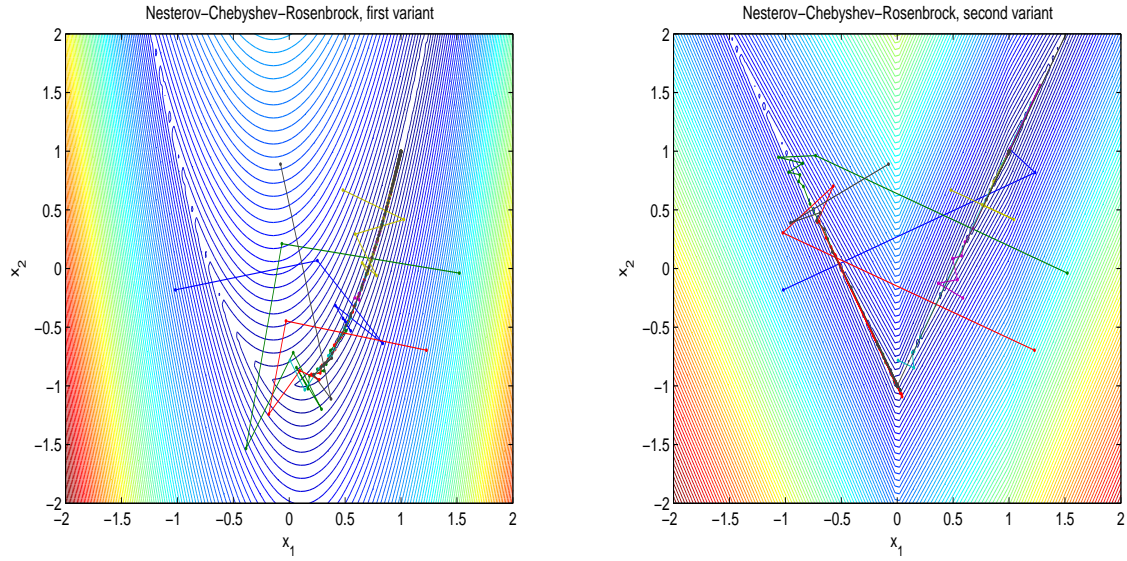


Figure 1: Contour plots for Nesterov’s first (left) and second (right) nonsmooth Chebyshev-Rosenbrock functions \hat{f} and f respectively, with $n = 2$. Points connected by line segments show the iterates generated by the BFGS method (see Section 3) initialized at 7 different randomly generated starting points. For the first variant \hat{f} , convergence always takes place to the only Clarke stationary point: the global minimizer $x^* = [1, 1]^T$. For the second variant f , some runs of BFGS generate iterates that approximate the nonminimizing Clarke stationary point $[0, -1]^T$ while others converge to the minimizer $[1, 1]^T$.

Again, the unique global minimizer is x^* . Consider the set

$$S = \{x : x_{i+1} - 2|x_i| + 1 = 0, \quad i = 1, \dots, n-1\}. \quad (3)$$

Minimizing f is equivalent to minimizing its first term on S . Like \mathcal{M} , the set S is highly oscillatory, but it has “corners”: it is not a manifold around any point x where any of the components x_1, \dots, x_{n-1} vanishes. For example, consider the case $n = 2$, for which a contour plot is shown on the right side of Figure 1. It is easy to verify that the point $[0, -1]^T$ is Clarke stationary (zero is in the convex hull of gradient limits at the point), but not a local minimizer ($[1, 2]^T$ is a direction of linear descent from $[0, -1]^T$). We will show in the next section that, in fact, f has 2^{n-1} Clarke stationary points, that the only local minimizer is the global minimizer x^* , and furthermore that the only stationary point in the sense of Mordukhovich is x^* .

In the next section, we define stationarity in both senses and present the main results. In Section 3, we report on numerical experiments showing the behavior of nonsmooth minimization algorithms on these functions.

2 Main Results

Before stating our main results, we will need the following well-known definitions. The Clarke subdifferential or generalized gradient [Cla83] of a locally Lipschitz function on a finite-dimensional space can be defined as follows [BL05, Theorem 6.2.5]. Let ∇ denote gradient.

Definition 1. (*Clarke subdifferential*) Consider a function $\phi : \mathbf{R}^n \rightarrow \mathbf{R}$ and a point $x \in \mathbf{R}^n$, and assume that ϕ is locally Lipschitz around x . Let $A \subset \mathbf{R}^n$ be an arbitrary set with measure zero. Then the Clarke subdifferential of ϕ at x is

$$\partial^C \phi(x) = \text{conv} \left\{ \lim_{m \rightarrow \infty} \nabla \phi(x^m) : x^m \rightarrow x, x^m \notin A \right\}. \quad (4)$$

Note that by Rademacher’s Theorem [EG92], locally Lipschitz functions are differentiable almost everywhere so A can be chosen as the set of points on which ϕ is not differentiable.

As expounded in [RW98], the Mordukhovich [Mor76] subdifferential is defined as follows.

Definition 2. (*Mordukhovich subdifferential*) Consider a function $\phi : \mathbf{R}^n \rightarrow \mathbf{R}$ and a point $x \in \mathbf{R}^n$. A vector $v \in \mathbf{R}^n$ is a regular subgradient of ϕ at x (written $v \in \hat{\partial} \phi(x)$) if

$$\liminf_{\substack{z \rightarrow x \\ z \neq x}} \frac{\phi(z) - \phi(x) - \langle v, z - x \rangle}{|z - x|} \geq 0,$$

where $\langle \cdot, \cdot \rangle$ is the usual inner product on \mathbf{R}^n . A vector $v \in \mathbf{R}^n$ is a Mordukhovich subgradient of ϕ at x (written $v \in \partial^M \phi(x)$) if there exist sequences x^m and v^m in \mathbf{R}^n satisfying

$$\begin{aligned} x^m &\rightarrow x \\ \phi(x^m) &\rightarrow \phi(x) \\ v^m &\in \hat{\partial} \phi(x^m) \\ v^m &\rightarrow v. \end{aligned}$$

We say that ϕ is *Clarke stationary* at x if $0 \in \partial^C \phi(x)$. Similarly, ϕ is *Mordukhovich stationary* at x if $0 \in \partial^M \phi(x)$. For a locally Lipschitz function ϕ , we have [RW98, Theorem 8.49]

$$\partial^C \phi(x) = \text{conv} \{ \partial^M \phi(x) \}. \quad (5)$$

We will need the concept of regularity (also known as subdifferential regularity or Clarke regularity) [RW98], which can be characterized for locally Lipschitz functions as follows [Lew06, Theorem 6.10].

Definition 3. (*regularity*) A locally Lipschitz function $\phi : \mathbf{R}^n \rightarrow \mathbf{R}$ is regular at a point x if and only if its ordinary directional derivative satisfies

$$\phi'(x; d) = \limsup_{z \rightarrow x} \langle \nabla \phi(z), d \rangle$$

for every direction $d \in \mathbf{R}^n$.

One consequence of regularity of ϕ at a point x is that $\partial^C \phi(x) = \partial^M \phi(x)$ [BLO02, Proposition 4.1(iii)] and another is that the Clarke stationarity condition $0 \in \partial^C \phi(x)$ is equivalent to the first-order optimality condition $\phi'(x, d) \geq 0$ for all directions d [SY06, Section 14.1].

A property that will be central in our analysis is *partial smoothness* [Lew03].

Definition 4. A function ϕ is partly smooth at x relative to a manifold \mathcal{X} containing x if

1. its restriction to \mathcal{X} , denoted by $\phi|_{\mathcal{X}}$, is twice continuously differentiable at x ,
2. at every point close to $x \in \mathcal{X}$, the function ϕ is regular and has a Mordukhovich subgradient,

3. par $\{\partial^M \phi(x)\}$, the subspace parallel to the affine hull of the subdifferential of ϕ at x , is the normal subspace to \mathcal{X} at x , and
4. the Mordukhovich subdifferential map $\partial^M \phi$ is continuous at x relative to \mathcal{X} .

We illustrate the definition by proving that \hat{f} is partly smooth.

Lemma 1. *Nesterov's first nonsmooth Chebyshev-Rosenbrock function \hat{f} , defined in (1), is partly smooth with respect to \mathcal{M} at all points in \mathcal{M} .*

Proof. For each $i \in \{1, \dots, n-1\}$, consider the function $h_i(x) = |x_{i+1} - 2x_i^2 + 1|$ and the manifold $\mathcal{M}_i := \{x : H_i(x) := x_{i+1} - 2x_i^2 + 1 = 0\}$. By the chain rule [RW98, Proposition 10.5], h_i is globally regular as a composition of two regular functions and we have

$$\partial^M h_i(x) = \nabla H_i(x) [\{\partial^M |\cdot|\} (x_{i+1} - 2x_i^2 + 1)]$$

for any $x \in \mathbf{R}^n$. We have the normal space $N_{\mathcal{M}_i} = \text{Range}(\nabla H_i(x))$ [RW98, Ex. 6.8] which is clearly parallel to the subdifferential $\partial^M h_i(x)$. Since $h_i|_{\mathcal{M}_i} = 0$ is smooth and $\partial^M h_i$ is continuous at x relative to \mathcal{M}_i , it follows from Definition 4 that h_i is partly smooth with respect to the manifold \mathcal{M}_i . We conclude from [Lew03, Corollary 4.6] and [Lew03, Corollary 4.7] that $\hat{f}(x) = \frac{1}{4}(x_1 - 1)^2 + \sum_{i=1}^{n-1} h_i(x)$ is partly smooth with respect to the manifold $\mathcal{M} = \cap_{i=1}^{n-1} \mathcal{M}_i$ at all points in \mathcal{M} . \square

It follows that \hat{f} has only one stationary point.

Theorem 1. *Nesterov's first nonsmooth Chebyshev-Rosenbrock function \hat{f} is Clarke stationary or Mordukhovich stationary only at the unique global minimizer $x^* = [1, 1, \dots, 1]^T$.*

Proof. If $x \notin \mathcal{M}$, then \hat{f} is smooth and nonstationary at x as the partial derivative of \hat{f} with respect to x_n is ± 1 . On the other hand, if $x \in \mathcal{M}$, then the restricted function $\hat{f}|_{\mathcal{M}} = \frac{1}{4}(x_1 - 1)^2$ is smooth and has a critical point only at the global minimizer x^* . If $x \in \mathcal{M}$ and $x \neq x^*$, it follows from [Lew03, Proposition 2.4] that $0 \notin \text{aff } \partial^M \hat{f}(x)$. Thus, $0 \notin \partial^M \hat{f}(x)$. By regularity, we have $\partial^C \hat{f}(x) = \partial^M \hat{f}(x)$ and the result follows. \square

The main results of the paper concern Nesterov's second nonsmooth example. For this we will need the usual "sign" function:

$$\text{sign}(x) = \begin{cases} 1 & : x > 0, \\ 0 & : x = 0, \\ -1 & : x < 0. \end{cases}$$

We start by stating a lemma whose proof is straightforward.

Lemma 2. *Let S be defined as in (3). There are $2^{n-1} - 1$ points in S such that $x_j = 0$ for some $j < n$. Let \bar{x} be such a point. Then \bar{x}_i takes non-integer values between -1 and 1 for $i < j$, $\bar{x}_i = 0$ for $i = j$, $\bar{x}_i = -1$ for $i = j + 1$ and $\bar{x}_i = 1$ if $n \geq i > j + 1$. In particular, $\bar{x}_1 < 1$ (with $\bar{x}_1 = 0$ if $j = 1$).*

Theorem 2. *Nesterov's second nonsmooth Chebyshev-Rosenbrock function f , defined in (2), is Clarke stationary at the $2^{n-1} - 1$ points in the set S with a vanishing x_j for some $j < n$.*

Proof. Let \bar{x} be such a point. Then, using Lemma 2, we see that around \bar{x} the function $\frac{|x_1-1|}{4}$ is equal to $\frac{1-x_1}{4}$ and furthermore $\bar{x}_i \neq 0$ if $i \neq j$. These observations allow us to write f in a simpler form eliminating most of the absolute values. We first prove the case $j = n - 1$. We will show that in an arbitrarily small neighborhood of \bar{x} the gradient vector (if defined) can take arbitrary signs in each coordinate. This will ensure that $0 \in \partial^C f(\bar{x})$ by (4).

Around \bar{x} , the function $f(x)$ may be rewritten as

$$f(x) = \frac{1-x_1}{4} + |x_2 + 2c_1x_1 + 1| + \dots + |x_{n-1} + 2c_{n-2}x_{n-2} + 1| + |x_n - 2|x_{n-1}| + 1| \quad (6)$$

where $c_i = -\text{sign}(\bar{x}_i)$, $i = 1, 2, \dots, n - 2$, depend only on the point \bar{x} and are fixed. (Note that $\bar{x}_i \neq 0$ for $i < j = n - 1$). Since $\bar{x} \in S$ and $\bar{x}_{n-1} = 0$, all the absolute value terms appearing in (6) are equal to 0 at \bar{x} . By the continuity of f at \bar{x} , it is possible to find points x arbitrarily close to \bar{x} such that each of the absolute value terms evaluated at x has arbitrary sign and at those points

$$\nabla f(x) = \left[-\frac{1}{4} + 2c_1d_1, d_1 + 2c_2d_2, \dots, d_{n-2} + 2c_{n-1}d_{n-1}, d_{n-1} \right]^T$$

where $c_{n-1} := -\text{sign}(x_{n-1})$ and each of d_1, d_2, \dots, d_{n-1} can be chosen to be $+1$ or -1 as desired. Hence, it is possible to have $\nabla f(x)$ in any of the 2^n quadrants of \mathbb{R}^n . Consequently, 0 lies in the convex combination of these gradient vectors and we conclude from (4) that $0 \in \partial^C f(\bar{x})$.

The case $j < n - 1$ is handled similarly. For a choice of x around \bar{x} , we get

$$\nabla f(x) = \left[-\frac{1}{4} + 2c_1d_1, d_1 + 2c_2d_2, \dots, d_{j-1} + 2c_jd_j, d_j + 2d_{j+1}, d_{j+1} - 2d_{j+2}, \dots, d_{n-1} \right]^T$$

where $c_i = -\text{sign}(x_i)$, $i = 1, 2, \dots, j - 1$, are fixed (when $j > 1$) and $c_j = -\text{sign}(x_j)$, d_1, d_2, \dots, d_{n-1} are free parameters to choose from $\{-1, 1\}$.

Suppose $d_j = d_j^0$, $d_{j+1} = d_{j+1}^0, \dots, d_{n-1} = d_{n-1}^0$ are fixed. By choosing c_j, d_1, \dots, d_{j-1} appropriately, the signs of the first j components of $\nabla f(x)$ vector can be chosen to be positive or negative. Thus, by convexity,

$$[0, \dots, 0, d_j^0 + 2d_{j+1}^0, d_{j+1}^0 - 2d_{j+2}^0, \dots, d_{n-1}^0]^T \in \partial^C f(\bar{x}).$$

Choosing $d_j = -d_j^0, d_{j+1} = -d_{j+1}^0, \dots, d_{n-1} = -d_{n-1}^0$, we have

$$[0, \dots, 0, -(d_j^0 + 2d_{j+1}^0), -(d_{j+1}^0 - 2d_{j+2}^0), \dots, -d_{n-1}^0]^T \in \partial^C f(\bar{x}).$$

and so by convexity $0 \in \partial^C f(\bar{x})$, completing the proof. \square

The following theorem characterizes all the stationary points of f in the sense of both subdifferentials.

Theorem 3. *Nesterov's second nonsmooth Chebyshev-Rosenbrock function f is Mordukhovich stationary only at the global minimizer $x^* = [1, 1, \dots, 1]^T$. Furthermore, f is Clarke stationary only at x^* and the $2^{n-1} - 1$ points in S with a vanishing x_j for some $j < n$. None of the Clarke stationary points of f except the global minimizer are local minimizers of f and there exists a direction of linear descent from each of these points.*

Proof. If $x \notin S$, f is smooth at x and we have $0 \notin \partial^M f(x) = \partial^C f(x) = \{\nabla f(x)\}$ since the partial derivative of f with respect to x_n at x is ± 1 .

When $x = x^* \in S$, we have $0 \in \hat{\partial}f(x) \subset \partial^M f(x) \subset \partial^C f(x)$. If $x \in S$, $x \neq x^*$ ($x_1 \neq 1$) and $x_j \neq 0$ for $j = 1, 2, \dots, n-1$, then the set S is a manifold around x . The function f is partly smooth with respect to S at x , with $f|_S(x) = \frac{|1-x_1|}{4}$, the restriction of f to S , smooth around x , and x is not a critical point of $f|_S$. It follows from [Lew03, Proposition 2.4] that $0 \notin \text{aff} \{\partial^M f(x)\}$. This implies directly that $0 \notin \partial^M f(x)$ and $0 \notin \partial^C f(x) = \text{conv} \{\partial^M f(x)\}$, using (5).

The remaining case is when $x \in S$ is such that $x_j = 0$ for some $j < n$. We have $x_1 < 1$. Let $\delta > 0$ be small and x^δ be the unique point near x such that $x^\delta \in S$ and $x_1^\delta = x_1 + \delta$. It follows from the definition of S that $x^\delta = x + \delta v$ where v is a fixed vector independent of $\delta > 0$ for δ sufficiently small. Since $f|_S(x) = \frac{1-x_1}{4}$ around x , we have $f(x^\delta) = f(x + \delta v) = f(x) - \frac{1}{4}\delta < f(x)$ which shows that v is a direction of linear descent. Furthermore, we have $0 \notin \hat{\partial}f(x)$ since the existence of the descent direction at \bar{x} implies

$$\liminf_{\substack{z \rightarrow \bar{x} \\ z \neq x}} \frac{f(z) - f(x)}{|z - x|} \leq \liminf_{\delta \downarrow 0} \frac{f(x + \delta v) - f(x)}{\delta |v|} = -\frac{1}{4|v|} < 0.$$

We want to prove that $0 \notin \partial^M f(x)$. This requires an investigation of the regular subdifferential $\hat{\partial}f(y)$ for y near x . Let y be a point near x , $y \neq x$. We have $x_j = 0$, so we distinguish two cases: $y \notin S$, $\{y \in S \text{ and } y_j \neq 0\}$. (If $y \in S$ and $y_j = 0$, then, for y to be near x , we would need $y = x$.)

1. $y \notin S$: $\nabla f(y)$ exists, we have $\hat{\partial}f(y) = \{\nabla f(y)\}$ and the n -th coordinate of $\nabla f(y)$ is ± 1 . This shows that there exists no sequence $y^m \rightarrow x$ such that $y^m \notin S$ for all m with $\hat{\partial}f(y^m) = \{\nabla f(y^m)\} \ni v^m \rightarrow 0$.
2. $y \in S$ and $y_j \neq 0$: We have, for y sufficiently close to x , that $y_k \neq 0$ for $k = 1, \dots, n$ and

$$S = \{x : F_i(x) = 0, \quad i = 1, \dots, n-1\}$$

where $F_i(x) = x_{i+1} - 2|x_i| + 1$ is smooth at y . Hence, S is a manifold around y and it is easy to see that f is partly smooth at y with respect to S . The restricted function $f|_S(x) = \frac{1-x_1}{4}$ is smooth at y and since $y_1 < 1$, y is not a critical point of f , so from [Lew03, Proposition 2.4] we conclude that $0 \notin \text{aff}\{\partial^M f(y)\}$ which leads to $0 \notin \hat{\partial}f(y)$. Furthermore, by [Lew03, Proposition 2.2], we have

$$\hat{\partial}f(y) \subset \nabla g(y) + N_S(y) \tag{7}$$

where $g(x) = \frac{1-x_1}{4}$ and $N_S(y)$ is the normal space to S at y . The normal space to S at y coincides with the normal cone to S at y so by [RW98, Ex. 6.8]

$$N_S(y) = \text{Range}(\nabla F)$$

where

$$\nabla F(y)^T = \left[\frac{\partial F_i}{\partial x_j}(y) \right]_{i,j=1}^{n-1,n} \in \mathbf{R}^{(n-1) \times n}$$

is the Jacobian matrix. We have

$$\nabla F(y) = \begin{bmatrix} -2\text{sign}(y_1) & & & & \\ 1 & -2\text{sign}(y_2) & & & \\ & 1 & \ddots & & \\ & & \ddots & -2\text{sign}(y_{n-1}) & \\ & & & 1 & \end{bmatrix} \in \mathbf{R}^{n \times (n-1)} \tag{8}$$

and $\nabla g(y) = [-1/4, 0, \dots, 0]^T$. From (7) and (8), we see that $0 \in \hat{\partial}f(y)$ is possible only if $[1/4, 0, \dots, 0]^T \in N_S(y)$. A straightforward calculation shows that this is impossible. We conclude that $0 \notin \hat{\partial}f(y)$.

The next step is to investigate the possible limits of $v^m \in \hat{\partial}f(y^m)$ as the sequence $y^m \in S$ approaches x . Let y^m be a sequence such that $y^m \rightarrow x$, $y^m \in S$ and $y^m \neq x$ for all m (this implies $y_j^m \neq 0$ for all m as before). Without loss of generality, assume $y_j^m > 0$ for all m . For fixed $k \in \{1, 2, \dots, n\}$, the quantity $\text{sign}(y_k^m)$ does not depend on m and is nonzero. Thus, $G := \nabla F(y^m)$ does not depend on m . Let $v \in \mathbf{R}^n$ be such that $\hat{\partial}f(y^m) \ni v^m \rightarrow v$. From (7), we have $v^m = [-1/4, 0, \dots, 0]^T + Gc^m$ for some $c^m \in \mathbf{R}^{n-1}$. Since $v^m \rightarrow v$ and G has full rank, we have $c^m \rightarrow c \in \mathbf{R}^{n-1}$ and $v = [-1/4, 0, \dots, 0]^T + Gc$. As previously, $v = 0$ is impossible.

We conclude that $0 \notin \partial^M f(x)$. Since we already know from Theorem 2 that $0 \in \partial^C f(x)$, this completes the proof of the theorem. \square

It follows immediately from Theorem 3 that f is not regular at the $2^{n-1} - 1$ non-locally-minimizing Clarke stationary points of f : see the comments after Definition 3.

3 Numerical Experiments

Nesterov has observed that Newton’s method with an inexact line search, when applied to minimize the *smooth* function \tilde{f} initiated at \hat{x} , takes many iterations to reduce $\tilde{f}(x)$ below a small tolerance ϵ . Indeed, the number of iterations is typically exponential in n . Our experimental observations are mainly obtained using the well known BFGS quasi-Newton method with a line search based on the Armijo and “weak Wolfe” conditions, which, as explained in [LO10], is applicable to the nonsmooth variants \hat{f} and f as well as to the smooth function \tilde{f} . For the results reported below, we used a publicly available MATLAB implementation¹ run using standard IEEE floating point arithmetic [Ove01].

For a smooth but nonconvex function such as \tilde{f} , there is no theorem known that guarantees that the BFGS iterates will converge to the only stationary point x^* , and pathological counterexamples have been constructed [Dai02, Mas04], although, unlike \tilde{f} , these are not analytic. However, it is widely accepted that BFGS generally provides sequences converging to local minimizers of smooth, nonconvex functions [LF01], so it is not surprising that this is the case for \tilde{f} , with superlinear convergence to x^* in the limit. As with Newton’s method, many iterations are required. For $n = 8$, starting

¹<http://www.cs.nyu.edu/overton/software/hanso>

at \hat{x} and with the initial inverse Hessian approximation H set to the identity matrix I , the BFGS method requires about 6700 iterations to reduce \tilde{f} below 10^{-15} , and for $n = 10$, nearly 50,000 iterations are needed.

As mentioned above, the BFGS algorithm is very effective for nonsmooth functions too [LO10], but convergence is generally much slower than on smooth functions, with asymptotic linear, not superlinear, convergence. To apply BFGS to Nesterov’s first nonsmooth variant \hat{f} , we cannot use \hat{x} for the initial point as the method immediately breaks down, \hat{f} being nondifferentiable at \hat{x} . Instead, we initialize x randomly, retaining the identity matrix for initializing H . The left panel of Figure 1 shows the path of the iterates generated by BFGS for the case $n = 2$ using 7 random starting points, plotted in 7 different colors. All sequences of iterates converge to the global minimizer $x^* = [1, 1]^T$. (Note that the points plotted later overwrite many of those plotted earlier near x^*). We find that we can also usually minimize \hat{f} reasonably accurately (reducing \hat{f} to about 10^{-8} before breakdown occurs) when $n = 3$, but not when $n = 4$, for which the method typically breaks down far from x^* . By breakdown we mean that the line search fails to return a point satisfying the Armijo and weak Wolfe conditions. We conjecture that the reason for this is rounding error, not a failure of the method to converge in theory. The final iterate x is very close to \mathcal{M} , but the method is unable to track \mathcal{M} to minimize \hat{f} .

The behavior of BFGS when applied to the second nonsmooth variant f is completely different. We find that depending on the initialization, the method generates iterates that approximate one of the 2^{n-1} Clarke stationary points, not necessarily the global minimizer x^* . The iterates for the case $n = 2$, again for multiple starting points, are shown in the right panel of Figure 1. Most of the runs converge to the minimizer $[1, 1]^T$, but some terminate near the Clarke stationary point $[0, -1]^T$. For $n \leq 6$, given enough randomly generated starting points, BFGS finds, that is approximates well, all 2^{n-1} Clarke stationary points. The left and right panels of Figure 2 plot final values of f found by 1000 runs of BFGS starting with random x and $H = I$, sorted into increasing order, for the cases $n = 5$ and $n = 6$ respectively. Most runs find either the minimizer or one of the $2^{n-1} - 1$ non-minimizing points described above, although a few runs break down earlier. For $n = 7$, the method usually breaks down far away from these points, again presumably because of the limitations of machine precision.

Experiments with the gradient sampling algorithm [BLO05] and Kiwiel’s bundle code [Kiw08] give similar results. Nonetheless, it remains an open question whether the nonminimizing Clarke stationary points are actually points of attraction for any of these algorithms. Typically, the computations

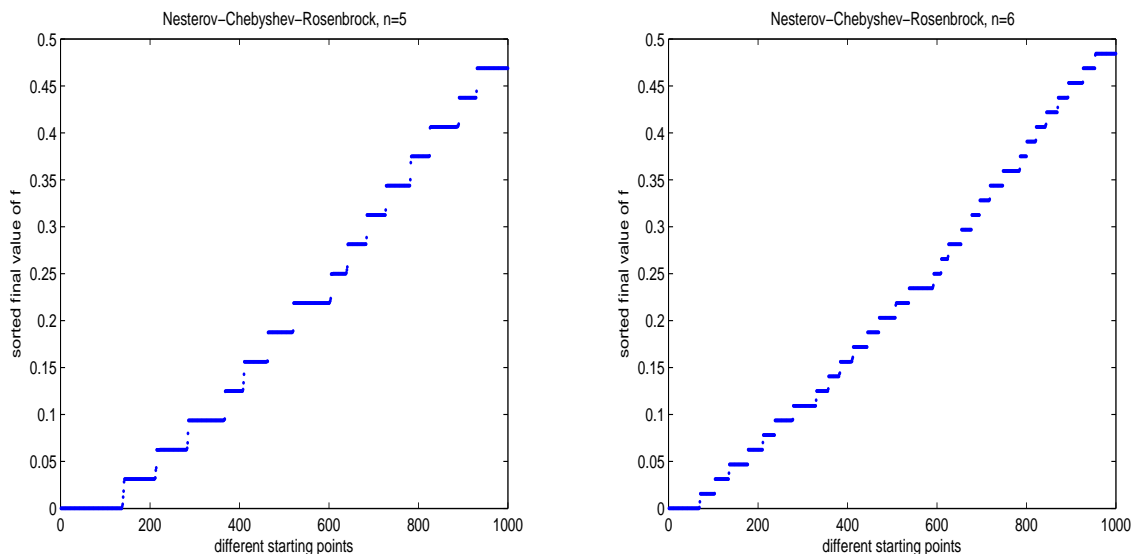


Figure 2: Left: sorted final values of f for 1000 randomly generated starting points, when $n = 5$: BFGS finds all 16 Clarke stationary points. Right: same with $n = 6$: BFGS finds all 32 Clarke stationary points.

terminate because, eventually, rounding error prevents the method from obtaining a lower point in the line search. But this does not establish whether, in exact arithmetic, the methods would actually generate sequences converging to the nonminimizing Clarke stationary points. Indeed, experiments with implementing BFGS at higher precision [Kak11] suggest that the higher the precision, the more likely BFGS is to eventually move away from the neighborhood of a nonminimizing Clarke stationary point, and eventually find a lower one, perhaps the minimizer.

Another observation was the difficulty of finding descent directions from the nonminimizing Clarke stationary points using random search. Although we know that such descent directions exist by Theorem 3, numerical experiments show that finding a descent direction by random search typically needs exponentially many trials in n . For example, when $n = 5$, usually 100,000 random trials do not suffice to find a descent direction. This illustrates the difficulty faced by an optimization method in moving away from these points.

4 Conclusion

Nesterov’s Chebyshev-Rosenbrock functions provide very interesting examples for optimization, both in theory and in practice. Specifically, the smooth function \tilde{f} , the first nonsmooth function \hat{f} and the second nonsmooth function f are very challenging nonconvex instances of smooth functions, partly smooth functions and non-regular functions respectively. As far as we know, Nesterov’s function f is the first documented case for which methods for nonsmooth optimization result in the approximation of Clarke stationary points from which there exist directions of linear descent. This observation is primarily due to Kiwiel [Kiw08]. It seems somewhat surprising since it is standard in the nonsmooth optimization literature to restrict convergence results to Clarke stationary points. Furthermore, since all first-order nonsmooth optimization methods, including bundle methods [Kiw85], the gradient sampling method [BLO05] and the BFGS method [LO10], are based on sampling gradient or subgradient information, the results given here for f suggest that limitation of convergence results to Clarke stationary points may be inherent, and that one may not in general be able to expect stronger results such as convergence only to Mordukhovich stationary points. Nonetheless, it remains an open question as to whether the nonminimizing Clarke stationary points of f are actually points of attraction for methods using exact arithmetic.

Acknowledgments. We thank Krzysztof Kiwiel, Adrian Lewis and Yuri Nesterov for various helpful comments and for their interest in this work.

References

- [BL05] J.M. Borwein and A.S. Lewis. *Convex Analysis and Nonlinear Optimization: Theory and Examples*. Springer, New York, second edition, 2005.
- [BLO02] J.V. Burke, A.S. Lewis, and M.L. Overton. Approximating sub-differentials by random sampling of gradients. *Math. Oper. Res.*, 27:567–584, 2002.
- [BLO05] J.V. Burke, A.S. Lewis, and M.L. Overton. A robust gradient sampling algorithm for nonsmooth, nonconvex optimization. *SIAM Journal on Optimization*, 15:751–779, 2005.

- [Cla83] F. H. Clarke. *Optimization and Nonsmooth Analysis*. John Wiley, New York, 1983. Reprinted by SIAM, Philadelphia, 1990.
- [Dai02] Y.-H. Dai. Convergence properties of the BFGS algorithm. *SIAM Journal on Optimization*, 13:693–701, 2002.
- [EG92] L. C. Evans and R. F. Gariepy. *Measure Theory and Fine Properties of Functions*. Studies in Advanced Mathematics. CRC Press, Boca Raton, FL, 1992.
- [GMW81] P.E. Gill, W. Murray, and M.H. Wright. *Practical Optimization*. Academic Press, New York and London, 1981.
- [Kak11] A. Kaku. Implementation of high precision arithmetic in the BFGS method for nonsmooth optimization. Master’s thesis, NYU, Jan 2011.
<http://www.cs.nyu.edu/overton/msthesis/kaku/msthesis.pdf>.
- [Kiw85] K.C. Kiwiel. *Methods of Descent for Nondifferentiable Optimization*. Lecture Notes in Mathematics 1133. Springer-Verlag, Berlin and New York, 1985.
- [Kiw08] K.C. Kiwiel, 2008. Private communication.
- [Lew03] A.S. Lewis. Active sets, nonsmoothness and sensitivity. *SIAM Journal on Optimization*, 13:702–725, 2003.
- [Lew06] A.S. Lewis. Eigenvalues and nonsmooth optimization. In *Foundations of Computational Mathematics, Santander 2005*, volume 331 of *London Math. Soc. Lecture Note Ser.*, pages 208–229. Cambridge Univ. Press, Cambridge, 2006.
- [LF01] D.-H. Li and M. Fukushima. On the global convergence of the BFGS method for nonconvex unconstrained optimization problems. *SIAM Journal on Optimization*, 11:1054–1064, 2001.
- [LO10] A.S. Lewis and M.L. Overton. Nonsmooth optimization via quasi-Newton methods, 2010. Submitted.
<http://www.cs.nyu.edu/overton/papers/pdf/nsoquasi.pdf>.
- [Mas04] W.F. Mascarenhas. The BFGS method with exact line searches fails for non-convex objective functions. *Mathematical Programming*, 99:49–61, 2004.

- [Mor76] B.S. Mordukhovich. Maximum principle in the problem of time optimal response with nonsmooth constraints. *Journal of Applied Mathematics and Mechanics*, 40:960–969, 1976.
- [Nes08] Y. Nesterov, 2008. Private communication.
- [NW06] J. Nocedal and S.J. Wright. *Nonlinear Optimization*. Springer, New York, second edition, 2006.
- [Ove01] M.L. Overton. *Numerical Computing with IEEE Floating Point Arithmetic*. SIAM, Philadelphia, 2001.
- [RW98] R.T. Rockafellar and R.J.B. Wets. *Variational Analysis*. Springer, New York, 1998.
- [SY06] W.-Y. Sun and Y.-X. Yuan. *Optimization Theory and Methods*, volume 1 of *Springer Optimization and Its Applications*. Springer, New York, 2006. Nonlinear programming.
- [Sze39] Gabor Szegő. *Orthogonal Polynomials*. American Mathematical Society, New York, 1939. American Mathematical Society Colloquium Publications, v. 23.