Foundations of Computational Mathematics
https://doi.org/10.1007/s10208-024-09653-y

FOUNDATIONS OF COMPUTATIONAL MATHEMATICS
The Journal of the Society for the Foundations of Computational Mathematics

A Local Nearly Linearly Convergent First-Order Method for Nonsmooth Functions with Quadratic Growth

Damek Davis[1] **Liwei Jiang**[1]

Received: 3 June 2022 / Revised: 15 July 2023 / Accepted: 21 November 2023

Abstract

Classical results show that gradient descent converges linearly to minimizers of smooth strongly convex functions. A natural question is whether there exists a locally nearly linearly convergent method for nonsmooth functions with quadratic growth. This work designs such a method for a wide class of nonsmooth and nonconvex locally Lipschitz functions, including max-of-smooth, Shapiro's decomposable class, and generic semialgebraic functions. The algorithm is parameter-free and derives from Goldstein's conceptual subgradient method.

Keywords Subgradient method · Goldstein subdifferential · Semialgebraic · Partial smoothness · VU-structure

1 Introduction

Slow sublinear convergence of first-order methods in nonsmooth optimization is often illustrated with the following simple strongly convex function:

$$f(x) = \max_{1 \le i \le m} x_i + \frac{1}{2}\|x\|^2 \quad \text{for some } m \le d \text{ and all } x \in \mathbb{R}^d. \tag{1.1}$$

For example, consider the subgradient method applied to f, which generates iterates x_k. Since f is strongly convex, classical results dictate that $f(x_k) - \inf f = O(k^{-1})$.

Communicated by Jérôme Bolte.

Research of Davis supported by an Alfred P. Sloan research fellowship and NSF DMS award 2047637.

B Damek Davis
dsd95@cornell.edu
http://people.orie.cornell.edu/dsd95/

Liwei Jiang
https://liwei-jiang97.github.io/

[1] School of Operations Research and Information Engineering, Cornell University, Ithaca, NY 14850, USA

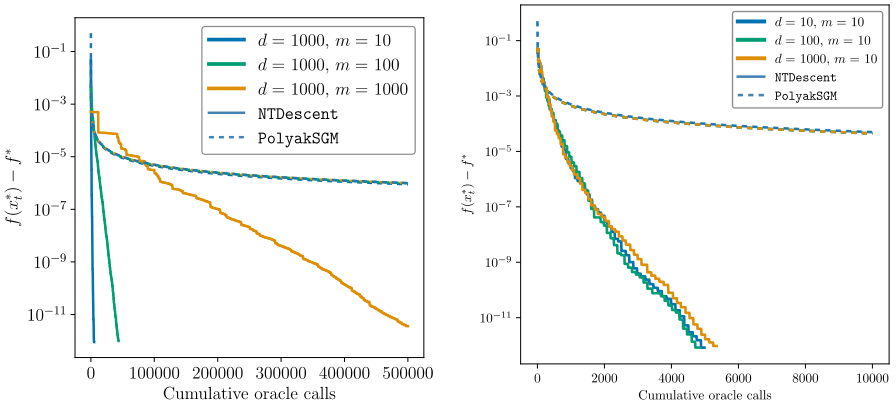

Fig. 1 Comparison of `NTDescent` with `PolyakSGM` on (1.1). Left: we fix d and vary m; Right: we fix m and vary d. For both algorithms, the value $f(x_t^*)$ denotes the best function value seen after t oracle evaluations

On the other hand, under proper initialization and an adversarial first-order oracle, there is a matching lower bound for the first m iterations: $f(x_k) - \inf f \geq (2m)^{-1}$ for all $k \leq m$; see [11, 43]. Beyond the subgradient method, the lower bound also holds for any algorithm whose kth iterate lies within the linear span of the initial iterate and past $k-1$ computed subgradients. Thus, one must make more than m first-order *oracle calls* to f, i.e., function and subgradient evaluations, before possibly seeing improved convergence behavior.

While such methods make little progress when $k \leq m$, this behavior may or may not continue for $k > m$. On one extreme, the subgradient method continues to converge slowly even when equipped with the popular Polyak stepsize (`PolyakSGM`) [47]; see dashed lines in Fig. 1. On the opposite extreme, more sophisticated algorithms such as the center of gravity method or the ellipsoid method converge linearly, but their complexity scales with the dimension of the problem, a necessary consequence of the linear rate of convergence; see the discussion in [11, Chapter 2].

A natural question is whether there exists a first-order method whose behavior lies in between these two extremes, at least for nonsmooth functions f satisfying regularity conditions at local minimizers. Regularity conditions often take the form of growth—linear or quadratic—away from minimizers. Well-known results show that subgradient methods converge linearly on nonsmooth functions with linear (also called *sharp*) growth [47]. On the other hand, in smooth convex optimization, quadratic growth entails linear convergence of gradient methods. However, to the best of our knowledge, no parallel result for nonsmooth functions with quadratic growth exists. Thus, in this work, we ask

> is there a locally nearly linearly convergent method for nonsmooth functions with quadratic growth whose rate of convergence and region of rapid local convergence solely depends on f?

Let us explain the qualifiers "nearly" and "solely depends on f." First, the qualifier "nearly" signifies that the method locally achieves a function gap of size ε using at most, say, $O(C_f \log^3(1/\varepsilon))$ first-order oracle evaluations of f, where C_f depends on f. Second, the qualifier "solely depends on the function," signifies that C_f and the size of the region of local convergence do not depend on the dimension of the problem, but instead depend only on the function f through intrinsic quantities, such as Lipschitz and quadratic growth constants.

In this work, we positively answer the above question for a class of nonsmooth optimization problems with quadratic growth. The method we develop is called *Normal Tangent Descent* (`NTDescent`). We formally describe `NTDescent` in Sect. 1.4. For now, we illustrate the performance of `NTDescent` on f from (1.1) in Fig. 1. In both plots, we see `NTDescent` improves on the performance of `PolyakSGM`, measured in terms of oracle calls. This is a fair basis for comparison since both `PolyakSGM` and `NTDescent` perform a similar amount of computation per oracle call. Figure 1b also shows that the performance of `NTDescent` is dimension independent. We highlight that this performance was achieved without any tuning of parameters for `NTDescent`. Indeed, our main theoretical guarantees for `NTDescent` (Theorem 1.1) do not require the user to set any parameters.

The problem class on which `NTDescent` succeeds consists of locally Lipschitz nonsmooth functions with quadratic growth and a certain *smooth substructure* at local minimizers. Importantly, we do not assume the problems under consideration are convex, though convexity entails improved guarantees. Two example classes with such smooth substructure include (i) "generic" semialgebraic functions and (ii) properly C^p decomposable loss functions satisfying strict complementarity and quadratic growth conditions [49]. A semialgebraic function is one whose graph is the finite union of intersections of polynomial inequalities. Semialgebraic functions (more generally *tame* [32] functions) model most problems of interest in applications. If f is semialgebraic, for a full Lebesgue measure set of $w \in \mathbb{R}^d$, we will use show that the tilted function $f_w : x \to f(x) + w^T x$ has quadratic growth and the desired smooth substructure at each local minimizer, explaining the qualifier "generic." We mention that this fact essentially follows from combining results of [19, 23]. On the other hand, a properly C^p decomposable function is one that decomposes near local minimizers as a composition of a positively homogeneous convex function with a smooth mapping that maps the minimizer to the origin. Decomposable functions appear often in practice, e.g., in eigenvalue and data fitting problems. An important subclass of decomposable functions consists of so-called "max-of-smooth" functions, which are the maximum of finitely many smooth functions that satisfy certain regularity conditions at minimizers, e.g., f in (1.1).

The precise smooth substructure used in this work was recently identified in [19], where it was shown to be available in decomposable and generic semialgebraic problems. Since it is available in many problems of interest, throughout this introduction we call this the combination of quadratic growth and smooth substructure *typical structure* and call functions possessing this combined structure *typical*. We present the formal structure in Sect. 3. At the heart of this structure is a distinguished smooth manifold M —called the *active manifold*—containing a local minimizer of interest. We formally define the active manifold concept in Definition 1.2, but at a high level, the

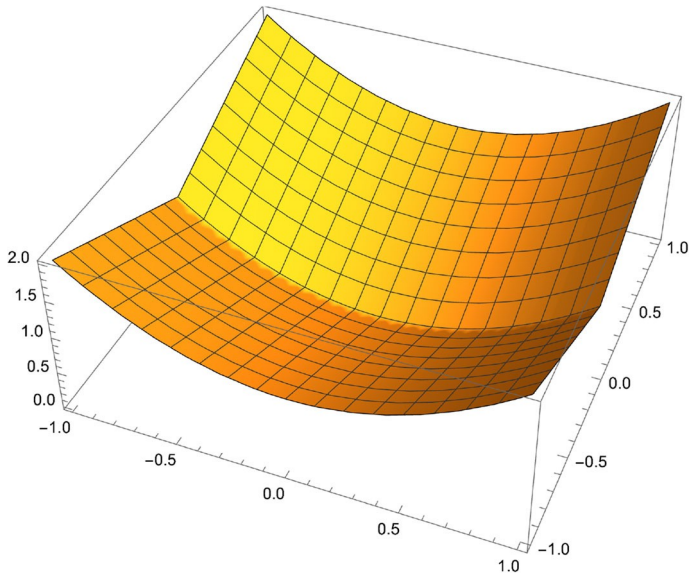

Fig. 2 The function $f(u, v) = u^2 + |v|$ has typical structure

two crucial characteristics are that (i) along the manifold, the function f is smooth and (ii) normal to the manifold, the function grows sharply. For example, Fig. 2 depicts the nonsmooth function $f(u, v) = u^2 + |v|$ for which the u-axis plays the role of M. In Sect. 1.3.2 we will examine this function and explain how we use its typical structure in `NTDescent`. This example also has the smooth substructure developed in several seminal works in the optimization literature, including those found in work on identifiable surfaces [52], partly smooth functions [38], VU-structures [36, 41], and minimal identifiable sets [26]. However, crucial to the analysis of `NTDescent` are two further properties introduced in [19], called *strong* (a)*-regularity* and $(b_{\leq})$*-regularity*. Strong (a)-regularity roughly states that the function is smooth in tangent directions to the manifold up to an error term which is linear in the distance to the manifold. On the other hand, $(b_{\leq})$-regularity is a one-sided uniform semismoothness [40] property that holds automatically when f is (weakly) convex. Both properties hold for the two variable example in Fig. 2 and for the function in (1.1), where the active manifold is the subspace in which the first m variables take on the same value: $M = \{x \in \mathbb{R}^d : x_1 = x_2 = \ldots x_m\}$.

Before turning to the description of `NTDescent`, we point out that similar smooth substructure has been used in the analysis first-order methods in nonsmooth optimization, most famously for functions with VU-structure [36, 41] and more recently for max-of-smooth functions.[1] For VU functions, so-called "bundle-methods," [35, 50] which possess an inner-outer loop structure, have been shown to converge superlinearly with respect to the number of outer-loop steps [41]; see also the survey [44]. These methods have excellent empirical performance, but a complete account of their

[1] Though they also benefit from smooth substructure, *proximal-methods* do not fall within the oracle model of first-order methods considered in this work. Thus, we omit them from our discussion.

inner-loop complexity remains elusive. On the other hand, in a recent work, Han and Lewis proposed a first-order method—Survey Descent—that converges linearly on certain strongly convex max-of-smooth objectives, stepping beyond the classical smooth setting [30]. The method shows favorable performance beyond the max-of-smooth class, e.g., on certain eigenvalue optimization problems, but no theoretical justification for this success is available. We discuss Survey Descent in more detail in Sect. 7.1. We now motivate `NTDescent`.

1.1 Motivation: Goldstein's Conceptual Subgradient Method

To motivate `NTDescent` and the role of smooth substructure, let us set the stage: consider the nonsmooth optimization problem:

$$\text{minimize}_{x \in \mathbf{R}^d} \; f(x),$$

where $f : \mathbf{R}^d \to \mathbf{R}$ is a locally Lipschitz function, which is not necessarily convex. The algorithm developed in this work assumes *first-order oracle access* to f [11, 42, 43]. In particular, at every $x \in \mathbf{R}^d$ we must be able to evaluate $f(x)$ and retrieve an element of the *Clarke subdifferential* $\partial f(x)$. Informally, the Clarke subdifferential is comprised of convex combinations of limits of gradients taken at nearby points; a formal definition appears in Sect. 1.7. The Clarke subdifferential reduces to the familiar objects in classical settings. For example, when f is C^1, the Clarke subdifferential reduces to the singleton mapping $\{\nabla f\}$. In addition, when f is convex, the Clarke subdifferential reduces to the subdifferential in the sense of convex analysis.

The starting point of this work is the classical conceptual subgradient method of Goldstein [29]. The core object in this method is the Goldstein subdifferential:

$$\partial_\sigma f(x) := \operatorname{conv}\left(\bigcup_{y \in \overline{B}\sigma(x)} \partial f(y) \right) \qquad \text{for all } x \in \mathbf{R}^d \text{ and } \sigma > 0. \tag{1.2}$$

This subdifferential is simply the convex hull of all Clarke subgradients of f taken at points inside the ball of radius σ. Its importance arises from the following descent property proved in [29]: fix $\sigma > 0$ and $x \in \mathbf{R}^d$ and let w denote the minimal norm element of $\partial_\sigma f(x)$. Then

$$f\left(x - \sigma \frac{w}{w}\right) \leq f(x) - \sigma w \qquad \text{if } w = 0. \tag{1.3}$$

This property motivates Goldstein's conceptual subgradient method, which simply iterates:

$$x_{k+1} = x_k - \sigma \frac{w_k}{w_k} \qquad \text{where} \qquad w_k = \operatorname*{argmin}_{w \in \partial_\sigma f(x_k)} w. \tag{1.4}$$

This algorithm is remarkable since it is provably a descent method for any Lipschitz function and even converges at a sublinear rate. Indeed, a quick appeal to (1.3) yields

$$\min_{k=0,\dots,K-1} w_k \le \varepsilon \qquad \text{holds when} \qquad K \ge \frac{f(x_0) - \min f}{\sigma\,\varepsilon}.$$

While this exact variant of the Goldstein method is not necessarily implementable, recent work has devised approximate versions of the method that have similar sublinear convergence properties [21, 55].

The algorithm introduced in this work approximately implements the method (1.4). The goal of this work is to prove that the method is locally nearly linearly convergent on typical nonsmooth functions. To develop such a method, we must resolve two issues for this problem class. First, we must develop rapidly convergent algorithms that approximately compute the minimal norm element of the Goldstein subdifferential. Second, we must devise an appropriate regularity property that ensures the proposed method converges nearly linearly. We will discuss both of these properties in turn, beginning with a regularity property that relates the decrement in (1.3) to the function gap.

1.2 Linear Convergence via a Gradient Inequality

Observe that if the bound

$$\sigma\, w_k \ge \eta\big(f(x_k) - \min f\big)$$

holds for some $\eta > 0$ and all $k > 0$, then the Goldstein method (1.4) converges linearly to a minimizer of f. A potential issue with this inequality is that the vector w_k is zero whenever σ is larger than the distance of x_k to the nearest critical point of f; thus the algorithm may stall whenever x_k is near enough to a minimizer. This suggests a simple relaxation of the property that allows σ to depend on x_k.

Indeed, we will provide conditions under which the following bound holds near a local minimizer $\bar{x}$ of f: there exists a constant $\eta > 0$ and a function $\sigma : \mathbb{R}^d \to \mathbb{R}_+$ such that for all x near $\bar{x}$, we have

$$\sigma(x)\operatorname{dist}(0, \partial_{\sigma(x)} f(x)) \ge \eta\big(f(x) - f(\bar{x})\big). \tag{1.5}$$

throughout, we will refer to this bound as a *gradient inequality*, due to its similarity to the Kurdyka-Łojasiewicz (KL) gradient inequality [7]. The KL inequality requires that a suitable nonlinear reparameterization $\psi : \mathbb{R} \to \mathbb{R}$ of the function gap is bounded by the minimal norm Clarke subgradient for all x near $\bar{x}$:

$$\operatorname{dist}(0, \partial f(x)) \ge \psi\big(f(x) - f(\bar{x})\big).$$

In recent years, the KL inequality has played a key role in establishing convergence and rates of convergence for proximal methods in nonsmooth optimization and in continuous time analogs of the subgradient method; see e.g., [3, 4, 7, 9, 53].

To illustrate, let us specialize to the semialgebraic setting, where the desingularization function ψ is known to take the form $\psi(r) = r^{\theta}$ for $\theta \in [0, 1)$. The work [2, Theorem 2] initiated the study of convergence of proximal methods in this setting, showing that the proximal point method asymptotically converges to its limit point, which is critical but not necessarily optimal. The method convergence in finitely many steps when $\theta = 0$, locally converges linearly when $\theta \in (0, 1/2]$, and locally converges at the rate $k^{\frac{-(1-\theta)}{2\theta-1}}$ when $\theta \in (1/2, 1)$. Further works such as [3, 9] generalized the techniques to related proximal methods. Passing to continuous time, one is interested in the convergence of the trajectory of subdifferential inclusion satisfying $\dot{x}(t) \in -\partial f(x(t))$ at almost every t. Here, the rates of convergence exactly parallel those in the proximal methods as shown in [6, Theorem 4.7]. [2] In contrast to the proximal and continuous-time settings, we do not know whether the KL inequality alone allows one to design a locally linearly convergent discrete-time subgradient method, except in the setting where $\theta = 0$ (i.e., f is *sharp*) and f is convex [47] or *weakly convex* [22]; weakly convex functions form a broad class of nonconvex functions that includes all compositions of Lipschitz convex functions with smooth mappings. When $\theta > 0$, to the best of our knowledge, the best rate proved in the literature for any subgradient type method is $k^{\frac{-(1-\theta)}{2\theta}}$ [33]; this result is only known to hold for convex functions. [3]

A well-known property of the KL inequality is its prevalence: it holds at each critical point of an arbitrary lower-semicontinuous semialgebraic function f [7]. We will show that the gradient inequality (1.5) is also prevalent in the sense that it holds for the aforementioned problems with typical structure. In this way, the conceptual method (1.4) with varying $\sigma_k := \sigma(x_k)$ will locally converge linearly on such problems. The reader may wonder whether we can or must find the precise value $\sigma(x_k)$. We will show that for typical problems, an appropriate σ_k may be found through a line search procedure.

1.3 Approximately Implementing Goldstein's Method

The gradient inequality ensures that the conceptual Goldstein method converges linearly, provided the stepsize σ is chosen adaptively. To move beyond the conceptual setting, we must develop strategies for approximating the minimal norm element of $\partial_\sigma f(x)$ for $\sigma > 0$ and $x \in \mathbb{R}^d$. Let us suppose we have such a method and denote it by $\texttt{MinNorm}(x, \sigma)$. Then the method of this work simply iterates:

$$x_{k+1} = x_k - \sigma_k \frac{w_k}{\|w_k\|} \quad \text{and} \quad w_k = \texttt{MinNorm}(x_k, \sigma_k) \tag{1.6}$$

for an appropriate sequence $\sigma_k > 0$. We will discuss and develop two different implementations of $\texttt{MinNorm}(x, \sigma)$ in this work. Given $x \in \mathbb{R}^d$ and $\sigma > 0$,

[2] These rates were shown only for "lower-C^2" semialgebraic losses, but extend to locally Lipschitz semialgebraic functions via the semialgebraic "chain rule" proved in [20].

[3] The results stated in [33] pertain to functions with Hölder growth; thus, to prove the results stated in the paragraph, we must use the following known fact: functions satisfying the KL inequality with exponent θ have Hölder growth with exponent $1/(1-\theta)$, which follows from the proof of [24, Theorem 3.7].

both methods iteratively construct a sequence of Clarke subgradients $g_0, \cdots, g_{T-1}$ taken at points in the ball $\overline{B}_\sigma(x)$ and then output a "small" convex combination $w \in \operatorname{conv}\{g_0, \cdots, g_{T-1}\}$, which satisfies the descent condition

$$f\left(x - \sigma \frac{w}{\|w\|}\right) \le f(x) - \frac{\sigma}{8}\|w\|. \tag{1.7}$$

The oracle complexity of $\texttt{MinNorm}(x, \sigma)$ is then T function/subgradient evaluations, and we hope to ensure that T is relatively small, for example, a constant or at most

$$T = O\left(\log \frac{1}{\delta_{x,\sigma}}\right) \qquad \text{where} \quad \delta_{x,\sigma} := \operatorname{dist}(0, \partial_\sigma f(x)).$$

Provided that T is on this order, that f satisfies the gradient inequality (1.5), and that σ_k is chosen appropriately, the iterate x_k will satisfy $f(x_k) - f(\bar{x}) \le \varepsilon$ after at most $O(\log^2(1/\varepsilon))$ iterations, a nearly linear rate of convergence. This complexity ignores the cost of choosing an appropriate stepsize σ_k, but we will show that in typical problems we can find appropriate σ_k with at most $O(\log(1/\varepsilon))$ function/subgradient evaluations.

We are aware of two `MinNorm` type methods in the literature, but their complexity is either too large or is useful only in low dimensions problems. For example, the works [21, 55] introduced such a method for general locally Lipschitz functions. However, the complexity of the method is $T = O(1/\delta_{x,\sigma})$—too large for our purposes. On the other hand, the work [21] also introduced a method tailored to low-dimensional weakly convex functions. However, the method is based on cutting plane techniques, so its complexity scales linearly with dimension: $T = O(d \log(1/\delta_{x,\sigma}))$.

While existing `MinNorm` methods are slow for general Lipschitz functions, we show that the aforementioned typical structure allows us to develop `MinNorm` methods that accelerate in a neighborhood of the minimizer. Our approach is based on a decomposition of a neighborhood of the minimizer into two regions: one where the method of [21, 55] is applicable, and another region where a novel `MinNorm` method may be applied.

1.3.1 The Normal and Tangent Regions

In this work, we use the active manifold M to split the space of (x, σ) for x nearby the minimizer $\bar{x} \in M$ into two sets where fast `MinNorm` methods are available. We call the first set the *normal region*. This region consists of points whose normal distance $\operatorname{dist}(x, M)$ is larger than a multiple of the squared tangential distance $\|P_M(x) - \bar{x}\|^2$, together with stepsizes σ proportional to a multiple of the normal distance:

$$\begin{aligned} &\tfrac{a_1}{2}\operatorname{dist}(x, M) \le \sigma \le a_1 \operatorname{dist}(x, M); \\ &a_2^2 \|P_M(x) - \bar{x}\|^2 \le \operatorname{dist}(x, M), \end{aligned}$$

for problem dependent constants $a_1, a_2 \in (0, 1)$; see Theorem 4.3 for more details. We will show that in this region, we have $\ldots_{x,\sigma} = \ldots 1)$, so the `MinNorm` method of [21, 55] terminates with descent in finitely many steps.

On the other hand, we call the second set the *tangent region*. This set consists of points whose squared tangential distance is larger than a multiple of the normal distance, together with stepsizes σ proportional to a multiple of the tangential distance:

$$\begin{aligned} &\tfrac{a_2}{2} \ldots P_M(x) - \bar{x} \le \sigma \le \ldots a_2 \ldots P_M(x) - \bar{x}; \\ &\tfrac{\operatorname{dist}(x, M)}{\sigma} \le 2a_2 \ldots P_M(x) - \bar{x}, \end{aligned}$$

where a_1 and a_2 are as in the normal region. For this region, we will propose a new `MinNorm` method, which terminates rapidly. We note that in both cases we provide a range of valid σ, rather than a single value since we aim to estimate σ with a line search.

1.3.2 A Simple Example

Before describing the methods in detail, let us illustrate the regions and the principles of the methods on the following simple function of two variables $f(u, v) = u^2 + |v|$, which has a unique minimizer at $\bar{x} = (0, 0)$. Here, the u-axis is the active manifold M. Along the manifold, f is smooth and grows quadratically, while off of the manifold f grows sharply; see Fig. 2 for a plot of the function and see Fig. 3 for the x-component of the normal and tangent regions for f (with $a_2 = 1/8$). The manifold M induces a decomposition of f into smooth $f_U(u, v) = u^2$ and nonsmooth $f_V(u, v) = |v|$ components. In particular, with variable $x = (u, v)$, we have

$$f(x) = f_U(x) + f_V(x) = \ldots P_M(x) - \bar{x} \ldots^2 + \operatorname{dist}(x, M). \tag{1.8}$$

From this decomposition, we see f_V is dominant in the normal region, while f_U is dominant in the tangent region. Likewise, as we will argue momentarily, the minimal norm Goldstein subgradient $w_\sigma \in \partial_\sigma f(x)$ satisfies $\ldots w_\sigma \ldots \ge \ldots \nabla f_V(x) \ldots$ in the normal region, while $\ldots w_\sigma \ldots = \ldots (\nabla f_U(x)) \ldots$ in the tangent region. This has several consequences. First, in the normal region, the `MinNorm` method of [21, 55] will terminate in finitely many steps, due to the lower bound $\ldots w_\sigma \ldots \ge 1$. On the other hand, in the tangent region, $\ldots w_\sigma \ldots$ can be much smaller, so we must introduce a new method to generate descent. Finally, assuming these approximations are accurate, the gradient inequality (1.6) quickly follows: in the normal region, we have

$$\sigma \ldots w_\sigma \ldots = \ldots \operatorname{dist}(x, M)) = \ldots (\ldots f(x)),$$

while in the tangent region, we have

$$\sigma \ldots w_\sigma \ldots = \ldots (P_M(x) - \bar{x} \nabla \ldots f_U(x)) = \ldots (\ldots P_M(x) - \bar{x} \ldots^2) = \ldots f(x)).$$

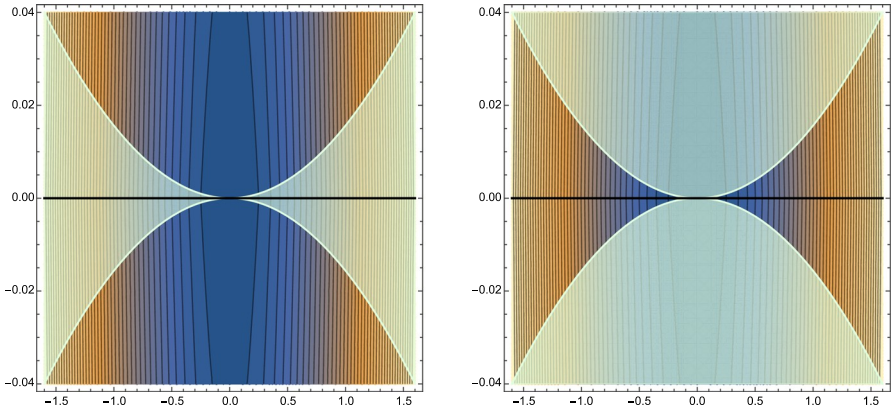

Fig. 3 Contour plots for $f(u, v) = u^2 + |v|$ together with M, shown in black. The x-component of the tangent (left) and normal (right) regions are overlaid in light green

Though it follows from immediate calculations in this example, in the more general setting, the following consequence of quadratic growth will be crucial in establishing a similar bound: $\nabla f_U(x) = (P_M(x) - \bar{x})$.

Now, to lower bound w_σ we use the following fact: $\nabla f_U(u, v)$ is tangent to M, while $\nabla f_V(u, v)$ is normal to M when $v = 0$. Thus, to lower bound w_σ in the normal region, we will simply lower bound the size of the normal component of w_σ. Indeed, since $\sigma < \text{dist}(x, M)$, all points $x \in B_\sigma(x)$ are on the same side of M. Therefore, the normal component of w_σ is an average of *identical* gradients $\nabla f_V(x) = \nabla f_V(x)$. Likewise, in the tangent region, we lower bound the tangent component of w_σ. Indeed, since $\sigma < \bar{x} - P_M(x)/8$, the projection onto M of all points $x \in B_\sigma(x)$ are on the same side of the origin. Thus, the tangent component of w_σ is an average of nearly identical gradients $\nabla f_U(x) \approx \nabla f_U(x)$, yielding the lower bound. We prove a more general form of both of these lower bounds Lemma 4.1 and Lemma 4.2, which follow from similar principles.

Turning to algorithms, we have so far noted that the `MinNorm` method of [21, 55] may be used in the normal region. In the tangent region, we are unsure how to design a method that can quickly recover w_σ. Instead of searching for w_σ directly, we take a slightly different perspective in the tangent region: we seek a vector $g \in \partial_\sigma f(x)$ with "small" normal component, meaning:

$$P_N(g) = O(\nabla f_U(x)^2)$$

where N is the normal space to M. Intuitively, when g has small normal component, the nonsmooth part f_V minimally changes along a gradient step. On the other hand, if g is sufficiently correlated with ∇f_U, the smooth part f_U decreases at an appropriate rate; we prove this in a more general setting in Lemma 5.2.

Why might one expect such a g to be available in the tangent region? The reason is that the gradient of the smooth component is itself a Goldstein subgradient. Indeed, for points near the origin and in the tangent region, the tangential distance is much larger

than the normal distance. Thus, the reflection of any point (u, v) across the manifold M is contained in $B_\sigma(x)$, which immediately implies gradient of the smooth component is an element of Goldstein subdifferential:

$$\nabla f_U(u, v) = \frac{1}{2}\nabla f(u, v) + \frac{1}{2}\nabla f(u, -v) \in \partial_\sigma f(x). \tag{1.9}$$

While the inclusion (1.9) illustrates one way to construct such a g, we cannot hope for perfect symmetry in general problems.

Instead, a central insight of this work is that a similar approximate reflection exists in problems with typical structure. To illustrate, consider Fig. 4. This figure depicts a point x in the tangent region together with the result of a normalized gradient step:

$$x_+ := x - \sigma \frac{\nabla f(x)}{\nabla f(x)}.$$

As can be seen from the figure, x_+ is an approximate reflection of x across the u-axis, which "flips the sign" of the nonsmooth component of ∇f: $\nabla f_V(x) = -\nabla f_V(x_+)$. Thus, in this setting, one may "cancel out" the nonsmooth component by a simple averaging:

$$\nabla f_U(x) \approx \frac{1}{2}\nabla f(x) + \frac{1}{2}\nabla f(x_+).$$

While seemingly crude, we will show this strategy generalizes to typical functions. An important distinction with the general setting is that a single averaging step alone will no longer suffice. Nevertheless, we show that by iterating this process, we can geometrically shrink the normal component of the Goldstein gradient, eventually yielding descent.

1.3.3 Two `MinNorm` Methods: `NDescent` and `TDescent`

To generalize the strategy outlined in the previous section, we will prove that the minimal norm Goldstein subgradients of typical problems similarly split into tangent and normal components just as in Sect. 1.3.2. Then, we introduce two `MinNorm` type methods for "normal" and "tangent" steps.

For (x, σ) in the normal region, we use a small modification of the `MinNorm` type method of [21]. We call this method *Normal Descent* (`NDescent`) and describe it in Algorithm 1. As in the simple example above, we will show that `NDescent` must terminate with an approximately minimal norm Goldstein subgradient in finitely many steps, provided σ lies within an appropriate range. We will show that this subgradient is a descent direction satisfying (1.7).

Algorithm 1 NDescent(x, g, σ, T)

Set $g_0 = g$ and $t = 0$.
while $T - 1 \geq t$, $g_t > 0$, and $\frac{\sigma}{8} g_t \geq f(x) - f\left(x - \sigma \frac{g_t}{g_t}\right)$ **do**
  Choose any r satisfying $0 < r < \sigma g_t$.
  Sample ζ_t uniformly from $\mathsf{B}_r(g_t)$.
  Choose y_t uniformly at random in the segment $\left[x, x - \sigma \frac{\zeta_t}{\zeta_t}\right]$.
  Choose $\hat{g}_t \in \partial f(y_t)$.
  $g_{t+1} = \operatorname{argmin}_{z \in [g_t, \hat{g}_t]} \|z\|_2$.
  $t = t + 1$.
end while
return g_t.

We illustrate the principle behind NDescent as follows. Suppose we are given a vector $g \in \partial_\sigma f(x)$ not satisfying the descent condition, i.e., with $u := \frac{g}{\|g\|}$, we have

$$f(x - \sigma u) - f(x) \geq -\frac{g}{8}.$$

Then by Lebourg mean value theorem [16, Theorem 2.4] (provided that f is differentiable along the line segment between $[x, x]$, which can be ensured by adding a small perturbation to g; we ignore this in our discussion), we may assume that

$$f(x - \sigma u) - f(x) = \sigma \int_0^1 -\langle \nabla f(x - \sigma t u), u \rangle \, dt = -\sigma \langle v, u \rangle,$$

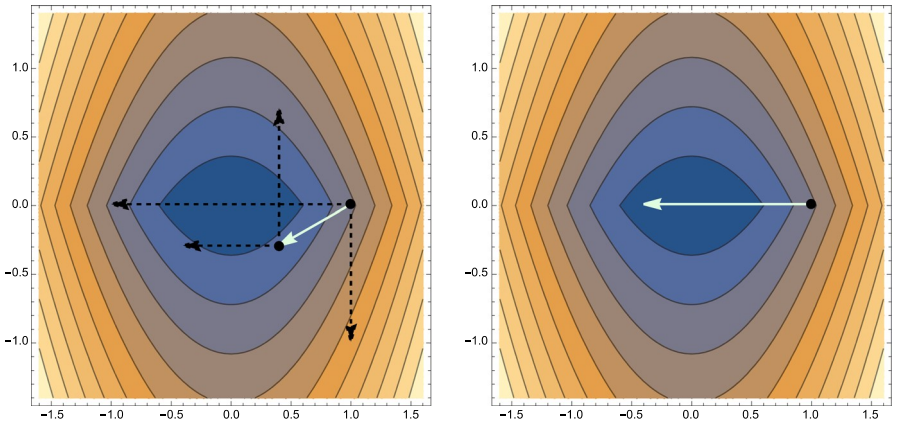

Fig. 4 Contour plots for $f(u, v) = u^2 + |v|$. Left: The point $x = (1, 1)$ together with the approximate reflection $x_+ = x - .3\frac{\nabla f(x)}{\|\nabla f(x)\|}$ across the u axis. The solid light green arrow is parallel to the negative gradient direction $-\nabla f(x)$. The dashed arrows denote the orthogonal decomposition of $-\nabla f(x)$, respectively $-\nabla f(x_+)$, into the vectors $-\nabla f_U(x)$ and $-\nabla f_V(x)$, respectively $-\nabla f_U(x_+)$ and $-\nabla f_V(x_+)$. From the plot, we see $\nabla f_V(x) = -\nabla f_V(x_+)$. Right: The point x with estimate $-\frac{1}{2}(\nabla f(x) + \nabla f(x_+))$ of the vector $-\nabla f_U(x)$

where $v := \int_0^1 \nabla f(x - tu)\, dt \in \partial_\sigma f(x)$. Consequently, $\langle v, g \rangle \le \|g\|^2/8$. While it is not possible to compute v, we can compute a *random* element of the Goldstein subdifferential, satisfying the same inequality in expectation. Indeed, defining $v = \nabla f(y)$ where y is uniformly sampled from the line segment $[x, x - \sigma u]$ (with end points x and $x - \sigma u$), we have $\mathbb{E}_y[\langle v \rangle], g \le \|g\|^2/8$. Based on this bound, a quick calculation shows that the minimal norm element g_+ of the line segment $[g, v]$ satisfies the bound

$$\mathbb{E}_y \|g_+\|^2 \le \|g\|^2 - \frac{\|g\|^4}{16L^2}.$$

Moreover $g_+ \in \partial_\sigma f(x)$. Thus, repeating this process yields a decreasing sequence of Goldstein subgradients which tend to zero as long as the descent condition is not met. In general, the norms of the subgradients generated by this process decay at a rate of $1/k$. However, we will prove that $\operatorname{dist}(0, \partial_\sigma f(x))$ is bounded below by a fixed constant when (x, σ) is in the normal region described in Sect. 1.3.1. Consequently, the loop must exist in finite time with descent (with high probability), for otherwise we will have found a subgradient norm strictly smaller than $\operatorname{dist}(0, \partial_\sigma f(x))$; see Proposition 5.1. The reader interested in the formal calculations may consult [21, 55].

On the other hand, for (x, σ) in the tangent region, we develop a new `MinNorm` type method, which likewise relies on an approximate reflection property. We call this method *Tangent Descent* (`TDescent`) and present it in Algorithm 2. Given an input point x, stepsize $\sigma > 0$, and initial subgradient $g_0 \in \partial f(x)$, `TDescent` repeats the following steps

$$\text{Choose: } \hat{g}_k \in \partial f\left(x - \sigma \frac{g_k}{\|g_k\|}\right);$$

$$\text{Update: } g_{k+1} = \operatorname*{argmin}_{g \in [g_k, \hat{g}_k]} \|g\|,$$

until it achieves descent $f\left(x - \sigma \frac{g_k}{\|g_k\|}\right) \le f(x) - \frac{\sigma}{8}\|g_k\|$ or runs over budget.

Algorithm 2 `TDescent`(x, g, σ, T)

Set $g_0 = g$ and $t = 0$.
while $T - 1 \ge t$, $\|g_t\| > 0$, and $\frac{\sigma}{8}\|g_t\|_2 \ge f(x) - f\left(x - \sigma \frac{g_t}{\|g_t\|}\right)$ **do**
  Choose $\hat{g}_t \in \partial f\left(x - \sigma \frac{g_t}{\|g_t\|}\right)$.
  $g_{t+1} = \operatorname{argmin}_{z \in [g_t, \hat{g}_t]} \|z\|$.
  $t = t + 1$.
end while
return g_t.

The motivation for this method is that for typical problems the step $x - \sigma \frac{g_k}{\|g_k\|}$ is locally an approximate reflection across M that "flips" the normal component of the Goldstein subgradient. Indeed, let $y := P_M(x)$ denote the projection of x onto M

and let $N := N_M(y)$ denote the normal space to M at y. Then we will prove that for all k, we have

$$P_N g_k, \hat{g}_k \leq -C \, P_N g_k + O(\|y - \bar{x}\|^2),$$

for some $C > 0$, provided σ lies within an appropriate range. This inequality ensures that each step of the `TDescent` geometrically decreases the "normal component" of g_k, until we arrive at a Goldstein subgradient with normal component on the order of $O(\|y - \bar{x}\|^2)$; see Sect. 5.2.2. Moreover, given $g \in \partial_\sigma f(x)$ satisfying

$$P_N(g) \leq C_3 \|y - \bar{x}\|^2$$

for a particular problem dependent constant $C_3 > 0$, we will prove the descent condition

$$f\left(x - \sigma \frac{g}{\|g\|}\right) \leq f(x) - \frac{\sigma \|g\|}{8}$$

holds; see Lemma 5.2. Combining these two facts shows that `TDescent` will rapidly terminate with descent.

1.4 The `NTDescent` Algorithm

We call the main algorithm of this work *Normal Tangent Descent* (`NTDescent`) and present it in Algorithm 4. At a high level the method is an approximate implementation of Goldstein's conceptual subgradient method as in (1.6), using `NDescent` and `TDescent` as `MinNorm` type methods. As input it takes three parameters: an initial point x; a sequence of grid-sizes $\{G_k\}$ for the line search on σ; and a sequence of budgets $\{T_k\}$ for the `MinNorm` type methods `NDescent` and `TDescent`. Later we will show that the user may simply set $T_k = G_k = k + 1$ for all $k \geq 0$.

Algorithm 3 `linesearch`(x, g, s, G, T)

1: **Set** $v_0 = g$.
2: **for** $i = 0, \ldots, G - 1$ **do**
3: $\quad \sigma_i = 2^{-(G-i)}$.
4: $\quad u_i =$ `TDescent`(x, v_i, σ_i, T).
5: $\quad v_{i+1} =$ `NDescent`(x, u_i, σ_i, T).
6: **end for**
7: $\tilde{x} := \operatorname{argmin}\{f(x) : x \in \{x\} \cup \{x - \sigma_i \frac{v_{i+1}}{\|v_{i+1}\|} : \sigma_i \leq \frac{\|v_{i+1}\|}{s}, i = 0, \ldots, G-1\}\}$.
8: **return** $\tilde{x}$.

Algorithm 4 NTDescent$(x, g, c_0, \{G_k\}, \{T_k\})$

Require: $sg = 0, c_0 \in (0, 1]$
1: **Set** $x_0 = x$ and $g_0 = g$.
2: **for** $k = 0, 1, \ldots$ **do**
3: $x_{k+1} =$ linesearch$(x_k, g_k, \max\{g_k, c_0 g_0\}, G_k, T_k)$.
4: Choose $g_{k+1} \in \partial f(x_{k+1})$.
5: **end for**

The workhorse of NTDescent is the line search procedure in Algorithm 3 (linesearch). Let us briefly comment on the structure of this method. Lines 2 through 6 of Algorithm 3 implement a line search on σ. Line 7 chooses the Goldstein subgradient that provides the most descent while enforcing the trust-region constraint $\sigma_i \le \frac{v_{i+1}}{s}$. Line 7 also ensures the NTDescent is a descent method. Within the line search procedure, we evaluate TDescent and NDescent a total of G times each. Not all of the calls to TDescent and NDescent will succeed with descent within the allotted budget T, but we will show that for typical problems, at least one will generate sufficient descent provided x_k is close enough to a local minimizer and T is sufficiently large. The reason at least one will succeed with descent is that given any x sufficiently near the solution and parameters G and T sufficiently large, linesearch will find a σ such that (x, σ) is in either the normal or tangent region described in Sect. 1.3.1. The line search allows the possibility that σ is as large as $1/2$, which might force x_{k+1} to leave the region surrounding the minimizer $\bar{x}$. This concern is what motivates the somewhat unusual structure of the line search method wherein the MinNorm-type methods are nested. Indeed, on the one hand, the nesting ensures the norms of the Goldstein subgradients v_{i+1} are decaying as σ_i increases. On the other hand, the trust region constraint ensures that σ_i is not chosen too large, which we need for two technical reasons in our analysis: (i) it prevents x_{k+1} from leaving a small neighborhood around the minimizer where our regularity assumptions hold; (ii) we can only ensure TDescent terminates quickly when $\sigma \le \delta_{\text{Grid}}$, for a certain radius δ_{Grid} defined in Lemma 5.7, which may be substantially smaller than $1/2$.

Computationally, it may at first seem desirable to drop the trust region constraint. Figure 5 shows this may not be the case. We suspect the reason is two-fold: First, the trust region constraint allows us to cut off a range of σ from our search, which might otherwise waste oracle calls; indeed, since v_{i+1} is nonincreasing in i, and σ_i is increasing, once the trust region is violated, it will be violated for all larger i. Second, although we may take longer steps by disabling the trust region constraint, the amount of descent we expect is on the order of $\sigma_i v_{i+1}(g)$. Thus, since the norms v_{i+1} are nonincreasing, larger stepsizes σ_i do not necessarily translate to larger descent.

Finally, we comment on our motivation for choosing the scaling $s_k = \max\{g_k, c_0 g_0\}$ in the trust region constraint. First, note that it is possible to prove, using identical techniques, that the NTDescent converges when one replaces s_k by any positive sequence that is bounded from above and below by positive constants. For our particular choice of s_k, the term $c_0 g_0$ ensures the sequence is bounded below, while the local Lipschitz continuity of f ensures that s_k is bounded above. Second, we wish for the trust region constraint to be unaffected by rescalings of f. Our particular choice of s_k guarantees scaling invariance, since the subgradients

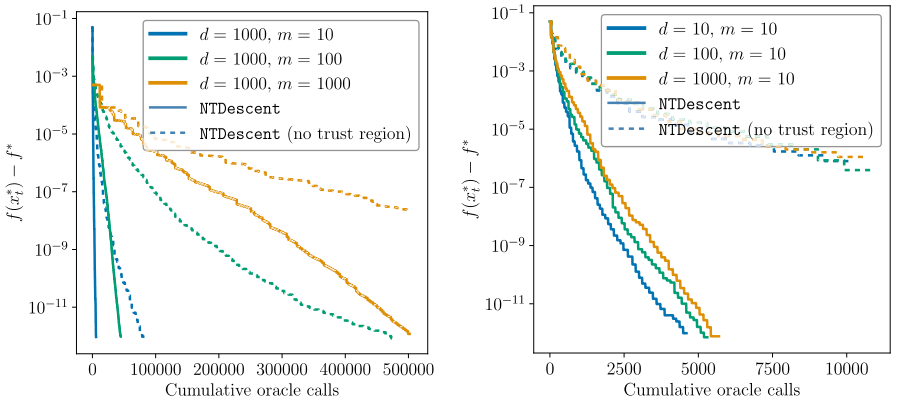

Fig. 5 Comparison of `NTDescent` on Problem (1.1) with the trust region constraint in Line 6 of Algorithm 3 removed. Left: we fix d and vary m; Right: we fix m and vary d. We invite the reader to compare these plots with Fig. 1

of $a f$ are simply the subgradients of f scaled by a for any positive constant a. One might introduce other schemes for choosing s, but we did not explore such strategies. Finally, we found that performance of `NTDescent` is relatively insensitive to the choice of $c_0 > 0$, and any $c_0 \in \{10^{-i} : i = 0, 2, 4, 6\}$ yielded adequate performance; see Fig. 6e, f.

1.5 Main Convergence Guarantees for `NTDescent`

The main contribution of this work is a local, nearly linear convergence rate for `NTDescent`. The local rate holds under a key structural assumption—Assumption A – which formalizes the concept of typical structure and mirrors the structure of the simple function considered in Sect. 1.3.2. While we formally describe Assumption A in Sect. 3, for now, we mention that it holds for max-of-smooth and properly C^p decomposable functions, provided the local minimizer $\bar{x}$ is a strong local minimizer that satisfies a strict complementarity condition; this class includes the max-of-smooth setting considered in [30]. Assumption A also holds for generic linear tilts of semi-algebraic functions: if f is semialgebraic, then for a full Lebesgue measure set of $w \in \mathbb{R}^d$, Assumption A holds at every local minimizer $\bar{x}$ of the tilted function $f_w : x \to f(x) + w \cdot x$. We now present the theorem.

Theorem 1.1 (Main convergence theorem) *Let $f : \mathbb{R}^d \to \mathbb{R}$ satisfy Assumption A at a local minimizer $\bar{x} \in \mathbb{R}^d$. Fix scalar $c_0 \in (0, 1]$, budget $\{T_k\}$ and grid size $\{G_k\}$ sequences satisfying*

$$\min\{T_k, G_k\} \geq k + 1 \quad \text{for all } k \geq 0.$$

Suppose that for initial point $x_0 \in \mathbb{R}^d$, there exists a subgradient $g_0 \in \partial f(x_0)$ such that $g_0 = 0$. Consider iterates $\{x_k\}$ generated by `NTDescent`$(x_0, g_0, c_0, \{G_k\}, \{T_k\})$.

For any $q, k_0, C > 0$, *let* $E_{k_0,q,C}$ *denote the event:*

$$f(x_k) - f(\bar{x}) \le \max\{(f(x_{k_0}) - f(\bar{x}))q^{k-k_0}, Cq^k\} \text{ for all } k \ge k_0.$$

Then there exists $q \in (0, 1)$, $C, C > 0$, *and a neighborhood* U *of* $\bar{x}$ *depending solely on* f *such that for any failure probability* $p \in (0, 1)$ *and all* $k_0 \ge C \max\{\log(1/p), 1\}$, *we have*

$$P(E_{k_0,q,C} \mid x_{k_0} \in U) \ge 1 - p,$$

provided $P(x_{k_0} \in U) > 0$. *Moreover, if* f *is convex, we have*

$$P(E_{k_0,q,C}) \ge 1 - p.$$

The theorem, which is justified in Theorems 6.3 and 6.5, bounds the function gap and distance by a quantity that geometrically decays in k. Let us examine the local complexity. Recall that each outer iteration of `NTDescent` requires at most $2T_kG_k$ first-order oracle evaluations. Thus, if $T_k = G_k = k + 1$ for all $k \ge 0$, the total number of oracle evaluations of K steps of `NTDescent` is at most $O(K^3)$. In other words, the local complexity of achieving an ε optimal solution is $O(\log^3(1/\varepsilon))$ for all sufficiently small $\varepsilon > 0$, where the big-O notation hides terms depending on the local conditioning of f; see Lemma 6.6. Therefore theorem establishes a local nearly linear rate of convergence for `NTDescent`.

1.6 Outline

The outline of this paper is as follows. In Sect. 1.7 we present notation and basic constructions. This section describes a key structure—the active manifold—and cannot be skipped. In Sect. 2, we present the sublinear convergence guarantees, which will be useful in the convex setting. This section also introduces key properties of the `NDescent` method, which will be used later in the work. In Sect. 3, we introduce our main structural assumption—Assumption A – and show that it is satisfied for the generic semialgebraic and decomposable problem classes. In Sect. 4, we show that Assumption A implies the gradient inequality (1.5). In Sect. 5 we show that the `TDescent` and `NDescent` methods terminate rapidly under appropriate conditions. In Sect. 6, we use the gradient inequality (1.5) and Assumption A to prove that `NTDescent` locally nearly linearly converges. Finally, in Sect. 7 we provide a brief numerical illustration.

1.7 Notation and Basic Constructions

We use standard convex analysis notation as set out in the monographs [16, 48]. Throughout, $\mathbb{R}^d$ denotes a d-dimensional Euclidean space with the inner product $\langle \cdot, \cdot \rangle$ and the induced norm $\|x\| = \sqrt{\langle x, x \rangle}$. We denote the open ball of radius $\varepsilon > 0$ around a point $x \in \mathbb{R}^d$ by the symbol $B_\varepsilon(x)$. We use the symbol $\bar{B}$ to denote the closed unit

ball at the origin. For any set $X \subseteq \mathbb{R}^d$, the *distance function* and the *projection map* are defined by

$$\operatorname{dist}(x, X) := \inf_{y \in X} \|y - x\| \qquad \text{and} \qquad P_X(x) := \operatorname*{argmin}_{y \in X} \|y - x\|,$$

respectively. Note that the function $\operatorname{dist}(\cdot, X)$ is 1-Lipschitz for any set X. For any set $X \subseteq \mathbb{R}^d$, all $\bar{x} \in X$, all $x \in \mathbb{R}^d$, and all $y \in P_X(x)$, we have

$$\|y - \bar{x}\| \leq 2\|x - \bar{x}\|.$$

We denote the diameter of a set X by

$$\operatorname{diam}(X) = \sup_{x, y \in X} \|x - y\|.$$

We call a function $h : \mathbb{R}^d \to \mathbb{R}$ *sublinear* if its epigraph is a closed convex cone, and in that case we define

$$\operatorname{Lin}(h) := \{x \in \mathbb{R}^d : h(x) = -h(-x)\}$$

to be its *lineality space.* Given a mapping $F : \mathbb{R}^d \to \mathbb{R}^m$ and a point $\bar{x} \in \mathbb{R}^d$, we define

$$\operatorname{lip}_F(\bar{x}) := \limsup_{\substack{x, x' \to \bar{x} \\ x \neq x'}} \frac{\|F(x) - F(x')\|}{\|x - x'\|}.$$

Given a mapping $F : \mathbb{R}^d \to \mathbb{R}^{m \times n}$ into the space of $m \times n$ matrices and a point $\bar{x} \in \mathbb{R}^d$ then we define

$$\operatorname{lip}_F^{\mathrm{op}}(\bar{x}) := \limsup_{\substack{x, x' \to \bar{x} \\ x \neq x'}} \frac{\|F(x) - F(x')\|_{\mathrm{op}}}{\|x - x'\|},$$

where $\|\cdot\|_{\mathrm{op}}$ denotes the operator norm defined on $\mathbb{R}^{m \times n}$.

Semialgebraicity. We call a set $X \subseteq \mathbb{R}^d$ *semialgebraic* if it is the union of finitely many sets defined by finitely many polynomial inequalities. Likewise, we call a function $f : \mathbb{R}^d \to \mathbb{R}$ semialgebraic if its graph $\operatorname{gph}(f) = \{(x, f(x)) : x \in \mathbb{R}^d\}$ is semialgebraic.

Subdifferentials. Consider a locally Lipschitz function $f : \mathbb{R}^d \to \mathbb{R}^d$ and a point $\bar{x} \in \mathbb{R}$. The Clarke subdifferential is the convex hull of limits of gradients evaluated at nearby points of differentiability:

$$\partial f(\bar{x}) = \operatorname{conv} \lim_{i \to \infty} \nabla f(x_i) : x_i \to \bar{x},$$

where $\mathbb{R}^d$ is the set of points at which f is differentiable (recall Radamacher's theorem). If f is L-Lipschitz on a neighborhood U, then

$$\text{for all } x \in U \text{ and } v \in \partial f(x), \text{ we have } \|v\| \leq L.$$

This fact will be used throughout the paper. A point $\bar{x}$ satisfying $0 \in \partial f(\bar{x})$ is said to be critical for f. The Goldstein subdifferential, which appears in (1.2), will be a central object throughout. An important fact is that $\partial_\sigma f(x)$ is a closed convex set for any $x \in \mathbb{R}^d$ and $\sigma > 0$.

Manifolds. We will need a few basic results about smooth manifolds, which can be found in the references [10, 34]. A set $M \subseteq \mathbb{R}^d$ is called a C^p-smooth manifold around $\bar{x}$ (with $p \geq 1$) if there exists a natural number m, an open neighborhood U of $\bar{x}$, and a C^p smooth mapping $F : U \to \mathbb{R}^m$ such that the Jacobian $\nabla F(x)$ is surjective and $M \cap U = F^{-1}(0)$. The tangent and normal spaces to M at $x \in M$ near $\bar{x}$ are defined to be $T_M(x) = \ker(\nabla F(x))$ and $N_M(x) = T_M(x)^\perp = \text{range}(\nabla F(x)^*)$, respectively. If M is a C^2-smooth manifold around a point $\bar{x}$, then there exists $C > 0$ such that $y - x \in T_M(x) + C\|y - x\|^2 B$ for all $x, y \in M$ near $\bar{x}$. We also have that $x - P_M(x) \in N_M(P_M(x))$ for all x near $\bar{x}$. Moreover, the projection mapping $P_M : \mathbb{R}^d \to \mathbb{R}^d$ is C^{p-1} smooth on a neighborhood of $\bar{x}$ and satisfies $\nabla P_M(x) = P_{T_M(x)}$ for all $x \in M$ near $\bar{x}$.

Covariant gradients and smooth extension Let $M \subset \mathbb{R}^d$ be a C^p-manifold around a point x for some $p \geq 1$. Then a function $f : M \to \mathbb{R}$ is called C^q-smooth (with $q \geq 1$) around the point x if there exists a C^q function $\hat{f} : U \to \mathbb{R}$ defined on an open neighborhood U of x and that agrees with f on $U \cap M$. In that case, the projection of $\nabla \hat{f}(x)$ onto $T_M(x)$ is independent of the choice of $\hat{f}$. We call this projection the *covariant gradient of f at x* and denote it by

$$\nabla_M f(x) := P_{T_M(x)}(\nabla \hat{f}(x)).$$

For example, the smooth extension

$$f_M := f \circ P_M$$

of f is $C^{\min\{p-1,q\}}$ smooth on a neighborhood of x and agrees with f along M. Thus, we will use the identification: $\nabla_M f(x) := \nabla f_M(x)$.

Active Manifolds. In this work, we will assume the local minimizer of interest lies on an *active manifold.* Informally, an active manifold is a smooth manifold along which the function varies smoothly and off of which the function varies sharply. We adopt the formal model of activity explicitly used in [26]. Related models exist, e.g., identifiable surfaces [52], manifolds of partial smoothness [38], VU-structures [36, 41], and $g \circ F$ decomposable functions [49].

Definition 1.2 *(Active manifold)* Consider a function $f : \mathbb{R}^d \to \mathbb{R}$ and fix a set $M \subseteq \mathbb{R}^d$ containing a point $\bar{x}$ satisfying $0 \in \partial f(\bar{x})$. Then M is called an *active*

C^p*-manifold around* $\bar{x}$ if there exists a neighborhood U of $\bar{x}$ such that the following are true:

- **(smoothness)** The set M is a C^p-smooth manifold near $\bar{x}$ and the restriction of f to M is C^p-smooth near $\bar{x}$.
- **(sharpness)** The lower bound holds:

$$\inf\{v : v \in \partial f(x),\ x \in U \setminus M\} > 0.$$

In the introduction, we provided two examples of functions that admit an active manifold at their minimizers. For example, function $f(u, v) = u^2 + |v|$ admits the active manifold $M := \{(u, 0) : u \in \mathbb{R}\}$ around the origin. On the other hand, the function f from (1.1) admits the active manifold $M = \{x \in \mathbb{R}^d : x_1 = x_2 = \ldots x_m\}$ around the origin. To draw a distinction with partial smoothness property of [38], the function $f(x, y) = \max\{x, 0\} + y^2$ is partly smooth along the x axis, but the x-axis is not an active manifold for f around the origin; indeed, f does not satisfy the sharpness condition at the origin. We now turn to sublinear convergence guarantees.

2 Global Sublinear Convergence of `NTDescent`

The main goal of this work is to show that `NTDescent` locally converges nearly linearly for "typical" nonsmooth optimization problems. A natural question is whether `NTDescent` also possesses global nonasymptotic convergence guarantees. In this section, we prove two such guarantees: First, for arbitrary Lipschitz functions, we analyze the rate at which $\operatorname{dist}(0, \partial_{\sigma_i} f(x_k))$ tends to zero as a function of k. Second, for convex Lipschitz functions, we analyze the rate at which $f(x_k)$ tends to $\inf f$.

In the proofs of this section, the `TDescent` loop is ignored as we can only prove it terminates with descent near the minimizer. Instead, the global convergence guarantees follow from the properties of `NDescent`. Thus, our analysis essentially follows that of [21], where a nearly identical `MinNorm` method was introduced. The main difference between the `NDescent` and the method of [21] lies in the perturbation radius in Line 3 of Algorithm 1: while the radius of `NDescent` can be computed with access only to σ_{g_t}, the radius in [21] requires knowledge of the Lipschitz constant of f, which we do not assume. Finally, we mention that [21] did not consider rates of convergence for convex problems.

Before stating the main result, we recall three key Lemmas, which underlie the proof. The first lemma shows that the vectors u_i and v_i generated by `linesearch` are Goldstein subgradients of decreasing norm.

Lemma 2.1 (Properties of `linesearch`) *Let* $f : \mathbb{R}^d \to \mathbb{R}$ *be a locally Lipschitz function. Fix* $x \in \mathbb{R}^d$, *subgradient* $g \in \partial f(x)$, *budget* T, *and grid size* G. *Let* u_i *and* v_i *be generated by* `linesearch`(x, g, G, T). *Then*

$$u_i, v_{i+1} \in \partial_{\sigma_i} f(x) \quad \text{and} \quad v_{i+1} \le u_i \le v_i \tag{2.1}$$

for all $i = 0, \ldots, G - 1$.

Proof The proof follows by induction. We prove the base case only, since the induction is straightforward. First note that the inclusion $v_0 \in \partial f(x)$ implies that $u_0 \in \partial_{\sigma_0} f(x)$, since `TDescent` constructs u_0 as a convex combinations of subgradients evaluated in the ball $\overline{B}_{\sigma_0}(\bar{x})$. Likewise, due to the argmin operation on line 4 of Algorithm 2, the subgradients generated by `TDescent` are decreasing in norm. Consequently, we have $\|u_0\| \leq \|v_0\|$. A similar argument shows that $v_1 \in \partial_{\sigma_0} f(x)$ and $\|v_1\| \leq \|u_0\|$. This completes the proof.

The next lemma shows that when f is convex, the minimal norm Goldstein subgradient may be used to bound the function values. We place the proof in Appendix A, since it follows from a standard argument.

Lemma 2.2 (Subgradient inequality) *Suppose that* $f : \mathbb{R}^d \to \mathbb{R}$ *is a continuous convex function. Let* $x, y \in \mathbb{R}^d$. *Let* L *denote a Lipschitz constant for* f *on the ball* $B_{2\sigma}(x)$. *Then*

$$f(x) - f(y) \leq \|x - y\| \operatorname{dist}(0, \partial_\sigma f(x)) + 2\sigma L.$$

The final lemma provides conditions under which `NDescent` terminates with descent with high probability. The result is closely related to [21, Corollary 2.6], but we take extra care to analyze the perturbation radius in Line 3 of Algorithm 1.

Lemma 2.3 (`NDescent` loop terminates with descent) *Let* f *be a locally Lipschitz function. Fix initial point* $x \in \mathbb{R}^d$, *radius* $\sigma > 0$, *subgradient* $g \in \partial_\sigma f(x)$, *and failure probability* $p \in (0, 1)$. *Furthermore, let* L *be a Lipschitz constant of* f *on the ball* $B_{2\sigma}(x)$. *Suppose that*

$$\sigma \leq \frac{\operatorname{dist}(0, \partial_\sigma f(x))}{\sqrt{128} L}; \quad \text{and} \quad T \geq \frac{64 L^2}{\operatorname{dist}^2(0, \partial_\sigma f(x))} 2 \log(1/p).$$

Define $g_+ := \texttt{NDescent}(x, g, \sigma, T)$. *Then* $g_+ = 0$ *and the point* $x_+ := x - \sigma \frac{g_+}{\|g_+\|}$ *satisfies*

$$f(x_+) \leq f(x) - \frac{\sigma \operatorname{dist}(0, \partial_\sigma f(x))}{8} \quad \text{with probability at least } 1 - p.$$

Proof First note that $g_+ \in \partial_\sigma f(x)$, so $\|g_+\| \geq \operatorname{dist}(0, \partial_\sigma f(x)) > 0$. Now, observe that `NDescent` is precisely [21, Algorithm 1] with a different bound on the perturbation radius r. Indeed, in [21, Algorithm 1], r must satisfy

$$r < \|g_t\| \sqrt{1 - \left(1 - \frac{\|g_t\|^2}{128 L^2}\right)^2}$$

for all $t \geq 0$. We now show that the constraint $r \leq \sigma \|g_t\| \sqrt{1 - (1 - \frac{a^2}{128 L^2})^2}$. Then h

is increasing in a for $a \leq L$. Moreover, for $a \in [0, L]$, we have $h(a) \geq \sqrt{\frac{a}{128L}}$. Consequently, since

$$\operatorname{dist}(0, \partial_\sigma f(x)) \leq g_t \leq L$$

for all $t \leq T$, we have

$$r < \sigma \quad g_t \leq \frac{\operatorname{dist}(0, \partial_\sigma f(x)) \quad g_t}{\sqrt{128L}} \leq h(\operatorname{dist}(0, \partial_\sigma f(x))) \quad g_t \leq h(g_t) \quad g_t.$$

Thus the proof is a direct application of [21, Corollary 2.6].

Given these lemmata, we are now ready to state and prove our main sublinear convergence guarantee.

Theorem 2.4 (Sublinear convergence) *Let $f : \mathbb{R}^d \to \mathbb{R}$ be a locally Lipschitz function. Fix initial point $x_0 \in \mathbb{R}^d$ and subgradient $g_0 \in \partial f(x_0)$. Assume that $g_0 = 0$. Let $L \in \mathbb{R} \cup \{+\infty\}$ be any Lipschitz constant of f over the widened sublevel set*

$$S := \{x + u : f(x) \leq f(x_0) \text{ and } u \in \overline{B}(0)\}.$$

Fix a scalar $c_0 \in (0, 1]$, budget sequence $\{T_k\}$, grid size sequence $\{G_k\}$, and failure probability $p \in (0, 1)$. Let $\{x_k\}$ be generated by NTDescent$(x, g, c_0, \{G_k\}, \{T_k\})$. *Then for all $K > 0$, the following holds with probability at least $1 - p$: Define $G := \min_{K \leq k \leq 2K-1} G_k$ and $T := \min_{K \leq k \leq 2K-1} T_k$. Then for all $i \leq G$, the following bound holds with $\sigma_i := 2^{-(G-i)}$:*

$$\min_{K \leq k \leq 2K-1} \operatorname{dist}(0, \partial_{\sigma_i} f(x_k)) \leq \max \left\{ \frac{8(f(x_K) - \inf f)}{\sigma_i K}, \frac{16L \sqrt{2\log(KG/p)}}{\sqrt{T}}, \sqrt{128L\sigma_i} \right\}.$$

Now suppose that f is convex and define $D := \operatorname{diam}(\{x \in \mathbb{R}^d : f(x) \leq f(x_0)\})$. Then

$$f(x_{2K-1}) - \inf f \leq \min_{i \leq G} D \max \left\{ \frac{8(f(x_K) - \inf f)}{\sigma_i K}, \frac{16L \sqrt{2\log(KG/p)}}{\sqrt{T}}, \sqrt{128L\sigma_i} \right\} + 2L\sigma_i. \tag{2.2}$$

Proof Let us assume that $L < +\infty$; otherwise the result is trivial. Fix $K > 0$ and $i \leq G$. Define

$$\epsilon_i := \max \left\{ \frac{16L \sqrt{2\log(KG/p)}}{\sqrt{T}}, \sqrt{128L\sigma_i} \right\}.$$

For every $K \le k \le 2K-1$, define

$$x_{k,i} := x_k - \sigma_i \frac{v_{i+1}}{\|v_{i+1}\|}, \qquad \text{where } v_{i+1} := \texttt{NDescent}(x_k, u_i, \sigma_i, T_k),$$

and u_i appear in the definition of $\texttt{linesearch}(x_k, g_k, \max\{\|g_k\|, c_0\|g_0\|\}, G_k, T_k)$; see Algorithm 3. Note that $v_{i+1} \in \partial_{\sigma_i} f(x_k)$ by Lemma 2.1. Thus, in the event $\{\operatorname{dist}(0, \partial_{\sigma_i} f(x_k)) \ge \varepsilon_i\}$, we have

1. $x_{k,i}$ is well-defined since $\|v_{i+1}\| \ne 0$;
2. the trust region constraint $\sigma_i \le \frac{\|v_{i+1}\|}{s}$ is satisfied for $s = \max\{\|g_k\|, c_0\|g_0\|\}$ (in Algorithm 3); indeed,

$$\frac{\|v_{i+1}\|}{s} \ge \frac{\operatorname{dist}(0, \partial_{\sigma_i} f(x_k))}{s} \ge \frac{\sqrt{128L\sigma_i}}{s} \ge \sigma_i,$$

where the final inequality follows from the bound $s \le L$, a consequence of the inclusion $x_0 \subseteq \operatorname{int} S$ and the Lipschitz continuity of f on S.

Finally, for every $K \le k \le 2K-1$, define

$$A_{k,i} := \left\{ f(x_{k,i}) - f(x_k) \ge -\frac{\sigma_i \operatorname{dist}(0, \partial_{\sigma_i} f(x_k))}{8} \right\} \cap \{\operatorname{dist}(0, \partial_{\sigma_i} f(x_k)) \ge \varepsilon_i\}.$$

Now we apply Lemma 2.3.

To that end, observe that since $f(x_k)$ is nonincreasing and $\sigma_i \le 1/2$, every iterate x_k satisfies $B_{2\sigma_i}(x_k) \subseteq S$. Consequently, L is a Lipschitz constant of f on $B_{2\sigma_i}(x_k)$. Therefore, by Lemma 2.3, for every $K \le k \le 2K-1$, we have

$$P(A_{k,i}) \le P(A_{k,i} \mid \operatorname{dist}(0, \partial_{\sigma_i} f(x_k)) \ge \varepsilon_i) \le \frac{p}{GK}. \tag{2.3}$$

Thus, by a union bound, with probability at least $1 - \frac{p}{G}$, at least one of the following must hold at every index $K \le k \le 2K-1$:

$$f(x_{k,i}) - f(x_k) \le -\frac{\sigma_i \operatorname{dist}(0, \partial_{\sigma_i} f(x_k))}{8} \qquad \text{or} \qquad \operatorname{dist}(0, \partial_{\sigma_i} f(x_k)) \le \varepsilon_i.$$

If $\operatorname{dist}(0, \partial_{\sigma_i} f(x_k)) \le \varepsilon_i$ for some k satisfying $K \le k \le 2K-1$, then the result follows. On the other hand, suppose that for all $K \le k \le 2K-1$, we have $\operatorname{dist}(0, \partial_{\sigma_i} f(x_k)) > \varepsilon_i$; in particular, we have $\operatorname{dist}(0, \partial_{\sigma_i} f(x_k)) > \sqrt{128L\sigma_i}$. Therefore, with probability at least $1 - \frac{p}{G}$, we must have

$$f(x_{k+1}) \le f(x_{k,i}) \le f(x_k) - \frac{\sigma_i \operatorname{dist}(0, \partial_{\sigma_i} f(x_k))}{8}, \qquad \text{for all } K \le k \le 2K-1.$$

where the first inequality follows since the trust region constraint is satisfied for $x_{k,i}$. Iterating this inequality, we have with probability at least $1 - \frac{p}{G}$, the bound

$$\min_{K \le k \le 2K-1} \operatorname{dist}(0, \partial_{\sigma_i} f(x_k)) \le \frac{1}{K} \sum_{k=K}^{2K-1} \operatorname{dist}(0, \partial_{\sigma_i} f(x_k)) \le \frac{8(f(x_K) - f(x_{2K}))}{\sigma_i K}.$$

This proves the result for i. Taking a union bound over i then yields the bound for minimal norm Goldstein subgradient for all $i \le G$.

To prove (2.2), fix an $i \le G$ and let k_i be the index that attains the minimum. Then

$$f(x_{2K-1}) - \inf f \le f(x_{k_i}) - \inf f \le \operatorname{dist}(x_{k_i}, X_\star) \min_{K \le k \le 2K-1} \operatorname{dist}(0, \partial_{\sigma_i} f(x_k)) + 2\sigma_i L,$$

where the first inequality follows since $f(x_k)$ is nonincreasing and the second inequality follows from Lemma 2.2. The proof then follows from the upper bound $\operatorname{dist}(x_{k_i}, X) \le D$.

theorem provides bounds on the minimal norm Goldstein subgradient within any window of indices $K \le k \le 2K - 1$. Let us briefly investigate the setting $T_k = k + 1$ for all $k \ge 0$. In this case, theorem implies that with probability at least $1 - p$, we have

$$\min_{K \le k \le 2K-1} \operatorname{dist}(0, \partial_{\sigma_i} f(x_k)) \le \max \left\{ \frac{8(f(x_K) - \inf f)}{\sigma_i K}, \frac{16L \sqrt{2\log(KG/p)}}{\sqrt{2K}}, \sqrt{128} L \sigma_i \right\}$$

for all $i \le G$. Let us now suppose G is large enough that there exists $i \le G$ satisfying $(1/2)K^{-1/2} \le \sigma_i \le K^{-1/2}$, e.g., we may assume $G_k = \log(k^{1/2}))$ for all $k > 0$. Then, we find that at most $O(KTG) = O(K^2 G)$ first-order oracle evaluations are needed to find a point x_k satisfying

$$\operatorname{dist}(0, \partial_{K^{-1/2}} f(x_k)) = \tilde{O}(K^{-1/2}),$$

where $\tilde{O}$ hides logarithmic terms in G, K and p. Let's consider two settings for G_k.

1. **Setting 1**: $G_k = O(\log(k^{1/2}))$. In this case, `NTDescent` finds a point x_k satisfying $\operatorname{dist}(0, \partial_\varepsilon f(x_k)) \le \varepsilon$ using at most $\tilde{O}(\varepsilon^{-4})$ first-order oracle evaluations.
2. **Setting 2**: $G_k = k + 1$. In this case, `NTDescent` finds a point x_k satisfying $\operatorname{dist}(0, \partial_\varepsilon f(x_k)) \le \varepsilon$ using at most $\tilde{O}(\varepsilon^{-6})$ first-order oracle evaluations.

The complexity of Setting 1 is more favorable than the complexity of Setting 2. Nevertheless, when we establish our local rapid convergence guarantees, we will work in Setting 2, which has more favorable local convergence properties. Before moving on, we note that the above guarantees likewise apply in the convex setting, namely

NTDescent finds a point x_k with $f(x_k) - f^* \le \varepsilon$ using at most $\tilde{O}(\varepsilon^{-4})$, respectively $\tilde{O}(\varepsilon^{-6})$, first-order oracle evaluations in Setting 1, respectively Setting 2.

In addition to the nonasymptotic guarantees of Theorem 2.4, the reader may wonder whether a given limit point $\bar{x}$ of NTDescent is Clarke critical, meaning $0 \in \partial f(\bar{x})$. We prove that this is indeed the case under a bounded sublevel set condition. We place the proof in Appendix C since it follows a similar line of reasoning as Theorem 2.4.

Corollary 2.5 (Limiting points are Clarke critical) *Let $f : \mathbb{R}^d \to \mathbb{R}$ be a locally Lipschitz function. Fix initial point $x_0 \in \mathbb{R}^d$ and subgradient $g_0 \in \partial f(x_0)$. Assume that $g_0 \neq 0$. Suppose the sublevel set $\{x : f(x) \le f(x_0)\}$ is bounded. Fix scalar $c_0 \in (0, 1]$, budget sequence $\{T_k\}$, grid size sequence $\{G_k\}$ such that $\{G_k\}$ tends to infinity and $T_k \ge k$. Let $\{x_k\}$ be generated by* NTDescent$(x, g, c_0, \{G_k\}, \{T_k\})$. *Then with probability one, all the limiting points of $\{x_k\}$ are Clarke critical.*

This concludes our sublinear convergence guarantees for NTDescent. In the following section, we describe the key structural assumptions needed to ensure that NTDescent locally rapidly converges.

3 Main Assumption, Examples, and Consequences

In this section, we introduce our key structural assumption—Assumption A. In Sect. 3.1 we show that Assumption A holds for generic semialgebraic functions and certain properly C^p decomposable functions. Then, in Sect. 3.2, we extract several key consequences of Assumption A. These consequences will be instrumental in proving the gradient inequality (1.5) and rapid convergence of NTDescent. We now turn to the assumption.

Assumption A Function $f : \mathbb{R}^d \to \mathbb{R}$ is locally Lipschitz with local minimizer $\bar{x} \in \mathbb{R}^d$.

(A1) **(Quadratic Growth)** There exists $\gamma > 0$ such that

$$f(x) - f(\bar{x}) \ge \frac{\gamma}{2} \|x - \bar{x}\|^2 \quad \text{for all } x \text{ near } \bar{x}.$$

(A2) **(Active Manifold)** Function f admits a C^4-smooth active manifold M around $\bar{x}$.

(A3) **(Strong-(a) regularity)** There exists $C_{(a)} > 0$ such that

$$\|P_{T_M(y)}(v - \nabla_M f(y))\| \le C_{(a)} \|x - y\|$$
$$\text{for all } x \in \mathbb{R}^d,\ v \in \partial f(x), \text{ and } y \in M \text{ near } \bar{x}.$$

(A4) **($(b_\le)$-regularity**) The following inequality holds

$$f(y) \ge f(x) + \langle v, y - x \rangle + o(\|y - x\|)$$
$$\text{as } y \xrightarrow{M} \bar{x} \text{ and } x \to \bar{x} \text{ with } v \in \partial f(x),$$

where $o(\cdot)$ is any univariate function satisfying $\lim_{t \to 0} o(t)/t = 0$.

Some comments are in order. Assumption (A1) is a classical regularity condition that ensures local linear convergence of gradient methods for smooth convex functions. Assumptions (A2), (A3), and (A4) describe the interaction of f and a distinguished smooth manifold M. Assumption (A2) requires M to be an active manifold for f around $\bar{x}$ in the sense of Definition 1.2. In particular, along the manifold M, the function f is C^4 smooth with covariant gradient $\nabla_M f$; see Sect. 1.7 for a definition. Assumption (A3) shows that in tangent directions the covariant gradient along the manifold approximates the subgradients of f up to a linear error. This property recently appeared in [5, 19], where it was used to study saddle avoidance properties of the subgradient method for nonsmooth optimization. Finally, Assumption (A4) is a restricted lower smoothness property, showing that linear models of f off the manifold are underapproximators of f on the manifold up to first-order. Note that the property is automatic if f is weakly convex, meaning the mapping $x \to f(x) + \frac{\rho}{2}\|x\|^2$ is convex for some $\rho \geq 0$. The weakly convex class is broad and contains all compositions of convex functions with smooth mappings that have Lipschitz Jacobians; see the survey [17] for an introduction. We mention that the name "$(b_{\leq})$-regularity" is motivated by "uniform semismoothness" property of [19], which was called the "(b)-regularity property."

In the following section, we provide examples of functions satisfying Assumption A.

3.1 Examples of Assumption A

In this section, we show that the aforementioned problems satisfy Assumption A. The most important example is the class of generic semialgebraic functions. The following theorem is essentially contained in [19, 23], but we provide a proof for completeness.

Theorem 3.1 (Generic semialgebraic functions) *Consider a locally Lipschitz semialgebraic function $f : \mathbb{R}^d \to \mathbb{R}$. Then for a full Lebesgue measure set of $w \in \mathbb{R}^d$, the tilted function $f_w : x \to f(x) + w^* x$ satisfies Assumption A at every local minimizer.*

Proof The proof is a consequence of [19, Theorem 3.31] and [23, Corollary 4.8, Theorem 4.16]. A combination of Corollary 4.8 and Theorem 4.16 in [23] shows that for a full Lebesgue measure set of $w \in \mathbb{R}^d$, the following hold: every local minimizer $\bar{x}$ of f_w lies on a C^4 active manifold M, verifying (A2); and the quadratic growth condition (A1) holds at $\bar{x}$. Next, [19, Theorem 3.31] shows that f_w also satisfies the strong (a) property (A3) along M; applying [19, Theorem 3.11 and Theorem 3.4], we deduce that f_w also satisfies the $(b_{\leq})$-regularity property (A4) along M at $\bar{x}$.

Turning to our second class, we introduce so-called *properly C^p decomposable* functions, originally proposed and analyzed in [49]. At a high-level, the class consists of functions that are locally the composition of a sublinear function with a smooth mapping, which together satisfy a transversality condition.

Definition 3.2 (Decomposable functions) A function $f : \mathbb{R}^d \to \mathbb{R}$ is called *properly C^p decomposable at $\bar{x}$ as $h \circ c$* if near $\bar{x}$ it can be written as

$$f(x) = f(\bar{x}) + h(c(x))$$

for some C^p-smooth mapping $c : \mathbb{R}^d \to \mathbb{R}^m$ satisfying $c(\bar{x}) = 0$ and some proper, closed sublinear function $h : \mathbb{R}^m \to \mathbb{R}$ satisfying the transversality condition:

$$\mathrm{lin}(h) + \mathrm{range}(\nabla c(\bar{x})) = \mathbb{R}^m.$$

The following theorem shows that decomposable functions satisfy Assumption A near local minimizers if they also satisfy a strict complementarity condition and a quadratic growth bound. The proof is a consequence of results found in works [19, 26, 38, 49].

Theorem 3.3 (Properly decomposable functions) *Consider a locally Lipschitz function $f : \mathbb{R}^d \to \mathbb{R}$. Let $\bar{x}$ be a local minimizer of f and suppose that f is properly C^4 decomposable at $\bar{x}$. Furthermore, suppose that*

1. **(Strict Complementarity)** *We have that $0 \in \mathrm{ri}\, \partial f(\bar{x})$.*
2. **(Quadratic growth)** *There exists $\gamma > 0$ such that*

$$f(x) - f(\bar{x}) \geq \frac{\gamma}{2} \|x - \bar{x}\|^2 \quad \text{for all } x \text{ near } \bar{x}.$$

Then f satisfies Assumption A at $\bar{x}$.

Proof To set the notation for the proof, recall that since f is properly C^4 decomposable, there exist functions h and c satisfying the conditions of Definition 3.2. The discussion in [49, p. 683-4] then shows that the set

$$M := c^{-1}(\mathrm{lin}(h))$$

is a so-called C^4 *manifold of partial smoothness* for f around $\bar{x}$ in the sense of Lewis [38]. Moreover, f is prox-regular at $\bar{x}$ for 0 in the sense of [46, Definition 1.1], since by definition it is *strongly amenable* [46, Definition 2.4] at $\bar{x}$; see [46, Proposition 2.5]. Thus, according to [31, Theorem 5.3], partial smoothness, prox-regularity, and strict complementarity ensure that the sharpness condition of Definition 1.2 holds. Consequently, M is a C^4 smooth active manifold around $\bar{x}$, verifying (A2). In addition, [19, Corollary 3.24] ensures that f satisfies the (A3) and (A4) properties along M.

A popular class of decomposable objectives arises from pointwise maxima of smooth functions that satisfy an affine independence property. For example, this class was considered in the work of Han and Lewis [30]. As an immediate corollary of Theorem 3.3, we show that such functions satisfy Assumption A.

Corollary 3.4 (Max-of-smooth functions) *Consider a locally Lipschitz function f and a family of C^4 smooth functions $f_i : \mathbb{R}^d \to \mathbb{R}$ indexed by a finite set $i \in I$. Fix a local*

minimizer $\bar{x}$ *of* f *and suppose the set* $\{\nabla f_i(\bar{x})\}_{i\in I}$ *is affinely independent. Suppose furthermore that* f *is locally expressible as*

$$f(x) := \max_{i\in I} f_i(x) \quad \text{for all } x \text{ near } \bar{x}.$$

Then provided the strict complementarity and quadratic growth conditions of Theorem 3.3 hold, the function f *satisfies Assumption A at* $\bar{x}$.

Proof To prove the result, note that the affine independence property is simply a restatement of the transversality condition of Definition 3.2 for the smooth mapping $x \to (f_i(x))_{i\in I}$ and the sublinear function $y \to \max_{i\in I} y_i$.

We now turn our attention to the key consequences of Assumption A.

3.2 Key Consequences of Assumption A

The following proposition summarizes the key consequences of Assumption A. The proof of the result is straightforward but technical, so we place it in Appendix B.

Proposition 3.5 (Consequences of Assumption A) *Suppose* f *satisfies Assumption A at* $\bar{x}$. *Then there exists* $\delta_A > 0$ *such that on the ball* $B_{2\delta_A}(\bar{x})$, *the projection operator* P_M *is* C^3 *with Lipschitz Jacobian and the smooth extension* $f_M := f \circ P_M$ *is* C^3 *with Lipschitz gradient. Moreover, the following bounds hold:*

1. **(Quadratic growth)** *The quadratic growth bound (A1) holds throughout* $B_{2\delta_A}(\bar{x})$.
2. *(***Smoothness of** P_M *) For all* $x \in B_{\delta_A}(\bar{x})$ *and* $x' \in B_{2\delta_A}(\bar{x})$, *we have*

$$P_M(x') - P_M(x) - P_{T_M(P_M(x))}(x' - x) \le C_M(\operatorname{dist}^2(x, M) + \|x' - x\|^2), \tag{3.1}$$

   where $C_M := 2\operatorname{lip}^{\mathrm{op}}_{\nabla P_M}(\bar{x})$.
3. *(***Bounds on** $\nabla_M f$ *) For all* $x \in B_{\delta_A}(\bar{x})$, *we have*

$$\frac{\gamma}{2}\|P_M(x) - \bar{x}\| \le \|\nabla_M f(P_M(x))\| \le \beta\|P_M(x) - \bar{x}\|, \tag{3.2}$$

   where $\beta := 2\operatorname{lip}_{\nabla f_M}(\bar{x})$.
4. *(***Consequence of strong** (a)*) For all* $x \in B_{\delta_A}(\bar{x})$ *and* $\sigma \le \delta_A$, *we have*

$$\sup_{g\in\partial_\sigma f(x)} \langle P_{T_M(P_M(x))}(g - \nabla_M f(P_M(x))), g\rangle \le C_{(a)}(\operatorname{dist}(x, M) + \sigma); \tag{3.3}$$

$$\sup_{g\in\partial_\sigma f(x)} \|P_{T_M(P_M(x))} g\| \le C_{(a)}(\operatorname{dist}(x, M) + \sigma) + \beta\|P_M(x) - \bar{x}\|; \tag{3.4}$$

$$\sup_{g,g'\in\partial_\sigma f(x)} \|P_{T_M(P_M(x))}(g - g')\| \le 2C_{(a)}(\operatorname{dist}(x, M) + \sigma). \tag{3.5}$$

5. **(Aiming)** *For all* $x \in B_{\delta_A}(\bar{x})$ *and all* $v \in \partial f(x)$, *we have*

$$\langle v, x - P_M(x) \rangle \geq \mu \operatorname{dist}(x, M), \tag{3.6}$$

where $\mu := \frac{1}{4} \liminf_{\substack{M \\ x \to \bar{x}}} \operatorname{dist}(0, \partial f(x))$.

6. **(Subgradient bound)** *For all* $x \in B_{\delta_A}(\bar{x})$ *and* $\sigma \leq \delta_A$, *we have*

$$\sup_{g \in \partial_\sigma f(x)} \|g\| \leq L,$$

where $L := 2\operatorname{lip}_f(\bar{x})$

7. **(Function gap)** *For all* $x \in B_{\delta_A}(\bar{x})$, *we have*

$$f(x) - f(\bar{x}) \leq L \operatorname{dist}(x, M) + \frac{\beta}{2} \|P_M(x) - \bar{x}\|^2. \tag{3.7}$$

Let us briefly comment on the result. Item 2 provides a crucial smoothness property of the projection operator of M. Item 3 shows that the Riemannian gradient of f is proportional to the distance of the projection y to $\bar{x}$. Item 4 shows how the Goldstein subgradients inherit the strong (a) property (A3) of Assumption A. Indeed, Equation (3.4) shows that Goldstein subgradients are "small" in tangent directions and Equation (3.5) shows Goldstein subgradients vary in an approximate Lipschitz fashion in tangent directions. Item 5 shows that the subgradients of f off of the manifold have a constant level of correlation with $x - P_M(x)$, i.e., the direction $-v$ "aims" towards the manifold. Note that $\mu > 0$ due to the active manifold Assumption (A2). The proof of Item 5 is based on Assumptions (A2) and (A4); a similar result appears in [18, Theorem D.2]. Item 6 provides a bound on the Goldstein subgradients of f near $\bar{x}$; we will appeal to this bound many times throughout the analysis without referencing this proposition. Finally, Item 7 decomposes the function gap into a sum of two terms: the distance to the manifold and the squared distance of the projection to the solution. The proof relies on the smoothness of f along the manifold. Note that the trivial upper bound $L\|x - \bar{x}\|$ for the gap can be weaker than (3.7).

This concludes our discussion of Assumption A. The following three sections establish further consequences: the gradient inequality (1.5) (Sect. 4); rapid local convergence of `NDescent` and `TDescent` (Sect. 5); and rapid local convergence of `NTDescent` (Sect. 6). In all three sections, we use the notation and results introduced in Proposition 3.5. Finally, the statements of the results in Sect. 4 and 5 contain several parameters/radii which we will use in Sect. 6 to determine the region of near linear convergence and the oracle complexity for `NTDescent`. For the readers' convenience, we have listed these parameters in Table 1.

Table 1 Parameters used throughout Sects. 4 and 5

Parameter	Definition
D_1	$\frac{\mu}{8(\mu + L)}$
D_2	$\frac{\mu}{2}$
C_1	$\frac{\gamma}{4}$
C_2	$\min\ \frac{\gamma}{8C_{(a)}}, \frac{\min\{1, 1/\delta_A\}}{2}$
C_3	$\frac{C_1^2}{8L}$
C_4	$\min\ \frac{\beta}{C_{(a)}(1+\delta_A)}, \frac{\min\{\mu/\delta_A, C_3 D_2/\beta\}}{4(1+(1+\delta_A)C_M)(\mu+L))}, \frac{1}{2}$
C_5	$\min\ \frac{\beta}{2C_{(a)}}, \frac{C_3 D_2}{32C_{(a)}\beta}, C_4, \frac{C_2}{4}$
δ_{GI}	$\min\ \frac{\delta_A}{4}, \frac{D_1}{C_M}$
δ_{ND}	$\min\ \delta_{GI}, \frac{D_2}{D_1 L}\sqrt{\frac{}{128}}$
δ_{Grid}	$\min\ \frac{\delta_A}{2}, \frac{1}{C_M\ (D_1^{-1}+1)}, \frac{\mu}{8(C_{(a)}+\beta)}$

4 Verifying the Gradient Inequality (1.5) Under Assumption A

In this section, we establish the gradient inequality (1.5) for functions satisfying Assumption A. Throughout the section, we assume that Assumption A is in force. We also use the notation set out in Proposition 3.5.

We present the formal statement and the gradient inequality (1.5) in Theorem 4.3, which appears at the end of this section. The proof is a consequence of the two lemmata. In the first lemma, we prove a constant-sized lower bound for $\mathrm{dist}(0, \partial_\sigma f(x))$, whenever σ is sufficiently small. The proof of this bound relies on the active manifold assumption (A2) and the aiming inequality (3.6). A consequence of the argument is that all elements of $\partial_\sigma f(x)$ are correlated with the normal direction $x - P_M(x) \in N_M(P_M(x))$. Later in Proposition 5.1 we will also show that Algorithm 1 (`NDescent`) terminates rapidly when σ is in the region, motivating the name Normal Descent. We now turn to the lemma.

Lemma 4.1 (Lower bound on Goldstein subgradients I) *Define*

$$D_1 := \frac{\mu}{8(\mu + L)}; \qquad D_2 := \frac{\mu}{2}; \qquad \text{and} \qquad \delta_{GI} := \min\ \frac{\delta_A}{4}, \frac{D_1}{C_M}\ .$$

Then for all $x \in B_{\delta_{GI}}(\bar{x})$ *and* $0 < \sigma \le D_1 \mathrm{dist}(x, M)$, *we have*

$$\mathrm{dist}(0, \partial_\sigma f(x)) \ge D_2.$$

Proof We begin with some preliminary bounds. Fix $x \in B_{\delta_{GI}}(\bar{x})$ and $\sigma > 0$ satisfying the lemma assumptions. We observe that

$$\sigma \le D_1 \mathrm{dist}(x, M) \le \mathrm{dist}(x, M) \le \|x - \bar{x}\| \le \delta_{GI},$$

where the second inequality follows since $D_1 \le 1$ and the third follows since $\bar{x} \in M$. Consequently,

$$\begin{aligned} LC_M(\sigma^2 + \operatorname{dist}^2(x, M)) &\le \delta_{GI} LC_M(\sigma + \operatorname{dist}(x, M)) \\ &\le 2L\delta_{GI} C_M \operatorname{dist}(x, M) \\ &\le 2L D_1 \operatorname{dist}(x, M), \end{aligned} \tag{4.1}$$

where the first inequality follows from the bound $\max\{\sigma, \operatorname{dist}(x, M)\} \le \delta_{GI}$ and the second follows from the bound $\sigma \le \operatorname{dist}(x, M)$. We now turn to the proof.

Now, let $x \in \overline{B}\sigma(x) \subseteq B\delta_A(\bar{x})$ and observe that by aiming condition (3.6),

$$\langle v, x - P_M(x)\rangle \ge \mu \operatorname{dist}(x, M) \qquad \text{for all } v \in \partial f(x).$$

We claim that $\langle v, x - P_M(x)\rangle \ge D_2 \operatorname{dist}(x, M)$ for all $v \in \partial f(x)$. Indeed, for all $v \in \partial f(x)$ we may upper bound the inner product as follows:

$$\begin{aligned} &\langle v, x - P_M(x)\rangle \\ &\le \langle v, x - P_M(x)\rangle + \langle v, (x - P_M(x)) - (x + P_M(x))\rangle \\ &\le \langle v, x - P_M(x)\rangle + L\langle (I - P_{T_M(P_M(x))})(x - x)\rangle + LC_M(\sigma^2 + \operatorname{dist}^2(x, M)) \\ &\le \langle v, x - P_M(x)\rangle + 3L D_1 \operatorname{dist}(x, M), \end{aligned}$$

where the second inequality follows from the bound $\|v\| \le L$ and Item 2 of Proposition 3.5; and the third inequality follows from $\|x - x\| \le \sigma \le D_1 \operatorname{dist}(x, M)$ and (4.1). Consequently, for all $v \in \partial f(x)$, we have

$$\begin{aligned} \langle v, x - P_M(x)\rangle &\ge \mu \operatorname{dist}(x, M) - 3L D_1 \operatorname{dist}(x, M) \\ &\ge \mu \operatorname{dist}(x, M) - \mu\sigma - 3L D_1 \operatorname{dist}(x, M) \\ &\ge \mu(1 - D_1(1 + 3L/\mu)) \operatorname{dist}(x, M) \\ &= D_2 \operatorname{dist}(x, M), \end{aligned} \tag{4.2}$$

where the second inequality follows from 1-Lipschitz continuity of $\operatorname{dist}(\cdot, M)$; and the final inequality follows from the bound $D_1 \le \frac{1}{2(1+3L/\mu)}$. This proves the claim.

Now, fix $g \in \partial_\sigma f(x)$. By definition of $\partial_\sigma f(x)$, there exists a family of coefficients $\lambda_i \in [0, 1]$, points $x_i \in \overline{B}\sigma(x) \subseteq B\delta_A(\bar{x})$, and subgradients $g_i \in \partial f(x_i)$ indexed by a finite set $i \in I$ such that $\sum_{i \in I} \lambda_i = 1$ and $g = \sum_{i \in I} \lambda_i g_i$. Thus, by (4.2), we have

$$\langle g, x - P_M(x)\rangle = \sum_{i \in I} \lambda_i \langle g_i, x - P_M(x)\rangle \ge D_2 \operatorname{dist}(x, M).$$

Therefore, we have

$$g \ge \frac{\langle g, x - P_M(x)\rangle}{\operatorname{dist}(x, M)} \ge D_2,$$

as desired.

In the second lemma, we provide a lower bound for $\operatorname{dist}(0, \partial_\sigma f(x))$ on the order of $\|P_M(x) - \bar{x}\|$, provided $\sigma = O(\|P_M(x) - \bar{x}\|)$. The proof of this bound relies on quadratic growth (A1) and strong (a)-regularity (A3). A consequence of the argument is that the minimal norm element of $\partial_\sigma f(x)$ is close to the tangent vector $\nabla_M f(P_M(x)) \in T_M(P_M(x))$. Later in Proposition 5.6 we will also show that Algorithm 2 (`TDescent`) terminates rapidly when σ is in the region, motivating the name Tangent Descent. We now turn to the lemma.

Lemma 4.2 (Lower bound on Goldstein subgradients II) *Define*

$$C_1 := \frac{\gamma}{4}; \qquad \text{and} \qquad C_2 := \min\left\{\frac{\gamma}{8C_{(a)}}, \frac{\min\{1, 1/\delta_A\}}{2}\right\}.$$

Then for all $x \in B_{\delta_A}(\bar{x})$ *and* $\sigma \geq 0$ *satisfying*

$$\max\{\operatorname{dist}(x, M), \sigma\} \leq C_2 \|P_M(x) - \bar{x}\|,$$

we have

$$\|P_{T_M(P_M(x))}(g)\| \geq C_1 \|P_M(x) - \bar{x}\| \qquad \text{for all } g \in \partial_\sigma f(x).$$

Proof For the purposes of this proof, the term $1/\delta_A$ in the definition of C_2 is unnecessary; however, it will be crucial in the proof of Theorem 4.3. Turning to the proof, fix $x \in B_{\delta_A}(\bar{x})$ and $\sigma \geq 0$ satisfying the lemma assumptions. Define $y = P_M(x)$. Note that

$$\sigma \leq C_2 \|y - \bar{x}\| \leq 2C_2 \|x - \bar{x}\| \leq \delta_A.$$

Thus, by (3.3), for all $g \in \partial_\sigma f(x)$, we have

$$\|P_{T_M(y)}(g - \nabla_M f(y))\| \leq C_{(a)}(\operatorname{dist}(x, M) + \sigma) \leq \frac{\gamma}{4}\|y - \bar{x}\|.$$

In addition, by (3.2), we have $\|\nabla_M f(y)\| \geq \frac{\gamma}{2}\|y - \bar{x}\|$. Therefore, for all $g \in \partial_\sigma f(x)$, we have

$$\|P_{T_M(y)}(g)\| \geq \|\nabla_M f(y)\| - C_{(a)}(\operatorname{dist}(x, M) + \sigma) \geq \frac{\gamma}{4}\|y - \bar{x}\|,$$

as desired.

Given these lemmata, we are now ready to establish the gradient inequality (1.5). The following theorem verifies the bound

$$\sigma \operatorname{dist}(0, \partial_\sigma f(x)) \geq \eta(f(x) - f(\bar{x})),$$

for some $\eta > 0$ provided x is sufficiently near $\bar{x}$ and (x, σ) lies within one of two regions, described in Item 1 and Item 2 of Theorem 4.3. Item 1 and Item 2 roughly correspond to the regions considered in Lemma 4.1 and Lemma 4.2, respectively. Comparing with the statement of the gradient inequality (1.5), we see that gradient inequality of Theorem 4.3 does not require knowledge of an explicit function $\sigma(x)$. Instead, we need only find some σ proportional to $D_1 \mathrm{dist}(x, M)$ or $C_2 \|P_M(x) - \bar{x}\|$ up to a factor of, say, 2. Later in Proposition 6.1 we show that this flexibility allows us to find an appropriate σ through the `linesearch` procedure.

Theorem 4.3 (Gradient inequality) *Suppose that function f satisfies Assumption A at $\bar{x} \in \mathbb{R}^d$. For any constants $a_1 \in (0, D_1]$ and $a_2 \in (0, C_2]$, we have*

$$\sigma \mathrm{dist}(0, \partial_\sigma f(x)) \geq \min\left\{ \frac{\gamma a_2}{8 \max\{4La_2^2, \beta\}}, \frac{\mu a_1}{4 \max\{2L, \beta/a_2^2\}} \right\} (f(x) - f(\bar{x})),$$

whenever $x \in B_{\delta_{\mathrm{GI}}}(\bar{x})$ and $\sigma > 0$ satisfy Item 1 or Item 2:

1. (a) $\frac{a_1}{2}\mathrm{dist}(x, M) \leq \sigma \leq a_1 \mathrm{dist}(x, M)$;
   (b) $a_2^2 \|P_M(x) - \bar{x}\|^2 \leq \mathrm{dist}(x, M)$.
2. (a) $\frac{a_2^2}{2}\|P_M(x) - \bar{x}\| \leq \sigma \leq a_2 \|P_M(x) - \bar{x}\|$;
   (b) $\frac{\mathrm{dist}(x, M)}{\sigma} \leq 2a_2 \|P_M(x) - \bar{x}\|$.

Moreover, for any $x \in B_{\delta_{\mathrm{GI}}}(\bar{x}) \backslash \{\bar{x}\}$, there exists $\sigma > 0$ such that Item 1 or Item 2 is satisfied.

Proof We first show that for any $x \in B_{\delta_{\mathrm{GI}}}(\bar{x}) \backslash \{\bar{x}\}$, there exists $\sigma > 0$ such that either Item 1 or Item 2 is satisfied. We consider two cases. First, suppose $x \in M$. Then Item 2 is trivially satisfied for $\sigma = a_2 \|P_M(x) - \bar{x}\|$. Second, suppose $x \notin M$ and Item 1 cannot be satisfied for any $\sigma > 0$. In this case, we have

$$\mathrm{dist}(x, M) \leq a_2^2 \|P_M(x) - \bar{x}\|^2 = 2a_2\sigma \|P_M(x) - \bar{x}\| \quad \text{with } \sigma := a_2 \|P_M(x) - \bar{x}\|/2.$$

Thus, Item 2 is satisfied.

Now we prove the gradient inequality is satisfied whenever σ satisfies Item 1 or Item 2. Let us suppose that Item 1 holds for some $x \in B_{\delta_{\mathrm{GI}}}(\bar{x})$ and $\sigma > 0$. From (3.7), we have the bound:

$$\begin{aligned} \frac{1}{\max\{2L, \beta/a_2^2\}}(f(x) - f(\bar{x})) &\leq \frac{1}{\max\{2L, \beta/a_2^2\}} L\mathrm{dist}(x, M) + \frac{\beta}{2}\|P_M(x) - \bar{x}\|^2 \\ &\leq \frac{1}{2}\mathrm{dist}(x, M) + a_2^2 \|P_M(x) - \bar{x}\|^2 \\ &\leq \mathrm{dist}(x, M). \end{aligned}$$

Now observe that the assumptions of Lemma 4.1 are satisfied since $x \in B_{\delta_{\mathrm{GI}}}(\bar{x})$, $a_1 \leq D_1$, and x and σ satisfy Item 1. Therefore, we have

$$\sigma \operatorname{dist}(0, \partial_\sigma f(x)) \geq \sigma D_2 \geq \frac{\mu a_1}{4} \operatorname{dist}(x, M) \geq \frac{\mu a_1}{4 \max\{2L, \beta/a_2^2\}} (f(x) - f(\bar{x})),$$

as desired.

Next, let us suppose that Item 2 holds for some $x \in B_{\delta_{\mathrm{GI}}}(\bar{x})$ and $\sigma > 0$. From (3.7), we have the bound:

$$\begin{aligned}
\frac{1}{\max\{4La_2^2, \beta\}} (f(x) - f(\bar{x})) &\leq \frac{1}{\max\{4La_2^2, \beta\}} L \operatorname{dist}(x, M) + \frac{\beta}{2} \|P_M(x) - \bar{x}\|^2 \\
&\leq \frac{1}{2} \frac{\operatorname{dist}(x, M)}{2a_2^2} + \|P_M(x) - \bar{x}\|^2 \\
&\leq \frac{1}{2} \frac{\operatorname{dist}(x, M) \|P_M(x) - \bar{x}\|}{2a_2\sigma} + \|P_M(x) - \bar{x}\|^2 \\
&\leq \|P_M(x) - \bar{x}\|^2.
\end{aligned}$$

Now observe that since $a_2 \leq C_2$ and x and σ satisfy Item 2, we have

$$\sigma \leq C_2 \|P_M(x) - \bar{x}\| \leq 2C_2\delta_{\mathrm{GI}} \leq (1/\delta_{\mathrm{A}})(\delta_{\mathrm{A}}/4) \leq 1,$$

where we use the bound $C_2 \leq 1/2\delta_{\mathrm{A}}$. Consequently, we have

$$\operatorname{dist}(x, M) \leq 2\sigma C_2 \|P_M(x) - \bar{x}\| \leq C_2 \|P_M(x) - \bar{x}\|.$$

Therefore, $\max\{\operatorname{dist}(x, M), \sigma\} \leq C_2 \|P_M(x) - \bar{x}\|$, so the conditions of Lemma 4.2 are satisfied (recall $\delta_{\mathrm{GI}} \leq \delta_{\mathrm{A}}$). Thus, let g denote the minimal norm element of $\partial_\sigma f(x)$ and let us apply Lemma 4.2:

$$\operatorname{dist}(0, \partial_\sigma f(x)) = \|g\| \geq \|P_{T_M(P_M(x))}(g)\| \geq \frac{\gamma}{4} \|P_M(x) - \bar{x}\|.$$

Consequently, we have

$$\sigma \operatorname{dist}(0, \partial_\sigma f(x)) \geq \frac{\sigma\gamma}{4} \|P_M(x) - \bar{x}\| \geq \frac{\gamma a_2}{8 \max\{4La_2^2, \beta\}} (f(x) - f(\bar{x})),$$

where the last inequality follows from $\sigma \geq \frac{a_2}{2} \|P_M(x) - \bar{x}\|$. This completes the proof.

Remark 1 Note that $a_1, a_2 \in (0, 1)$ as claimed in Sect. 1.3.1, where we introduced the *normal and tangent regions* appearing in the statement of Theorem 4.3.

This concludes the proof of the gradient inequality (1.5) under Assumption A. In Sect. 6, we will use the gradient inequality to establish rapid local convergence of `NTDescent`. Before proving that, the following section analyzes `TDescent` and `NDescent` methods.

5 Rapid Termination of `NDescent` and `TDescent` Under Assumption A

In this section, we analyze the `NDescent` and `TDescent` methods, showing that both methods rapidly terminate with descent in appropriate regions. Throughout the section, we assume that Assumption A is in force. We also use the results and notation of Proposition 3.5, Table 1, Lemma 4.1, and Lemma 4.2.

The main results of this section are Propositions 4.1 and 5.6, which analyze `NDescent` and `TDescent`, respectively. Proposition 5.1 shows that `NDescent` terminates with descent in a constant number of iterations within the region considered in Item 1 of Theorem 4.3. Proposition 5.6 shows that `TDescent` either terminates with descent in $O(\log^{-1}(f(x) - f(\bar{x})))$ iterations or $f(x) - f(\bar{x})$ is already exponentially small in T within the region considered in Item 2 of Theorem 4.3. These lemmata will be the basis of our main convergence theorem—Theorem 6.3—appearing in Sect. 6.

5.1 Analysis of `NDescent`

The following proposition shows that `NDescent` locally terminates in finitely many iterations whenever σ is sufficiently small. The result is a simple consequence of Lemmas 2.3 and 4.1.

Proposition 5.1 (`NDescent` loop terminates with descent) *Define a radius*

$$\delta_{\mathrm{ND}} := \min\left\{\delta_{\mathrm{GI}}, \frac{D_2}{D_1 L\sqrt{128}}\right\}.$$

Then for all $x \in B_{\delta_{\mathrm{ND}}}(\bar{x})$, *radii* $\sigma > 0$ *with* $\sigma \le D_1 \operatorname{dist}(x, M)$, *subgradients* $g \in \partial_\sigma f(x)$, *failure probabilities* $p \in (0, 1)$ *and budgets* $T > 0$ *satisfying*

$$T \ge \frac{64L^2}{D_2^2} 2\log(1/p),$$

the point $x_+ := \texttt{NDescent}(x, g, \sigma, T)$ *satisfies*

$$f(x_+) \le f(x) - \frac{\sigma \operatorname{dist}(0, \partial_\sigma f(x))}{8} \quad \text{with probability at least } 1 - p.$$

Proof Fix $x \in B_{\delta_{\mathrm{ND}}}(\bar{x})$ and $\sigma > 0$ satisfying the lemma assumptions. Observe that

$$\sigma \le D_1 \operatorname{dist}(x, M) \le D_1 \delta_{\mathrm{ND}} \le \min\left\{\delta_{\mathrm{GI}}, \frac{D_2}{L\sqrt{128}}\right\},$$

where the final inequality follows from the bound $D_1 \leq 1$; see Lemma E.1. Thus, by Lemma 4.1, we have $\operatorname{dist}(0, \partial_\sigma f(x)) \geq D_2$ (recall $\delta_{\mathrm{ND}} \leq \delta_{\mathrm{GI}}$). Consequently,

$$\sigma \leq \frac{D_2}{L\sqrt{128}} \leq \frac{\operatorname{dist}(0, \partial_\sigma f(x))}{L\sqrt{128}}.$$

Therefore, σ and T satisfy the assumptions of Lemma 2.3. Hence, the desired descent condition is guaranteed with probability at least $1 - p$.

We now turn to the analysis of the `TDescent` step.

5.2 Analysis of `TDescent`

In this section, we analyze `TDescent`, proving two main results. First, we prove Proposition 5.6, which shows that `TDescent` terminates rapidly. Second, in Lemma 5.8 we show that the trust region constraint in Line 7 of Algorithm 3 (`linesearch`) prevents long steps. Thus, once the method enters a sufficiently small neighborhood of $\bar{x}$, it cannot leave.

We begin with descent Proposition 5.6, which relies on four technical lemmata that analyze the structure of Goldstein subgradients when σ is sufficiently small and x is sufficiently near $\bar{x}$: Lemma 5.2 states that elements of Goldstein subdifferential with small normal components are descent directions. Lemmas 5.3 and 5.4 show that normalized subgradient steps approximately reflect points across the active manifold. Lemma 5.5 uses the approximate reflection property to show that `TDescent` geometrically decreases the normal component of the input subgradient, ensuring that we rapidly find a descent direction. We now turn to the Lemmata.

5.2.1 Descent with Small Normal Part

The first lemma shows that Goldstein subgradients with small normal components are descent directions.

Lemma 5.2 (Descent with small normal part) *Define*

$$C_3 := \frac{C_1^2}{8L}.$$

Then for all $x \in B_{\delta_A}(\bar{x})$, $\sigma > 0$, *and* $g \in \partial_\sigma f(x) \setminus \{0\}$ *satisfying*

1. $\max\{\operatorname{dist}(x, M), \sigma\} \leq \frac{C_2}{4} \|P_M(x) - \bar{x}\|$;
2. $\|P_{N_M(P_M(x))}(g)\| \leq C_3 \|P_M(x) - \bar{x}\|^2$,

we have

$$f\left(x - \sigma \frac{g}{\|g\|}\right) \leq f(x) - \frac{\sigma \|g\|}{8}.$$

Proof We begin with preliminary notation and bounds. We fix $x \in B_{\delta_A}(\bar{x})$ and subgradient $g \in \partial_\sigma f(x) \backslash \{0\}$. We define $y := P_M(x)$, $T := T_M(y)$, and $N := N_M(y)$. We observe that

$$\sigma \leq \frac{C_2}{4}\|y - \bar{x}\| \leq \frac{C_2}{2}\|x - \bar{x}\| \leq C_2\delta_A \leq \delta_A,$$

where the final inequality follows since $C_2 \leq 1$; see Lemma 4.2. We now turn to the proof.

The starting point of the proof is Lebourg's mean value Theorem [16, Theorem 2.4], which ensures that there exists $v \in \partial_\sigma f(x)$ such that

$$f\left(x - \sigma\frac{g}{\|g\|}\right) - f(x) = \left\langle v, -\sigma\frac{g}{\|g\|}\right\rangle = -\frac{\sigma}{\|g\|}\langle v, P_T(g)\rangle - \frac{\sigma}{\|g\|}\langle v, P_N(g)\rangle.$$

In what follows, we will show that the first term satisfies $\langle v, P_T(g)\rangle \geq \frac{3}{8}\|g\|^2$, while the second term satisfies $|\langle v, P_N(g)\rangle| \leq \frac{1}{8}\|g\|^2$, yielding the result.

Indeed, beginning with $|\langle v, P_N(g)\rangle|$, we note that

$$\|P_N(g)\| \leq C_3\|P_M(x) - \bar{x}\|^2 \leq \frac{C_3}{C_1^2}\|g\|^2 = \frac{1}{8L}\|g\|^2, \tag{5.1}$$

where the second inequality follow from Lemma 4.2. Consequently, we have the bound $|\langle v, P_N(g)\rangle| \leq L\|P_N(g)\| \leq \frac{1}{8}\|g\|^2$, where we first inequality relies on the estimate $\|v\| \leq L$; see Item 6 of Proposition 3.5.

Next, we prove a lower bound on $\langle v, P_T(g)\rangle$. Since $v \in \partial_\sigma f(x)$,

$$\|P_T(v - g)\| \leq 2C_{(a)}(\operatorname{dist}(x, M) + \sigma) \leq C_2C_{(a)}\|P_M(x) - \bar{x}\| \leq \frac{C_2C_{(a)}}{C_1}\|g\| \leq \frac{1}{2}\|g\|.$$

where the first inequality follows from (3.5); the second by assumption; the third follows from Lemma 4.2; and the fourth follows from the bound $\sigma \leq \frac{C_1}{2C_{(a)}}$. Therefore,

$$\|P_T(v)\| - \|g\| \leq \|P_T(v - g)\| + \|P_N(g)\| \leq \frac{1}{2}\|g\| + \frac{1}{8L}\|g\|^2 \leq \frac{5}{8}\|g\|,$$

where the second inequality follows from (5.1) and the third follows from the bound $\|g\| \leq L$. Consequently, we have the bound

$$\langle v, P_T(g)\rangle = \|P_T(v)\|\|g\| \geq \|g\|^2 - \|P_T(v) - g\|\|g\| \geq \frac{3}{8}\|g\|^2.$$

This completes the proof.

Note that the proof implies a slightly stronger bound than claimed, namely that we have $f(x - \sigma g/\|g\|) \leq f(x) - \sigma\|g\|/4$. For the sake of maintaining symmetry with Proposition 5.1, however, we use the constant $1/8$ throughout.

5.2.2 The Approximate Reflection Property

The next two lemmata prove the approximate reflection property that was described in the introduction. The lemmas roughly show that normalized subgradient steps approximately "flip the sign" of the normal component of the subgradient nearby the manifold; see Sect. 1.3.2 for more intuition. The first lemma proves the approximate reflection property up to a tolerance depending on the distance to the manifold and σ. This lemma will be used again in the proofs of Lemma 5.4 and Lemma 5.7.

Lemma 5.3 (Approximate reflection inequality, general case) *For all* $x \in B_{\delta_A/2}(\bar{x})$, $\sigma \in (0, \delta_A/2]$, $g \in \partial_\sigma f(x) \setminus \{0\}$ *and* $\hat{g} \in \partial f\left(x - \sigma \frac{g}{g}\right)$, *we have*

$$\begin{aligned} &\langle P_{N_M(P_M(x))}(\hat{g}), g\rangle \\ &\le -\mu \left\| P_{N_M(P_M(x))} g \right\| + \frac{(\mu + L)\, g\, \operatorname{dist}(x, M)}{\sigma} \\ &\quad + \frac{(\mu + L)\, g\, C_M\left(\operatorname{dist}^2(x, M) + \sigma^2\right)}{\sigma}. \end{aligned} \tag{5.2}$$

Proof We begin with preliminary notation and bounds. We fix $x \in B_{\delta_A/2}(\bar{x})$ and subgradient $g \in \partial_\sigma f(x) \setminus \{0\}$. We define $y := P_M(x)$, $T := T_M(y)$, and $N := N_M(y)$. Finally, define $u := \frac{g}{g}$. Note that since $x \in B_{\delta_A/2}(\bar{x})$ and $\sigma \le \delta_A/2$, we have $x - \sigma u \in B_{\delta_A}(\bar{x})$.

Therefore, by the aiming inequality (3.6), we have

$$\underbrace{\left\langle \hat{g}, x - \sigma u - P_M(x - \sigma u)\right\rangle}_{=: A} \ge \mu \underbrace{\left\| x - \sigma u - P_M(x - \sigma u)\right\|}_{=: B}$$

We aim to simplify this inequality with (3.1). To that end, first note that

$$\begin{aligned} &(x - \sigma u - P_M(x - \sigma u)) - (x - P_M(x) - \sigma P_N(u)) \\ &= P_M(x - \sigma u) - P_M(x) + \sigma P_T(u) \\ &\le C_M\left(\operatorname{dist}^2(x, M) + \sigma^2\right). \end{aligned} \tag{5.3}$$

Consequently, we have

$$A \ge B \ge \mu \left\| x - P_M(x) - \sigma P_N(u)\right\| - \mu C_M\left(\operatorname{dist}^2(x, M) + \sigma^2\right) \ge \sigma\mu \left\| P_N(u)\right\| - \mu S,$$

where $S := \operatorname{dist}(x, M) + C_M\left(\operatorname{dist}^2(x, M) + \sigma^2\right)$. In addition, by (5.3) we have

$$\langle \hat{g}, (x - \sigma u - P_M(x - \sigma u)) + \sigma P_N u\rangle \le L S.$$

Therefore, we have

$$\langle \hat{g}, \sigma P_N(u)\rangle = -A + \langle \hat{g}, (x - \sigma u - P_M(x - \sigma u)) + \sigma P_N u\rangle$$

$$\leq -\sigma\mu \quad P_N(u) + (\mu + \quad L)S. \tag{5.4}$$

Inequality (5.2) then follows by multiplying both sides of inequality (5.4) by g/σ .

The second lemma is an application of Lemma 5.3 nearby the manifold.

Lemma 5.4 (Approximate reflection inequality near the manifold) *Define*

$$C_4 := \min \quad \frac{\beta}{C_{(a)}(1+\delta_A)}, \frac{\min\{\mu/\delta_A, C_3D_2/\beta\}}{4(1+(1+\delta_A)C_M)(\mu + \quad L)}, \frac{1}{2} \quad .$$

Then for all $x \in B\delta_{A/2}(\bar{x})$, $\sigma > 0$, *and* $g \in \partial\sigma f(x)\backslash\{0\}$ *satisfying*

$$\max \quad \frac{\text{dist}(x, M)}{\sigma}, \sigma \quad \leq C_4 \quad P_M(x) - \bar{x},$$

we have

$$P_{N_M(P_M(x))}(\hat{g}), g \leq - D_2 \quad P_{N_M(P_M(x))}g + \quad \frac{C_3D_2}{2} \quad P_M(x) - \bar{x}^2$$
$$\text{for all } \hat{g} \in \partial f \quad x - \sigma\frac{g}{g} \quad .$$

Proof We begin with preliminary notation and bounds. We fix $x \in B\delta_{A/2}(\bar{x})$ and subgradient $g \in \partial_\sigma f(x)\backslash\{0\}$. We define $y := P_M(x)$, $T := T_M(y)$, and $N := N_M(y)$. We observe that

$$\sigma \leq C_4 \quad y - \bar{x} \leq \quad 2C_4 \quad x - \bar{x} \leq \quad C_4\delta_A \leq \delta_A/2.$$

Finally, we have

$$S := \text{dist}(x, M) + \quad C_M(\text{dist}^2(x, M) + \sigma^2) \leq \sigma C_4(1 + C_M(1+\delta_A)) \quad y - \bar{x}. \tag{5.5}$$

where the inequality follows from the bound $\text{dist}(x, M) \leq \quad x - \bar{x} \leq \delta_A$.

We now apply inequality (5.2):

$$\begin{aligned} P_N\hat{g}, g &\leq -\mu \quad P_N g + \quad \frac{(\mu + L) \quad g \quad S}{\sigma} \\ &\leq -\mu \quad P_N g + (1 + (1+\delta_A)C_M)(\mu + L)C_4 \quad g \quad y - \bar{x} \\ &\leq -\mu \quad P_N g + (1 + (1+\delta_A)C_M)(\mu + L)C_4(P_T(g) + \quad P_N(g)) \quad y - \bar{x} \\ &\leq -\frac{\mu}{2} \quad P_N g + \quad \frac{C_3D_2}{4\beta} \quad P_T(g) \quad y - \bar{x}, \end{aligned}$$

where the second inequality follows from (5.5); the third inequality follows from triangle inequality; and the fourth inequality follows from the bound

$$(1 + (1 + \delta_A)C_M)(\mu + L)C_4 \|y - \bar{x}\| \le \frac{\mu/\delta_A}{4} \cdot (2\|x - \bar{x}\|) \le \frac{\mu/\delta_A}{4}\delta_A \le \mu/2.$$

The proof will be complete if we can show that

$$P_T(g) \le 2\beta \|y - \bar{x}\|.$$

To that end, we have

$$\begin{aligned} P_T(g) &\le C_{(a)}(\operatorname{dist}(x, M) + \sigma) + \beta \|y - \bar{x}\| \\ &\le (C_4 C_{(a)}(1 + \delta_A) + \beta) \|y - \bar{x}\| \le 2\beta \|y - \bar{x}\|, \end{aligned}$$

where the first inequality follows from (3.4); the second inequality follows from the lemma assumptions and the bound $\operatorname{dist}(x, M) \le C_4\sigma \|y - \bar{x}\| \le C_4\delta_A \|y - \bar{x}\|$; and the third inequality follows from the bounds on C_4. This completes the proof.

5.2.3 The Normal Component Shrinks Geometrically

The following lemma shows that every step of `TDescent` geometrically shrinks the normal component of the subgradient, up to a tolerance of $O(\|P_M(x) - \bar{x}\|^2)$.

Lemma 5.5 (Normal component shrinks geometrically) *Define*

$$C_5 := \min\left\{\frac{\beta}{2C_{(a)}}, \frac{C_3 D_2}{32 C_{(a)}\beta}, C_4, \frac{C_2}{4}\right\}.$$

Then for all $x \in B_{\delta_A/2}(\bar{x})$, $\sigma > 0$, $g \in \partial_\sigma f(x) \setminus \{0\}$, *and* $\hat{g} \in \partial f(x - \sigma \frac{g}{g}) \setminus \{0\}$ *satisfying*

1. $\|P_{N_M(P_M(x))} g\| \ge C_3 \|P_M(x) - \bar{x}\|^2$;
2. $\max\left\{\frac{\operatorname{dist}(x, M)}{\sigma}, \sigma\right\} \le C_5 \|P_M(x) - \bar{x}\|$,

the vector $g = \operatorname{argmin}_{h \in [g, \hat{g}]} \|h\|$ *satisfies:*

$$\|P_{N_M(P_M(x))}(g)\|^2 \le \left(1 - \frac{3D_2^2}{64L^2}\right) \|P_{N_M(P_M(x))} g\|^2.$$

Proof We begin with preliminary notation and bounds. We fix $x \in B_{\delta_A/2}(\bar{x})$ and subgradient $g \in \partial_\sigma f(x) \setminus \{0\}$. We define $y := P_M(x)$, $T := T_M(y)$, and $N := N_M(y)$. We observe two bounds. First, we have

$$\sigma \le C_5 \|y - \bar{x}\| \le 2C_5 \|x - \bar{x}\| \le C_5\delta_A \le 1.$$

where the final inequality follows since $C_5 \le C_2/4 \le 1/(8\delta_A)$. Second, we have

$$\operatorname{dist}(x, M) \le C_5\sigma\ \|y - \bar{x}\| \le C_5\ \|y - \bar{x}\|, \tag{5.6}$$

since $\sigma \le 1$. We now turn to the proof.

Consider the optimal weight $\lambda := \operatorname{argmin}_{\lambda\in[0,1]}\ \|g + \lambda(\hat{g} - g)\|$. By definition we have $g = g + \lambda\ (\hat{g} - g)$. Moreover, a quick calculation shows that

$$\lambda = \max\ \min\ -\frac{\langle g, \hat{g} - g\rangle}{\|\hat{g} - g\|^2}, 1\ , 0\ .$$

We claim that the following bound holds on λ:

$$\underbrace{-\frac{\langle P_N(g), \hat{g} - g\rangle}{8L^2}}_{=:\lambda_1} \le \lambda \le \underbrace{-\frac{3\ \langle P_N(g), \hat{g} - g\rangle}{2\ \|P_N(\hat{g} - g)\|^2}}_{=:\lambda_2}. \tag{5.7}$$

Note that (5.7) is an immediate consequence of the following bound:

$$0 \le -\frac{1}{2}\ \langle P_N(g), \hat{g} - g\rangle \le -\ \langle g, \hat{g} - g\rangle \le -\frac{3}{2}\ \langle P_N(g), \hat{g} - g\rangle\ . \tag{5.8}$$

Indeed, if (5.8) holds, then $\lambda = \min\ -\frac{\langle g, \hat{g} - g\rangle}{\|\hat{g} - g\|^2}, 1$. Thus, we obtain the upper bound

$$\lambda \le -\frac{\langle g, \hat{g} - g\rangle}{\|\hat{g} - g\|^2} \le -\frac{3}{2}\frac{\langle P_N(g), \hat{g} - g\rangle}{\|g - \hat{g}\|^2} \le -\frac{3}{2}\frac{\langle P_N(g), \hat{g} - g\rangle}{\|P_N(g - \hat{g})\|^2} = \lambda_2.$$

Likewise, we obtain the lower bound

$$\begin{aligned}\lambda &= \min\ -\frac{\langle g, \hat{g} - g\rangle}{\|\hat{g} - g\|^2}, 1\ \ge \min\ -\frac{\langle g, \hat{g} - g\rangle}{4L^2}, 1\\ &= -\frac{\langle g, \hat{g} - g\rangle}{4L^2} \ge -\frac{\langle P_N(g), \hat{g} - g\rangle}{8L^2} = \lambda_1,\end{aligned}$$

where the first inequality follows from the bound $\|\hat{g} - g\|^2 \le 2(\|\hat{g}\|^2 + \|g\|^2) \le 4L^2$; and the second equality follows from the bound $|\langle g, \hat{g} - g\rangle| \le \|g\|\,\|\hat{g} - g\| \le 2L^2$. Thus, we now prove (5.8).

To that end, note that (5.8) is equivalent to the following bound:

$$\left|\langle P_T(g), \hat{g} - g\rangle\right| \le \frac{-\langle P_N(g), \hat{g} - g\rangle}{2}. \tag{5.9}$$

Therefore, we first bound $\left|\langle P_T(g), \hat{g} - g\rangle\right|$:

$$\left|\langle P_T(g), \hat{g} - g\rangle\right| \le \|P_T(g)\|\ \|P_T(\hat{g} - g)\|$$

$$\leq 2C_{(a)}(\mathrm{dist}(x, M) + \sigma)(C_{(a)}(\mathrm{dist}(x, M) + \sigma) + \beta \quad y - \bar{x})$$
$$\leq 4C_{(a)}C_5(2C_{(a)}C_5 + \beta) \quad y - \bar{x}\ ^2$$
$$\leq \frac{C_3D_2}{4} \quad y - \bar{x}\ ^2,$$

where the second inequality follows from (3.4) and (3.5); the third inequality follows from (5.6) and the bound $\sigma \leq C_5 \quad y - \bar{x}$; and the fourth inequality follows from the definition of C_5. To complete the proof of (5.9), we show that $\frac{C_3D_2}{4} \quad y - \bar{x}\ ^2 \leq -\frac{1}{2} \quad P_N(g), \hat{g} - g$:

$$\begin{aligned}\frac{C_3D_2}{2} \quad y - \bar{x}\ ^2 &\leq D_2 \quad P_N(g) - \frac{C_3D_2}{2} \quad y - \bar{x}\ ^2 \\ &\leq - \quad P_N(g), \hat{g} \\ &\leq - \quad P_N(g), \hat{g} - g\ ,\end{aligned} \tag{5.10}$$

where the first inequality follows from the assumption $\frac{D_2}{2} \quad P_N(g) \geq \frac{C_3D_2}{2} \quad y - \bar{x}\ ^2$; the second inequality follows from Lemma 5.4 (recall $C \leq C_4$ and $x \in B\delta_{A/2}(\bar{x})$); and the third inequality follows from $P_N(g), g = P_N(g)\ ^2 \geq 0$. Thus, the equivalent bounds (5.9) and (5.8) hold. Consequently, Equation (5.7) holds.

Now we turn to the contraction argument. Consider the function $r : \mathbb{R} \to \mathbb{R}$ satisfying

$$r(\lambda) = \quad P_N(g)\ ^2 + 2\lambda \quad P_N(g), \hat{g} - g\ + \lambda^2 \quad P_N(\hat{g} - g)\ ^2 \qquad \text{for all } \lambda \in \mathbb{R}.$$

Observe that

$$P_N(g\ )\ ^2 = \quad P_N(g)\ ^2 + 2\lambda \quad P_N(g), \hat{g} - g\ + (\lambda\ )^2 \quad P_N(\hat{g} - g)\ ^2 = r(\lambda\ ).$$

Therefore, by convexity of r and (5.7), we have

$$P_N(g\ )\ ^2 = r(\lambda\ ) \leq \max_{\lambda \in [\lambda_1, \lambda_2]} r(\lambda) \leq \max\{r(\lambda_1), r(\lambda_2)\}.$$

To complete the proof, we show each term in the "max" is bounded by $1 - \frac{3D_2^2}{64L^2} \quad P_N(g)\ ^2$.

To show this, we will use the following consequence of (5.10):

$$- \quad P_N(g), \hat{g} - g \geq D_2 \quad P_N(g) - \frac{C_3D_2}{2} \quad y - \bar{x}\ ^2 \geq \frac{D_2}{2} \quad P_N(g), \tag{5.11}$$

where the final inequality follows from the assumption $\frac{C_3D_2}{2} \quad y - \bar{x}\ ^2 \leq \frac{D_2}{2} \quad P_N(g)$. Indeed, first observe that

$$r(\lambda_2) = \quad P_N(g)\ ^2 - \frac{3}{4}\frac{P_N(g), \hat{g} - g\ ^2}{P_N(\hat{g} - g)\ ^2}$$

$$\leq \left(1 - \frac{3D_2^2}{16\ P_N(\hat{g} - g)\ ^2}\right) P_N(g)\ ^2$$
$$\leq \left(1 - \frac{3D_2^2}{64L^2}\right) P_N(g)\ ^2,$$

where the first inequality from (5.11) and the second inequality follows from the bound $P_N(\hat{g} - g)\ ^2 \leq \hat{}g - g\ ^2 \leq 4L^2$. Likewise, observe that

$$r(\lambda_1) = P_N(g)\ ^2 - \frac{P_N(g), \hat{g} - g\ ^2}{4L^2} + \frac{P_N(g), \hat{g} - g\ ^2\ P_N(\hat{g} - g)\ ^2}{64L^4}$$
$$\leq P_N(g)\ ^2 - \frac{P_N(g), \hat{g} - g\ ^2}{4L^2} + \frac{P_N(g), \hat{g} - g\ ^2}{16L^2}$$
$$\leq \left(1 - \frac{3D_2^2}{64L^2}\right) P_N(g)\ ^2,$$

where the first inequality follows from the bound $P_N(\hat{g} - g)\ ^2 \leq \hat{}g - g\ ^2 \leq 4L^2$; and the second inequality follows from (5.11). Therefore, the proof is complete.

5.2.4 `TDescent` Terminates with Descent

The following proposition is the main result of this section. It shows that `TDescent` must either terminate with descent or $f(x) - f(\bar{x})$ is already exponentially small in T.

Proposition 5.6 (`TDescent` loop terminates with descent) *Fix* $T \in \mathbb{N}$. *Then for all* $x \in B\delta_A/2(\bar{x})$, $v \in \partial_\sigma f(x)$, *and* $\sigma > 0$ *satisfying*

$$\max\left(\frac{\text{dist}(x, M)}{\sigma}, \sigma\right) \leq C_5\ P_M(x) - \bar{x},$$

at least one of the following holds:

1. *we have*

$$f(x) - f(\bar{x}) \leq \frac{(C_5^2 L + \beta)\ L}{C_3}\left(1 - \frac{3\mu^2}{256L^2}\right)^{T/2};$$

2. *the vector* $g :=$ `TDescent`(x, v, σ, T) *satisfies* $g > 0$ *and*

$$f\left(x - \sigma\frac{g}{g}\right) \leq f(x) - \frac{\sigma\,\text{dist}(0, \partial_\sigma f(x))}{8}.$$

Proof We begin with preliminary notation and bounds. We fix $x \in B_{\delta_A/2}(\bar{x})$ and subgradient $v \in \partial_\sigma f(x)$. We define $y := P_M(x)$, and $N := N_M(y)$. Observe that

$$\sigma \le C_5 \|y - \bar{x}\| \le 2C_5 \|x - \bar{x}\| \le C_5 \delta_A \le 1,$$

where the final inequality follows by definition of $C_5 \le C_2/4 \le 1/(8\delta_A)$. In addition, since $C_5 \le C_2/4$, we have $\sigma \le (C_2/4)\|y - \bar{x}\|$ and

$$\operatorname{dist}(x, M) \le \sigma \|C_5\| y - \bar{x}\| \le \frac{C_2}{4}\|y - \bar{x}\|,$$

where the final inequality follows from $\sigma \le 1$. Consequently,

$$\max\{\sigma, \operatorname{dist}(x, M)\} \le \frac{C_2}{4}\|y - \bar{x}\|. \tag{5.12}$$

We now turn to the proof.

Turning to the proof, note that since $x \in B_{\delta_A/2}(\bar{x})$, Lemma 4.2 and (5.12) ensure that

$$\operatorname{dist}(0, \partial_\sigma f(x)) \ge C_1 \|y - \bar{x}\| > 0.$$

Thus, if $\texttt{TDescent}(x, v, \sigma, T)$ terminates at $t < T$, then Item 2 must hold. For the remainder of the proof, we suppose that $\texttt{TDescent}(x, v, \sigma, T)$ terminates at the final iteration $t = T$ and that Item 2 does not hold. In this case, Lemma 5.2 and (5.12) ensure that the iterates g_t of $\texttt{TDescent}(x, v, \sigma, T)$ satisfy $\|P_N(g_t)\| > C_3\|y - \bar{x}\|^2$ for all $0 \le t \le T - 1$. Therefore, since $x \in B_{\delta_A/2}(\bar{x})$, $\max\{\operatorname{dist}(x, M)/\sigma, \sigma\} \le C_5\|y - \bar{x}\|$, and $\|P_N(g_t)\| > C_3\|y - \bar{x}\|^2$, Lemma 5.5, yields the contraction:

$$\|P_N(g_{t+1})\|^2 \le \left(1 - \frac{3D_2^2}{64L^2}\right)\|P_N(g_t)\|^2, \quad \text{for all } 0 \le t \le T - 1.$$

Unfolding this contraction, we see that g_T is an exponentially small Goldstein subgradient:

$$\|P_N(g_T)\| \le \left(1 - \frac{3D_2^2}{64L^2}\right)^{T/2}\|P_N(g_0)\|.$$

As a result, the projection y is nearby $\bar{x}$:

$$\|y - \bar{x}\|^2 \le \frac{\|P_N(g_T)\|}{C_3} \le \frac{\|P_N(g_0)\|}{C_3}\left(1 - \frac{3D_2^2}{64L^2}\right)^{T/2} \le \frac{L}{C_3}\left(1 - \frac{3D_2^2}{64L^2}\right)^{T/2}. \tag{5.13}$$

Consequently,

$$\begin{aligned} f(x) - f(\bar{x}) &\le L\,\mathrm{dist}(x, M) + \beta \|y - \bar{x}\|^2 \\ &\le (C_5^2 L + \beta) \|y - \bar{x}\|^2 \\ &\le \frac{(C_5^2 L + \beta) L}{C_3} \left(1 - \frac{3D_2^2}{64L^2}\right)^{T/2}, \end{aligned}$$

where the first inequality follows from (3.7) (recall $x \in B_{\delta_A/2}(\bar{x})$); the second inequality follows since $\mathrm{dist}(x, M) \le \sigma \|C_5 (y - \bar{x})\| \le C_5^2 \|y - \bar{x}\|^2$; and the third inequality follows from (5.13). The proof then follows from the identity $D_2 = \frac{\mu}{2}$.

5.2.5 The "Trust Region" Constraint Prevents Long Steps

Before ending this section, we must establish one final technical result for `TDescent`. Namely, in Lemma 5.8, we show that for appropriate σ, `TDescent` eventually generates small subgradients on the order of $O(\|x - \bar{x}\|)$. This property is intuitive because $\mathrm{dist}(0, \partial_\sigma f(x)) = 0$ whenever $\sigma \ge \|x - \bar{x}\|$. This property will help us ensure that the iterates of `NTDescent` (Algorithm 4) cannot leave sufficiently small neighborhoods of $\bar{x}$. Indeed, since the subgradients v_{i+1} generated by Algorithm 3 (`linesearch`) are decreasing in norm, we will show that the trust region constraint $\sigma_i \le \frac{v_{i+1}}{s}$ in Line 7 of Algorithm 3 must eventually be violated for large i. This ensures large σ_i are never chosen.

To prove this claim, we first establish a refinement of the approximate reflection property in Lemma 5.4. Compared to Lemma 5.4, the following lemma deals with a different range of parameters. We place the proof in Appendix D as it follows from a similar line of reasoning as Lemma 5.4.

Lemma 5.7 (Approximate reflection across manifold, large steps) *Define*

$$\delta_{\mathrm{Grid}} := \min\left\{\frac{\delta_A}{2}, \frac{1}{C_M (D_1^{-1} + 1)}, \frac{\mu}{8(C_{(a)} + \beta)}\right\}.$$

Then for all $x \in B_{\delta_{\mathrm{Grid}}}(\bar{x})$, $\sigma > 0$, *and* $g \in \partial_\sigma f(x) \setminus \{0\}$ *satisfying*

$$D_1^{-1}\mathrm{dist}(x, M) \le \sigma \le \delta_{\mathrm{Grid}},$$

we have

$$\left\langle \hat{g}, g \right\rangle \le -D_2 \|g\| + 2D_2 \|P_{T_M(P_M(x))}(g)\| \quad \text{for all } \hat{g} \in \partial f\left(x - \sigma \frac{g}{\|g\|}\right).$$

Finally, we prove that `TDescent` eventually generates small subgradients.

Lemma 5.8 (`TDescent` yields small subgradients) *Fix* $T \in \mathbb{N}$. *Then for all* $x \in B\delta_{\mathrm{Grid}}(\bar{x}), \sigma > 0$, *and* $g \in \partial_\sigma f(x) \backslash \{0\}$ *satisfying*

$$D_1^{-1}\mathrm{dist}(x, M) \le \sigma \le \delta_{\mathrm{Grid}},$$

the vector $g := \mathtt{TDescent}(x, g, \sigma, T)$ *satisfies*

$$g \le \max\left\{ \left(1 - \frac{\mu^2}{64L^2}\right)^{T/2} g,\ 4C_{(a)}\sigma + 4(C_{(a)} + 2\beta)\|x - \bar{x}\|,\ \frac{8(f(x) - f(\bar{x}))}{\sigma} \right\}.$$

Proof We begin with preliminary notation and bounds. We fix $x \in B\delta_{\mathrm{Grid}}(\bar{x})$ and subgradient $g \in \partial_\sigma f(x) \backslash \{0\}$. We define $y := P_M(x)$ and $T := T_M(y)$. We also define $c := C_{(a)}(\mathrm{dist}(x, M) + \sigma) + \beta \|y - \bar{x}\|$. We have the following two bounds: First, we have

$$c \le C_{(a)}(\|x - \bar{x}\| + \sigma) + 2\beta\|x - \bar{x}\| \le C_{(a)}\sigma + (C_{(a)} + 2\beta)\|x - \bar{x}\|. \tag{5.14}$$

Second, by (3.4), we have

$$P_T(v) \le c \quad \text{for all } v \in \partial_\sigma f(x). \tag{5.15}$$

We now turn to the proof.

Note that the result holds automatically if $g = 0$. Thus, we first consider the case where `TDescent` terminates in descent, meaning

$$f(x_+) - f(x) \le -\frac{\sigma\|g\|}{8} \qquad \text{where } x_+ := x - \sigma\frac{g}{\|g\|}.$$

Since $\sigma \le \delta_{\mathrm{Grid}} \le \delta_{\mathrm{A}}/2$ and $x \in B\delta_{\mathrm{A}/2}(\bar{x})$, it follows that $x_+ \in B\delta_{\mathrm{A}}(\bar{x})$. Thus, by Item 1 of Proposition 3.5, we have

$$f(x_+) \ge f(\bar{x}) + \frac{\gamma}{2}\|x - \bar{x}\|^2 \ge f(\bar{x}).$$

Consequently, we have

$$f(\bar{x}) - f(x) \le -\frac{\sigma\|g\|}{8}.$$

Rearranging then gives the upper bound $g \le \frac{8(f(x) - f(\bar{x}))}{\sigma}$, as desired.

Let us now suppose that `TDescent` does not terminate with descent or with $g = 0$. In this case, the iterates $g_0, \ldots, g_T$ of $\mathtt{TDescent}(x, g, \sigma, T)$ exist and satisfy $g_t \in \partial_\sigma f(x)$ for all $t \le T$. We consider two cases.

Case 1. Now suppose $g_t \le 4c$ for some t satisfying $0 \le t \le T$. Since g_t is a decreasing sequence, it follows that $g = g_T \le 4c$. Recalling (5.14), yields the bound

$$g \le 4c \le 4C_{(a)}\sigma + 4(C_{(a)} + 2\beta)\|x - \bar{x}\|,$$

as desired.

Case 2. Next suppose that for all $0 \le t \le T$ we have $4c < g_t$. In this case, Lemma 5.7 shows that for all $t \le T$, we have

$$\langle \hat{g}_t, g_t \rangle \le -\frac{\mu}{2}\|g_t\| + \mu \quad \|P_T g_t\| \le -\frac{\mu}{2}\|g_t\| + \mu \|c\| \le -\frac{\mu}{4}\|g_t\|. \tag{5.16}$$

We now use this bound to prove a one-step geometric improvement bound for $\|g_t\|^2$. To that end, fix any $t \le T - 1$ and define the weight $\lambda := \frac{\mu\|g_t\|}{16L^2}$ and the vector $g_\lambda := g_t + \lambda(\hat{g}_t - g_t)$. Notice that $\lambda \in [0, 1]$, since

$$\lambda = \frac{\mu\|g_t\|}{16L^2} \le \frac{\mu}{16L} \le 1,$$

where the first equation follows since $g_t \in \partial_\sigma f(x)$ and the second follows since $L \ge \mu$; see Lemma E.1. Thus

$$\begin{aligned}
\|g_{t+1}\|^2 \le \|g_\lambda\|^2 &= \|g_t\|^2 + 2\lambda\langle g_t, \hat{g}_t - g_t\rangle + \lambda^2\|\hat{g}_t - g_t\|^2 \\
&\le \|g_t\|^2 + 2\lambda\langle g_t, \hat{g}_t\rangle - 2\lambda\|g_t\|^2 + 4L^2\lambda^2 \\
&\le \|g_t\|^2 - \frac{\lambda\mu}{2}\|g_t\| + 4L^2\lambda^2 \\
&= \left(1 - \frac{\mu^2}{64L^2}\right)\|g_t\|^2,
\end{aligned}$$

where the first inequality follows by definition of g_{t+1}; the second inequality follows from the fact that L is a local Lipschitz constant of f near $\bar{x}$; and the third inequality follows from (5.16). Thus, to complete the proof, simply unfold this recursion to get the bound

$$g = g_T \le \left(1 - \frac{\mu^2}{64L^2}\right)^{T/2}\|g_0\|^2,$$

as desired.

6 Rapid Local Convergence of NTDescent

In this Section, we present our main convergence guarantees for the NTDescent method under Assumption A. The main results of the section are Theorem 6.3 and Theorem 6.5, which analyze the nonconvex and convex settings respectively. In the

Table 2 Parameters used throughout Sect. 6; see also Table 1

Parameter	Definition
s_{lb}	$c_0\ g_0$
a_1	$\min\{D_1, D_2/L\}$
a_2	$\frac{\min\{C_1/L, C_5\}}{2}$
δ_{LS}	$\min\ \frac{\delta_{\mathrm{A}}}{2}, \delta_{\mathrm{GI}}, \delta_{\mathrm{ND}}, \delta_{\mathrm{Grid}}, \frac{1}{2(a_1+2a_2)}, \frac{\gamma D_1^2 \min\{\delta_{\mathrm{Grid}}/2, 1/4\}^2}{2L}, 1$
C_6	$\max\ 1, \frac{8(C_{(a)}+2\beta+2C_{(a)}D_1^{-1})}{s_{\mathrm{lb}}}, 2D_1^{-1}, \frac{4\gamma D_1}{s_{\mathrm{lb}}}$
$_{1,T}$	$\max\ \frac{(C_5^2 L+\beta)\ L}{C_3}\ \ 1-\frac{3\mu^2}{256L^2}^{T/2}, \ 1-\frac{\mu^2}{64L^2}^{T/2}\ L$
$_{2,G}$	$\max\ \frac{L}{\min\{1,a_1\}}+\frac{\beta}{2\min\{1,a_1\}a_2^2}, 8C_{(a)}, L\ \ 2^{-G}$
ρ	$1-\frac{1}{8}\min\ \frac{\gamma a_2}{8\max\{4La_2^2,\beta\}}, \frac{\mu a_1}{4\max\{2L,\beta/a_2^2\}}$

nonconvex setting, we prove that iterates of `NTDescent` locally nearly linearly converge, provided some iterate reaches a sufficiently small neighborhood of $\bar{x}$. In the convex setting, we strengthen this guarantee, showing that for any initial starting point x_0 and any failure probability p, there exists some index K_p after which `NTDescent` nearly converges linearly with probability at least $1-p$. Both results are a consequence of the local one-step improvement bound of Proposition 6.1. This proposition shows that with high probability, the following hold locally for `linesearch`: its output is nearby its input; and the function gap geometrically decreases whenever it is larger than a quantity that is exponentially small in the inner loop budget and the grid size. The former property will help ensure that the iterates of `NTDescent` do not escape a local neighborhood of $\bar{x}$.

6.1 Assumptions and Notation

Throughout this section, we assume the following assumptions and notations are in force. We assume that

1. the budget T_k and grid size G_k satisfy $\min\{T_k, G_k\} \geq k+1$ for all $k \geq 0$.
2. We fix an initial we an initial point $x_0 \in \mathbb{R}^d$ and $g_0 \in \partial f(x_0)$. We assume that $g_0 = 0$. We assume Assumption A is in force at a point $\bar{x} \in \mathbb{R}^d$ and use the notation of Proposition 3.5 throughout. We let $\{x_k\}$ denote the sequence of iterates generated by $\texttt{NTDescent}(x_0, g_0, c_0, \{G_k\}, \{T_k\})$ when applied to f.

Turning to notation, we now summarize in Table 2 the main constants used in this section.

In the following, we lower and upper bound the trust region parameter in `linesearch`:

$$s_{\mathrm{lb}} \leq \max\{\ g_k,\ \ c_0\ g_0\} \leq\ \ L, \tag{6.1}$$

where the lower bound follows by definition, and the upper bound follows from Part 6 of Proposition 3.5. In addition, we apply Theorem 4.3 with the constants a_1, a_2. These constants are derived from the parameters D_1, D_2, C_1, and C_5 which are defined in Lemmas 4.1, 4.2, and 5.5 respectively. We also define a neighborhood $B_{\delta_{\mathrm{LS}}}(\bar{x})$ for which `linesearch` results in geometric improvement. Here, the radius δ_{LS} is derived from the parameters δ_{A}, δ_{GI}, δ_{ND}, δ_{Grid}, and γ which appear in Proposition 3.5 and Lemmas 4.1, 5.1, 5.7, and 5.8. In addition, the constant C_6 will appear in an upper bound on the steplength of `linesearch`.

We then define three terms $\epsilon_{1,T}$, $\epsilon_{2,G}$, and ρ which appear in our convergence rate analysis. These terms are defined for all $T, G > 0$ and are derived from the parameters C_5, C_3, a_1, a_2, L, β, $C_{(a)}$, γ, and μ which appear in Lemma 5.2, Lemma 5.5, Proposition 3.5, and Assumption A.

Finally, in the following propositions, the constant $\rho \in (0, 1)$ plays the role of a local contraction factor, while the terms $\epsilon_{1,T}$ and $\epsilon_{2,G}$ are upper bounds for function gap of `NTDescent`.

We now turn to the one-step improvement argument.

6.2 One Step Improvement

The following proposition presents our one-step improvement bound.

Proposition 6.1 (One step improvement) *Assume the assumptions of Sect. 6.1 are satisfied. Recall the notation in Table 2. Then the following holds for all* $x \in B_{\delta_{\mathrm{LS}}}(\bar{x})$, *subgradients* $g \in \partial f(x)$, *and grid sizes* $G > \log_2(1/\delta_{\mathrm{Grid}})$ *: Fix a scalar* $s \in [s_{\mathrm{lb}}, L]$, *a failure probability* $p \in (0, 1)$ *and budget* T *satisfying*

$$T \geq \frac{256L^2}{\mu^2} 2\log(1/p) .$$

Then with probability at least $1 - p$, *the point* $\tilde{x} = \texttt{linesearch}(x, g, s, G, T)$ *satisfies*

1. $f(\tilde{x}) - f(\bar{x}) \leq \max\{\rho(f(x) - f(\bar{x})), \epsilon_{1,T}, \epsilon_{2,G}\}$;
2. $\tilde{x} - x \leq C_6 \max\left\{\epsilon_{1,T}/s_{\mathrm{lb}}, \epsilon_{2,G}/s_{\mathrm{lb}}, \sqrt{2(f(x) - f(\bar{x}))/\min\{s_{\mathrm{lb}}, \gamma\}}\right\}$.

Proof We fix $x \in B_{\delta_{\mathrm{LS}}}(\bar{x})$, define $y := P_M(x)$, and choose a subgradient $g \in \partial f(x)$. Throughout we may freely use the results of Proposition 3.5 since $\delta_{\mathrm{LS}} \leq \delta_{\mathrm{A}}$. We will first establish the first item of the Proposition. To that end, let us assume that

$$f(x) - f(\bar{x}) > \max\{\epsilon_{1,T}, \epsilon_{2,G}\};$$

otherwise the proof is trivial. In this case, we claim that x must satisfy either Item 1 or Item 2 of Theorem 4.3 for at least one σ_i with $i \leq G - 1$. To derive a contradiction, suppose that both items are not satisfied for x with any choice of σ_i with $i = 0, \ldots, G - 1$. We will show that neither Item 1b nor its complement can be satisfied, leading to a contradiction.

Throughout the following argument, we will use the following bound:

$$\max\{a_1 \mathrm{dist}(x, M),\ a_2\ y - \bar{x}\} \leq\ (a_1 + 2a_2)\,\delta_{\mathrm{LS}} \leq \frac{1}{2} = \sigma_{G-1}.$$

Now suppose that Item 1b holds, i.e., $\frac{a_2}{2}\ y - \bar{x}\ ^2 \leq \mathrm{dist}(x, M)$. Then by assumption, Item 1a must fail for any σ_i. We claim that this failure ensures that $\sigma_0 > a_1 \mathrm{dist}(x, M)$. Indeed, if $\sigma_0 \leq a_1 \mathrm{dist}(x, M)$, we must have

$$\sigma_0 \leq (a_1/2)\mathrm{dist}(x, M) \leq\ a_1 \mathrm{dist}(x, M) \leq \sigma_{\ G-1},$$

since σ_0 cannot satisfy Item 1a. Thus, there exists some $j \leq\ G - 1$ such that $\sigma_j = 2^j \sigma_0$ satisfies Item 1a, a contradiction. Therefore, we have

$$\sigma_0 > a_1 \mathrm{dist}(x, M) \geq\ a_1 a_2^2\ y - \bar{x}\ ^2.$$

In this case, by (3.7), we have

$$f(x) - f(\bar{x}) \leq L\mathrm{dist}(x, M) + \frac{\beta}{2}\ y - \bar{x}\ ^2 \leq \frac{L}{a_1} + \frac{\beta}{2a_2^2 a_1}\ \sigma_0 \leq\ _{2,G},$$

which is a contradiction. Therefore, Item 1b cannot hold, so we have $\frac{a_2}{2}\ y - \bar{x}\ ^2 > \mathrm{dist}(x, M)$.

Next, for the sake of contradiction, suppose that there exists σ_i satisfying Item 2a. In this case, since $\sigma_i \geq (a_2/2)\ y - \bar{x}\ $, we have

$$\mathrm{dist}(x, M) <\ a_2^2\ y - \bar{x}\ ^2 \leq 2a_2\sigma_i\ y - \bar{x},$$

i.e., σ_i also satisfies Item 2b, which is a contradiction. Therefore no σ_i satisfies Item 2a. We claim that this ensures $\sigma_0 > a_2\ y - \bar{x}\ $. Indeed, if $\sigma_0 \leq a_2\ y - \bar{x}\ $, we must have

$$\sigma_0 \leq (a_2/2)\ y - \bar{x} \leq\ a_0\ y - \bar{x} \leq \sigma_{\ G-1},$$

since σ_0 cannot satisfy Item 2a. Thus, there exists some $j \leq\ G - 1$ such that $\sigma_j = 2^j \sigma_0$ satisfies Item 2a, a contradiction. Therefore, we have

$$\sigma_0 > a_2\ y - \bar{x} \geq\ \overline{\mathrm{dist}(x, M).}$$

In this case, by (3.7), we have

$$f(x) - f(\bar{x}) \leq L\mathrm{dist}(x, M) + \frac{\beta}{2}\ y - \bar{x}\ ^2 \leq\ L + \frac{\beta}{2a_2^2}\ \sigma_0^2 \leq\ _{2,G},$$

which is a contradiction. Therefore, there must exist σ_i satisfying either Item 1 or Item 2 of Theorem 4.3.

Let us now fix a σ_i satisfying either Item 1 or Item 2 of Theorem 4.3. Then, by Theorem 4.3, we have the bound

$$\sigma_i \operatorname{dist}(0, \partial_{\sigma_i} f(x)) \geq 8(1-\rho)(f(x) - f(\bar{x})).$$

In what follows, we will use the above bound to prove that with probability at least $1-p$, we have $f(\tilde{x}) - f(\bar{x}) \leq \rho(f(x) - f(\bar{x}))$ whenever $f(x) - f(\bar{x}) > \max\{\varepsilon_{1,T}, \varepsilon_{2,G}\}$.

Contraction case 1: normal step. We first suppose that there exists σ_i satisfying Item 1. In the interest of analyzing $v_{i+1} \in \partial_{\sigma_i} f(x)$, let us show that x, σ_i, and T satisfy the conditions of Proposition 5.1: First $x \in B_{\delta_{\mathrm{ND}}}(\bar{x})$ since $\delta_{\mathrm{LS}} \leq \delta_{\mathrm{ND}}$. Second, by Item 1a of Theorem 4.3, we have

$$0 < \sigma_i \leq a_1 \operatorname{dist}(x, M) \leq D_1 \operatorname{dist}(x, M). \tag{6.2}$$

Finally, from the definition $D_2 = \mu/2$, it follows that T satisfies the conditions of Proposition 5.1. Therefore, with probability at least $1-p$, we have

$$f\left(x - \sigma_i \frac{v_{i+1}}{\|v_{i+1}\|}\right) - f(\bar{x}) \leq f(x) - f(\bar{x}) - \frac{\sigma_i}{8} \operatorname{dist}(0, \partial_{\sigma_i} f(x)) \leq \rho(f(x) - f(\bar{x})).$$

Next, we show that v_{i+1} and σ_i satisfy the trust region condition $\sigma_i \leq \frac{\|v_{i+1}\|}{s}$. To that end, note that the conditions of Lemma 4.1 are met: We have $x \in B_{\delta_{\mathrm{GI}}}(\bar{x})$ since $\delta_{\mathrm{LS}} \leq \delta_{\mathrm{GI}}$. We also have bound $\sigma_i \leq D_1 \operatorname{dist}(x, M)$ from (6.2). Therefore, it follows that the minimal norm Goldstein subgradient is lower bounded: $\operatorname{dist}(0, \partial_{\sigma_i} f(x)) \geq D_2$. Consequently, we have

$$\sigma_i \leq a_1 \operatorname{dist}(x, M) \leq \frac{D_2 \delta_{\mathrm{LS}}}{s} \leq \frac{\operatorname{dist}(0, \partial_{\sigma_i} f(x)) \delta_{\mathrm{LS}}}{s} \leq \frac{\|v_{i+1}\|}{s},$$

where the second inequality follows from the definition of in Table 2 and the inequality $s \leq L$; and the fourth inequality follows from the bound $\delta_{\mathrm{LS}} \leq 1$. Therefore, since the trust region constraint $\sigma_i \leq \frac{\|v_{i+1}\|}{s}$ is satisfied, the following holds with probability at least $1-p$:

$$f(\tilde{x}) - f(\bar{x}) \leq f\left(x - \sigma_i \frac{v_{i+1}}{\|v_{i+1}\|}\right) - f(\bar{x}) \leq \rho(f(x) - f(\bar{x})).$$

Thus, the first item of the proposition follows.

Contraction case 2: tangent step. Next, we suppose that there exists σ_i satisfying Item 2 of Theorem 4.3. In the interest of analyzing $u_i \in \partial_{\sigma_i} f(x)$, let us show that x, σ_i, and T satisfy the conditions of Proposition 5.6: $x \in B_{\delta_{\mathrm{A}}/2}(\bar{x})$ since $\delta_{\mathrm{LS}} \leq \delta_{\mathrm{A}}/2$. Second, by Item 2a of Theorem 4.3, we have

$$\sigma_i \leq a_2 \|y - \bar{x}\| \leq C_5 \|y - \bar{x}\|.$$

Finally, by Item 2b of Theorem 4.3, we have

$$\operatorname{dist}(x, M)/\sigma_i \le 2a_2 \|y - \bar{x}\| \le C_5 \|y - \bar{x}\|.$$

Therefore, since $f(x) - f(\bar{x}) > \kappa_{1,T}$, Proposition 5.6 implies that

$$f\left(x - \sigma_i \frac{u_i}{\|u_i\|}\right) - f(\bar{x}) \le f(x) - f(\bar{x}) - \frac{\sigma_i}{8}\operatorname{dist}(0, \partial_{\sigma_i} f(x)) \le \rho(f(x) - f(\bar{x})).$$

Next, we show that u_i and σ_i satisfy the trust region condition $\sigma_i \le \frac{\|u_i\|}{s}$. To show this, we first note that σ_i and x satisfy the conditions of Lemma 4.2: First $x \in B_{\delta_A/2}(\bar{x})$ since $\delta_{LS} \le \delta_A/2$. Second, by Item 2a of Theorem 4.3, we have

$$\sigma_i \le a_2 \|y - \bar{x}\| \le C_2 \|y - \bar{x}\|.$$

Finally, by Item 2 of Theorem 4.3, we have

$$\operatorname{dist}(x, M) \le 2a_2\sigma_i \|y - \bar{x}\| \le 2a_2^2 \|y - \bar{x}\|^2 \le C_2 \|y - \bar{x}\|,$$

where the third inequality follows from the bounds $\|y - \bar{x}\| \le 2\delta_{LS} \le 1/a_2$ and $a_2 \le C_2/2$ (recall that $C_5 \le C_2$). Therefore, by Lemma 4.2 we have $\|u_i\| \ge \|P_{T_M(y)} u_i\| \ge C_1 \|y - \bar{x}\|$. Consequently, we have

$$\sigma_i \le a_2 \|y - \bar{x}\| \le \frac{C_1 \|y - \bar{x}\|}{s} \le \frac{\|u_i\|}{s},$$

where the second inequality follows from the definition of a_2 in Table 2 and the inequality $s \le L$. To complete the proof, observe that $v_{i+1} = u_i$: since the sufficient descent condition is met, namely $f(x - \sigma_i u_i/\|u_i\|) \le f(x) - \sigma\|u_i\|$, `NDescent` terminates at the first iteration. Therefore, we must have

$$f(\tilde{x}) - f(\bar{x}) \le f\left(x - \sigma_i \frac{v_{i+1}}{\|v_{i+1}\|}\right) - f(\bar{x}) \le \rho(f(x) - f(\bar{x})),$$

as desired.

Having proved the desired contraction $f(\tilde{x}) - f(\bar{x}) \le \rho(f(x) - f(\bar{x}))$, we now turn to the bound on $\|\tilde{x} - x\|$.

Stepsize bound. We now no longer assume that $f(x) - f(\bar{x}) > \max\{\kappa_{2,G}, \kappa_{1,T}\}$. We claim that we have

$$\max_{0 \le i \le G-1} \{\sigma_i : \sigma_i \le \|v_{i+1}\|/s\} \le C_6 \max\left\{\kappa_{1,T}/s_{lb}, \kappa_{2,G}/s_{lb}, \sqrt{2(f(x) - f(\bar{x}))/\min\{s_{lb}, \gamma\}}\right\}. \tag{6.3}$$

Note that inequality (6.3) immediately yields the second item of the proposition, since

$$\|\tilde{x} - x\| \le \max_{0 \le i \le G-1} \{\sigma_i : \sigma_i \le \nu\, \|u_{i+1}\|/s\}.$$

To prove (6.3), we will apply Lemma 5.8.

To that end, first note that $x \in B_{\delta_{\mathrm{Grid}}}(\bar{x})$ since $\delta_{\mathrm{LS}} \le \delta_{\mathrm{Grid}}$. Next, we verify that there exists an index i such that σ_i satisfies a slightly stronger version of the assumptions of Lemma 5.8. Indeed, recall that by the quadratic growth condition (A1), we have the bound

$$\mathrm{dist}(x, M) \le \|x - \bar{x}\| \le \sqrt{2(f(x) - f(\bar{x}))/\gamma}. \tag{6.4}$$

Thus, to satisfy the assumptions of Lemma 5.8, we prove that there exists i such that

$$R_x := \|D_1^{-1}\| \sqrt{2(f(x) - f(\bar{x}))/\gamma} \le \sigma_i \le \delta_{\mathrm{Grid}}. \tag{6.5}$$

Indeed, first notice that $\sigma_0 \le \delta_{\mathrm{Grid}}$ since $G \ge \log_2(1/\delta_{\mathrm{Grid}})$. Thus, if $\sigma_0 \ge R_x$, the bound (6.5) holds for σ_0. If instead $\sigma_0 < R_x$, we have

$$\sigma_0 < R_x \le \|D_1^{-1}\| \sqrt{2L\delta_{\mathrm{LS}}/\gamma} \le \min\{\delta_{\mathrm{Grid}}/2, 1/4\} \le \min\{\delta_{\mathrm{Grid}}, 1/2\} \le 1/2 = \sigma_{G-1},$$

where the second inequality follows since $\|x - \bar{x}\| < \delta_{\mathrm{LS}}$ and f is $L-$Lipschitz continuous on $B_{\delta_{\mathrm{LS}}}(\bar{x})$; and the third inequality follows since $\delta_{\mathrm{LS}} \le \gamma\, \|D_1^2\| \min\{\delta_{\mathrm{Grid}}/2, 1/4\}^2/(2L)$. Thus, there exists i such that $\sigma_i \in [\min\{\delta_{\mathrm{Grid}}/2, 1/4\}, \min\{\delta_{\mathrm{Grid}}, 1/2\}]$. Since $\min\{\delta_{\mathrm{Grid}}/2, 1/4\} \ge R_x$, inequality (6.5) follows.

Now let i_* be the minimal such index such that (6.5) is satisfied for $i = i_*$. If $i_* = 0$, the bound $\sigma_{i_*-1} \le R_x$ holds. In particular, $\sigma_{i_*} \le 2R_x$. Therefore, considering the cases $i_* = 0$ and $i_* = 0$ separately, we have

$$R_x \le \sigma_{i_*} \le \max\{\sigma_0, 2R_x\}. \tag{6.6}$$

Now we bound the step length $\|x - \tilde{x}\|$ by considering two cases.

First suppose that $\sigma_{i_*} > \|u_{i_*}\|/s$. In this case, (2.1) ensures $\sigma_{i_*} > \nu\, \|u_{i_*+1}\|/s$. Then, since σ_i is increasing in i, we have

$$\begin{aligned}
&\max_{0 \le i \le G-1} \{\sigma_i : \sigma_i \le \nu\, \|u_{i+1}\|/s\} \le \sigma_{i_*} \\
&\le \max\{\sigma_0, 2R_x\} \\
&\le C_6 \max\left\{\|u_{2,G}\|/s_{\mathrm{lb}}, \sqrt{2(f(x) - f(\bar{x}))/\min\{s_{\mathrm{lb}}, \gamma\}}\right\} \\
&\le C_6 \max\left\{\|u_{1,T}\|/s_{\mathrm{lb}}, \|u_{2,G}\|/s_{\mathrm{lb}}, \sqrt{2(f(x) - f(\bar{x}))/\min\{s_{\mathrm{lb}}, \gamma\}}\right\},
\end{aligned}$$

which verifies (6.3). We now consider the alternative case.

Next suppose that $\sigma_{i_*} \le u_{i_*}/s$. We consider two subcases. First suppose that the following bound also holds:

$$u_{i_*} \le \frac{8(f(x) - f(\bar{x}))}{\sigma_{i_*}}. \tag{6.7}$$

Then, since $\sigma_{i_*} \ge R_x$, we have

$$u_{i_*} \le \sqrt{32\gamma D_1^2 (f(x) - f(\bar{x}))}.$$

Second, suppose that (6.7) does not hold. Let us apply Lemma 5.8 to $\sigma = \sigma_{i_*}$:

$$\begin{aligned}
u_{i_*} &\le \max\left\{1 - \frac{\mu^2}{64L^2}\right\}^{T/2} L, 4C_{(a)} \max\{\sigma_0, 2R_x\} + 4(C_{(a)} + 2\beta)\|x - \bar{x}\| \\
&\le \max\left\{1 - \frac{\mu^2}{64L^2}\right\}^{T/2} \cdot, 8C_{(a)}\sigma_0, 8(C_{(a)} + 2\beta)\|x - \bar{x}\| + 16C_{(a)}R_x \\
&\le \max\{1 \cdot \Delta_{1,T}, 1 \cdot \Delta_{2,G}, 8(C_{(a)} + 2\beta + 2C_{(a)}D_1^{-1})\sqrt{2(f(x) - f(\bar{x}))/\gamma}} \\
&\le sC_6 \max\{\Delta_{1,T}/s_{\mathrm{lb}}, \Delta_{2,G}/s_{\mathrm{lb}}, \sqrt{2(f(x) - f(\bar{x}))/\gamma}\},
\end{aligned}$$

where first inequality follows from Lemma 5.8 and bound (6.6); the second inequality follows from the bound: $\max\{a, b\} + c \le a + b + c \le 2\max\{a, b + c\}$ for all $a, b, c \ge 0$; the third inequality follows by definition of $\Delta_{1,T}$, $\Delta_{2,G}$ and R_x, and (6.4); and the last inequality follows since $C_6 \ge \max\{1, 8(C_{(a)} + 2\beta + 2C_{(a)}D_1^{-1})/s_{\mathrm{lb}}\}$.

Therefore, as long as $\sigma_{i_*} \le u_{i_*}/s$, we have

$$\begin{aligned}
&u_{i_*}/s \\
&\le \max\left\{C_6 \max\{\Delta_{1,T}/s_{\mathrm{lb}}, \Delta_{2,G}/s_{\mathrm{lb}}, \sqrt{2(f(x) - f(\bar{x}))/\gamma}\}, \sqrt{32\gamma D_1^2 (f(x) - f(\bar{x}))/s_{\mathrm{lb}}^2}\right\} \\
&\le C_6 \max\left\{\Delta_{1,T}/s_{\mathrm{lb}}, \Delta_{2,G}/s_{\mathrm{lb}}, \sqrt{2(f(x) - f(\bar{x}))/\gamma}\right\} \\
&\le C_6 \max\left\{\Delta_{1,T}/s_{\mathrm{lb}}, \Delta_{2,G}/s_{\mathrm{lb}}, \sqrt{2(f(x) - f(\bar{x}))/\min\{s_{\mathrm{lb}}, \gamma\}}\right\},
\end{aligned}$$

where second inequality follows from the bound $C_6 \ge 4\gamma D_1/(s_{\mathrm{lb}})$; and the third inequality follows from the bound (6.4). To complete the proof of (6.3), recall that by (2.1), for all $j > i_*$, we have $v_j \le u_{i_*}$. Consequently,

$$\begin{aligned}
&\max_{0 \le i \le G-1} \{\sigma_i : \sigma_i \le v_{i+1}/s\} \\
&\le \max\{\sigma_{i_*}, u_{i_*}/s\} \\
&= u_{i_*}/s \\
&\le C_6 \max\left\{\Delta_{1,T}/s_{\mathrm{lb}}, \Delta_{2,G}/s_{\mathrm{lb}}, \sqrt{2(f(x) - f(\bar{x}))/\min\{s_{\mathrm{lb}}, \gamma\}}\right\},
\end{aligned}$$

Table 3 Parameters used throughout Sect. 6.3; see also Tables 1 and 2

Parameter	Definition
C	$\frac{2048L^2}{\mu^2}$
C	$\max\left\{\frac{(C_5^2L+\beta)\,L}{C_3}, L, \frac{L}{\min\{1,a_1\}} + \frac{\beta}{2\min\{1,a_1\}a_2^2}, 8C_{(a)}\right\}$
q	$\max\left\{\rho, 1-\frac{3\mu^2}{256L^2}, \frac{1}{2}\right\}$
δ_{NTD}	$\min\left\{\frac{\delta_{\text{LS}}}{4}, \frac{\delta_{\text{LS}}^2\min\{s_{\text{lb}},\gamma\}(1-q^{1/2})^2}{32LC_6^2}, \frac{\delta_{\text{LS}}s_{\text{lb}}(1-q)}{4LC_6}\right\}$
K_0	$\max\left\{\log_q \frac{\delta_{\text{LS}}^2\min\{s_{\text{lb}},\gamma\}(1-q^{1/2})^2}{32CC_6^2}, \log_q \frac{\delta_{\text{LS}}s_{\text{lb}}(1-q)}{4CC_6}, \log_2 \frac{1}{\delta_{\text{Grid}}}\right\}$

which verifies (6.3).

6.3 Main Convergence Theorems

We are now ready to prove the main results of this work. The goal of this section is to prove that an event of the following form occurs with high probability.

Definition 6.2 ($E_{k_0,q,C}$) For any $k_0 > 0$, $q \in (0,1)$ and $C > 0$, let $E_{k_0,q,C}$ denote the event that for all $k \geq k_0$, we have the following two bounds:

$$f(x_k) - f(\bar{x}) \leq \max\{(f(x_{k_0}) - f(\bar{x}))q^{k-k_0}, Cq^k\};$$
$$\|x_k - \bar{x}\|^2 \leq \frac{2}{\gamma}\max\{(f(x_{k_0}) - f(\bar{x}))q^{k-k_0}, Cq^k\}.$$

We will lower bound the probability of the event $E_{k_0,q,C}$ in both nonconvex and convex settings for a particular choice of k_0, q, and C. In the nonconvex setting, our result will lower bound the conditional probability of $E_{k_0,q,C}$, given that iterate x_{k_0} enters a sufficiently small neighborhood of $\bar{x}$. To prove the result, we will simply iterate the one-step improvement bound of Proposition 6.1. In the convex setting, we will lower bound the unconditional probability of $E_{k_0,q,C}$. To prove this result, we will combine the conditional result with the sublinear convergence guarantee of Theorem 2.4.

Before turning to the proofs, we introduce the main parameters that are common to both the nonconvex and convex settings.

6.3.1 The Nonconvex Setting

The following theorem is our main convergence theorem in the nonconvex setting.

Theorem 6.3 (Main Theorem: Nonconvex Setting) *Assume the assumptions outlined at the start of Sect. 6 are satisfied. Recall the notation of Table 3. Fix a failure probability* $p \in (0,1)$ *and an index* $k_0 \geq \max\{K_0, C\log(C/p)\}$. *Suppose*

$P(x_{k_0} \in B\delta_{\mathrm{NTD}}(\bar{x})) > 0$. *Then,*

$$P(E_{k_0,q,C} \mid x_{k_0} \in B\delta_{\mathrm{NTD}}(\bar{x})) \geq 1 - p.$$

Proof We begin with preliminary notation and bounds.

Fix $k_0 \geq \max\{K_0, C \log(C/p)\}$ and for all $k \geq k_0$, define the quantity

$$R_k := \max\{(f(x_{k_0}) - f(\bar{x}))q^{k-k_0}, Cq^k\}.$$

Note that whenever $x_{k_0} \in B\delta_{\mathrm{NTD}}(\bar{x})$ we have the bound

$$R_k \leq \max\{L\delta_{\mathrm{NTD}}q^{k-k_0}, Cq^k\}, \tag{6.8}$$

since f is L-Lipschitz continuous on $B\,\delta(\bar{x})$.

Next, we prove that

$$\max\{\ _{1,T_k}, \ _{2,G_k}\} \leq R_{k+1} \quad \text{for all } k \geq 0. \tag{6.9}$$

Indeed, beginning with $_{1,T_k}$, we have

$$\begin{aligned} {}_{1,T_k} &= \max\left\{\frac{(C_5^2 L + \beta)\,L}{C_3}\left(1 - \frac{3\mu^2}{256L^2}\right)^{T_k/2}, \left(1 - \frac{\mu^2}{64L^2}\right)^{T_k/2} L\right\} \\ &\leq C\max\left\{\left(1 - \frac{3\mu^2}{256L^2}\right)^{\frac{T_k}{2}}, \left(1 - \frac{\mu^2}{64L^2}\right)^{\frac{T_k}{2}}\right\} \\ &\leq Cq^{k+1} \leq R_{k+1}, \end{aligned}$$

where the first and second inequalities follow from the definitions of C and q together with the lower bound $T_k \geq k + 1$. Turning to $_{2,G_k}$, we have

$$\begin{aligned} {}_{2,G_k} &= \max\left\{\frac{L}{\min\{1, a_1\}} + \frac{\beta}{2\min\{1, a_1\}a_2^2}, 8C_{(a)}, L\right\} 2^{-G_k} \leq C2^{-G_k} \\ &\leq Cq^{k+1} \leq R_{k+1}, \end{aligned}$$

where the first and second inequalities follow from the definition of C and q together with the lower bound $G_k \geq k + 1$. Thus (6.9) holds.

Finally, we analyze the quantity

$$D_{k_0,\delta_{\mathrm{NTD}}} := \sum_{k=k_0}^{\infty} C_6 \max\left\{\sqrt{2R_k/\gamma}, R_{k+1}/s_{\mathrm{lb}}\right\} \quad \text{where } \gamma := \min\{s_{\mathrm{lb}}, \gamma\}.$$

We claim in particular that

$$D_{k_0,\delta_{\mathrm{NTD}}} + \delta_{\mathrm{NTD}} \le \delta_{\mathrm{LS}}/2. \tag{6.10}$$

Since $\delta_{\mathrm{NTD}} \le \delta_{\mathrm{LS}}/4$, it suffices to prove $D_{k_0,\delta_{\mathrm{NTD}}} \le \delta_{\mathrm{LS}}/4$. To that end, we have

$$\begin{aligned}
D_{k_0,\delta_{\mathrm{NTD}}} &= \sum_{k=k_0}^{\infty} C_6 \max\left\{\sqrt{2R_k/\gamma},\, R_{k+1}/s_{\mathrm{lb}}\right\} \\
&\le \sum_{k=k_0}^{\infty} C_6 \max\left\{\sqrt{2\max\{L\delta_{\mathrm{NTD}}q^{k-k_0}, Cq^k\}/\gamma},\, \max\{L\delta_{\mathrm{NTD}}q^{k-k_0}/\gamma,\, Cq^k\}/s_{\mathrm{lb}}\right\} \\
&\le C_6 \max\left\{\frac{\sqrt{2L\delta_{\mathrm{NTD}}}}{\sqrt{\gamma}(1-q^{1/2})}, \frac{\sqrt{2Cq^{k_0}}}{\sqrt{\gamma}(1-q^{1/2})}, \frac{L\delta_{\mathrm{NTD}}}{s_{\mathrm{lb}}(1-q)}, \frac{Cq^{k_0}}{s_{\mathrm{lb}}(1-q)}\right\} \\
&\le \frac{\delta_{\mathrm{LS}}}{4},
\end{aligned}$$

where the first inequality follows from the bounds (6.8) and the bound $R_{k+1} \le R_k$; the second inequality follows by summing the infinite series; and the third inequality follows from the definitions of K_0 and δ_{NTD} together with the bound $k_0 \ge K_0$. This proves (6.10).

We now turn to the proof. Consider the following sequence defined for all $k \ge k_0$:

$$b_k := \delta_{\mathrm{NTD}} + \sum_{j=k_0}^{k-1} C_6 \max\left\{\sqrt{2R_k/\gamma},\, R_{k+1}/s_{\mathrm{lb}}\right\}.$$

Note that (6.10) ensures that $b_k \le \delta_{\mathrm{LS}}/2$ for all $k \ge k_0$. Now, define the event

$$F_{k_0} := \{x_{k_0} \in B_{\delta_{\mathrm{NTD}}}(\bar{x})\}.$$

In addition, define the following decreasing sequence of events

$$A_k := \bigcap_{j=k_0}^{2^k} \left\{f(x_j) - f(\bar{x}) \le R_j \text{ and } \|x_j - \bar{x}\| \le b_j\right\}.$$

We claim that

$$P(A_{k+1} \mid A_k \cap F_{k_0}) \ge 1 - \exp(-T_k/C) \quad \text{for all } k \ge k_0. \tag{6.11}$$

Indeed, Proposition 6.1 implies that conditioned on $A_k \cap F_{k_0}$, the following four inequalities are satisfied with probability at least $1 - \exp(-T_k/C)$:

1. $f(x_k) - f(\bar{x}) \le R_k$
2. $\|x_k - \bar{x}\| \le b_k$;

3. $x_{k+1} - x_k \le C_6 \max\{\langle 1, T_k\rangle / s_{\mathrm{lb}}, \langle 2, G_k\rangle / s_{\mathrm{lb}}, \sqrt{2(f(x_k) - f(\bar{x}))/\gamma}\}$;
4. $f(x_{k+1}) - f(\bar{x}) \le \max\{\rho(f(x_k) - f(\bar{x})), \langle 1, T_k\rangle, \langle 2, G_k\rangle\}$.

(Note that in applying the Proposition 6.1, we use the scalar $s = \max\{g_k, c_0 g_0\}$ and the inclusion $s \in [s_{\mathrm{lb}}, L]$, which was proved (6.1).) Thus, the bound (6.11) will follow by induction if we can prove that whenever the above four inequalities hold, we have $x_{k+1} - \bar{x} \le b_{k+1}$ and $f(x_{k+1}) - f(\bar{x}) \le R_{k+1}$.

To that end, we first prove $x_{k+1} - \bar{x} \le b_{k+1}$. Indeed,

$$\begin{aligned} x_{k+1} - \bar{x} &\le \|x_{k+1} - x_k\| + \|x_k - \bar{x}\| \\ &\le C_6 \max\{\langle 1, T_k\rangle / s_{\mathrm{lb}}, \langle 2, G_k\rangle / s_{\mathrm{lb}}, \sqrt{2(f(x_k) - f(\bar{x}))/\gamma}\} + b_k \\ &\le C_6 \max\{\sqrt{2R_k/\gamma}, R_{k+1}/s_{\mathrm{lb}}\} + b_k = b_{k+1}, \end{aligned} \tag{6.12}$$

where the second inequality follows from Proposition 6.1; and the third inequality follows from the bound (6.9). Next, we prove the bound on $f(x_{k+1}) - f(\bar{x}) \le R_{k+1}$. Indeed,

$$\begin{aligned} f(x_{k+1}) - f(\bar{x}) &\le \max\{\rho(f(x_k) - f(\bar{x})), \langle 1, T_k\rangle, \langle 2, G_k\rangle\} \\ &\le \max\{\rho \max\{(f(x_{k_0}) - f(\bar{x}))q^{k-k_0}, Cq^k\}, \langle 1, T_k\rangle, \langle 2, G_k\rangle\} \\ &\le R_{k+1}, \end{aligned}$$

where the final inequality follows from (6.9) and the bound $\rho \le q$. Consequently, the bound (6.11) holds. Moreover, due to the bound $T_k \ge k + 1$, we have

$$P(A_{k+1} \mid A_k \cap F_{k_0}) \ge 1 - \exp(-T_k/C) \ge 1 - \exp(-(k+1)/C). \tag{6.13}$$

Now we relate A_k to $E_{k_0,q,C}$. To that end, by the conditional law of total probability, for all $k \ge k_0$, we have

$$P(A_{k+1} \mid F_{k_0}) \ge P(A_{k+1} \mid A_k \cap F_{k_0}) P(A_k \mid F_{k_0}) \ge P(A_k \mid F_{k_0}) - \exp(-(k+1)/C).$$

Therefore, for all $k \ge k_0$, we have

$$P(A_k \mid F_{k_0}) \ge P(A_{k_0} \mid F_{k_0}) - \sum_{j=k_0+1}^{\infty} \exp(-j/C) = 1 - \frac{\exp\left(-\frac{k_0+1}{C}\right)}{1 - \exp\left(-\frac{1}{C}\right)} \ge 1 - p,$$

where the equality follows since $P(A_{k_0} \mid F_{k_0}) = 1$; and the final inequality follows by definition of $k_0 \ge C \log(C/p)$. Now recall that $\sup_{k \ge k_0} b_k \le \delta_{\mathrm{LS}}/2$. Therefore, defining the event

$$E_{k_0,q,C} := \{f(x_k) - f(\bar{x}) \le R_k \text{ for all } k \ge k_0 \text{ and } x_k \in B\delta_{\mathrm{LS}}(\bar{x})\},$$

we have

$$P(E_{k_0,q,C} \mid F_{k_0}) \geq \lim_{k\to\infty} P(A_k \mid F_{k_0}) \geq 1 - p.$$

Next, recall that since $\delta_{\mathrm{LS}} \leq \delta_{\mathrm{A}}$, the quadratic growth bound (A1)

$$\|x_k - \bar{x}\|^2 \leq \frac{2}{\gamma}(f(x_k) - f(\bar{x})) \leq \frac{2}{\gamma} R_k$$

holds for every $k \geq k_0$ within the event $E_{k_0,q,C}$. Thus, $E_{k_0,q,C} \supseteq E_{k_0,q,C}$. Therefore, we have

$$P(E_{k_0,q,C} \mid F_{k_0}) \geq P(E_{k_0,q,C} \mid F_{k_0}) \geq 1 - p,$$

as desired.

6.3.2 The Convex Setting

Now we turn to the convex setting. Our goal is to prove a lower bound on $P(E_{k_0,q,C})$ for q and C chosen as in Table 3 and all sufficiently large k_0. Before stating the result, we recall a simple fact about convex functions satisfying Assumption A. A similar result appears in [8, Section 2.4], but for completeness we provide a proof in Appendix F.

Lemma 6.4 *In addition to the assumption set out at the start of the section, suppose that function f is convex. Then for all $a > 0$, we have*

$$\{x \in \mathbb{R}^d : f(x) - f(\bar{x}) \leq a\} \subseteq \overline{B}_{r_a}(\bar{x}) \qquad \text{where } r_a := \max\left\{\frac{2a}{\gamma\,\delta_A}, \sqrt{\frac{2a}{\gamma}}\right\}.$$

In particular, f has bounded sublevel sets.

We now turn to our main theorem.

Theorem 6.5 (Main Theorem: Convex setting) *Assume the assumptions of Sect. 6.1 are satisfied. Recall the notation of Table 3. In addition, suppose that function f is convex. Consider the bounded set*

$$S := \{x + u : f(x) \leq f(x_0) \text{ and } u \in \overline{B}(x)\}.$$

Let L be a Lipschitz constant of f on S. Define the constants

$$a := \min\left\{\frac{\gamma\,\delta_{\mathrm{A}}\,\delta_{\mathrm{NTD}}}{4}, \frac{\gamma\,\delta_{\mathrm{NTD}}^2}{8}\right\}; \qquad \text{and} \qquad b := \inf_{\alpha\in(0,1)} \frac{64L\left(\frac{2}{\bar{\alpha}}\right)^{\frac{2}{(1-\alpha)}}}{\left(\frac{a}{\operatorname{diam}(S)}\right)^{\frac{2\alpha}{(1-\alpha)}}}.$$

Finally, define

$$K_1 := \frac{4\mathrm{diam}^2(S)}{a^2} \min\left\{16^2(f(x_0) - \inf f)^2, \frac{b}{4}, 2048L^2 \log\left(\frac{2}{p}\right), 128(L)^2 + \frac{(4L)^2}{a^2}\right\}.$$

Then, for every failure probability $p \in (0, 1)$, we have

$$P(E_{k_0,q,C}) \geq 1 - p \quad \text{for all } k_0 \geq \max\left\{K_0, C \log\left(\frac{2C}{p}\right), 2K_1 - 1\right\}.$$

Proof Theorem 6.3 shows that

$$P(E_{k_0,q,C} \mid x_{k_0} \in B\delta_{\mathrm{NTD}}(\bar{x})) \geq 1 - p/2 \quad \text{for all } k_0 \geq \max\left\{K_0, C \log\left(\frac{2C}{p}\right)\right\}. \tag{6.14}$$

We claim that

$$P(x_{k_0} \in B\delta_{\mathrm{NTD}}(\bar{x})) \geq 1 - p/2 \quad \text{for all } k_0 \geq 2K_1 - 1. \tag{6.15}$$

Note that this yields the proof, since in that case

$$P(E_{k_0,q,C}) \geq P(E_{k_0,q,C} \mid x_{k_0} \in U)P(x_{k_0} \in B\delta_{\mathrm{NTD}}(\bar{x})) \geq 1 + p^2/4 - p \geq 1 - p,$$

for all $k_0 \geq \max\left\{K_0, C \log\left(\frac{2C}{p}\right), 2K_1 - 1\right\}$.

Observe that (6.15) will follow if

$$P(f(x_{k_0}) - f(\bar{x}) \leq a) \geq 1 - p/2 \quad \text{for all } k_0 \geq 2K_1 - 1. \tag{6.16}$$

Indeed, by Lemma 6.4, we have.

$$\{x \in \mathbb{R}^d : f(x) - f(\bar{x}) \leq a\} \subseteq \overline{B}\delta_{\mathrm{NTD}/2}(\bar{x}) \subseteq B\delta_{\mathrm{NTD}}(\bar{x}).$$

To prove (6.16), we apply Theorem 2.4. To that end, note that $\{x \in \mathbb{R}^d : f(x) \leq f(x_0)\}$ and the widened sublevel set S are indeed bounded, due to Lemma 6.4. Therefore D and the Lipschitz constant L of f on S are finite. Now observe $G := \min_{K_1 \leq k \leq 2K_1 - 1}\{G_k\} \geq K_1$ since $G_k \geq k + 1$ for all k. Thus, there exists $i \leq G$ such that

$$(1/2)K_1^{-1/2} \leq \sigma_i \leq K_1^{-1/2}.$$

Therefore, applying Theorem 2.4 (in particular (2.2)) with this σ_i, we have

$$f(x_{2K_1 - 1}) - f(\bar{x})$$

$$\leq D \max \left\{ \frac{16(f(x_{K_1}) - \inf f)}{K_1^{1/2}}, \frac{16L\sqrt{2\log(2K_1^2/p)}}{K_1^{1/2}}, \frac{\sqrt{128}L}{K_1^{1/2}} \right\} + \frac{2L}{K_1^{1/2}} \tag{6.17}$$

with probability at least $1 - p/2$. Thus, to complete the proof, we show that the left-hand side of (6.17) is smaller than a. Indeed, it is straightforward to check that

$$\max \left\{ \frac{2L}{K_1^{1/2}}, \frac{16D(f(x_{K_1}) - \inf f)}{K_1^{1/2}}, \frac{\sqrt{128}DL}{K_1^{1/2}} \right\} \leq \frac{a}{2}.$$

Thus, the proof will follow if

$$\frac{16DL\sqrt{2\log(2K_1^2/p)}}{K_1^{1/2}} \leq \frac{a}{2}. \tag{6.18}$$

We perform this calculation in Appendix G. Thus, the proof is complete.

6.3.3 Local Oracle Complexity

Thus, we have established a local nearly linear convergence rate for `NTDescent`. To understand the overall complexity of the method, we must derive an upper bound on the contraction factor q. The following lemma, which is proved in Appendix H, provides one that depends on a worst-case condition number of f.

Lemma 6.6 *Suppose that without loss of generality that* $\delta_A \leq 1$. *Define the condition number*

$$\kappa = \frac{\max\{L, \beta, C_{(a)}\}}{\min\{\gamma, \mu\}}.$$

Then there exists a universal constant $\eta > 0$ *independent of* f *such that*

$$q \leq 1 - \frac{\eta}{\kappa^8(1 + C_M)^2}.$$

where q is defined as in Table 3.

With this upper bound on q, it is straightforward to derive a local complexity estimate for `NTDescent`: the method locally produces a point $\hat{x}$ satisfying $f(\hat{x}) - f(\bar{x}) \leq \varepsilon$ with at most

$$O\left(\kappa^8(1 + C_M)^2 \log(1/\varepsilon)\right)^3,$$

first-order oracle evaluations. This bound may be pessimistic since we did not attempt to optimize the constants C_i or a_i. We leave the improvement of this complexity as an intriguing open question.

Before moving to a brief numerical illustration, we explain how Theorem 1.1 from the introduction follows from the above results.

Remark 2 (Establishing Theorem 1.1) Theorem 1.1 from the introduction immediately follows from Theorems 6.3 and 6.5. Indeed, first the event $E_{k_0,q,C}$ from Theorems 6.3 and 6.5 contains the corresponding event $E_{k_0,q,C}$ from Theorem 1.1 for particular q and C, which depend solely on f. Second, from the statement of theorems, we see that the neighborhood of local nearly linear convergence, $B_{\delta_{\mathrm{NTD}}}(\bar{x})$, depends solely on f.

7 Numerical Illustration

In this section, we briefly illustrate the numerical performance of `NTDescent` on two nonsmooth objective functions, borrowed from [1, 12, 37, 39]. In both experiments, we compare `NTDescent` to the subgradient method with the popular Polyak stepsize (`PolyakSGM`) [47], which iterates

$$x_{k+1} = x_k - \frac{f(x_k) - \inf f}{w_k^2} w_k \qquad \text{for some } w_k \in \partial f(x_k).$$

In the first example, $\inf f$ is known, in the second, we estimate $\inf f$ from multiple runs of `NTDescent`. We choose to compare against the subgradient method because it is a simple first-order method with strong convergence guarantees in convex [47] and nonconvex settings [20]. Importantly, `PolyakSGM` accesses the objective solely through function and subgradient evaluations. Thus, we compare the accuracy achieved by `PolyakSGM` and `NTDescent` after a fixed number of oracle calls, i.e., evaluations of ∂f.

Let us comment on the implementation of `NTDescent`. First, in all experiments, unless otherwise noted, we do not tune parameters of `NTDescent`. Instead, we simply choose scaling constant $c_0 = 10^{-6}$ and loop size parameters

$$T_k = k + 1 \qquad \text{and} \qquad G_k = \min\{k + 1,\ \log_2(10^{16})\} \qquad \text{for all } k \geq 0.$$

Second, we attempt to save first-order oracle calls by breaking the loop on Lines 2 through 6 of Algorithm 3 whenever we find that $\sigma_i > v_{i+1}/s$. Since σ_i is increasing in i and v_{i+1} is nonincreasing in i, this does not affect the iterates x_k of `NTDescent`; see Lemma 2.1. Finally, in all problems, we initialize `NTDescent` and `PolyakSGM` at a random vector az where z is sampled from the uniform distribution on the unit sphere. For all problems, we use $a = 1$ unless otherwise noted. Note that in the problems of Sect. 7.1 and 7.2, the solution is known, while in the problem of Sect. 7.3, the solution is unknown.

The purpose of this section is not to argue that `NTDescent` is a substitute for standard subgradient methods in most problems. Instead, we only wish to point out

some scenarios where standard first-order methods are known to perform poorly, yet NTDescent asymptotically accelerates. We are also not arguing that NTDescent has fast global rates: indeed, we previously mentioned that the NTDescent's global rate is $O(\epsilon^{-6})$ which is much worse than PolyakSGM's $O(\epsilon^{-2})$ rate for general convex problems. In practice, one could devise schemes that couple NTDescent with PolyakSGM, eventually switching to NTDescent when it begins to outperform PolyakSGM. While we leave a more thorough numerical study to future work, the reader may download and run our PyTorch [45] implementation of NTDescent at the following url: https://github.com/COR-OPT/ntd.py

We now turn to the examples.

7.1 A Max-of-Smooth Function

In this example, f takes the following form

$$f(x) = \max_{i=1,\ldots,m} \; g_i^{\top} x + \frac{1}{2} x^{\top} H_i x \quad , \tag{7.1}$$

where we generate a random vector $\lambda \in \mathbb{R}^m$ in $\{\lambda > 0 : \sum_{i=1}^m \lambda_i = 0\}$, a random positive semi-definite matrix H_i, and a random vector g_i satisfying that $\sum_{i=1}^m \lambda_i g_i = 0$. In this case, one can show that with probability 1, f satisfies Assumption A at its unique minimizer 0.

In Fig. 6 we plot the performance of NTDescent and PolyakSGM for multiple pairs of (d, m), varying initialization scale, a slight modification of NTDescent that allows longer steps, and varying scales c_0. We begin with Fig. 6a, b. Figure 6a shows that the performance of NTDescent depends on m. On the other hand, Fig. 6b shows NTDescent performance is independent of d, as expected. Both plots show that NTDescent asymptotically outperforms PolyakSGM. Turning to initialization, Fig. 6c shows the result of initializing NTDescent at a random vector az, where z is uniformly drawn from the sphere and a is a scale parameter satisfying $a \in \{1, 10, 100\}$. Clearly, NTDescent is affected by the initialization scale but surpasses PolyakSGM after 30,000 oracle calls. While we expect NTDescent to converge slowly when far from minimizers, we introduce a simple strategy to mitigate this behavior.

Adaptive grid strategy. Briefly, suppose we run linesearch the full G steps without exiting (via the violation of the trust region constraint). Then we simply continue the linesearch loop trying $\sigma_{-1} = 10\sigma_0$, $\sigma_{-2} = 10\sigma_{-1}$, and so on, until we violate the trust region constraint or σ_{-i} exceeds a predefined threshold.

Figure 6d shows the result of this strategy with a predefined threshold ∞, showing that it compensates for poor initialization quality. Finally in Fig. 6e, f, we show the effect of changing the c_0 input to NTDescent. It appears NTDescent is relatively insensitive to c_0 and smaller choices generally result in better performance. This motivates our default choice $c_0 = 10^{-6}$ in the remainder of the experiments.

Before turning to our second experiment, we briefly mention two alternative methods—Prox-linear [13, 27, 28, 51, 54] and Survey Descent [30]—which could

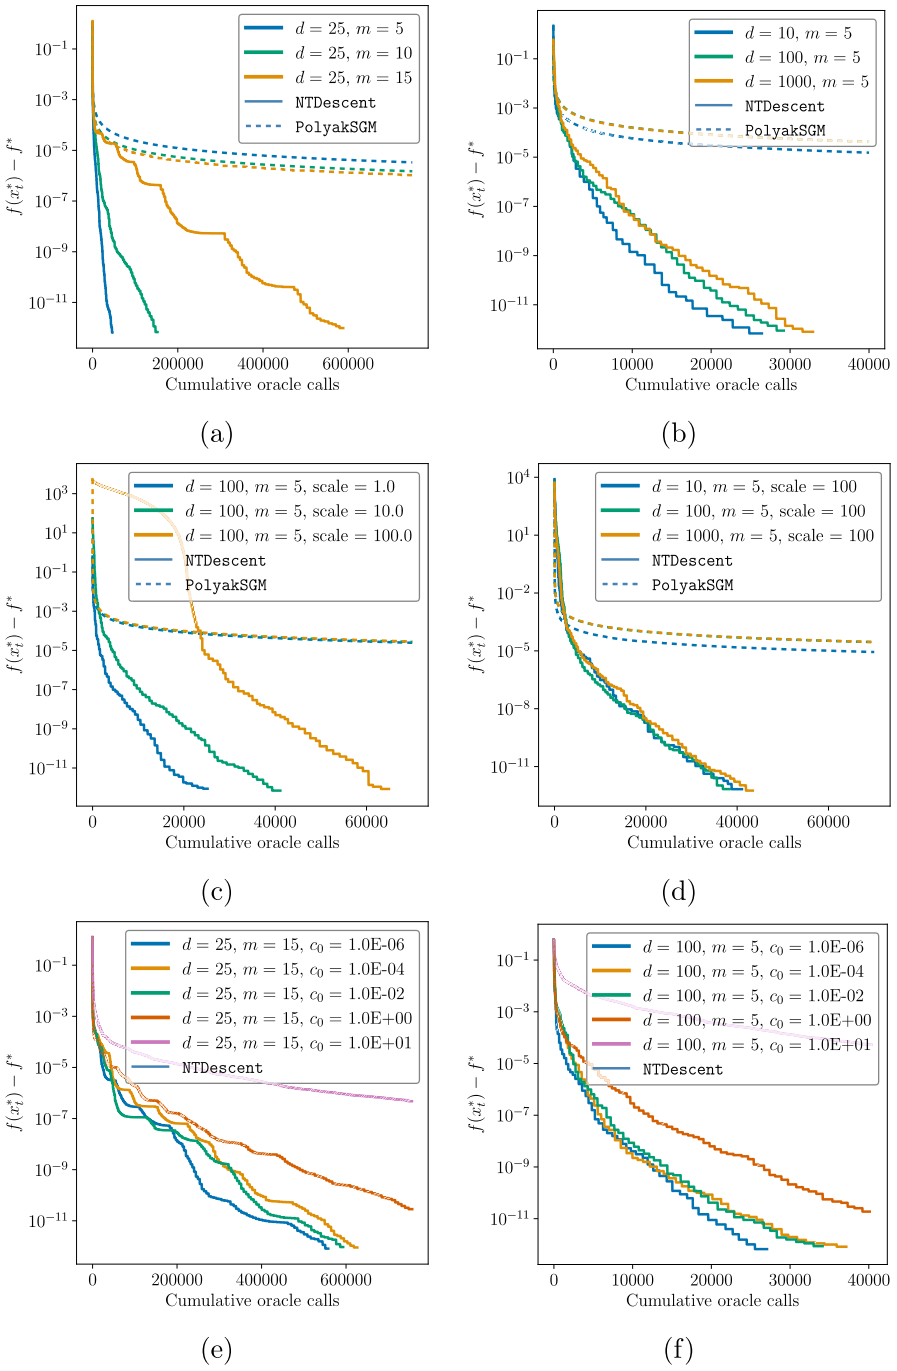

Fig. 6 Comparison of `NTDescent` with `PolyakSGM` on (7.1). For both algorithms, the value $f(x_t^*)$ denotes the best function seen after t oracle evaluations. See text for description

be applied to this problem. In order to explain these algorithms, let us write $f = \max_{i=1,\dots,m}\{f_i\}$, where the f_i are the quadratic function from (7.1).

Prox-linear method. Given a point $x \in \mathbb{R}^d$, the Prox-linear update x_+ solves

$$x_+ = \operatorname*{argmin}_{y\in\mathbb{R}^d} \max_{i=1,\dots,m} \{f_i(x) + \langle \nabla f_i(x), y - x\rangle\} + \frac{\rho}{2}\|y - x\|^2.$$

One may show that x_+ geometrically improves on x; see [25]. However, in contrast to `NTDescent`, the prox-linear method requires that the components f_i are known. This is stronger than the first-order oracle model considered in this work. Thus, we do not compare `NTDescent` with prox-linear.

Survey Descent The Survey Descent method is a multi-point generalization of gradient descent, designed for max-of-smooth functions. Rather than maintaining a single iterate sequence, the Survey Descent maintains a *survey* S of points, meaning a collection of points $\{s_i\}_{i=1}^m$ at which f is differentiable. A single iteration of the Survey Descent method then aims to produce a new survey $S^+ = \{s_i^+\}_{i=1}^m$ satisfying

$$\begin{aligned} s_i^+ := \operatorname*{argmin}_{x\in\mathbb{R}^d} & \left\| x - \left(s_i - \frac{1}{L}\nabla f(s_i)\right)\right\|^2 \\ \text{subject to: } & f(s_j) + \langle \nabla f(s_j), x - s_j\rangle \\ & + \frac{L}{2}\|x - s_j\|^2 \le f(s_i) + \langle \nabla f(s_i), x - s_i\rangle \quad \forall j = i. \end{aligned}$$

Here, L is an upper bound on the Lipschitz constant of ∇f_i for all $i = 1, \dots, m$. In [30], Han and Lewis study linear convergence of Survey Descent on max-of-smooth functions under the conditions of Corollary 3.4. Given a survey S, they show that the updated survey S^+ geometrically improves on S (in an appropriate sense) whenever the following conditions are satisfied: (i) all elements of the survey S are near $\bar{x}$; (i) the survey S is *valid*, meaning there exists a permutation a on $[m]$ such that

$$f_{a(i)}(s_i) = f(s_i) \quad \text{and} \quad \partial f(s_i) = \{\nabla f_{a(i)}(s_i)\} \quad \text{for all } i = 1, \dots, m.$$

To estimate the number of components m and find a valid initial survey S sufficiently close to $\bar{x}$, Han and Lewis suggest an empirical procedure based on running a nonsmooth variant of BFGS [39] for several iterations. After running BFGS, they suggest to (i) compute an estimate $\hat{m}$ of m from a singular value decomposition of the computed gradients, and (ii) build the survey from $\hat{m}$ past iterates in such a way that the computed gradients form an affine independent set. From the numerical illustration in [30], Survey Descent performs well on several small problems. However, since the initialization procedure and implementation of Survey Descent are somewhat sophisticated, we leave a detailed comparison between `NTDescent` and Survey Descent and to future work.

7.2 A Matrix Sensing Problem

In this example, f takes the following form

$$f(X) = \frac{1}{n}\mathcal{A}(XX^T) - \mathcal{A}(M)\ _1,$$

where $M \in \mathbb{R}^{N\times N}$ is an unknown positive semidefinite matrix of rank r that we wish to recover from known linear measurements $\mathcal{A}(M)$; the linear operator $\mathcal{A} : \mathbb{R}^{N\times N} \to \mathbb{R}^n$ takes the form $Y \to (a_i^T Y a_i - b_i^T Y b_i)_{i=1}^n$, for $n \in N$, where $a_i, b_i \in \mathbb{R}^d$ are random vectors sampled from a standard multivariate normal distribution; and the decision variable is a tall and skinny matrix $X \in \mathbb{R}^{N\times r}$, where in general we allow $r = r$. This optimization problem appears in various signal processing applications and is known as *quadratic sensing* [15]. Note that this objective does not satisfy Assumption A, since the solution set is not isolated.

We consider two settings in this section: the exact setting $r = r$ and the overparameterized setting $r > r$. In the exact setting [14] showed that if $n = Nr$(), the objective f is sharp, meaning $f(x) = \mathrm{dist}(x, \operatorname{argmin} f))$ and that `PolyakSGM` converges linearly whenever the initial iterate is sufficiently close to the set of minimizers. In the overparameterized setting, we are not aware of similar guarantees. Note that in practice, r is unknown, so the overparameterized setting is likely to be encountered.

In Fig. 7 we plot the performance of `NTDescent` and `PolyakSGM` in two experiments. In Fig. 7a we use base dimensions $N = 100$, optimal rank $r = 5$, and varying overparameterization $r \in \{r, r + 2, r + 5\}$. In Fig. 7b we use base dimensions $N = 100$, varying optimal rank $r \in \{5, 10, 15\}$, and fixed overparameterization $r = r + 5$. In both experiments, we fix $n = 4Nr$. Note that the dimension of the decision variable X is varying across each run since $d = Nr$ and r are varying. As can be seen from the plot, `PolyakSGM` outperforms `NTDescent` in the exact setting. This is to be expected, since f is a sharp function on which `PolyakSGM` is known to perform well. On the other hand, when $r > r$, we find that both methods slow down. However, `NTDescent` continues to converge nearly linearly, while `PolyakSGM` converges sublinearly.

7.3 An Eigenvalue Product Function

In this example, we aim to optimize a function $\tilde{f}$ that takes the following form

$$\tilde{f}(X) = \log E_K(A \,"\, XX),$$

where A is a fixed positive semi-definite data matrix, $E_K(Y)$ denotes the product of K largest eigenvalues of a symmetric matrix $Y \in \mathbb{S}^N$, and " denotes the Hadamard (entrywise) matrix product, subject to the constraint that X is positive semi-definite and its diagonal entries are 1. This example is a nonconvex relaxation of an entropy minimization problem arising in an environmental application [1, 12]. In our experiments, we choose A as in [1]: A is the leading $N \times N$ submatrix of a 63×63 covariance

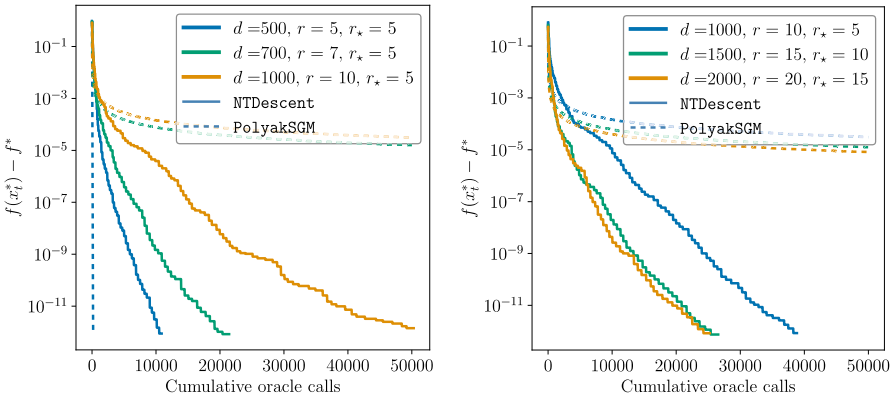

Fig. 7 Comparison of `NTDescent` with `PolyakSGM` on (7.1). In both plots, the base dimension is $N = 100$. Left: fixed optimal rank $r = 5$ and varying overparameterization $r \in \{5, 7, 10\}$; Right: varying $r \in \{5, 10, 15\}$, fixed overparameterization $r = r + 5$. For both algorithms, the value $f(x_t^*)$ denotes the best function seen after t oracle evaluations

matrix, scaled so that the largest entry is 1. As suggested by [12], we reformulate this problem as an unconstrained optimization problem using a Burer-Monteiro type factorization

$$\min_{V \in \mathbb{R}^{N \times N}} f(V) = \tilde{f}(c(V)c(V)^*), \tag{7.2}$$

where $c : \mathbb{R}^{N \times N} \to \mathbb{S}^N$ satisfies $c(V) = \mathrm{Diag}([\mathrm{diag}(VV^*)]^{-1/2})V$ for all $V \in \mathbb{R}^{N \times N}$. Here, the mapping $\mathrm{diag}(\cdot)$ takes a matrix an $N \times N$ matrix A to the N dimensional vector with i th entry A_{ii}. On the other hand, the mapping $\mathrm{Diag}(\cdot)$ takes an N dimensional vector v to the $N \times N$ diagonal matrix with i th diagonal entry v_i. A formula for the subgradient of f may be found [12]. We do not attempt to verify that f satisfies the full Assumption A. Instead, we point out that under a "transversality condition," function f admits an active manifold at local minimizers [39].

Turning to the experiment, we consider the case where $N = 14$ and $K = 7$. In this example, the optimal function value $\inf f$ is not known. Thus, we run `NTDescent` from four random initial starting points. We terminate each run of `NTDescent` when a certain "optimality gap" R_k satisfies $R_k \le 10^{-12}$. We denote the minimal function value achieved across all four runs by $\bar{f}$. Let us now define and motivate the optimality gap. For iteration k in Algorithm 4, define

$$R_k := \min \max\{\sigma_i^{(k)}, \|v_{i+1}^{(k)}\|^2\} : \sigma_i^{(k)} \le \|v_{i+1}^{(k)}\|,$$

where $\sigma_i^{(k)}$ and $v_{i+1}^{(k)}$ are computed in Lines 2 through 6 of Algorithm 3 at iteration k. Provided that x_k is sufficiently close to a point $\bar{x}$ at which function f satisfies Assumption A, it is possible to show that R_k satisfies $f(x_k) - f(\bar{x}) \le R_k$. This is illustrated in Fig. 8a: there, the optimality gap closely tracks the estimated function

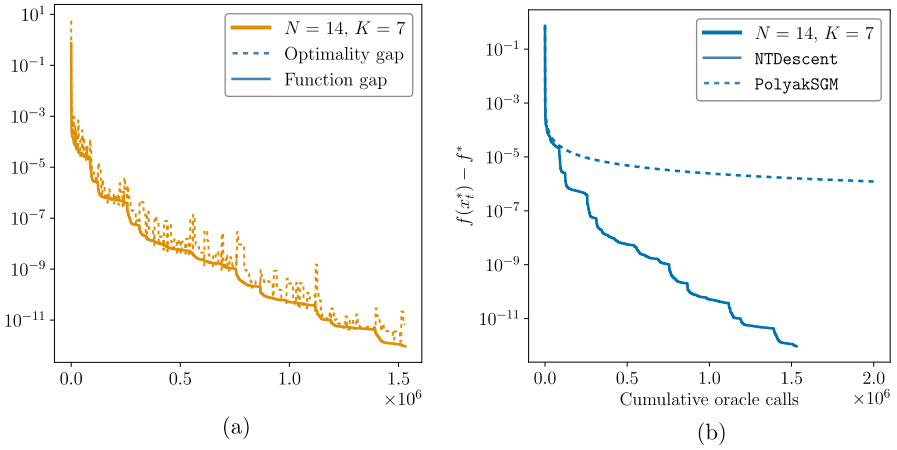

Fig. 8 Numerical performance on (7.2). Left: the close relationship between "optimality gap" and function gap; Right: comparison of `PolyakSGM` and `NTDescent` from three initial starting points. For both algorithms, the value $f(x_t^*)$ denotes the best function seen after t oracle evaluations. See text for detail

gap, when approximating by $\inf f$ by f^*. In Fig. 8b, we compare the performance of `NTDescent` on the three runs which did not achieve function value f^* before termination. In all three cases, we see similar performance. Next, for each run of `NTDescent`, we also run `PolyakSGM` from the same initial starting point, estimating $\inf f$ by f^*. We see that `NTDescent` outperforms `PolyakSGM`.

A Proof of Lemma 2.2

Let g denote the minimal norm element of $\partial_\sigma f(x)$. Write g as a convex combination of subgradients: $g = \sum_{i=1}^{n} \lambda_i g_i$ where $\sum_{i=1}^{n} \lambda_i = 1$ and $g_i \in \partial f(x_i)$ for some $x_i \in B_\sigma(x)$ and $n > 0$. Then

$$\begin{aligned} f(x) &\leq f\left(\sum_{i=1}^{n} \lambda_i x_i + \sum_{i=1}^{n} \lambda_i (x - x_i)\right) \\ &\leq \sum_{i=1}^{n} \lambda_i f(x_i) + L\sigma \\ &\leq f(y) + \sum_{i=1}^{n} \lambda_i \langle g_i, x_i - y \rangle + L\sigma \\ &\leq f(y) + \langle g, x - y \rangle + \sum_{i=1}^{n} \lambda_i \langle g_i, x_i - x \rangle + L\sigma \\ &\leq f(y) + \operatorname{dist}(0, \partial_\sigma f(x)) \|x - y\| + 2L\sigma, \end{aligned}$$

as desired.

B Proof of Proposition 3.5

We begin with preliminary notation and bounds. First, since M is C^4 smooth, the projection P_M is C^3 smooth near $\bar{x}$. Second, since f is C^3 smooth along M near $\bar{x}$, the composition $f_M := f \circ P_M$ is also C^3 smooth near $\bar{x}$. Third, the constant μ is positive due to the active manifold assumption. Fourth, choose $\delta > 0$ small enough that the following hold:

1. ∇P_M is C_M -Lipschitz on $B_\delta(\bar{x})$;
2. ∇f_M is β-Lipschitz on $B_\delta(\bar{x})$;
3. $\nabla^2 f_M$ is ρ-Lipschitz on $B_\delta(\bar{x})$ in the operator norm, where $\rho := 2\mathrm{lip}^{\mathrm{op}}_{\nabla^2 f_M}(\bar{x})$;
4. f is L-Lipschitz on $B_\delta(\bar{x})$;
5. the quadratic growth bound (A1) holds:

$$f(x) - f(\bar{x}) \geq \frac{\gamma}{2}\|x - \bar{x}\|^2 \qquad \text{for all } x \in \bar{B}_\delta(\bar{x});$$

6. the strong (a) bound (A3) holds:

$$P_{T_M(y)}(v - \nabla_M f(y)) \leq C_{(a)}\|x - y\| \tag{B.1}$$

for all $x \in \bar{B}_\delta(\bar{x})$, $v \in \partial f(x)$, and $y \in M \cap \bar{B}_\delta(\bar{x})$.

7. the $(b_\leq)$ regularity bound (A4) holds:

$$f(x') \geq f(x) + \langle v, x' - x\rangle - \frac{\mu}{2}\|x' - x\|^2 \tag{B.2}$$

for all $x \in B_\delta(\bar{x})$, $v \in \partial f(x)$, and $x' \in B_\delta(\bar{x}) \cap M$.

8. the sharpness condition holds:

$$\mathrm{dist}(0, \partial f(x)) > 2\mu \qquad \text{for all } x \in B_\delta(\bar{x}) \backslash M.$$

Given these bounds, let us define

$$\delta_A := \frac{1}{2}\min\left\{\delta, \frac{9\gamma}{16\rho}, \frac{\mu}{2(C_{(a)} + 2\beta + 2C_M L)}\right\}.$$

For this choice of δ_A, Item 1 holds automatically. We now prove the remaining items.

B.1 Item 2: Smoothness of P_M

Fix $x \in B_{2\delta_A}(\bar{x})$ and $x' \in B_{\delta_A}(x)$. Observe that $P_M(x) \in B_{2\delta_A}(\bar{x})$ and we have the inclusion $x - P_M(x) \in N_M(P_M(x))$. Consequently, we have

1. $P_{T_M(P_M(x))}(x) = P_{T_M(P_M(x))}(P_M(x))$;
2. $P_M(x) = P_M(P_M(x))$;
3. $\nabla P_M(P_M(x)) = P_{T_M(P_M(x))}$.

Therefore, we have

$$\begin{aligned} & P_M(x) - P_M(\bar{x}) - P_{T_M(P_M(\bar{x}))}(x - \bar{x}) \\ &= P_M(x) - P_M(P_M(\bar{x})) - \nabla P_M(P_M(\bar{x}))(x - P_M(\bar{x})) \\ &\le \frac{C_M}{2} \|x - P_M(\bar{x})\|^2 \\ &\le C_M (\|x - \bar{x}\|^2 + \operatorname{dist}^2(x, M)), \end{aligned}$$

where the first inequality follows from Lipschitz continuity of ∇P_M on $B_{2\delta_A}(\bar{x}) \subseteq B_\delta(\bar{x})$.

B.2 Item 3: Bounds on $\nabla_M f$

Recall that $P_M(x) \in B_{2\delta_A}(\bar{x})$ whenever $x \in B_{\delta_A}(\bar{x})$. Thus, below we prove that

$$\frac{\gamma}{2}\|y - \bar{x}\| \le \|\nabla f_M(y)\| \le \beta \|y - \bar{x}\| \qquad \text{for all } y \in B_{2\delta_A}(\bar{x}) \cap M.$$

This is equivalent to the claimed bound since $\nabla f_M(y) = \nabla_M f(y)$ for all $y \in B_{2\delta_A}(\bar{x}) \cap M$.

Let us first prove the claimed upper bound. Due to the inequality,

$$f_M(x) - f_M(\bar{x}) \ge \frac{\gamma}{2}\|P_M(x) - \bar{x}\|^2 \qquad \text{for all } x \in B_\delta(\bar{x}),$$

it follows that $\bar{x}$ is a local minimizer of f_M. Consequently, $\nabla f_M(\bar{x}) = 0$. Thus, since β is a local Lipschitz constant of ∇f_M on $B_\delta(\bar{x})$, we have

$$\|\nabla f_M(y)\| \le \beta \|y - \bar{x}\| \qquad \text{for all } y \in B_\delta(\bar{x}) \cap M.$$

Since $2\delta_A \le \delta$, this proves the claimed upper bound.

Next, we prove the claimed lower bound. It suffices to establish the following convexity inequality:

$$f_M(y) + \langle \nabla f_M(y), \bar{x} - y \rangle \le f_M(\bar{x}) \qquad \text{for all } y \in B_{2\delta_A}(\bar{x}) \cap M. \tag{B.3}$$

Indeed, if this inequality holds, we have

$$\langle \nabla f_M(y), y - \bar{x} \rangle \ge f_M(y) - f_M(\bar{x}) \ge \frac{\gamma}{2}\|y - \bar{x}\|^2 \qquad \text{for all } y \in B_{2\delta_A}(\bar{x}) \cap M,$$

and the desired result follows from Cauchy-Schwarz.

To that end, observe that since $\nabla f_M(\bar{x}) = 0$ and $\nabla^2 f_M$ is ρ-Lipschitz in $B_{2\delta_A}(\bar{x})$, we have

$$f_M(y) \le f_M(\bar{x}) + \frac{1}{2}\langle \nabla^2 f_M(\bar{x})(y - \bar{x}), y - \bar{x} \rangle + \frac{\rho}{6}\|y - \bar{x}\|^3 \qquad \text{for all } y \in B_{2\delta_A}(\bar{x}).$$

Consequently, we have the lower bound on the quadratic form: for all $y \in B_{2\delta_A}(\bar{x}) \cap M$, we have

$$\begin{aligned}\frac{1}{2}\left\langle \nabla^2 f_M(\bar{x})(y-\bar{x}), (y-\bar{x})\right\rangle &\geq f_M(y) - f_M(\bar{x}) - \frac{\rho}{6}\|y-\bar{x}\|^3 \\ &\geq \frac{\gamma}{2}\|y-\bar{x}\|^2 - \frac{\rho}{6}\|y-\bar{x}\|^3 \\ &\geq \frac{3\gamma}{8}\|y-\bar{x}\|^2, \end{aligned} \tag{B.4}$$

where the second inequality follows from the quadratic growth bound and the third follows from the bound $\|y-\bar{x}\| \leq 2\delta_A \leq \frac{3\gamma}{4\rho}$. Therefore, for all $y \in M \cap B_{2\delta_A}(\bar{x})$, we have

$$\begin{aligned} f_M(\bar{x}) &\geq f_M(y) + \langle \nabla f_M(y), \bar{x}-y\rangle + \frac{1}{2}\langle \nabla^2 f_M(y)(\bar{x}-y), (\bar{x}-y)\rangle - \frac{\rho}{6}\|y-\bar{x}\|^3 \\ &\geq f_M(y) + \langle \nabla f_M(y), \bar{x}-y\rangle + \frac{1}{2}\langle \nabla^2 f_M(\bar{x})(\bar{x}-y), (\bar{x}-y)\rangle - \frac{2\rho}{3}\|y-\bar{x}\|^3 \\ &\geq f_M(y) + \langle \nabla f_M(y), \bar{x}-y\rangle + \frac{3\gamma}{8}\|y-\bar{x}\|^2 - \frac{2\rho}{3}\|y-\bar{x}\|^3 \\ &\geq f_M(y) + \langle \nabla f_M(y), \bar{x}-y\rangle, \end{aligned}$$

where the first and second inequalities follow by Lipschitz continuity of $\nabla^2 f_M$; the third inequality follows from (B.4); and the fourth inequality follows from the bound $\|y-\bar{x}\| \leq 2\delta_A \leq \frac{9\gamma}{16\rho}$. This completes the proof.

B.3 Item 4: Consequences of Strong (a)-Regularity

Fix $x \in B_{\delta_A}(\bar{x})$ and $\sigma \leq \delta_A$. Recall that $y := P_M(x) \in B_{2\delta_A}(\bar{x})$ since $x \in B_{\delta_A}(\bar{x})$. Fix $g \in \partial_\sigma f(x)$. By definition of $\partial_\sigma f(x)$, there exists a family of coefficients $\lambda_i \in [0,1]$, points $x_i \in \overline{B}_\sigma(x) \subseteq \overline{B}_\delta(\bar{x})$, and subgradients $g_i \in \partial f(x_i)$ indexed by a finite set $i \in I$ such that $\sum_{i\in I}\lambda_i = 1$ and $g = \sum_{i\in I}\lambda_i g_i$. Therefore, by averaging the strong (a) bound (B.1) over g_i, we find that

$$\begin{aligned} P_{T_M(y)}(g - \nabla_M f(y)) &\leq \sum_{i\in I}\lambda_i \, P_{T_M(y)}(g_i - \nabla_M f(y)) \\ &\leq \sum_{i\in I}\lambda_i C_{(a)}\|x_i - y\|. \\ &\leq C_{(a)}(\operatorname{dist}(x, M) + \sigma). \end{aligned}$$

Since g was arbitrary, it follows that for all $x \in B_{\delta_A}(\bar{x})$ and $\sigma \leq \delta_A$, we have

$$\sup_{g\in\partial_\sigma f(x)} \|P_{T_M(y)}(g - \nabla_M f(y))\| \leq C_{(a)}(\operatorname{dist}(x, M) + \sigma). \tag{B.5}$$

Now we apply this bound to establish the two remaining inequalities.

Indeed, first observe that for all $x \in B_{\delta_A}(\bar{x})$ and $\sigma \le \delta_A$, we have

$$\begin{aligned}\sup_{g \in \partial_\sigma f(x)} \|P_{T_M(y)} g\| &\le \|\nabla_M f(y)\| + C_{(a)}(\mathrm{dist}(x, M) + \sigma) \\ &\le \beta \|y - \bar{x}\| + C_{(a)}(\mathrm{dist}(x, M) + \sigma),\end{aligned}$$

where the first inequality follows from (B.5) and the second inequality follows from Item 3. This proves the first claimed bound. Second, observe that for all $x \in B_{\delta_A}(\bar{x})$ and $\sigma \le \delta_A$, we have

$$\begin{aligned}\sup_{g, g' \in \partial_\sigma f(x)} \|P_{T_M(y)}(g - g')\| &\le \sup_{g \in \partial_\sigma f(x)} \|P_{T_M(y)}(g - \nabla_M f(y))\| \\ &\quad + \sup_{g' \in \partial_\sigma f(x)} \|P_{T_M(y)}(g' - \nabla_M f(y))\| \\ &\le 2C_{(a)}(\mathrm{dist}(x, M) + \sigma).\end{aligned}$$

where the second inequality follows from (B.5). This completes the proof.

B.4 Item 5: Aiming Inequality

Consider a point $x \in B_{\delta_A}(\bar{x})$, let $\kappa = 2\mu$, and define

$$\hat{x} \in \operatorname*{argmin}_{x' \in \overline{B}_{2\delta_A}(\bar{x})} f(x') + \kappa \|x' - x\|^2.$$

We claim that $\hat{x} \in M \cap B_{2\delta_A}(\bar{x})$. Indeed, first note that by definition of $\hat{x}$ and the inclusion $\hat{x} \in \overline{B}_{2\delta_A}(\bar{x})$, we have

$$\|\hat{x} - x\| \le \frac{f(\bar{x}) - f(\hat{x})}{\kappa} + \|\bar{x} - x\| \le \|\bar{x} - x\| < \delta_A,$$

where the second inequality follows since $\bar{x}$ is a minimizer of f on $B_{2\delta_A}(\bar{x})$, a consequence of quadratic growth. Thus, by the triangle inequality, we have $\hat{x} \in B_{2\delta_A}(\bar{x})$. By Fermat's rule, we, therefore, have the inclusion:

$$0 \in \partial(f + \kappa \|\cdot - x\|^2)(\hat{x}) \subseteq \partial f(\hat{x}) + \kappa \overline{B}.$$

If $\hat{x} \notin M$, then $\mathrm{dist}(0, \partial f(\hat{x})) > \kappa$, contradicting the above inclusion. Therefore, we have $\hat{x} \in M \cap B_{2\delta_A}(\bar{x})$.

Turning to the aiming inequality, apply the $(b_\le)$-regularity bound (B.2) to $\hat{x}$:

$$f(\hat{x}) \ge f(x) + \langle v, \hat{x} - x \rangle - \varepsilon \|x - \hat{x}\|^2 \ge f(\hat{x}) + \langle v, \hat{x} - x \rangle + (\kappa - \varepsilon) \|x - \hat{x}\|^2,$$

where we define $\varepsilon := \mu/2$. Consequently, we have

$$\langle v, x - P_M(x) \rangle \ge (\kappa - \varepsilon) \|x - \hat{x}\|^2 + \langle v, \hat{x} - P_M(x) \rangle \quad \text{for all } v \in \partial f(x). \tag{B.6}$$

We now bound the term $\langle v, \hat{x} - P_M(x) \rangle$: By the conclusion of Item 2, we have

$$\begin{aligned} &P_M(\hat{x}) - P_M(x) - P_{T_M(P_M(x))}(\hat{x} - x) \\ &\leq C_M(\|x - \hat{x}\|^2 + \operatorname{dist}^2(x, M)) \leq 2C_M\|x - \hat{x}\|^2, \end{aligned}$$

where the second inequality follows since $\hat{x} \in M$. Thus, we have

$$\begin{aligned} |\langle v, \hat{x} - P_M(x)\rangle| &\leq |\langle v, P_{T_M(P_M(x))}(\hat{x} - x)\rangle| + 2C_M\|v\|\|x - \hat{x}\|^2 \\ &\leq \|P_{T_M(P_M(x))}v\|\|x - x\| + 2C_M L\|x - \hat{x}\|^2 \\ &\leq (C_{(a)}\operatorname{dist}(x, M) + \beta\|P_M(x) - \bar{x}\|)\|x - x\| + 2C_M L\|x - \hat{x}\|^2 \\ &\leq (C_{(a)}\delta_A + 2\beta\delta_A + 2C_M L\delta_A)\|x - x\| \\ &\leq \varepsilon\|x - x\|. \end{aligned}$$

where the second inequality follows from Item 4 and the third inequality follows from the inclusion $P_M(x) \in B_{2\delta_A}(\bar{x})$. Therefore, plugging this bound into (B.6), we arrive at

$$\langle v, x - P_M(x)\rangle \geq (\kappa - 2\varepsilon)\|x - \hat{x}\| \geq \mu \operatorname{dist}(x, M),$$

as desired.

B.5 Item 6: Bounding Subgradients

Fix $x \in B_{\delta_A}(\bar{x})$, $\sigma \leq \delta_A$, and $g \in \partial_\sigma f(x)$. By definition of $\partial_\sigma f(x)$, there exists a family of coefficients $\lambda_i \in [0, 1]$, points $x_i \in \overline{B}_\sigma(x) \subseteq \overline{B}_\delta(\bar{x})$, and subgradients $g_i \in \partial f(x_i)$ indexed by a finite set $i \in I$ such that $\sum_{i\in I}\lambda_i = 1$ and $g = \sum_{i\in I}\lambda_i g_i$. Recall that by Lipschitz continuity of f on $B_\delta(\bar{x})$, we have $\|g_i\| \leq L$ for $i \in I$. Therefore,

$$\|g\| \leq \sum_{i\in I}\lambda_i\|g_i\| \leq L,$$

as desired.

B.6 Item 7: Bounding the Function Gap

Fix a point $x \in B_{\delta_A}(\bar{x})$ and recall that $P_M(x) \in B_{2\delta_A}(\bar{x})$. Then by Lipschitz continuity of f on $B_\delta(\bar{x})$, we have

$$f(x) - f(P_M(\bar{x})) \leq L\operatorname{dist}(x, M).$$

Next, arguing as in the proof of Item 3, we find that $\nabla f_M(\bar{x}) = 0$. Thus, since ∇f_M is β-Lipschitz on $B_\delta(\bar{x})$, we have

$$f(P_M(x)) - f(\bar{x}) = f_M(P_M(x)) - f(\bar{x}) \le \langle \nabla f_M(\bar{x}), P_M(x) - \bar{x} \rangle + \frac{\beta}{2}\|P_M(x) - \bar{x}\|^2 = \frac{\beta}{2}\|P_M(x) - \bar{x}\|^2.$$

By putting both bounds together, we have

$$f(x) - f(\bar{x}) = f(x) - f(P_M(x)) + f(P_M(x)) - f(\bar{x}) \le L\,\mathrm{dist}(x, M) + \frac{\beta}{2}\|P_M(x) - \bar{x}\|^2,$$

as desired.

C Proof of Corollary 2.5

We begin with the following known Lemma, which immediately follows from [29, Proposition 2.8]

Lemma C.1 *Let $f : \mathbb{R}^d \to \mathbb{R}$ be a locally Lipschitz function. Suppose that there exists sequences $x_k \to \bar{x}, \tau_k \to 0$, and $g_k \in \partial_{\tau_k} f(x_k)$ with $g_k \to 0$. Then $\bar{x}$ is a Clarke critical point.*

Now we turn to the proof of the Corollary. Since f has bounded initial sublevel set, the following widened sublevel set is bounded:

$$S := \{x + u : f(x) \le f(x_0) \text{ and } u \in \overline{B}(x)\}.$$

Thus, there exists $L > 0$ such that f is L-Lipschitz on S. In addition, ∂f is uniformly bounded by L on int S.

We begin with a claim.

Claim: Fix $i > 0$ and define $\tau_i := 2^{-i}$. Let $g_k := \max\{\|g_k\|, c_0 g_0\}$ be the trust region parameter used in Algorithm 3 and define $\varepsilon_{i,k} := \frac{g_k}{128 L \tau_i}$. Then with probability one, the event

$$E_k^{(i)} = \mathrm{dist}(0, \partial_{\tau_i} f(x_k)) > \varepsilon_{i,k} \text{ and } f(x_{k+1}) > f(x_k) - \frac{\tau_i \mathrm{dist}(0, \partial_{\tau_i} f(x_k))}{8}$$

cannot happen infinitely often, i.e.,

$$P\left(\cap_{T=1}^{\infty} \cup_{k=T}^{\infty} E_k^{(i)}\right) = 0.$$

Proof We prove that $P(E_k^{(i)})$ is summable in k. Indeed, first, note that $P(E_k^{(i)}) = 0$ when $P(\mathrm{dist}(0, \partial_{\tau_i} f(x_k)) > \varepsilon_{i,k}) = 0$. On the other hand, suppose

$P(\operatorname{dist}(0, \partial_{\tau_i} f(x_k)) > \quad_{i,k}) > 0$. Now we upper bound $P\ (E_k^{(i)})$ for all G_k satisfying $G_k \geq i$. For such $G := G_k$, the radius $\tau_i = \sigma_{G-i}$ is among those considered in Algorithm 3. Moreover, since $s_k \leq L$ (recall $x_k \in \operatorname{int}(S)$), the radius satisfies the trust region constraint: $\sigma_{G-i} = \tau_i \leq \quad_{i,k}/s_k \leq \operatorname{dist}(0, \partial_{\sigma_{G-i}} f(x))/s_k$. Therefore, if `NDescent` terminates with descent at the $(G - i)$-th level in Algorithm 3, it follows that

$$f(x_{k+1}) > f(x_k) - \frac{\tau_i \operatorname{dist}(0, \partial_{\tau_i} f(x_k))}{8}.$$

We estimate the probability of this success with Lemma 2.3: there exist $\in 0$ depending on $\quad_{i,k}$ and for all $k \geq i$, we have

$$P(E_k^{(i)}) \leq P \quad f(x_{k+1}) > f(x_k) - \frac{\tau_i \operatorname{dist}(0, \partial_{\tau_i} f(x_k))}{8} \| \operatorname{dist}(0, \partial_{\tau_i} f(x_k)) > \quad_{i,k}$$
$$\leq \exp(-Ck).$$

Therefore, $P(E_k^{(i)})$ is summable in k. The result then follows from Borel–Cantelli lemma.

By the claim and a union bound, we know that with probability one, for any fixed i, $E_k^{(i)}$ cannot happen infinitely often. Now, suppose that a subsequence $\{x_{k_l}\}$ (where $k_l \geq l$ is strictly increasing in l) converges to a point $\bar{x}$. We note that the sequence $\{f(x_k)\}$ is bounded below: Indeed, since x_{k_l} converges and f is continuous, it follows $\{f(x_{k_l})\}$ is bounded below by a constant $c \in \mathbb{R}$. Consequently, since $\{f(x_k))\}$ is nonincreasing and $k_l \geq l$, it follows that $c \leq f(x_{k_l}) \leq f(x_l)$ and for every $l > 0$, as desired. As a result, the following inequalities cannot be valid simultaneously infinitely often:

$$\operatorname{dist}(0, \partial_{\tau_i} f(x_{k_l})) > \quad_{i,k} \text{ and } f\ (x_{k_l+1}) \leq f(x_{k_l}) - \frac{\tau_i \operatorname{dist}(0, \partial_{\tau_i} f(x_{k_l}))}{8}.$$

Therefore, $\operatorname{dist}(0, \partial_{\tau_i} f(x_{k_l})) > \quad_{i,k}$ cannot happen infinitely often. Consequently, we can find a sequence of increasing indices j_i such that

$$\operatorname{dist}(0, \partial_{\tau_i} f(x_{j_i})) \leq \quad_{i,k} \qquad \text{and } x_{j_i} \to \bar{x}.$$

Since $\quad_{i,k} \to 0$ as $k \to \infty$, Lemma C.1, shows that $\bar{x}$ is Clarke critical.

D Proof of Lemma 5.7

We begin with preliminary notation and bounds. We fix $x \in B_{\delta_{\text{Grid}}}(\bar{x})$ and subgradient $g \in \partial_\sigma f(x) \backslash \{0\}$. We define $y := P_M(x)$, $T := T_M(y)$, and $N := N_M(y)$. We have the following two bounds: First, we have

$$(\mu + L) C_M (D_1 \operatorname{dist}(x, M) + \sigma) \leq (\mu + \quad L) C_M \delta_{\text{Grid}} (D_1 + 1)$$

$$= \frac{\mu}{8} C_M \left(D_1^{-1} + 1\right)\delta_{\text{Grid}} \leq \frac{\mu}{8}. \tag{D.1}$$

Second, we have

$$C_{(a)}(\operatorname{dist}(x, M) + \sigma) + \beta \quad y - \bar{x} \leq \quad 2C_{(a)}\delta_{\text{Grid}} + 2\beta\delta_{\text{Grid}} \leq \frac{\mu}{4}. \tag{D.2}$$

We now turn to the proof.

By Lemma 5.3 (which is applicable since $x \in B_{\delta_A/2}(\bar{x})$ and $\sigma \leq \delta_{\text{Grid}} \leq \delta_A/2$), we have

$$\hat{g}, \sigma \frac{P_N g}{g}^{!} \leq -\sigma\mu \quad \frac{P_N g}{g} + (\mu + \quad L)\operatorname{dist}(x, M) + (\mu + \quad L)C_M \left(\operatorname{dist}^2(x, M) + \sigma^{\ 2}\right).$$

Rearranging, we find that

$$\begin{aligned} P_N\hat{g}, g \leq -\mu \quad P_N g + \quad & \frac{(\mu + \ L)\ g\ \operatorname{dist}(x, M)}{\sigma} + \frac{(\mu + \ L)\ g\ C_M\ (\operatorname{dist}^2(x, M) + \sigma^{\ 2})}{\sigma} \\ \leq -\mu \quad P_N g + \quad & \frac{\mu}{8}\ g + (\mu + \quad L)C_M\ (D_1 \operatorname{dist}(x, M) + \sigma) \cdot \quad g \\ \leq -\mu \quad P_N g + \quad & \frac{\mu}{4}\ g, \end{aligned}$$

where the second inequality follows from the assumption $D_1^{-1}\operatorname{dist}(x, M) \leq \sigma$ and the third follows from (D.1). Now observe that

$$P_T\hat{g}, g \leq \quad P_T\hat{g} \quad g \leq (\ C_{(a)}(\operatorname{dist}(x, M) + \sigma) + \beta \quad y - \bar{x}) \cdot \quad g \leq \quad \frac{\mu}{4}\ g,$$

where second inequality follows from (3.4) and the third inequality follows from (D.2). Therefore,

$$\hat{g}, g \ = \ P_N\hat{g}, g \ + \ P_T\hat{g}, g \leq -\mu \quad P_N g + \ \frac{\mu}{2}\ g \leq - \ \frac{\mu}{2}\ g + \mu \quad P_T(g),$$

as desired.

E Proof that $\leq L$

Lemma E.1 *We have that* $\mu \leq L$.

Proof Indeed,

$$\mu = \ \frac{1}{4} \liminf_{x \xrightarrow{M} \bar{x}} \operatorname{dist}(0, \partial f(x)) \leq \ \limsup_{x \to \bar{x}} \operatorname{dist}(0, \partial f(x)) \leq \ L.$$

by Proposition 3.5.

F Proof of Lemma 6.4

We fix $a > 0$. Note that the claimed inclusion is a consequence of the following bound:

$$f(x) - f(\bar{x}) \geq \frac{\gamma}{2} \min\{\delta_A, \|x - \bar{x}\|\} \|x - \bar{x}\| \qquad \text{for all } x \in \mathbb{R}^d. \tag{F.1}$$

Here we provide a proof for completeness.

To that end, we remind the reader that Assumption A is in force. Consequently, by Item 1 of Proposition 3.5, we have:

$$f(x) - f(\bar{x}) \geq \frac{\gamma}{2} \|x - \bar{x}\|^2 \qquad \text{for all } x \in \overline{B}_{\delta_A}(\bar{x}).$$

Thus, if $x \in B_{\delta_A}(\bar{x})$, bound (F.1) is immediate. On the other hand, suppose that we have $x \in \mathbb{R}^d \setminus B_{\delta_A}(\bar{x})$. Define the curve $x_t : t \to (1-t)x + t\bar{x}$. Choose $t_0 \in [0, 1]$ such that $x_{t_0} \in \operatorname{bdry} B_{\delta_A}(\bar{x})$. Then by Jensen's inequality, we have

$$\begin{aligned}(1 - t_0) f(x) &\geq f(x_{t_0}) - t_0 f(\bar{x}) \geq (1 - t_0) f(\bar{x}) + \frac{\gamma}{2} \|x_{t_0} - \bar{x}\|^2 \\ &= (1 - t_0) f(\bar{x}) + \frac{\gamma (1 - t_0)}{2} \|x - \bar{x}\| \|x_{t_0} - \bar{x}\|.\end{aligned}$$

Consequently, since $\|x_{t_0} - \bar{x}\| = \delta_A$, we have

$$f(x) - f(\bar{x}) \geq \frac{\gamma \delta_A}{2} \|x - \bar{x}\| \geq \frac{\gamma}{2} \min\{\delta_A, \|x - \bar{x}\|\} \|x - \bar{x}\|,$$

as desired. This completes the proof.

G Proof of (6.18)

Let us expand the left-hand-side of (6.18):

$$\frac{16L\sqrt{2\log(2K_1^2/p)}}{K_1^{1/2}} \leq \underbrace{\frac{16DL\sqrt{2\log(K_1^2)}}{K_1^{1/2}}}_{=:A} + \underbrace{\frac{16DL\sqrt{2\log(2/p)}}{K_1^{1/2}}}_{=:B}.$$

Note that $B \leq a/4$ by definition of K_1. Consequently, the proof will follow if $A \leq a/4$. To that end, for any $\alpha \in (0, 1)$, we have

$$A = \frac{16DL\sqrt{2\log(K_1^2)}}{K_1^{1/2}} = \frac{16DL\sqrt{2\log(K_1^{2\alpha})/\alpha}}{K_1^{1/2}} \leq \frac{16DL\sqrt{2/\alpha}}{K_1^{(1-\alpha)/2}},$$

Therefore, we have $A \leq a/4$ whenever

$$K_1 \geq \inf_{\alpha\in(0,1)} \frac{64DL^{\frac{2}{\alpha}}}{a^{\frac{2}{(1-\alpha)}}} = \frac{D^2}{a^2} \inf_{\alpha\in(0,1)} \frac{64L^{\frac{2}{\alpha}}}{0_{\frac{a}{D}}1^{\frac{2\alpha}{(1-\alpha)}}} = \frac{D^2}{a^2} b.$$

This lower bound holds by definition of K_1. Consequently $A \leq a/4$. Therefore, the proof is complete.

H Proof of Lemma 6.6

Throughout this section, we use the symbol $a \lesssim b$ to mean that $a \leq \eta b$ for a fixed numerical constant η that is independent of f. In addition, we use the bound on the condition number: $\kappa \geq 1$, since $\mu \leq L$; see Lemma E.

Turning to the bound, we wish to upper bound q.

$$q = \max\left\{\rho, \sqrt{1 - \frac{3\mu^2}{256L^2}}, \frac{1}{2}\right\}.$$

First note that

$$1 - \sqrt{1 - \frac{3\mu^2}{256L^2}} \leq \frac{\mu^2}{L^2} \geq \frac{1}{\kappa^2}.$$

Next, we upper bound ρ. To that end, we must bound the constants a_1 and a_2, which rely on the somewhat involved constants C_4 and C_5. Thus, we first lower bound C_4:

$$\begin{aligned} C_4 &= \min\left\{\frac{\beta}{C_{(a)}(1+\delta_A)}, \frac{\min\{\mu/\delta_A, C_3D_2/\beta\}}{4(1+(1+\delta_A)C_M)(\mu+L))}, \frac{1}{2}\right. \\ &\quad \left.\min\left\{\frac{\beta}{C_{(a)}}, \frac{\mu}{L(1+C_M)}, \frac{\gamma^2\mu}{L^2\beta(1+C_M)}\right\}\right. \\ &\geq \frac{1}{\kappa^3(1+C_M)}, \end{aligned}$$

where we use the bounds $\mu \leq L$, $C_3\gamma^2/L$, and $D_2\mu$. Turning to C_5, we have:

$$\begin{aligned} C_5 &= \min\left\{\frac{\beta}{2C_{(a)}}, \frac{C_3D_2}{32C_{(a)}\beta}, C_4, \frac{C_2}{4}\right. \\ &\quad \left.\min\left\{\frac{\beta}{C_{(a)}}, \frac{\gamma^2\mu}{LC_{(a)}\beta}, \frac{1}{\kappa^3(1+C_M)}, \frac{\gamma}{C_{(a)}}\right\}\right. \\ &\geq \frac{1}{\kappa^3(1+C_M)}, \end{aligned}$$

where we again use $C_3 \quad \gamma \quad ^2/L$, and $D_2 \quad \mu \quad$. Therefore, we have the lower bound for a_2:

$$a_2 = \frac{\min\{C_1/L, C_5\}}{2} \quad \min \quad \frac{\gamma^2}{L^2}, \frac{1}{\kappa^3(1 + C_M)} \qquad \frac{1}{\kappa^3(1 + C_M)}.$$

In addition, we have the upper bound:

$$a_2 = \frac{\min\{C_1/L, C_5\}}{2} \leq C_4/2 \leq 1/4.$$

Finally to lower bound a_1, we have

$$a_1 = \min\{D_1, D_2/L\} \quad \min \quad \frac{\mu}{L}, \frac{\gamma}{L} \qquad \frac{1}{\kappa},$$

where we use the bound $D_1 \quad \mu/ \quad L$ and $D_2/L \quad \mu/ \quad L$.

Now we upper bound ρ by providing a lower bound on $1 - \rho$.

$$\begin{aligned} 1 - \rho = \quad & \frac{1}{8} \min \quad \frac{\gamma a_2}{8 \max\{4La_2^2, \beta\}}, \frac{\mu a_1}{4 \max\{2L, \beta/a_2^2\}} \\ & \min \quad \frac{\gamma}{La_2}, \frac{\gamma a_2}{\beta}, \frac{\mu a_1}{L}, \frac{\mu a_1 a_2^2}{\beta} \\ & \min \quad \frac{\gamma}{L}, \frac{\gamma}{\kappa^3(1 + C_M)\beta}, \frac{\mu}{\kappa L}, \frac{\mu}{\beta\kappa^7(1 + C_M)^2} \\ & \frac{1}{\kappa^8(1 + C_M)^2} \end{aligned}$$

Putting all these bounds together, we find that:

$$1 - q \quad \min \quad \frac{1}{\kappa^8(1 + C_M)^2}, \frac{1}{\kappa^2} \quad \geq \quad \frac{1}{\kappa^8(1 + C_M)^2},$$

as desired.

References

1. K. M. Anstreicher and J. Lee. A masked spectral bound for maximum-entropy sampling. In *mODa 7–Advances in Model-Oriented Design and Analysis*, pages 1–12. Springer, 2004.
2. H. Attouch and J. Bolte. On the convergence of the proximal algorithm for nonsmooth functions involving analytic features. *Mathematical Programming*, 116:5–16, 2009.
3. H. Attouch, J. Bolte, P. Redont, and A. Soubeyran. Proximal alternating minimization and projection methods for nonconvex problems: An approach based on the Kurdyka-Łojasiewicz inequality. *Mathematics of operations research*, 35(2):438–457, 2010.

4. H. Attouch, J. Bolte, and B. F. Svaiter. Convergence of descent methods for semi-algebraic and tame problems: proximal algorithms, forward–backward splitting, and regularized gauss–seidel methods. *Mathematical Programming*, 137(1):91–129, 2013.
5. P. Bianchi, W. Hachem, and S. Schechtman. Stochastic subgradient descent escapes active strict saddles. *arXiv preprint* arXiv:2108.02072, 2021.
6. J. Bolte, A. Daniilidis, and A. Lewis. The Łojasiewicz inequality for nonsmooth subanalytic functions with applications to subgradient dynamical systems. *SIAM J. Optim.*, 17(4):1205–1223 (electronic), 2006.
7. J. Bolte, A. Daniilidis, A. Lewis, and M. Shiota. Clarke subgradients of stratifiable functions. *SIAM Journal on Optimization*, 18(2):556–572, 2007.
8. J. Bolte, T. P. Nguyen, J. Peypouquet, and B. W. Suter. From error bounds to the complexity of first-order descent methods for convex functions. *Mathematical Programming*, 165:471–507, 2017.
9. J. Bolte, S. Sabach, and M. Teboulle. Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Mathematical Programming*, 146(1):459–494, 2014.
10. N. Boumal. An introduction to optimization on smooth manifolds. *Available online, Aug*, 2020.
11. S. Bubeck et al. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
12. S. Burer and J. Lee. Solving maximum-entropy sampling problems using factored masks. *Mathematical Programming*, 109(2):263–281, 2007.
13. J. V. Burke. Descent methods for composite nondifferentiable optimization problems. *Mathematical Programming*, 33(3):260–279, 1985.
14. V. Charisopoulos, Y. Chen, D. Davis, M. Díaz, L. Ding, and D. Drusvyatskiy. Low-rank matrix recovery with composite optimization: good conditioning and rapid convergence. *Foundations of Computational Mathematics*, 21(6):1505–1593, 2021.
15. Y. Chen, Y. Chi, and A. J. Goldsmith. Exact and stable covariance estimation from quadratic sampling via convex programming. *IEEE Transactions on Information Theory*, 61(7):4034–4059, 2015.
16. F. H. Clarke, Y. S. Ledyaev, R. J. Stern, and P. R. Wolenski. *Nonsmooth analysis and control theory*, volume 178. Springer Science & Business Media, 2008.
17. D. Davis and D. Drusvyatskiy. Subgradient methods under weak convexity and tame geometry. *SIAG/OPT Views and News*, 28(1):1–10, 2020.
18. D. Davis, D. Drusvyatskiy, and V. Charisopoulos. Stochastic algorithms with geometric step decay converge linearly on sharp functions. *arXiv preprint* arXiv:1907.09547, 2019.
19. D. Davis, D. Drusvyatskiy, and L. Jiang. Active manifolds, stratifications, and convergence to local minima in nonsmooth optimization. *arXiv preprint* arXiv:2108.11832, 2023.
20. D. Davis, D. Drusvyatskiy, S. Kakade, and J. D. Lee. Stochastic subgradient method converges on tame functions. *Foundations of computational mathematics*, 20(1):119–154, 2020.
21. D. Davis, D. Drusvyatskiy, Y. T. Lee, S. Padmanabhan, and G. Ye. A gradient sampling method with complexity guarantees for general lipschitz functions. *arXiv preprint* arXiv:2112.06969, 2022.
22. D. Davis, D. Drusvyatskiy, K. J. MacPhee, and C. Paquette. Subgradient methods for sharp weakly convex functions. *Journal of Optimization Theory and Applications*, 179:962–982, 2018.
23. D. Drusvyatskiy, A. D. Ioffe, and A. S. Lewis. Generic minimizing behavior in semialgebraic optimization. *SIAM Journal on Optimization*, 26(1):513–534, 2016.
24. D. Drusvyatskiy, A. D. Ioffe, and A. S. Lewis. Nonsmooth optimization using taylor-like models: error bounds, convergence, and termination criteria. *Mathematical Programming*, 185:357–383, 2021.
25. D. Drusvyatskiy and A. Lewis. Error bounds, quadratic growth, and linear convergence of proximal methods. *To appear in Math. Oper. Res.*, arXiv:1602.06661, 2016.
26. D. Drusvyatskiy and A. S. Lewis. Optimality, identifiability, and sensitivity. *Mathematical Programming*, 147(1):467–498, 2014. Citations refer to long version arXiv:1207.6628 .
27. P. J. Enright and B. A. Conway. Discrete approximations to optimal trajectories using direct transcription and nonlinear programming. *Journal of Guidance, Control, and Dynamics*, 15(4):994–1002, 1992.
28. R. Fletcher. A model algorithm for composite nondifferentiable optimization problems. In *Nondifferential and Variational Techniques in Optimization*, pages 67–76. Springer, 1982.
29. A. Goldstein. Optimization of lipschitz continuous functions. *Mathematical Programming*, 13(1):14–22, 1977.
30. X. Han and A. S. Lewis. Survey descent: A multipoint generalization of gradient descent for nonsmooth optimization. *arXiv preprint* arXiv:2111.15645, 2021.

31. W. L. Hare and A. S. Lewis. Identifying active constraints via partial smoothness and prox-regularity. *Journal of Convex Analysis*, 11(2):251–266, 2004.
32. A. Ioffe. An invitation to tame optimization. *SIAM J. Optim.*, 19(4):1894–1917, 2009.
33. P. R. Johnstone and P. Moulin. Faster subgradient methods for functions with hölderian growth. *Mathematical Programming*, 180(1-2):417–450, 2020.
34. J. M. Lee. Smooth manifolds. In *Introduction to Smooth Manifolds*, pages 1–31. Springer, 2013.
35. C. Lemarechal. An extension of davidon methods to non differentiable problems. In *Nondifferentiable optimization*, pages 95–109. Springer, 1975.
36. C. Lemaréchal, F. Oustry, and C. Sagastizábal. The $\mathcal{U}$-lagrangian of a convex function. *Transactions of the American mathematical Society*, 352(2):711–729, 2000.
37. A. Lewis and C. Wylie. A simple newton method for local nonsmooth optimization. *arXiv preprint* arXiv:1907.11742, 2019.
38. A. S. Lewis. Active sets, nonsmoothness, and sensitivity. *SIAM Journal on Optimization*, 13(3):702–725, 2002.
39. A. S. Lewis and M. L. Overton. Nonsmooth optimization via quasi-newton methods. *Mathematical Programming*, 141(1):135–163, 2013.
40. R. Mifflin. Semismooth and semiconvex functions in constrained optimization. *SIAM J. Control Optim.*, 15(6):959–972, 1977.
41. R. Mifflin and C. Sagastizábal. $\mathcal{VU}$-algorithm for convex minimization. *Mathematical Programming*, 104(2):583–608, 2005.
42. A. Nemirovsky and D. Yudin. *Problem complexity and method efficiency in optimization*. A Wiley-Interscience Publication. John Wiley & Sons, Inc., New York, 1983. Translated from the Russian and with a preface by E. R. Dawson, Wiley-Interscience Series in Discrete Mathematics.
43. Y. Nesterov. *Introductory lectures on convex optimization*, volume 87 of *Applied Optimization*. Kluwer Academic Publishers, Boston, MA, 2004. A basic course.
44. W. D. Oliveira and C. Sagastizábal. Bundle methods in the xxist century: A bird's-eye view. *Pesquisa Operacional*, 34:647–670, 2014.
45. A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
46. R. Poliquin and R. Rockafellar. Prox-regular functions in variational analysis. *Transactions of the American Mathematical Society*, 348(5):1805–1838, 1996.
47. B. T. Polyak. Minimization of unsmooth functionals. *USSR Computational Mathematics and Mathematical Physics*, 9(3):14–29, 1969.
48. R. Rockafellar and R.-B. Wets. *Variational Analysis*. Grundlehren der mathematischen Wissenschaften, Vol 317, Springer, Berlin, 1998.
49. A. Shapiro. On a class of nonsmooth composite functions. *Mathematics of Operations Research*, 28(4):677–692, 2003.
50. P. Wolfe. A method of conjugate subgradients for minimizing nondifferentiable functions. In *Nondifferentiable optimization*, pages 145–173. Springer, 1975.
51. S. J. Wright. Convergence of an inexact algorithm for composite nonsmooth optimization. *IMA journal of numerical analysis*, 10(3):299–321, 1990.
52. S. J. Wright. Identifiable surfaces in constrained optimization. *SIAM Journal on Control and Optimization*, 31(4):1063–1079, 1993.
53. Y. Xu and W. Yin. A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion. *SIAM Journal on imaging sciences*, 6(3):1758–1789, 2013.
54. Y.-X. Yuan. On the superlinear convergence of a trust region algorithm for nonsmooth optimization. *Mathematical Programming*, 31(3):269–285, 1985.
55. J. Zhang, H. Lin, S. Jegelka, S. Sra, and A. Jadbabaie. Complexity of finding stationary points of nonconvex nonsmooth functions. *Proceedings of Machine Learning Research*, pages 11173–11182, 2020.