

The complexity of first-order optimization methods from a metric perspective

A.S. Lewis ^{*} Tonghua Tian [†]

May 4, 2023

Abstract

A central tool for understanding first-order optimization algorithms is the Kurdyka-Lojasiewicz inequality. Standard approaches to such methods rely crucially on this inequality to leverage sufficient decrease conditions involving gradients or subgradients. However, the KL property fundamentally concerns not subgradients but rather “slope”, a purely metric notion. By highlighting this view, and avoiding any use of subgradients, we present a simple and concise complexity analysis for first-order optimization algorithms on metric spaces. This subgradient-free perspective also frames a short and focused proof of the KL property for nonsmooth semi-algebraic functions.

Key words: nonsmooth optimization, first-order algorithms, slope, KL property, complexity, semi-algebraic

AMS Subject Classification: 90C48, 49J52, 65Y20, 14P10

1 Introduction

Over the last two decades, beginning with pioneering works such as [1, 6], the Kurdyka-Lojasiewicz inequality has become a central tool for the convergence analysis of first-order optimization algorithms. Viewed as a subgradient property [12, 13], several factors — the profound reasons underlying the inequality’s validity for a large class of nonsmooth functions (semi-algebraic, for example), and the sometimes involved arguments applying the inequality to optimization complexity analysis — can obscure the essential simplicity of the technique. However, as remarked in [14],

^{*}ORIE, Cornell University, Ithaca, NY. people.orie.cornell.edu/aslewis
Research supported in part by National Science Foundation Grant DMS-2006990.

[†]ORIE, Cornell University, Ithaca, NY. tt543@cornell.edu

the inequality is fundamentally a metric space property. This view can be powerful and illuminating, both in non-Euclidean applications like [11, 26], and — the philosophy we pursue here — in understanding and proving the inequality, and leveraging it in simple complexity analysis. The original motivation for the KL property in analytic [32], more general smooth [29], and nonsmooth [12] settings was to bound the length of *continuous*-time descent trajectories, a topic outside our current scope, but fundamental in general metric spaces [3].

In this self-contained exposition we present some basic and broadly applicable complexity consequences of the KL property, with short, elementary proofs. Some ingredients are familiar from scattered sources, but by reorganizing the proof sequence and avoiding any use of subgradients or other traditional nonsmooth analysis, we condense and unify the development. To begin, after introducing the KL inequality, we immediately prove a rudimentary complexity result about objective value convergence, valid in any metric space. This new observation, startling in its breadth and simplicity, illustrates our two central claims.

- Contemporary interest in optimization algorithms for non-Euclidean settings such as manifolds [2, 17, 43] invites a subgradient-free complexity analysis.
- A purely metric development of KL-based complexity analysis promotes not just generality, but also brevity and transparency.

Thus motivated, we present a fresh proof of a nonsmooth Kurdyka-Łojasiewicz property for semi-algebraic (and more general “definable”) functions on Euclidean spaces. By focusing on the property’s metric rather than subdifferential nature, we simplify and clarify earlier developments with no essential loss in applicability.

Returning to our development of complexity results, we study sequences of iterates satisfying a fundamental *slope-descent* condition, combining a traditional notion of sufficient decrease with a metric slope property. We show the satisfiability of this condition in metric spaces, under mild conditions, and derive local and global convergence results for the resulting iterates. As an illustration of our framework, we conclude by analyzing majorization-minimization methods for optimization on Riemannian manifolds and more general geodesic spaces.

2 The Kurdyka-Łojasiewicz property

Our setting is a metric space X . In practice, X may be a Euclidean or Hilbert space, or a Riemannian manifold, but our presentation is purely metric. Laying aside the semi-algebraic setting for our proof of the Kurdyka-Łojasiewicz inequality in Section 4, we make no assumptions akin to finite-dimensionality until Section 7, and until the final examples we use no inner products: our techniques are thus entirely primal in nature.

We seek to minimize a lower semicontinuous function $f: X \rightarrow \bar{\mathbf{R}} = [-\infty, +\infty]$ that is proper, meaning it never takes the value $-\infty$ and has nonempty domain

$$\text{dom } f = \{x \in X : f(x) < +\infty\}.$$

In practice, f is often regular in some sense — perhaps smooth, or convex, or some composition thereof — but this development requires no such assumptions.

Denote the metric in the space X by $d(\cdot, \cdot)$ and the distance function to any point $x \in X$ by $d_x(\cdot)$. Following [20], the *slope* $|\nabla f|(x)$ of f at any point x with finite value $f(x)$ is the infimum of those values $\epsilon > 0$ for which x locally minimizes the perturbed function $f + \epsilon d_x$. At points x where $f(x) = +\infty$, the slope is defined to be $+\infty$. The notation is motivated by the following standard example.

Example 2.1 (Submanifolds in Euclidean spaces)

In the Euclidean space \mathbf{R}^n , consider any $\mathcal{C}^{(1)}$ -smooth submanifold \mathcal{M} , endowed with the Euclidean metric, and any differentiable function $f: \mathcal{M} \rightarrow \mathbf{R}$. At any point $x \in \mathcal{M}$, the covariant gradient of f is an element of the subspace of \mathbf{R}^n tangent to \mathcal{M} at x , and its norm is the slope of f . In particular, when $\mathcal{M} = \mathbf{R}^n$, we have the formula

$$|\nabla f|(x) = |\nabla f(x)|.$$

Definition 2.2 A *desingularizer* is a continuous function $\phi: \mathbf{R}_+ \rightarrow \mathbf{R}_+$ with value $\phi(0) = 0$ and continuous, strictly positive derivative on $(0, +\infty)$. The desingularizer ϕ has *moderate growth* if its derivative satisfies $\phi'(\tau) = o(\frac{1}{\tau})$ as $\tau \downarrow 0$.

The following definition follows the idea of [14]. The strategy is to use a desingularizer to induce a type of “sharpness” by rescaling the *gap*:

$$g(\cdot) = f(\cdot) - \inf f.$$

Definition 2.3 Consider a lower semicontinuous proper function $f: X \rightarrow \bar{\mathbf{R}}$ that is bounded below. We say that f has the *Kurdyka-Lojasiewicz property (for minimization)* if there exists a desingularizer ϕ such that all points $x \in X$ with small, strictly positive gap $g(x)$ satisfy $|\nabla(\phi \circ g)|(x) \geq 1$.

We can rewrite the basic inequality using a simple chain rule [9, Lemma 4.1]:

$$|\nabla(\phi \circ g)|(x) = \phi'(g(x)) \cdot |\nabla g|(x)$$

whenever $0 < g(x) < +\infty$. The proof is a simple application of the classical Mean Value Theorem.

More robust than the notion of slope is the *limiting slope*, which, at any point $\bar{x} \in \text{dom } f$, is the quantity

$$|\bar{\nabla} f|(\bar{x}) = \liminf_{\substack{x \rightarrow \bar{x} \\ f(x) \rightarrow f(\bar{x})}} |\nabla f|(x).$$

It is easy to verify that the definition of the KL property above could equivalently be written with the limiting slope replacing the slope, as could the subsequent chain rule. When the metric space X is Euclidean, the following remarkable identity ([27, Proposition 8.5], for example) always relates the limiting slope and the subdifferential ∂f [40]:

$$\text{dist}(0, \partial f(x)) = |\bar{\nabla} f|(x),$$

leading to an equivalent subdifferential version of the KL property. We emphasize, however, that we make no use of subgradients in this exposition.

When the metric space X is complete, the KL property is equivalent to an “error bound”: for all points x with small gap $g(x)$ and all values r satisfying $0 < r < g(x)$, the distance from x to the level set $g^{-1}[0, r]$ is no larger than $\phi(g(x)) - \phi(r)$. We make no use of this observation (see [8, Theorem 2.1]), but it offers some compelling intuition for the power of the inequality.

Results in [12, 13] guarantee that, on a Euclidean space, any lower semicontinuous semi-algebraic (or, more generally, “globally subanalytic”) function with a bounded set of minimizers has the KL property for minimization, and some corresponding desingularizer has moderate growth. Indeed, there exists such a desingularizer of the form of a concave function $\phi(\tau) = \kappa\tau^{1-\theta}$, where the exponent θ lies in the interval $[0, 1)$ and $\kappa > 0$. Such desingularizers, for example, feature in the original Lojasiewicz proof of the KL property for real-analytic functions [32], and its smooth subanalytic extension in [29]. We will often assume desingularizers to be concave in convergence proofs: for a discussion, see [5].

The rather involved original arguments for the nonsmooth KL property in [12, 13], along with a subsequent proof (due to Drusvyatskiy) in [27] that follows more closely the path laid out in [29], all depend heavily on the geometry of subgradients. In line with our philosophy in this work, and inspired by another recent proof [41, Theorem 2.9], in Section 4 we derive the nonsmooth result from the smooth version in [29] using just a simple slope-based argument.

3 Objective value convergence

We next follow a theme with roots in [33], taken up in [1], and promoted in [6, 15, 16, 38]. Many classical optimization algorithms — the first-order methods considered in [15] such as gradient descent, the proximal point method, the proximal gradient method, variable metric and quasi-Newton methods, trust region methods, alternating projections, Gauss-Seidel schemes and so forth, majorization-minimization procedures considered in [16], including sequential quadratic programming schemes and the moving balls method, and some bundle methods [38] — generate sequences in X with a certain descent property. Specifically, for some constant $\delta > 0$, succes-

sive iterates x, x_+ always satisfy

$$(3.1) \quad f(x) - f(x_+) \geq \delta(|\bar{\nabla}f|(x_+))^2.$$

Algorithms outside the Euclidean or Hilbert setting may also guarantee the same condition: as remarked in [30], a basic example is the proximal point method [10,28].

Theorem 3.2 (Value convergence) *On a metric space X , consider any lower semicontinuous proper function $f: X \rightarrow \bar{\mathbf{R}}$ that is bounded below and satisfies the KL property for minimization, and any sequence (x_k) in X satisfying the descent property (3.1) and with initial value $f(x_0)$ near $\inf f$. Then the values $f(x_k)$ decrease to $\inf f$. If, furthermore, the corresponding desingularizer has moderate growth, then the gap satisfies*

$$f(x_k) - \inf f = o\left(\frac{1}{k}\right).$$

Proof Consider the gap $\tau_k = g(x_k) \geq 0$ at iteration k . By assumption, we have

$$\tau_{k-1} - \tau_k \geq \delta(|\bar{\nabla}f|(x_k))^2,$$

so τ_k is nonincreasing. If $\tau_k = 0$ at any iteration k then there is nothing to prove, so we can suppose that τ_k is small and nonzero for all k .

If the result fails, then τ_k decreases to some limit $\bar{\tau} > 0$. Using the chain rule, the KL property ensures

$$\phi'(\tau_k) \cdot |\bar{\nabla}f|(x_k) \geq 1$$

for all k , and hence

$$(3.3) \quad \tau_{k-1} - \tau_k \geq \frac{\delta}{(\phi'(\tau_k))^2}.$$

But as $k \rightarrow \infty$, the left-hand side converges to zero, while the right-hand side converges to the strictly positive limit $\delta/(\phi'(\bar{\tau}))^2$, which is a contradiction.

Suppose now that the desingularizer ϕ has moderate growth. Given any constant $\mu > 0$, since $0 < \tau_k \downarrow 0$, we know $\tau_k < \mu$ and $\tau_k \phi'(\tau_k) < \sqrt{\mu\delta}$ for all large k . In particular, there exists an integer m such that whenever $k \geq m$,

$$\tau_{k-1} - \tau_k > \frac{1}{\mu} \tau_k^2,$$

and hence

$$\frac{1}{\tau_{k-1}} < \frac{1}{\tau_k} - \frac{1}{\mu + \tau_k} < \frac{1}{\tau_k} - \frac{1}{2\mu}.$$

Induction shows

$$\frac{1}{\tau_k} \geq \frac{1}{\tau_m} + \frac{k-m}{2\mu}$$

so

$$k\tau_k \leq \frac{2\mu\tau_m k}{2\mu + \tau_m(k - m)}.$$

We deduce

$$\limsup_{k \rightarrow \infty} k\tau_k \leq 2\mu.$$

Since the constant $\mu > 0$ was arbitrary, the result follows. \square

The $o(\frac{1}{k})$ rate is reminiscent of the behavior of the proximal point method for convex optimization [25, Theorem 3.1]. In the case of the standard desingularizer $\phi(\tau) = \kappa\tau^{1-\theta}$ for $\kappa > 0$ and $0 \leq \theta < 1$, we can refine this result slightly. The key inequality (3.3) becomes

$$\tau_{k-1} - \tau_k \geq \mu\tau_k^{2\theta}$$

for some constant $\mu > 0$. An elementary induction then shows four different scenarios (cf. [4]) for the rate at which the gap converges to zero: finite convergence ($\theta = 0$), superlinear ($0 < \theta < \frac{1}{2}$), linear ($\theta = \frac{1}{2}$), and finally, for $\frac{1}{2} < \theta < 1$,

$$f(x_k) - \inf f = O\left(\left(\frac{1}{k}\right)^{\frac{1}{2\theta-1}}\right).$$

The first case is immediate, and the second and third follow from the inequality

$$\frac{\tau_{k-1}}{\tau_k} \geq 1 + \mu\tau_k^{2\theta-1}$$

for all large k . In the second case, the right-hand side grows indefinitely with k , and in the third it equals $1 + \mu$. Suppose the last case fails. For any constant $M \geq 1$ suppose

$$\tau_k > M\left(\frac{1}{k}\right)^{\frac{1}{2\theta-1}}$$

for some $k \geq 2$ and yet the same inequality fails when we replace k by $k - 1$. We obtain

$$M\left(\frac{1}{k}\right)^{\frac{1}{2\theta-1}} + \mu M^{2\theta}\left(\frac{1}{k}\right)^{\frac{2\theta}{2\theta-1}} < \tau_k + \mu\tau_k^{2\theta} \leq \tau_{k-1} \leq M\left(\frac{1}{k-1}\right)^{\frac{1}{2\theta-1}}$$

and then convexity implies a contradiction for sufficiently large M :

$$1 + \frac{\mu M^{2\theta-1}}{k} < \left(1 - \frac{1}{k}\right)^{-\frac{1}{2\theta-1}} \leq \left(1 - \frac{2}{k}\right) + \frac{2}{k}\left(\frac{1}{2}\right)^{-\frac{1}{2\theta-1}} = 1 + \frac{2(2^{\frac{1}{2\theta-1}} - 1)}{k}.$$

The induction follows.

4 A slope-based proof of the KL property

Viewed through the lens of the slope, the fact that nonsmooth semi-algebraic functions satisfy a Kurdyka-Łojasiewicz property is a simple consequence of the original smooth results due to Łojasiewicz [32] and Kurdyka [29]. Appealing to subgradients, as in the original nonsmooth arguments [12, 13], obscures this simplicity. In the brief proof we present here, like [41], we focus entirely on the slope.

Our exposition in fact applies in a setting much broader than that of semi-algebraic geometry, using instead the property known as “definability in an o-minimal structure”. Readers unfamiliar with definable sets and functions can comfortably focus on the semi-algebraic special case. The accessible reference [19] gives a concise introduction to definability.

We fix an o-minimal structure \mathcal{O} , with respect to which the various functions and sets we consider are definable. We start with a smooth result (cf. [13, Theorem 11]), generalizing [29, Theorem 1].

Theorem 4.1 *In the Euclidean space \mathbf{R}^n , consider any bounded $\mathcal{C}^{(1)}$ -smooth definable submanifold \mathcal{M} , and any differentiable definable function $f: \mathcal{M} \rightarrow \mathbf{R}$. Then there exists a concave definable desingularizer ϕ and a constant $\rho > 0$ such that*

$$(4.2) \quad |\nabla_{\mathcal{M}}(\phi \circ f)(x)| \geq 1 \quad \text{whenever } 0 < f(x) < \rho,$$

where $\nabla_{\mathcal{M}}$ denotes the covariant gradient on \mathcal{M} .

Proof Consider the definable function $\psi: (0, \infty) \rightarrow \bar{\mathbf{R}}$ defined as

$$\psi(t) = \inf\{|\nabla_{\mathcal{M}}f(x)| : f(x) = t\}.$$

(As usual, the infimum over the empty set is $+\infty$.) We lose no generality in supposing $\liminf_{t \searrow 0} \psi(t) = 0$, since otherwise, for all sufficiently small $\rho > 0$, the desingularizer defined by $\phi(t) = \frac{1}{\rho}t$ for $t \geq 0$ has the desired property. A standard property of definable functions ([19, Theorem 2.1]) then implies $\lim_{t \searrow 0} \psi(t) = 0$.

Consider the definable set

$$S = \{(x, t) \in \mathcal{M} \times (0, +\infty) : f(x) = t, |\nabla_{\mathcal{M}}f(x)| < \psi(t) + t^2\}.$$

Since the manifold \mathcal{M} is bounded, an easy argument shows the existence of a point $\bar{x} \in \mathbf{R}^n$ satisfying $(\bar{x}, 0) \in \text{cl } S$. Further standard properties (see [19, Theorems 2.1 and 3.2]) now guarantee that the point $(\bar{x}, 0)$ is accessible by a smooth curve in S . More precisely, there exists a $\mathcal{C}^{(1)}$ -smooth curve $\gamma: (0, 1) \rightarrow \mathcal{M}$ such that the composition $g = f \circ \gamma$ satisfies $(\gamma(s), g(s)) \in S$ for $0 < s < 1$, and as $s \searrow 0$, we have $(\gamma(s), g(s)) \rightarrow (\bar{x}, 0)$ and $(\gamma'(s), g'(s))$ converges to some limit $(v, r) \in \mathbf{R}^n \times \mathbf{R}_+$. We can furthermore assume that the derivative g' is either constant or strictly monotone.

Whenever $0 < s < 1$ we have

$$g'(s) = \langle \nabla_{\mathcal{M}} f(\gamma(s)), \gamma'(s) \rangle \leq |\nabla_{\mathcal{M}} f(\gamma(s))| \cdot |\gamma'(s)|.$$

Furthermore, we observe

$$\begin{aligned} r &= \lim_{s \searrow 0} \langle \nabla_{\mathcal{M}} f(\gamma(s)), \gamma'(s) \rangle \\ &\leq \lim_{s \searrow 0} \left(|\gamma'(s)| (\psi(g(s)) + g(s)^2) \right) = 0, \end{aligned}$$

so in fact $r = 0$. Consequently, the derivative g' must be strictly increasing, and so g is also strictly increasing and satisfies $g(s) < sg'(s)$ for all $s \in (0, 1)$. Take any constant ρ in the interval $(0, \lim_{s \nearrow 1} g(s))$ and let $\xi: (0, \rho] \rightarrow (0, 1)$ denote the restriction of the inverse function g^{-1} , which is also strictly increasing.

Now consider any point $x \in \mathcal{M}$ satisfying $0 < f(x) < \rho$. Setting $t = f(x)$ and $s = \xi(t)$, then providing ρ is sufficiently small, we have

$$\begin{aligned} |\nabla_{\mathcal{M}}(\xi \circ f)(x)| &= \xi'(f(x)) \cdot |\nabla_{\mathcal{M}} f(x)| \geq \xi'(t) \cdot \psi(t) = \frac{\psi(g(s))}{g'(s)} \\ &> \frac{|\nabla_{\mathcal{M}} f(\gamma(s))| - g(s)^2}{g'(s)} > \frac{1}{|\gamma'(s)|} - sg(s) > \frac{1}{1 + |v|}. \end{aligned}$$

Note that the derivative $\xi' = \frac{1}{g' \circ \xi} > 0$ is strictly decreasing, and $\lim_{t \searrow 0} \xi(t) = 0$. Consequently, the function $\phi: \mathbf{R}_+ \rightarrow \mathbf{R}_+$ defined by

$$\phi(t) = \begin{cases} 0 & (t = 0) \\ (1 + |v|) \cdot \xi(t) & (0 < t < \rho) \\ (1 + |v|) \cdot \xi(\rho) & (t \geq \rho) \end{cases}$$

has the desired property. \square

Using a simple stratification argument, we arrive at the following result of Kurdyka-Lojasiewicz type for nonsmooth definable functions.

Theorem 4.3 *In the Euclidean space \mathbf{R}^n , consider any bounded definable set U , and any definable function $f: U \rightarrow \mathbf{R}$. Then there exists a concave definable desingularizer ϕ and a constant $\rho > 0$ such that*

$$(4.4) \quad |\nabla(\phi \circ f)|(x) \geq 1 \quad \text{whenever } 0 < f(x) < \rho.$$

Proof A standard stratification [42] partitions U into finitely-many bounded $\mathcal{C}^{(1)}$ -smooth definable submanifolds $\mathcal{M}_1, \dots, \mathcal{M}_k$ of \mathbf{R}^n such that each restriction $f|_{\mathcal{M}_i}$

is a differentiable definable function. Applying Theorem 4.1 to each restriction $f|_{\mathcal{M}_i}$ gives a concave definable desingularizer ϕ_i and a constant $\rho > 0$ such that

$$|\nabla_{\mathcal{M}_i}(\phi_i \circ f|_{\mathcal{M}_i})(x)| \geq 1 \quad \text{whenever } x \in \mathcal{M}_i \text{ with } 0 < f(x) < \rho.$$

Applying [19, Theorem 2.1] to the collection of functions $\phi_i - \phi_j$ with $i \neq j$, by reducing ρ if necessary, we deduce the existence of a particular index j such that $\phi'_j \geq \phi'_i$ on the interval $(0, \rho)$ for each index i . Take $\phi = \phi_j$. Now consider any point $x \in U$ with $0 < f(x) < \rho$. For some i we have $x \in \mathcal{M}_i$, and then

$$\begin{aligned} |\nabla(\phi \circ f)|(x) &\geq |\nabla(\phi_i \circ f)|(x) \geq |\nabla(\phi_i \circ f|_{\mathcal{M}_i})(x) \\ &= |\nabla_{\mathcal{M}_i}(\phi_i \circ f|_{\mathcal{M}_i})(x)| \geq 1 \end{aligned}$$

where the first inequality follows from the inequality $\phi' \geq \phi'_i$ and the chain rule, and the second inequality follows from the definition of the slope. \square

Note that when U and f are subanalytic, by the Puiseux expansion we know the desingularizer can take the form $\phi(\tau) = \kappa\tau^{1-\theta}$ for some exponent $\theta \in [0, 1)$ and constant $\kappa > 0$ (see [29]).

The proof we present here is more concise than the original development in [13], which relies heavily on the geometry of subgradients. This brevity, however, comes with a cost. As we remarked earlier, the KL inequality (4.4) can be written equivalently using the limiting slope, and hence as the following property:

Subgradients of the function $\phi \circ f$ at the point x all have norm at least 1.

By contrast, [13] proves a stronger result, namely the same property for *Clarke* subgradients. In particular, when the function f is Lipschitz, this ensures that the property above holds for all *convex combinations* of subgradients. By relying on the weaker, limiting-slope-based KL property, we potentially narrow the scope of our convergence analysis: the basic descent condition (3.1) at the iterate x amounts to requiring the existence at the next iterate x_+ of a subgradient of norm no larger than $\frac{1}{\sqrt{\delta}}\sqrt{f(x) - f(x_+)}$, whereas using the stronger KL property is more flexible, allowing Clarke subgradients instead. On the other hand, concrete methods such as majorization-minimization algorithms do indeed enforce condition (3.1), so this distinction has no impact. For the convergence results presented here, we are unaware of any practical impact deriving from the stronger KL property in [13]: the simpler slope-based version proved here seems adequate.

5 Convergence to a minimizer

As we have seen, convergence in value follows from the basic descent condition (3.1). In practice, this condition often follows from two *slope-descent conditions*:

there exists constants $\alpha > 0$ and $\beta > 0$ such that successive iterates x, x_+ always satisfy

$$(5.1) \quad f(x) - f(x_+) \geq \alpha d^2(x, x_+),$$

$$(5.2) \quad d(x, x_+) \geq \beta |\bar{\nabla} f|(x_+).$$

We first observe that these conditions, which originated with the seminal analysis of [33], are always satisfiable. We now assume that the underlying metric space X is complete.

Proposition 5.3 (Existence of slope descent sequences) *On a complete metric space X , consider any lower semicontinuous function $f: X \rightarrow \bar{\mathbf{R}}$ that is bounded below. Then for any point $x \in \text{dom } f$ and constants $\alpha, \beta > 0$ satisfying $\alpha\beta < \frac{1}{2}$, there exists a point $x_+ \in X$ satisfying the slope-descent conditions (5.1) and (5.2).*

Proof In fact we prove a stronger property, with the limiting slope replaced by the slope. The function $h = f + \alpha d_x^2$ is also lower semicontinuous and bounded below. We claim that, at all points $y \in X$, its slope satisfies

$$|\nabla h|(y) \geq |\nabla f|(y) - 2\alpha d(x, y).$$

To prove this, we can restrict attention to a point $y \in \text{dom } h$ where $|\nabla h|(y)$ is finite. Consider any constants $\epsilon' > \epsilon > |\nabla h|(y)$. Since y locally minimizes the function $h + \epsilon d_y$, we deduce that all points $z \in X$ sufficiently close to y satisfy

$$\begin{aligned} f(y) &\leq f(z) + \alpha(d(x, z)^2 - d(x, y)^2) + \epsilon d(y, z) \\ &\leq f(z) + \alpha(2d(x, y)d(y, z) + d(y, z)^2) + \epsilon d(y, z) \\ &\leq f(z) + (2\alpha d(x, y) + \epsilon')d(y, z), \end{aligned}$$

Consequently y also locally minimizes the function

$$f + (2\alpha d(x, y) + \epsilon')d_y$$

and we deduce $|\nabla f|(y) \leq 2\alpha d(x, y) + \epsilon'$. Letting $\epsilon' \downarrow |\nabla h|(y)$ proves the claim.

If x minimizes the function h , then $|\nabla h|(x) = 0$ and hence $|\nabla f|(x) = 0$, so we can choose $x_+ = x$. Otherwise, define $t = \frac{1}{2}(h(x) + \inf h)$ and note that the distance δ from x to the level set $\{y \in X : h(y) \leq t\}$ is strictly positive. The Ekeland principle [24] implies the existence of a point $x_+ \in X$ satisfying both the inequalities $h(x_+) \leq t$ and $|\nabla h|(x_+) \leq (\frac{1}{\beta} - 2\alpha)\delta$. We deduce $d(x, x_+) \geq \delta$ and

$$|\nabla f|(x_+) \leq |\nabla h|(x_+) + 2\alpha d(x, x_+) \leq \frac{1}{\beta} d(x, x_+),$$

so inequality (5.2) follows. Inequality (5.1) follows from $h(x_+) \leq t < h(x)$. \square

The argument above shows in particular that any *proximal point* x_+ (a point that minimizes the function $f + \alpha d_x^2$) must satisfy the slope descent conditions (5.1) and (5.2) assuming $\alpha\beta \leq \frac{1}{2}$. Existence (and uniqueness) of proximal points for lower semicontinuous convex functions f is guaranteed on a Hilbert space X by Minty's Theorem [35], and also in more general metric spaces [10, 28]. However, in general, unless the metric space X is *proper* (meaning that all closed balls are compact), proximal points may not exist.

The following result and its proof condense arguments in [6]. We assume concavity of the desingularizer, and use lower semicontinuity of the objective, but rely on no further continuity properties.

Theorem 5.4 (Iterate convergence) *On a complete metric space X , consider any lower semicontinuous function $f: X \rightarrow \bar{\mathbf{R}}$ that is bounded below and satisfies the KL property for minimization, with a concave desingularizer ϕ , and any sequence (x_k) in X satisfying the slope-descent conditions (5.1) and (5.2) with initial value $f(x_0)$ near $\inf f$. Then, x_k converges to a minimizer of f .*

Proof We can assume that no iterate minimizes f , since otherwise the first inequality would ensure that the sequence is constant thereafter. The objective values $f(x_k)$ decrease monotonically, so are finite, and the KL inequality ensures that the gap g satisfies

$$1 \leq |\bar{\nabla}(\phi \circ g)|(x_k) = \phi'(g(x_k)) \cdot |\bar{\nabla}g|(x_k).$$

We deduce that each limiting slope $|\bar{\nabla}g|(x_k) = |\bar{\nabla}f|(x_k)$ is nonzero, and so the distances $d_k = d(x_{k-1}, x_k)$ are nonzero for all $k > 0$.

Using the concavity of the desingularizer ϕ , we observe, for all $k > 0$,

$$\begin{aligned} \phi(g(x_k)) - \phi(g(x_{k+1})) &\geq \phi'(g(x_k))(g(x_k) - g(x_{k+1})) \\ &\geq \frac{f(x_k) - f(x_{k+1})}{|\bar{\nabla}f|(x_k)} \geq \frac{\alpha d_{k+1}^2}{\beta d_k}, \end{aligned}$$

so

$$\frac{\beta}{\alpha} \left(\phi(g(x_k)) - \phi(g(x_{k+1})) \right) \geq \frac{2d_k d_{k+1} - d_k^2}{d_k} \geq 2d_{k+1} - d_k.$$

In terms of the numbers

$$\lambda_k = \frac{\beta}{\alpha} \phi(g(x_k)) + d_k > 0 \quad \text{for } k > 0,$$

we have proved

$$d(x_k, x_{k+1}) \leq \lambda_k - \lambda_{k+1}.$$

It follows that the sequence (x_k) is Cauchy, and hence converges to some limit x^* . But Theorem 3.2 shows that $f(x_k)$ decreases to $\inf f$, so lower semicontinuity now ensures that x^* minimizes f . \square

In general, we know nothing about how fast the error $d(x_k, x^*)$ converges to zero, even in the canonical case of the proximal point method. However, with some knowledge of the desingularizer ϕ , the proof shows more. Following [4, 6], we have

$$d(x_k, x^*) \leq \sum_{j=k}^{m-1} d(x_j, x_{j+1}) + d(x_m, x^*) \leq \lambda_k - \lambda_m + d(x_m, x^*).$$

Letting $m \rightarrow \infty$ shows

$$d(x_k, x^*) \leq \lambda_k = \frac{\beta}{\alpha} \phi(g(x_k)) + d_k \leq \frac{\beta}{\alpha} \phi(g(x_k)) + \sqrt{\frac{g(x_{k-1})}{\alpha}},$$

and complexity estimates follow. Following the discussion in the earlier section, in the case of the standard desingularizer $\phi(\tau) = \kappa\tau^{1-\theta}$ for $\kappa > 0$ and $0 \leq \theta < 1$, we obtain four different scenarios in terms of the convergence rate: finite ($\theta = 0$), superlinear ($0 < \theta < \frac{1}{2}$), linear ($\theta = \frac{1}{2}$), and, for the case $\frac{1}{2} < \theta < 1$, the scenario

$$d(x_k, x^*) = O\left(\left(\frac{1}{k}\right)^{\frac{1-\theta}{2\theta-1}}\right).$$

6 Convergence to a critical point

The previous convergence results rely on the initial value of the objective function f being close to its infimum. Without that assumption, in the absence of any global convex-like assumptions on f , we instead resort to seeking a point $x^* \in X$ that is *critical*, meaning that the limiting slope $|\bar{\nabla}f|(x^*)$ is zero. For smooth functions on manifolds, for example, this idea coincides with the classical notion. Clearly, any local minimizer is a critical point.

In this section, we present conditions guaranteeing a *global* convergence property: any sequence satisfying the slope-descent conditions (5.1) and (5.2) must converge to a critical point. The approach essentially follows [6], in particular relying first on the following property.

Definition 6.1 Consider a lower semicontinuous proper function $f: X \rightarrow \bar{\mathbf{R}}$. We call f a *KL-function* if, for all points $\bar{x} \in X$ with finite value, there exists a concave desingularizer ϕ and a constant $\rho > 0$ such that

$$(6.2) \quad d(x, \bar{x}) < \rho \quad \text{and} \quad f(x) - \bar{f} \in (0, \rho) \quad \Rightarrow \quad |\nabla(\phi \circ (f - \bar{f}))|(x) \geq 1,$$

where $\bar{f} = f(\bar{x})$.

Some remarks about this definition are in order. First, the desired property is automatic at points \bar{x} that are not critical, as is easy to verify, so the notion really

concerns critical points. Secondly, we note that all lower semicontinuous semi-algebraic (or globally subanalytic) functions are KL-functions [12, 13, 41]. Finally, as before, we can equivalently replace the slope in the definition by the limiting slope.

From Section 4, we deduce the following result.

Corollary 6.3 *All lower semicontinuous semi-algebraic (or, more generally, subanalytic) functions are KL-functions.*

An important feature of any KL-function f is the possibility of fixing a *uniform* desingularizer over any compact set $\bar{X} \subset X$ on which f is constant. More precisely, there exists a desingularizer ϕ and constant $\rho > 0$ such that the KL property (6.2) holds (in terms of both the slope and limiting slope) for all points $\bar{x} \in \bar{X}$. The proof is a routine compactness argument: using the desingularizers $\phi_{\bar{x}}$ and constants $\rho_{\bar{x}}$ for each point \bar{x} , the corresponding neighborhoods cover \bar{X} , so we can select a subcover corresponding to a finite set of points $\{x_i\}$, and then set the constant $\rho = \min_i \rho_{x_i}$ and the desingularizer $\phi = \sum_i \phi_{x_i}$. As a consequence of this uniformity property, any KL function with a compact set of minimizers has the Kurdyka-Łojasiewicz property for minimization, in the sense of Definition 2.3.

We would not be exaggerating much to assert that all optimization objectives in practice are KL-functions. The second condition we assume is less universal, but still common (cf. [3, Theorem 2.3.1] and [22]).

Definition 6.4 A proper function $f: X \rightarrow \bar{\mathbf{R}}$ is *continuous on slope-bounded sets* if, for all points $\bar{x} \in \text{dom } f$ and sequences $x_k \rightarrow \bar{x}$, if the values $f(x_k)$ and slopes $|\nabla f|(x_k)$ are uniformly bounded, then $f(x_k) \rightarrow f(\bar{x})$.

Clearly this property holds if f is continuous, or more generally, continuous on its domain, but it also holds for lower semicontinuous functions on a Euclidean space that are convex or, more generally, “subdifferentially continuous” [40].

Theorem 6.5 (Global convergence) *On a complete metric space X , consider any KL-function $f: X \rightarrow \bar{\mathbf{R}}$ that is bounded below and continuous on slope-bounded sets. Then any relatively compact sequence (x_k) in X satisfying the slope-descent conditions (5.1) and (5.2) must converge to a critical point of f .*

Proof The value $\mu = \inf_k f(x_k)$ is finite. Note $f(x_k) - f(x_{k+1}) \rightarrow 0$, because the values $f(x_k)$ decrease monotonically to μ . The slope-descent conditions (5.1) and (5.2) now imply

$$(6.6) \quad d(x_k, x_{k+1}) \rightarrow 0 \quad \text{and} \quad |\bar{\nabla} f|(x_k) \rightarrow 0.$$

Furthermore, if $f(x_k) = \mu$ ever holds, then $|\bar{\nabla} f|(x_k) = 0$ and $x_j = x_k$ for all $j \geq k$, so there is nothing to prove. Thus we can assume $f(x_k) > \mu$ for all k .

The set \bar{X} of the limit points of the sequence (x_k) is closed, and is contained in the closure of the sequence, which by assumption is compact. Hence \bar{X} is nonempty and compact. An elementary argument implies $\text{dist}_{\bar{X}}(x_k) \rightarrow 0$.

Consider any subsequence $K \subset \mathbf{N}$ for which the sequence $(x_k)_{k \in K}$ converges. Denote the limit by \bar{x} . As $k \rightarrow \infty$ in K , we know $f(x_k) \rightarrow \mu$ and $|\bar{\nabla} f|(x_k) \rightarrow 0$, so by the definition of the limiting slope, there exist corresponding points x'_k for $k \in K$ such that $x'_k \rightarrow \bar{x}$, $f(x'_k) \rightarrow \mu$, and $|\nabla f|(x'_k) \rightarrow 0$ as $k \rightarrow \infty$ in K . Thus \bar{x} is critical, and continuity on slope-bounded sets implies $f(\bar{x}) = \mu$. Thus f is constant on the set \bar{X} , which furthermore consists entirely of critical points.

Following the discussion of uniformity after Definition 6.1, since the set \bar{X} is compact, there exists a concave desingularizer ϕ and a constant $\rho > 0$ such that property (6.2) holds with $\bar{f} = \mu$ for all points $\bar{x} \in \bar{X}$. For all large k , we know

$$\text{dist}_{\bar{X}}(x_k) < \rho \quad \text{and} \quad f(x_k) - \mu \in (0, \rho),$$

so

$$|\bar{\nabla}(\phi \circ (f - \mu))|(x_k) \geq 1.$$

Following exactly the same argument as in the proof of Theorem 5.4, we deduce that the whole sequence (x_k) converges, and the limit is critical since it lies in \bar{X} . \square

7 Majorization-minimization methods

To illustrate the philosophy of this work, we present a purely metric approach to the convergence analysis of majorization-minimization (MM) algorithms, conceptually simply iterative techniques for minimizing an objective function $f: X \rightarrow \mathbf{R}$ over a set X . At each point $x \in X$ we associate a model subset $D(x)$ of X that contains x , and a *majorizing* model function h_x : that is, $h_x \geq f$ with equality at the point x . The simplest examples have $D(x)$ identically equal to X . We then consider the corresponding *envelope* $F: X \rightarrow \bar{\mathbf{R}}$ and set-valued *iteration map* $p: X \rightrightarrows X$ defined at each point $x \in X$ by

$$F(x) = \inf_{D(x)} h_x \quad \text{and} \quad p(x) = \underset{D(x)}{\text{argmin}} h_x$$

and *MM sequences*, in which successive iterates x, x_+ always satisfy $x_+ \in p(x)$. Majorization-minimization algorithms have been widely used in practice, originally in the statistics literature [21].

A fundamental initial question is whether an MM sequence (x_k) is a *minimizing sequence* for the objective f , by which we simply mean $f(x_k) \rightarrow \inf f$. We first collect some elementary observations.

Proposition 7.1 Consider an MM method for an objective f over a set X . Then f majorizes the corresponding envelope F , and $\inf f = \inf F$. An MM sequence is a minimizing sequence for f if and only if it is a minimizing sequence for F . Furthermore, the iteration map p satisfies

$$p(\operatorname{argmin} F) \subset \operatorname{argmin} f \subset \operatorname{argmin} F,$$

with equality if p is everywhere nonempty-valued and the model satisfies $h_x(y) > f(y)$ for all distinct points $x \in X$ and $y \in D(x)$.

Proof All $x \in X$ satisfy $F(x) \leq h_x(x) = f(x)$, so majorization follows. We also have

$$\inf_{y \in X} F(y) \leq \inf_{y \in X} f(y) \leq \inf_{y \in D(x)} h_x(y) = \inf_x F(x),$$

so f and F have the same infimum and furthermore $\operatorname{argmin} f \subset \operatorname{argmin} F$. To deduce the claim about MM sequences (x_k) , notice $f(x_k) \geq F(x_k) \geq f(x_{k+1})$, because $F(x_k) = h_{x_k}(x_{k+1})$. Finally, consider any minimizer x for F . If there exists a point $y \in p(x)$, then we have

$$\inf F = F(x) = h_x(y) \geq f(y) \geq \inf f = \inf F.$$

Hence equality holds throughout, so y minimizes f . If we have the extra strict domination condition, then the above relationship implies $y = x$. \square

The final claim may fail if the iteration map is possibly empty-valued. An example is the objective $f(x) = e^x$ for $x \in X = \mathbf{R}$, with model subset $D(x) = \mathbf{R}$ for all x and model function $h_x(y) = (1 + |x - y|)e^y$ for $y \in \mathbf{R}$. The envelope F is identically zero and the iteration map p is always empty-valued.

The proposition above motivates our subsequent analysis, an approach inspired by [16] but once again avoiding any variational-analytic tools such as subgradients or normal cones. Rather than applying convergence results like those in Section 3 directly to the objective f , Proposition 7.1 opens up the possibility of working instead with the envelope F .

We return to the setting where X is a metric space, and consider bivariate functions $h: X^2 \rightarrow \mathbf{R}$ that are locally Lipschitz, or equivalently, locally Lipschitz with respect to each variable separately. The novelty in our analysis revolves around the following strengthening of the Lipschitz property with respect to the second variable, an apparently new notion.

Definition 7.2 Given a metric space X , an *approximator* is a locally Lipschitz function $h: X^2 \rightarrow \mathbf{R}$ such that for each point $\bar{x} \in X$ there exists a constant $\gamma \geq 0$ so all points $x, y, z \in X$ near \bar{x} satisfy the inequality

$$(7.3) \quad |h(x, y) - h(z, y)| \leq \gamma d(x, z)(d(x, y) + d(z, y)).$$

Here are some simple examples.

Example 7.4 (Squared distance) The function defined by $h(x, y) = d^2(x, y)$ for $x, y \in X$ is an approximator, being locally Lipschitz and satisfying

$$\begin{aligned} |h(x, y) - h(z, y)| &= |d^2(x, y) - d^2(z, y)| = |d(x, y) - d(z, y)|(d(x, y) + d(z, y)) \\ &\leq d(x, z)(d(x, y) + d(z, y)) \end{aligned}$$

for all $z \in X$.

Example 7.5 (Threshold) For any radius $\rho > 0$, the function defined by

$$h(x, y) = \max\{d(x, y) - \rho, 0\} = \text{dist}_{B_\rho(x)}(y) \quad (x, y \in X)$$

is an approximator. Being a maximum of Lipschitz functions, it too is Lipschitz. Furthermore, for any point $z \in X$, if $d(x, y) \geq \rho$ and $d(z, y) \geq \rho$, then

$$h(x, y) - h(z, y) = d(x, y) - d(z, y) \leq d(x, z) \leq \frac{1}{2\rho}d(x, z)(d(x, y) + d(z, y)),$$

while if $d(x, y) \geq \rho \geq d(z, y)$, then

$$\begin{aligned} h(x, y) - h(z, y) &= d(x, y) - \rho \leq d(x, y) - d(z, y) \leq d(x, z) \\ &\leq \frac{1}{\rho}d(x, z)(d(x, y) + d(z, y)). \end{aligned}$$

We deduce

$$h(x, y) - h(z, y) \leq \frac{1}{\rho}d(x, z)(d(x, y) + d(z, y)),$$

the case $d(x, y) \leq \rho$ being trivial. The desired property now follows by symmetry.

The following example is particularly useful.

Example 7.6 (Linear approximation) In the Euclidean space $X = \mathbf{R}^n$, consider a function $g: \mathbf{R}^n \rightarrow \mathbf{R}$ whose gradient ∇g is L -Lipschitz. Then the function defined from the linear approximation by

$$h(x, y) = g(x) + \langle \nabla g(x), y - x \rangle \quad (x, y \in X)$$

is an approximator.

The following observation is simple to verify.

Proposition 7.7 *Sums and scalar multiples of approximators are also approximators.*

Our development relies on some strong structural assumptions about the underlying feasible region X . We assume in particular that, as a metric space, X is *proper*, meaning that all closed balls are compact. In particular, X is therefore complete and locally compact.

Proposition 7.8 (Uniform approximators) *Consider a proper metric space X and a function $h: X^2 \rightarrow \mathbf{R}$. Then h is an approximator if and only if for all bounded subsets $X' \subset X$, the function h is Lipschitz on $(X')^2$ and there exists a constant $\gamma > 0$ such that the approximator inequality (7.3) holds uniformly for all points $x, y, z \in X'$.*

Proof The given property clearly implies that h is an approximator, so we focus on the converse. Assuming that h is an approximator, the uniform Lipschitz property on bounded sets is a standard compactness argument, and the following proof of inequality (7.3) is similar. If the property fails, then there exist sequences (x_k) , (y_k) and (z_k) in the bounded set X' and constants $\gamma_k \rightarrow +\infty$ satisfying, for all $k = 1, 2, 3, \dots$, the inequality

$$(7.9) \quad |h(x_k, y_k) - h(z_k, y_k)| > \gamma_k d(x_k, z_k)(d(x_k, y_k) + d(z_k, y_k)).$$

Since X is proper, we can take subsequences and suppose $x_k \rightarrow \bar{x}$, $y_k \rightarrow \bar{y}$, and $z_k \rightarrow \bar{z}$. Denoting by λ a Lipschitz constant for h on $X' \times X'$, we deduce

$$\lambda > \gamma_k (d(x_k, y_k) + d(z_k, y_k)),$$

so both $d(x_k, y_k)$ and $d(z_k, y_k)$ converge to zero. Consequently we have $\bar{x} = \bar{y} = \bar{z}$, and so inequality (7.9) contradicts the definition of an approximator. \square

Lemma 7.10 (Distance functions as approximators) *Consider a proper metric space X and a set-valued mapping $D: X \rightrightarrows X$ such that for each point $x \in X$, the image $D(x)$ is closed and contains x . Then the function defined by*

$$h(x, y) = \text{dist}_{D(x)}(y)$$

is an approximator if and only if h is locally Lipschitz and for each point $\bar{x} \in X$ there exists a constant $\gamma \geq 0$ such that all points $x, y, z \in X$ near \bar{x} with $y \in D(z)$ satisfy the inequality

$$\text{dist}_{D(x)}(y) \leq \gamma d(x, z)(d(x, y) + d(z, y)).$$

Proof The direct implication is immediate, so we focus on the converse. For any point $x \in X$, the function $H(x, \cdot)$, being a distance function, is 1-Lipschitz. Turning

to the second variable and the approximator inequality, first note that for all points $x, y, z \in X$, choosing any $u \in \text{Proj}_{D(z)}(y)$ and $w \in \text{Proj}_{D(x)}(u)$, we have

$$\text{dist}_{D(x)}(y) - \text{dist}_{D(z)}(y) = \text{dist}_{D(x)}(y) - d(u, y) \leq d(w, y) - d(u, y) \leq d(w, u)$$

so

$$(7.11) \quad \text{dist}_{D(x)}(y) - \text{dist}_{D(z)}(y) \leq \text{dist}_{D(x)}(u).$$

Since $z \in D(z)$, we deduce $d(u, y) \leq d(z, y)$. So when x, y, z are all near \bar{x} , u is also near \bar{x} , and our assumption implies

$$\begin{aligned} \text{dist}_{D(x)}(y) - \text{dist}_{D(z)}(y) &\leq \text{dist}_{D(x)}(u) \leq \gamma d(x, z)(d(x, u) + d(z, u)) \\ &\leq \gamma d(x, z)(d(x, y) + d(u, y) + d(z, y) + d(u, y)) \\ &\leq \gamma d(x, z)(d(x, y) + 3d(z, y)) \\ &\leq 3\gamma d(x, z)(d(x, y) + d(z, y)). \end{aligned}$$

The approximator property (7.3) follows by symmetry, with the constant 3γ . \square

Example 7.5 illustrates exactly this kind of approximator.

Remark 7.12 The Lipschitz continuity of the function h defined in Lemma 7.10 amounts to *pseudo-Lipschitz continuity* for the mapping D : see [39].

In the development that follows, in addition to assuming that the feasible region X is a proper metric space, we suppose that it is a closed subset of a possibly larger metric space \bar{X} , with the inherited metric. The role of \bar{X} is purely to deploy some convexity arguments. We use the idea of a *geodesic* (of unit speed), which is an isometry from a compact interval of \mathbf{R} into \bar{X} ; the *endpoints* and *midpoint* are the images of the endpoints of the interval and its midpoint respectively. In our discussions, we often identify a geodesic with its image. We call \bar{X} a *geodesic space* if every pair of points are the endpoints of a geodesic. A subset D is *geodesically convex* if D contains all geodesics whose endpoints lie in D . In that case, a function $g: D \rightarrow \mathbf{R}$ is *geodesically convex* if its composition with every geodesic in D is convex, and is *geodesically μ -strongly convex* (for $\mu > 0$) if its composition with every geodesic in D is μ -strongly convex.

A geodesic space \bar{X} is a CAT(0) *space* when, for all points $x \in \bar{X}$, the function $\frac{1}{2}d^2(x, \cdot)$ is geodesically 1-strongly convex. This latter property is crucial for the proximal point iteration [10, 28]. In particular, it implies that balls in \bar{X} are geodesically convex. Complete, locally compact geodesic spaces are proper as a consequence of a suitable version of the Hopf-Rinow Theorem [36, Theorem 7.2] and [18, Proposition I.3.7]. Examples of such spaces that are also CAT(0) include all finite-dimensional Riemannian manifolds that are complete, simply connected, and having nonpositive sectional curvature.

For convenience, we collect our assumptions below.

Assumption 7.13 *The feasible region X is a proper metric subspace of a geodesic space, the objective function $f: X \rightarrow \mathbf{R}$ is bounded below, and the optimization problem $\inf_X f$ admits the following majorization-minimization tools.*

- An **approximating mapping** $D: X \rightrightarrows X$ that assigns to each feasible point $x \in X$ a closed and geodesically convex subset $D(x)$ that contains x .
- A constant $\mu \geq 0$ and a **modeling function** $h: X^2 \rightarrow \mathbf{R}$ assigns to each $x \in X$ a geodesically μ -strongly convex function $h_x: D(x) \rightarrow \mathbf{R}$ satisfying

$$h(x, y) = h_x(y) \geq f(y) \quad \text{for } y \in D(x), \quad \text{with equality if } y = x.$$

Furthermore, both h and the function defined by

$$(x, y) \mapsto \text{dist}_{D(x)}(y) \quad (x, y \in X)$$

are approximators in the sense of Definition 7.2.

Before proceeding, we make some observations. Implicit in the assumption (as a consequence of the approximator properties) is that the feasible region X is closed and the objective f is locally Lipschitz. However, we emphasize that neither need be geodesically convex. The assumption furthermore ensures that the approximation mapping D has closed graph, and in fact is continuous, since the function $(x, y) \mapsto \text{dist}_{D(x)}(y)$ is continuous. We also see that each approximating set $D(x)$ is a closed second-order approximation of X , and each model h_x is a continuous second-order approximation of f : for points $y \in X$ near x we have

$$\text{dist}_{D(x)}(y) = O(d^2(x, y)) \quad \text{and} \quad (h_x - f)(y) = O(d^2(x, y)).$$

Example 7.14 (Proximal Point Method) Consider a proper CAT(0) space X and a geodesically convex, locally Lipschitz function $f: X \rightarrow \mathbf{R}$ that is bounded below. Let $D(x) = X$ for all points $x \in X$, and for any constant $\mu > 0$, define the associated *quadratic penalty modeling function*

$$(7.15) \quad h(x, y) = f(y) + \frac{\mu}{2}d^2(x, y).$$

Then Assumption 7.13 holds. The associated envelope is the Moreau-Yosida approximation of f , the iteration map is the metric version of the proximal point iteration [10], and MM sequences are just proximal point sequences.

In any proper CAT(0) space X , a nontrivial example of an approximating mapping is defined, for any radius $\rho > 0$, by $D(x) = B_\rho(x)$, as seen in Example 7.5.

MM sequences are just fixed point iterations associated with the iteration map p . We first study the fixed point property $x \in p(x)$. Clearly, every minimizer of the objective f is a fixed point of p . Moreover, we have the following converse result.

Proposition 7.16 (Fixed points and zero slope) *If Assumption 7.13 holds, then at any fixed point of the iteration map, the objective has slope zero.*

Proof Consider any fixed point x . If the objective f has strictly positive slope, then there exists a constant $\sigma > 0$ and a sequence of points $x_k \neq x$ converging to x in X and all satisfying

$$\frac{f(x) - f(x_k)}{d(x, x_k)} > \sigma.$$

By our assumptions, there exist points $y_k \in D(x)$ satisfying $d(y_k, x_k) = o(d(x_k, x))$. The fixed point property ensures $f(x) \leq h_x(y_k)$, so for all large k we have

$$\frac{h_x(x_k) - f(x_k)}{d(x, x_k)} = \frac{h_x(y_k) - f(x_k)}{d(x, x_k)} > \sigma,$$

using Lipschitz property of the model h_x . But this contradicts the implication from our assumptions that $h_x(x_k) - f(x_k) = o(d(x_k, x))$. \square

Proposition 7.17 *On a proper geodesic space, if a geodesically strongly convex function is continuous and bounded below, then it has a unique minimizer.*

Proof Denote the given function by g , and the strong convexity constant by μ . Fix any point $\bar{x} \in X$. At any point $x \in X$ satisfying $g(x) \leq g(\bar{x})$, consider a geodesic with endpoints x and \bar{x} . The midpoint x' satisfies

$$\inf g \leq g(x') \leq \frac{1}{2}g(x) + \frac{1}{2}g(\bar{x}) - \frac{\mu}{8}d^2(x, \bar{x}) \leq g(\bar{x}) - \frac{\mu}{8}d^2(x, \bar{x}),$$

so x lies in the closed ball defined by $d^2(x, \bar{x}) \leq \frac{8}{\mu}(g(\bar{x}) - \inf g)$. By continuity, g has a minimizer \hat{x} over this compact set, and \hat{x} must minimize g globally. The preceding inequality shows uniqueness. \square

This result holds more generally if the space is CAT(0) [34, Lemma 1.3].

Proposition 7.18 *Suppose that Assumption 7.13 holds. Then the iteration map is single-valued and continuous. Furthermore, for any associated MM sequence that is bounded, there exists a constant $\beta \geq 0$ such that the envelope F at all successive triples of iterates x_-, x, x_+ satisfies the descent conditions*

$$(7.19) \quad F(x_-) - F(x) \geq \frac{\mu}{2}d^2(x, x_+),$$

$$(7.20) \quad |\nabla F|(x) \leq \beta d(x, x_+).$$

Proof The previous result shows that the iteration map p is single-valued. We next prove the quadratic growth condition

$$(7.21) \quad h_x(y) - h_x(p(x)) \geq \frac{\mu}{2}d^2(y, p(x)) \quad \text{whenever } y \in D(x).$$

To see this, for any point $y \neq p(x)$ in the geodesically convex set $D(x)$, consider the distance $\delta = d(p(x), y) > 0$ and a geodesic $\gamma: [0, \delta] \rightarrow D(x)$ satisfying $\gamma(0) = p(x)$ and $\gamma(\delta) = y$. By μ -strong convexity, for $\tau \in [0, \delta]$ we have $\gamma(t) \in D(x)$ and

$$h_x(\gamma(\tau)) - \frac{\mu}{2}\tau^2 \leq \left(1 - \frac{\tau}{\delta}\right)h_x(p(x)) + \frac{\tau}{\delta}\left(h_x(y) - \frac{\mu}{2}\delta^2\right).$$

Since $p(x)$ minimizes h_x on $D(x)$, for small $\tau > 0$ we deduce

$$0 \leq h_x(\gamma(\tau)) - h_x(p(x)) \leq \frac{\tau}{\delta}\left(h_x(y) - h_x(p(x)) - \frac{\mu}{2}\delta^2\right) + \frac{\mu}{2}\tau^2.$$

Dividing by τ and taking the limit as $\tau \downarrow 0$ proves the growth condition (7.21).

Consider any sequence (x_k) in X converging to \bar{x} . Since $p(\bar{x}) \in D(\bar{x})$, the continuity of D ensures the existence of points $y_k \in D(x_k)$ satisfying $y_k \rightarrow p(\bar{x})$. The growth condition (7.21) implies, for all k , the inequality

$$h_{x_k}(y_k) - h_{x_k}(p(x_k)) \geq \frac{\mu}{2}d^2(y_k, p(x_k)).$$

Taking the lim sup of both sides implies $d(y_k, p(x_k)) \rightarrow 0$, and hence $p(x_k) \rightarrow p(\bar{x})$. Thus p is continuous.

At any point $x \in X$, inequality (7.21) implies

$$\begin{aligned} F(x) - F(p(x)) &= h_x(p(x)) - h_{p(x)}(p(p(x))) \\ &\geq h_{p(x)}(p(x)) - h_{p(x)}(p(p(x))) \geq \frac{\mu}{2}d^2(p(x), p(p(x))). \end{aligned}$$

This proves the descent condition (7.19).

To prove the final descent condition (7.20), we can assume $|\nabla F|(x) > 0$, since otherwise there is nothing to prove, so there exists a sequence of points $x_k \rightarrow x$ in X with

$$|\nabla F|(x) = \lim_k \frac{F(x) - F(x_k)}{d(x, x_k)} = \lim_k \frac{h(x, p(x)) - h(x_k, p(x_k))}{d(x, x_k)}.$$

Consider the sequence

$$\alpha_k = d(x, p(x_k)) + d(x_k, p(x_k)).$$

Since $p(x_k) \in D(x_k)$ and the function $(x, y) \mapsto \text{dist}_{D(x)}(y)$ is an approximator, we know for some constant $\gamma \geq 0$ the inequality

$$\text{dist}_{D(x)}(p(x_k)) \leq \gamma \alpha_k d(x, x_k)$$

for all k , so there exists a bounded sequence of points $y_k \in D(x)$ satisfying

$$d(y_k, p(x_k)) \leq \gamma \alpha_k d(x, x_k).$$

Increasing γ if necessary, we can also ensure, for all k the inequalities

$$\begin{aligned} h(x, p(x_k)) - h(x_k, p(x_k)) &\leq \gamma \alpha_k d(x, x_k) \\ h(x, y_k) - h(x, p(x_k)) &\leq \gamma d(y_k, p(x_k)), \end{aligned}$$

since h is also an approximator. Using our assumptions and the continuity of p , we deduce

$$\begin{aligned} |\nabla F|(x) &\leq \liminf_k \frac{h(x, y_k) - h(x_k, p(x_k))}{d(x, x_k)} \\ &\leq \liminf_k \frac{h(x, p(x_k)) + \gamma d(y_k, p(x_k)) - h(x_k, p(x_k))}{d(x, x_k)} \\ &\leq \gamma(\gamma + 1) \liminf_k \alpha_k = 2\gamma(\gamma + 1) d(x, p(x)). \end{aligned}$$

This completes the proof. □

Following the general framework promoted in [23], this result serves to justify the natural “small step” stopping criterion for MM algorithms, since if two successive iterates x, x_+ are close, then the slope of the envelope at x must be small.

Theorem 7.22 (MM sequence convergence) *Suppose that Assumption 7.13 holds. If the corresponding envelope, when restricted to closed balls, always satisfies the KL property for minimization, with a concave desingularizer, then any bounded MM sequence (x_k) with initial value $f(x_0)$ near $\inf f$ converges to a minimizer of f .*

Proof The proof follows similar techniques in [16, 38] in imitating that of Theorem 5.4. By our previous results, successive triples of iterates $\dots x_-, x, x_+ \dots$ satisfy

$$(7.23) \quad F(x_-) - F(x) \geq \frac{\mu}{2} d^2(x, x_+) \quad \text{and} \quad d(x, x_+) \geq \beta |\nabla F|(x).$$

We can assume that no iterate minimizes F , since otherwise the first inequality would ensure that the sequence is constant thereafter. Define the gap function $g = F - \inf F$. The envelope values $F(x_k)$ decrease monotonically, and the KL inequality ensures that the gap g satisfies

$$1 \leq |\nabla(\phi \circ g)|(x_k) = \phi'(g(x_k)) \cdot |\nabla g|(x_k).$$

We deduce that each slope $|\nabla g|(x_k) = |\nabla F|(x_k)$ is nonzero, and so the distances $d_k = d(x_{k-1}, x_k)$ are nonzero for all $k > 0$.

Using the concavity of the desingularizer ϕ , we observe, for all $k > 0$,

$$\phi(g(x_{k-1})) - \phi(g(x_k)) \geq \phi'(g(x_{k-1}))(g(x_{k-1}) - g(x_k)) \geq \frac{F(x_{k-1}) - F(x_k)}{|\nabla F|(x_{k-1})} \geq \frac{\mu d_{k+1}^2}{2\beta d_k},$$

so

$$\frac{2\beta}{\mu} \left(\phi(g(x_{k-1})) - \phi(g(x_k)) \right) \geq \frac{2d_k d_{k+1} - d_k^2}{d_k} \geq 2d_{k+1} - d_k.$$

We deduce that the positive numbers $\lambda_k = \frac{2\beta}{\mu} \phi(g(x_{k-1})) + d_k$ satisfy the key inequality $d(x_k, x_{k+1}) \leq \lambda_k - \lambda_{k+1}$, so the sequence (x_k) , being Cauchy, converges to some limit x^* . Theorem 3.2 shows that (x_k) is a minimizing sequence for F and hence also for f , so by continuity we deduce that x^* minimizes f , as required. \square

Remark 7.24 This result implies the convergence of MM sequences for semi-algebraic optimization in broad generality. More precisely, suppose that the feasible region is a semi-algebraic subset of a Euclidean space, and that the approximating mapping and modeling function are also semi-algebraic. In that case the envelope F is also semi-algebraic, and so satisfies the condition.

Theorem 7.25 (Global convergence of MM sequences) *Suppose that Assumption 7.13 holds. If the corresponding envelope is a KL function with a concave desingularizer, as holds in particular if the feasible region, approximating mapping and modeling function are all semi-algebraic, then any bounded MM sequence converges to a point where the objective has zero slope.*

Proof We follow the proof of Theorem 6.5, but instead using the inequalities (7.23) as in the previous proof. If the objective f has bounded level sets, then any MM sequence must be bounded because the corresponding objective values decrease monotonically. The proof shows that the sequence generated by the iteration $x \leftarrow p(x)$ converges to some limit x^* . Since the iteration map p is continuous, x^* must be a fixed point, from which we deduce $|\nabla f|(x^*) = 0$ by Proposition 7.16. \square

Note that MM sequences are always bounded if the objective has bounded level sets. We end with two illustrations from the literature.

Convex-composite optimization [31]

Given Euclidean spaces X and W , a $\mathcal{C}^{(2)}$ -smooth map $G: X \rightarrow W$, a convex function $q: W \rightarrow \mathbf{R}$ and any constant $\mu > 0$, we consider the composite objective $f: X \rightarrow \mathbf{R}$,

a function $H: X^2 \rightarrow W$ describing the linear approximation to G , and the “prox-linearized” function $h: X^2 \rightarrow \mathbf{R}$, defined for points x, y in X by

$$\begin{aligned} f(x) &= q(G(x)), \\ H(x, y) &= G(x) + DG(x)[y - x], \\ h(x, y) &= q(H(x, y)) + \frac{\mu}{2}|x - y|^2, \end{aligned}$$

where the linear map $DG(x): X \rightarrow W$ is the derivative for G at x . With the trivial approximating mapping defined by $D(x) = X$ for all points $x \in X$, we show under reasonable conditions that Assumption 7.13 holds.

We suppose that f is bounded below. We assume too that q is L -Lipschitz, and that the second derivative of G , the bilinear map $D^2G(x): X^2 \rightarrow W$, has norm bounded by a constant M for all x . (For simplicity of exposition, these are global assumptions; a more refined development allows the constants to depend on a given bounded domain.) For any point $x \in X$ and unit vector $w \in W$, consider the function $e: X \rightarrow \mathbf{R}$ defined by

$$e(y) = \langle w, H(x, y) - G(y) \rangle \quad (y \in X).$$

Notice for all vectors $u \in X$ we have

$$De(y)[u] = \langle w, (DG(x) - DG(y))[u] \rangle.$$

By the Mean Value Theorem, there exists a point $z \in [x, y]$ such that

$$e(y) = e(x) + De(z)[y - x] = \langle w, (DG(x) - DG(z))[y - x] \rangle,$$

and furthermore a point $u \in [x, z]$ such that

$$(DG(x) - DG(z))[y - x] = D^2G(u)[y - x, x - z].$$

Hence

$$|G(x) + DG(x)[y - x] - G(y)| \leq M|y - x||x - z|,$$

so

$$\left| q(G(y)) - q(G(x) + DG(x)[y - x]) \right| \leq LM|x - y|^2.$$

Providing $\mu > 2LM$ we deduce the majorization property.

Turning to the remaining assumptions, fix a point $y \in X$. We first prove that all points $x, z \in X$ satisfy

$$(7.26) \quad |H(x, y) - H(z, y)| \leq M|x - z|(|x - y| + |z - y|).$$

Fix any unit vector $w \in W$, and define a function $c: X \rightarrow \mathbf{R}$ by $c(x) = \langle w, H(x, y) \rangle$. Then, by the Mean Value Theorem, there exists a point $u \in [x, z]$ such that

$$c(x) - c(z) = Dc(u)[x - z] = \langle w, D^2G(u)[y - u, x - z] \rangle.$$

Since w was arbitrary, our claim now follows, since $|y - u| \leq |x - y| + |z - y|$.

Using inequality (7.26), we deduce

$$\begin{aligned} |h(x, y) - h(z, y)| &\leq LM|x - z|(|x - y| + |z - y|) + \frac{\mu}{2}||x - y|^2 - |z - y|^2| \\ &= LM|x - z|(|x - y| + |z - y|) + \mu|\langle x - z, x + z - 2y \rangle| \\ &\leq \left(LM + \frac{\mu}{2}\right)|x - z|(|x - y| + |z - y|). \end{aligned}$$

Assumption 7.13 follows.

Suppose in addition that the functions q and G are semi-algebraic. Then for all sufficiently large constants $\mu > 0$, any bounded sequence generated by the prox-linear iteration

$$x \leftarrow \operatorname{argmin}_y \left\{ q(G(x) + DG(x)[y - x]) + \frac{\mu}{2}|x - y|^2 \right\}$$

converges to a point where the composite function $q \circ G$ has slope zero.

Moving balls method [7, 16]

For differentiable functions $f, f_1, f_2, \dots, f_m: \mathbf{R}^n \rightarrow \mathbf{R}$ with Lipschitz gradients, we consider the problem of minimizing the objective f over the feasible region

$$X = \{x \in \mathbf{R}^n : f_i(x) \leq 0 \text{ for each } i\}.$$

We assume that this classical nonlinear program is bounded below, and that the Mangasarian-Fromovitz constraint qualification (MFCQ) holds: in other words, for all points $x \in X$, there exists a vector $d \in \mathbf{R}^n$ satisfying

$$\langle \nabla f_i(x), d \rangle < 0 \quad \text{for all } i \in I(x) = \{i : f_i(x) = 0\}.$$

With the aim of constructing an optimization method that maintains feasibility, our majorization-minimization tools consist of an approximating mapping D and a modeling function h defined by

$$\begin{aligned} h(x, y) &= f(x) + \langle \nabla f(x), y - x \rangle + \frac{L}{2}\|y - x\|^2, \\ h_i(x, y) &= f_i(x) + \langle \nabla f_i(x), y - x \rangle + \frac{L_i}{2}\|y - x\|^2 \quad \text{for each } i, \\ D(x) &= \bigcap_i D_i(x) \quad \text{with } D_i(x) = \{y \in \mathbf{R}^n : h_i(x, y) \leq 0\} \text{ for each } i, \end{aligned}$$

where L, L_1, \dots, L_m are the Lipschitz moduli of $\nabla f, \nabla f_1, \dots, \nabla f_m$ respectively. For each point $x \in X$, the set $D_i(x)$ is a closed ball centered at the point $x - \frac{1}{L_i}\nabla f_i(x)$ and with radius

$$\rho_i(x) = \sqrt{\frac{1}{L_i^2}\|\nabla f_i(x)\|^2 - \frac{2}{L_i}f_i(x)}.$$

Hence the name of the method. Our aim is to show that Assumption 7.13 holds, ensuring convergence of the method under reasonable conditions.

It suffices to verify that both the functions h and $(x, y) \mapsto \text{dist}_{D(x)}(y)$ (for points $x, y \in \mathbf{R}^n$) are approximators, since the other conditions easily follow from the standard quadratic upper bound for L -smooth functions $g: \mathbf{R}^n \rightarrow \mathbf{R}$:

$$g(y) \leq g(x) + \langle \nabla g(x), y - x \rangle + \frac{L}{2} \|y - x\|^2.$$

The fact that h is an approximator follows from Examples 7.4 and 7.6, along with Proposition 7.7. In fact, the following inequality is easy to verify directly that all points $x, y, z \in \mathbf{R}^n$ satisfy the inequality

$$(7.27) \quad |h(x, y) - h(z, y)| \leq \frac{3L}{2} |x - z| (|x - y| + |z - y|)$$

along with the corresponding inequality for the functions h_i .

To show that the function $(x, y) \mapsto \text{dist}_{D(x)}(y)$ is also an approximator, we need the following elementary fact [37]: for closed convex sets $S_1, S_2, \dots, S_m \subset \mathbf{R}^n$, each containing the ball $B_\delta(0)$, all points $x \in \mathbf{R}^n$ satisfy

$$(7.28) \quad \text{dist}_{\cap_i S_i}(x) \leq \frac{\|x\|}{\delta} \max_i \{\text{dist}_{S_i}(x)\}.$$

Fix a point $\bar{x} \in X$. By the MFCQ condition, there exists a unit vector $d \in \mathbf{R}^n$ satisfying

$$\langle \nabla f_i(\bar{x}), d \rangle < 0 \quad \text{for all } i \in I(\bar{x}).$$

For any constant $\epsilon \in (0, \min_i 2L_i)$ sufficiently small, all points $x \in B_\epsilon(\bar{x})$ satisfy the inequalities

$$(7.29) \quad f_i(x) \leq -\epsilon \quad \text{for all } i \notin I(\bar{x}) \quad \text{and} \quad \nabla f_i(x)^\top d \leq -\epsilon \quad \text{for all } i \in I(\bar{x}).$$

and, for sufficiently large $M \geq \max_i L_i$, the inequalities $\|\nabla f_i(x)\| \leq M$ for each $i = 1, 2, \dots, m$.

Define constants $t = \min\{1, \frac{\epsilon}{4M}\}$ and $\delta = \min\{1, \frac{\epsilon}{2M}\}t$. For any point $x \in B_\epsilon(\bar{x}) \cap X$ and any vector $e \in B_\delta(0)$, consider the quantity

$$h_i(x, x + td + e) = f_i(x) + \langle \nabla f_i(x), td + e \rangle + \frac{L_i}{2} \|td + e\|^2.$$

For all $i \notin I(\bar{x})$, the left-hand side is no larger than $-\epsilon + 4Mt$, and for all $i \in I(\bar{x})$ it is no larger than $-\frac{\epsilon}{2}t + 2Mt^2$, so in either case it is nonpositive. We deduce

$$B_\delta(x + td) \subset \bigcap_i D_i(x) = D(x).$$

Property (7.28) therefore implies that all points $y \in \mathbf{R}^n$ satisfy the inequality

$$(7.30) \quad \text{dist}_{D(x)}(y) \leq \frac{\|y - x - td\|}{\delta} \max_i \{\text{dist}_{D_i(x)}(y)\}.$$

Consider points $x \in B_\epsilon(\bar{x}) \cap X$, $z \in X$ and $y \in D(z)$. If $y \notin D_i(x)$, then

$$\begin{aligned} \text{dist}_{D_i(x)}(y) &= \left\| y - x + \frac{1}{L_i} \nabla f_i(x) \right\| - \rho_i(x) \\ &= \frac{2}{L_i} \cdot \frac{h_i(x, y)}{\left\| y - x + \frac{1}{L_i} \nabla f_i(x) \right\| + \rho_i(x)} \\ &\leq \frac{2}{L_i} \cdot \frac{h_i(x, y) - h_i(z, y)}{\left\| y - x + \frac{1}{L_i} \nabla f_i(x) \right\| + \rho_i(x)} \\ &\leq \frac{h_i(x, y) - h_i(z, y)}{\sqrt{\|\nabla f_i(x)\|^2 - 2L_i f_i(x)}} \leq \frac{1}{\epsilon} (h_i(x, y) - h_i(z, y)) \\ &\leq \frac{3L_i}{2\epsilon} |x - z| (|x - y| + |z - y|), \end{aligned}$$

the penultimate inequality following from the inequalities (7.29). Property (7.30) now implies

$$\text{dist}_{D(x)}(y) \leq \frac{3M\|y - x - td\|}{2\epsilon\delta} |x - z| (|x - y| + |z - y|).$$

Since the above inequality holds for all $y \in D(z)$, it along with (7.11) implies the local Lipschitz continuity of the function $(x, y) \mapsto \text{dist}_{D(x)}(y)$. Hence, by Lemma 7.10, the function $(x, y) \mapsto \text{dist}_{D(x)}(y)$ is an approximator, completing our verification of Assumption 7.13.

References

- [1] P.-A. Absil, R. Mahony, and B. Andrews. Convergence of the iterates of descent methods for analytic cost functions. *SIAM J. Optim.*, 6(2):531–547, 2005.
- [2] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, 2008.
- [3] L. Ambrosio, N. Gigli, and G. Savaré. *Gradient Flows in Metric Spaces and in the Space of Probability Measures*. Birkhäuser-Verlag, Basel, 2nd edition, 2008.
- [4] H. Attouch and J. Bolte. On the convergence of the proximal algorithm for nonsmooth functions involving analytic features. *Math. Program.*, 116:5–16, 2009.

- [5] H. Attouch, J. Bolte, P. Redont, and A. Soubeyran. Proximal alternating minimization and projection methods for nonconvex problems: an approach based on the Kurdyka-Lojasiewicz inequality. *Mathematics of Operations Research*, 35:438–457, 2010.
- [6] 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. *Math. Program.*, 137(1-2, Ser. A):91–129, 2013.
- [7] A. Auslender, R. Shefi, and M. Teboulle. A moving balls approximation method for a class of smooth constrained minimization problems. *SIAM Journal on Optimization*, 20(6):3232–3259, 10 2010.
- [8] D. Azé and J.-N. Corvellec. Characterizations of error bounds for lower semicontinuous functions on metric spaces. *ESAIM: COCV*, 10:409–425, 2004.
- [9] D. Azé and J.-N. Corvellec. Nonlinear error bounds via a change of function. *Journal of Optimization Theory and Applications*, 172:9–32, 2017.
- [10] M. Bačák. The proximal point algorithm in metric spaces. *Isr. J. Math.*, 194:689–701, 2013.
- [11] A. Blanchet and J. Bolte. A family of functional inequalities: Lojasiewicz inequalities and displacement convex functions. *Journal of Functional Analysis*, 275:1650–1673, 2018.
- [12] J. Bolte, A. Daniilidis, and A.S. Lewis. The Łojasiewicz inequality for nonsmooth subanalytic functions with applications to subgradient dynamical systems. *SIAM J. Optim.*, 17:1205–1223, 2006.
- [13] J. Bolte, A. Daniilidis, A.S. Lewis, and M. Shiota. Clarke subgradients of stratifiable functions. *SIAM J. Optimization*, 18(2):556–572, 2007.
- [14] J. Bolte, A. Daniilidis, O. Ley, and L. Mazet. Characterizations of Lojasiewicz inequalities: subgradient flows, talweg, convexity. *Trans. Amer. Math. Soc.*, 362(6):3319–3363, 2010.
- [15] J. Bolte, T.P. Nguyen, J. Peypouquet, and B. Suter. From error bounds to the complexity of first order descent methods for convex functions. *Math. Programming*, 165:471–507, 2017.
- [16] J. Bolte and E. Pauwels. Majorization-minimization procedures and convergence of SQP methods for semi-algebraic and tame programs. *Mathematics of Operations Research*, 41:442–465, 2016.
- [17] N. Boumal. *An Introduction to Optimization on Smooth Manifolds*. Cambridge University Press, Cambridge, 2022.
- [18] M.R. Bridson and A. Haefliger. *Metric Spaces of non-positive curvature*. Springer-Verlag, Berlin, 1999.

- [19] M. Coste. *An Introduction to o-minimal Geometry*. RAAG Notes, 81 pages, Institut de Recherche Mathématiques de Rennes, November 1999.
- [20] E. De Giorgi, A. Marino, and M. Tosques. Problems of evolution in metric spaces and maximal decreasing curve. *Atti Accad. Naz. Lincei Rend. Cl. Sci. Fis. Mat. Nat. (8)*, 68:180–187, 1980.
- [21] A.P. Dempster, N.M. Laird, and D.B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *J. Royal Statist. Soc., Ser. B*, 39:1–38, 1977.
- [22] D. Drusvyatskiy, A.D. Ioffe, and A.S. Lewis. Curves of descent. *SIAM J. Control Optim.*, 53(1):114–138, 2015.
- [23] 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.
- [24] I. Ekeland. On the variational principle. *J. Math. Anal. Appl.*, 47:324–353, 1974.
- [25] O. Güler. On the convergence of the proximal point algorithm for convex minimization. *SIAM Journal on Control and Optimization*, 29:403–419, 1991.
- [26] D. Hauer and J.M. Mazón. Kurdyka–Lojasiewicz–Simon inequality for gradient flows in metric spaces. *Transactions of the American Mathematical Society*, 372:4917–4976, 2019.
- [27] A.D Ioffe. *Variational Analysis of Regular Mappings*. Springer International Publishing, 2017.
- [28] J. Jost. Convex functionals and generalized harmonic maps into spaces of non positive curvature. *Comment. Math. Helvetici*, 70:659–673, 1995.
- [29] K. Kurdyka. On gradients of functions definable in o-minimal structures. *Ann. Inst. Fourier (Grenoble)*, 48(3):769–783, 1998.
- [30] A.S. Lewis and Tonghua Tian. Identifiability, the KL property in metric spaces, and subgradient curves. [arXiv:2205.02868](https://arxiv.org/abs/2205.02868), 2022.
- [31] A.S. Lewis and S.J. Wright. A proximal method for composite minimization. *Math. Program.*, pages 1–46, 2015.
- [32] S. Lojasiewicz. Ensembles semi-analytiques. Preprint, IHES, 1965.
- [33] Z.-Q. Luo and P. Tseng. Error bounds and convergence analysis of feasible descent methods: a general approach. *Ann. Oper. Res.*, 46/47(1-4):157–178, 1993. Degeneracy in optimization problems.
- [34] U.F. Mayer. Gradient flows on nonpositively curved metric spaces and harmonic maps. *Communications in Analysis and Geometry*, 6:199–253, 1998.

- [35] G.J. Minty. Monotone (nonlinear) operators in Hilbert space. *Duke Mathematical Journal*, 29:341–346, 1962.
- [36] S.B. Myers. Arcs and geodesics in metric spaces. *Transactions of the American Mathematical Society*, 57:217–227, 1945.
- [37] A. Nedić. Random projection algorithms for convex set intersection problems. In *49th IEEE Conference on Decision and Control (CDC)*, pages 7655–7660, 2010.
- [38] D. Noll. Convergence of non-smooth descent methods using the Kurdyka–Lojasiewicz inequality. *Journal of Optimization Theory and Applications*, 160:553–572, 2014.
- [39] R.T. Rockafellar. Lipschitzian properties of multifunctions. *Nonlinear Analysis*, 9:867–885, 1985.
- [40] R.T. Rockafellar and R.J-B. Wets. *Variational Analysis*. Grundlehren der mathematischen Wissenschaften, Vol 317, Springer, Berlin, 1998.
- [41] A Valette. Łojasiewicz inequality at singular points. *Proc. Amer. Math. Soc.*, 147:1109–1117, 2019.
- [42] L. van den Dries. *Tame topology and o-minimal structures*, volume 248 of *LMS Lecture Note Series*. Cambridge University Press, Cambridge, 1998.
- [43] Hongyi Zhang and S. Sra. First-order methods for geodesically convex optimization. *Journal of Machine Learning Research*, 49:1–22, 2016.