

Survey Descent: A Multipoint Generalization of Gradient Descent for Nonsmooth Optimization

X.Y. Han*

A.S. Lewis†

Abstract

For strongly convex objectives that are smooth, the classical theory of gradient descent ensures linear convergence relative to the number of gradient evaluations. An analogous nonsmooth theory is challenging. Even when the objective is smooth at every iterate, the corresponding local models are unstable and the number of cutting planes invoked by traditional remedies is difficult to bound, leading to convergences guarantees that are sublinear relative to the cumulative number of gradient evaluations. We instead propose a multipoint generalization of the gradient descent iteration for local optimization. While designed with general objectives in mind, we are motivated by a “max-of-smooth” model that captures the subdifferential dimension at optimality. We prove linear convergence when the objective is itself max-of-smooth, and experiments suggest a more general phenomenon.

1 Introduction

To approach our target—a fast local iteration for nonsmooth optimization—let us recall, as motivation, the classical foundations of gradient descent (GD). For some function h , denote the linear approximation of h centered at some differentiable iterate \tilde{x} by

$$\ell_{\tilde{x}}^h(x) \equiv h(\tilde{x}) + \nabla h(\tilde{x})^T(x - \tilde{x}),$$

where $\nabla h(\cdot)$ is the gradient of h . When h is smooth, given a step-size $\frac{1}{L}$, the canonical GD step $\tilde{x}^+ = \tilde{x} - \frac{1}{L}\nabla h(\tilde{x})$ produces the global minimum of the following local model.

$$h_{\tilde{x}}^{\text{GD}}(x) = \ell_{\tilde{x}}^h(x) + \frac{L}{2} \|x - \tilde{x}\|_2^2. \quad (1)$$

The model $h_{\tilde{x}}^{\text{GD}}$ incorporates the local linear behavior observed around \tilde{x} through the gradient $\nabla h(\tilde{x})$ as well as the prior belief that the gradients themselves are L -Lipschitz functions of x . After iterating, one then builds a new model around \tilde{x}^+ and repeats the procedure. Indeed—on convex, L -smooth objectives—canonical results (as described in standard texts such as Beck (2017, Chapter 10)) guarantee that every GD iteration will both reduce the objective value and move closer to the true global minimizer of h .

When the Hessian is available, one can alternatively consider the Newton model

$$h_{\tilde{x}}^{\text{Newton}}(x) = \ell_{\tilde{x}}^h(x) + \frac{1}{2} (x - \tilde{x})^T \nabla^2 h(\tilde{x}) (x - \tilde{x}). \quad (2)$$

For first-order methods, however, the Hessian $\nabla^2 h(\tilde{x})$ is inaccessible. This motivates popular quasi-Newton approaches such as the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm (as described in standard texts such as Nocedal and Wright (2006, Chapter 6)) that minimizes an approximation to (2) after estimating $\nabla^2 h(\tilde{x})$ using only gradient differences and then performs a line-search.

For nonsmooth objectives, both of the above smoothness-hypothesizing quadratic models—and their variants—are inadequate for generating theoretically-guaranteed improvement at every step: They fail to

*ORIE, Cornell University, Ithaca, NY. xh332@cornell.edu

†ORIE, Cornell University, Ithaca, NY. people.orie.cornell.edu/aslewis

Research supported in part by National Science Foundation Grant DMS-2006990

capture the discontinuity of objective gradients. Thus, some algorithms work instead with a richer class of model functions. For example, one class of multipoint methods (discussed later in Section 1.4) adaptively refines a lower cutting-plane model of the objective through a series of “null steps” around each iterate until it achieves descent in a “serious step”. Such methods work well under reasonable conditions, finitely-many null steps always sufficing to construct a successful serious step. However, analyzing convergence relative to all steps (null and serious) is challenging. Current bounds on the cumulative number of steps remain sublinear due to the difficulty of uniformly bounding the number of “in-between” null-steps.

1.1 The Survey Descent Iteration

In this work, we will propose and analyze a new local *Survey Descent* iteration for nonsmooth objectives h .

Definition 1. (Survey Descent Iteration $\{(P_i^{\mathcal{S}})\}_{i=1}^k$) Given a *survey* of k points, $\mathcal{S} = \{s_i\}_{i=1}^k$, at which the nonsmooth objective h is differentiable and a step-control parameter L , for each $i = 1, \dots, k$, define the i -th *subproblem* $(P_i^{\mathcal{S}})$ as follows:

$$\min_x \left\| x - \left(s_i - \frac{1}{L} \nabla h(s_i) \right) \right\|_2^2 \quad (3)$$

$$\text{s.t. } \ell_{s_j}^h(x) + \frac{L}{2} \|x - s_j\|_2^2 \leq \ell_{s_i}^h(x) \quad \forall j \neq i. \quad (4)$$

We refer to the solving of all subproblems $\{(P_i^{\mathcal{S}})\}_{i=1}^k$ as a *Survey Descent iteration*.

When $(P_i^{\mathcal{S}})$ is feasible, we denote its optimal solution as s_i^+ . When all subproblems are feasible, we say that the entire Survey Descent iteration is *feasible* and call $\mathcal{S}^+ = \{s_i^+\}_{i=1}^k$ the *outputs* of the iteration; otherwise, we say the entire iteration is *infeasible*. After a feasible iteration, we would update $\mathcal{S} \leftarrow \mathcal{S}^+$ and repeat the Survey Descent iteration on the updated survey.

Observe that, if $k=1$, Survey Descent reduces to GD since there is then only one subproblem $(P_1^{\mathcal{S}})$ and survey point s_1 , the constraints (4) are then empty, and the objective (3) of $(P_1^{\mathcal{S}})$ is then minimized by $s_1 - \frac{1}{L} \nabla h(s_1)$, which is the GD-step from s_1 . Thus, we can consider Survey Descent a generalization of GD.

1.1.1 Main Results

When the function h is a maximum of k smooth functions, we prove local linear convergence of the Survey Descent iteration to the minimizer under reasonable conditions (Theorems 17-19). More precisely, we prove the following local properties of Survey Descent given input surveys sufficiently close to the minimizer of h :

- Survey Descent iterations are always feasible.
- Survey Descent outputs are always unique, and h remains differentiable at these outputs.
- Surveys converge Q-linearly to $\arg\min_x h(x)$ when Survey Descent is applied repeatedly.
- Survey function-values converge R-linearly to $\min_x h(x)$ when Survey Descent is applied repeatedly.

Here, we follow the terminology of Ortega and Rheinboldt (2000, Chapter 9). A major assumption for our development and any practical implementation is the availability of a workable choice of the survey size, k , a question associated with the “active structure” of h at its minimizer. We discuss this choice throughout the exposition and present some empirical heuristics in Remark 20.

1.2 Linear Convergence and Nonsmooth Objectives

In Figure 1, we illustrate a simple experiment on a max-of-smooth function objective suggesting linear convergence. Figure 2 suggests that this behavior persists in higher dimensions. Moreover, Figures 3 and 4, which apply Survey Descent on objectives not expressible as the maximum of smooth functions, suggest that Survey Descent may retain this linear convergence behavior even on more general nonsmooth functions, hinting at its potential as a local nonsmooth, minimization technique.

Survey Descent on Simple Max-Function Example

$$h_{\max}(x, y) = |x - y^2| + x^2 + 2y^2$$




Figure 1: **Survey Descent Iteration on Max-Function.** We use Survey Descent iterations to minimize the nonsmooth objective $h_{\max}(x, y)$ that achieves minimum value 0 at unique minimizer $\bar{x} = (0, 0)$. We use step-control parameter $L = 10$ and initialize with survey points $s_1 = (0.9, 1)$ and $s_2 = (1.1, 1)$. Panel (a) visualizes h_{\max} . The Survey Descent iterates (blue and orange dots) are shown in the xy -plane along with the $x = y^2$ curve (red) on which h_{\max} is non-differentiable. Panel (b) shows the location of iterates in the xy -plane where darker colors correspond to later iterations. Panel (c) shows the step-size (the magnitude of the difference between two consecutive iterates) at each iteration. Panel (d) shows the function value of the iterates. Iterates and function values converge linearly to global minimizer and minimum, respectively. Both survey points, s_1 and s_2 , remain on the same smooth piece of the objective throughout the optimization. Sections 2-4 will theoretically derive these behaviors on objectives that are the maximum of smooth functions—of which h_{\max} is a simple example.

Survey Descent in Higher Dimensions

$$h_n(x) = \max_{i=1, \dots, k} (a_i x + x^T A_i x) \quad x \in \mathbf{R}^n \quad k = n/5$$

for random vectors a_i and positive semidefinite matrices A_i




Figure 2: Panels plot the best objective values over 25 iterations, on 5 instances, for dimension $n = 100, 250, 500$. The initialization heuristic always chose the true number of components, k , as the survey size without *a priori* information. The data $A_i = C_i^T C_i$ and a_i were generated from square matrices C_i and vectors with standard Gaussian entries, the set $\{a_i\}$ being adjusted to ensure that its convex hull contains zero, giving the optimal solution $x^* = 0$.

Many popular first-order methods exhibit empirical linear convergence on nonsmooth objectives. For example, our Figures 5 and 6 illustrate the linear convergence of BFGS when minimizing two simple nonsmooth objectives, and Lewis and Overton (2008, 2013) explicitly discuss this aspect of the algorithm. Similarly, the experiments of Teo et al. (2010) show that first-order multipoint methods also display linear convergence when applied to a variety of nonsmooth machine learning tasks.

However, these behaviors are thought-provoking because modern nonsmooth, convex optimization *theory* has typically only proven that canonical first-order methods converge *sublinearly* when smoothness assumptions are absent—even in the presence of desirable properties such as strongly convex objectives: For example,

Survey Descent (with BFGS init.) on *Non-Max-of-Smooth* Objective

$$h_{\text{ME}}(x, y, z, w) = \text{MaxEigenvalue} \left(\begin{bmatrix} x & y & z \\ y & -x & w \\ z & w & -1 \end{bmatrix} \right)$$




Figure 3: **Survey Descent of *Non-Max-of-Smooth* Objective.** We use Survey Descent iterations with step-control parameter $L = 10$ to minimize h_{ME} with minimum value 0 at minimizer $\bar{x} = (0, 0, 0, 0)$. h_{ME} is nonsmooth and *not* expressible as a maximum of smooth functions, as is clear by noting $h_{\text{ME}}(x, y, 0, 0) = \sqrt{x^2 + y^2}$. To create an initializing survey, we run 20 preliminary iterations of BFGS exactly as described in (Lewis and Overton, 2008) starting from $(1, 1, 1, 1)$. We then chose both the size of a initializing survey—3 points—as well as the survey points themselves through an empirical heuristic based on the dimensionality of the BFGS iterate gradients. The left and right panels plot the objective values and step-sizes (the magnitude of the difference between two consecutive iterates), respectively, of each survey point during the Survey Descent iterations. After some initial fluctuations, we see the objective value and iterates of all Survey Descent survey-iterates converge at stably linear rates to the optimum. Thus, Survey Descent displays desirable performance even on an objective not captured by the theory in Sections 2-4.

see Nesterov (2003, Chapter 3.2) or Beck (2017, Chapter 8) that analyzes the subgradient method as well as Juditsky and Nemirovski (2011) that analyze mirror descent. Even for popular multipoint methods designed for nonsmooth optimization, previous theoretical guarantees have remained generally sublinear (discussed later in Section 1.4). For comparison, in smooth settings, GD variants possess well-recognized *linear* convergence guarantees of on L -smooth and δ -strongly convex objectives (Beck, 2017, Theorem 10.29). Our derivation of local linear convergence for Survey Descent on nonsmooth objectives will directly connect to these smooth GD results.

1.3 Motivation of Survey Descent

Before the formal analysis, we first describe and motivate some characteristic features within Survey Descent’s design. First, note that the subproblems are simple quadratic programs with Euclidean ball constraints. Thus, they are efficiently solvable using second-order conic solvers. When all constraints moreover hold with equality, as occurs in our local convergence analysis, solving them further simplifies to routine linear algebra.

Next, the differentiability of input survey points is inspired by the common “differentiability of every iterate” behavior exhibited by a variety of popular first-order methods¹ when applied to nonsmooth objectives. To illustrate, Figures 5c and 6c show that the BFGS iterates minimizing the figures’ nonsmooth objectives are always differentiable. This behavior also manifests in Lewis and Overton (2008), for example, when minimizing the nonsmooth Rosenbrock and Chebyshev-Rosenbrock objectives. Recent works have begun

¹Bundle methods (discussed in Section 1.4) are a notable exception. Their iterates often land exactly on the nonsmooth points of their associated objectives as a consequence of minimizing their underlying cutting-plane models.

Survey Descent on Larger Eigenvalue Optimization Problems

$$h_{\text{ME2}}(x) = \text{MaxEigenvalue} \left(\begin{bmatrix} 0_{3 \times 3} & 0_{3 \times 7} \\ 0_{7 \times 3} & -I_{7 \times 7} \end{bmatrix} + \sum_1^{40} x_i A_i \right) \quad x \in \mathbf{R}^{40}$$

for random symmetric $A_i \in \mathbf{R}^{10 \times 10}$

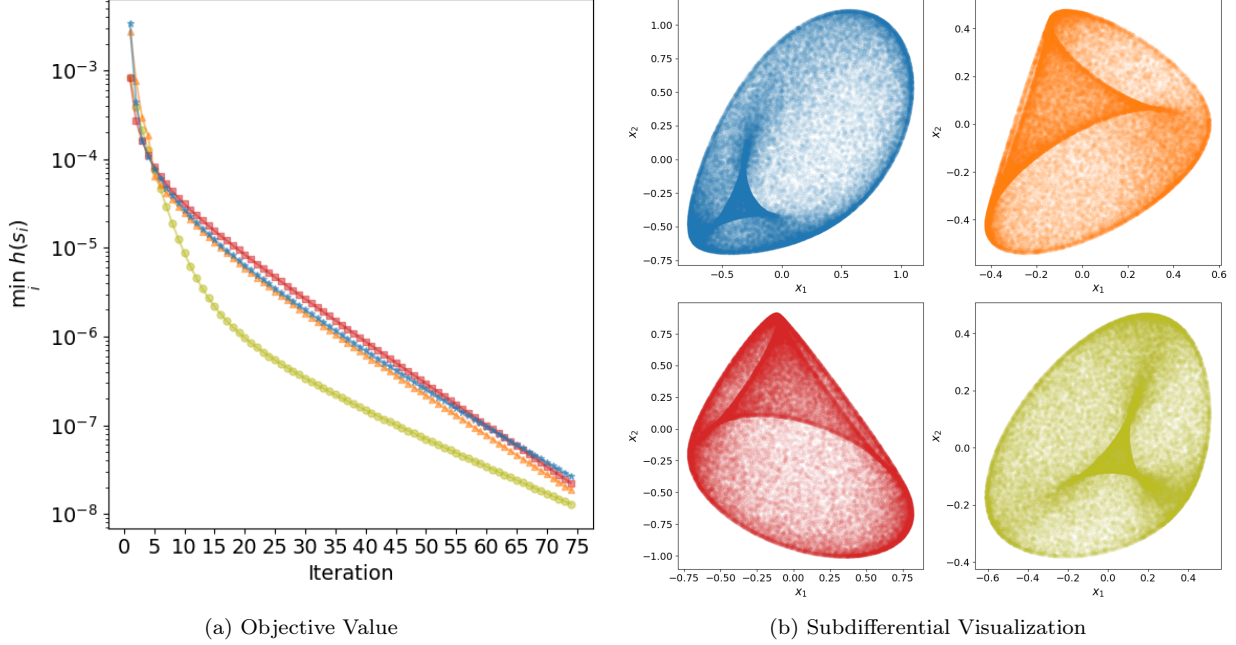


Figure 4: **Survey Descent on Larger Eigenvalue Optimization Problems.** We use Survey Descent to minimize four instances of h_{ME2} . The set $\{A_i\} = \{C_i^T + C_i\}$ was generated from square matrices C_i with standard Gaussian entries, shifted to ensure linear independence and an optimal solution at zero. The initialization heuristic always chose $\dim(\partial h_{\text{ME2}}(0)) + 1 = 6$ as the survey size without *a priori* information. Panel (a) plots best objective values over 75 iterations. Panel (b) shows two-dimensional projections of $\partial h_{\text{ME2}}(0)$ formed by sampling 10^5 gradients within a ball of radius 10^{-8} . The outputs do not cluster around finitely-many points, confirming heuristically that the instances are *not* max-of-smooth.

developing theoretical explanations for this occurrence. In particular, Bianchi et al. (2022) proved that, under standard assumptions, (stochastic) gradient descent iterates are differentiable with probability-one² even when objectives are nonconvex. Thus, we can easily imagine creating a survey of points at which h is differentiable by running BFGS or GD for a few initializing iterates and collecting some subset of those iterates. (For further discussion of implementation, see Section 5.)

Lastly, the size k of the survey itself relates to the dimension of the subdifferential of h at its global minimizer \bar{x} (assuming it exists) while the individual points $\{s_i\}_{i=1}^k$ “survey the landscape” of h near \bar{x} . The next two subsections elaborate on these intuitions.

1.3.1 Dimension of the Subdifferential

The dimensionality of subdifferentials captures key nonsmoothness properties. For example, smooth functions possess everywhere zero-dimensional subdifferentials since they contain only one point: the gradient. In comparison, Figure 5 shows an objective possessing a one-dimensional subdifferential at its minimizer

²We distinguish the notion of “almost all iterates of a first-order method are differentiable” from the also common notion of “almost all points in the domain are differentiable”—describing how all locally Lipschitz functions are non-differentiable on an at most measure-zero set in their domain. The latter result is often called Rademacher’s Theorem.

Simple Max-Function Example with BFGS

$$h_{\max}(x, y) = |x - y^2| + x^2 + 2y^2$$

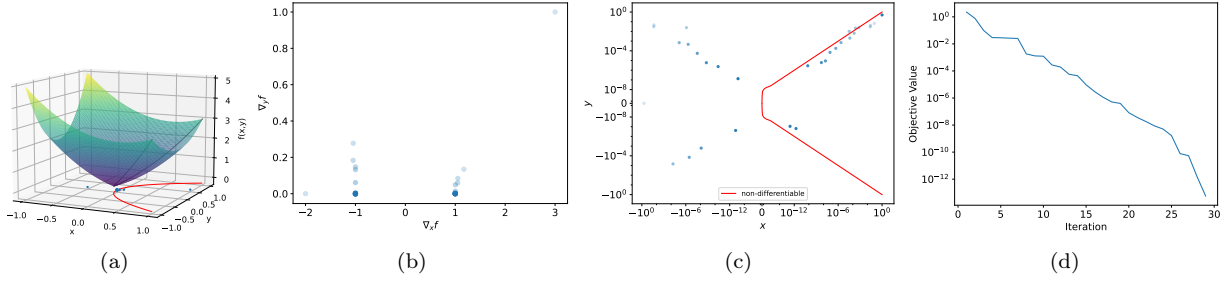


Figure 5: **Differentiability and Linear Convergence on Nonsmooth Objectives.** We use BFGS to minimize the nonsmooth objective $h_{\max}(x, y)$ that achieves minimum value 0 at unique minimizer $\bar{x} = (0, 0)$. BFGS is implemented exactly as described in Lewis and Overton (2008), initialized at $(1, 0.5)$, and ran for 30 iterations. Panel (a) visualizes h_{\max} . It shows the iterates (blue dots) in the xy -plane as well as the $x = y^2$ curve (red) on which h_{\max} is non-differentiable. Panels (b) and (c), respectively, show the gradients at and locations of each iterate with darker colors indicating later iterations. Panel (b) shows gradients form a convex hull that empirically resembles $\partial h_{\max}(0, 0)$'s one-dimensional structure. Panel (c) (in symmetric-log scale) shows that iterates lie within smooth subregions of the objective. Panel (d) records the linear convergence of the objective value.

characterized by

$$\partial h_{\max}(\bar{x}) = \text{conv} \left(\left\{ \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \end{bmatrix} \right\} \right), \quad (5)$$

where $\text{conv}(\cdot)$ denotes the convex hull. The objective is actually “partly smooth”: It has a “smooth edge” in the $[0, 1]^T$ -direction orthogonal to $\partial h_{\max}(\bar{x})$. Partial smoothness is formalized and explored in detail by Lewis (2002) and follow-up works.

Similarly, Figure 6 shows an objective where the subdifferential at its minimizer is a two-dimensional

Simple Elliptical-Norm Example with BFGS

$$h_{\text{ellipse}}(x, y) = \sqrt{x^2 + 2y^2}$$

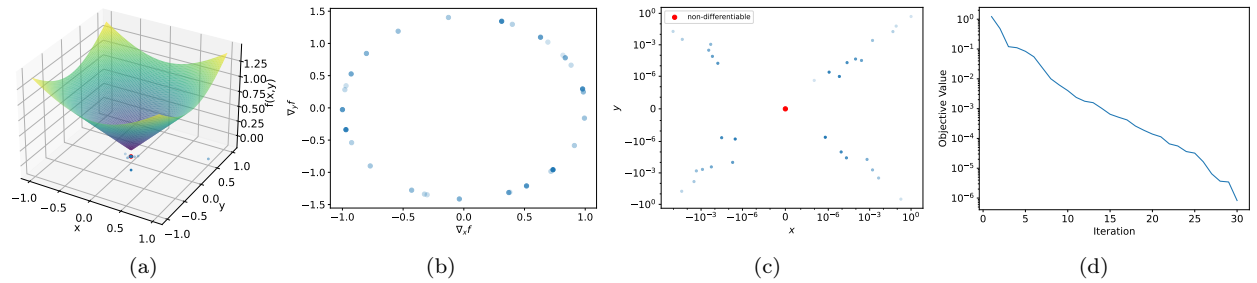


Figure 6: **Differentiability and Linear Convergence on Nonsmooth Objectives.** We use BFGS to minimize the nonsmooth objective $h_{\text{ellipse}}(x, y)$ that achieves minimum value 0 at unique minimizer $\bar{x} = (0, 0)$. BFGS is implemented exactly as described in Lewis and Overton (2008), initialized at $(1, 0.5)$, and ran for 30 iterations. Panel (a) visualizes h_{ellipse} . It shows the iterates (blue dots) in the xy -plane as well as the origin (red) at which h_{\max} is non-differentiable. Panels (b) and (c), respectively, show the gradients at and locations of each iterate with darker colors indicating later iterations. Panel (b) shows gradients forming a convex hull that empirically resembles $\partial h_{\text{ellipse}}(0, 0)$'s two-dimensional structure. Panel (c) (in symmetric-log scale) shows all iterates lie within differentiable regions of the domain. Panel (d) records the linear convergence of the objective value.

ellipse in the \mathbf{R}^2 domain

$$\partial h_{\text{ellipse}}(\bar{x}) = \left\{ (x, y) : x^2 + \frac{1}{2}y^2 \leq 1 \right\}. \quad (6)$$

The “full-dimensionality” of $\partial h_{\text{ellipse}}(\bar{x})$ captures the fact that h_{ellipse} is non-differentiable in all \mathbf{R}^2 directions.

Our focus on subdifferentials is also empirically motivated: When employing first-order methods to minimize some objective h , the gradients of iterates near the global minimizer \bar{x} typically have a convex hull whose dimension empirically coincides with the dimension of the subdifferential $\partial h(\bar{x})$. Figures 5a,b and 6a,b demonstrate this behavior. Moreover, this phenomenon occurs despite none of the iterates actually achieving the global minimum (Figures 5c and 6c). This is unsurprising since $\partial h(\bar{x})$ is the convex hull of the limit of gradients at nearby points (Clarke, 1990, Theorem 2.5.1). This motivates the creation of a “survey” of points near \bar{x} such that their gradients capture the dimensionality of $\partial h(\bar{x})$. Introducing the max-of-smooth model clarifies this intuition.

1.3.2 Max-of-Smooth Model

As is well-known, Fenchel conjugacy allows us to represent any continuous, convex function as the supremum of a family of affine functions (see, for example, Borwein and Lewis (2010, Chapter 3.3)). More generally, continuous, convex functions are instances of *lower- \mathcal{C}^1 functions*, which are functions representable as maximums of smooth and compactly parameterized functions.

Definition 2. (Lower- \mathcal{C}^1 Functions; Spingarn (1981), Daniilidis and Georgiev (2004, Cor. 3))

A locally Lipschitz function $h : \mathbf{R}^n \rightarrow \mathbf{R}$ is called *lower- \mathcal{C}^1* if, for every $\bar{x} \in \mathbf{R}^n$, there exists a compact parameterizing set \mathbf{T} , a neighborhood \mathbf{X} around \bar{x} , and functions $g : \mathbf{T} \times \mathbf{X} \rightarrow \mathbf{R}$ such that $g(t, x)$ and $\nabla_x g(t, x)$ are jointly continuous in t and x ; and

$$h(x) = \max_{t \in \mathbf{T}} g(t, x). \quad (7)$$

Lower- \mathcal{C}^1 functions are exactly the locally-Lipschitz and approximately convex functions, where “approximately convex” relaxes the canonical notion of convexity and is defined in Daniilidis and Georgiev (2004).

Lower- \mathcal{C}^1 functions subsume the entire class of continuous, convex objectives. Definition 2 implies that, for any global minimizer \bar{x} of the objective h , there exists a family of functions $\{g(t, \cdot)\}_{t \in \mathbf{T}}$ satisfying (7) at $x = \bar{x}$. This family also allows us to describe h ’s subdifferential at its minimizer:

$$\partial h(\bar{x}) = \text{conv} \left\{ \nabla_x g(t, \bar{x}) : t \in \mathbf{T}' \right\}, \quad (8)$$

where $\mathbf{T}' = \{t : g(t, \bar{x}) = h(\bar{x})\}$. For more details on (8), see Rockafellar and Wets (2009, Theorem 10.31). This characterization inspires us to consider a structurally revealing approximating model for convex, nonsmooth objectives where \mathbf{T} is *finite* and $\mathbf{T} = \mathbf{T}'$. In this setting, (7) reduces to a *max-of-smooth* function:

$$f(x) = \max_{i=1, \dots, k} f_i(x), \quad (9)$$

where k is a finite integer and each f_i is a differentiable function. Just like in classical analyses of GD (see, for example, Beck (2017, Chapter 10.6)), we will assume that each f_i is a \mathcal{C}^2 function and strongly convex.

When the objective is a max-of-smooth function (9), we could “survey” the landscape of f by somehow obtaining a point, s_i , from each region $\mathcal{R}_i \equiv \{x : f_i(x) > f_j(x) \forall j \neq i\}$ for all i . If the gradients $\{\nabla f_i(\bar{x})\}_{i=1}^k$ are moreover affinely independent—meaning \bar{x} is a “nondegenerate” minimizer—and the s_i ’s are sufficiently close to \bar{x} , the dimension of the convex hull of the survey gradients, $\text{conv} \left(\{\nabla f_i(s_i)\}_{i=1}^k \right)$, will coincide with the dimension of the objective’s subdifferential at its minimizer, $\partial f(\bar{x}) = \text{conv} \left(\{\nabla f_i(\bar{x})\}_{i=1}^k \right)$. Moreover, the dimension of $\partial f(\bar{x})$ is then $k-1$. In other words, it is exactly one less than the size of the survey.

Next, observe from Definition 1 that each Survey Descent subproblem simply performs one projected-GD step onto the feasible region defined by (4). The max-of-smooth objective intuitively motivates this region. In this case, when Survey Descent is equipped with an initializing survey consisting of one point in each \mathcal{R}_i as discussed above, $f(s_i) = f_i(s_i)$ and $\nabla f(s_i) = \nabla f_i(s_i)$ for all i . Then, since the components f_i are

L -smooth, the constraints for the i -th Survey Descent subproblem (4) restrict solutions to a region where the linear lower bound of f_i is at least the quadratic upper bounds of the remaining f_j , $j \neq i$. Therefore, when the i -th subproblem is feasible, its output s_i^+ would necessarily remain within \mathcal{R}_i . Section 2 formalizes these ideas in detail, but their geometry is simple: Figure 1 shows the behavior of Survey Descent on a simple max-of-smooth objective on \mathbf{R}^2 while Figure 7 presents an abstract illustration of the above intuitions.




Figure 7: Intuition of Survey Descent. Abstract depiction of Survey Descent on max-of-smooth objectives (9) with $k=2$ components and using two survey points (circular dots). Each survey point, s_i , is associated with one f_i -component of the objective for $i \in \{1, 2\}$. In the i -th Survey Descent subproblem, a gradient step (arrow) from the i -th survey point is projected onto the subproblem’s constraints (illustrated with dashed-line boundaries). The constraints (4) prevent subproblem outputs (stars) from crossing the “nonsmooth boundary” (red line). In other words, the i -th output must remain in the subregion of the domain where the f_i -component is “active”.

Despite our discussions of $\{f_i\}_{i=1}^k$, note the very important feature that *performing the Survey Descent iteration on max-of-smooth function objectives does not require access to individual f_i ’s*. It only involves function-value and gradient evaluations of the overall objective f . Moreover—although Survey Descent is intuitively motivated by the max-of-smooth function objective and possesses a structurally revealing local linear convergence theory on these objectives (Sections 2-4)—it is broadly implementable on *any* objective function. Indeed, Survey Descent still displays promising behavior even on objectives *not* expressible as the maximum of finitely-many smooth functions. For instance, Figures 3 and 4 show the visibly linear convergence of Survey Descent on max-matrix-eigenvalue objectives whose eigenvalues are nonsmooth functions of the input variables.

1.4 Related Works

A well-known example of Nemirovski and Yudin (1983) shows sublinear worst-case global complexity for any first-order method for arbitrary nonsmooth convex objectives. That example focuses on an initial sequence of iterates of length less than the dimension of the domain. In contrast, we prove that Survey Descent achieves linear convergence, but only locally and asymptotically, and on an interesting subclass of objectives.

Also worth noting for comparison is classical work on minimizing the maximum of finitely-many *given* smooth functions f_i , where one can access each component f_i rather than just their pointwise maximum, as in Survey Descent. Such problems are standard in classical optimization, easily solved by reduction to classical nonlinear programs. Kiwiel (1983) developed a more sophisticated two-phase method, first identifying a search direction by solving a quadratic program built from the function values f_i and gradients ∇f_i , and then performing a line search. More recently, Nesterov (2003, Scheme 2.3.13) computes a linearization of each f_i -component, $\{\ell_{\tilde{x}}^{f_i}(x)\}_{i=1}^k$, at some iterate \tilde{x} ; performs a proximal-point iteration on $\max_i \ell_{\tilde{x}}^{f_i}(x)$; and then determines the next iteration using a carefully-designed momentum step. Both Kiwiel (1983) and Nesterov (2003, Scheme 2.3.13) achieve global, linear convergence on strongly convex, max-of-smooth function objectives.

At large scale, motivated by modern ML applications, Carmon et al. (2021) describes a more sophisticated approach. Using function-value and gradient evaluations of the f_i -components, Carmon et al. (2021) implements a Ball Regularized Optimization Oracle (BROO) that returns the minimizer of the objective—subject to a proximal-regularizer—within a ball around a queried point. Using this BROO, Carmon et al. (2021) then designs an efficient minimization procedure by repeatedly updating an iterate using subroutines of line searches and BROO calls. Carmon et al. (2021) prove global, sublinear convergence rates for this procedure on max-of-smooth function objectives.

Within the derivative-free setting, researchers have also studied max-of-smooth objectives (Hare and Nutini, 2013) as well as generalizations into (potentially nonconvex) objectives that are the composition of

multiple component functions (Khan et al., 2018). By leveraging careful line-searches and (approximations to) active set information at individual iterates, these methods guarantee convergence using only function-value evaluations of individual objective components without any need for gradient evaluations. Convergence rates are challenging to derive in derivative-free settings and are generally sublinear where they do exist (for example, in Garmanjani et al. (2016)). For more detailed discussions on works investigating such derivative-free methods, see Khan et al. (2018).

Kiwiel (1983), Nesterov (2003), Carmon et al. (2021), and the derivative-free methods above differ from Survey Descent, since they crucially rely on access to the individual components of a max-function objective. In comparison, Survey Descent only assumes a first-order oracle that returns the function value and gradient of f *without* access to any component f_i ’s. Thus, Survey Descent is implementable on *any* objective, even those without max-of-smooth structure, suggesting a promising future research direction. Also notable is the distinction in computational style: Survey Descent subproblems consist of k parallel projected-GD steps, updating an entire collection of survey points rather than only one iterate as in the works above.

On the other hand, more sophisticated structural oracles can lead even to superlinear convergence on nonsmooth objectives, as remarked in several prior works. For eigenvalue optimization, Oustry (2000) and Noll and Apkarian (2005) develop quadratically convergent procedures by incorporating in particular a Hessian oracle within a nearby \mathcal{U} -subspace (Lemaréchal et al., 2000). On general nonsmooth objectives, using \mathcal{U} -Hessians and proximal operators on the objective, Mifflin and Sagastizábal (2005) present a \mathcal{VU} -algorithm whose serious steps converge superlinearly. In contrast, although Survey Descent conceptually relies on a max-of-smooth model, its implementation is general, requiring only function and gradient evaluations.

Among multipoint methods, one popular class for nonsmooth optimization are bundle methods (Lemaréchal, 1975; Wolfe, 1975), which possess a long, successful history (described in Mifflin and Sagastizábal (2012)). In its most transparent proximal form, they rely on a multipoint collection (a “bundle”) to build a piecewise-linear, cutting-plane model that is minimized in each step with a proximal operator around a current iterate (a “center”). If the proximal-step outputs a point that sufficiently decreases the function-value, the method takes a “serious step” updating the center with the outputted point; otherwise, the method takes a “null step” which does not update the center but adds the outputted point to the running point collection to improve the cutting-plane model.

The related “level-bundle” approach (Lemaréchal et al. (1995); Kiwiel (1995); de Oliveira and Sagastizábal (2014a)) also shares some similarities with Survey Descent. Each iteration projects the current center onto a sublevel set for the current cutting plane model, the update either being accepted if the objective value decreases satisfactorily—a serious step—or rejected, in which case the cutting plane model is updated. The projection ingredient is similar in both methods, but Survey Descent solves k *parallel* subproblems at each iteration whereas level-bundle methods perform projections sequentially, enforcing objective decrease for serious steps.

Bundle methods work well in practice (Borghetti et al., 2003; Emiel and Sagastizábal, 2010; Sagastizábal, 2012; van Ackooij et al., 2014) and enjoy a robust global convergence theory: see de Oliveira and Sagastizábal (2014b) for a comprehensive survey. In particular, relative to the number of *serious steps*, bundle methods converge linearly on convex objectives whose subdifferentials satisfy certain growth conditions away from the minimizer, as a consequence of a more general framework of Robinson (1999) and developed further in Atenas et al. (2021).³

However, relative to both null and serious steps, prior published literature (Kiwiel, 2000; Du and Ruszczyński, 2017; Diaz and Grimmer, 2021) have only derived sublinear convergence guarantees, even when the objective is strongly convex (although Davis and Jiang (2022) is a recent promising advance). Two recent works (Liang and Monteiro, 2021a,b) unify and present optimal iteration complexities for proximal bundle methods, all of which converge sublinearly, even on strongly convex objectives. In contrast, we show in this paper that the Survey Descent iteration—at least in the case of strongly convex, max-of-smooth function objectives—achieves a *local, linear* convergence rate.

Finally, we remark that Survey Descent surveys and bundle method bundles serve different functional

³Two independent related works (Atenas et al., 2021; Davis and Jiang, 2022) were announced recently. The first is a flexible, unified analysis of linearly convergent descent methods on weakly convex objectives, in particular covering the serious step sequence for bundle methods. A transparent analysis of Survey Descent in this framework is not immediately apparent; the original approach we present here is direct. The second announcement presents an interesting blackbox randomized first-order method that is nearly linear convergent with high probability.

roles. In bundle methods, null steps query points sequentially to improve the objective model around the center. In Survey Descent, there are neither auxiliary null steps nor a center. Instead, a fixed-size survey of points is cautiously updated *in parallel* trying to mimic the steady progress of GD.

2 Survey Descent on the Max-of-Smooth Objectives

We are particularly interested in convex objective functions f with a *nondegenerate minimizer* \bar{x} :

$$\bar{x} = \arg \min_{x \in \mathbf{R}^n} f(x) \text{ and } 0 \in \text{relint}(\partial f(\bar{x})), \quad (10)$$

where $\text{relint}(\cdot)$ denotes the relative interior. We further assume the following structure around \bar{x} .

Definition 3. A function $f : \mathbf{R}^n \rightarrow \mathbf{R}$ is a *strong \mathcal{C}^2 max function* if it is locally expressible near \bar{x} as

$$f(x) = \max_{i=1, \dots, k} f_i(x), \quad (11)$$

where k is some finite number, the *components* $\{f_i\}_{i=1}^k$ are \mathcal{C}^2 -functions satisfying $f_i(\bar{x}) = f(\bar{x})$ for all i , their gradients $\{\nabla f_i(\bar{x})\}_{i=1}^k$ are affinely independent, and their Hessians $\nabla^2 f_i(\bar{x})$ are positive definite. As a consequence, there exists constants $\delta, L > 0$ such that their Hessians satisfy

$$\delta I \preceq \nabla^2 f_i(x) \preceq LI \quad \forall i = 1, \dots, k, \quad (12)$$

for x near \bar{x} —where I is the $n \times n$ identity matrix.

Our analysis in this paper is entirely local, but we assume for the rest of this paper that f is a strong \mathcal{C}^2 max function with nondegenerate minimizer \bar{x} . For simplicity, we also assume that the properties in Definition 3 hold *globally* for all $x \in \mathbf{R}^n$. In this setting, we refer to k —which is necessarily unique—as the *degree* of f . Moreover, note that (12) implies all components $\{f_i\}_{i=1}^k$ are L -smooth, both f and its components are δ -strongly convex, and \bar{x} is f 's *unique* minimizer.

To study Survey Descent (Definition 1) on the objective f , we choose the above-presented L as the step-control parameter and define the following notions of *valid* and *minimizing* surveys.

Definition 4. ($\mathcal{S}, \bar{\mathcal{X}}$ -Surveys and Validity) For the objective f , define a *survey* as a matrix $\mathcal{S} \equiv [s_1, \dots, s_k] \in \mathbf{R}^{n \times k}$ consisting of k columns⁴ (the “survey points”), where k is the degree of f ; *valid surveys* as surveys satisfying $f_i(s_i) > f_j(s_i)$ for all $i \neq j$; and the *minimizing survey* as $\bar{\mathcal{X}} \equiv [\bar{x}, \dots, \bar{x}]$ whose survey points are all equal to \bar{x} .

On the space of surveys, we adopt the norm

$$\|\mathcal{S}\| \equiv \|[s_1, \dots, s_k]\|_{2, \infty} = \max_{i=1, \dots, k} \|s_i\|_2, \quad (13)$$

which allows us to quantify the distance between a survey \mathcal{S} from $\bar{\mathcal{X}}$ using $\|\mathcal{S} - \bar{\mathcal{X}}\|$, an intuitive measure of the distance to the furthest survey point. The precise choice of norm is immaterial to our computations and theory.

Given a valid \mathcal{S} sufficiently close to $\bar{\mathcal{X}}$, Survey Descent exhibits many desirable properties when minimizing f . As a preliminary, note that (10) and Definition 3 implies there exists unique *critical weights* $\{\bar{\lambda}_i\}_{i=1}^k > 0$ such that

$$\sum_{i=1}^k \bar{\lambda}_i = 1, \quad (14)$$

$$\sum_{i=1}^k \bar{\lambda}_i \nabla f_i(\bar{x}) = 0. \quad (15)$$

⁴We can equivalently consider \mathcal{S} as a set $\{s_i\}_{i=1}^k$. We will use the matrix and set interpretations interchangeably.

As we will see, the Lagrange multipliers of Survey Descent subproblems will lie “close” to $\{\bar{\lambda}_i\}_{i=1}^k > 0$, while the subproblem outputs $\{s_i^+\}_{i=1}^k$ will lie near \bar{x} . To characterize the order-of-magnitude of these distances, we adopt the “Big-Oh” notation. In particular, for a mapping $g : \mathbf{E} \rightarrow \mathbf{F}$ between two Euclidean spaces and letting $\|\cdot\|^p$ denote an arbitrary Euclidean norm raised to the p -th power, we use the notation $g(x) = O(\|x\|^p)$ to indicate the property that there exists a constant $K > 0$ such that $\|g(x)\| \leq K\|x\|^p$ holds for all small x . By itself, we let $O(\|x\|^p)$ denote an element of the class of all functions with this property.

With this terminology, we present our primary theoretical setting as well as our first result.

Setting A. (Local Analysis of Survey Descent) Consider an iteration of Survey Descent on a strong C^2 max function objective, f , with a survey \mathcal{S} that is valid and sufficiently close to $\bar{\mathcal{X}}$. For $i=1, \dots, k$, denote the output of the i -th Survey Descent subproblem by s_i^+ .

Theorem 5. (Local Feasibility, Uniqueness, Tightness, and Smoothness of Survey Descent) Assume Setting A. Then for all i , the Survey Descent subproblem $(P_i^{\mathcal{S}})$ is feasible and has a unique solution s_i^+ , which satisfies all the constraints with equality, with unique associated Lagrange multipliers λ_j^i for the j -th constraint within the i -th subproblem (for $j \neq i$). The solutions and multipliers depend smoothly on the input survey \mathcal{S} , and satisfy

$$s_i^+(\mathcal{S}) = \bar{x} + O(\|\mathcal{S} - \bar{\mathcal{X}}\|) \quad \text{and} \quad \lambda_j^i(\mathcal{S}) = \bar{\lambda}_j + O(\|\mathcal{S} - \bar{\mathcal{X}}\|) \quad \forall j \neq i. \quad (16)$$

Proof. Follows from a routine analysis of the subproblem’s first-order optimality conditions combined with the implicit function theorem. \square

We refer to Appendix B for a more direct and elementary proof that is computationally revealing. In particular, it shows we can compute $s_i^+(\mathcal{S})$ and $\lambda_j^i(\mathcal{S})$ entirely with linear algebra routines and a scalar square-root.

Theorem 5 plays a key role in deducing many theoretical results in this paper. Most immediately, under the assumptions of the theorem, the output survey $\mathcal{S}^+ = \{s_i^+\}_{i=1}^k$ of a Survey Descent iteration will also be valid. First, note that f is differentiable at every valid survey point and the associated gradients satisfy

$$\nabla f(s_i) = \nabla f_i(s_i) \quad \forall i = 1, \dots, k, \quad (17)$$

which leads to the following observation.

Observation 6. (Survey Descent Iterations on Max-of-Smooth Objectives) Assume that \mathcal{S} is a valid survey for the objective f . Then, for all i , the i -th subproblem $(P_i^{\mathcal{S}})$ of Survey Descent (Definition 1) is equivalent to

$$\begin{aligned} \min_x \quad & \left\| x - \left(s_i - \frac{1}{L} \nabla f_i(s_i) \right) \right\|_2^2 \\ \text{s.t.} \quad & \ell_{s_j}^{f_j}(x) + \frac{L}{2} \|x - s_j\|_2^2 \leq \ell_{s_i}^{f_i}(x) \quad \forall j \neq i. \end{aligned} \quad (18)$$

Using Theorem 5 and Observation 6, we deduce the aforementioned validity of \mathcal{S}^+ .

Theorem 7. (Preservation of Validity) Assuming Setting A, the output survey \mathcal{S}^+ of a Survey Descent iteration is valid.

Proof. Consider any fixed i . If $s_i^+ = s_i$, then the fact that $f_i(s_i^+) > f_j(s_i^+)$ for all $j \neq i$ immediately follows since Theorem 5 assumes \mathcal{S} is valid.

Now, consider the $s_i^+ \neq s_i$ case. By Theorem 5, s_i^+ satisfies the constraints of $(P_i^{\mathcal{S}})$ with equality. Using Observation 6, we express this as

$$\ell_{s_i}^{f_i}(s_i^+) = \ell_{s_j}^{f_j}(s_i^+) + \frac{L}{2} \|s_i^+ - s_j\|_2^2 \quad \forall j \neq i. \quad (19)$$

The δ -strong convexity of f_i implies

$$f_i(s_i^+) \geq \ell_{s_i}^{f_i}(s_i^+) + \frac{\delta}{2} \|s_i^+ - s_i\|_2^2 > \ell_{s_i}^{f_i}(s_i^+),$$

where the strict inequality follows from $s_i^+ \neq s_i$. Next, by the L -smoothness of $\{f_j\}_{j=1}^k$, the right-hand side of (19) is a s_j -centered quadratic upper bound of f_j evaluated at s_i^+ . Thus,

$$\ell_{s_j}^{f_j}(s_i^+) + \frac{L}{2} \|s_i^+ - s_j\|_2^2 \geq f_j(s_i^+) \quad \forall j \neq i.$$

Therefore,

$$f_i(s_i^+) > \ell_{s_i}^{f_i}(s_i^+) = \ell_{s_j}^{f_j}(s_i^+) + \frac{L}{2} \|s_i^+ - s_j\|_2^2 \geq f_j(s_i^+) \quad \forall j \neq i.$$

This completes the proof. \square

3 Connecting Survey Descent and Gradient Descent

We will build a local convergence theory by connecting Survey Descent to projected GD-steps on the smooth components $\{f_i\}_{i=1}^k$ of the objective f . To form this connection, we first identify the affine $\bar{\mathcal{U}}$ -subspace possessing the following equivalent characterizations:

$$\bar{\mathcal{U}} \equiv \bar{x} + \text{span} \{ \nabla f_i(\bar{x}) - \nabla f_j(\bar{x}) : 1 \leq i, j \leq k \}^\perp \quad (20)$$

$$= \bar{x} + \left\{ \sum_i^k \gamma_i \nabla f_i(\bar{x}) : \sum_i^k \gamma_i = 0 \right\}^\perp \quad (21)$$

$$= \bar{x} + \{ x \in \mathbf{R}^n : \nabla f_i(\bar{x})^T x = \nabla f_j(\bar{x})^T x \quad \forall 1 \leq i, j \leq k \}. \quad (22)$$

$\bar{\mathcal{U}}$ is so-named because the \bar{x} -centered, linear subspace $\{\bar{\mathcal{U}} - \bar{x}\}$ is commonly called the “ \mathcal{U} -subspace” (Lemaréchal et al., 2000) and captures the directions in which f is smooth around \bar{x} . Theorem 11 will show that the outputs of Survey Descent subproblems are approximately—up to an $O(\|S - \bar{\mathcal{X}}\|^2)$ term—a convex combination of $\bar{\mathcal{U}}$ -projected GD-steps on the smooth components $\{f_i\}_{i=1}^k$.

To prove Lemmas 8-10 leading to Theorem 11, we adopt the following notation for the gradients and Hessians of the components of f (Definition 3):

$$a_i \equiv \nabla f_i(\bar{x}) \text{ and } A_i \equiv \nabla^2 f_i(\bar{x}) \quad \forall i = 1, \dots, k. \quad (23)$$

Additionally, assuming $\bar{x}=0$ without loss of generality will simplify many of our proofs. In this case, the second-order Taylor expansion of f_i around zero is

$$f_i(x) = f(0) + a_i^T x + \frac{1}{2} x^T A_i x + r_i(x) \quad \forall i = 1, \dots, k, \quad (24)$$

where $r_i(x)$ is a residual function.

Lemma 8. *For all i , the residual function $r_i(x)$ in (24) satisfies $r_i(x) = O(\|x\|^2)$ and $\nabla r_i(x) = O(\|x\|)$.*

Proof. $r_i(x) = O(\|x\|^2)$ follows directly from Taylor’s Theorem. Next, observe that

$$r_i(x) = f_i(x) - \left(f(0) + a_i^T x + \frac{1}{2} x^T A_i x \right).$$

Thus, since f_i is \mathcal{C}^2 , the residual r_i must also be \mathcal{C}^2 , so its gradient and Hessian are well-defined. This implies $\nabla r_i(x)$ is then \mathcal{C}^1 . Taylor expanding ∇r_i around zero then gives the desired $\nabla r_i(x) = O(\|x\|)$ after observing that $\nabla r_i(0)$ is the zero-vector and $\nabla^2 r_i(0)$ is the zero-matrix (since r_i is the residual of the Taylor expansion of f_i around zero). \square

We now prove and present the main results of this section.

Lemma 9. *Assume Setting A and suppose $\bar{x}=0$. Then, s_i^+ satisfies*

$$\mathcal{P}_{\bar{\mathcal{U}}}[s_i^+] = \frac{1}{2 - \bar{\lambda}_i} \mathcal{P}_{\bar{\mathcal{U}}} \left[s_i - \frac{\bar{\lambda}_i}{L} \nabla f_i(s_i) \right] + \sum_{j \neq i} \frac{\bar{\lambda}_j}{2 - \bar{\lambda}_i} \mathcal{P}_{\bar{\mathcal{U}}} \left[s_j - \frac{1}{L} \nabla f_j(s_j) \right] + O(\|S\|^2) \quad \forall i = 1, \dots, k.$$

where $\mathcal{P}_{\bar{\mathcal{U}}}$ is the projection operator onto the $\bar{\mathcal{U}}$ -subspace (20).

Proof. Without loss of generality, assume $i = k$. Scale the objective of problem (18) by $\frac{L}{2}$ and consider the equivalent problem

$$\begin{aligned} \min_x \quad & \frac{L}{2} \left\| x - \left(s_k - \frac{1}{L} \nabla f_k(s_k) \right) \right\|_2^2 \\ \text{s.t.} \quad & \ell_{s_j}^{f_j}(x) + \frac{L}{2} \|x - s_j\|_2^2 \leq \ell_{s_k}^{f_k}(x) \quad \forall j \neq k. \end{aligned} \quad (25)$$

By Theorem 5 and Observation 6, there exists unique Lagrange multipliers $\{\lambda_j\}_{j=1}^{k-1}$ (we omit the k -superscript for simplicity) satisfying the following first-order optimality condition for (25):

$$0 = L \left(s_k^+ - \left(s_k - \frac{1}{L} \nabla f_k(s_k) \right) \right) + \sum_{j=1}^{k-1} \lambda_j (\nabla f_j(s_j) + L(s_k^+ - s_j) - \nabla f_k(s_k))$$

After rearranging this gives

$$0 = L \left(1 + \sum_{j=1}^{k-1} \lambda_j \right) s_k^+ - L \left(s_k + \sum_{j=1}^{k-1} \lambda_j s_j \right) + \nabla f_k(s_k) + \sum_{j=1}^{k-1} \lambda_j (\nabla f_j(s_j) - \nabla f_k(s_k)). \quad (26)$$

Define the following scalars:

$$\lambda_k \equiv 1 - \sum_{j=1}^{k-1} \lambda_j, \quad (27)$$

$$\mu_j \equiv \bar{\lambda}_j - \lambda_j \quad \forall j = 1, \dots, k. \quad (28)$$

For $j \neq k$, Theorem 5 gives $\mu_j = O(\|\mathcal{S}\|)$. For $j = k$, combining (27)-(28) with equation (14) implies

$$\mu_k = \bar{\lambda}_k - \left(1 - \sum_{j=1}^{k-1} \lambda_j \right) = \bar{\lambda}_k - \left(1 - \sum_{j=1}^{k-1} \bar{\lambda}_j \right) + O(\|\mathcal{S}\|) = O(\|\mathcal{S}\|).$$

Thus,

$$\mu_j = O(\|\mathcal{S}\|) \quad \forall j = 1, \dots, k. \quad (29)$$

Substituting (28) into (26), applying identities (14)-(15), collecting $O(\|\mathcal{S}\|^2)$ terms, and solving for s_k^+ gives

$$s_k^+ = \frac{1}{2 - \bar{\lambda}_k} \left[s_k - \frac{\bar{\lambda}_k}{L} \nabla f_k(s_k) \right] + \sum_{j=1}^{k-1} \frac{\bar{\lambda}_j}{2 - \bar{\lambda}_k} \left[s_j - \frac{1}{L} \nabla f_j(s_j) \right] + \frac{1}{L(2 - \bar{\lambda}_k)} \sum_{j=1}^k \mu_j \nabla f_j(s_j) + O(\|\mathcal{S}\|^2). \quad (30)$$

Adopting the notation (23) and applying the Taylor expansion (24), observe

$$\sum_{j=1}^k \mu_j \nabla f_j(s_j) = \sum_{j=1}^k \mu_j (a_j + A_j s_j + \nabla r_j(s_j)) = \sum_{j=1}^k \mu_j a_j + O(\|\mathcal{S}\|^2), \quad (31)$$

where last equality follows from (29) and Lemma 8. Substituting (31) into (30) then leads to

$$s_k^+ = \frac{1}{2 - \bar{\lambda}_k} \left[s_k - \frac{\bar{\lambda}_k}{L} \nabla f_k(s_k) \right] + \sum_{j=1}^{k-1} \frac{\bar{\lambda}_j}{2 - \bar{\lambda}_k} \left[s_j - \frac{1}{L} \nabla f_j(s_j) \right] + \underbrace{\frac{1}{L(2 - \bar{\lambda}_k)} \sum_{j=1}^k \mu_j a_j}_{(*)} + O(\|\mathcal{S}\|^2). \quad (32)$$

Next, again combine (27)-(28) with equation (14) to deduce

$$\sum_{j=1}^k \mu_j = \sum_{j=1}^k (\bar{\lambda}_j - \lambda_j) = 0.$$

Then, it follows from (21) that $\mathcal{P}_{\bar{U}}[(*)] = 0$. The desired result then follows from applying $\mathcal{P}_{\bar{U}}$ to (32). \square

Lemma 10. Assume Setting A, and let $\bar{x}=0$. Then, s_i^+ satisfies

$$\mathcal{P}_{\bar{\mathcal{U}}^\perp} [s_i^+] = \mathcal{O} \left(\|\mathcal{S}\|^2 \right) \quad \forall i = 1, \dots, k.$$

where $\mathcal{P}_{\bar{\mathcal{U}}^\perp}$ is the projection operator onto the orthogonal complement of the $\bar{\mathcal{U}}$ -subspace (20).

Proof. Consider any fixed i . By Theorem 5, the constraints of (P_i^S) are tight. Thus, by Observation 6, s_i^+ satisfies

$$f_j(s_j) + \nabla f_j(s_j)^T (s_i^+ - s_j) + \frac{L}{2} \|s_i^+ - s_j\|_2^2 = f_i(s_i) + \nabla f_i(s_i)^T (s_i^+ - s_i) \quad \forall j \neq i.$$

Using equation (24) in the above gives, for all $j \neq i$,

$$\begin{aligned} f(0) + a_j^T s_j + \frac{1}{2} s_j^T A_j s_j + r_j(s_j) + (a_j + A_j s_j + \nabla r_j(s_j))^T (s_i^+ - s_j) + \frac{L}{2} \|s_i^+ - s_j\|_2^2 \\ = f(0) + a_i^T s_i + \frac{1}{2} s_i^T A_i s_i + r_i(s_i) + (a_i + A_i s_i + \nabla r_i(s_i))^T (s_i^+ - s_i). \end{aligned}$$

After simplification and rearrangement, we obtain

$$(a_j - a_i)^T s_i^+ = -(A_j s_j - A_i s_i + \nabla r_j(s_j) - \nabla r_i(s_i))^T s_i^+ - \frac{L}{2} \|s_i^+ - s_j\|_2^2 - c_j(s_j) + c_i(s_i) \quad \forall j \neq i, \quad (33)$$

where we define the functions

$$c_j(x) \equiv -\frac{1}{2} x^T A_j x + r_j(x) - \nabla r_j(x)^T x \quad \forall j = 1, \dots, i.$$

By Lemma 8, $c_j(x)$ is $\mathcal{O}(\|x\|^2)$. Since we assume $\bar{x}=0$, by Theorem 5, $s_i^+ = \mathcal{O}(\|\mathcal{S}\|)$. We then deduce that all terms on the right-hand side of (33) are $\mathcal{O}(\|\mathcal{S}\|^2)$:

$$(a_j - a_i)^T s_i^+ = \mathcal{O} \left(\|\mathcal{S}\|^2 \right) \quad \forall j \neq i. \quad (34)$$

By (20),

$$\bar{\mathcal{U}}^\perp = \text{span} \{a_i - a_j : 1 \leq i, j \leq k\} = \text{span} \{a_j - a_i : j \neq i\}.$$

Combined with the affine independence of $\{a_j\}_{j=1}^k$ (Definition 3), the above implies $\{a_j - a_i\}_{j \neq i}$ forms a basis for $\bar{\mathcal{U}}^\perp$. Therefore, we can express the projection onto $\bar{\mathcal{U}}^\perp$ in terms $\{a_j - a_i\}_{j \neq i}$. The result now follows from equation (34) and standard linear algebra. \square

Theorem 11. (Gradient Descent Approximation of Survey Updates) Assume Setting A. Then the solution of the i -th Survey Descent subproblem satisfies, for all i ,

$$s_i^+ = \tilde{s}_i + \mathcal{O} \left(\|\mathcal{S} - \bar{\mathcal{X}}\|^2 \right) \quad (35)$$

where we define

$$\tilde{s}_i \equiv \frac{1}{2 - \bar{\lambda}_i} \mathcal{P}_{\bar{\mathcal{U}}} \left[s_i - \frac{\bar{\lambda}_i}{L} \nabla f_i(s_i) \right] + \sum_{j \neq i} \frac{\bar{\lambda}_j}{2 - \bar{\lambda}_i} \mathcal{P}_{\bar{\mathcal{U}}} \left[s_j - \frac{1}{L} \nabla f_j(s_j) \right], \quad (36)$$

and $\mathcal{P}_{\bar{\mathcal{U}}}$ is the projection operator onto the $\bar{\mathcal{U}}$ -subspace (20).

Proof. Without loss of generality, assume that $\bar{x}=0$. Then, the desired result follows from decomposing $s_k^+ = \mathcal{P}_{\bar{\mathcal{U}}} [s_k^+] + \mathcal{P}_{\bar{\mathcal{U}}^\perp} [s_k^+]$ and applying Lemmas 9 and 10. We can deduce the general $\bar{x} \neq 0$ case with the change of variables $x \leftarrow (x - \bar{x})$. \square

4 Local Linear Convergence of Surveys

Observe that the constituent terms of \tilde{s}_i in (36) are $\bar{\mathcal{U}}$ -projected gradient steps on the L -smooth and δ -strongly convex component functions $\{f_i\}_{i=1}^k$ (Definition 3). Thus, we can apply classical theory on projected-GD described in standard texts such as Beck (2017, Theorem 10.29). The idea of tracking iterates relative to their projections onto some nearby “active” manifold is familiar in nonsmooth optimization. Recent examples include Burke et al. (2020), Duchi and Ruan (2021), and Davis et al. (2021).

Theorem 12. (Projected-GD Behavior) *Assume that \mathcal{S} is a valid survey for the strong \mathcal{C}^2 max function objective f . Then, the following holds for all i and any constant $\tilde{L} \geq L$:*

$$\begin{aligned} \left\| \mathcal{P}_{\bar{\mathcal{U}}} \left[s_i - \frac{1}{\tilde{L}} \nabla f_i(s_i) \right] - \bar{x} \right\|_2^2 &\leq \left(1 - \frac{\delta}{\tilde{L}} \right) \|s_i - \bar{x}\|_2^2, \\ f_i \left(\mathcal{P}_{\bar{\mathcal{U}}} \left[s_i - \frac{1}{\tilde{L}} \nabla f_i(s_i) \right] \right) - f_i(\bar{x}) &\leq \frac{\tilde{L}}{2} \left(1 - \frac{\delta}{\tilde{L}} \right) \|s_i - \bar{x}\|_2^2. \end{aligned}$$

Proof. Consider any fixed i . It is easy to check that $\bar{x} = \arg \min_{x \in \bar{\mathcal{U}}} f_i(x)$ using first-order optimality conditions combined with (14)-(15) and (22). The desired result then follows from classical projected-GD convergence results (given in standard texts such as Beck (2017, Theorem 10.29)) after observing that $\bar{\mathcal{U}}$ is convex and f_i is L -smooth and δ -strongly convex. \square

To prove local linear convergence, we show below that the survey points contract towards \bar{x} . The corresponding contraction ratios will depend on the minimum critical weight defined as follows:

$$\bar{\lambda}_{\min} \equiv \min_{i=1, \dots, k} \bar{\lambda}_i > 0. \quad (37)$$

Lemma 13. *Assume Setting A, and suppose $\bar{x}=0$. Then, the following must hold:*

$$\|\tilde{s}_i\|_2^2 \leq \left(1 - \frac{\bar{\lambda}_{\min} \delta}{L} \right) \max_{j=1, \dots, k} \|s_j\|_2^2 \quad \forall i = 1, \dots, k,$$

where \tilde{s}_i and $\bar{\lambda}_{\min}$ are as defined in equations (36) and (37), respectively.

Proof. Consider any fixed i . Then,

$$\begin{aligned} \|\tilde{s}_i\|_2^2 &= \left\| \frac{1}{2 - \bar{\lambda}_i} \mathcal{P}_{\bar{\mathcal{U}}} \left[s_i - \frac{\bar{\lambda}_i}{L} \nabla f_i(s_i) \right] + \sum_{j \neq i} \frac{\bar{\lambda}_j}{2 - \bar{\lambda}_i} \mathcal{P}_{\bar{\mathcal{U}}} \left[s_j - \frac{1}{L} \nabla f_j(s_j) \right] \right\|_2^2 \quad (\text{Substitute (36)}) \\ &\leq \frac{1}{2 - \bar{\lambda}_i} \left\| \mathcal{P}_{\bar{\mathcal{U}}} \left[s_i - \frac{\bar{\lambda}_i}{L} \nabla f_i(s_i) \right] \right\|_2^2 + \sum_{j \neq i} \frac{\bar{\lambda}_j}{2 - \bar{\lambda}_i} \left\| \mathcal{P}_{\bar{\mathcal{U}}} \left[s_j - \frac{1}{L} \nabla f_j(s_j) \right] \right\|_2^2 \quad (\text{Convexity}) \\ &\leq \frac{1}{2 - \bar{\lambda}_i} \left(1 - \frac{\bar{\lambda}_i \delta}{L} \right) \|s_i\|_2^2 + \left(1 - \frac{\delta}{L} \right) \sum_{j \neq i} \frac{\bar{\lambda}_j}{2 - \bar{\lambda}_i} \|s_j\|_2^2 \quad (\text{Theorem 12 where } \tilde{L} = \frac{L}{\bar{\lambda}_i}) \\ &\leq \left(1 - \frac{\bar{\lambda}_i \delta}{L} \right) \left[\frac{1}{2 - \bar{\lambda}_i} \|s_i\|_2^2 + \sum_{j \neq i} \frac{\bar{\lambda}_j}{2 - \bar{\lambda}_i} \|s_j\|_2^2 \right] \\ &\leq \left(1 - \frac{\bar{\lambda}_{\min} \delta}{L} \right) \max_{j=1, \dots, k} \|s_j\|_2^2. \end{aligned}$$

This completes the proof. \square

Theorem 14. (Survey-Norm Contraction) *Assume Setting A. Then the following must hold:*

$$\max_{j=1, \dots, k} \|s_j^+ - \bar{x}\|_2^2 \leq \left(1 - \frac{\bar{\lambda}_{\min} \delta}{L} \right) \max_{j=1, \dots, k} \|s_j - \bar{x}\|_2^2 + \mathcal{O}(\|\mathcal{S} - \bar{\mathcal{X}}\|^3),$$

where $\bar{\lambda}_{\min}$ is defined as in (37).

Proof. Without loss of generality, assume $\bar{x} = 0$. Consider any fixed i . By Lemma 13, $\tilde{s}_i = O(\|\mathcal{S}\|)$. Then,

$$\begin{aligned}\|s_i^+\|_2^2 &= \left\| \tilde{s}_i + O(\|\mathcal{S}\|^2) \right\|_2^2 \quad (\text{Theorem 11}) \\ &= \|\tilde{s}_i\|_2^2 + 2 \left\langle \tilde{s}_i, O(\|\mathcal{S}\|^2) \right\rangle + O(\|\mathcal{S}\|^4) \\ &\leq \left(1 - \frac{\bar{\lambda}_{\min}\delta}{L}\right) \max_{j=1,\dots,k} \|s_j\|_2^2 + O(\|\mathcal{S}\|^3) \quad (\text{Lemma 13}).\end{aligned}$$

The right-hand side is independent of i , so taking the max over i on the left-hand side leads to the desired result. We can deduce the general $\bar{x} \neq 0$ case with the change of variables $x \leftarrow (x - \bar{x})$. \square

Theorem 14 implies that for input surveys \mathcal{S} , sufficiently close to $\bar{\mathcal{X}}$, Survey Descent outputs a survey \mathcal{S}^+ that is strictly *closer* to $\bar{\mathcal{X}}$. As a result, the feasibility and validity guarantees of Theorems 5 and 7 will continue to hold in repeated applications of Survey Descent iterations allowing us to repeat the procedure indefinitely. Letting \mathbf{N} denote the non-negative natural numbers, we formalize these ideas below.

Procedure B. (Repeated Iterations of Survey Descent) For the strong \mathcal{C}^2 max function objective f , assume \mathcal{S}^0 is a valid initial survey sufficiently close to $\bar{\mathcal{X}}$, and initialize $t = 0$. Iterate the following steps:

1. Solve the Survey Descent iteration $\left\{ \left(P_i^{\mathcal{S}^t} \right) \right\}_{i=1}^k$ and denote the output $\mathcal{S}^{t+1} \equiv (\mathcal{S}^t)^+$.
2. Increment $t \leftarrow t + 1$.

Corollary 15. (Well-Definedness of Survey Descent Repetitions) Procedure B is well-defined: for each $t = 0, 1, 2, \dots$, the Survey Descent iteration $\left\{ \left(P_i^{\mathcal{S}^t} \right) \right\}_{i=1}^k$ is feasible and produces a valid output survey.

Proof. For \mathcal{S}^0 sufficiently close to $\bar{\mathcal{X}}$, Theorem 14 inductively implies that $\|\mathcal{S}^{t+1} - \bar{\mathcal{X}}\| < \|\mathcal{S}^t - \bar{\mathcal{X}}\|$ for all $t \in \mathbf{N}$. In other words, output surveys will always move closer to $\bar{\mathcal{X}}$, and Theorem 5 and 7 continue to hold for Survey Descent iterations $\left\{ \left(P_i^{\mathcal{S}^t} \right) \right\}_{i=1}^k$ for all $t \in \mathbf{N}$. Thus, Theorem 5 implies the feasibility all Survey Descent iterations while Theorem 7 implies the validity of output surveys. \square

Observation 16. The proof of Corollary 15 more generally shows that, for all $t \in \mathbf{N}$, all surveys \mathcal{S}^t will remain sufficiently close to $\bar{\mathcal{X}}$. Thus, the conclusions from Theorems 5-14 apply to every iteration of Procedure B.

Having established that the repeated application of Survey Descent is locally well-defined, we now deduce the local Q-linear convergence induced by this procedure.

Theorem 17. (Q-Linear Convergence of Survey Points) Procedure B satisfies the following property:

$$\limsup_{t \rightarrow \infty} \frac{\|\mathcal{S}^{t+1} - \bar{\mathcal{X}}\|^2}{\|\mathcal{S}^t - \bar{\mathcal{X}}\|^2} \leq \left(1 - \frac{\bar{\lambda}_{\min}\delta}{L}\right), \quad (38)$$

where $\bar{\lambda}_{\min}$ is defined as in (37).

Proof. By Theorem 14 and Observation 16, there exists $K > 0$ such that

$$\|\mathcal{S}^{t+1} - \bar{\mathcal{X}}\|^2 \leq \left(1 - \frac{\bar{\lambda}_{\min}\delta}{L}\right) \|\mathcal{S}^t - \bar{\mathcal{X}}\|^2 + K \|\mathcal{S}^t - \bar{\mathcal{X}}\|^3 \quad \forall t \in \mathbf{N}.$$

Corollary 15 implies that \mathcal{S}^t is a valid survey. By Definition 4, $\bar{\mathcal{X}}$ is not a valid survey because $f_i(\bar{x}) \neq f_j(\bar{x})$ for all i, j . Thus, $\mathcal{S}^t \neq \bar{\mathcal{X}}$, so $\|\mathcal{S}^t - \bar{\mathcal{X}}\|^2$ is nonzero and we can divide both sides by it. We deduce

$$\frac{\|\mathcal{S}^{t+1} - \bar{\mathcal{X}}\|^2}{\|\mathcal{S}^t - \bar{\mathcal{X}}\|^2} \leq \left(1 - \frac{\bar{\lambda}_{\min}\delta}{L}\right) + K \|\mathcal{S}^t - \bar{\mathcal{X}}\| \quad \forall t \in \mathbf{N}.$$

For \mathcal{S}^0 sufficiently close to $\bar{\mathcal{X}}$, this inequality guarantees that $\|\mathcal{S}^t - \bar{\mathcal{X}}\|$ monotonically decreases with t and, in particular, $\|\mathcal{S}^t - \bar{\mathcal{X}}\| < \frac{\bar{\lambda}_{\min}\delta}{LK}$ for all $t \in \mathbf{N}$. Thus, $\frac{\|\mathcal{S}^{t+1} - \bar{\mathcal{X}}\|}{\|\mathcal{S}^t - \bar{\mathcal{X}}\|} < 1$ implying $\lim_{t \rightarrow \infty} \|\mathcal{S}^t - \bar{\mathcal{X}}\| = 0$. Taking the lim-sup with $t \rightarrow \infty$ on both sides of the above inequality then gives the desired result. \square

Combining Theorem 17 with the L -smoothness assumptions (Definition 3), we can immediately establish a loose R -linear convergence result on the objective values.

Corollary 18. (*R-Linear Convergence of Objective Values, Weak Version*) Consider Procedure B. Then, for any $K > \max_j \|\nabla f_j(\bar{x})\|_2$,

$$f(s_i^t) - f(\bar{x}) \leq K \|\mathcal{S}^t - \bar{\mathcal{X}}\| \quad \forall i = 1, \dots, k, \text{ and } t \in \mathbf{N}.$$

Consequently, survey objective values $\{f(s_i^t)\}_{t \in \mathbf{N}}$ converge R -linearly to $f(\bar{x})$ for all i .

Proof. Follows from the L -smoothness of the f_i -components and a routine application of the Cauchy-Schwarz inequality. R -linear convergence follows from the Q -linear convergence of $\{\|\mathcal{S}^t - \bar{\mathcal{X}}\|\}_{t \in \mathbf{N}}$ (Theorem 17). \square

Compared to classical GD convergence guarantees on smooth objectives (Beck, 2017, Theorem 10.29), Corollary 18 is rather weak: Firstly, the upper-bounding sequence is $\{\|\mathcal{S}^t - \bar{\mathcal{X}}\|\}_{t \in \mathbf{N}}$ rather than the tighter $\{\|\mathcal{S}^t - \bar{\mathcal{X}}\|^2\}_{t \in \mathbf{N}}$ seen in the smooth GD case; secondly, the constant K depends on $\max_j \|\nabla f_j(\bar{x})\|_2$, which does not appear in the smooth GD case. A stronger function-value convergence guarantee for Survey Descent—more structurally similar to the aforementioned GD results—is indeed possible.

Theorem 19. (*R-Linear Convergence of Objective Values, Strong Version*) Consider Procedure B. Fix $\bar{K} > \frac{5L}{\bar{\lambda}_{\min}} \left(1 - \frac{\bar{\lambda}_{\min}\delta}{L}\right)$. Then,

$$f(s_i^{t+1}) - f(\bar{x}) \leq \bar{K} \|\mathcal{S}^t - \bar{\mathcal{X}}\|^2 \quad \forall i = 1, \dots, k, \text{ and } t \in \mathbf{N}.$$

Consequently, survey objective values $\{f(s_i^t)\}_{t \in \mathbf{N}}$ converge R -linearly to $f(\bar{x})$ for all i .

Proof. Follows from combining Theorem 11, Theorem 12, and the L -smoothness and δ -strong convexity assumptions on the components $\{f_i\}_{i=1}^k$. We defer a detailed description to Appendix A. \square

Note that the rates of linear convergence presented in Theorems 17 and 19 are conservative. Experiments suggest that these rate bounds are far from tight.

5 Implementing Survey Descent

Stepping back from our theoretical analysis on strong \mathcal{C}^2 max functions f , we conclude with two remarks on the potential implementation of Survey Descent on general nonsmooth objectives h (Definition 1).

Remark 20. (*Informal Heuristic for Survey Initialization*) We sketch an informal idea. First, run some initializing iterations of a standard method such as BFGS or a subgradient or proximal bundle method, and collect the gradients at these iterates. (As discussed, these iterates are typically points of differentiability.) In practice, the iterates explore the nonsmooth landscape effectively. In particular, on max-functions, they “bounce” between the active functions near the minimizer: for BFGS, see our Figure 5 and Lewis and Overton (2008, Fig. 5.5,6), and for the subgradient method, see the figures and discussion on “oscillations” in Bolte et al. (2020). As discussed, the convex hulls of nearby iterate gradients thus approximate the subdifferential at optimality, so we can estimate its dimension d via singular value decomposition. An initialization heuristic could then select a survey of $k = d + 1$ iterates with robustly affinely independent gradients. Implementing such heuristics is intricate: we defer further discussion to future work.

Remark 21. (*Possible Acceleration of Survey Descent*) First, note that we can solve the k constituent subproblems of Survey Descent (Definition 1) in parallel. Second, when all subproblems are feasible, we found in exploratory experiments that the following adjustment expedites Survey Descent’s empirical convergence:

For all i , only update $s_i \leftarrow s_i^+$ when $h(s_i^+) < h(s_i)$; otherwise, keep s_i the same in the subsequent survey.

This also tends to make Survey Descent compatible with a wider variety of initializing heuristics. For simplicity, we omit this additional enhancement from Definition 1 and do not apply it in this paper’s experiments.

Acknowledgements

Thanks are due to the Associate Editor and three anonymous referees for many helpful suggestions that significantly improved the initial manuscript.

References

- F. Atenas, C. Sagastizábal, P. J. Silva, and M. Solodov. A unified analysis of descent sequences in weakly convex optimization, including convergence rates for bundle methods. Technical report, Optimization Online, 2021. URL http://www.optimization-online.org/DB_FILE/2021/12/8727.pdf.
- A. Beck. *First-Order Methods in Optimization*. SIAM, 2017.
- P. Bianchi, W. Hachem, and S. Schechtman. Convergence of constant step stochastic gradient descent for non-smooth non-convex functions. *Set-Valued and Variational Analysis*, pages 1–31, 2022.
- J. Bolte, E. Pauwels, and R. Rios-Zertuche. Long term dynamics of the subgradient method for lipschitz path differentiable functions. *arXiv preprint arXiv:2006.00098*, 2020.
- A. Borghetti, A. Frangioni, F. Lacalandra, and C. A. Nucci. Lagrangian heuristics based on disaggregated bundle methods for hydrothermal unit commitment. *IEEE Transactions on Power Systems*, 18(1):313–323, 2003.
- J. Borwein and A. S. Lewis. *Convex Analysis and Nonlinear Optimization: Theory and Examples*. Springer Science & Business Media, 2010.
- J. V. Burke, F. E. Curtis, A. S. Lewis, M. L. Overton, and L. E. Simões. Gradient sampling methods for nonsmooth optimization. *Numerical Nonsmooth Optimization*, pages 201–225, 2020.
- Y. Carmon, A. Jambulapati, Y. Jin, and A. Sidford. Thinking inside the ball: Near-optimal minimization of the maximal loss. In *Conference on Learning Theory 2021*, Proceedings of Machine Learning Research, pages 866–882. PMLR, 2021.
- F. H. Clarke. *Optimization and Nonsmooth Analysis*. SIAM, 1990.
- A. Daniilidis and P. Georgiev. Approximate convexity and submonotonicity. *Journal of Mathematical Analysis and Applications*, 291(1):292–301, 2004.
- D. Davis and L. Jiang. A nearly linearly convergent first-order method for nonsmooth functions with quadratic growth. *arXiv preprint arXiv:2205.00064*, 2022.
- D. Davis, D. Drusvyatskiy, and L. Jiang. Subgradient methods near active manifolds: saddle point avoidance, local convergence, and asymptotic normality. *arXiv preprint arXiv:2108.11832*, 2021.
- W. de Oliveira and C. Sagastizábal. Level bundle methods for oracles with on-demand accuracy. *Optimization Methods and Software*, 29(6):1180–1209, 2014a.
- W. de Oliveira and C. Sagastizábal. Bundle methods in the XXIst century: A bird’s-eye view. *Pesquisa Operacional*, 34(3):647–670, 2014b.
- M. Diaz and B. Grimmer. Optimal convergence rates for the proximal bundle method. *arXiv preprint arXiv:2105.07874*, 2021.

- Y. Du and A. Ruszczyński. Rate of convergence of the bundle method. *Journal of Optimization Theory and Applications*, 173(3):908–922, 2017.
- J. C. Duchi and F. Ruan. Asymptotic optimality in stochastic optimization. *The Annals of Statistics*, 49(1):21–48, 2021.
- G. Emiel and C. Sagastizábal. Incremental-like bundle methods with application to energy planning. *Computational Optimization and Applications*, 46(2):305–332, 2010.
- R. Garmanjani, D. Júdice, and L. N. Vicente. Trust-region methods without using derivatives: worst case complexity and the nonsmooth case. *SIAM Journal on Optimization*, 26(4):1987–2011, 2016.
- W. Hare and J. Nutini. A derivative-free approximate gradient sampling algorithm for finite minimax problems. *Computational Optimization and Applications*, 56(1):1–38, 2013.
- N. J. Higham. *Accuracy and Stability of Numerical Algorithms*. SIAM, 2002.
- A. Juditsky and A. Nemirovski. First order methods for nonsmooth convex large-scale optimization, i: general purpose methods. *Optimization for Machine Learning*, 30(9):121–148, 2011.
- K. A. Khan, J. Larson, and S. M. Wild. Manifold sampling for optimization of nonconvex functions that are piecewise linear compositions of smooth components. *SIAM Journal on Optimization*, 28(4):3001–3024, 2018.
- K. Kiwiel. A phase I—phase II method for inequality constrained minimax problems. *Control and Cybernetics*, 12(1-2):55–75, 1983.
- K. C. Kiwiel. Proximal level bundle methods for convex nondifferentiable optimization, saddle-point problems and variational inequalities. *Mathematical Programming*, 69(1):89–109, 1995.
- K. C. Kiwiel. Efficiency of proximal bundle methods. *Journal of Optimization Theory and Applications*, 104(3):589–603, 2000.
- C. Lemarechal. An extension of Davidon methods to non differentiable problems. In *Nondifferentiable Optimization*, pages 95–109. Springer, 1975.
- C. Lemaréchal, A. Nemirovskii, and Y. Nesterov. New variants of bundle methods. *Mathematical Programming*, 69(1):111–147, 1995.
- 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.
- A. S. Lewis. Active sets, nonsmoothness, and sensitivity. *SIAM Journal on Optimization*, 13(3):702–725, 2002.
- A. S. Lewis and M. L. Overton. Nonsmooth optimization via BFGS. Technical report, Optimization Online, 2008. URL http://www.optimization-online.org/DB_FILE/2008/12/2172.pdf.
- A. S. Lewis and M. L. Overton. Nonsmooth optimization via quasi-Newton methods. *Mathematical Programming*, 141(1):135–163, 2013.
- J. Liang and R. D. Monteiro. A proximal bundle variant with optimal iteration-complexity for a large range of prox stepsizes. *SIAM Journal on Optimization*, 31(4):2955–2986, 2021a.
- J. Liang and R. D. Monteiro. A unified analysis of a class of proximal bundle methods for solving hybrid convex composite optimization problems. *arXiv preprint arXiv:2110.01084*, 2021b.
- R. Mifflin and C. Sagastizábal. A \mathcal{VU} -algorithm for convex minimization. *Mathematical Programming*, 104(2):583–608, 2005.
- R. Mifflin and C. Sagastizábal. A science fiction story in nonsmooth optimization originating at IIASA. *Documenta Mathematica*, 2012.

- A. S. Nemirovski and D. B. Yudin. Problem complexity and method efficiency in optimization. 1983.
- Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*, volume 87. Springer Science & Business Media, 2003.
- J. Nocedal and S. Wright. *Numerical Optimization*. Springer Science & Business Media, 2006.
- D. Noll and P. Apkarian. Spectral bundle methods for non-convex maximum eigenvalue functions: second-order methods. *Mathematical Programming*, 104(2):729–747, 2005.
- J. M. Ortega and W. C. Rheinboldt. *Iterative Solution of Nonlinear Equations in Several Variables*. SIAM, 2000.
- F. Oustry. A second-order bundle method to minimize the maximum eigenvalue function. *Mathematical Programming*, 89(1):1–33, 2000.
- S. M. Robinson. Linear convergence of epsilon-subgradient descent methods for a class of convex functions. *Mathematical Programming*, 86(1):41–50, 1999.
- R. T. Rockafellar and R. J.-B. Wets. *Variational Analysis*, volume 317. Springer Science & Business Media, 2009.
- C. Sagastizábal. Divide to conquer: decomposition methods for energy optimization. *Mathematical Programming*, 134(1):187–222, 2012.
- J. E. Spingarn. Submonotone subdifferentials of Lipschitz functions. *Transactions of the American Mathematical Society*, 264(1):77–89, 1981.
- C. H. Teo, S. Vishwanathan, A. Smola, and Q. V. Le. Bundle methods for regularized risk minimization. *Journal of Machine Learning Research*, 11(1), 2010.
- W. van Ackooij, R. Henrion, A. Möller, and R. Zorgati. Joint chance constrained programming for hydro reservoir management. *Optimization and Engineering*, 15(2):509–531, 2014.
- P. Wolfe. A method of conjugate subgradients for minimizing nondifferentiable functions. In *Nondifferentiable Optimization*, pages 145–173. Springer, 1975.

Appendix

A Stronger R-Linear Convergence of Function Values (Theorem 19)

In Section 4, we analyzed the distance of the Survey Descent updates, $\mathcal{S}^+ = \{s_i^+\}_{i=1}^k$, from the objective minimizer by examining the set of reference points $\{\hat{s}_i\}_{i=1}^k$ —which coincides with \mathcal{S}^+ up to $O(\|\mathcal{S}\|^2)$ terms and is defined in (36). This appendix will prove Theorem 19—showing R-linear function value convergence—through an analogous strategy. To simplify the derivations, we adopt the following notation:

$$\hat{s}_i = \mathcal{P}_{\bar{\mathcal{U}}} \left[s_i - \frac{1}{L} \nabla f_i(s_i) \right] \text{ and } \hat{s}_i = \mathcal{P}_{\bar{\mathcal{U}}} \left[s_i - \frac{\bar{\lambda}_i}{L} \nabla f_i(s_i) \right] \quad \forall i = 1, \dots, k. \quad (39)$$

In this notation, equation (36) simplifies to

$$\tilde{s}_i \equiv \frac{1}{2 - \bar{\lambda}_i} \hat{s}_i + \sum_{j \neq i} \frac{\bar{\lambda}_j}{2 - \bar{\lambda}_i} \hat{s}_j. \quad (40)$$

For any fixed i , we will prove Theorem 19 in three steps:

1. **Lemma 22:** Bounding the deviation between $f_i(\hat{s}_i)$ and $f(\hat{s}_i)$ because \hat{s}_i does *not* necessarily satisfy $f_i(\hat{s}_i) = f(\hat{s}_i)$. Similarly, we also bound the difference between $f_i(\hat{s}_i)$ and $f(\hat{s}_i)$.
2. **Lemma 23:** Using convexity and Theorem 12 to bound $f(\tilde{s}_i)$.
3. **Lemma 24:** Bounding the difference between $f(s_i^+)$ and $f(\tilde{s}_i)$.

We will again need to invoke the Taylor expansion (24) to construct our bounds. However, we require a stronger version of Taylor's Theorem than that in Lemma 8—using “Little-Oh” rather than “Big-Oh”. In particular, for a mapping $g : \mathbf{E} \rightarrow \mathbf{F}$ between two Euclidean spaces and letting $\|\cdot\|^p$ denote an arbitrary Euclidean norm raised to the p -th power, we use the notation $g(u) = o(\|u\|^p)$ to indicate the property of g that—given *any* constant $K > 0$ —there exists $U_K > 0$ such that $|g(u)| \leq K \|u\|^p$ for all $\|u\| \leq U_K$. By itself, we let $o(\|u\|^p)$ denote an element of the class of all functions with this property. Then, an identical argument as that in Lemma 8 shows that the residual function in (24) satisfies

$$r_i(x) = o(\|x\|^2) \text{ and } \nabla r_i(x) = o(\|x\|). \quad (41)$$

We now prove our main result.

Lemma 22. *Assume Setting A, and suppose $\bar{x}=0$. Then, the following inequalities hold:*

$$\begin{aligned} |f_i(\hat{s}_i) - f(\hat{s}_i)| &\leq \frac{L}{2} \left(1 - \frac{\delta}{L}\right) \|s_i\|_2^2 + o(\|s_i\|_2^2), \\ |f_i(\hat{s}_i) - f(\hat{s}_i)| &\leq \frac{L}{2\bar{\lambda}_i} \left(1 - \frac{\bar{\lambda}_i \delta}{L}\right) \|s_i\|_2^2 + o(\|s_i\|_2^2). \end{aligned}$$

Proof. Using the Taylor expansion (24), observe for all pairs i, j that

$$\begin{aligned} |f_i(\hat{s}_i) - f_j(\hat{s}_i)| &= \left| (a_i - a_j)^T \hat{s}_i + \frac{1}{2} \hat{s}_i^T (A_i - A_j) \hat{s}_i + r_i(\hat{s}_i) - r_j(\hat{s}_i) \right| \\ &= \left| \frac{1}{2} \hat{s}_i^T (A_i - A_j) \hat{s}_i + r_i(\hat{s}_i) - r_j(\hat{s}_i) \right| \quad (\text{By } \hat{s}_i \in \bar{U} \text{ and (20)}) \\ &\leq \frac{L}{2} \|\hat{s}_i\|_2^2 + o(\|\hat{s}_i\|_2^2) \quad (\text{By } L\text{-smoothness and (41)}) \\ &\leq \frac{L}{2} \left(1 - \frac{\delta}{L}\right) \|s_i\|_2^2 + o(\|s_i\|_2^2) \quad (\text{Theorem 12}). \end{aligned}$$

Observe the right-hand side is independent of j , so the above still holds after replacing f_j with $f = \max_j f_j$, which gives our desired result. An analogous argument for \hat{s}_i , which replaces $\frac{1}{L}$ with $\frac{\bar{\lambda}_i}{L}$, shows

$$|f_i(\hat{s}_i) - f(\hat{s}_i)| \leq \frac{L}{2\bar{\lambda}_i} \left(1 - \frac{\bar{\lambda}_i \delta}{L}\right) \|s_i\|_2^2 + o(\|s_i\|_2^2).$$

This completes the proof. \square

Lemma 23. *Assume Setting A, and suppose $\bar{x}=0$ and $f(\bar{x})=0$. Then, the following inequality holds:*

$$f(\tilde{s}_i) \leq \frac{L}{\bar{\lambda}_{\min}} \left(1 - \frac{\delta \bar{\lambda}_{\min}}{L}\right) \max_{j=1, \dots, k} \|s_j\|_2^2 + o(\|S\|^2), \quad (42)$$

where \tilde{s}_i and $\bar{\lambda}_{\min}$ are as defined in equations (36) and (37) respectively.

Proof. Using (40), Theorem 11, and convexity, we deduce

$$f(\tilde{s}_i) = f\left(\frac{1}{2 - \bar{\lambda}_i} \hat{s}_i + \sum_{j \neq i} \frac{\bar{\lambda}_j}{2 - \bar{\lambda}_i} \hat{s}_j\right) \leq \frac{1}{2 - \bar{\lambda}_i} f(\hat{s}_i) + \sum_{j \neq i} \frac{\bar{\lambda}_j}{2 - \bar{\lambda}_i} f(\hat{s}_j). \quad (43)$$

Define the quantity

$$\Delta \equiv \frac{L}{2\bar{\lambda}_{\min}} \left(1 - \frac{\bar{\lambda}_{\min}\delta}{L} \right) \left(\frac{1}{2 - \bar{\lambda}_i} \|s_i\|_2^2 + \sum_{j \neq i} \frac{\bar{\lambda}_j}{2 - \bar{\lambda}_i} \|s_j\|_2^2 \right).$$

Applying Lemma 22 to (43) combined with the facts that $\bar{\lambda}_{\min} \leq \bar{\lambda}_j \leq 1$ for all j gives

$$f(\tilde{s}_i) \leq \underbrace{\frac{1}{2 - \bar{\lambda}_i} f_i(\hat{s}_i) + \sum_{j \neq i} \frac{\bar{\lambda}_j}{2 - \bar{\lambda}_i} f_j(\hat{s}_j)}_{(*)} + \Delta + o(\|\mathcal{S}\|^2). \quad (44)$$

Theorem 12 implies the following bound on $(*)$:

$$(*) \leq \frac{1}{2 - \bar{\lambda}_i} \frac{L}{2\bar{\lambda}_i} \left(1 - \frac{\delta\bar{\lambda}_i}{L} \right) \|s_i\|^2 + \frac{L}{2} \left(1 - \frac{\delta}{L} \right) \sum_{j \neq i} \frac{\bar{\lambda}_j}{2 - \bar{\lambda}_i} \|s_j\|_2^2 \leq \Delta.$$

Substituting the above into (44) gives the desired

$$\begin{aligned} f(\tilde{s}_i) &\leq 2\Delta + o(\|\mathcal{S}\|^2) = \frac{L}{\bar{\lambda}_{\min}} \left(1 - \frac{\delta\bar{\lambda}_{\min}}{L} \right) \left[\frac{1}{2 - \bar{\lambda}_i} \|s_i\|^2 + \sum_{j \neq i} \frac{\bar{\lambda}_j}{2 - \bar{\lambda}_i} \|s_j\|_2^2 \right] + o(\|\mathcal{S}\|^2) \\ &\leq \frac{L}{\bar{\lambda}_{\min}} \left(1 - \frac{\delta\bar{\lambda}_{\min}}{L} \right) \max_{j=1, \dots, k} \|s_j\|_2^2 + o(\|\mathcal{S}\|^2), \end{aligned}$$

which is our desired result. \square

Lemma 24. *Assume Setting A, and suppose $\bar{x}=0$ and $f(\bar{x})=0$. Then, the following inequality holds:*

$$f(s_i^+) - f(\tilde{s}_i) \leq \frac{4L}{\bar{\lambda}_{\min}} \left(1 - \frac{\bar{\lambda}_{\min}\delta}{L} \right) \max_{j=1, \dots, k} \|s_i\|_2^2 + O(\|\mathcal{S}\|^3),$$

where \tilde{s}_i and $\bar{\lambda}_{\min}$ are as defined in equations (36) and (37) respectively.

Proof. Consider any fixed i . By Theorem 5, all constraints of $(P_i^{\mathcal{S}})$ are tight, so s_i^+ satisfies the constraints of Observation 6 with equality for all $j \neq i$:

$$\ell_{s_j}^{f_j}(s_i^+) + \frac{L}{2} \|s_i^+ - s_j\|_2^2 = \ell_{s_i}^{f_i}(s_i^+). \quad (45)$$

Define $\xi_i \equiv s_i^+ - \tilde{s}_i$. In other words, ξ_i is the $O(\|\mathcal{S}\|^2)$ term in equation (35). Consider any fixed j . Substituting $s_i^+ = \tilde{s}_i + \xi_i$ into (45) and expanding the quadratic term gives

$$\ell_{s_j}^{f_j}(\tilde{s}_i) + \nabla f_j(s_j)^T \xi_i + \frac{L}{2} \|\tilde{s}_i - s_j\|_2^2 + L \langle \tilde{s}_i - s_j, \xi_i \rangle + \frac{L}{2} \|\xi_i\|^2 = \ell_{s_i}^{f_i}(\tilde{s}_i) + \nabla f_i(s_i)^T \xi_i.$$

Applying the Taylor expansion (24) to $\nabla f_i(s_i)$ and $\nabla f_j(s_j)$ gives

$$\begin{aligned} \ell_{s_j}^{f_j}(\tilde{s}_i) + (a_j + A_j s_j + \nabla r_j(s_j))^T \xi_i + \frac{L}{2} \|\tilde{s}_i - s_j\|_2^2 + L \langle \tilde{s}_i - s_j, \xi_i \rangle + \frac{L}{2} \|\xi_i\|^2 \\ = \ell_{s_i}^{f_i}(\tilde{s}_i) + (a_i + A_i s_i + \nabla r_i(s_i))^T \xi_i. \end{aligned}$$

Note that the terms $\nabla r_j(s_j)$, $\nabla r_i(s_i)$, and \tilde{s}_i are all $O(\|\mathcal{S}\|)$ by Lemma 8 and (36) and so is s_i by definition. Thus, collecting $O(\|\mathcal{S}\|^3)$ terms simplifies the above to

$$\ell_{s_j}^{f_j}(\tilde{s}_i) + a_j^T \xi_i + \frac{L}{2} \|\tilde{s}_i - s_j\|_2^2 + O(\|\mathcal{S}\|^3) = \ell_{s_i}^{f_i}(\tilde{s}_i) + a_i^T \xi_i. \quad (46)$$

Multiplying (46) by $\bar{\lambda}_j$, summing over $j \neq i$, and using identities (14) and (15) gives

$$\sum_{j \neq i} \bar{\lambda}_j \ell_{s_j}^{f_j}(\tilde{s}_i) - \bar{\lambda}_i a_i^T \xi_i + \frac{L}{2} \sum_{j \neq i} \left(\bar{\lambda}_j \|\tilde{s}_i - s_j\|_2^2 \right) + O(\|\mathcal{S}\|^3) = (1 - \bar{\lambda}_i) \ell_{s_i}^{f_i}(\tilde{s}_i) + (1 - \bar{\lambda}_i) a_i^T \xi_i.$$

Adding $\bar{\lambda}_i \ell_{s_i}^{f_i}(\tilde{s}_i)$ to both sides and again applying Lemma 8 leads to

$$\sum_{j=1}^k \bar{\lambda}_j \ell_{s_j}^{f_j}(\tilde{s}_i) + \frac{L}{2} \sum_{j \neq i} \left(\bar{\lambda}_j \|\tilde{s}_i - s_j\|_2^2 \right) + O(\|\mathcal{S}\|^3) = \ell_{s_i}^{f_i}(\tilde{s}_i) + a_i^T \xi_i = \ell_{s_i}^{f_i}(s_i^+) - \underbrace{(A_i s_i + \nabla r_i(s_i))^T \xi_i}_{O(\|\mathcal{S}\|^3)}.$$

Again collecting $O(\|\mathcal{S}\|^3)$ terms and adding a non-negative $\frac{\bar{\lambda}_i L}{2} \|\tilde{s}_i - s_i\|_2^2$ term to the left-hand side leads to the following inequality:

$$\underbrace{\sum_{j=1}^k \bar{\lambda}_j \ell_{s_j}^{f_j}(\tilde{s}_i) + \frac{L}{2} \sum_{j=1}^k \left(\bar{\lambda}_j \|\tilde{s}_i - s_j\|_2^2 \right) + O(\|\mathcal{S}\|^3)}_{(\star)} \geq \ell_{s_i}^{f_i}(s_i^+) \quad (47)$$

For (\star) , the assumed δ -strong convexity on f_j (Definition 3) implies

$$\sum_{j=1}^k \bar{\lambda}_j \ell_{s_j}^{f_j}(\tilde{s}_i) \leq \sum_{j=1}^k \bar{\lambda}_j \left(f_j(\tilde{s}_i) - \frac{\delta}{2} \|\tilde{s}_i - s_j\|_2^2 \right) \leq \underbrace{\max_{j=1, \dots, k} f_j(\tilde{s}_i)}_{=f(\tilde{s}_i)} - \frac{\delta}{2} \sum_{j=1}^k \bar{\lambda}_j \|\tilde{s}_i - s_j\|_2^2, \quad (48)$$

where the right-most inequality follows from (14). Substituting (48) into (47) leads to

$$f(\tilde{s}_i) + \frac{L}{2} \left(1 - \frac{\delta}{L} \right) \sum_{j=1}^k \left(\bar{\lambda}_j \|\tilde{s}_i - s_j\|_2^2 \right) + O(\|\mathcal{S}\|^3) \geq \ell_{s_i}^{f_i}(s_i^+) \geq f_i(s_i^+) - \frac{L}{2} \|s_i^+ - s_i\|_2^2, \quad (49)$$

where the second inequality follows from the L -smoothness of f_i (Definition 3). By Theorem 7, s_i^+ is a valid survey point, so $f_i(s_i^+) = f(s_i^+)$. Substituting this into (49) and rearranging gives

$$f(s_i^+) \leq f(\tilde{s}_i) + \underbrace{\frac{L}{2} \left(1 - \frac{\delta}{L} \right) \sum_{j=1}^k \left(\bar{\lambda}_j \|\tilde{s}_i - s_j\|_2^2 \right)}_{(\dagger)} + \underbrace{\frac{L}{2} \|s_i^+ - s_i\|_2^2}_{(\star)} + O(\|\mathcal{S}\|^3) \quad (50)$$

We use Jensen's inequality to bound (\star) :

$$(\star) \leq \max_{j=1, \dots, k} \|s_i^+ - s_j\|_2^2 = \max_{j=1, \dots, k} \left\| \frac{1}{2} (2s_i^+) + \frac{1}{2} (-2s_j) \right\|_2^2 \leq \frac{1}{2} \|2s_i^+\|_2^2 + \frac{1}{2} \max_{j=1, \dots, k} \|(-2s_j)\|_2^2.$$

Hence,

$$\begin{aligned} (\star) &\leq 2 \left(\|s_i^+\|_2^2 + \max_{j=1, \dots, k} \|s_j\|_2^2 \right) \\ &= 2 \left(2 - \frac{\bar{\lambda}_{\min} \delta}{L} \right) \max_{j=1, \dots, k} \|s_j\|_2^2 + O(\|\mathcal{S}\|^3) \quad (\text{Theorem 14}) \\ &\leq 4 \max_{j=1, \dots, k} \|s_j\|_2^2 + O(\|\mathcal{S}\|^3). \end{aligned}$$

An analogous argument (using Lemma 13 on \tilde{s}_i instead of Theorem 14 on s_i^+) gives

$$(\dagger) \leq \max_{j=1, \dots, k} \|\tilde{s}_i - s_j\|_2^2 \leq 4 \max_{j=1, \dots, k} \|s_j\|_2^2.$$

Substituting these bounds on (†) and (★) into (50), we deduce

$$\begin{aligned} f(s_i^+) &\leq f(\tilde{s}_i) + 2L \left(1 - \frac{\delta}{L}\right) \max_{j=1,\dots,k} \|s_i\|_2^2 + 2L \max_{j=1,\dots,k} \|s_i\|_2^2 + O(\|\mathcal{S}\|^3) \\ &= f(\tilde{s}_i) + 4L \left(1 - \frac{\delta}{2L}\right) \max_{j=1,\dots,k} \|s_i\|_2^2 + O(\|\mathcal{S}\|^3). \end{aligned}$$

If $k = 1$, Survey Descent reduces to GD and the result is trivial from Beck (2017, Theorem 10.29) or Nesterov (2003, Theorem 2.1.15). When $k \geq 2$, the property (14) implies $\bar{\lambda}_{\min} \leq \frac{1}{k} \leq \frac{1}{2}$. Thus, $(1 - \frac{\delta}{2L}) \leq (1 - \frac{\bar{\lambda}_{\min}\delta}{L})$. We then deduce

$$f(s_i^+) \leq f(\tilde{s}_i) + \frac{4L}{\bar{\lambda}_{\min}} \left(1 - \frac{\bar{\lambda}_{\min}\delta}{L}\right) \max_{j=1,\dots,k} \|s_i\|_2^2 + O(\|\mathcal{S}\|^3).$$

This completes the proof. \square

Theorem 19. (*R-Linear Convergence of Objective Values, Strong Version*) Consider Procedure B. Fix $\bar{K} > \frac{5L}{\bar{\lambda}_{\min}} \left(1 - \frac{\bar{\lambda}_{\min}\delta}{L}\right)$. Then,

$$f(s_i^{t+1}) - f(\bar{x}) \leq \bar{K} \|\mathcal{S}^t - \bar{\mathcal{X}}\|^2 \quad \forall i = 1, \dots, k, \text{ and } t \in \mathbf{N}.$$

Consequently, survey objective values $\{f(s_i^t)\}_{t \in \mathbf{N}}$ converge R -linearly to $f(\bar{x})$ for all i .

Proof. Without loss of generality, assume $\bar{x}=0$ and $f(\bar{x})=0$. Consider any fixed i and t . For simplicity, adopt the notation $s_i = s_i^t$ and $s_i^+ = s_i^{t+1}$. Adding the inequalities in Lemmas 23 and 24 gives

$$f(s_i^+) \leq \frac{5L}{\bar{\lambda}_{\min}} \left(1 - \frac{\bar{\lambda}_{\min}\delta}{L}\right) \max_{j=1,\dots,k} \|s_j\|_2^2 + \underbrace{o(\|\mathcal{S}\|^2) + O(\|\mathcal{S}\|^3)}_{=o(\|\mathcal{S}\|^2)}.$$

Then, for any $K > 0$ and for all \mathcal{S} sufficiently close to $\bar{\mathcal{X}}$, we have

$$f(s_i^+) \leq \frac{5L}{\bar{\lambda}_{\min}} \left(1 - \frac{\bar{\lambda}_{\min}\delta}{L}\right) \|\mathcal{S}\|^2 + K \|\mathcal{S}\|^2 = \left(\frac{5L}{\bar{\lambda}_{\min}} \left(1 - \frac{\bar{\lambda}_{\min}\delta}{L}\right) + K\right) \|\mathcal{S}\|^2.$$

Denote $\bar{K} \equiv \frac{5L}{\bar{\lambda}_{\min}} \left(1 - \frac{\bar{\lambda}_{\min}\delta}{L}\right) + K > 0$. Since K can be any strictly positive constant, choosing K is equivalent to choosing an arbitrary $\bar{K} > \frac{5L}{\bar{\lambda}_{\min}} \left(1 - \frac{\bar{\lambda}_{\min}\delta}{L}\right)$. Under such a choice,

$$f(s_i^+) \leq \bar{K} \|\mathcal{S}\|^2,$$

which is our desired result. The general case follows from applying the change of variables $x \leftarrow (x - \bar{x})$ and function-value shift $f(\cdot) \leftarrow (f(\cdot) - f(\bar{x}))$. \square

B Geometric Proof of Theorem 5

Although Theorem 5 can be proven by verifying first-order optimality conditions and applying the implicit function theorem, we present here a more direct, computationally revealing approach that highlights the underlying geometry of Survey Descent. In particular, after change-of-variables and tightness-of-constraints arguments (presented later), Observation 6 reduces to a simple, geometric *multiple balls problem*.

Definition 25. (Multiple Balls Problem, MBP(y, Z, r)) For some target vector $y \in \mathbf{R}^n$, ball centers $Z = [z_1, \dots, z_m] \in \mathbf{R}^{n \times m}$ where $m \leq n$, and non-negative radii $r = [r_1, \dots, r_m]^T \in \mathbf{R}_+^m$, define the *Multiple Balls Problem (MBP)* as

$$\min_{x \in \mathbf{R}^n} \|x - y\|_2^2$$

$$\text{s.t. } \|x - z_j\|_2^2 = r_j^2 \quad \forall j = 1, \dots, m. \quad (51)$$

Denote the problem as MBP (y, Z, r) and its *feasible region* as $\mathcal{B}(Z, r)$. We will refer to the tuple (y, Z, r) as the *parameters* and the sub-tuple (Z, r) as the *ball-specifications* of the MBP.

Thus, MBP simply projects y onto the boundary-intersection, $\mathcal{B}(Z, r)$, of m balls of radii $\{r_j\}_{j=1}^m$ —each centered at the corresponding $\{z_j\}_{j=1}^m$. This is illustrated in Figure 8a.

Since y , Z , and r all reside in Euclidean spaces, we can also consider (Z, r) and (y, Z, r) elements of some product space. Let $\|\cdot\|$ denote an arbitrary Euclidean norm. This will allow us to formalize the notion of distance in of parameters and ball-specifications in their respective spaces. We are particularly interested in the behavior of MBP (y, Z, r) where (y, Z, r) is near some set of *base parameters*.

Definition 26. (Base Parameters) We refer to a set of parameters as *base parameters* and denote them as $(\bar{y}, \bar{Z}, \bar{r})$ if they satisfy the following properties:

- The columns of $\bar{Z} = [\bar{z}_1, \dots, \bar{z}_m]$ are linearly independent.
- $\bar{r}_j^2 = \|\bar{z}_j\|_2^2$.
- There exist strictly-positive weights $\check{\lambda} = [\check{\lambda}_1, \dots, \check{\lambda}_m]^T \in \mathbf{R}_{++}^m$ such that $\bar{y} = -\sum_{j=1}^m \check{\lambda}_j \bar{z}_j$.

Note that $\check{\lambda}$ is necessarily unique due to the linear independence of \bar{Z} , and their strict-positivity are connected to the nondegenerate first-order optimality that we will later show for MBP $(\bar{y}, \bar{Z}, \bar{r})$.

The remainder of Appendix B is organized as follows. For parameters (y, Z, r) near some base parameters $(\bar{y}, \bar{Z}, \bar{r})$, Section B.1 derives an explicit form solution for MBP (y, Z, r) , Section B.2 shows that the solution is the same if one were to solve MBP (y, Z, r) with inequality rather than equality constraints, and Section B.3 connects the MBPs to Survey Descent.

B.1 Local, Explicit-Form Solution to the Multiple Balls Problem

In Theorem 32, we will build the explicit-form solution for MBP (y, Z, r) —for parameters near $(\bar{y}, \bar{Z}, \bar{r})$ —from $\mathcal{B}(Z, r)$ as well as the following geometric objects (whose well-definedness and other properties will be derived in Lemmas 27-31):

- The *central affine subspace*

$$\mathcal{A}(Z, r) \equiv \left\{ x \in \mathbf{R}^n : 2z_j^T x - \|z_j\|_2^2 + r_j^2 \text{ equal for all } j = 1, \dots, m \right\}. \quad (52)$$

- The *core*

$$c(Z, r) \equiv \mathcal{P}_{\mathcal{A}(Z, r)}(z_j) \quad \text{for any } j = 1, \dots, m, \quad (53)$$

where $\mathcal{P}_{\mathcal{A}(Z, r)}$ denotes the projection onto $\mathcal{A}(Z, r)$.

- The *margin*

$$\rho(Z, r) \equiv \sqrt{r_j^2 - \|z_j - c(Z, r)\|_2^2} \quad \text{for any } j = 1, \dots, m. \quad (54)$$

They are collectively visualized in Figure 8b.

Lemma 27. (Local Properties of the Central Affine Subspace) For any ball-specifications (Z, r) sufficiently close to base ball-specifications (\bar{Z}, \bar{r}) , the following properties hold:

1. $\mathcal{A}(Z, r)$ is non-empty with dimension $n - m + 1$.
2. For all $i \neq j$, the direction $(z_i - z_j)$ is orthogonal to the linear subspace parallel to $\mathcal{A}(Z, r)$.
3. For $w \in \mathbf{R}^n$, the projection $\mathcal{P}_{\mathcal{A}(Z, r)}(w)$ is well-defined and varies smoothly with (Z, r) and w .




Figure 8: **Visualization of the Multiple Balls Problem.** Given a target $y \in \mathbf{R}^n$, ball centers $Z = [z_1, \dots, z_m] \in \mathbf{R}^{n \times m}$, and corresponding radii $r = [r_1, \dots, r_m]^T \in \mathbf{R}^m$, the above displays the geometries explored in Appendix B for $n = 3$ and $m = 2$. Panel (a) shows the Multiple Balls Problem in which the point y is projected onto the boundary-intersection of the m -balls specified by (Z, r) . Panel (b) illustrates the four fundamental structures explored in Appendix B: the central affine subspace $\mathcal{A}(Z, r)$, the core $c(Z, r)$, the margin $\rho(Z, r)$, and the feasible region $\mathcal{B}(Z, r)$.

Proof. Observe that we can rewrite (52) as the following $(m-1) \times n$ linear system:

$$\mathcal{A}(Z, r) = \{x \in \mathbf{R}^n : M(Z, r)x = b(Z, r)\}, \quad (55)$$

where

$$M(Z, r) \equiv 2 \begin{bmatrix} (z_2 - z_1)^T \\ \vdots \\ (z_m - z_1)^T \end{bmatrix} \text{ and } b(Z, r) \equiv \begin{bmatrix} (\|z_2\|_2^2 - r_2^2) - \|z_1\|_2^2 + r_1^2 \\ \vdots \\ (\|z_m\|_2^2 - r_m^2) - \|z_1\|_2^2 + r_1^2 \end{bmatrix}.$$

Consider $\mathcal{A}(\bar{Z}, \bar{r})$. Since $\{\bar{z}_j\}_{j=1}^m$ are linearly independent, $\{\bar{z}_j - \bar{z}_1\}_{j=2}^m$ —which are the rows of $M(\bar{Z}, \bar{r})$ —are also linearly independent. For (Z, r) sufficiently close to (\bar{Z}, \bar{r}) , the rows of $M(Z, r)$ will remain linearly independent. In this case, the linear system (55) is then necessarily solvable: $\mathcal{A}(Z, r)$ is non-empty. Let x_0 be such a solution. Then, $\mathcal{A}(Z, r)$ consists of all vectors $x \in \mathbf{R}^n$ such that $x - x_0$ is orthogonal to the $m-1$ linearly independent rows of $M(Z, r)$. Thus, $\dim \{\mathcal{A}(Z, r)\} = n - m + 1$.

The vectors $\{(z_i - z_j)\}_{j \neq i}$ clearly all lie in the rowspace of $M(Z, r)$. Hence, the vectors $\{(z_i - z_j)\}_{j \neq i}$ are orthogonal to all direction in $\mathcal{A}(Z, r)$.

Finally, write M instead of $M(Z, r)$ for simplicity. Since M has independent rows, the projection of w onto $\mathcal{A}(Z, r)$ is explicitly given by

$$\mathcal{P}_{\mathcal{A}(Z, r)}(w) = \left(I - M(M^T M)^{-1} M^T\right) w + M(M^T M)^{-1} b,$$

which varies smoothly with (Z, r) and w . □

Lemma 28. (Local Properties of Core) For any ball-specifications (Z, r) sufficiently close to base ball-specifications (\bar{Z}, \bar{r}) , the core in (53) is well-defined and varies smoothly with (Z, r) .

Proof. By the second fact in Lemma 27, for any pair $i \neq j$, the difference $(z_i - z_j)$ is orthogonal to the linear subspace parallel to $\mathcal{A}(y, Z, r)$. Therefore, $\mathcal{P}_{\mathcal{A}(y, Z, r)}(z_i)$ and $\mathcal{P}_{\mathcal{A}(y, Z, r)}(z_j)$ are the same. Hence, (53) is well-defined. The smoothness of the $c(Z, r)$ with respect to (Z, r) then follows from (53) and the third fact of Lemma 27. \square

Lemma 29. (The Base Case) *At base ball-specifications, the central affine subspace is a linear subspace and the core is non-zero.*

Proof. By Definition 26, $\bar{r}_j^2 = \|\bar{z}_j\|_2^2$ for all j . Thus, (52) simplifies to

$$\mathcal{A}(\bar{Z}, \bar{r}) = \{x \in \mathbf{R}^n : 2\bar{z}_j^T x \text{ equal for all } j = 1, \dots, m\}.$$

Clearly, the origin satisfies the RHS conditions above, so $0 \in \mathcal{A}(\bar{Z}, \bar{r})$.

Next, assume for contradiction that $0 = c(\bar{Z}, \bar{r})$. By definition, $c(\bar{Z}, \bar{r})$ is the projection of \bar{z}_j onto $\mathcal{A}(\bar{Z}, \bar{r})$ for any j , so the vectors $\{\bar{z}_j - c(\bar{Z}, \bar{r})\}_{j=1}^m = \{\bar{z}_j\}_{j=1}^m$ must be orthogonal to $\mathcal{A}(\bar{Z}, \bar{r})$. The linear independence of $\{\bar{z}_j\}_{j=1}^m$ implies $\dim\{\mathcal{A}(\bar{Z}, \bar{r})\} \leq n - m$, which contradicts the first fact of Lemma 27. Thus, $0 \neq c(\bar{Z}, \bar{r})$. \square

Lemma 30. (Local Properties of the Margin) *For ball-specifications (Z, r) near some base ball-specifications (\bar{Z}, \bar{r}) , the margin in (54) is well-defined, strictly-positive, and varies smoothly with (Z, r) .*

Proof. It suffices to show that the square of the margin,

$$r_j^2 - \|z_j - c(Z, r)\|_2^2, \quad (56)$$

does not depend on the index j , is strictly-positive, and varies smoothly with (y, Z, r) . First, consider the core and central affine subspace exactly at base ball-specifications (\bar{Z}, \bar{r}) . By Lemma 29, since $0 \in \mathcal{A}(\bar{Z}, \bar{r})$ and $\mathcal{P}_{\mathcal{A}(\bar{Z}, \bar{r})}(\bar{z}_j) = c(\bar{Z}, \bar{r}) \neq 0$ for all j , the definition of a projection gives

$$\|\bar{z}_j - c(\bar{Z}, \bar{r})\|_2 < \|\bar{z}_j - 0\|_2 = \bar{r}_j \quad \forall j = 1, \dots, m.$$

Thus, (56) is strictly positive, which implies the argument of the square-root in (54) is strictly positive. Next, observe that

$$r_j^2 - \|z_j - c(Z, r)\|_2^2 = -\|c(Z, r)\|_2^2 + \underbrace{2z_j^T c(Z, r) - \|z_j\|_2^2 + r_j^2}_{(\star)_j} \quad \forall j = 1, \dots, m.$$

Since $c(Z, r) \in \mathcal{A}(Z, r)$, the term $(\star)_j$ is equal for all j . Therefore, (56) is j -independent, and $\rho(Z, r)$ in (54) is well-defined. The smoothness of $\rho(Z, r)$ with respect to (Z, r) then follows directly from (54) and the smoothness of the core (Lemma 28). \square

Lemma 31. (Local Properties of the Feasible Region) *For ball-specifications (Z, r) near some base ball-specifications (\bar{Z}, \bar{r}) , the feasible region is nonempty and satisfies*

$$\mathcal{B}(Z, r) = \{x \in \mathcal{A}(Z, r) : \|x - c(Z, r)\|_2 = \rho(Z, r)\}. \quad (57)$$

Proof. Equivalence of (51) and (57): Expanding and rearranging (51) gives

$$\begin{aligned} \|x - z_j\|_2 = r_j \quad \forall j &\iff \|x\|_2^2 = 2z_j^T x - \|z_j\|_2^2 + r_j^2 \quad \forall j \\ &\implies 2z_j^T x - \|z_j\|_2^2 + r_j^2 = 2z_i^T x - \|z_i\|_2^2 + r_i^2 \quad \forall i \neq j. \end{aligned}$$

Hence, $\mathcal{B}(Z, r) \subseteq \mathcal{A}(Z, r)$ implying

$$\mathcal{B}(Z, r) = \{x \in \mathcal{A}(Z, r) : \|x - z_j\|_2 = r_j \text{ for all } j\}.$$

Since $c(r, Z) = \mathcal{P}_{\mathcal{A}(Z, r)}(z_i)$, for any $x \in \mathcal{A}(Z, r)$, the vectors $(x - c(r, Z))$ and $(z_j - c(Z, r))$ are orthogonal. Then, applying the Pythagorean Theorem to the equality condition in (51) gives

$$r_j^2 = \|x - z_i\|_2^2 = \|x - c(r, Z)\|_2^2 + \|z_j - c(Z, r)\|_2^2 \quad \forall x \in \mathcal{A}(Z, r) \text{ and } j = 1, \dots, k.$$

We then deduce (57) by subtracting $\|z_j - c(Z, r)\|_2^2$ from the left and right sides, and applying the definition of $\rho(Z, r)$ in (54).

Non-Emptiness of $\mathcal{B}(Z, r)$: By the first fact of Lemma 27, the affine subspace $\mathcal{A}(Z, r)$ is $(n-m+1)$ -dimensional, so $\mathcal{A}(Z, r)$ is a non-empty space. Thus, there exists $a \in \mathcal{A}(Z, r)$ such that $a \neq c(Z, r)$. Choose any such a and consider the direction

$$d = \frac{1}{\|a - c(Z, r)\|} (a - c(Z, r))$$

along with the vector

$$v = c(Z, r) + \rho(Z, r) d.$$

Then, by construction,

$$\|v - c(Z, r)\|_2 = \rho(Z, r).$$

Since $\mathcal{A}(Z, r)$ contains both $c(Z, r)$ and d , we deduce $v \in \mathcal{A}(Z, r)$. Thus, v satisfies all conditions in (57), so $v \in \mathcal{B}(Z, r)$, and hence $\mathcal{B}(Z, r)$ is non-empty. \square

Theorem 32. (Local, Explicit Solution of MBP (y, Z, r)) For parameters (y, Z, r) near some base parameters $(\bar{y}, \bar{Z}, \bar{r})$, the optimal solution to MBP (y, Z, r) is unique, varies smoothly with (y, Z, r) , and is given by

$$x(y, Z, r) = c(Z, r) + \rho(Z, r) u(y, Z, r), \quad (58)$$

where

$$u(y, Z, r) = \frac{1}{\|\mathcal{P}_{\mathcal{A}(Z, r)}(y) - c(Z, r)\|_2} (\mathcal{P}_{\mathcal{A}(Z, r)}(y) - c(Z, r)). \quad (59)$$

In particular, $x(\bar{y}, \bar{Z}, \bar{r}) = 0$.

Proof. The vector $u(y, Z, r)$ is well-defined: To show the well-definedness of the expression in (59), observe the following at base parameters $(\bar{y}, \bar{Z}, \bar{r})$:

$$\begin{aligned} \mathcal{P}_{\mathcal{A}(\bar{Z}, \bar{r})}(\bar{y}) - c(\bar{Z}, \bar{r}) &= \mathcal{P}_{\mathcal{A}(\bar{Z}, \bar{r})} \left(- \sum_{j=1}^m \check{\lambda}_j \bar{z}_j \right) - c(\bar{Z}, \bar{r}) && \text{(By Definition 26)} \\ &= - \sum_{j=1}^m \check{\lambda}_j \mathcal{P}_{\mathcal{A}(\bar{Z}, \bar{r})}(\bar{z}_j) - c(\bar{Z}, \bar{r}) && \text{(By linearity of } \mathcal{A}(\bar{Z}, \bar{r}) \text{ from Lemma 29)} \\ &= - \left(\sum_{j=1}^m \check{\lambda}_j + 1 \right) c(\bar{Z}, \bar{r}) && \text{(By definition of core (53))} \\ &\neq 0. && \text{(By Definition 26 and Lemma 29).} \end{aligned}$$

By the smoothness of the projection (third fact of Lemma 27), we deduce $\mathcal{P}_{\mathcal{A}(Z, r)}(y) - c(Z, r) \neq 0$ for (y, Z, r) sufficiently close to $(\bar{y}, \bar{Z}, \bar{r})$ as well. Thus, (59) is well-defined. Moreover, the smoothness of the projection and (59) further shows that $u(y, Z, r)$ is smooth with respect to (y, Z, r) .

Derivation of $x(y, Z, r)$: By Lemma 31, $\mathcal{B}(Z, r) \subseteq \mathcal{A}(Z, r)$ so

$$x(y, Z, r) = \mathcal{P}_{\mathcal{B}(Z, r)}(y) = \mathcal{P}_{\mathcal{B}(Z, r)}(\mathcal{P}_{\mathcal{A}(Z, r)}(y)).$$

Lemma 31 also states that $\mathcal{B}(Z, r)$ is simply the $c(Z, r)$ -centered ball with radius $\rho(Z, r)$ restricted to $\mathcal{A}(Z, r)$. Thus, $\mathcal{P}_{\mathcal{B}(Z, r)}(\mathcal{P}_{\mathcal{A}(Z, r)}(y))$ must be the (necessarily unique) point that

1. lies on the straight line from $c(Z, r)$ to $\mathcal{P}_{\mathcal{A}(Z, r)}(y)$, and
2. is exactly the distance $\rho(Z, r)$ from $c(Z, r)$.

Formalizing the above two criteria leads to the explicit-form

$$x(y, Z, r) = c(Z, r) + \frac{\rho(Z, r)}{\|\mathcal{P}_{\mathcal{A}(Z, r)}(y) - c(Z, r)\|_2} (\mathcal{P}_{\mathcal{A}(Z, r)}(y) - c(Z, r)) = c(Z, r) + \rho(Z, r) u(y, Z, r).$$

The smoothness of $x(y, Z, r)$ follows directly from the smoothness of the core (Lemma 28), the margin (Lemma 30), and $u(y, Z, r)$ (shown above).

Demonstration of $x(\bar{y}, \bar{Z}, \bar{r}) = 0$: Observe that the Lagrangian of MBP $(\bar{y}, \bar{Z}, \bar{r})$ is

$$\mathcal{L}(x, \lambda) = \|x - \bar{y}\|_2^2 + \sum_{j=1}^m \lambda_j \|x - \bar{z}_j\|_2^2,$$

where λ_j is the j -th element of a Lagrange multipliers vector $\lambda \in \mathbf{R}^m$. By Definitions 25-26, the zero-vector is a feasible solution to MBP $(\bar{y}, \bar{Z}, \bar{r})$. A feasible vector x^* is the solution of MBP $(\bar{y}, \bar{Z}, \bar{r})$ if and only if there exists λ^* such that x^* minimizes $\mathcal{L}(\cdot, \lambda^*)$. It is easy to check using Definition 26 and first-order optimality conditions that $x^* = 0$ is the unique minimizer of $\mathcal{L}(\cdot, \lambda)$. Thus, $x(\bar{y}, \bar{Z}, \bar{r}) = 0$ is the unique solution to MBP $(\bar{y}, \bar{Z}, \bar{r})$. \square

B.2 Lagrange Multipliers and Inequality Constraints

In this section, we will show that the optimal solutions of MBPs are unchanged when the equality constraints (51) are relaxed to inequalities. In particular, regardless of whether MBP (y, Z, r) is defined with equality or inequality constraints, note that the stationarity conditions remain the same. Therefore, there must exist some Lagrange multipliers, $\lambda(y, Z, r) \in \mathbf{R}^m$, such that

$$x(y, Z, r) - y + \sum_{j=1}^m \lambda_j(y, Z, r) [x(y, Z, r) - z_j] = 0, \quad (60)$$

where $\lambda_j(y, Z, r)$ denotes the j -th element of $\lambda(y, Z, r)$. The only difference is that the inequality version possesses additional dual feasibility conditions requiring the non-negativity of $\lambda(y, Z, r)$. Lemma 33 will show that the $\lambda(y, Z, r)$ are unique and strictly positive.

Then, for (y, Z, r) sufficiently close to $(\bar{y}, \bar{Z}, \bar{r})$, Theorem 34 will show that the KKT conditions are sufficient to certify the optimality of $(x(y, Z, r), \lambda(y, Z, r))$ for inequality-constrained MBPs as well.

Lemma 33. (Properties of MBP Lagrange Multipliers) *Suppose the parameters (y, Z, r) are sufficiently close to base parameters $(\bar{y}, \bar{Z}, \bar{r})$ with corresponding base weights vector $\check{\lambda}$. Consider the optimal solution $x(y, Z, r)$ of MBP (y, Z, r) . Then, the following hold:*

1. *There exists a unique Lagrange multiplier vector, $\lambda(y, Z, r) \in \mathbf{R}_{++}^m$, satisfying the stationarity condition (60).*
2. *The vector $\lambda(y, Z, r)$ varies smoothly with (y, Z, r) .*
3. *At the base parameters, $\lambda(\bar{y}, \bar{Z}, \bar{r}) = \check{\lambda}$.*

Proof. Since $x(y, Z, r)$ is optimal for MBP (y, Z, r) , there must exist a Lagrange multipliers vector $\lambda(y, Z, r) \in \mathbf{R}^m$ satisfying (60). We first show $\lambda(y, Z, r)$ is unique. Finding Lagrange multipliers satisfying (60) is equivalent to finding minimizers of the function

$$F(\lambda) = \|x(y, Z, r) - y + X\lambda\|_2^2,$$

where we define the matrix

$$X = [x(y, Z, r) - z_1, \dots, x(y, Z, r) - z_m] \in \mathbf{R}^{n \times m}.$$

Then, a vector λ minimizes F if and only if it satisfies the below $m \times m$ linear system:

$$X^T (x(y, Z, r) - y + X\lambda) = 0 \iff X^T X\lambda = X^T (y - x(y, Z, r)).$$

By Theorem 32, for (y, Z, r) close to $(\bar{y}, \bar{Z}, \bar{r})$, the MBP-solution $x(y, Z, r)$ is close to zero. Therefore, columns of X are close to $\{-\bar{z}_j\}_{j=1}^m$, which are linearly independent (Definition 26). This implies $X^T X$ is invertible, and the minimizer of F is unique and explicitly given by

$$\lambda(y, Z, r) = \underset{\lambda}{\operatorname{argmin}} F(\lambda) = (X^T X)^{-1} X^T (y - x(y, Z, r)).$$

The smoothness of $\lambda(y, Z, r)$ with respect to (y, Z, r) follows from the smoothness of $x(y, Z, r)$ (Theorem 32) and the smoothness of the above matrix operations.

Finally, it follows directly from Definition 26 and Theorem 32 that

$$x(\bar{y}, \bar{Z}, \bar{r}) - \bar{y} + \sum_{j=1}^m \check{\lambda}_j [x(\bar{y}, \bar{Z}, \bar{r}) - \bar{z}_j] = 0,$$

which implies the Lagrange multiplier vector at the base parameters is uniquely $\lambda(\bar{y}, \bar{Z}, \bar{r}) = \check{\lambda}$. \square

Theorem 34. (Inequality-Constrained MBP) Consider the following inequality-constrained MBP:

$$\begin{aligned} \min_{x \in \mathbf{R}^n} \quad & \|x - y\|_2^2 \\ \text{s.t.} \quad & \|x - z_j\|_2^2 \leq r_j^2 \quad \forall j = 1, \dots, m, \end{aligned}$$

and denote it as $\text{MBP}_{\leq}(y, Z, r)$. For parameters (y, Z, r) sufficiently close to $(\bar{y}, \bar{Z}, \bar{r})$, $\text{MBP}_{\leq}(y, Z, r)$ is feasible and possesses the same unique optimal solution, $x(y, Z, r)$, and unique Lagrange multipliers, $\lambda(y, Z, r)$, as $\text{MBP}(y, Z, r)$. Consequently, all constraints of $\text{MBP}_{\leq}(y, Z, r)$ are tight. Moreover, the optimal solution and multipliers satisfy

$$x(y, Z, r) = O(\|(y, Z, r) - (\bar{y}, \bar{Z}, \bar{r})\|) \quad \text{and} \quad \lambda(y, Z, r) = \check{\lambda} + O(\|(y, Z, r) - (\bar{y}, \bar{Z}, \bar{r})\|). \quad (61)$$

Proof. The feasibility of $\text{MBP}_{\leq}(y, Z, r)$ follows directly from Lemma 31 since $\mathcal{B}(y, Z, r)$ is a subset of the feasible region of $\text{MBP}_{\leq}(y, Z, r)$.

$\text{MBP}_{\leq}(y, Z, r)$ possesses a unique minimizer since it minimizes a strongly convex objective on a convex feasible region. Moreover, checking the KKT conditions is sufficient for demonstrating optimality because the $\|x - z_j\|_2^2$ -term in the constraints is convex and differentiable for all j . Therefore—denoting the unique pair of optimal solutions and Lagrange multipliers of $\text{MBP}(y, Z, r)$ as $(x(y, Z, r), \lambda(y, Z, r))$ —it suffices to check that $(x(y, Z, r), \lambda(y, Z, r))$ satisfies the KKT conditions of $\text{MBP}_{\leq}(y, Z, r)$:

1. **(Stationarity)** Follows from $\text{MBP}_{\leq}(y, Z, r)$ and $\text{MBP}(y, Z, r)$ having the same stationarity condition (60).
2. **(Primal Feasibility)** Follows from $\mathcal{B}(y, Z, r)$ being a subset of the feasible region of $\text{MBP}_{\leq}(y, Z, r)$.
3. **(Dual Feasibility)** The inequality constraints of $\text{MBP}_{\leq}(y, Z, r)$ require the non-negativity of all Lagrange multipliers. By Definition 26, the elements of $\check{\lambda}$ are strictly-positive. Thus, by Lemma 33, the elements of $\lambda(y, Z, r)$ are also strictly-positive for (y, Z, r) sufficiently close to $(\bar{y}, \bar{Z}, \bar{r})$.
4. **(Complementary Slackness)** By construction, $x(y, Z, r)$ satisfies all the inequality constraints of $\text{MBP}_{\leq}(y, Z, r)$ with equality. Correspondingly, $\lambda(y, Z, r)$ are all strictly positive.

Therefore, $(x(y, Z, r), \lambda(y, Z, r))$ is the unique pair of optimal solutions and Lagrange multipliers for $\text{MBP}_{\leq}(y, Z, r)$ as well. The characterizations in (61) directly follows from Theorem 32 and Lemma 33 as well as recalling that smooth functions are locally Lipschitz. \square

B.3 Survey Descent as a Inequality Constrained MBP

We can reformulate Survey Descent subproblems into inequality-constrained MBPs. In particular, for all i , we can rewrite the i -th Survey Descent subproblem in Observation 6 as

$$\begin{aligned}
& \min_x \| (x - s_i) - y_i(\mathcal{S}) \|_2^2 \\
& \text{s.t. } \| (x - s_i) - z_{ij}(\mathcal{S}) \|_2^2 \leq r_{ij}^2(\mathcal{S}) \quad \forall j \neq i,
\end{aligned} \tag{62}$$

where, for all i and j , we define the following mappings using the components $\{f_i\}_{i=1}^k$ from Definition 3:

$$\begin{aligned}
y_i(\mathcal{S}) &\equiv -\frac{1}{L} \nabla f_i(s_i), \\
z_{ij}(\mathcal{S}) &\equiv s_j - s_i + y_j(\mathcal{S}) - y_i(\mathcal{S}), \\
r_{ij}^2(\mathcal{S}) &\equiv \frac{2}{L} \left[f_i(s_i) - f_j(s_j) + \nabla f_i(s_i)^T (s_j - s_i) \right] + \|y_j(\mathcal{S}) - y_i(\mathcal{S})\|_2^2.
\end{aligned} \tag{63}$$

Observation 35. (Survey Descent as Inequality-Constrained MBP) Consider any k -point survey \mathcal{S} and any $i=1, \dots, k$. After the changes-of-variables $x \leftarrow x - s_i$, the minimization problem in (62) is an inequality-constrained MBP with $m = k-1$ and parameters

$$(y, Z, r) = \left(y_i(\mathcal{S}), [z_{ij}(\mathcal{S})]_{j \neq i}, [r_{ij}(\mathcal{S})]_{j \neq i}, [s_j]_{j \neq i} \right),$$

where, for any j -indexed scalars $\{v_j\}_{j=1}^k$, we use the notation $[v_j]_{j \neq i} \equiv [v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_k]^T$ to denote the length $k-1$ vector excluding v_i . Furthermore, the mappings (63) evaluated at the base survey⁵ $\bar{\mathcal{X}}$ (Definition 4) form a set of base parameters

$$(\bar{y}, \bar{Z}, \bar{r}) = \left(y_i(\bar{\mathcal{X}}), [z_{ij}(\bar{\mathcal{X}})]_{j \neq i}, [r_{ij}(\bar{\mathcal{X}})]_{j \neq i}, [\bar{x}]_{j \neq i} \right), \tag{64}$$

with corresponding weights $\{\bar{\lambda}_j\}_{j \neq i}$ from (14)-(15), and where $[\bar{x}]_{j \neq i}$ is the $k-1$ length vector with all \bar{x} entries.

We then deduce our desired theorem.

Theorem 5. (Local Feasibility, Uniqueness, Tightness, and Smoothness of Survey Descent) Assume Setting A. Then for all i , the Survey Descent subproblem (P_i^S) is feasible and has a unique solution s_i^+ , which satisfies all the constraints with equality, with unique associated Lagrange multipliers λ_j^i for the j -th constraint within the i -th subproblem (for $j \neq i$). The solutions and multipliers depend smoothly on the input survey \mathcal{S} , and satisfy

$$s_i^+(\mathcal{S}) = \bar{x} + O(\|\mathcal{S} - \bar{\mathcal{X}}\|) \quad \text{and} \quad \lambda_j^i(\mathcal{S}) = \bar{\lambda}_j + O(\|\mathcal{S} - \bar{\mathcal{X}}\|) \quad \forall j \neq i. \tag{16}$$

Proof. It is easy to check from the mapping definitions (63) and the Taylor expansion (24) that

$$\left\| \left(y_i(\mathcal{S}), [z_{ij}(\mathcal{S})]_{j \neq i}, [r_{ij}(\mathcal{S})]_{j \neq i}, [s_j]_{j \neq i} \right) - \left(y_i(\bar{\mathcal{X}}), [z_{ij}(\bar{\mathcal{X}})]_{j \neq i}, [r_{ij}(\bar{\mathcal{X}})]_{j \neq i}, [\bar{x}]_{j \neq i} \right) \right\| = O(\|\mathcal{S} - \bar{\mathcal{X}}\|).$$

The result then follows directly from Observation 6, Observation 35, and Theorem 34. \square

Observe that the proofs of Lemmas 27-33 explicitly construct both $x(y, Z, r)$ and $\lambda(y, Z, r)$ using simple linear algebra operations and a scalar square-root. Thus, combining these lemmas with Observation 35, we deduce the below computationally-efficient procedure for solving Survey Descent subproblems within the setting of Theorem 5. Moreover, recall that $\nabla f(s_i) = \nabla f_i(s_i)$ since \mathcal{S} is valid, so the procedure requires only a first-order oracle on f without need to access its $\{f_i\}_{i=1}^k$ components.

Observation 36. (Survey Descent with Linear Algebra) Assume Setting A. Consider $i=k$ without loss of generality. We can compute the solution, $s_k^+(\mathcal{S})$, and Lagrange multipliers, $\{\lambda_j(\mathcal{S})\}_{j=1}^{k-1}$, of the Survey Descent subproblem (P_k^S) using the following procedure:

⁵Evaluating of the mappings (63) at $\bar{\mathcal{X}}$ is a technical step. $\bar{\mathcal{X}}$ is not a valid survey (Definition 4), so it does not correspond to an actual Survey Descent iteration.

1. Compute

$$\begin{aligned}
y_j &\leftarrow -\frac{1}{L} \nabla f(s_j) & \forall j = 1, \dots, k, \\
z_j &\leftarrow s_j - s_k + y_j - y_k & \forall j = 1, \dots, k-1, \\
r_j^2 &\leftarrow \frac{2}{L} \left[f(s_k) - f(s_j) + \nabla f(s_k)^T (s_j - s_k) \right] + \|y_j - y_k\|_2^2 & \forall j = 1, \dots, k-1, \\
M &\leftarrow 2 \begin{bmatrix} (z_2 - z_1)^T \\ \vdots \\ (z_{k-1} - z_1)^T \end{bmatrix}, \\
b &\leftarrow \begin{bmatrix} \left(\|z_2\|_2^2 - r_2^2 \right) - \|z_1\|_2^2 + r_1^2 \\ \vdots \\ \left(\|z_{k-1}\|_2^2 - r_{k-1}^2 \right) - \|z_1\|_2^2 + r_1^2 \end{bmatrix}.
\end{aligned}$$

2. Define the projection function

$$\mathcal{P}(w) \equiv \left(I - M (M^T M)^{-1} M^T \right) w + M (M^T M)^{-1} b.$$

We would implement the above projection with a QR-factorization instead of explicitly taking the matrix inverse, which can be numerically unstable (see, for example, Higham (2002, Chapter 14.1)).

3. Choose and fix any \bar{j} from $1, \dots, k-1$.

4. Compute

$$\begin{aligned}
c &\leftarrow \mathcal{P}(z_{\bar{j}}), \\
\rho &\leftarrow \sqrt{r_{\bar{j}}^2 - \|z_{\bar{j}} - c\|_2^2}, \\
x &\leftarrow c + \frac{\rho}{\|\mathcal{P}(y_k) - c\|_2} (\mathcal{P}(y_k) - c), \\
X &\leftarrow [x - z_1, \dots, x - z_{k-1}], \\
\lambda' &\leftarrow (X^T X)^{-1} X^T (y_k - x).
\end{aligned}$$

Similar to step 2, we would compute λ' using a QR-factorization for numerical stability.

5. Output

$$\begin{aligned}
s_k^+(\mathcal{S}) &\leftarrow x + s_k, \\
\lambda_j(\mathcal{S}) &\leftarrow \lambda'_j \quad \forall j = 1, \dots, k-1,
\end{aligned}$$

where λ'_j denotes the j -th element of the vector λ' .

The general procedure for solving (P_i^S) , for any $i = 1, \dots, k$, follows from re-indexing.

The variable x in steps 4 and 5 is directly related to Observation 35, which formulates the i -th Survey Descent subproblem as an inequality-constrained MBP after the variable-shift, $x \leftarrow x - s_i$. In step 4, we compute the solution, x , of the inequality-constrained MBP using the shifted-variables. Step 5 then reverses the variable shift to obtained the survey point update.