

A unifying framework for the analysis of projection-free first-order methods under a sufficient slope condition

Francesco Rinaldi* Damiano Zeffiro†

Abstract

The analysis of projection-free first order methods is often complicated by the presence of different kinds of "good" and "bad" steps. In this article, we propose a unifying framework for projection-free methods, aiming to simplify the converge analysis by getting rid of such a distinction between steps. The main tool employed in our framework is the Short Step Chain (SSC) procedure, which skips gradient computations in consecutive short steps until proper stopping conditions are satisfied. This technique allows us to give a unified analysis and converge rates in the general smooth non convex setting, as well as convergence rates under a Kurdyka-Łojasiewicz (KL) property, a setting that, to our knowledge, has not been analyzed before for the projection-free methods under study. In this context, we prove local convergence rates comparable to those of projected gradient methods under the same conditions.

Our analysis relies on a sufficient slope condition, ensuring that the directions selected by the methods have the steepest slope possible up to a constant among feasible directions. This condition is satisfied, among others, by several Frank-Wolfe (FW) variants on polytopes, and by some projection-free methods on convex sets with smooth boundary.

Keywords: Nonconvex optimization, First-order optimization, Projection-free optimization, Frank-Wolfe variants, Kurdyka-Łojasiewicz inequality.

AMS subject classification: 46N10, 65K05, 90C06, 90C25, 90C30

1 Introduction

Projection-free first-order methods aim to provide a valid alternative to projected gradient approaches for the constrained optimization of a smooth objective $f : \mathbb{R}^n \rightarrow \mathbb{R}$, in settings where projecting on the feasible set may be unpractical. Popular methods of this kind are the FW algorithm [17] and its variants (see, e.g., [18], [30] and references therein), which have many applications in sparse and structured optimization (see, e.g., [18], [22], [24], [39] and references therein).

In this paper, we present a specific algorithmic framework that can easily embed a number of projection-free (deterministic) methods proposed in the literature. This enables us to give a unified theoretical analysis for the methods at hand and new convergence results in the non convex case. More specifically, the framework we propose aims to overcome an annoying issue affecting the analysis of projection-free methods, that is the presence of "bad iterations", i.e., iterations where we cannot show good progress. This happens when we are forced to take a short step along the search direction to guarantee feasibility of the iterate. The number of short steps typically needs to be upper bounded in the convergence analysis with "ad hoc" arguments (see, e.g., [18] and [30]). The main idea behind our method is to chain several short steps by skipping gradient updates until proper stopping conditions are met. This has the additional benefit to rule out unpractical convergence rates due to large numbers of short steps in the analysis of the methods, like, for instance, in the PFW algorithm on the $N - 1$ dimensional simplex, where the best known upper bound on the number of consecutive short steps is

*Dipartimento di Matematica "Tullio Levi-Civita", Università di Padova, Italy (rinaldi@math.unipd.it)

†Dipartimento di Matematica "Tullio Levi-Civita", Università di Padova, Italy (zeffiro@math.unipd.it)

$3N!$ [30].

The projection-free first order methods we consider satisfy a sufficient slope condition related to the slope along descent directions. Essentially, for any direction d_x selected by the method in x and any direction d feasible in x we must have

$$-\langle \nabla f(x), \frac{d_x}{\|d_x\|} \rangle \geq -\tau \langle \nabla f(x), \frac{d}{\|d\|} \rangle$$

for a fixed $\tau > 0$ (independent from x and ∇f). Thus any direction selected by the method must be, up to a fixed constant, the steepest descent direction among feasible directions. The sufficient slope condition rules out the well known zig zagging behaviour of the classic FW method, whose directions can be almost orthogonal to the gradient on the interior of the feasible domain, leading to a slow convergence of $O(1/k)$ even for smooth strongly convex objectives [14, 17, 30, 43]. We show that this condition applies to the AFW, the PFW and the FW method with in face directions (FDFW) (see, e.g., [18] and [20]) on polytopes, to a retraction based approach on convex sets with smooth boundary as the one proposed in [33], to the FDFW method on sublevel sets of smooth strongly convex functions, and to tailored approaches over product domains.

We analyze convergence rates for non convex smooth objectives both in the general case and under the Kurdyka-Lojasiewicz property ([3], [8], [9]; see also Section 8.2 for a review of the essential definitions). This property holds for semialgebraic functions and more in general for functions definable in an \mathcal{o} -minimal class (see, e.g., [3], [8]). While there are many works on the convergence rates of proximal subgradient type methods under KL properties (see, e.g., [3], [4], [10], [42], [44]), we are not aware of previous analyses including projection-free methods for convex sets and smooth non convex objectives. In the convex setting, FW variants have been analyzed on polytopes [28] and uniformly convex sets (see, e.g., [29], [45]) under Hölderian error bound conditions, to interpolate between the well known $O(1/k)$ rate for general convex objectives and the faster rates for strongly convex ones proved in [19] and [30]. These Hölderian error bound conditions are equivalent to certain KL properties for convex objectives, as proved for a more general class of error bounds in [10]. As for the smooth non convex case, in [35] a KL property leading to linear convergence of some methods using retractions is proved for a quadratic problem on the Stiefel manifold; in [7] the convergence of a method using orthographic retractions is analyzed on proximally smooth surfaces under a KL property (see also Remark 8.2).

When dealing with the FW method and its variants, which represent a specific class of projection-free first-order methods, the SSC procedure has the additional advantage that it skips the computation of the solutions of linear programming subproblems. Different approaches to avoid computations of those solutions have been discussed in [12], [11] for convex optimization over polytopes. The blended conditional gradient introduced in [12] uses projections on the convex hull of the current active set as an alternative to pairwise or away FW steps, but its analysis still allows for short steps. Other works in projection-free optimization (e.g. [21], [32]) rely on the classic FW method to approximately compute the projection in accelerated projected gradient methods for convex objectives. To our knowledge this approach, called conditional gradient sliding, has not yet been combined effectively with faster FW variants, and it does not currently lead to any improvement on the $O(1/\varepsilon)$ linear oracle complexity of the FW method. Our methods using the SSC procedure instead have a complexity of $O(\ln \frac{1}{\varepsilon})$ both in gradients and linear oracles for non convex objectives satisfying a KL property implied e.g. by the Luo Tseng error bound [36] under some mild separability conditions for the stationary points [34, Theorem 4.1]. This error bound is known to hold in a variety of convex and non convex settings (see references in [34]).

The structure of the paper is as follows. In Section 2, we define some notation and state some preliminary results from convex analysis. In Section 3, we introduce a sufficient slope property for first-order projection free methods. We define the SSC procedure in Section 4, and prove it can be applied to generate some subgradient descent-like sequences (see, e.g., [4], [10], [36]) adapted to our particular setting. In Section 5, we state the main convergence results. In Section 6, we show examples of methods fitting our framework. Finally, some proofs, technical lemmas and background material can be found in the appendix.

2 Notation and preliminaries

We consider the following constrained optimization problem:

$$\min \{f(x) \mid x \in \Omega\}. \quad (2.1)$$

In the rest of the article Ω is a compact and convex set, unless specified otherwise, and $f \in C^1(\Omega)$ with L -Lipschitz gradient:

$$\|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\| \text{ for all } x, y \in \Omega.$$

We define D as the diameter of Ω , $\hat{c} = c/\|c\|$ for $c \in \mathbb{R}^n/\{0\}$ and $\hat{c} = 0$ for $c = 0$. For sequences we write $\{x_k\}$ instead of $\{x_k\}_{k \in I}$ when I is clear from the context, with $I = \mathbb{N}_0$ by default and $[1 : m] = \{1, \dots, m\}$. Given a function \tilde{f} with values in \mathbb{R} we use the notation

$$[\eta_1 < \tilde{f} < \eta_2] = \{x \in \mathbb{R}^n \mid \eta_1 < \tilde{f}(x) < \eta_2\}$$

for $-\infty \leq \eta_1 < \eta_2 \leq +\infty$ and we use the same notation for non strict inequalities.

For subsets C, D of \mathbb{R}^n we define $\text{dist}(C, D)$ as

$$\text{dist}(C, D) = \inf\{\|y - z\| \mid z \in C, y \in D\},$$

$B_R(C)$ as the neighborhood $\{x \in \mathbb{R}^n \mid \text{dist}(C, x) < R\}$ of C of radius R and in particular $B_R(x)$ as the open euclidean ball of radius R and center x . When C is closed and convex we define as $\pi(C, \cdot)$ the projection on C . If C is a cone then we denote with C^* its polar.

We now state some elementary properties related to the tangent and the normal cones, where for $\bar{x} \in \Omega$ we denote with $T_\Omega(\bar{x})$ and $N_\Omega(\bar{x})$ the tangent and the normal cone to Ω in \bar{x} respectively. The next proposition (from [41], Theorem 6.9) characterizes these cones for closed convex subsets of \mathbb{R}^n .

Proposition 2.1. *Let Ω be a closed convex set. For every point $\bar{x} \in \Omega$ we have*

$$\begin{aligned} T_\Omega(\bar{x}) &= \text{cl}\{w \mid \exists \lambda > 0 \text{ with } \bar{x} + \lambda w \in \Omega\}, \\ \text{int}(T_\Omega(\bar{x})) &= \{w \mid \exists \lambda > 0 \text{ with } \bar{x} + \lambda w \in \text{int}(\Omega)\}, \\ N_\Omega(\bar{x}) &= T_\Omega(\bar{x})^* = \{v \in \mathbb{R}^n \mid (v, y - \bar{x}) \leq 0 \ \forall y \in \Omega\}. \end{aligned}$$

We have the following formula connecting the maximal "slope" of a linear function along feasible directions to the tangent and the normal cone:

Proposition 2.2. *If Ω is a closed convex subset of \mathbb{R}^n , $\bar{x} \in \Omega$ then for every $g \in \mathbb{R}^n$*

$$\max \left\{ 0, \sup_{h \in \Omega \setminus \{\bar{x}\}} \left(g, \frac{h - \bar{x}}{\|h - \bar{x}\|} \right) \right\} = \text{dist}(N_\Omega(\bar{x}), g) = \|\pi(T_\Omega(\bar{x}), g)\|.$$

This property is a consequence of the Moreau-Yosida decomposition [41] and we refer the reader to Section 8.1 in the Appendix for a detailed proof. In the rest of the article to simplify notations we often use $\pi_{\bar{x}}(g)$ as a shorthand for $\|\pi(T_\Omega(\bar{x}), g)\|$, so that by Proposition 2.2 first order stationarity conditions in \bar{x} for the gradient $-g$ become equivalent to $\pi_{\bar{x}}(g) = 0$.

3 Sufficient slope condition

Let \mathcal{A} be a first-order optimization method defined for smooth functions on a closed subset Ω of \mathbb{R}^n . We assume that given first-order information $(x_k, \nabla f(x_k))$ the method always selects x_{k+1} along a feasible descent direction, so that for $(x, g) \in \Omega \times \mathbb{R}^n$ we can define

$$\mathcal{A}(x, g) \subset T_\Omega(x) \cap \{y \in \mathbb{R}^n \mid \langle g, y \rangle > 0\} \cup \{0\}$$

as the possible descent directions selected by \mathcal{A} when $x = x_k$, $g = -\nabla f(x_k)$ for some k (see Algorithm 1). When x is first-order stationary, we set $\mathcal{A}(x, g) = \{0\}$, otherwise we always assume $0 \notin \mathcal{A}(x, g)$.

Algorithm 1: First-order method

Initialization. $x_0 \in \Omega$, $k := 0$.

1. If x_k is stationary, then STOP
 2. select a descent direction $d_k \in \mathcal{A}(x_k, -\nabla f(x_k))$
 3. set $x_{k+1} = x_k + \alpha_k d_k$ for some stepsize $\alpha_k \in [0, \alpha_k^{\max}]$
 4. set $k := k + 1$, go to Step 1.
-

We want to formulate a "sufficient slope" condition for the descent directions selected by \mathcal{A} with respect to the steepest slope among feasible directions. In order to do that, we define the directional slope lower bound as

$$\text{DSB}_{\mathcal{A}}(\Omega, x, g) = \inf_{d \in \mathcal{A}(x, g)} \frac{\langle g, d \rangle}{\pi_x(g) \|d\|}$$

if $0 \notin \mathcal{A}(x, g)$. Otherwise by our assumptions x is stationary for $-g$, $\pi_x(g) = 0$ and we set $\text{DSB}_{\mathcal{A}}(\Omega, x, g) = 1$. With these definitions for any $x \in \Omega$, $g \in \mathbb{R}^n$

$$\text{DSB}_{\mathcal{A}}(\Omega, x, g) \leq 1 \tag{3.1}$$

since if $\pi_x(g) \neq 0$ then for any d feasible descent direction $\frac{\langle g, d \rangle}{\|d\|} \leq \pi_x(g)$ by Proposition 2.2, and if $\pi_x(g) = 0$ by definition $\text{DSB}_{\mathcal{A}}(\Omega, x, g) = 1$. Given a subset P of Ω we can finally define the slope lower bound

$$\text{SB}_{\mathcal{A}}(\Omega, P) = \inf_{\substack{g \in \mathbb{R}^n \\ x \in P}} \text{DSB}_{\mathcal{A}}(\Omega, x, g) = \inf_{\substack{g: \pi_x(g) \neq 0 \\ x \in P}} \text{DSB}_{\mathcal{A}}(\Omega, x, g),$$

where in the second equality we used that (3.1) holds with equality if $\pi_x(g) = 0$. For simplicity if $P = \Omega$ we write $\text{SB}_{\mathcal{A}}(\Omega)$ instead of $\text{SB}_{\mathcal{A}}(\Omega, \Omega)$.

In Section 6, we show a few examples of methods and sets satisfying a sufficient slope condition, i.e., cases where the slope lower bound is strictly greater than zero.

4 First order projection free methods with SSC procedure

In this section, we briefly recall the subgradient descent sequence conditions used in the analysis of a variety of descent methods (see, e.g., [4], [10], [36]), with a few simple adaptations to our smooth constrained projection-free setting. We then introduce the SSC procedure, and show how it allows us to obtain our adapted subgradient descent sequences.

4.1 Subgradient descent sequences for projection-free methods

Let $\tilde{f}: \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ be a proper lower semicontinuous function and let $\partial \tilde{f}$ be its limiting subdifferential. We refer the reader to Section 8.2 of the appendix for a brief review of the Fréchet subdifferential and the related KL property as defined in [3], [4]. In this work we are only interested in the case $\tilde{f} = f + i_{\Omega}$ with i_{Ω} the indicator function for Ω . We anyway notice that the abstract convergence result we give in Section 5 holds for general lower semicontinuous functions.

Subgradient descent sequences as defined in [10] satisfy the following properties for some given $a, b > 0$:

$$\tilde{f}(x_{k+1}) + a \|x_{k+1} - x_k\|^2 \leq \tilde{f}(x_k), \tag{H1}$$

$$\exists w_{k+1} \in \partial \tilde{f}(x_{k+1}) \text{ such that } \|w_{k+1}\| \leq b \|x_{k+1} - x_k\|. \tag{H2}$$


Since $\partial \tilde{f}(x_k)$ is closed, condition (H2) is true iff $\partial \tilde{f}(x_{k+1}) \neq \emptyset$ and

$$\text{dist}(\partial \tilde{f}(x_{k+1}), 0) \leq b \|x_{k+1} - x_k\|. \tag{4.2}$$

In our smooth constrained setting, that is when $\tilde{f} = f + i_\Omega$, by (8.8) we can rewrite (4.2) as

$$\|\pi(T_\Omega(x_{k+1}), -\nabla f(x_{k+1}))\| \leq b\|x_{k+1} - x_k\|. \quad (4.3)$$

However, for projection-free methods equation (4.2) and in particular (4.3) do not hold in general, even in linear convergence settings.



(a) FW method (6.3) does a full step from the boundary of a sphere. $s_k \in \operatorname{argmin}_{s \in \Omega} \langle \nabla f(x), s \rangle$.

(b) AFW method (6.2) does a short step near the boundary of a polytope.

Figure 1

Intuitively, for these methods the first obstacle for condition (4.3) or variants is that if $x_k \in \partial\Omega$ and $x_{k+1} \in \operatorname{int}(\Omega)$ with $\|\pi(T_\Omega(x_k) - \nabla f(x_k))\| \ll \|\nabla f(x_k)\|$ we have, for typical choices of the stepsize (see figure 1a):

$$\|x_{k+1} - x_k\| \approx \|\pi(T_\Omega(x_k), -\nabla f(x_k))\| \ll \|\nabla f(x_k)\| \approx \|\nabla f(x_{k+1})\| = \|\pi(T_\Omega(x_{k+1}), -\nabla f(x_{k+1}))\|.$$

This leads us to introduce the weaker condition

$$\tilde{f}(x_k) - \tilde{f}(x_{k+1}) \geq \frac{a}{b^2} \operatorname{dist}(\partial\tilde{f}(\tilde{x}_k), 0)^2, \quad (\text{H}'2)$$

for some auxiliary point \tilde{x}_k s.t. $\partial\tilde{f}(\tilde{x}_k) \neq \emptyset$ and

$$\tilde{f}(x_{k+1}) \leq \tilde{f}(\tilde{x}_k) \leq \tilde{f}(x_k) - a\|\tilde{x}_k - x_k\|^2. \quad (4.4)$$

In particular in the setting of figure 1a a typical choice is $\tilde{x}_k = x_k$, as we shall see more formally in the next section. In the rest of this article (H'2) always comes with (4.4) for \tilde{x}_k .

As anticipated, condition (H'2) is weaker than (H2) if one is already assuming (H1). Indeed if (H1) and (H2) hold then

$$\tilde{f}(x_k) - \tilde{f}(x_{k+1}) \geq a\|x_k - x_{k+1}\|^2 \geq \frac{a}{b^2} \operatorname{dist}(\partial\tilde{f}(x_{k+1}), 0)^2,$$

where the second inequality directly follows from (H2) in the form (4.2). One can then satisfy (H'2) and (4.4) by taking $\tilde{x}_k = x_{k+1}$.

In our setting, that is when $\tilde{f} = f + i_\Omega$, we use the smooth version of (H'2):

$$f(x_k) - f(x_{k+1}) \geq \frac{a}{b^2} \|\pi(T_\Omega(\tilde{x}_k), -\nabla f(\tilde{x}_k))\|^2, \quad (\text{H}'2a)$$

with analogous conditions on \tilde{x}_k .

4.2 The SSC procedure

The second (and arguably most important) issue for applying conditions such as (4.3) to projection-free methods is the presence of short steps. Given a projection-free first-order method \mathcal{A} , we say that the step k is short if the stepsize α_k is maximal, that is $\alpha_k = \alpha_k^{\max}$. Roughly speaking, when x_k is close enough to $\partial\Omega$ and d_k points toward $\partial\Omega$ we have (see figure 1b with $p_{k+1} = \pi(T_\Omega(x_{k+1}), -\nabla f(x_{k+1}))$)

$$\|x_{k+1} - x_k\| = \alpha_k \|d_k\| = \alpha_k^{\max} \|d_k\| \ll \|\pi(T_\Omega(x_{k+1}), -\nabla f(x_{k+1}))\|.$$

An analogous issue affects also weaker variants of (4.3), including (H'2a). In order to address this issue, we apply the SSC procedure to first order methods (Algorithm 2). This procedure (Algorithm 3) chains consecutive short steps, skipping updates for the gradient and possibly for related information, like linear minimizers, until proper stopping conditions are met. After the SSC returns a candidate point x_{k+1}^{tr} , the method may still choose another point in Step 4, selecting a x_{k+1} with a smaller objective value than x_{k+1}^{tr} and satisfying (H1):

$$f(x_{k+1}) \leq \min(f(x_{k+1}^{tr}), f(x_k) - \frac{L}{2} \|x_{k+1} - x_k\|^2). \quad (4.5)$$

The method can always take x_{k+1}^{tr} , since, as we prove later in this section, the following inequality holds

$$f(x_{k+1}^{tr}) \leq f(x_k) - \frac{L}{2} \|x_{k+1}^{tr} - x_k\|^2.$$

Algorithm 2: First-order method with SSC

Initialization. $x_0 \in \Omega$, $k = 0$.

1. **while** x_k is not stationary:
 2. $g = -\nabla f(x_k)$
 3. $x_{k+1}^{tr} = \text{SSC}(x_k, g)$
 4. select $x_{k+1} \in \Omega$ satisfying (4.5)
 5. $k = k + 1$.
-

Algorithm 3: SSC(\bar{x}, g)

Initialization. $y_0 = \bar{x}$, $j = 0$.

Phase I

1. select $d_j \in \mathcal{A}(y_j, g)$, $\alpha_{\max}^{(j)} \in \alpha_{\max}(y_j, d_j)$
2. **if** $d_j = 0$ **then:**
3. **return** y_j

Phase II

4. compute β_j with (4.7)
5. let $\alpha_j = \min(\alpha_{\max}^{(j)}, \beta_j)$
6. **if** $\alpha_j = 0$ **then:**
7. **return** y_j

Phase III

8. $y_{j+1} = y_j + \alpha_j d_j$
 9. **if** $\alpha_j = \beta_j$ **then:**
 10. **return** y_{j+1}
 11. $j = j + 1$, go to Step 1.
-

Given that the gradient $-g$ is constant during the SSC, this procedure can be viewed as an application of \mathcal{A} for the minimization of the linearized objective $f_g(z) = \langle -g, z - \bar{x} \rangle + f(\bar{x})$ with peculiar stepsizes and stopping criterion. Fundamentally, these must ensure that the linearized objective is still a sufficiently accurate approximation of f on the generated sequence, and that the directions selected are descent directions not only for f_g but also for the true objective f .

We can split the SSC into three phases. First, the SSC computes a feasible descent direction $d_j \in \mathcal{A}(y_j, g)$ if it exists, otherwise $d_j = 0$ and the algorithm returns y_j . Notice that $d_j = 0$ iff y_j is stationary w.r.t. the approximated gradient $-g$. Second, the stepsize α_j is computed by taking the minimum between the maximal stepsize $\alpha_{\max}^{(j)}$ and an auxiliary stepsize β_j . The maximal stepsize $\alpha_{\max}^{(j)}$ may depend on boundary conditions but also on additional information available like the active set in the AFW and in the PFW. In the theoretical analysis, we only need to assume termination of the procedure and $\alpha_{\max}^{(j)} > 0$. In particular, $\alpha_j = 0$ in Phase II iff $\beta_j = 0$, or equivalently if d_j is not feasible from y_j for the auxiliary set Ω_j (see (4.6)). As we shall see later in this section, when $d_0 \neq 0$ we always have $\beta_0 > 0$, and therefore Phase II can never return y_0 (which could cause an infinite loop in Algorithm 2). In the third and final phase, the new point of the SSC $y_{j+1} = y_j + \alpha_j d_j$ is computed. Then the procedure terminates if $\alpha_j = \beta_j$, otherwise it continues without updating the gradient.

We now define the auxiliary step size β_j , which allows us to retrieve the descent sequence conditions (H1) and (H'2a). Let

$$\Omega_j = \bar{B}_{\|g\|/2L}(\bar{x} + \frac{g}{2L}) \cap \bar{B}_{\langle g, \hat{d}_j \rangle / L}(\bar{x}). \quad (4.6)$$

We define

$$\beta_j = \begin{cases} 0 & \text{if } y_j \notin \Omega_j, \\ \beta_{\max}(\Omega_j, y_j, d_j) & \text{if } y_j \in \Omega_j. \end{cases} \quad (4.7)$$

where $\beta_{\max}(\Omega_j, y_j, d_j) = \max\{\beta \in \mathbb{R}_{\geq 0} \mid y_j + \beta d_j \in \Omega_j\}$ is the maximal feasible stepsize in the direction d_j starting from y_j with respect to Ω_j . Since Ω_j is the intersection of two balls there is a simple closed form expression for β_j . Given that $y_0 = \bar{x}$ it is immediate to check that if $d_0 \neq 0$

$$\beta_0 = \frac{\langle g, \hat{d}_0 \rangle}{L \|d_0\|},$$

which corresponds to the constant stepsize $1/L$ for gradient descent (when $d_0 = g$), and where $\beta_0 > 0$ since $d_0 \neq 0$ is by assumption a descent direction for $-g$.

The second ball appearing in the definition of Ω_j does not depend on j and can be rewritten as

$$\bar{B}_{\|g\|/2L}(\bar{x} + \frac{g}{2L}) = \{z \in \mathbb{R}^n \mid L \|z - \bar{x}\|^2 - \langle z - \bar{x}, g \rangle \leq 0\}. \quad (4.8)$$

Then for any $z \in \bar{B}_{\|g\|/2L}(\bar{x} + \frac{g}{2L})$ we have

$$f(z) \leq f(\bar{x}) - \langle g, z - \bar{x} \rangle + \frac{L}{2} \|z - \bar{x}\|^2 \leq f(\bar{x}) - \frac{L}{2} \|\bar{x} - z\|^2, \quad (4.9)$$

where the first inequality is the standard descent lemma and the second follows from (4.8). Let now $\bar{B} = \bar{B}_{\|g\|/2L}(\bar{x} + \frac{g}{2L})$. Since $y_0 = x_k \in \bar{B}$, and if $y_j \in \bar{B}$ then $y_{j+1} \in \bar{B}$ by definition of $\beta_j \geq \alpha_j$, by induction we obtain $y_j \in \bar{B}$ for every $j \in \{0, \dots, T\}$. Therefore thanks to (4.9) for every $j \in \{0, \dots, T\}$ we have the sufficient decrease condition

$$f(y_j) \leq f(x_k) - \frac{L}{2} \|x_k - y_j\|^2. \quad (4.10)$$

We now show how the condition $y_{j+1} \in \Omega_j \subset \bar{B}_{\langle g, \hat{d}_j \rangle / L}(\bar{x})$ guarantees that the true objective f is monotone decreasing during the SSC.

Lemma 4.1. Assume $y_j \in \bar{B}_{\langle g, \hat{d}_j \rangle / L}(\bar{x})$. Then for every $\beta \in [0, \beta_j]$ we have

$$\frac{d}{d\beta} f(y_j + \beta d_j) \leq 0,$$

and thus in particular $f(y_j + \beta_j d_j) \leq f(y_j)$.

Proof. We have

$$\begin{aligned} \frac{d}{d\beta} f(y_j + \beta d_j) &= \|d_j\| \langle \nabla f(y_j + \beta d_j), \hat{d}_j \rangle \\ &= \|d_j\| \langle (\nabla f(y_j + \beta d_j) + g) - g, \hat{d}_j \rangle = \|d_j\| (\langle \nabla f(y_j + \beta d_j) + r, \hat{d}_j \rangle - \langle g, \hat{d}_j \rangle) \\ &\leq \|d_j\| (L \|\bar{x} - y_j - \beta d_j\| - \langle g, \hat{d}_j \rangle) \leq 0, \end{aligned}$$

where we used $g = -\nabla f(\bar{x})$ and the Lipschitzianity of ∇f in the first inequality and

$$y_j + \beta d_j \in \bar{B}_{\langle g, \hat{d}_j \rangle / L}(\bar{x})$$

in the second. □

We can now prove the descent sequence conditions (H1) and (H'2a) for a sequence $\{x_k\}$ generated by Algorithm 2.

Proposition 4.1. Let $\tau = SB_{\mathcal{A}}(\Omega) > 0$. Assume the SSC is given in input (x_k, g) and let x_{k+1}^{tr} be the output. Then

$$f(x_k) - f(x_{k+1}^{tr}) \geq \frac{L}{2} \|x_k - x_{k+1}^{tr}\|^2, \quad (4.11)$$

$$\|x_k - x_{k+1}^{tr}\| \geq K \|\pi(T_{\Omega}(\tilde{x}_k), -\nabla f(\tilde{x}_k))\|. \quad (4.12)$$

for some \tilde{x}_k such that $f(x_{k+1}^{tr}) \leq f(\tilde{x}_k) \leq f(x_k) - \frac{L}{2} \|x_k - \tilde{x}_k\|^2$ and for $K = \tau / (L(1 + \tau))$. Furthermore, conditions (H1) and (H'2a) are satisfied with $b = K$ and $a = L/2$.

Proof. Let $B_j = \bar{B}_{\langle g, \hat{d}_j \rangle / L}(x_k)$ and let T be such that $x_{k+1}^{tr} = y_T$.

Inequality (4.10) applied with $j = T$ gives (4.11). Moreover, by taking $\tilde{x}_k = y_{\tilde{T}}$ for some $\tilde{T} \in \{0, \dots, T\}$ the conditions

$$f(x_{k+1}^{tr}) \leq f(\tilde{x}_k) \leq f(x_k) - \frac{L}{2} \|x_k - \tilde{x}_k\|^2 \quad (4.13)$$

are satisfied by Lemma 4.1 and (4.10).

Let now $p_j = \|\pi(T_{\Omega}(y_j), -\nabla f(y_j))\|$ and $\tilde{p}_j = \|\pi(T_{\Omega}(y_j), g)\| = \|\pi(T_{\Omega}(y_j), -\nabla f(x_k))\|$. We have

$$\begin{aligned} |p_j - \tilde{p}_j| &= \left| \|\pi(T_{\Omega}(y_j), -\nabla f(y_j))\| - \|\pi(T_{\Omega}(y_j), -\nabla f(x_k))\| \right| \\ &\leq \|-\nabla f(y_j) + \nabla f(x_k)\| \leq L \|y_j - x_k\|, \end{aligned} \quad (4.14)$$

where we used the 1-Lipschitzianity of the projections on convex sets in the first inequality. We now distinguish four cases according to how the SSC terminates.

Case 1: $T = 0$ or $d_T = 0$. Since there are no descent directions $x_{k+1}^{tr} = y_T$ must be stationary for the gradient g . Equivalently, $\tilde{p}_T = \|\pi(T_{\Omega}(x_{k+1}^{tr}), g)\| = 0$. We can now write

$$\|x_{k+1}^{tr} - x_k\| \geq \frac{1}{L} (|p_T - \tilde{p}_T|) = \frac{p_T}{L} > K p_T,$$

where we used (4.14) in the first inequality and $\tilde{p}_T = 0$ in the equality. We now prove that if $T = 0$ then necessarily $d_0 = 0$. This is clear if y_0 is stationary for $-g$. Otherwise $d_0 \neq 0$ is a feasible descent direction, so that $\alpha_{\max}^{(0)} > 0$ and also $\beta_0 > 0$. But then $\alpha_0 > 0$ and the SSC can't terminate in Phase I

or Phase II, in contradiction with $T = 0$.

Before examining the remaining cases we remark that if the SSC terminates in Phase III then $\alpha_{T-1} = \beta_{T-1}$ must be maximal w.r.t. the conditions $y_T \in B_{T-1}$ or $y_T \in \bar{B}$, which imply $y_T \in \partial B_{T-1}$ (case 3) and $y_T \in \partial \bar{B}$ (case 4) respectively. On the other hand if $\alpha_T = 0$ then necessarily $y_T \in \text{int}(\Omega_T)^c$. In this case we cannot have $y_T \in \partial \bar{B}$, otherwise the SSC would terminate in Phase III of the previous cycle. Therefore necessarily $y_T \in \text{int}(B_T)^c$ (Case 2).

Case 2: $y_T \in \text{int}(B_T)^c$. We can rewrite the condition as

$$\langle g, \hat{d}_T \rangle \leq L \|y_T - x_k\|. \quad (4.15)$$

Thus

$$p_T \leq \tilde{p}_T + L \|y_T - x_k\| \leq \frac{1}{\tau} \langle g, \hat{d}_T \rangle + L \|y_T - x_k\| \leq \left(\frac{L}{\tau} + L \right) \|y_T - x_k\|, \quad (4.16)$$

where the first inequality follows from (4.14), the second from $\frac{\langle g, \hat{d}_T \rangle}{\tilde{p}_T} \geq \text{DSB}_{\mathcal{A}}(\Omega, y_T, g) \geq \text{SB}_{\mathcal{A}}(\Omega) = \tau$, and the third from (4.15). Then $\tilde{x}_k = x_{k+1}^{tr} = y_T$ satisfies the desired conditions.

Case 3: $y_T = y_{T-1} + \beta_{T-1} d_{T-1}$ and $y_T \in \partial B_{T-1}$. Then from $y_{T-1} \in B_{T-1}$ it follows

$$L \|y_{T-1} - x_k\| \leq \langle g, \hat{d}_{T-1} \rangle, \quad (4.17)$$

and $y_T \in \partial B_{T-1}$ implies

$$\langle g, \hat{d}_{T-1} \rangle = L \|y_T - x_k\|. \quad (4.18)$$

Combining (4.17) with (4.18) we obtain

$$L \|y_{T-1} - x_k\| \leq L \|y_T - x_k\|. \quad (4.19)$$

Thus

$$p_{T-1} \leq \tilde{p}_{T-1} + L \|y_{T-1} - x_k\| \leq \frac{1}{\tau} \langle g, \hat{d}_{T-1} \rangle + L \|y_{T-1} - x_k\| \leq \left(\frac{L}{\tau} + L \right) \|y_T - x_k\|,$$

where we used (4.18), (4.19) in the last inequality and the rest follows reasoning as for (4.16). In particular we can take $\tilde{x}_k = y_{T-1}$.

Case 4: $y_T = y_{T-1} + \beta_{T-1} d_{T-1}$ and $y_T \in \partial \bar{B}$.

By the characterization of \bar{B} in (4.8) the condition $x_{k+1}^{tr} = y_T \in \bar{B}$ can be rewritten as

$$L \|x_{k+1}^{tr} - x_k\|^2 - \langle g, x_{k+1}^{tr} - x_k \rangle = 0. \quad (4.20)$$

For every $j \in \{0, \dots, T\}$ we have

$$x_{k+1}^{tr} = y_j + \sum_{i=j}^{T-1} \alpha_i d_i. \quad (4.21)$$

We now want to prove that for every $j \in \{0, \dots, T\}$

$$\|x_{k+1}^{tr} - x_k\| \geq \|y_j - x_k\|. \quad (4.22)$$

Indeed, we have

$$\begin{aligned} L \|x_{k+1}^{tr} - x_k\|^2 &= \langle g, x_{k+1}^{tr} - x_k \rangle = \langle g, y_j - x_k \rangle + \sum_{i=j}^{T-1} \alpha_i \langle g, d_i \rangle \\ &\geq \langle g, y_j - x_k \rangle \geq L \|y_j - x_k\|^2, \end{aligned}$$

where we used (4.20) in the first equality, (4.21) in the second, $\langle g, d_j \rangle \geq 0$ for every j in the first inequality and $y_j \in \bar{B}$ in the second inequality.

We also have

$$\begin{aligned} \frac{\langle g, x_{k+1}^{tr} - x_k \rangle}{\|x_{k+1}^{tr} - x_k\|} &= \frac{\langle g, \sum_{j=0}^{T-1} \alpha_j d_j \rangle}{\|\sum_{j=0}^{T-1} \alpha_j d_j\|} \geq \frac{\langle g, \sum_{j=0}^{T-1} \alpha_j d_j \rangle}{\sum_{j=0}^{T-1} \alpha_j \|d_j\|} \\ &\geq \min \left\{ \frac{\langle g, d_j \rangle}{\|d_j\|} \mid 0 \leq j \leq T-1 \right\}. \end{aligned} \quad (4.23)$$

Thus for $\tilde{T} \in \operatorname{argmin} \left\{ \frac{\langle g, d_j \rangle}{\|d_j\|} \mid 0 \leq j \leq T-1 \right\}$

$$\langle g, \hat{d}_{\tilde{T}} \rangle \leq \frac{\langle g, x_{k+1}^{tr} - x_k \rangle}{\|x_{k+1}^{tr} - x_k\|} = L \|x_{k+1}^{tr} - x_k\|, \quad (4.24)$$

where we used (4.23) in the first inequality and (4.20) in the second. We finally have

$$p_{\tilde{T}} \leq \tilde{p}_{\tilde{T}} + L \|y_{\tilde{T}} - x_k\| \leq \frac{1}{\tau} \langle g, \hat{d}_{\tilde{T}} \rangle + L \|y_{\tilde{T}} - x_k\| \leq \left(\frac{L}{\tau} + L \right) \|x_{k+1}^{tr} - x_k\|,$$

where we used (4.22), (4.24) in the last inequality and the rest follows reasoning as for (4.16). In particular $\tilde{x}_k = y_{\tilde{T}}$ satisfies the desired properties. It remains to prove that x_{k+1} satisfies (H1) and (H'2a). The first condition (H1) clearly holds by (4.5). As for (H'2a), we have

$$f(x_k) - f(x_{k+1}) \geq f(x_k) - f(x_{k+1}^{tr}) \geq \frac{L}{2} \|x_k - x_{k+1}^{tr}\|^2 \geq \frac{2K^2}{L} \|\pi(T_{\Omega}(\tilde{x}_k), -\nabla f(\tilde{x}_k))\|^2,$$

where the first inequality follows from (4.5), the second from (4.11), and the third from (4.12). Finally, since $f(x_{k+1}) \leq f(x_{k+1}^{tr})$, by (4.13) the conditions (4.4) on \tilde{x}_k are satisfied. \square

As an immediate corollary we have convergence to the set of stationary points with a rate of $O(\frac{1}{\sqrt{k}})$ for the norm of the projected gradient $\pi(T_{\Omega}(\tilde{x}_i), -\nabla f(\tilde{x}_i))$. While the norm of the projected gradient may not be available, by (4.11) we have the readily available upper bound $\|x_k - x_{k+1}^{tr}\|/K$, which has an analogous convergence rate.

Corollary 4.1. *Under the assumptions of Proposition 4.1, $\{f(x_k)\}$ is decreasing, $f(x_k) \rightarrow f^* \in \mathbb{R}$ and the limit points of $\{x_k\}$ are stationary. Furthermore, for any sequence $\{\tilde{x}_k\}$ satisfying (4.12),*

$$\min_{0 \leq i \leq k} \|\pi(T_{\Omega}(\tilde{x}_i), -\nabla f(\tilde{x}_i))\| \leq \min_{0 \leq i \leq k} \frac{\|x_{i+1}^{tr} - x_i\|}{K} \leq \sqrt{\frac{2(f(x_0) - \tilde{f})}{K^2 L(k+1)}}. \quad (4.25)$$

Proof. The sequence $\{f(x_k)\}$ is decreasing by (4.5) and (4.11). Thus by compactness $f(x_k) \rightarrow \tilde{f} \in \mathbb{R}$ and in particular $f(x_k) - f(x_{k+1}) \rightarrow 0$. Since $f(x_{k+1}^{tr}) - f(x_k) \leq f(x_k) - f(x_{k+1})$ we also have $f(x_{k+1}^{tr}) - f(x_k) \rightarrow 0$ and by (4.11) $x_{k+1}^{tr} - x_k \rightarrow 0$. Let $\{x_{k(i)}\} \rightarrow \tilde{x}^*$ be any convergent subsequence of $\{x_k\}$. For $\{\tilde{x}_k\}$ chosen as in the proof of Proposition 4.1 we have $\|\tilde{x}_k - x_k\| \leq \|x_k^{tr} - x_k\|$ because $\tilde{x}_k = y_T = x_k^{tr}$ in case 1 and case 2, by (4.19) in case 3, and by (4.22) in case 4. Therefore

$$\|\tilde{x}_{k(i)} - x_{k(i)}\| \leq \|x_{k(i)} - x_{k(i)+1}^{tr}\| \rightarrow 0.$$

Furthermore, $\|\pi(T_{\Omega}(\tilde{x}_{k(i)}), -\nabla f(\tilde{x}_{k(i)}))\| \leq \frac{\|x_{k(i)} - x_{k(i)+1}^{tr}\|}{K} \rightarrow 0$ again by Proposition 4.1, so that $\tilde{x}_{k(i)} \rightarrow \tilde{x}^*$ with $\|\pi(T_{\Omega}(\tilde{x}_{k(i)}), -\nabla f(\tilde{x}_{k(i)}))\| \rightarrow 0$. Then $\|\pi(T_{\Omega}(\tilde{x}^*), -\nabla f(\tilde{x}^*))\| = 0$ and \tilde{x}^* is stationary.

The first inequality in (4.25) follows directly from (4.12). As for the second, we have

$$\begin{aligned} & \frac{k+1}{K^2} \left(\min_{0 \leq i \leq k} \|x_{i+1}^{tr} - x_i\| \right)^2 = \frac{k+1}{K^2} \min_{0 \leq i \leq k} \|x_{i+1}^{tr} - x_i\|^2 \\ & \leq \frac{1}{K^2} \sum_{i=0}^k \|x_i - x_{i+1}^{tr}\|^2 \leq \frac{2}{LK^2} \sum_{i=0}^k (f(x_{i+1}) - f(x_i)) \leq \frac{2(f(x_0) - \tilde{f})}{LK^2}, \end{aligned}$$

where we used (4.11) together with (4.5) in the second inequality, $f(x_i) \rightarrow \tilde{f}$ decreasing in the second and (4.25) follows by rearranging. \square

Remark 4.1. The proof of Proposition gives an easy way to retrieve \tilde{x}_k . In particular it shows that one can take $\tilde{x}_k = y_{\tilde{T}}$ with $\tilde{T} = T$ in case 1 and case 2, $\tilde{T} = T - 1$ in case 3 and

$$\tilde{T} \in \operatorname{argmin}\{\langle g, \hat{d}_i \rangle \mid i \in [0 : T]\}$$

in case 4.

Remark 4.2. For any $d \in \mathcal{A}(\tilde{x}_k, -\nabla f(\tilde{x}_k))$ we have

$$\langle -\nabla f(\tilde{x}_k), \hat{d} \rangle \leq \|\pi(T_\Omega(\tilde{x}_i), -\nabla f(\tilde{x}_i))\| \leq \frac{\langle -\nabla f(\tilde{x}_k), \hat{d} \rangle}{\tau},$$

where the first inequality follows from Proposition 2.2 and the second by the assumption on $\text{SB}_{\mathcal{A}}(\Omega)$. Therefore at the price of computing $-\nabla f(\tilde{x}_k)$ and a related descent directions d we have upper and lower bounds separated by a factor of τ for $\|\pi(T_\Omega(\tilde{x}_i), -\nabla f(\tilde{x}_i))\|$.

5 Convergence properties

In this section, we use the descent sequence properties proved in the previous section in combination with a KL property for smooth functions to obtain convergence rates with respect to the objectives and the tail length of the iterate sequence. First, we give an abstract convergence lemma for proper lower semicontinuous functions. We then apply this lemma in combination with Proposition 4.1 to prove convergence rates for Algorithm 2 under finite termination of the SSC, and thus in particular for the examples discussed in Section 6.

For the abstract lemma with respect to the foundational work [4] we only have the weaker condition (H'2) instead of (H2), where the proof of the convergence lemma [4, Lemma 2.6] relies on the stronger condition. We therefore present a different proof technique in Section 8.3 of the appendix based on Karamata's inequality ([25], [27]).

Before stating the main convergence result we need to introduce a one dimensional worst case sequence related to the desingularizing function φ , which we use bound the error on the objective. With respect to the (different) worst case sequence introduced in [10] we do not require neither the moderate behavior hypothesis on φ nor $(\varphi')^{-1}$ to be Lipschitz regular.

Let $\alpha : [0, \eta) \rightarrow \mathbb{R}_{>0} \cup \{\infty\}$ be a continuous monotone non increasing function finite in $(0, \eta)$. For $t \in [0, \eta)$ we define

$$\sigma_\alpha(t) = \max \left\{ s \in [0, t] \mid \frac{1}{\alpha(s)^2} \leq t - s \right\},$$

with the convention $\max(\emptyset) = 0$. Since $s + 1/\alpha(s)^2$ is continuous and strictly increasing if

$$\left\{ s \in [0, t] \mid \frac{1}{\alpha(s)^2} \leq t - s \right\}$$

is non empty then it has a maximum and

$$t - \sigma_\alpha(t) = \frac{1}{\alpha(\sigma_\alpha(t))^2}. \quad (5.1)$$

Otherwise

$$t - \sigma_\alpha(t) = t \leq \frac{1}{\alpha(\sigma_\alpha(t))^2} = \frac{1}{\alpha(0)^2}. \quad (5.2)$$

Let $\sigma_\alpha^{(k)}$ be equal to σ_α applied k times, with the convention $\sigma_\alpha^{(0)}(t) = t$. Then the worst sequence related to an initial error $f_0 \in [0, \eta)$ is $\{\sigma_\alpha^{(k)}(f_0)\}$. Given that by continuity α is upper bounded in every interval $[\varepsilon, \eta)$ with $\varepsilon > 0$, we have the limit

$$\lim_{k \rightarrow \infty} \sigma_\alpha^{(k)}(f_0) \rightarrow 0 \quad (5.3)$$

for every $f_0 \in [0, \eta)$.

We can now state the main convergence result for proper lower semicontinuous functions.

Proposition 5.1. *Let $\tilde{f}: \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$ be a proper l.s.c. function satisfying the KL inequality w.r.t. x^* in $B_\delta(x^*) \cap [\tilde{f}(x^*) < \tilde{f} < \eta]$. Let then $\{x_k\}$ be a sequence satisfying (H1) and (H'2) such that*

$$\tilde{f}(x^*) \leq \tilde{f}(x_0) < \eta, \quad (5.4)$$

$$x_k \in B_\delta(x^*) \Rightarrow \tilde{f}(x_{k+1}) \geq \tilde{f}(x^*). \quad (5.5)$$

Let $\alpha(t) = \frac{b}{\sqrt{a}}\varphi'(t)$. Assume also that

$$\frac{b}{a}\varphi(\tilde{f}(x_0) - \tilde{f}(x^*)) + 2\sqrt{\frac{\tilde{f}(x_0) - \tilde{f}(x^*) - \sigma_\alpha(\tilde{f}(x_0) - \tilde{f}(x^*))}{a}} + \|x_0 - x^*\| < \delta. \quad (5.6)$$

Then $B_\delta(x^*) \supset \{x_k\} \rightarrow \tilde{x}^*$ with $f(\tilde{x}^*) \leq f(x^*)$, and

$$\tilde{f}(x_k) - \tilde{f}(x^*) \leq \sigma_\alpha^{(k)}(\tilde{f}(x_0) - \tilde{f}(x^*)), \quad (5.7)$$

$$\|x_k - \tilde{x}^*\| \leq \sum_{i=k}^{\infty} \|x_i - x_{i+1}\| \leq \frac{b}{a}\varphi(\tilde{f}(x_k) - \tilde{f}(x^*)) + 2\sqrt{\frac{\tilde{f}(x_k) - \tilde{f}(x^*) - \sigma_\alpha(\tilde{f}(x_k) - \tilde{f}(x^*))}{a}}. \quad (5.8)$$

Condition (5.6) is a slight refinement of [4, condition (4)]. It implies of course $x_0 \in B_\delta(x^*)$, and the additional terms on the LHS are sufficient to ensure $\{x_k\}, \{\tilde{x}_k\} \subset B_\delta(x^*)$. It is not difficult to see from the proofs that it can be replaced with $\{x_k\}, \{\tilde{x}_k\} \subset B_\delta(x^*)$. When the desingularizing function is of the form $\varphi(t) = \frac{M}{\theta}t^\theta$ with $\theta \in (0, \frac{1}{2}]$ we obtain more explicit convergence rates on the length and the objective gap. For convex functions this KL property is equivalent to Hölderian error bounds, as proved more in general in [10, Theorem 5].

Corollary 5.1. *Under the assumptions of Proposition 5.1, let $\varphi(t) = \frac{M}{\theta}t^\theta$ for some $\theta \in (0, \frac{1}{2}]$. Then*

$$\tilde{f}(x_k) - \tilde{f}(x^*) = \begin{cases} O(1/k^{\frac{1}{1-2\theta}}) & \text{for } 0 < \theta < \frac{1}{2}, \\ O((1 + \frac{a}{b^2M^2})^{-k}) & \text{for } \theta = \frac{1}{2}. \end{cases} \quad (5.9)$$

Moreover

$$\sum_{i=k}^{\infty} \|x_i - x_{i+1}\| = \begin{cases} O(1/k^{\frac{\theta}{1-2\theta}}) & \text{for } 0 < \theta < \frac{1}{2}, \\ O((1 + \frac{a}{b^2M^2})^{-\frac{k}{2}}) & \text{for } \theta = \frac{1}{2}. \end{cases} \quad (5.10)$$

Finally, the implicit constants can be taken as functions of M, θ, a, b, η , and in particular independent from x_0 .

The proofs of Proposition 5.1 and Corollary 5.1 are included in Section 8.3.

As a corollary of Proposition 5.1, by Proposition 4.1 we have the following results on the convergence rates of Algorithm 2:

Corollary 5.2. *Let us consider Problem (2.1) with $f \in C^1(\Omega)$ satisfying the KL inequality w.r.t. $x^* \in \Omega$ in $B_\delta(x^*) \cap [f(x^*) < f < \eta]$. Assume that $\{x_k\}$ is generated by Algorithm 2, and that the method \mathcal{A} in the SSC satisfies the following conditions:*

- the SSC procedure always terminates in a finite number of steps;
- $\text{SB}_{\mathcal{A}}(\Omega) = \tau > 0$.

Assume that (5.4), (5.5) and (5.6) hold with $a = \frac{L}{2}$ and $b = \tau/(L(1+\tau))$. Then the sequence converges to a stationary point and the convergence rates (5.7) and (5.8) hold.

Proof. Under our general assumptions on \mathcal{A} by Proposition 4.1 the sequence $\{x_k\}$ satisfies (H1) and (H'2) in the form (H'2a). It then suffices to apply Proposition 5.1 to $\tilde{f} = f + i_\Omega$ to obtain the convergence rates (5.7), (5.8). Finally, \tilde{x}^* is stationary by Corollary 4.1 and $f(\tilde{x}^*) = f(x^*)$ by continuity. \square

Remark 5.1. If $x_{k+1} = x_{k+1}^{tr}$ in Algorithm 2 the property (5.5) holds in particular when x^* is a global minimum in the connected component C_0 containing x_k of the sublevel set $[-\infty < f \leq f(x_k)]$. Indeed x_{k+1}^{tr} cannot be outside C_0 since

$$x_{k+1}^{tr} \in B_{\frac{\|g\|}{2L}}(x_k + \frac{g}{2L}) \subset C_0,$$

where $g = -\nabla f(x_k)$ and $B_{\frac{\|g\|}{2L}}(x_k + \frac{g}{2L}) \subset C_0$ because it is connected and contained in

$$[-\infty < f \leq f(x_k)]$$

by (4.9).

6 Analysis of some projection-free first-order methods

Here, we report some relevant examples of projection-free first-order methods and show that those methods satisfy the assumptions given in Corollary 5.2. More specifically, we show finiteness of the SSC procedure and that a sufficient slope condition is satisfied when the method is applied over a specific class of feasible sets. The reported examples include the AFW, PFW, FDFW on polytopes, the FDFW on sublevel sets of strongly convex smooth functions, a method based on orthographic retractions for convex sets with smooth boundary, as well as combination of these methods on product domains. At the end of the section we show how to apply our framework together with a KL property of (non convex) quadratic programming problems to obtain an asymptotic linear convergence rate for some FW variants.

Before giving the examples we prove a broad SSC finite termination criterion, showing that the procedure ends when using a first order method with mild convergence properties for linear objectives.

Lemma 6.1. *Assume that the method \bar{A} applied to any linear function $L_g(x) = -\langle g, x \rangle$ and with every stepsize maximal always generates a (possibly finite) sequence $\{y_j\}$ such that*

$$\liminf \pi_{y_j}(g) = 0. \tag{6.1}$$

Then the SSC with the method \bar{A} always terminates in a finite number of steps.

Proof. Assume by contradiction that the SSC generates an infinite sequence $\{y_j\}$. In this case the method \bar{A} applied in the SSC with gradient $-g$ always does maximal steps, since the SSC terminates as soon as $\alpha_j < \alpha_{\max}^{(j)}$ so that $\alpha_j = \beta_j$. Let $p_j = \pi_{y_j}(g)$. Then by (6.1) we can take $\{j(k)\}$ subsequence of indexes such that $p_{j(k)} \rightarrow 0$. We claim that $y_{j(k)} \rightarrow y_0$, in contradiction with the strict monotonicity of $j \rightarrow \langle g, y_j \rangle$. Indeed we have

$$\|y_{j(k)} - y_0\| \leq \frac{\langle \hat{d}_{j(k)}, g \rangle}{L} \leq \frac{p_{j(k)}}{L} \rightarrow 0,$$

by definition of $\Omega_{j(k)}$ in the first inequality and by Proposition 2.2 in the second. \square

6.1 PFW, AFW, FDFW directions

The AFW and PFW depend from a set of "elementary atoms" A such that $\Omega = \text{conv}(A)$. Given A , for a base point $x \in \Omega$ we can define

$$S_x = \{S \subset A \mid x \text{ is a proper convex combination of all the elements in } S\},$$

the family of possible active sets for x . In the rest of the article A is always clear from the context and for simplicity we write PFW, AFW instead of PFW_A , AFW_A . For $x \in \Omega$, $S \in S_x$, d^{PFW} is a PFW direction with respect to the active set S and gradient $-g$ iff

$$d^{\text{PFW}} = s - q \text{ with } s \in \operatorname{argmax}_{s \in \Omega} \langle s, g \rangle \text{ and } q \in \operatorname{argmin}_{q \in S} \langle q, g \rangle.$$

Similarly, given $x \in \Omega$, $S \in S_x$, d^{AFW} is an AFW direction with respect to the active set S and gradient $-g$ iff

$$d^{\text{AFW}} \in \operatorname{argmax} \{ \langle g, d \rangle \mid d \in \{d^{\text{FW}}, x - q\} \} \text{ with } q \in \operatorname{argmin}_{q \in S} \langle q, g \rangle, \quad (6.2)$$

where d^{FW} is a classic Frank Wolfe direction

$$d^{\text{FW}} = s - x \text{ with } s \in \operatorname{argmax}_{s \in \Omega} \langle s, g \rangle. \quad (6.3)$$

If the FW method is applied to a linear objective $L_g(x) = -\langle g, x \rangle$ it clearly terminates after at most one maximal step, with $y_1 \in \operatorname{argmin}_{x \in \Omega} L_g(x)$. As for the AFW or the PFW, they generate a sequence of active sets $\{S^{(j)}\}$ with y_j proper convex combination of the elements in $S^{(j)}$. After a FW step or a PFW step the linear minimizer s_j is always in the active set, and we assume that no other point can be added at the step j . In particular, $S^{(j+1)} \subset S^{(j)} \cup \{s_j\}$. Now on the one hand, all the linear minimizers added to the active set can't be dropped from the active set, since the gradient doesn't change. On the other hand, after every maximal away or PFW step the element q_j corresponding to the away direction is dropped from the active set. Furthermore, as the FW method also the AFW terminates after a maximal FW step. Then the AFW and the PFW applied to linear objectives terminate in at most $|S^{(0)}|$ steps. Applying Lemma 6.1 to obtain finite SSC termination we have the following:

Proposition 6.1. *For the FW, the PFW and the AFW applied to a linear objective the number of maximal steps is bounded, provided that no points beside linear minimizers are added to the active set. In particular, under this assumption for these methods the SSC terminates.*

The FDFW from [18] relies only on the current point x and the current gradient $-g$ to choose a descent direction and, unlike the AFW and the PFW, does not need to keep track of the active set. This method makes use of a linear minimization oracle

$$\text{LMO}_C(-g) \in \operatorname{argmin}_{x \in C} \langle -g, x \rangle,$$

for C varying among the faces of Ω . Let $\mathcal{F}(x)$ be the minimal face of Ω containing x . The in face direction is defined as

$$d^A = x_k - x_A \text{ with } x_A \in \operatorname{argmin} \{ \langle g, y \rangle \mid y \in \mathcal{F}(x) \}.$$

The selection criterion is then analogous to the one used by the AFW:

$$d^{\text{FD}} \in \operatorname{argmax} \{ \langle g, d \rangle \mid d \in \{d^A, d^{\text{FW}}\} \}. \quad (6.4)$$

For the FDFW we assume that the maximal stepsize is given by feasibility conditions as in [18]:

$$\alpha_{\max}(x, d) = \max \{ \alpha \in \mathbb{R}_{\geq 0} \mid x + \alpha d \in \Omega \}. \quad (6.5)$$

Then after a maximal in face step from y_j we have $\dim \mathcal{F}(y_{j+1}) < \dim \mathcal{F}(y_j)$ because y_{j+1} lies on the boundary of $\mathcal{F}(y_j)$. Whence there can only be a finite number of consecutive such steps. Furthermore, we've seen that after a maximal FW step every method applied to a linear objective terminates. Applying Lemma 6.1 to obtain finite SSC termination we have the following:

Proposition 6.2. *The FDFW on any compact and convex set Ω does at most $\dim(\Omega) + 1$ consecutive maximal steps. In particular, for this method the SSC terminates.*

We write $\text{SB}_{\text{FD}}, \text{DSB}_{\text{FD}}$ instead of $\text{SB}_{\text{FD,FW}}, \text{DSB}_{\text{FD,FW}}$ in the rest of the paper.

When Ω is a polytope and $|A| < \infty$ the sufficient slope conditions hold for the directions we introduced. Before stating a lower bound for $\text{SB}_{\mathcal{A}}(\Omega)$ in this setting we need to recall the definition of pyramidal width $\text{PWidth}(A)$ as it was given in [30]. We refer the reader to [40] for a discussion of various properties of this parameter.

For a given $g \in \mathbb{R}^n \setminus \{0\}$ the pyramidal directional width is defined as

$$\text{PdirW}(A, g, x) = \min_{S \in \mathcal{S}_x} \max_{\substack{a \in A \\ s \in S}} \langle \frac{g}{\|g\|}, a - s \rangle, \quad (6.6)$$

and the pyramidal width is defined as

$$\text{PWidth}(A) = \min_{\substack{F \in \text{faces}(\text{conv}(A)), x \in F \\ g \in \text{cone}(F-x) \setminus \{0\}}} \text{PdirW}(F \cap A, g, x).$$

Here we use one key property of $\text{PWidth}(A)$ which relates it to the slope along the PFW direction. We have the following lower bound (see [30, equation (12)])

$$\frac{\langle g, d^{\text{PFW}} \rangle}{\langle g, \hat{e} \rangle} \geq \text{PWidth}(A) > 0, \quad (6.7)$$

where e is any direction in $\Omega - \{x\} = T_{\Omega}(x)$ which is also a descent direction¹ for $-g$. Another relevant property is that by (6.6) we have that $\text{PWidth}(A)$ is monotone decreasing in A , so that if $V(\Omega)$ is the set of vertexes of Ω we always have $\text{PWidth}(A) \leq \text{PWidth}(V(\Omega))$.

Proposition 6.3. *Let $\text{diam}(\Omega) = D$. Then*

$$\begin{aligned} \text{SB}_{\text{PFW}}(\Omega) &\geq \frac{\text{PWidth}(A)}{D}, \\ \text{SB}_{\text{AFW}}(\Omega) &\geq \frac{\text{PWidth}(A)}{2D}, \\ \text{SB}_{\text{FD}}(\Omega) &\geq \frac{\text{PWidth}(V(\Omega))}{2D}. \end{aligned} \quad (6.8)$$

Proof. Let g be such that $\pi_x(g) \neq 0$. Then there exists descent directions for $-g$ feasible for Ω from x , and

$$0 < \max_{e \in \Omega - \{x\}} \langle g, \hat{e} \rangle,$$

so that

$$\min_{\substack{e \in \Omega - \{x\}, \\ \langle g, \hat{e} \rangle > 0}} \frac{1}{\langle g, \hat{e} \rangle} = \frac{1}{\max_{e \in \Omega - \{x\}} \langle g, \hat{e} \rangle} = \frac{1}{\sup_{h \in \Omega \setminus \{\bar{x}\}} \left(g, \frac{h - \bar{x}}{\|h - \bar{x}\|} \right)} = \frac{1}{\|\pi(T_{\Omega}(x), g)\|}, \quad (6.9)$$

where we used Proposition 2.2 in the last equality. Thanks to (6.9) taking the min on all the feasible descent directions e in the LHS of (6.7) we obtain

$$\frac{\langle g, d^{\text{PFW}} \rangle}{\|\pi(T_{\Omega}(x), g)\|} \geq \text{PWidth}(A).$$

We now have

$$\text{DSB}_{\text{PFW}}(\Omega, x, g) = \inf_{d^{\text{PFW}} \in \text{PFW}(x, g)} \frac{\langle g, d^{\text{PFW}} \rangle}{\|d^{\text{PFW}}\| \|\pi(T_{\Omega}(x), g)\|} \geq \frac{\langle g, d^{\text{PFW}} \rangle}{D \|\pi(T_{\Omega}(x), g)\|} \geq \frac{\text{PWidth}(A)}{D}.$$

¹In [30] the direction e is defined as a possible error direction $e = x^* - x$, where x^* is an optimal point of a convex objective f with $\nabla f(x) = -g$. However, this definition is equivalent to ours. Indeed if $e = x^* - x$ then by convexity it must be a feasible descent direction for $-g$. Conversely, every feasible descent direction is always an error direction as defined above for some choice of f , i.e. consider $f(y) = \frac{1}{2} \langle y - x^*, Q(y - x^*) \rangle$ with Q positive definite such that $\nabla f(x) = Q(x - x^*) = -g$.

and the first part of (6.8) follows by taking the inf on the LHS for $x \in \Omega$ and g such that $\pi_x(g) \neq 0$. Inequality (6.8) for the AFW method is a corollary since

$$\langle g, d^{\text{AFW}} \rangle \geq \frac{1}{2} \langle g, d^{\text{PFW}} \rangle,$$

as it follows immediately from the definitions (see also [30, equation (6)]).

For the FDFW, first observe that for any $S \in \mathcal{S}_x$:

$$\begin{aligned} \langle g, d^{\text{FW}} + d^A \rangle &= \max\{\langle g, s - q \rangle \mid s \in \Omega, q \in \mathcal{F}(x)\} \geq \max\{\langle g, s - q \rangle \mid s \in \Omega, q \in S\} \\ &= \langle g, d^{\text{PFW}} \rangle, \end{aligned} \quad (6.10)$$

where we used $\mathcal{F}(x) \supset S$ in the inequality and the last equality follows immediately from the definition of d^{PFW} . We then have

$$\text{DSB}_{\text{FD}}(\Omega, x, g) = \frac{\langle g, d^{\text{FD}} \rangle}{\pi_x(g) \|d^{\text{FD}}\|} \geq \frac{\langle g, d^{\text{FW}} + d^A \rangle}{2D\pi_x(g)} \geq \frac{\langle g, d^{\text{PFW}} \rangle}{2\pi_x(g) \|d^{\text{FD}}\|} \geq \frac{\text{PWidth}(A)}{2D}, \quad (6.11)$$

where the first inequality follows from the choice criterion (6.4) together with $\|d^{\text{FD}}\| \leq D$, we used (6.10) in the second inequality, and (6.7) in the third. The thesis follows by taking the inf on x and g in the LHS and the sup on A in the RHS. \square

6.2 FW and FDFW on sublevel sets of strongly convex smooth functions

In spite of the zig-zagging behaviour discussed in the introduction, in any fixed point of a polytope the FW directions do satisfy the sufficient slope condition, e.g. by Proposition 6.3 since they are selected by the AFW and the PFW when the active set is a singleton. This property can be extended (uniformly) to the boundaries of sublevel sets of smooth strongly convex functions. As a consequence, using that the FDFW selects FW directions on the boundary of strictly convex sets, we also obtain a global slope bound for this method.

Proposition 6.4. *Let $h: \mathbb{R}^n \rightarrow \mathbb{R}$ be μ_h -strongly convex and with L_h -Lipschitz gradient, $a > \min_{x \in \mathbb{R}^n} h(x)$, $\Omega = \{x \in \mathbb{R}^n \mid h(x) \leq a\}$. Then:*

$$\text{SB}_{\text{FW}}(\Omega, \partial\Omega) \geq \frac{u_h}{2L_h}, \quad (6.12)$$

$$\text{SB}_{\text{FDFW}}(\Omega) \geq \frac{u_h}{2L_h}. \quad (6.13)$$

Proof. Since for every $x \in \partial\Omega$

$$N_\Omega(\bar{x}) = \text{cone}(\{\nabla h(x)\}),$$

we can apply Lemma 8.1 with $J(x) = \nabla h(x)$. We obtain

$$\text{SB}_{\text{FW}}(\Omega, \partial\Omega) \geq \inf_{\substack{h(y) = h(\bar{x}) = a \\ \bar{x} \neq y}} \frac{\langle \nabla h(y), y - \bar{x} \rangle}{\|\nabla h(\bar{x}) - \nabla h(y)\| \|y - \bar{x}\|}. \quad (6.14)$$

First, by L_h -smoothness we have:

$$\|\nabla h(\bar{x}) - \nabla h(y)\| \leq L_h \|\bar{x} - y\|. \quad (6.15)$$

Second, by μ_h strong convexity

$$\langle \nabla h(y), y - \bar{x} \rangle \geq h(\bar{x}) - h(y) + \frac{\mu_h}{2} \|y - \bar{x}\|^2 = \frac{\mu_h}{2} \|y - \bar{x}\|^2, \quad (6.16)$$

where we used $h(y) = h(\bar{x}) = a$ in the equality. The first part of the thesis follows applying (6.15) and (6.16) to the RHS of (6.14).

By strong convexity for every $y \in \partial\Omega$

$$\frac{\mu_h}{2} \|y - x^*\|^2 \leq h(y) - h^* = a - h^*,$$

and therefore $\|y - x^*\| \leq \sqrt{2(a - h^*)/\mu_h}$. This clearly implies

$$D \leq 2\sqrt{2(a - h^*)/\mu_h}. \quad (6.17)$$

Analogously, by L_h smoothness for every $y \in \partial\Omega$

$$\frac{L_h}{2} \|y - x^*\|^2 \geq h(x) - h^* = a - h^*, \quad (6.18)$$

and therefore $\|y - x^*\| \geq \sqrt{2(a - h^*)/L_h}$. We can now give a slope bound for $x \in \text{int}(\Omega)$:

$$\begin{aligned} \text{DSB}_{\text{FD}}(\Omega, x, g) &\geq \frac{\langle g, d^{\text{FW}} + d^A \rangle}{2D\pi_x(g)} = \frac{\langle \hat{g}, d^{\text{FW}} + d^A \rangle}{2D} \geq \frac{1}{4} \sqrt{\frac{\mu_h}{2(a - h^*)}} \max_{s, q \in \Omega} \langle \hat{g}, s - q \rangle \\ &\geq \frac{1}{4} \sqrt{\frac{\mu_h}{2(a - h^*)}} 2\sqrt{\frac{2(a - h^*)}{L_h}} = \frac{1}{2} \sqrt{\frac{\mu_h}{L_h}} \geq \frac{\mu_h}{2L_h}, \end{aligned} \quad (6.19)$$

where we used (6.11) in the first inequality, $\pi_x(g) = \|g\|$ in the first equality, (6.17) in the second inequality, $x^* \pm \sqrt{2(a - h^*)/L_h}\hat{g} \in \Omega$ by (6.18) in the third inequality, and $\mu_h/L_h \leq 1$ in the last one. But on the boundary the FDFW selects FW directions so that (6.12) holds also for the FDFW, and in particular (6.13) follows by combining (6.12) with (6.19) for points on the interior. \square


Remark 6.1. The slope bound (6.13) clearly holds for any method selecting FW directions on the boundary and other directions on the interior satisfying the sufficient slope condition for $\mu_h/2L_h$ (e.g. the negative gradient).

6.3 Orthographic retractions for convex sets with smooth boundary

In this section we show an example of how the SSC can be employed effectively to deal with short steps also in non linearly constrained settings. More specifically, we consider here the case of Ω compact, convex, full dimensional and with smooth boundary of class C^1 (e.g., the unitary ball w.r.t. $\|\cdot\|_p$, with $p \in (1, \infty)$). In this case for $x \in \partial\Omega$ we can define $J(x)$ as the outward pointing normal and $T_x\Omega$ as the tangent hyperplane to $\partial\Omega$ in x . For $x \in \partial\Omega$, $u \in T_x\Omega$, we then define $P(x, u)$ as the first intersection between Ω and $\{x + u\} + \text{line}(-J(x))$. By smoothness $P(x, u)$ is always defined for $\|u\|$ small enough. This transformation is called an orthographic retraction, and can be easily extended to manifolds with arbitrary codimension. Orthographic retractions are widely used in manifold optimization for manifolds of class at least C^2 (see e.g. [1], [2], [26], [46]), in which case it can be proved that they are differentiable retractions [1]. In the codimension 1 case orthographic retractions were used in [7] for optimization over proximally smooth surfaces, and in [33] for optimization on sublevel sets of convex functions with Lipschitz gradient. In these works the additional regularity assumptions ensure that the manifold is sufficiently flat so that steps with length proportional to $\|g_\Omega(x)\|$ are effective. Here we do not need additional regularity assumptions, thanks to the SSC which allows our method to choose shorter steps while skipping gradient computations.

For $g \in \mathbb{R}^n$ we denote with $g_J(x)$ and $g_\Omega(x)$ the components of g along $J(x)$ and $T_x\Omega$, so that $g = g_J(x) + g_\Omega(x)$, with $\|g_\Omega(x)\| = \pi_x(g)$. We can finally define the Short Orthographic Retractions (SOR) method directions $\text{SOR}(x, g)$ simply as g if $g \in \text{int}(T_\Omega(x))$, and using orthographic projections when (x, g) is in

$$\bar{T}(\Omega) = \{(x, g) \in \partial\Omega \times \mathbb{R}^n \mid g \notin \text{int}(T_\Omega(x)) \cup N_\Omega(x)\}.$$



(a) Projection on $T_x\Omega$ and orthographic projection from $T_x\Omega$.

(b) The directions in $\text{SOR}(x, g)$ point in the arc between x and $P(\bar{v}) = P(x, \bar{v}(x, g)g_\Omega(x))$ for $\bar{\tau} = \frac{1}{2}$.

Figure 2

More precisely, in this case we define

$$\text{SOR}(x, g) = \{P(x, \lambda g_\Omega(x)) \mid \lambda \in (0, \bar{v}(x, g)]\},$$

where $\bar{v}(x, g) > 0$ is a proper "shrinking coefficient", ensuring that the orthographic projection is taken close to x enough to satisfy the usual slope condition for a fixed $\bar{\tau} \in (0, 1)$:

$$\bar{v}(x, g) = \sup \left\{ \lambda \in \mathbb{R}_{\geq 0} \mid \frac{\langle g, P(x, \lambda g_\Omega(x)) - x \rangle}{\pi_x(g) \|P(x, \lambda g_\Omega(x)) - x\|} \geq \bar{\tau} \right\}. \quad (6.20)$$

By smoothness we have $\bar{v}(x, g) > 0$ whenever $(x, g) \in \bar{T}(\Omega)$. It is also immediate to check from the definition that for $\bar{\lambda} > 0$ we have $\bar{\lambda} \bar{v}(x, \bar{\lambda}g) = \bar{v}(x, g) = \bar{v}(x, \hat{g})/\|g\|$. We now prove the anticipated slope lower bound:

Proposition 6.5. *In the setting introduced above $\text{SB}_{\text{SOR}}(\Omega) \geq \bar{\tau}$.*

Proof. When $g \in \text{int}(T_\Omega(x))$ we obviously have $\text{DSB}_{\text{SOR}}(\Omega, x, g) = 1 > \bar{\tau}$ since $\text{SOR}(x, g) = \{g\}$. By the definition of \bar{v} if $(x, g) \in \bar{T}(\Omega)$ then $P(x, \bar{v}(x, g)g_\Omega(x)) \in \partial\Omega$ and (see also figure 2b)

$$v = P(x, \bar{v}(x, g)g_\Omega(x)) - x \in \text{int}(T_\Omega(x)).$$

In particular $P(x, \lambda \bar{v}(x, g)g_\Omega(x))$ is in $\text{conv}(\{x + \lambda \bar{v}(x, g)g_\Omega(x), x + \lambda v\})$ strictly below $x + \lambda v$ for every $\lambda \in (0, 1)$. Therefore, for $v_\lambda = P(x, \lambda \bar{v}(x, g)g_\Omega(x)) - x$

$$\frac{\langle g, v_\lambda \rangle}{\pi_x(g) \|v_\lambda\|} > \frac{\langle g, v \rangle}{\pi_x(g) \|v\|} \geq \bar{\tau}, \quad (6.21)$$

where the second inequality follows by the inequality in the definition (6.20) of $\bar{v}(x, g)$. The equation above directly implies the desired bound

$$\begin{aligned} \text{DSB}_{\text{SOR}}(\Omega, x, g) &= \inf_{\lambda \in (0, 1]} \frac{\langle g, v_\lambda \rangle}{\pi_x(g) \|v_\lambda\|} \\ &= \frac{\langle g, v \rangle}{\pi_x(g) \|v\|} \geq \bar{\tau}. \end{aligned}$$

and taking into account the previous case we can conclude $\text{SB}_{\text{SOR}}(\Omega) \geq \bar{\tau}$. \square

By definition of $\text{SOR}(x, g)$ we only need a lower bound on $\bar{\nu}(x, u)$ to compute a direction $d \in \text{SOR}(x, g)$. While in practice such bounds can be obtained with an adaptive strategy, for sublevel sets simple lower bounds depending from a smoothness modulus are sometimes available (see Proposition 8.2 and Remark 8.1 in the appendix for some examples). The general principle in this case is that for every $x \in \partial\Omega$ a "simple" smooth full dimensional convex set $\bar{B}(x) \ni x$ (e.g. a ball) is contained in Ω , with the same tangent hyperplane in x . We can therefore bound $\bar{\nu}$ with the shrinking coefficient relative to $\bar{B}(x)$, which is much easier to control.

The SOR method directions are well suited for the SSC, given that they satisfy the sufficient slope condition, but in general they are too short to ensure descent sequence like conditions in a single step. Still, even with the SSC some lower bound on the norm of the directions chosen is needed to obtain finite termination results. In other words, it is necessary to "throw away" the shortest directions selected by the SOR. Here for simplicity we consider the method

$$\overline{\text{SOR}}(\Omega, x, g) = \{P(x, \lambda g_\Omega(x)) \mid \min(\hat{\nu}, \hat{\nu}\bar{\nu}(x, g)) \leq \lambda \leq \bar{\nu}(x, g)\},$$

where $\hat{\nu} \in (0, 1]$ is a fixed constant in the rest of the paper. It is clear that

$$\text{SB}_{\overline{\text{SOR}}}(\Omega) \geq \text{SB}_{\text{SOR}}(\Omega) \geq \bar{\tau},$$

given that the directions selected by the $\overline{\text{SOR}}$ are a subset of the ones selected by the SOR. In the next lemma we prove that also for the $\overline{\text{SOR}}$ we have convergence for linear objectives, and in particular finite SSC termination thanks to Lemma 6.1.

Lemma 6.2. *If the method $\overline{\text{SOR}}$ applied to the linear function $L_g(x) = -\langle g, x \rangle$ generates a sequence $\{y_j\}$ with stepsizes maximal w.r.t. boundary conditions (as in (6.5)), then*

$$\lim \pi_{y_j}(g) = 0. \tag{6.22}$$

In particular, the SSC with the $\overline{\text{SOR}}$ always terminates in a finite number of steps.

Proof. We just need to prove (6.22) since thanks to Lemma 6.1 finite SSC termination immediately follows. Furthermore, if the sequence is finite, terminating in a stationary point for $-g$, (6.22) is obvious, so that from now on we always assume that $\{y_j\}$ is an infinite sequence. Let

$$I = \{j \in \mathbb{N}_0 \mid (y_j, g) \in \bar{T}(\Omega)\},$$

and let $\{\lambda_j\}_{j \in I}$ be such that $d_j = P(y_j, \lambda_j g_\Omega(y_j)) - y_j$. We first claim that if the sequence doesn't stop in a stationary point then $\mathbb{N} \subset I$. To start with, since the stepsizes must always be maximal, $y_j \in \partial\Omega$ for every $j \geq 1$. Now if $g \in \text{int}(T_\Omega(y_j))$ and also $g \in \text{int}(T_\Omega(y_{j+1})) = \text{int}(T_\Omega(y_j + \alpha_j g))$ then necessarily $\text{int}(\Omega) \ni y_j + \alpha_j g = y_{j+1}$, a contradiction. On the other hand if $(y_j, g) \in \bar{T}(\Omega)$ then $y_j + d_j \in \partial\Omega$ with $y_j + d_j - \lambda J(x) \notin \Omega$ for every $\lambda > 0$. Then by convexity it is easy to see

$$y_j + \alpha_j d_j + \lambda_1 g_\Omega(\Omega) + \lambda_2 J(x) \notin \Omega$$

for every $\lambda_2 \geq 0, \lambda_1 > 0$, and thus in particular

$$y_{j+1} + \lambda g = y_j + \alpha_j d_j + \lambda g_\Omega(y_j) + \lambda \|g_J(y_j)\| J(x) \notin \Omega$$

for every $\lambda > 0$. Equivalently, $g \notin \text{int}(T_\Omega(y_{j+1}))$ as desired.

Let $\{y_{j(k)}\} \rightarrow \bar{y}$ be a convergent subsequence of $\{y_j\}$, and assume by contradiction

$$\lim_{k \rightarrow \infty} \pi_{y_{j(k)}}(g) = \pi_{\bar{y}}(g) \neq 0,$$

so that $(\bar{y}, g) \in \bar{T}(\Omega)$ and $\bar{\nu}(\bar{y}, g) > 0$. The sequence $f_j = \langle g, y_j \rangle$ is strictly decreasing with limit $\bar{f} \in \mathbb{R}$ by compactness. In particular we have

$$\begin{aligned}
\langle g, y_1 \rangle - \bar{f} &= \sum_{j=1}^{\infty} \langle g, y_j - y_{j+1} \rangle \geq \sum_{k=1}^{\infty} \langle g, y_{j(k)} - y_{j(k+1)} \rangle \\
&\geq \sum_{k=1}^{\infty} \langle g, P(y_{j(k)}, \lambda_{j(k)} g_{\Omega}(y_{j(k)})) - y_{j(k)} \rangle \geq \bar{\tau} \sum_{k=1}^{\infty} \|P(y_{j(k)}, \lambda_{j(k)} g_{\Omega}(y_{j(k)})) - y_{j(k)}\| \pi_{y_{j(k)}}(g) \quad (6.23) \\
&\geq \bar{\tau} \sum_{k=1}^{\infty} \lambda_{j(k)} \pi_{y_{j(k)}}(g)^2 \geq \bar{\tau} \hat{\nu} \sum_{k=1}^{\infty} \pi_{y_{j(k)}}(g)^2 \min(\bar{\nu}(y_{j(k)}, g), 1),
\end{aligned}$$

where we used $\alpha_{j(k)} = \alpha_{j(k)}^{\max} \geq 1$ in the second inequality, $\text{SB}_{\overline{\text{SOR}}}(\Omega) \geq \bar{\tau}$ in the third one,

$$\|P(y_{j(k)}, \lambda_{j(k)} g_{\Omega}(y_{j(k)})) - y_{j(k)}\| \geq \lambda_{j(k)} \pi_{y_{j(k)}}(g)$$

in the fourth one, and $\lambda_j \geq \min(\hat{\nu} \bar{\nu}(y_{j(k)}, g), \hat{\nu})$ in the last one. By construction $\pi_{y_{j(k)}}(g)$ converges to $\pi_{\bar{y}}(g) \neq 0$, and therefore (6.23) implies $\bar{\nu}(y_{j(k)}, g) \rightarrow 0$, since the last sum is bounded by the LHS. But by the lower semicontinuity property proved in Lemma 8.2 this is in contradiction with $\bar{\nu}(\bar{y}, g) > 0$. \square

Remark 6.2. If instead of P we have a retraction $R(x, u)$, defined in a neighborhood U of $\partial\Omega \times \{0\} \subset T\partial\Omega$, such that

$$R(x, tu) = x + tu + o(t), \quad (6.24)$$

then we can always define a shrinking coefficient $\bar{\nu}_R$ analogous to $\bar{\nu}$ and a method \mathcal{A}_R analogous to the SOR. The essential property needed to have finite termination is that for $x_j \rightarrow \bar{x}$ such that $(\bar{x}, g) \in \bar{T}(\Omega)$ and $d_j \in \mathcal{A}_R(x_j, g)$

$$\liminf_{j \rightarrow \infty} \|d_j\| \neq 0.$$

This is satisfied by the $\overline{\text{SOR}}$ thanks to Lemma 8.2 and can be satisfied by analogous variants e.g. if the remainder $o(t)$ in (6.24) is uniform for x, u varying in U .

6.4 Directions on product domains

Assume that the feasible set is block separable, that is $\Omega = \Omega_{(1)} \times \dots \times \Omega_{(m)}$ with $\Omega_{(i)} \subset \mathbb{R}^{n_i}$ compact and convex for $1 \leq i \leq m$. The next proposition states that we can select directions with positive slope bound for Ω by properly weighting directions selected by algorithms with positive slope bound in each of the factors. For a block vector $x \in \mathbb{R}^n = \mathbb{R}^{n_1} \times \dots \times \mathbb{R}^{n_m}$ we denote with $x_{(i)} \in \mathbb{R}^{n_i}$ the component corresponding to the i -th block, so that $x = (x_{(1)}, \dots, x_{(m)})$.

Proposition 6.6. *For Ω defined as above and $1 \leq i \leq m$. let \mathcal{A}^i be such that $\text{SB}_{\mathcal{A}^i}(\Omega_{(i)}) = \tau_i > 0$, and let $\tau^* = \min_{1 \leq i \leq m} \tau_i$.*

1. Assume that $d \in \mathcal{A}(x, g)$ if and only if for every $i \in [1 : m]$ we have

$$d_{(i)} = w_i \hat{c}_i, \text{ with } c_i \in \mathcal{A}^i(x_{(i)}, g_{(i)}), w_i = \langle g, \hat{c}_i \rangle. \quad (6.25)$$

Then $\text{SB}_{\mathcal{A}}(\Omega) = \tau^*$.

2. Assume that $d \in \mathcal{A}(x, g)$ if and only if for every $i \in [1 : m]$ we have

$$d_{(i)} = w_i \hat{c}_i, \text{ with } c_i \in \mathcal{A}^i(x_{(i)}, g_{(i)}), w_i \geq 0 \quad (6.26)$$

for a weight vector $w \in \mathbb{R}^m \setminus \{0\}$ such that

$$\text{argmax}_{i \in [1 : m]} w_i \cap \text{argmax}_{i \in [1 : m]} \langle g, \hat{c}_i \rangle \neq \emptyset.$$

Then $\text{SB}_{\mathcal{A}}(\Omega) \geq \frac{\tau^*}{m}$.

Proof. For Ω product of closed convex sets the normal cone is the product of the normal cones of the factors (see e.g. [5, Table 4.3])

$$N_{\Omega}(x) = N_{\Omega_{(1)}}(x_{(1)}) \times \dots \times N_{\Omega_{(m)}}(x_{(m)}). \quad (6.27)$$

As a consequence of (6.27) for $g \in \mathbb{R}^n$

$$\text{dist}(g, N_{\Omega}(x)) = \sqrt{\sum_{i=1}^m \text{dist}(g_{(i)}, N_{\Omega_{(i)}}(x_{(i)}))^2}. \quad (6.28)$$

For $d \in \mathcal{A}(x, g)$ we have

$$\|d\| = \sqrt{\sum_{i=1}^m \|d_{(i)}\|^2} = \sqrt{\sum_{i=1}^m w_i^2} \quad (6.29)$$

and

$$\langle g, d \rangle = \sum_{i=1}^m \langle g_{(i)}, d_{(i)} \rangle = \sum_{i=1}^m w_i \langle g_{(i)}, \hat{c}_i \rangle, \quad (6.30)$$

with $c_i \in \mathcal{A}^i(x_{(i)}, g_{(i)})$ for $1 \leq i \leq m$.

Let $\pi_x^i(g) = \|\pi(T_{\Omega_{(i)}}(x_{(i)}), g_{(i)})\|$. Then we get

$$\pi_x(g) = \sqrt{\sum_{i=1}^m \pi_x^i(g)^2} \leq \frac{1}{\tau^*} \sqrt{\sum_{i=1}^m \langle g_{(i)}, \hat{c}_i \rangle^2} \quad (6.31)$$

by applying Proposition 2.2 to the RHS of (6.28) in the equality, and using

$$\langle g_{(i)}, \hat{c}_i \rangle \geq \text{DSB}_{\mathcal{A}^i}(\Omega_{(i)}, x_{(i)}, g_{(i)}) \pi_x^i(g) \geq \text{SB}_{\mathcal{A}^i}(\Omega_{(i)}) \pi_x^i(g) = \tau_i \pi_x^i(g) \geq \tau^* \pi_x^i(g)$$

in the inequality.

1. For $d \in \mathcal{A}(x, g)$ as in (6.25) we have

$$\frac{\langle g, d \rangle}{\pi_x(g) \|d\|} = \frac{\sum_{i=1}^m \langle g_{(i)}, \hat{c}_i \rangle^2}{\sqrt{\sum_{i=1}^m \pi_x^i(g)^2} \sqrt{\sum_{i=1}^m \langle g_{(i)}, \hat{c}_i \rangle^2}} = \frac{\sqrt{\sum_{i=1}^m \langle g_{(i)}, \hat{c}_i \rangle^2}}{\sqrt{\sum_{i=1}^m \pi_x^i(g)^2}} \geq \tau^*,$$

where we used (6.31), (6.29) and (6.30) in the first equality, and the second part of (6.31) in the inequality. Then

$$\text{DSB}_{\mathcal{A}}(\Omega, x, g) = \inf_{d \in \mathcal{A}(x, d)} \frac{\langle g, d \rangle}{\pi_x(g) \|d\|} \geq \tau^*,$$

and the lower bound $\text{SB}_{\mathcal{A}}(\Omega) \geq \tau^*$ follows by taking the inf on x, g .

Let \bar{m} such that $\tau_{\bar{m}} = \tau^*$. It is immediate to obtain the opposite inequality $\text{SB}_{\mathcal{A}}(\Omega) \leq \tau^*$ by considering the block gradients of the form $g = (0, \dots, g_{(\bar{m})}, \dots, 0)$ with all zeros outside the block \bar{m} , where for every such g we have

$$\text{DSB}_{\mathcal{A}}(\Omega, x, g) = \text{DSB}_{\mathcal{A}^{\bar{m}}}(\Omega_{(\bar{m})}, x_{(\bar{m})}, g_{(\bar{m})}),$$

by the definition of \mathcal{A} .

2. Let $l \in \text{argmax}_{i \in [1:m]} w_i \cap \text{argmax}_{i \in [1:m]} \langle g, \hat{c}_i \rangle$. Then for $d \in \mathcal{A}(x, g)$ as in (6.25)

$$\frac{\langle g, d \rangle}{\pi_x(g) \|d\|} = \frac{\sum_{i=1}^m w_i \langle g_{(i)}, \hat{c}_i \rangle}{\sqrt{\sum_{i=1}^m \pi_x^i(g)^2} \sqrt{\sum_{i=1}^m w_i^2}} \geq \frac{1}{\sqrt{m}} \frac{\langle g_{(l)}, \hat{c}_l \rangle}{\sqrt{\sum_{i=1}^m \pi_x^i(g)^2}} \geq \frac{\tau^*}{\sqrt{m}} \frac{\langle g_{(l)}, \hat{c}_l \rangle}{\sqrt{\sum_{i=1}^m \langle g_{(i)}, \hat{c}_i \rangle^2}} \geq \frac{\tau^*}{m},$$

where we used (6.29) and (6.30) in the first equality, $l \in \operatorname{argmax}_{i \in [1:m]} w_i$ in the first inequality, (6.31) in the second inequality, and $l \in \operatorname{argmax}_{i \in [1:m]} \langle g, \hat{c}_i \rangle$ in the last one. \square

The directions given in (6.25) have components normalized and then multiplied by the slopes in the corresponding blocks. Thus in particular \mathcal{A} does longer steps in the blocks with greater slope, which seems a rather intuitive way to obtain direction satisfying the sufficient slope condition. As an example, consider the case of a product of simplices

$$\Omega = \Delta_{n_1-1} \times \dots \times \Delta_{n_m-1},$$

relevant e.g. for structural SVMs optimization [31] and suitable for projection-free methods. Given $x \in \Omega$, $g \in \mathbb{R}^n$ it can be readily seen that the direction computed by the standard PFW is of the form $d_{(i)} = e_{s(i)} - e_{q(i)}$ for $i \in [1:m]$, with $s(i) \in \operatorname{argmax}_{j \in [1:n_i]} \langle g_{(i)}, e_j \rangle$ and $q(i)$ depending on the active set. Meanwhile, the direction suggested by Proposition 6.25 when applying the PFW to every block is of the form $d_{(i)} = w_i \hat{d}^{\text{PFW}_i}$ for $i \in [1:m]$, with $w_i = \langle g_{(i)}, \hat{d}^{\text{PFW}_i} \rangle$ and $\hat{d}^{\text{PFW}_i} \in \text{PFW}(x_{(i)}, g_{(i)})$.

We now state a termination result for the SSC procedure on product domains. Essentially, while finite termination holds for a reasonable choice of the maximal stepsize even for products of blocks with the AFW or the PFW, the blocks with the FDFW or the $\overline{\text{SOR}}$ must be dealt with one at a time. The result in particular implies finite termination of the SSC for the $\overline{\text{SOR}}$.

Proposition 6.7. *Let \mathcal{A} be as in (6.26), with \mathcal{A}^i the AFW, the PFW, the FDFW or the $\overline{\text{SOR}}$ for $i \in [1:m]$, and let $I \subset [1:m]$ the set of blocks with the FDFW or the $\overline{\text{SOR}}$. Assume that during the SSC only linear minimizers are added to the active sets in blocks with the AFW or the PFW. If*

- for every maximal step of \mathcal{A} in the SSC a maximal step is done in \mathcal{A}^i for some $i \in [1:m]$,
- whenever a maximal step is done by \mathcal{A}^l with $l \in I$ then the directions in the other blocks in I are 0 and

$$l \in \operatorname{argmax}_{i \in [1:m]} \langle \hat{c}_i, g_{(i)} \rangle,$$

then the SSC terminates.

Proof. Assume by contradiction that the SSC generates an infinite sequence $\{y_j\}$, leading to an infinite sequence of maximal steps. For every $i \in I^c$ by Proposition 6.1 the method $\mathcal{A}^{(i)}$ can do limited number of maximal steps (consecutive or not). Then for $j \geq \bar{l}$ large enough every maximal step must take place in a block in I . By assumption then the blocks in I must be changed one at a time and with maximal steps, so that each \mathcal{A}^i for $i \in I$ does consecutive maximal steps in its block. In particular, for some $l \in I$ infinite consecutive maximal steps are done in the l -th block. Since the FDFW can do a finite number of consecutive maximal steps, necessarily \mathcal{A}^l must be the $\overline{\text{SOR}}$. Let now $y_{ji} = (y_j)_{(i)}$, $d_{ji} = (d_j)_{(i)}$ and $p_{ji} = \pi_{y_{ji}}(g_{(i)})$. Let $j(k)$ be the subsequence of indexes corresponding to maximal steps in the l -th block. Then $p_{j(k)l} \rightarrow 0$ by Lemma 6.2. Whence $p_{j(k)} \rightarrow 0$, since

$$p_{j(k)} \leq \frac{1}{\tau^*} \sqrt{\sum_{i=1}^m \langle \hat{d}_{j(k)i}, g_{(i)} \rangle^2} \leq \frac{\sqrt{m}}{\tau^*} \langle \hat{d}_{j(k)l}, g_{(l)} \rangle \leq \frac{\sqrt{m}}{\tau^*} p_{j(k)l},$$

where we used (6.31) in the first inequality, $l \in \operatorname{argmax}_{i \in [1:m]} \langle \hat{d}_{j(k)i}, g_{(i)} \rangle$ in the second, and Proposition 2.2 in the third one. We therefore have all the hypotheses to apply Lemma 6.1. \square

We notice that the assumption on the maximal steps can be satisfied by simply taking α_{\max} as the minimum among the maximal stepsizes in the blocks. As for the second assumption, it is always compatible with (6.26), where e.g. both are satisfied by directions with only one weight different from 0, located in a block with maximal slope.

6.5 Linear convergence for quadratic objectives on polytopes with FW variants

In this section, we first prove that a uniform KL property together with the usual slope condition on \mathcal{A} imply an asymptotic linear convergence rate for a sequence generated by Algorithm 2. This result together with a KL property of quadratic programming problems proved in [16] and then more in general in [34] imply, in turn, asymptotic linear convergence rates for several FW variants.

Proposition 6.8. *Let us consider the set \mathcal{X} of stationary points of Problem (2.1). Assume that every point $x^* \in \mathcal{X}$ satisfies the KL inequality in $B_{\delta(x^*)}(x^*) \cap [f(x^*) < f < \eta(x^*)]$ for $\varphi_{x^*}(t) = 2M_{x^*}t^{\frac{1}{2}}$ and suitable $\delta(x^*) > 0$, $\eta(x^*) > f(x^*)$.*

Let $\{z_k\}$ be a sequence generated by Algorithm 2, with the method \mathcal{A} in the SSC satisfying the following:

- *the SSC procedure always terminates in a finite number of steps;*
- $\text{SB}_{\mathcal{A}}(\Omega) = \tau > 0$.

Then $z_k \rightarrow \tilde{x}^ \in \mathcal{X}$ at an asymptotic rate*

$$\|z_k - \tilde{x}^*\| = O\left(\left(1 + \frac{a}{b^2 M^2}\right)^{-\frac{k}{2}}\right),$$

where $a = L/2$, $b = \tau/(L(1+\tau))$, and M is a constant depending on Ω, f , but not z_0 .

Proof. By continuity for every $x^* \in \mathcal{X}$ we can define $\hat{\delta}(x^*) \in (0, \delta(x^*))$ so that conditions (5.4) and (5.6) hold for every $x_0 \in B_{\hat{\delta}(x^*)}(x^*) \cap [f(x^*) < f]$. By compactness we can then take a finite subset \mathcal{X}^* of stationary points such that

$$\bigcup_{\tilde{x} \in \mathcal{X}^*} B_{\hat{\delta}(\tilde{x})}(\tilde{x}) \supset \mathcal{X}.$$

Furthermore, by Corollary 4.1 the sequence $\{f(z_k)\}$ is decreasing and converges to $f^* \in \mathbb{R}$, and $\{z_k\}$ converges to \mathcal{X} . Then for some \bar{k} large enough and some $x^* \in \mathcal{X}^*$ such that $f(x^*) = f^*$ we must have $z_{\bar{k}} \in B_{\hat{\delta}(x^*)}(x^*)$. Given that $f(z_k)$ is decreasing and converges to f^* , also the assumption (5.5) is satisfied for $\{x_k\} = \{z_{\bar{k}+k}\}$. By Corollary 5.2 applied to $\{x_k\} = \{z_{\bar{k}+k}\}$ we obtain

$$\{x_k\} \rightarrow \tilde{x}^* \in B_{\hat{\delta}(x^*)}(x^*) \cap \mathcal{X},$$

with

$$\|x_k - \tilde{x}^*\| = O\left(\left(1 + \frac{a}{b^2 M_{x^*}^2}\right)^{-\frac{k}{2}}\right) \tag{6.32}$$

by Corollary 5.1, where the implicit constant can be taken independent from $z_{\bar{k}}$. To conclude,

$$\|z_k - \tilde{x}^*\| = O\left(\left(1 + \frac{a}{b^2 M_{x^*}^2}\right)^{-\frac{k-\bar{k}}{2}}\right) = O\left(\left(1 + \frac{a}{b^2 M_{x^*}^2}\right)^{-\frac{k}{2}}\right) = O\left(\left(1 + \frac{a}{b^2 M^2}\right)^{-\frac{k}{2}}\right)$$

for $M = \max\{M_{\tilde{x}} \mid \tilde{x} \in \mathcal{X}^*\}$ independent from z_0 , where we used (6.32) in the first equality. \square

We can now easily prove the result for quadratic programming problems with compact domain.

Corollary 6.1. *The conclusion of Proposition 6.8 holds in particular if $f(x) = \frac{1}{2}x^\top Qx - b^\top x$, Ω is a polytope, and the method \mathcal{A} used in the SSC is the AFW, the PFW or the FDFW, with no points beside linear minimizers added in the active set during the SSC for the AFW and the PFW.*

Proof. By [34, Corollary 5.2] the KL property needed is satisfied in particular when f is quadratic and Ω is a polytope. Moreover, when \mathcal{A} is the AFW, the PFW or the FDFW and $\Omega = \text{conv}(A)$ with $|A| < +\infty$ is a polytope then $\text{SB}_{\mathcal{A}}(\Omega) > 0$ by Proposition 6.3. In this setting we have also the finite termination property by Proposition 6.1 and Proposition 6.2. Therefore we have all the hypotheses to apply Proposition 6.8. \square

Remark 6.3. If for every $x \in \mathcal{X}$ there exists δ such that $f(y) = f(x)$ for every $y \in \mathcal{X} \cap B_{\delta}(x)$ then by [34, Theorem 4.1] the condition of Proposition 6.8 is implied by the Luo Tseng error bound. When instead the condition does not hold, i.e. we do not have the same desingularizing function for every stationary point (even up to a constant), we can of course still give local results. As an example, if \mathcal{A} is the FDFW or the SOR, then by applying Algorithm 2 to the trust region subproblem and using the KL property proved in [23] we can obtain local converge rates around global minimizers. Specifically, we obtain the $O(1/k^2)$ rate given in Corollary 5.1 for $\theta = 1/4$ in the (rather technical) critical case condition [23, equation (24)], and the linear rate given in Corollary 5.1 for $\theta = 1/2$ otherwise.

7 Conclusions

Projection-free first-order methods rely on the choice of good feasible descent directions, for which there needs to be a trade-off between slope and maximal stepsize. To adress this issue we proposed the SSC procedure, which allowed us to give a unified analysis for non convex objectives of several methods satisfying a sufficient slope condition. With this analysis we proved convergence rates independent from the number of short steps, and showed how our SSC procedure easily adapts to different choices of descent directions. Finally, our work is the first to analyze projection-free methods on the class of smooth non convex objectives satisfying KL inequalities.

Future research directions include generalizing our framework to constrained stochastic optimization, as well as applications for the solution of real-world data science problems.

8 Appendix

8.1 Proofs for some elementary properties

We need the following result to prove Proposition 2.2:

Proposition 8.1. *Let C be a closed convex cone. For every $y \in \mathbb{R}^n$*

$$\text{dist}(C^*, y) = \sup_{c \in C} \langle \hat{c}, y \rangle.$$

As stated in [13] this is an immediate consequence of the Moreau-Yosida decomposition:

$$y = \pi(C, y) + \pi(C^*, y).$$

Proposition 2.2. We first prove that

$$\sup_{h \in \Omega \setminus \{\bar{x}\}} \left(g, \frac{h - \bar{x}}{\|h - \bar{x}\|} \right) = \sup_{h \in T_{\Omega}(\bar{x}) \setminus \{0\}} (g, \hat{h}). \quad (8.1)$$

Let $h \in T_{\Omega}(\bar{x}) \setminus \{0\}$. Then there exists sequences $\{\lambda_i\}$ and $\{h_i\}$ in $\mathbb{R}_{>0}$ and Ω respectively such that $\lambda_i(h_i - \bar{x}) \rightarrow h$. In particular $\|\lambda_i(h_i - \bar{x})\| \rightarrow \|h\|$ so that we also have $\lambda_i(h_i - \bar{x}) / \|\lambda_i(h_i - \bar{x})\| = (h_i - \bar{x}) / \|h_i - \bar{x}\| \rightarrow \hat{h}$. Hence

$$\text{cl} \left(\left\{ \frac{h - \bar{x}}{\|h - \bar{x}\|} \mid h \in \Omega \setminus \{0\} \right\} \right) = \{ \hat{h} \mid h \in T_{\Omega}(\bar{x}) \setminus \{0\} \},$$

and (8.1) follows immediately by the continuity of (g, \cdot) .

Since $N_\Omega(\bar{x}) = T_\Omega(\bar{x})^*$ the first equality is exactly the one of Proposition 8.1 if $g \notin N_\Omega(\bar{x})$, and it is trivial since both terms are clearly 0 if $g \in N_\Omega(\bar{x})$.

It remains to prove

$$\text{dist}(N_\Omega(\bar{x}), g) = \|\pi(T_\Omega(\bar{x}), g)\|,$$

which is true by the Moreau - Yosida decomposition. \square

In the next lemma we lower bound $\text{SB}_{\text{FW}}(\Omega, \partial\Omega)$ in terms of a relatively more manageable expression in the strictly convex smooth case.

Lemma 8.1. *Let Ω be a strictly convex smooth subset of \mathbb{R}^n , and let $J : \partial\Omega \rightarrow \mathbb{R}^n$ be such that $J(\bar{x}) \in N_\Omega(\bar{x}) \setminus \{0\}$ for every $\bar{x} \in \partial\Omega$ (so that in particular $N_\Omega(\bar{x}) = \text{cone}(J(\bar{x}))$). Then*

$$\text{SB}_{\text{FW}}(\Omega, \partial\Omega) \geq \inf_{\substack{y, \bar{x} \in \partial\Omega \\ \bar{x} \neq y}} \frac{\langle J(y), y - \bar{x} \rangle}{\|J(\bar{x}) - J(y)\| \|y - \bar{x}\|}.$$

Proof. Given $\bar{x} \in \Omega$, $g \in \mathbb{R}^n \setminus N_\Omega(\bar{x})$, by strict convexity $g = \gamma J(y)$ for some $\gamma > 0$ and

$$\{y\} = \text{argmax}_{x \in \Omega} \langle g, x \rangle.$$

In particular, $\text{FW}(\bar{x}, g) = y - \bar{x}$. Therefore

$$\text{DSB}_{\text{FW}}(\Omega, \bar{x}, g) = \frac{\langle g, y - \bar{x} \rangle}{\pi_{\bar{x}}(g) \|y - \bar{x}\|} = \frac{\langle J(y), y - \bar{x} \rangle}{\pi_{\bar{x}}(J(y)) \|y - \bar{x}\|} \geq \frac{\langle J(y), y - \bar{x} \rangle}{\|J(\bar{x}) - J(y)\| \|y - \bar{x}\|},$$

where we used $g = kJ(y)$ in the second equality and $N_\Omega(\bar{x}) = \text{cone}(J(\bar{x}))$ in the inequality. The lemma follows taking the inf in g and \bar{x} , where for g varying in $\mathbb{R}^n \setminus N_\Omega(\bar{x})$ the point y varies in all $\partial\Omega \setminus \bar{x}$ by strict convexity. \square

We now prove a lower semicontinuity property for the shrinking coefficient $\bar{\nu}$ introduced for orthographic projections on convex sets with smooth boundary in Section 6.3.

Lemma 8.2. *Let Ω be a full dimensional convex set with C^1 regular boundary, $\{x_j\} \subset \partial\Omega$ be convergent to x . Let $g \in \mathbb{R}^n$ such that $\{(x_j, g)\} \subset \bar{T}(\Omega)$ and $(x, g) \in \bar{T}(\Omega)$. Then*

$$\liminf \bar{\nu}(x_j, g) \geq \bar{\nu}(x, g).$$

Proof. Fix $\lambda \in (0, 1)$, and let $v = P(x, \bar{\nu}(x, g)g_\Omega(x)) - x$, $v_\lambda = P(x, \lambda\bar{\nu}(x, g)g_\Omega(x)) - x$. Recall from the proof of Proposition 6.5 that $x + v_\lambda$ is in $\text{conv}(x + \lambda\bar{\nu}(x, g)g_\Omega(x), x + \lambda v)$ minus $x + \lambda v \in \text{int}(\Omega)$. Therefore, for some $\varepsilon_\lambda > 0$

$$x + v_\lambda + \gamma J(x) \in \begin{cases} \text{int}(\Omega) & \text{if } \gamma \in (0, \varepsilon_\lambda), \\ \Omega^c & \text{if } \gamma < 0. \end{cases} \quad (8.2)$$

Since $J(x_j) \rightarrow J(x)$, $g_\Omega(x_j) \rightarrow g_\Omega(x)$ by smoothness, thanks to (8.2) we have

$$P(x_j, \lambda\bar{\nu}(x, g)g_\Omega(x_j)) \rightarrow P(x, \lambda\bar{\nu}(x, g)g_\Omega(x)) = x + v_\lambda.$$

Thus

$$\frac{\langle P(x_j, \lambda\bar{\nu}(x, g)g_\Omega(x_j)) - x_j, g \rangle}{\pi_{x_j}(g) \|P(x_j, \lambda\bar{\nu}(x, g)g_\Omega(x_j)) - x_j\|} \rightarrow \frac{\langle v_\lambda, g \rangle}{\pi_x(g) \|v_\lambda\|},$$

and for j large enough by (6.21) the LHS must be greater than $\bar{\nu}$. In particular, $\bar{\nu}(x_j, g_\Omega(x_j)) \geq \lambda\bar{\nu}(x, g)$. The thesis follows since $\lambda \in (0, 1)$ is arbitrary. \square

We conclude this section by giving a practical example of lower bound for the shrinking coefficient $\bar{\nu}$ introduced in Section 6.3.

Proposition 8.2. *Let h be a uniformly smooth convex function with minimum h^* , compact set of minimizers X and with L_h -Lipschitz gradient. If given $a > h^*$ the sublevel set $\Omega = \{x \mid h(x) \leq a\}$ is compact, and $\bar{\tau} = 1/2$ then*

$$\bar{\nu}(x, g) \geq \frac{\|\nabla h(x)\|}{\|g\|L_h} \quad (8.3)$$

for every $(x, g) \in \bar{T}(\Omega)$.

Proof. Fix $\bar{x} \in \partial\Omega$ so that $h(\bar{x}) = a$, and let $\bar{h} = \|\nabla h(\bar{x})\|/L$. By the standard descent Lemma for any $y \in \mathbb{R}^n$

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

and in particular $\bar{B}(\bar{x}) = \text{cl}(B_{\bar{x} - \nabla h(\bar{x})/L}(\bar{h})) \subset \Omega$. Let now $P_{\bar{B}(\bar{x})}$ be the orthographic retraction on $\bar{B}(\bar{x})$, and for g such that $(\bar{x}, g) \in \bar{T}(\Omega)$, let $\bar{g} = \|g\|$.

In order to prove (8.3) it suffices to prove

$$\frac{\langle g, v_g \rangle}{\pi_{\bar{x}}(g)\|v_g\|} \geq \frac{1}{2}. \quad (8.4)$$

for $v_g = P(\bar{x}, \frac{\bar{h}}{\bar{g}}g_\Omega(\bar{x})) - \bar{x} = P(\bar{x}, \bar{h}\hat{g}_\Omega(\bar{x})) - \bar{x}$. First, we have

$$\frac{\langle g, v_g \rangle}{\pi_{\bar{x}}(g)\|v_g\|} \geq \frac{\langle g, \bar{v}_g \rangle}{\pi_{\bar{x}}(g)\|\bar{v}_g\|} \quad (8.5)$$

for $\bar{v}_g = P_{\bar{B}(\bar{x})}(\bar{x}, \bar{h}\hat{g}_\Omega(\bar{x})) - \bar{x}$, given that $\bar{B}(\bar{x}) \subset \Omega$. The RHS of (8.5) can be bounded with a tedious but straightforward computation. Let $\bar{g}_\Omega = \|\hat{g}_\Omega(\bar{x})\|$ and $\bar{g}_J = \|\hat{g}_J(\bar{x})\|$, so that $\pi_x(g) = \bar{g}_\Omega\bar{g}$, $\bar{g}_\Omega^2 + \bar{g}_J^2 = 1$, $\bar{g}_\Omega > 0$. Of course we can explicitly compute the orthographic projection on $\bar{B}(\bar{x})$ and in particular we have $\bar{v}_g = \bar{h}\hat{g}_\Omega(\bar{x}) - \gamma(\bar{h}\bar{g}_\Omega)J(\bar{x})$, for $\gamma(r) = \bar{h} - \sqrt{\bar{h}^2 - r^2}$. Then $\langle g, \bar{v}_g \rangle = \bar{g}_\Omega\bar{g}\bar{h} - \bar{g}_J\bar{g}\gamma(\bar{h}\bar{g}_\Omega)$, $\pi_{\bar{x}}(g) = \bar{g}_\Omega\bar{g}$, and $\|\bar{v}_g\| = \sqrt{\bar{h}^2\bar{g}_\Omega^2 + \gamma(\bar{h}\bar{g}_\Omega)^2}$. Therefore

$$\frac{\langle g, \bar{v}_g \rangle}{\pi_{\bar{x}}(g)\|\bar{v}_g\|} = \frac{\bar{g}_\Omega\bar{h} - \bar{g}_J\gamma(\bar{g}_\Omega\bar{h})}{\bar{g}_\Omega\sqrt{\bar{h}^2\bar{g}_\Omega^2 + \gamma(\bar{h}\bar{g}_\Omega)^2}},$$

and we can finally write

$$\inf_{g: (\bar{x}, g) \in \bar{T}(\Omega)} \frac{\langle g, \bar{v}_g \rangle}{\pi_{\bar{x}}(g)\|\bar{v}_g\|} = \inf \left\{ \frac{\bar{g}_\Omega\bar{h} - \bar{g}_J\gamma(\bar{g}_\Omega\bar{h})}{\bar{g}_\Omega\sqrt{\bar{h}^2\bar{g}_\Omega^2 + \gamma(\bar{h}\bar{g}_\Omega)^2}} \mid \bar{g}_\Omega > 0, \bar{g}_J \geq 0, \bar{g}_J^2 + \bar{g}_\Omega^2 = 1 \right\} = \frac{1}{2},$$

which combined with (8.5) proves (8.4). \square

Remark 8.1. The above proposition can be generalized for any function h with a convex increasing smoothness modulus $\bar{\eta}$ (see e.g. [6] for some related properties) such that

$$h(y) \leq h(x) + \langle \nabla h(x), y - x \rangle + \bar{\eta}(\|x - y\|)$$

for every $x, y \in \mathbb{R}^n$. In this case $\bar{\nu}(x, \hat{g})$ can be lower bounded by a function of $\bar{\eta}/\|\nabla h(x)\|$ and $\|\hat{g}_\Omega(x)\|$, decreasing in both arguments. In particular when h is q -uniformly smooth with $\bar{\eta}(t) \approx t^q$ for some $q \in (1, 2]$ we have

$$\bar{\nu}(x, \hat{g}) \gtrsim \|\hat{g}_\Omega(x)\|^{\frac{2-q}{q-1}}$$

for $\hat{g}_\Omega(x) \rightarrow 0$ and uniformly in x using that by compactness $\|\nabla h(x)\|$ is bounded away from 0 in $\partial\Omega$.

8.2 Fréchet subdifferential and KL property

In this section we report the definition of the Kurdyka-Łojasiewicz property as it is presented in [3], [4]. We first need to recall a few additional definitions from variational analysis, the standard references being [15], [37] and [41].

Definition 8.1. Given a lower semicontinuous function $\tilde{f} : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ with $\text{dom}(\tilde{f}) \neq \emptyset$ the vector z is said to be in the Fréchet subdifferential of \tilde{f} at x , written $\hat{\partial}\tilde{f}(x)$, iff

$$\tilde{f}(y) \geq \tilde{f}(x) + \langle z, y - x \rangle + o(\|y - x\|) \quad \text{for all } y \in \mathbb{R}^n.$$

The limiting subdifferential of \tilde{f} at $x \in \text{dom}(\tilde{f})$, written $\partial\tilde{f}(x)$, is the set

$$\partial\tilde{f}(x) := \{z \in \mathbb{R}^n : \exists x_n \rightarrow x, \tilde{f}(x_n) \rightarrow \tilde{f}(x), z_n \in \hat{\partial}\tilde{f}(x_n) \rightarrow z\}.$$

It is clear from the definition that the limiting subdifferential is always closed.

We can now define the KL property.

Definition 8.2. A function \tilde{f} as above is said to have the KL property at $x^* \in \text{dom}(\partial\tilde{f})$ if there exists $\eta \in (0, +\infty)$, a neighborhood U of x^* and a continuous concave function $\varphi : [0, \eta] \rightarrow \mathbb{R}_{\geq 0}$ (called *desingularizing function*) such that:

- (i) $\varphi(0) = 0$,
- (ii) φ is C^1 on $(0, \eta)$,
- (iii) for all $t \in (0, \eta)$, $\varphi'(t) > 0$,
- (iv) for all $x \in U \cap [\tilde{f}(x^*) < \tilde{f} < \tilde{f}(x^*) + \eta]$, the KL inequality holds

$$\varphi'(\tilde{f}(x) - \tilde{f}(x^*)) \text{dist}(\partial\tilde{f}(x), 0) \geq 1. \quad (8.6)$$

When $\tilde{f} = f + i_\Omega$ with $f \in C^1(\Omega)$, Ω convex and closed, the KL inequality (8.6) can be rewritten as

$$\varphi'(f(x) - f(x^*)) \|\pi(T_\Omega(x), -\nabla f(x))\| \geq 1. \quad (8.7)$$

Indeed, by the assumptions on Ω (closedness and convexity), we can write

$$\partial i_\Omega(x) = N_\Omega(x) \quad \forall x \in \Omega,$$

and we have the sum rule (by e.g. [38, Proposition 2.2])

$$\partial\tilde{f}_\Omega(x) = \partial i_\Omega(x) + \nabla f(x) = N_\Omega(x) + \nabla f(x).$$

Finally,

$$\text{dist}(\partial\tilde{f}_\Omega(x), 0) = \text{dist}(N_\Omega(x) + \nabla f(x), 0) = \text{dist}(N_\Omega(x), -\nabla f(x)) = \|\pi(T_\Omega(x), -\nabla f(x))\|, \quad (8.8)$$

where we used Proposition 2.2 in the last equality and the equivalence is proved.

The study of bounds on the desingularizing function φ for smooth functions constrained to convex sets is still very much an open area of research. In [34] it was proved that a KL property with desingularizing function of the form $\varphi(t) = Mt^{1/2}$ is implied by the Luo-Tseng error bound. For other examples see Section 6.5.

Remark 8.2. Consider the LPL condition (see [7], Definition 1 for the details)

$$\|\pi(T_x\Omega, -\nabla f(x))\|^\alpha \geq \mu(f(x) - f^*),$$

with $f^* = \min_{x \in \Omega} f(x)$, $T_x\Omega$ the tangent plane to Ω in x equal to $T_\Omega(x)$, f smooth and Ω a C^1 surface. For $\alpha > 1$ this condition is a particular case of the KL condition. Indeed the KL inequality can still be written in the form (8.7) when Ω is a surface, with $T_\Omega(x) = T_x\Omega$, as it can be proved applying the same reasoning we have seen for convex sets. Then if $\alpha > 1$ we retrieve the LPL condition considering $\varphi(t) = \mu^{-\frac{1}{\alpha}} \frac{\alpha}{\alpha-1} t^{1-1/\alpha}$ and any $x \in \text{argmin}_{x \in \Omega} f(x)$ in (8.7).

8.3 Proofs for Section 5

In this section we prove Proposition 5.1 and Corollary 5.1. We start by recalling Karamata's inequality for concave functions. Given $A, B \in \mathbb{R}^N$ it is said that A majorizes B , written $A \succ B$, if

$$\begin{aligned} \sum_{i=1}^j A_i &\geq \sum_{i=1}^j B_i \text{ for } j \in [1 : N], \\ \sum_{i=1}^N A_i &= \sum_{i=1}^N B_i. \end{aligned}$$

If h is concave and $A \succ B$ by Karamata's inequality

$$\sum_{i=1}^N h(A_i) \leq \sum_{i=1}^N h(B_i).$$

We now prove a technical lemma. It gives an upper bound on a sum of square roots which is the key to upper bound the length of the tail of the sequence under the descent conditions (H1), (H'2a).

Lemma 8.3. *Let $f_0 \in [0, \eta]$ and assume*

$$I = \int_0^\eta \alpha(t) dx < +\infty.$$

Then

$$\sum_{s=0}^{\infty} \sqrt{\sigma_\alpha^{(s)}(f_0) - \sigma_\alpha^{(s+1)}(f_0)} \leq I + 2\sqrt{f_0 - \sigma_\alpha(f_0)}. \quad (8.9)$$

Proof. Let $S = \inf\{s \in \mathbb{N}_0 \mid \sigma_\alpha^{(s)}(f_0) = 0\}$. If $S \in \{0, 1\}$ then (8.9) is trivially satisfied. Otherwise for every $0 \leq s \leq S-1$ we have

$$\sqrt{\sigma_\alpha^{(s)}(f_0) - \sigma_\alpha^{(s+1)}(f_0)} \leq \frac{1}{\alpha(\sigma_\alpha^{(s+1)}(f_0))} \quad (8.10)$$

thanks to (5.2) and (5.1), with equality for $0 \leq s \leq S-2$.

Let $2 \leq P \leq S$. Then by (8.10) for $0 \leq s \leq S-2$

$$\sum_{s=0}^{P-2} \sqrt{\sigma_\alpha^{(s)}(f_0) - \sigma_\alpha^{(s+1)}(f_0)} = \sum_{s=0}^{P-2} \frac{1}{\alpha(\sigma_\alpha^{(s+1)}(f_0))}. \quad (8.11)$$

From now on for simplicity we write $\bar{\alpha}^{(s)}(f_0)$ to denote $\alpha(\sigma_\alpha^{(s)}(f_0))$.

We have

$$I = \int_0^\eta \alpha(t) dx \geq \sum_{s=0}^{P-2} \bar{\alpha}^{(s)}(f_0) (\sigma_\alpha^{(s)}(f_0) - \sigma_\alpha^{(s+1)}(f_0)) = \sum_{s=0}^{P-2} \frac{\bar{\alpha}^{(s)}(f_0)}{\bar{\alpha}^{(s+1)}(f_0)^2}, \quad (8.12)$$

where in the first inequality the RHS is a Riemann sum and we used that α is decreasing, and in the second equality we used again (5.1).

We now have

$$\begin{aligned} &\sum_{s=0}^{P-2} \frac{1}{\bar{\alpha}^{(s+1)}(f_0)} - \sum_{s=0}^{P-2} \frac{\bar{\alpha}^{(s)}(f_0)}{\bar{\alpha}^{(s+1)}(f_0)^2} \\ &= \frac{1}{\bar{\alpha}^{(P-1)}(f_0)} - \frac{\alpha(f_0)}{\bar{\alpha}(f_0)^2} + \sum_{s=1}^{P-2} \bar{\alpha}^{(s)}(f_0) \left(\frac{1}{\bar{\alpha}^{(s)}(f_0)^2} - \frac{1}{\bar{\alpha}^{(s+1)}(f_0)^2} \right). \end{aligned} \quad (8.13)$$

The sum on the RHS turns out to be easily lower bounded by a telescopic sum

$$\begin{aligned}
& \sum_{s=1}^{P-2} \bar{\alpha}^{(s)}(f_0) \left(\frac{1}{\bar{\alpha}^{(s)}(f_0)^2} - \frac{1}{\bar{\alpha}^{(s+1)}(f_0)^2} \right) \\
&= \sum_{s=1}^{P-2} \bar{\alpha}^{(s)}(f_0) \left(\frac{1}{\bar{\alpha}^{(s)}(f_0)} - \frac{1}{\bar{\alpha}^{(s+1)}(f_0)} \right) \left(\frac{1}{\bar{\alpha}^{(s)}(f_0)} + \frac{1}{\bar{\alpha}^{(s+1)}(f_0)} \right) \\
&\leq 2 \sum_{s=1}^{P-2} \left(\frac{1}{\bar{\alpha}^{(s)}(f_0)} - \frac{1}{\bar{\alpha}^{(s+1)}(f_0)} \right) = \frac{2}{\bar{\alpha}(f_0)} - \frac{2}{\bar{\alpha}^{(P-1)}(f_0)},
\end{aligned} \tag{8.14}$$

where in the inequality we used

$$\bar{\alpha}^{(s)}(f_0) \left(\frac{1}{\bar{\alpha}^{(s)}(f_0)} + \frac{1}{\bar{\alpha}^{(s+1)}(f_0)} \right) \leq 2.$$

We therefore have

$$\begin{aligned}
& \sum_{s=0}^{P-2} \frac{1}{\bar{\alpha}^{(s+1)}(f_0)} - I \leq \sum_{s=0}^{P-2} \frac{1}{\bar{\alpha}^{(s+1)}(f_0)} - \sum_{s=0}^{P-2} \frac{\bar{\alpha}^{(s)}(f_0)}{\bar{\alpha}^{(s+1)}(f_0)^2} \\
&\leq \frac{1}{\bar{\alpha}^{(P-1)}(f_0)} - \frac{\alpha(f_0)}{\bar{\alpha}(f_0)^2} + \frac{2}{\bar{\alpha}(f_0)} - \frac{2}{\bar{\alpha}^{(P-1)}(f_0)} \\
&\leq \frac{2}{\bar{\alpha}(f_0)} - \frac{1}{\bar{\alpha}^{(P-1)}(f_0)},
\end{aligned} \tag{8.15}$$

where we used (8.12) in the first inequality, (8.13) together with (8.14) in the second. Then if $S = +\infty$ taking the limit for $P \rightarrow +\infty$ in (8.15) we obtain

$$\sum_{s=0}^{\infty} \frac{1}{\bar{\alpha}^{(s+1)}(f_0)} - I \leq \frac{2}{\bar{\alpha}(f_0)}, \tag{8.16}$$

and the thesis follows from (8.11) with $P = +\infty$.

If $S < \infty$ then for $P = S$ we obtain

$$\sum_{s=0}^{S-1} \frac{1}{\bar{\alpha}^{(s+1)}(f_0)} - I = \frac{1}{\bar{\alpha}^{(S)}(f_0)} + \sum_{s=0}^{P-2} \frac{1}{\bar{\alpha}^{(s+1)}(f_0)} - I \leq \frac{2}{\bar{\alpha}(f_0)} + \frac{1}{\bar{\alpha}^{(S)}(f_0)} - \frac{1}{\bar{\alpha}^{(S-1)}(f_0)} \leq \frac{2}{\bar{\alpha}(f_0)}, \tag{8.17}$$

where we used (8.15) in the first inequality and that $1/\bar{\alpha}^{(s)}(f_0)$ is decreasing in s in the last inequality.

We can now write

$$\sum_{s=0}^{\infty} \sqrt{\sigma_{\alpha}^{(s)}(f_0) - \sigma_{\alpha}^{(s+1)}(f_0)} = \sum_{s=0}^{S-1} \sqrt{\sigma_{\alpha}^{(s)}(f_0) - \sigma_{\alpha}^{(s+1)}(f_0)} \leq \sum_{s=0}^{S-1} \frac{1}{\bar{\alpha}^{(s+1)}(f_0)} \leq I + \frac{2}{\bar{\alpha}(f_0)}, \tag{8.18}$$

where we used $\sigma^{(s)} = 0$ for $s \geq S$ in the equality, (8.10) in the first inequality, (8.17) and (8.16) in the second for the cases $S < +\infty$ and $S = +\infty$ respectively. By the assumption $S \geq 2$ we have

$$2/\bar{\alpha}(f_0) = 2\sqrt{f_0 - \sigma_{\alpha}(f_0)}$$

by (8.10) and applying this inequality to the RHS of (8.18) the thesis is proved. \square

To prove Proposition 5.1 we first need another technical lemma where we use Karamata's inequality in combination with condition (H1) to upper bound the length of the tails of $\{x_k\}$.

Lemma 8.4. Let \tilde{f} be as in Proposition 5.1, $\{x_k\}_{k \in I}$ be a sequence indexed by $I = [0, m] \cap \mathbb{N}_0$ for some $m \in \mathbb{N} \cup \{\infty\}$. Assume that $\{x_k\}_{k \in I}$ satisfies conditions (H1) and (H'2) for every $k \in I$ and that \tilde{f} satisfies the KL inequality with respect to x^* in $\{\tilde{x}_k\}_{k \in I \setminus \{m\}} \cap [\tilde{f}(x^*) < \tilde{f} < \eta]$ with $\tilde{f}(x_0) < \eta$. Let $\alpha(t) = \frac{b}{\sqrt{a}}\varphi'(t)$. Then

$$\begin{aligned} \tilde{f}(x_k) - \tilde{f}(x^*) &\leq \sigma_\alpha^{(k)}(\tilde{f}(x_0) - \tilde{f}(x^*)) \\ \sum_{i=k}^{m-1} \|x_i - x_{i+1}\| + \sqrt{\frac{\tilde{f}(x_m) - \tilde{f}(x^*)}{a}} &\leq \frac{b}{a}\varphi(\tilde{f}(x_k) - \tilde{f}(x^*)) + 2\sqrt{\frac{\tilde{f}(x_k) - \tilde{f}(x^*) - \sigma_\alpha(\tilde{f}(x_k) - \tilde{f}(x^*))}{a}}, \end{aligned} \quad (8.19)$$

for all $k \leq m-1$ where we set $\tilde{f}(x_m) = \tilde{f}(x^*)$ for $m = \infty$.

Proof. First observe that $\{\tilde{f}(x_k)\}$ is decreasing by condition (H1), so that for every $k \in I$ we have $\tilde{f}(\tilde{x}_k) \leq \tilde{f}(x_k) < \eta$. Therefore for every $k \in I$ if $\tilde{f}(\tilde{x}_k) > \tilde{f}(x^*)$ then the KL inequality holds in \tilde{x}_k . Let $f_i = \tilde{f}(x_i) - \tilde{f}(x^*) \geq 0$ for $i \in I$, $f_i = 0$ for $i > m$. If $f_{i+1} = 0$ then trivially

$$f_{i+1} = 0 \leq \sigma_\alpha(f_i).$$

If $f_{i+1} > 0$ then

$$\begin{aligned} \tilde{f}(x_i) - \tilde{f}(x_{i+1}) &\geq \frac{a}{b^2} \text{dist}(\partial \tilde{f}(\tilde{x}_{i+1}), 0)^2 \\ &\geq \frac{a}{b^2 \varphi'(\tilde{f}(\tilde{x}_i) - \tilde{f}(x^*))^2} \geq \frac{1}{\alpha(f_{i+1})^2}, \end{aligned} \quad (8.20)$$

where the first inequality is (H'2), the second the KL property which holds in \tilde{x}_i by (4.4) since $\tilde{f}(\tilde{x}_i) \geq \tilde{f}(x_{i+1}) > \tilde{f}(x^*)$, and the third follows from the monotonicity of α . We can rewrite (8.20) as

$$f_{i+1} + \frac{1}{\alpha(f_{i+1})^2} \leq f_i,$$

and from the definition of σ_α also in this case $f_{i+1} \leq \sigma_\alpha(f_i)$. Now by induction we obtain

$$f_i \leq \sigma_\alpha^{(i-l)}(f_l) \quad (8.21)$$

for every $0 \leq l \leq i$, and to prove the first part of (8.19) it suffices to take $k = i$ and $l = 0$.

By the monotonicity of α equation (5.2) together with (5.1) imply that the sequence $\sigma_\alpha^{(i-l)}(f_l) - \sigma_\alpha^{(i-l+1)}(f_l)$ is decreasing in i . Let now $i \geq k+1$, and let

$$n(i, k) = \min\{j \mid f_{j+1} \leq \sigma_\alpha^{(i-k)}(f_k)\}.$$

Then by (8.21) with $k = l$ we have $n(i, k) \leq i-1$. Consider the vectors v_{ik}, w_{ik}^* in $\mathbb{R}_{\geq 0}^{i-k}$ defined by

$$\begin{aligned} v_{ik} &= (f_k - \sigma_\alpha(f_k), \dots, \sigma_\alpha^{(i-k-1)}(f_k) - \sigma_\alpha^{(i-k)}(f_k)), \\ w_{ik}^* &= (f_k - f_{k+1}, \dots, f_{n(i,k)} - \sigma_\alpha^{(i-k)}(f_k), 0, \dots, 0). \end{aligned}$$

so that in particular v_{ik} is in decreasing order.

For $n(i, k) - k + 1 \leq u \leq i - k$ we have

$$\sum_{j=1}^u v_{ik}(j) = f_k - \sigma_\alpha^{(u)}(f_k) \leq f_k - \sigma_\alpha^{(i-k)}(f_k) = \sum_{j=1}^u w_{ik}^*(j), \quad (8.22)$$

where the inequality follows by monotonicity of $\sigma_\alpha^{(u)}$ in u and we have equality if $u = i - k$. For $1 \leq u \leq n(i, k) - k$ we have

$$\sum_{j=1}^u v_{ik}(j) = f_k - \sigma_\alpha^{(u)}(f_k) \leq f_k - f_{u+k} = \sum_{j=1}^u w_{ik}^*(j). \quad (8.23)$$

where the inequality follows by (8.21) with $l = k$ and $u = i - l$. Combining (8.22) and (8.23) we obtain that if w_{ik} is the permutation in decreasing order of w_{ik}^* then $w_{ik} \succ v_{ik}$ so that by Karamata's inequality

$$\sqrt{f_{n(i,k)} - \sigma_\alpha^{(i-k)}(f_k)} + \sum_{j=k+1}^{n(i,k)} \sqrt{f_{j-1} - f_j} = \sum_{j=1}^{i-k} \sqrt{w_{ik}^*(j)} \leq \sum_{j=1}^{i-k} \sqrt{v_{ik}(j)} = \sum_{j=1}^{i-k} \sqrt{\sigma_\alpha^{(j-1)}(f_k) - \sigma_\alpha^{(j)}(f_k)}.$$

Sending i to infinity since $f_{n(i,k)} - \sigma_\alpha^{(i-k)}(f_k) \rightarrow 0$ (because both terms converge to 0) we obtain

$$\sum_{j=k}^{\infty} \sqrt{f_j - f_{j+1}} \leq \sum_{j=0}^{\infty} \sqrt{\sigma_\alpha^{(j)}(f_k) - \sigma_\alpha^{(j+1)}(f_k)}. \quad (8.24)$$

We now apply Lemma 8.3 to the RHS:

$$\begin{aligned} \sum_{j=0}^{\infty} \sqrt{\sigma_\alpha^{(j)}(f_k) - \sigma_\alpha^{(j+1)}(f_k)} &\leq \int_0^{f_k} \alpha(t) dt + 2\sqrt{\tilde{f}(x_k) - \tilde{f}(x^*) - \sigma_\alpha(\tilde{f}(x_k) - \tilde{f}(x^*))} \\ &= \frac{b}{\sqrt{a}} \varphi(\tilde{f}(x_k) - \tilde{f}(x^*)) + 2\sqrt{\tilde{f}(x_k) - \tilde{f}(x^*) - \sigma_\alpha(\tilde{f}(x_k) - \tilde{f}(x^*))}, \end{aligned} \quad (8.25)$$

where we used $\alpha(t) = \frac{b}{\sqrt{a}} \varphi'(t)$ in the equality.

We can now prove the second part of (8.19):

$$\begin{aligned} \sum_{i=k}^{m-1} \|x_i - x_{i+1}\| + \sqrt{\frac{\tilde{f}(x_m) - \tilde{f}(x^*)}{a}} &\leq \frac{1}{\sqrt{a}} \sum_{i=k}^{m-1} \sqrt{f_i - f_{i+1}} + \sqrt{\frac{\tilde{f}(x_m) - \tilde{f}(x^*)}{a}} = \frac{1}{\sqrt{a}} \sum_{i=k}^{\infty} \sqrt{f_i - f_{i+1}} \\ &\leq \frac{b}{a} \varphi(\tilde{f}(x_k) - \tilde{f}(x^*)) + 2\sqrt{\frac{\tilde{f}(x_k) - \tilde{f}(x^*) - \sigma_\alpha(\tilde{f}(x_k) - \tilde{f}(x^*))}{a}}, \end{aligned}$$

where in the first inequality we applied (H1), the equality follows from $f_i = \tilde{f}(x_i) - \tilde{f}(x^*) = 0$ for $i > m$, and the second inequality follows from (8.24) together with (8.25). \square

We can now prove the main result.

Proposition 5.1. We show by induction $\{x_k\}, \{\tilde{x}_k\} \subset B_\delta(x^*)$ with $f(x_k) \geq f(x^*)$ for every k . By (5.6) we have $\|x_0 - x^*\| < \delta$. Now if we know $\{x_k\}_{k \leq u} \subset B_\delta(x^*)$, $\min_{k \leq u} f(x_k) \geq f(x^*)$ and $\{\tilde{x}_k\}_{k \leq u-1} \subset B_\delta(x^*)$ for some $u \in \mathbb{N}_0$ then by condition (5.5) we have $f(x_{u+1}) \geq f(x^*)$. We also have

$$\begin{aligned} \|x_{u+1} - x^*\| &\leq \|x_{u+1} - x_u\| + \|x_u - x^*\| \\ &\leq \sqrt{\frac{f(x_u) - f(x_{u+1})}{a}} + \|x_u - x^*\| \\ &\leq \sqrt{\frac{f(x_u) - f(x_{u+1})}{a}} + \sum_{k=0}^{u-1} \|x_{k+1} - x_k\| + \|x_0 - x^*\| \\ &\leq \sqrt{\frac{f(x_u) - f(x^*)}{a}} + \sum_{k=0}^{u-1} \|x_{k+1} - x_k\| + \|x_0 - x^*\| \\ &\leq \frac{b}{a} \varphi(f(x_0) - f(x^*)) + 2\sqrt{\frac{f(x_0) - f(x^*) - \sigma_\alpha(f(x_0) - f(x^*))}{a}} + \|x_0 - x^*\| < \delta, \end{aligned} \quad (8.26)$$

where we used (H1) in the first inequality, condition (5.5) in the third inequality, Lemma 8.4 in the fourth one and assumption (5.6) in the last one. Finally,

$$\begin{aligned} \|\tilde{x}_u - x^*\| &\leq \|\tilde{x}_u - x_u\| + \|x_u - x^*\| \leq \sqrt{\frac{f(x_u) - f(\tilde{x}_u)}{a}} + \|x_u - x^*\| \\ &\leq \sqrt{\frac{f(x_u) - f(x_{u+1})}{a}} + \|x_u - x^*\| < \delta, \end{aligned}$$

where in the second and in the third inequality we used (4.4), while the last inequality follows from (8.26). This completes the induction. Applying Lemma 8.4 with $m = +\infty$ we obtain (5.7), the second inequality of (5.8) and in particular $f(x_k) \rightarrow f(x^*)$ by (5.3) together with the condition $f(x_k) \geq f(x^*)$ for every $k \in \mathbb{N}_0$. Finally, given that $\{x_k\}$ has finite length it has a limit \tilde{x}^* with $f(\tilde{x}^*) \leq f(x^*)$ by lower semicontinuity, and by the triangular inequality we have that the first part of (5.8) is also satisfied. \square

The proof of Corollary 5.1 given Proposition 5.1 is based on a straightforward computation of the asymptotic behaviour of the worst case sequence for $\alpha(t) = Ut^\epsilon$, with $U > 0$ and $\epsilon \in (-1, -\frac{1}{2}]$. For the sake of completeness we include here a detailed version of the proof.

Corollary 5.1. By Proposition 5.1 we have

$$\tilde{f}(x_k) - \tilde{f}(x^*) \leq \sigma_\alpha^{(k)}(\tilde{f}(x_0) - \tilde{f}(x^*)), \quad (8.27)$$

with $\alpha(t) = Mb\sqrt{\frac{1}{a}}t^{\theta-1}$. Now if $\theta = \frac{1}{2}$ then equation (5.1) becomes

$$t - \sigma_\alpha(t) = \frac{a}{b^2M^2}\sigma_\alpha(t),$$

which has always the solution

$$\sigma_\alpha(t) = t \left(1 + \frac{a}{b^2M^2}\right)^{-1}.$$

The bound (5.9) for $\theta = \frac{1}{2}$ then follows by induction.

In this case we also have

$$\frac{b}{a}\varphi(\tilde{f}(x_k) - \tilde{f}(x^*)) = O\left(\left(1 + \frac{a}{b^2M^2}\right)^{-\frac{k}{2}}\right) \quad (8.28)$$

and

$$2\sqrt{\frac{\tilde{f}(x_k) - \tilde{f}(x^*) - \sigma_\alpha(\tilde{f}(x_k) - \tilde{f}(x^*))}{a}} = O(\sqrt{\tilde{f}(x_k) - \tilde{f}(x^*)}) = O\left(\left(1 + \frac{a}{b^2M^2}\right)^{-\frac{k}{2}}\right). \quad (8.29)$$

We can now prove (5.10) for $\theta = \frac{1}{2}$:

$$\begin{aligned} \sum_{i=k}^{\infty} \|x_i - x_{i+1}\| &\leq \frac{b}{a}\varphi(\tilde{f}(x_k) - \tilde{f}(x^*)) + 2\sqrt{\frac{\tilde{f}(x_k) - \tilde{f}(x^*) - \sigma_\alpha(\tilde{f}(x_k) - \tilde{f}(x^*))}{a}} \\ &= O\left(\left(1 + \frac{a}{b^2M^2}\right)^{-\frac{k}{2}}\right), \end{aligned}$$

where we used Corollary 5.2 in the first inequality, and we summed (8.29), (8.28) in the second. Let $0 < \theta < \frac{1}{2}$, $g = \frac{1}{1-2\theta}$, $f_0 = \tilde{f}(x_0) - \tilde{f}(x^*)$ and let

$$P = \max\left(f_0, \left(\frac{2^{r+2}rb^2M^2}{a}\right)^r\right). \quad (8.30)$$

We now prove by induction

$$\sigma_\alpha^{(k)}(f_0) \leq \frac{P}{(k+1)^r}, \quad (8.31)$$

so that the bound (5.9) follows by

$$\tilde{f}(x_k) - \tilde{f}(x^*) \leq \sigma_\alpha^{(k)}(f_0) \leq \frac{P}{(k+1)^r},$$

where we used (8.27) in the first inequality and (8.31) in the second.

For $k = 0$ (8.31) is true by definition of P . Assume now $\sigma_\alpha^{(k)}(f_0) = \frac{P}{J^r}$ for some $J \geq k + 1$. Then on the one hand by (5.1)

$$0 = \sigma_\alpha^{(k)}(f_0) - \sigma_\alpha^{(k+1)}(f_0) - \frac{1}{\alpha(\sigma_\alpha^{(k+1)}(f_0))^2} = \frac{P}{J^r} - \sigma_\alpha\left(\frac{P}{J^r}\right) - \frac{1}{\alpha(\sigma_\alpha(\frac{P}{J^r}))^2}, \quad (8.32)$$

while on the other hand

$$\frac{P}{J^r} - \frac{P}{(J+1)^r} - \frac{1}{\alpha(\frac{P}{(J+1)^r})^2} = \frac{P}{J^r} - \frac{P}{(J+1)^r} - \frac{aP^{2-2\theta}}{b^2M^2(J+1)^{r+1}} < \frac{rP}{J^{r+1}} - \frac{aP^{2-2\theta}}{b^2M^2(J+1)^{r+1}} \leq 0, \quad (8.33)$$

where the last inequality follows from the definition (8.30) of P . Now the function

$$q(s) = \frac{P}{J^r} - s - \frac{1}{\alpha(s)^2}$$

is decreasing in s and $q(\frac{P}{(J+1)^r}) < q(\sigma_\alpha(\frac{P}{J^r}))$ by (8.32) and (8.33). It follows

$$\sigma_\alpha^{(k+1)}(f_0) = \sigma_\alpha(\sigma_\alpha^{(k)}(f_0)) = \sigma_\alpha\left(\frac{P}{J^r}\right) < \frac{P}{(J+1)^r} \leq \frac{P}{(k+2)^r},$$

where we used $J \geq k + 1$ in the last inequality and the induction is complete.

Having established $\tilde{f}(x_k) - \tilde{f}(x^*) = O(1/k^{\frac{1}{1-2\theta}})$, we have

$$\frac{b}{a}\varphi(\tilde{f}(x_k) - \tilde{f}(x^*)) = O(1/k^{\frac{\theta}{1-2\theta}}).$$

Whence

$$\begin{aligned} & 2\sqrt{\frac{\tilde{f}(x_k) - \tilde{f}(x^*) - \sigma_\alpha(\tilde{f}(x_k) - \tilde{f}(x^*))}{a}} = O\left(\frac{1}{\alpha(\sigma_\alpha(\tilde{f}(x_k) - \tilde{f}(x^*)))}\right) \\ & = O\left(\frac{1}{\alpha(\tilde{f}(x_k) - \tilde{f}(x^*))}\right) = O(1/k^{\frac{1-\theta}{1-2\theta}}) = O(1/k^{\frac{\theta}{1-2\theta}}), \end{aligned}$$

where we used (5.1) in the first equality, $\sigma_\alpha(\tilde{f}(x_k) - \tilde{f}(x^*)) < \tilde{f}(x_k) - \tilde{f}(x^*)$ in the second and $1 - \theta > \theta$ in the last one.

Reasoning exactly as for $\theta = \frac{1}{2}$ we obtain the bound (5.10) for $\theta \in (0, \frac{1}{2})$.

Finally, in the above proof the dependence from f_0 of the implicit constants is always monotone increasing, so that it can be eliminated by replacing $f_0 = \tilde{f}(x_0) - \tilde{f}(x^*)$ with $\eta \geq f_0$ in the bounds (8.27) and (8.31). □

References

- [1] P-A Absil and Jérôme Malick. Projection-like retractions on matrix manifolds. *SIAM Journal on Optimization*, 22(1):135–158, 2012.
- [2] P-A Absil and Ivan V Oseledets. Low-rank retractions: a survey and new results. *Computational Optimization and Applications*, 62(1):5–29, 2015.
- [3] Hédÿ Attouch, Jérôme Bolte, Patrick Redont, and Antoine Soubeyran. Proximal alternating minimization and projection methods for nonconvex problems: An approach based on the Kurdyka-Lojasiewicz inequality. *Mathematics of Operations Research*, 35(2):438–457, 2010.
- [4] Hedy Attouch, Jérôme Bolte, and Benar Fux Svaiter. Convergence of descent methods for semi-algebraic and tame problems: proximal algorithms, forward–backward splitting, and regularized Gauss–Seidel methods. *Mathematical Programming*, 137(1-2):91–129, 2013.
- [5] Jean-Pierre Aubin and Hélène Frankowska. *Set-Valued Analysis*. Birkhauser Boston, 2009.
- [6] Dominique Aze and Jean-Paul Penot. Uniformly convex and uniformly smooth convex functions. *Annales de la Faculté des sciences de Toulouse : Mathématiques*, Ser. 6, 4(4):705–730, 1995.
- [7] M. V. Balashov, B. T. Polyak, and A. A. Tremba. Gradient projection and conditional gradient methods for constrained nonconvex minimization. *Numerical Functional Analysis and Optimization*, 41(7):822–849, 2020.
- [8] Jérôme Bolte, Aris Daniilidis, Adrian Lewis, and Masahiro Shiota. Clarke subgradients of stratifiable functions. *SIAM Journal on Optimization*, 18(2):556–572, 2007.
- [9] Jérôme Bolte, Aris Daniilidis, Olivier Ley, and Laurent Mazet. Characterizations of lojasiewicz inequalities: subgradient flows, talweg, convexity. *Transactions of the American Mathematical Society*, 362(6):3319–3363, 2010.
- [10] Jérôme Bolte, Trong Phong Nguyen, Juan Peypouquet, and Bruce W Suter. From error bounds to the complexity of first-order descent methods for convex functions. *Mathematical Programming*, 165(2):471–507, 2017.
- [11] Gábor Braun, Sebastian Pokutta, and Daniel Zink. Lazifying conditional gradient algorithms. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML’17, pages 566–575. JMLR.org, 2017.
- [12] Gábor Braun, Sebastian Pokutta, Dan Tu, and Stephen Wright. Blended conditional gradients. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 735–743, Long Beach, California, USA, 09–15 Jun 2019. PMLR.
- [13] James V Burke and Jorge J Moré. On the identification of active constraints. *SIAM Journal on Numerical Analysis*, 25(5):1197–1211, 1988.
- [14] Michael D Canon and Clifton D Cullum. A tight upper bound on the rate of convergence of Frank-Wolfe algorithm. *SIAM Journal on Control*, 6(4):509–516, 1968.
- [15] Francis H Clarke, RJ Stern, and PR Wolenski. Proximal smoothness and the lower-C2 property. *J. Convex Anal*, 2(1-2):117–144, 1995.
- [16] M. Forti, P. Nistri, and M. Quincampoix. Convergence of neural networks for programming problems via a nonsmooth lojasiewicz inequality. *IEEE Transactions on Neural Networks*, 17(6):1471–1486, 2006.

- [17] Marguerite Frank and Philip Wolfe. An algorithm for quadratic programming. *Naval research logistics quarterly*, 3(1-2):95–110, 1956.
- [18] Robert M Freund, Paul Grigas, and Rahul Mazumder. An extended Frank–Wolfe method with in-face directions, and its application to low-rank matrix completion. *SIAM Journal on Optimization*, 27(1):319–346, 2017.
- [19] Dan Garber and Elad Hazan. Faster rates for the Frank-Wolfe method over strongly-convex sets. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37*, ICML’15, pages 541–549. JMLR.org, 2015.
- [20] Jacques Guelat and Patrice Marcotte. Some comments on Wolfe’s away step. *Mathematical Programming*, 35(1):110–119, 1986.
- [21] Elad Hazan and Haipeng Luo. Variance-reduced and projection-free stochastic optimization. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1263–1271, New York, New York, USA, 20–22 Jun 2016. PMLR.
- [22] Martin Jaggi. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28*, ICML’13, pages I–427–I–435. JMLR.org, 2013.
- [23] Rujun Jiang and Xudong Li. Holderian error bounds and Kurdyka-łojasiewicz inequality for the trust region subproblem. *arXiv preprint arXiv:1911.11955*, 2019.
- [24] Armand Joulin, Kevin Tang, and Li Fei-Fei. Efficient image and video co-localization with Frank-Wolfe algorithm. In *European Conference on Computer Vision*, pages 253–268. Springer, 2014.
- [25] Zoran Kadelburg, Dusan Dukic, Milivoje Lukic, and Ivan Matic. Inequalities of Karamata, Schur and Muirhead, and some applications. *The Teaching of Mathematics*, 8(1):31–45, 2005.
- [26] Tetsuya Kaneko, Simone Fiori, and Toshihisa Tanaka. Empirical arithmetic averaging over the compact stiefel manifold. *IEEE Transactions on Signal Processing*, 61(4):883–894, 2012.
- [27] Jovan Karamata. Sur une inégalité relative aux fonctions convexes. *Publications de l’Institut Mathématique*, 1(1):145–147, 1932.
- [28] Thomas Kerdreux, Alexandre d’Aspremont, and Sebastian Pokutta. Restarting Frank-Wolfe. In Kamalika Chaudhuri and Masashi Sugiyama, editors, *Proc. Mach. Learn. Res. (PMLR)*, volume 89 of *Proceedings of Machine Learning Research*, pages 1275–1283. PMLR, 16–18 Apr 2019.
- [29] Thomas Kerdreux, Alexandre d’Aspremont, and Sebastian Pokutta. Projection-free optimization on uniformly convex sets. *arXiv preprint arXiv:2004.11053*, 2020.
- [30] Simon Lacoste-Julien and Martin Jaggi. On the global linear convergence of Frank-Wolfe optimization variants. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, NIPS’15, pages 496–504, Cambridge, MA, USA, 2015. MIT Press.
- [31] Simon Lacoste-Julien, Martin Jaggi, Mark Schmidt, and Patrick Pletscher. Block-coordinate Frank-Wolfe optimization for structural SVMs. In Sanjoy Dasgupta and David McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 53–61, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR.
- [32] Guanghui Lan and Yi Zhou. Conditional gradient sliding for convex optimization. *SIAM Journal on Optimization*, 26(2):1379–1409, 2016.

- [33] Kfir Levy and Andreas Krause. Projection free online learning over smooth sets. In Kamalika Chaudhuri and Masashi Sugiyama, editors, *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pages 1458–1466. PMLR, 16–18 Apr 2019.
- [34] Guoyin Li and Ting Kei Pong. Calculus of the exponent of Kurdyka–Łojasiewicz inequality and its applications to linear convergence of first-order methods. *Foundations of computational mathematics*, 18(5):1199–1232, 2018.
- [35] Huikang Liu, Weijie Wu, and Anthony Man-Cho So. Quadratic optimization with orthogonality constraints: Łojasiewicz exponent and linear convergence of line-search methods. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ICML’16, pages 1158–1167. JMLR.org, 2016.
- [36] Zhi-Quan Luo and Paul Tseng. Error bounds and convergence analysis of feasible descent methods: a general approach. *Annals of Operations Research*, 46(1):157–178, 1993.
- [37] Boris S Mordukhovich. *Variational analysis and generalized differentiation I: Basic theory*, volume 330. Springer Science & Business Media, 2006.
- [38] Boris S Mordukhovich, Nguyen Mau Nam, and ND Yen. Fréchet subdifferential calculus and optimality conditions in nondifferentiable programming. *Optimization*, 55(5-6):685–708, 2006.
- [39] Anton Osokin, Jean-Baptiste Alayrac, Isabella Lukasewitz, Puneet Dokania, and Simon Lacoste-Julien. Minding the gaps for block Frank-Wolfe optimization of structured SVMs. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 593–602, New York, New York, USA, 20–22 Jun 2016. PMLR.
- [40] Javier Peña and Daniel Rodriguez. Polytope conditioning and linear convergence of the Frank–Wolfe algorithm. *Math. Oper. Res.*, 44(1):1–18, 2018.
- [41] R Tyrrell Rockafellar and Roger J-B Wets. *Variational analysis*, volume 317. Springer Science & Business Media, 2009.
- [42] Yu Wang, Wotao Yin, and Jinshan Zeng. Global convergence of admm in nonconvex nonsmooth optimization. *Journal of Scientific Computing*, 78(1):29–63, 2019.
- [43] Philip Wolfe. Convergence theory in nonlinear programming. *Integer and nonlinear programming*, pages 1–36, 1970.
- [44] Yangyang Xu and Wotao Yin. A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion. *SIAM Journal on imaging sciences*, 6(3):1758–1789, 2013.
- [45] Yi Xu and Tianbao Yang. Frank-Wolfe method is automatically adaptive to error bound condition. *arXiv preprint arXiv:1810.04765*, 2018.
- [46] Teng Zhang and Yi Yang. Robust pca by manifold optimization. *The Journal of Machine Learning Research*, 19(1):3101–3139, 2018.