

The condition of a function relative to a polytope

David H. Gutman* Javier F. Peña†

February 1, 2018

Abstract

The *condition number* of a smooth convex function, namely the ratio of its smoothness to strong convexity constants, is closely tied to fundamental properties of the function. In particular, the condition number of a quadratic convex function is precisely the square of the *diameter-to-width* ratio of a canonical ellipsoid associated to the function. Furthermore, the condition number of a function bounds the linear rate of convergence of the gradient descent algorithm for unconstrained minimization.

We propose a condition number of a smooth convex function relative to a *reference polytope*. This relative condition number is defined as the ratio of a *relative smooth constant* to a *relative strong convexity constant* of the function, where both constants are relative to the reference polytope. The relative condition number extends the main properties of the traditional condition number. In particular, we show that the condition number of a quadratic convex function relative to a polytope is precisely the square of the *diameter-to-facial-distance* ratio of a scaled polytope for a canonical scaling induced by the function. Furthermore, we illustrate how the relative condition number of a function bounds the linear rate of convergence of first-order methods for minimization of the function over the polytope.

1 Introduction

Let $f : \mathbb{R}^m \rightarrow \mathbb{R} \cup \{\infty\}$ be a convex function. The *condition number* of f is the ratio $\frac{L_f}{\mu_f}$ where L_f and μ_f are respectively the *smoothness* and *strong convexity* constants of the function f as detailed in Definition 1 below. The condition number $\frac{L_f}{\mu_f}$ is closely tied to a number of fundamental properties of the function f . In the special case when f is a quadratic convex function the condition number has the following geometric insight. Suppose $f(u) = \frac{1}{2} \langle Qu, u \rangle + \langle b, u \rangle$ where Q is symmetric and positive definite. Then the condition number $\frac{L_f}{\mu_f}$ is precisely the square of the aspect ratio of the ellipsoid $Q^{1/2}\mathbb{B} := \{Q^{1/2}u : u \in \mathbb{R}^m, \|u\| \leq 1\}$.

The condition number $\frac{L_f}{\mu_f}$ also bounds the linear convergence rate of the gradient descent algorithm for the unconstrained minimization problem

$$f^* = \min_{u \in \mathbb{R}^m} f(u).$$

*Department of Mathematical Sciences, Carnegie Mellon University, USA, dgutman@andrew.cmu.edu

†Tepper School of Business, Carnegie Mellon University, USA, jfp@andrew.cmu.edu

More precisely, for a suitable choice of step sizes the iterates u_k , $k = 0, 1, \dots$ generated by the gradient descent algorithm satisfy

$$f(u_k) - f^* \leq \left(1 - \frac{\mu_f}{L_f}\right)^k (f(u_0) - f^*).$$

The references [3–5, 8–11], among others, discuss the above type of linear convergence and a number of interesting related developments. In particular, Necoara, Nesterov and Glineur [9] establish linear convergence properties for a wide class of first-order methods under assumptions that are weaker than strong convexity.

Let $f : \mathbb{R}^m \rightarrow \mathbb{R} \cup \{\infty\}$ and $A \in \mathbb{R}^{m \times n}$ be such that $\text{conv}(A) \subseteq \text{dom}(f)$. We propose a relative smoothness constant $L_{f,A}$ and a relative strong convexity constant $\mu_{f,A}$ of the function f relative to the reference polytope $\text{conv}(A)$. See Definition 2 below for details. Our main results highlight the tight connection between the relative constants and geometric features of the polytope $\text{conv}(A)$. In particular, we establish some interesting relationship between the relative smoothness and strong convexity constants and the diameter and *facial distance* of the polytope. The facial distance was introduced by Peña and Rodríguez [12] albeit in a more restricted context. These relationships in turn enables us to show that the relative condition number $\frac{L_{f,A}}{\mu_{f,A}}$ extends some of the main properties of the classical condition number $\frac{L_f}{\mu_f}$. In particular, we provide the following interesting geometric insight on the relative condition number when f is quadratic. Suppose $f(u) = \frac{1}{2} \langle Qu, u \rangle + \langle b, u \rangle$ where Q is symmetric and positive definite. As we detail in Section 2 below, in this case the relative condition number $\frac{L_{f,A}}{\mu_{f,A}}$ is precisely the square of the ratio of the diameter to the facial distance of the polytope $\text{conv}(Q^{1/2}A)$. For a general convex function, we show that the relative condition number $\frac{L_{f,A}}{\mu_{f,A}}$ can be bounded above by the product of the classical condition number $\frac{L_f}{\mu_f}$ and the square of the ratio of the diameter to the facial distance of $\text{conv}(A)$. The latter quantity can be seen as a kind of condition number of the polytope $\text{conv}(A)$.

We also illustrate how the relative condition number bounds the linear convergence rate of first-order methods for the minimization problem

$$f^* = \min_{u \in \text{conv}(A)} f(u). \tag{1}$$

More precisely, we show that the iterates generated by both the Frank-Wolfe algorithm with away steps and a version of the projected gradient algorithm have objective values that converge linearly to f^* with linear convergence rate bounded by $\frac{L_{f,A}}{\mu_{f,A}}$. We should note that the linear convergence of the Frank-Wolfe algorithm with away steps and the projected gradient algorithm, as well as of other first-order methods had been previously established in [1, 6, 9, 11, 12] under various kinds of assumptions. Our approach based on the relative condition number yields a proof of linear convergence for the Frank-Wolfe algorithm with away steps that is significantly shorter and simpler than the ones previously presented in [1, 6, 12]. Our approach also reveals some simple ideas at the root of the proofs of linear convergence properties of these first-order algorithms.

The relative constants $L_{f,A}, \mu_{f,A}$ are defined *globally*. In particular, they do not depend on any specific point in $\text{conv}(A)$. We also consider a version of the *quadratic functional growth constant* $\mu_{f,A}^*$ following the construction of Necoara, Nesterov and Glineur [9, Definition 4]. Unlike $\mu_{f,A}$, the constant $\mu_{f,A}^*$ depends explicitly on the set of minimizers of f on $\text{conv}(A)$. The constant $\mu_{f,A}^*$ can be seen as a refinement of $\mu_{f,A}$. It is always the case that $\mu_{f,A}^*$ is larger, and can be quite a bit larger, than $\mu_{f,A}$. Indeed, we show that for some important classes of non-strongly convex functions the constant $\mu_{f,A}^*$ is positive while $\mu_{f,A}$ may not be. (See Theorem 1.)

Our work draws on and relates to the recent articles [1, 6, 7, 9, 12]. Our construction of $L_{f,A}$ and $\mu_{f,A}$ is in the spirit introduced by Lu, Freund, and Nesterov [7]. Lu et al. [7] extend the concepts of smoothness and strong convexity constants by considering them *relative* to a *reference* function, see [7, Definition 1.1 and 1.2]. Our construction of $L_{f,A}$ and $\mu_{f,A}$ is also related to the *curvature constant* C_f^A and *geometric strong convexity constant* μ_f^A proposed by Lacoste-Julien and Jaggi in [6, Appendix C]. The quadratic functional growth constant, as well as other more restrictive growth constants, were proposed by Necoara, Nesterov, and Glineur [9] to give conditions that ensure the linear convergence of first-order methods. A similar quadratic growth approach was also used by Beck and Shtern [1] to establish the linear convergence of a conditional gradient algorithm with away steps for non-strongly convex functions. In contrast to the approaches in [1, 6, 9, 12], our construction of the relative constants applies to any choice of norm in \mathbb{R}^n . Our results reveal interesting geometric insights when this norm is the ℓ_1 norm. Our construction of the relative constants $L_{f,A}, \mu_{f,A}, \mu_{f,A}^*$ and all of our results concerning them scale appropriately, that is, they scale by λ whenever the objective function f is replaced by $\tilde{f} = \lambda f$ for some constant $\lambda > 0$. In particular, the relative condition number $\frac{L_{f,A}}{\mu_{f,A}}$ and all of our bounds on it are invariant under positive scaling of f .

The remainder of the paper is organized as follows. Section 2, the main section of the paper, presents our central ideas and results, namely the construction of relative smoothness and strong convexity constants and their main properties. Section 3 illustrates how the relative condition number bounds the linear rate of convergence of the Frank-Wolfe algorithm with away steps and of the projected gradient algorithm for problem (1).

2 Condition number relative to a polytope

This section presents the central ideas of this paper. We introduce the relative smoothness and relative strong convexity of a function relative to a polytope and establish their main properties. We will use the following notation. Let $\Delta_{n-1} \subseteq \mathbb{R}^n$ denote the standard simplex, that is, $\Delta_{n-1} := \{x \in \mathbb{R}_+^n : \|x\|_1 = 1\}$. For convenience, we will make the following slight abuse of notation. For $A \in \mathbb{R}^{m \times n}$ we will also write A to denote the set of columns of A . The precise meaning of A will be clear from the context.

For $A \in \mathbb{R}^{m \times n}$ consider the following polytope generated by A

$$\text{conv}(A) := \{Ax : x \in \Delta_{n-1}\}.$$

For $u \in \text{conv}(A)$ let $Z(u) := \{z \in \Delta_{n-1} : Az = u\}$. Suppose \mathbb{R}^n is endowed with a norm

$\|\cdot\|$. For $x \in \Delta_{n-1}$ and $u \in \text{conv}(A)$ define

$$\text{dist}(x, Z(u)) := \min_{z \in Z(u)} \|x - z\|.$$

2.1 The relative Lipschitz and strong convexity constants

To motivate our main construction we first recall the classical notion of smoothness and strong convexity constants.

Definition 1. Suppose \mathbb{R}^m is endowed with a norm $\|\cdot\|$ and $f : \mathbb{R}^m \rightarrow \mathbb{R} \cup \{\infty\}$ is a differentiable convex function.

(a) The function f is L_f -smooth on $S \subseteq \text{dom}(f)$ for some constant $L_f > 0$ if for all $u, v \in S$

$$f(v) \leq f(u) + \langle \nabla f(u), v - u \rangle + \frac{L_f}{2} \cdot \|v - u\|^2.$$

(b) The function f is μ_f -strongly convex on $S \subseteq \text{dom}(f)$ for some constant $\mu_f \geq 0$ if for all $u, v \in S$

$$f(v) \geq f(u) + \langle \nabla f(u), v - u \rangle + \frac{\mu_f}{2} \cdot \|v - u\|^2.$$

Next, we present our main construction.

Definition 2. Let \mathbb{R}^m be endowed with a norm $\|\cdot\|$, $A \in \mathbb{R}^{m \times n}$ have at least two different columns, and $f : \mathbb{R}^m \rightarrow \mathbb{R} \cup \{\infty\}$ be a differentiable convex function such that $\text{conv}(A) \subseteq \text{dom}(f)$.

(a) Define the smoothness constant $L_{f,A}$ of f relative to A as follows

$$L_{f,A} := \sup_{\substack{u \in \text{conv}(A) \\ x \in \Delta_{n-1} \setminus Z(u)}} \frac{2(f(Ax) - f(u) - \langle \nabla f(u), Ax - u \rangle)}{\text{dist}(x, Z(u))^2}.$$

(b) Define the strong convexity constant $\mu_{f,A}$ of f relative to A as follows

$$\mu_{f,A} := \inf_{\substack{u \in \text{conv}(A) \\ x \in \Delta_{n-1} \setminus Z(u)}} \frac{2(f(Ax) - f(u) - \langle \nabla f(u), Ax - u \rangle)}{\text{dist}(x, Z(u))^2}.$$

The relative constants are natural extensions of the classical ones. Observe that a differentiable function $f : \mathbb{R}^m \rightarrow \mathbb{R} \cup \{\infty\}$ is L_f -smooth and μ_f -strongly convex on $S \subseteq \text{dom}(f)$ if and only if for all $u, v \in S$

$$\frac{\mu_f}{2} \cdot \|v - u\|^2 \leq f(v) - f(u) - \langle \nabla f(u), v - u \rangle \leq \frac{L_f}{2} \cdot \|v - u\|^2.$$

Likewise, a differentiable function $f : \mathbb{R}^m \rightarrow \mathbb{R} \cup \{\infty\}$ is $L_{f,A}$ -smooth and $\mu_{f,A}$ -strongly convex on $\text{conv}(A) \subseteq \text{dom}(f)$ for $A \in \mathbb{R}^{m \times n}$ if and only if for all $u \in \text{conv}(A)$ and $x \in \Delta_{n-1}$

$$\frac{\mu_{f,A}}{2} \cdot \text{dist}(x, Z(u))^2 \leq f(Ax) - f(u) - \langle \nabla f(u), Ax - u \rangle \leq \frac{L_{f,A}}{2} \cdot \text{dist}(x, Z(u))^2.$$

The relative smoothness and strong convexity constants $L_{f,A}$ and $\mu_{f,A}$ are closely related to the *curvature constant* C_f^A and *geometric strong convexity constant* μ_f^A proposed by Lacoste-Julien and Jaggi in [6, Appendix C]. However, the construction in [6, Appendix C] follows a fairly different path. In particular, the definition of μ_f^A is tied to some variants of the Frank-Wolfe algorithm for problem (1). By contrast, our construction of $L_{f,A}$ and $\mu_{f,A}$ depends only on the pair (f, A) , applies to any norm in \mathbb{R}^n , and does not depend on any particular algorithm. As we discuss in Section 3, the relative condition number $\frac{L_{f,A}}{\mu_{f,A}}$ bounds the linear rates of convergence of the projected gradient algorithm and of the Frank-Wolfe algorithm with away steps for problem (1). We also note that in the special case when \mathbb{R}^n is endowed with the Euclidean norm ℓ_2 , the strong convexity constant $\mu_{f,A}$ is related to the quadratic gradient growth condition defined in [9].

2.2 Geometric properties of the relative constants

We next present some geometric properties of the constants $L_{f,A}$ and $\mu_{f,A}$. The properties below show that these constants are finite and positive when f is L_f -smooth and μ_f -strongly convex for some $L_f, \mu_f > 0$. The properties below also yield a nice analogy between the relative condition number $\frac{L_{f,A}}{\mu_{f,A}}$ and the usual condition number $\frac{L_f}{\mu_f}$.

Our results rely on the concept of facial distance introduced by Peña and Rodríguez [12]. Let \mathbb{R}^m be endowed with a norm $\|\cdot\|$. For $A \in \mathbb{R}^{m \times n}$ with at least two different columns the *facial distance* $\Phi(A)$ is defined as follows

$$\Phi(A) := \min_{\substack{F \in \text{faces}(\text{conv}(A)) \\ \emptyset \neq F \neq \text{conv}(A)}} \text{dist}(F, \text{conv}(A \setminus F)).$$

Here $\text{faces}(\text{conv}(A))$ denotes the set of faces of $\text{conv}(A)$ and $\text{dist}(F, G) = \inf_{u \in F, v \in G} \|u - v\|$ for nonempty $F, G \subseteq \mathbb{R}^m$. Observe that $\Phi(A) > 0$ for all $A \in \mathbb{R}^{m \times n}$ with at least two different columns.

The following example illustrates the facial distance for two canonical polytopes, namely the standard simplex and the ℓ_1 unit ball.

Example 1. Suppose $m > 1$ and \mathbb{R}^m is endowed with the Euclidean norm $\|\cdot\|_2$.

- (a) For $A = I_m \in \mathbb{R}^{m \times m}$ we have $\text{conv}(A) = \Delta_{m-1}$. In this case $\Phi(A)$ is attained at any face of $\text{conv}(A)$ of dimension $k := \lfloor \frac{m}{2} \rfloor$. In particular for $F = \text{conv}\{e_1, \dots, e_k\}$ we get

$$\begin{aligned} \Phi(A) &= \text{dist}(F, \text{conv}(A \setminus F)) = \left\| \frac{e_1 + \dots + e_k}{k} - \frac{e_{k+1} + \dots + e_m}{m-k} \right\|_2 \\ &= \sqrt{\frac{m}{k(m-k)}} \\ &= \begin{cases} \frac{2}{\sqrt{m}} & \text{if } m \text{ is even} \\ \frac{2}{\sqrt{m-\frac{1}{m}}} & \text{if } m \text{ is odd.} \end{cases} \end{aligned}$$

- (b) For $A = \begin{bmatrix} I_m & -I_m \end{bmatrix}$ we have $\text{conv}(A) = \{u \in \mathbb{R}^m : \|u\|_1 \leq 1\}$. In this case $\Phi(A)$ is attained at any face of $\text{conv}(A)$ of dimension $(m-2)$. In particular, for $F = \text{conv}\{e_1, \dots, e_{m-1}\}$ we get

$$\Phi(A) = \text{dist}(F, \text{conv}(A \setminus F)) = \left\| \frac{e_1 + \dots + e_{m-1}}{m-1} - 0 \right\|_2 = \frac{1}{\sqrt{m-1}}.$$

Some of the results below are stated in terms of the *diameter* of a set defined as follows. For $A \subseteq \mathbb{R}^m$

$$\text{diam}(A) := \sup_{u, v \in A} \|u - v\|.$$

The alternative characterizations of $\text{diam}(A)$ and $\Phi(A)$ in the following proposition provide the crux for the geometric properties of $L_{f,A}$ and $\mu_{f,A}$. We defer the proof of Proposition 1 to Section 2.4 since it relies on Lemma 2 below.

Proposition 1. *Suppose \mathbb{R}^n is endowed with the ℓ_1 norm $\|\cdot\|_1$. Then for all $A \in \mathbb{R}^{m \times n}$ with at least two different columns*

$$\text{diam}(A) = \max_{\substack{u \in \text{conv}(A) \\ x \in \Delta_{n-1} \setminus Z(u)}} \frac{2\|Ax - u\|}{\text{dist}(x, Z(u))}, \quad \Phi(A) = \min_{\substack{u \in \text{conv}(A) \\ x \in \Delta_{n-1} \setminus Z(u)}} \frac{2\|Ax - u\|}{\text{dist}(x, Z(u))}.$$

Corollary 1. *Suppose \mathbb{R}^n and \mathbb{R}^m are respectively endowed with the ℓ_1 norm $\|\cdot\|_1$ and the ℓ_2 norm $\|\cdot\|_2$, and $f(u) = \frac{1}{2} \langle Qu, u \rangle + \langle b, u \rangle$ for some $b \in \mathbb{R}^m$ and $Q \in \mathbb{R}^{m \times m}$ symmetric and positive definite. Then for all $A \in \mathbb{R}^{m \times n}$ with at least two different columns*

$$L_{f,A} = \frac{\text{diam}(Q^{1/2}A)^2}{4}, \quad \mu_{f,A} = \frac{\Phi(Q^{1/2}A)^2}{4}.$$

In particular

$$\frac{L_{f,A}}{\mu_{f,A}} = \frac{\text{diam}(Q^{1/2}A)^2}{\Phi(Q^{1/2}A)^2}. \quad (2)$$

Corollary 1 yields the following analogy between the relative condition number $\frac{L_{f,A}}{\mu_{f,A}}$ and the usual condition number $\frac{L_f}{\mu_f}$ of a strongly convex quadratic function f . Under the assumptions of Corollary 1 it readily follows that

$$L_f = \lambda_{\max}(Q) = \lambda_{\max}(Q^{1/2})^2, \quad \mu_f = \lambda_{\min}(Q) = \lambda_{\min}(Q^{1/2})^2.$$

Furthermore, observe that $2\lambda_{\max}(Q^{1/2})$ and $2\lambda_{\min}(Q^{1/2})$ are respectively the diameter (length of longest principal axis) and the width (length of shortest principal axis) of the ellipsoid $\{Q^{1/2}u : \|u\|_2 \leq 1\} = Q^{1/2}\mathbb{B}$, where $\mathbb{B} := \{u \in \mathbb{R}^m : \|u\|_2 \leq 1\}$. Therefore,

$$L_f = \frac{\text{diam}(Q^{1/2}\mathbb{B})^2}{4}, \quad \mu_f = \frac{\text{width}(Q^{1/2}\mathbb{B})^2}{4}.$$

In particular,

$$\frac{L_f}{\mu_f} = \frac{\text{diam}(Q^{1/2}\mathbb{B})^2}{\text{width}(Q^{1/2}\mathbb{B})^2}. \quad (3)$$

Observe the striking resemblance between (2) and (3).

Corollary 2. Let \mathbb{R}^n be endowed with the ℓ_1 norm $\|\cdot\|_1$. Suppose $A \in \mathbb{R}^{m \times n}$ has at least two different columns and $f : \mathbb{R}^m \rightarrow \mathbb{R} \cup \{\infty\}$ is L_f -smooth and μ_f -strongly convex on $\text{conv}(A)$. Then

$$L_{f,A} \leq \frac{L_f \cdot \text{diam}(A)^2}{4}, \quad \mu_{f,A} \geq \frac{\mu_f \cdot \Phi(A)^2}{4}.$$

In particular,

$$\frac{L_{f,A}}{\mu_{f,A}} \leq \frac{L_f}{\mu_f} \cdot \frac{\text{diam}(A)^2}{\Phi(A)^2}.$$

The following proposition gives an identity and bound similar to those in Proposition 1 for the general case when \mathbb{R}^n is endowed with an arbitrary norm. We defer the proof of Proposition 2 to Section 2.4.

Proposition 2. Let \mathbb{R}^n and \mathbb{R}^m be endowed with arbitrary norms and $A \in \mathbb{R}^{m \times n}$ have at least two different columns. Then

$$\max_{\substack{w \in \mathbb{R}^n \setminus \{0\} \\ \langle 1, w \rangle = 0}} \frac{\|Aw\|}{\|w\|} = \max_{\substack{u \in \text{conv}(A) \\ x \in \Delta_{n-1} \setminus Z(u)}} \frac{\|Ax - u\|}{\text{dist}(x, Z(u))}, \quad (4)$$

and

$$\frac{\Phi(A)}{\max_{i=1, \dots, n} \|e_i\|} \leq \min_{\substack{u \in \text{conv}(A) \\ x \in \Delta_{n-1} \setminus Z(u)}} \frac{2\|Ax - u\|}{\text{dist}(x, Z(u))}. \quad (5)$$

Corollary 3. Let \mathbb{R}^n and \mathbb{R}^m be endowed with arbitrary norms. Suppose $A \in \mathbb{R}^{m \times n}$ has at least two different columns and $f : \mathbb{R}^m \rightarrow \mathbb{R} \cup \{\infty\}$ is L_f -smooth and μ_f -strongly convex on $\text{conv}(A)$. Then

$$L_{f,A} \leq L_f \cdot \max_{\substack{w \in \mathbb{R}^n \setminus \{0\} \\ \langle 1, w \rangle = 0}} \frac{\|Aw\|^2}{\|w\|^2}, \quad \mu_{f,A} \geq \frac{\mu_f \cdot \Phi(A)^2}{4 \max_{i=1, \dots, n} \|e_i\|^2},$$

In particular,

$$\frac{L_{f,A}}{\mu_{f,A}} \leq \frac{L_f}{\mu_f} \cdot \frac{4 \max_{i=1, \dots, n} \|e_i\|^2}{\Phi(A)^2} \cdot \max_{\substack{w \in \mathbb{R}^n \setminus \{0\} \\ \langle 1, w \rangle = 0}} \frac{\|Aw\|^2}{\|w\|^2}.$$

2.3 A refinement of the relative strong convexity constant

The construction of the constants $L_{f,A}$ and $\mu_{f,A}$ is *global* as it depends on f and the entire set $\text{conv}(A)$. We next describe a *local* refinement of the strong convexity constant.

Definition 3. Let \mathbb{R}^m be endowed with a norm $\|\cdot\|$, $A \in \mathbb{R}^{m \times n}$ have at least two different columns, and $f : \mathbb{R}^m \rightarrow \mathbb{R} \cup \{\infty\}$ be a differentiable convex function such that $\text{conv}(A) \subseteq \text{dom}(f)$. Let $f^* := \min_{u \in \text{conv}(A)} f(u) = \min_{x \in \Delta_{n-1}} f(Ax)$, $Z^* := \{z \in \Delta_{n-1} : f(Az) = f^*\}$. Following [9], define the quadratic functional growth constant $\mu_{f,A}^*$ as

$$\mu_{f,A}^* := \inf_{x \in \Delta_{n-1} \setminus Z^*} \frac{2(f(Ax) - f^*)}{\text{dist}(x, Z^*)^2}. \quad (6)$$

The convexity of f readily implies $\mu_{f,A}^* \geq \mu_{f,A} \geq 0$. Furthermore, as we next discuss, for an important class of functions $\mu_{f,A}^*$ is positive while $\mu_{f,A}$ may not be. Suppose $f : \mathbb{R}^m \rightarrow \mathbb{R} \cup \{\infty\}$ is defined as $f(u) = g(Eu) + \langle b, u \rangle$ where $g : \mathbb{R}^p \rightarrow \mathbb{R} \cup \{\infty\}$ is a strongly convex function, $b \in \mathbb{R}^m$, and $E \in \mathbb{R}^{p \times m}$. A function f of this form is typically not strongly convex and the relative strong convexity constant $\mu_{f,A}$ may be zero as illustrated in Example 2 below. On the other hand, for a function f of this form it is always the case that $\mu_{f,A}^* > 0$ as Theorem 1 below shows.

Theorem 1 gives a lower bound for $\mu_{f,A}^*$ similar in spirit to the lower bound for $\mu_{f,A}$ in Corollary 2. The statement and proof of Theorem 1 rely on the concept of local facial distance (an extension of the facial distance) introduced in [12]. Suppose \mathbb{R}^m is endowed with a norm $\|\cdot\|$. For $v \in \mathbb{R}^m$ define $\|\cdot\|_v : \mathbb{R}^{m+1} \rightarrow \mathbb{R}$ as follows. For $\bar{u} = \begin{bmatrix} u \\ u_{m+1} \end{bmatrix} \in \mathbb{R}^{m+1}$ let

$$\|\bar{u}\|_v := \sqrt{\|u\|^2 + |\langle v, u \rangle + u_{m+1}|}. \quad (7)$$

Observe that $\|\bar{u}\|_v > 0$ if $\bar{u} \neq 0$. For nonempty $F, G \subseteq \mathbb{R}^{m+1}$ let

$$\text{dist}_v(F, G) = \min_{\bar{u} \in F, \bar{w} \in G} \|\bar{u} - \bar{w}\|_v.$$

For $\bar{A} \in \mathbb{R}^{(m+1) \times n}$ with at least two different columns and $v \in \mathbb{R}^m$ let

$$F(v) := \underset{\bar{u} \in \text{conv}(\bar{A})}{\text{Argmin}} \left\langle \begin{bmatrix} v \\ 1 \end{bmatrix}, \bar{u} \right\rangle \in \text{faces}(\text{conv}(\bar{A})).$$

The *local facial distance* $\Phi_v(\bar{A})$ is defined as follows

$$\Phi_v(\bar{A}) := \min_{\substack{G \in \text{faces}(F(v)) \\ \emptyset \neq G \neq \text{conv}(\bar{A})}} \text{dist}_v(G, \text{conv}(\bar{A} \setminus G)).$$

Observe that $\Phi_v(\bar{A}) > 0$ for all $v \in \mathbb{R}^m$ and $\bar{A} \in \mathbb{R}^{(m+1) \times n}$ with at least two different columns. Furthermore, the facial distance can be recovered a special case of the local facial distance: Given $A \in \mathbb{R}^{m \times n}$ take $\bar{A} = \begin{bmatrix} A \\ 0 \end{bmatrix} \in \mathbb{R}^{(m+1) \times n}$ and $v = 0 \in \mathbb{R}^m$. In this case, $F(v) = \text{conv}(\bar{A}) = \text{conv}(A) \times \{0\}$, $\text{faces}(\text{conv}(\bar{A})) = \{F \times \{0\} : F \in \text{faces}(\text{conv}(A))\}$, and $\|\bar{u}\|_v = \|u\|$ for all $\bar{u} = \begin{bmatrix} u \\ 0 \end{bmatrix} \in \text{conv}(\bar{A})$. Therefore,

$$\Phi_v(\bar{A}) = \min_{\substack{G \in \text{faces}(F(v)) \\ \emptyset \neq G \neq \text{conv}(\bar{A})}} \text{dist}_v(G, \text{conv}(\bar{A} \setminus G)) = \min_{\substack{F \in \text{faces}(\text{conv}(A)) \\ \emptyset \neq F \neq \text{conv}(A)}} \text{dist}(F, \text{conv}(A \setminus F)) = \Phi(A).$$

Theorem 1. *Suppose \mathbb{R}^n is endowed with the ℓ_1 norm. Let $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, and $E \in \mathbb{R}^{p \times m}$ be such that $\begin{bmatrix} EA \\ b^\top A \end{bmatrix}$ has at least two different columns. Let $f : \mathbb{R}^m \rightarrow \mathbb{R} \cup \{\infty\}$ be defined by $f(u) = g(Eu) + \langle b, u \rangle$ where $g : \mathbb{R}^p \rightarrow \mathbb{R} \cup \{\infty\}$ is μ_g -strongly convex on $\text{conv}(EA)$ for some $\mu_g > 0$. Then $v = \underset{\mu_g}{\frac{2}{\mu_g}} \nabla g(Eu^*)$ is the same for all $u^* \in \underset{u \in \text{conv}(A)}{\text{Argmin}} f(u)$ and*

$$\mu_{f,A}^* \geq \frac{\mu_g \cdot \Phi_v(\bar{A})^2}{4} > 0, \quad (8)$$

$$\text{for } \bar{A} := \begin{bmatrix} EA \\ \frac{2}{\mu_g} b^\top A \end{bmatrix}.$$

Observe that the bound in Theorem 1 scales appropriately in the following sense. Suppose we replace $f(u) = g(Eu) + \langle b, u \rangle$ by $\tilde{f}(u) := \lambda f(u) = \lambda g(Eu) + \langle \lambda b, u \rangle =: \tilde{g}(Eu) + \langle \tilde{b}, u \rangle$ for some $\lambda > 0$. Then $\mu_{\tilde{f}, A}^* = \lambda \mu_{f, A}^*$, $\mu_{\tilde{g}} = \lambda \mu_g$, $\nabla \tilde{g} = \lambda \nabla g$ and thus v and \bar{A} are unchanged. Therefore all terms in inequality (8) scale exactly by λ .

The following corollary specializes Theorem 1 to the special case when the objective function f is a convex quadratic function. In that case $\mu_{f, A}^*$ is positive regardless of the strong convexity of f .

Corollary 4. *Let $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$ and $Q \in \mathbb{R}^{m \times m}$ be such that Q is symmetric positive semidefinite and $\begin{bmatrix} Q^{1/2} A \\ b^\top A \end{bmatrix}$ has at least two different columns. Let $f : \mathbb{R}^m \rightarrow \mathbb{R}$ be defined by $f(u) = \frac{1}{2} \langle Qu, u \rangle + \langle b, u \rangle$. Then $v = 2Q^{1/2}u^*$ is the same for all $u^* \in \underset{u \in \text{conv}(A)}{\text{Argmin}} f(u)$ and*

$$\mu_{f, A}^* \geq \frac{\Phi_v(\bar{A})^2}{4} > 0,$$

$$\text{for } \bar{A} := \begin{bmatrix} Q^{1/2} A \\ 2b^\top A \end{bmatrix}.$$

The next example describes a simple case when $\mu_{f, A}^* > \mu_{f, A} = 0$.

Example 2. Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ be and $A \in \mathbb{R}^{2 \times 3}$ be as follows

$$f(s, t) = \frac{1}{2}s^2 + t, \quad A = \begin{bmatrix} 1 & -1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

For $u = [0 \ 0]^\top$ and $x = [0 \ 0 \ 1]^\top$ we have $f(Ax) - f(u) - \langle \nabla f(u), Ax - u \rangle = 0$ and $\text{dist}(x, Z(u)) = 2$. Hence $\mu_{f, A} = 0$. On the other hand, $u^* = [0 \ 0]^\top$ and hence $Z^* = \{[\frac{1}{2} \ \frac{1}{2} \ 0]^\top\}$. Corollary 4 and some straightforward calculations yield the lower bound

$$\mu_{f, A}^* \geq \frac{\Phi_0(\bar{A})^2}{4} = \frac{1}{2},$$

for $\bar{A} = \begin{bmatrix} 1 & -1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 2 \end{bmatrix}$. A more detailed calculation shows that indeed $\mu_{f, A}^* = \frac{1}{2}$.

The proof of Theorem 1 relies on Lemma 1 and Lemma 2 below. These lemmas in turn use the following notation. For $x \in \Delta_{n-1}$, let $I(x) := \{i \in \{1, \dots, n\} : \langle e_i, x \rangle > 0\}$.

Lemma 1. *Let $\bar{A} \in \mathbb{R}^{(m+1) \times n}$ have at least two different columns and $v \in \mathbb{R}^m$. Suppose $\bar{u} \in F(v)$ and $x \in \Delta_{n-1}$ are such that $\bar{A}x \neq \bar{u}$. Then for $d := \frac{\bar{A}x - \bar{u}}{\|\bar{A}x - \bar{u}\|_v}$*

$$\Phi_v(\bar{A}) \leq \max\{\lambda : \exists y, z \in \Delta_{n-1}, I(y) \subseteq I(x), \bar{A}(y - z) = \lambda d\}. \quad (9)$$

Proof. This proof is a modification of the proof of [12, Proposition 1]. Let $I := I(x)$ and $J := \{j \in \{1, \dots, n\} : \bar{a}_j \in F(v)\}$. We will prove the following inequality that evidently implies (9):

$$\Phi_v(\bar{A}) \leq \max\{\lambda : \exists y, z \in \Delta_{n-1}, I(y) \subseteq I, I(z) \subseteq J, \bar{A}(y - z) = \lambda d\}. \quad (10)$$

To that end, observe that the right-hand side in (10) can be computed via the following primal-dual pair of linear programs

$$\begin{aligned} \min_{y_I, z_J, \lambda} \quad & \lambda \\ & \bar{A}_I y_I - \bar{A}_J z_J - \lambda d = 0 \\ & \mathbf{1}_I^\top y_I = 1 \\ & \mathbf{1}_J^\top z_J = 1 \\ & y_I, z_J \geq 0. \end{aligned} \quad (11)$$

and

$$\begin{aligned} \max_{p, t, \tau} \quad & t + \tau \\ & \bar{A}_I^\top p \leq t \mathbf{1}_I \\ & \bar{A}_J^\top p \geq -\tau \mathbf{1}_J \\ & \langle d, p \rangle = 1. \end{aligned} \quad (12)$$

Observe that (11) is feasible because in particular the triple (y_I, z_J, λ) defined by taking $y_I = x_I$, any $z_J \geq 0$ such that $\bar{u} = \bar{A}_J z_J$ and $\mathbf{1}_J^\top z_J = 1$, and $\lambda = \|\bar{A}x - \bar{u}\|$ satisfies the constraints in (11). Furthermore, (11) is evidently bounded since any feasible (y_I, z_J, λ) must have y_I and z_J bounded and $\bar{A}_I y_I - \bar{A}_J z_J = \lambda d$ with $d \neq 0$. Therefore both (11) and (12) attain their finite optimal values. Let $(\hat{y}_I, \hat{z}_J, \hat{\lambda})$ and $(\hat{p}, \hat{t}, \hat{\tau})$ be optimal solutions to (11) and (12) respectively. Then $\hat{\lambda} = \|\bar{A}_I \hat{y}_I - \bar{A}_J \hat{z}_J\|$. By complementary slackness, $\hat{z}_j > 0$ if and only if $\bar{a}_j \in \text{Argmin}_{\ell \in J} \langle \bar{a}_\ell, \hat{p} \rangle$. Likewise, $\hat{w}_i > 0$ if and only if $\bar{a}_i \in \text{Argmax}_{\ell \in I} \langle \bar{a}_\ell, \hat{p} \rangle$. Therefore $\bar{A}_J \hat{z}_J \in G := \text{Argmin}_{\bar{u} \in F} \langle \hat{p}, \bar{u} \rangle \in \text{faces}(F)$ and $\bar{A}_I \hat{y}_I \in \text{conv}(\bar{A} \setminus G)$. To finish, observe that

$$\begin{aligned} \Phi_v(\bar{A}) &\leq \text{dist}_v(G, \bar{A} \setminus G) \\ &\leq \|\bar{A}_I \hat{y}_I - \bar{A}_J \hat{z}_J\|_v \\ &= \max\{\lambda : \exists y, z \in \Delta_{n-1}, I(y) \subseteq I, I(z) \subseteq J, \bar{A}(y - z) = \lambda d\}. \end{aligned}$$

□

Lemma 2. *Let $\bar{A} \in \mathbb{R}^{(m+1) \times n}$ and $v \in \mathbb{R}^m$. Then for all $\bar{u} \in F(v)$ and $x \in \Delta_{n-1}$*

$$\min\{\|x - z\|_1 : z \in \Delta_{n-1}, \bar{A}z = \bar{u}\} \leq \frac{2\|\bar{A}x - \bar{u}\|_v}{\Phi_v(\bar{A})}. \quad (13)$$

Proof. Suppose $\bar{A}x \neq \bar{u}$ as otherwise there is nothing to show. To prove (13) we proceed by contradiction. Assume

$$z \in \Delta_{n-1}, \bar{A}z = \bar{u} \Rightarrow \|x - z\|_1 > \frac{2\|\bar{A}x - \bar{u}\|_v}{\Phi_v(\bar{A})}. \quad (14)$$

Let $d := \frac{\bar{A}x - \bar{u}}{\|\bar{A}x - \bar{u}\|_v}$ and consider the following linear program

$$\begin{aligned} \max_{w,t} \quad & t \\ & \bar{A}w = td \\ & x - w \in \Delta_{n-1} \\ & \|w\|_1 \leq \frac{2t}{\Phi_v(\bar{A})}. \end{aligned} \tag{15}$$

By (9) there exist $y, z \in \Delta_{n-1}$ with $I(y) \subseteq I(x)$ and $\bar{A}(y - z) = \Phi_v(\bar{A})d$. Thus for $\delta > 0$ sufficiently small the linear program (15) has a feasible solution $(w, t) = \delta \cdot (y - z, \Phi_v(\bar{A}))$ with $t = \delta \cdot \Phi_v(\bar{A}) > 0$. Assumption (14) thus implies that (15) has an optimal solution (\hat{w}, \hat{t}) with $0 < \hat{t} < \|\bar{A}x - \bar{u}\|_v$. Let $\hat{x} := x - \hat{w} \in \Delta_{n-1}$. Observe that $\bar{A}\hat{x} - \bar{u} = \bar{A}x - \bar{u} - \bar{A}\hat{w} = (\|\bar{A}x - \bar{u}\|_v - \hat{t})d \neq 0$ and hence $d = \frac{\bar{A}\hat{x} - \bar{u}}{\|\bar{A}\hat{x} - \bar{u}\|_v}$. Consider the modification of (15) obtained by replacing x with \hat{x} :

$$\begin{aligned} \max_{w,t} \quad & t \\ & \bar{A}w = td \\ & \hat{x} - w \in \Delta_{n-1} \\ & \|w\|_1 \leq \frac{2t}{\Phi_v(\bar{A})}. \end{aligned} \tag{16}$$

Proceeding as above, it follows that (16) has a feasible solution (w', t') with $t' > 0$. This implies that $(\hat{w} + w', \hat{t} + t')$ is feasible for (15) and $\hat{t} + t' > \hat{t}$ which contradicts the optimality of (\hat{w}, \hat{t}) for (15). \square

Proof of Theorem 1. The optimality conditions for $\min_{u \in \text{conv}(A)} f(u)$ imply that for all $u^* \in \text{Argmin } f(u)$ and all $u \in \text{conv}(A)$

$$\langle E^\top \nabla g(Eu^*) + b, u - u^* \rangle \geq 0. \tag{17}$$

Therefore if $u^*, u' \in \text{Argmin}_{u \in \text{conv}(A)} f(u)$, the strong convexity of g and (17) imply

$$\begin{aligned} \mu_g \|Eu^* - Eu'\|^2 &\leq \langle \nabla g(Eu^*) - \nabla g(Eu'), Eu^* - Eu' \rangle \\ &= \langle E^\top \nabla g(Eu^*) - E^\top \nabla g(Eu'), u^* - u' \rangle \\ &\leq 0. \end{aligned}$$

Hence $Eu^* = Eu'$ whenever $u^*, u' \in \text{Argmin}_{u \in \text{conv}(A)} f(u)$. In particular, $v = \nabla g(Eu^*)$ is the same for all $u^* \in \text{Argmin}_{u \in \text{conv}(A)} f(u)$. Furthermore, from (17) it follows that $\bar{u} := \bar{A}z$ is the same for all $z \in Z^*$ and $\bar{u} \in F(v)$. Lemma 2 implies that for all $x \in \Delta_{n-1}$

$$\text{dist}(x, Z^*) = \min\{\|x - z\| : z \in \Delta_{n-1}, \bar{A}z = \bar{u}\} \leq \frac{2\|\bar{A}x - \bar{u}\|_v}{\Phi_v(\bar{A})}. \tag{18}$$

Next, observe that the strong convexity of g and (17) imply that for all $x \in \Delta_{n-1}$ and $z \in Z^*$

$$\begin{aligned}
f(Ax) - f^* &= g(EAx) - g(EAz) + \langle b, Ax - Az \rangle \\
&\geq \langle \nabla g(EAz), EAx - EAz \rangle + \frac{\mu_g}{2} \|EAx - EAz\|^2 + \langle b, Ax - Az \rangle \\
&= \frac{\mu_g}{2} \left(\|EAx - EAz\|^2 + \frac{2}{\mu_g} \langle E^\top \nabla g(EAz) + b, Ax - Az \rangle \right) \\
&= \frac{\mu_g}{2} \left(\|EAx - EAz\|^2 + \left| \langle v, EAx - EAz \rangle + \left\langle \frac{2}{\mu_g} b, Ax - Az \right\rangle \right| \right) \\
&= \frac{\mu_g}{2} \|\bar{A}x - \bar{u}\|_v^2.
\end{aligned} \tag{19}$$

Putting together (18) and (19) we get

$$\mu_{f,A}^* = \min_{x \in \Delta_{n-1} \setminus Z^*} \frac{2(f(Ax) - f^*)}{\text{dist}(x, Z^*)^2} \geq \min_{\substack{x \in \Delta_{n-1} \\ Ax \neq \bar{u}}} \frac{\mu_g \cdot \|\bar{A}x - \bar{u}\|_v^2}{\text{dist}(x, Z^*)^2} \geq \frac{\mu_g \cdot \Phi_v(\bar{A})^2}{4}.$$

□

2.4 Proofs of Proposition 1 and Proposition 2

Proof of Proposition 1. For $\text{diam}(A)$ observe that

$$\begin{aligned}
\text{diam}(A) &= \max_{x, y \in \Delta_{n-1}} \|A(x - y)\| = \max_{\substack{x, y \in \Delta_{n-1} \\ \|x - y\|_1 = 2}} \frac{2\|A(x - y)\|}{\|x - y\|_1} = \max_{\substack{x, y \in \Delta_{n-1} \\ x \neq y}} \frac{2\|A(x - y)\|}{\|x - y\|_1} \\
&= \max_{\substack{u \in \text{conv}(A) \\ x \in \Delta_{n-1} \setminus Z(u)}} \frac{2\|Ax - u\|}{\text{dist}(x, Z(u))}.
\end{aligned}$$

For $\Phi(A)$ we prove the two inequalities separately. From Lemma 2 applied to $\bar{A} = \begin{bmatrix} A \\ 0 \end{bmatrix}$ and $v = 0$ it follows that for all $u \in \text{conv}(A)$ and $x \in \Delta_{n-1}$

$$\text{dist}(x, Z(u)) = \min\{\|x - z\| : z \in \Delta_{n-1}, Az = u\} \leq \frac{2\|Ax - u\|}{\Phi(A)}.$$

Therefore $\Phi(A) \leq \min_{\substack{u \in \text{conv}(A) \\ x \in \Delta_{n-1} \setminus Z(u)}} \frac{2\|Ax - u\|}{\text{dist}(x, Z(u))}$. For the reverse inequality, let $F \in \text{faces}(\text{conv}(A))$

be such that $\emptyset \neq F \neq \text{conv}(A)$ and $\Phi(A) = \text{dist}(F, \text{conv}(A \setminus F))$. Then $\Phi(A) = \|A\hat{x} - \hat{u}\|$ for some $\hat{u} \in F$ and $\hat{x} \in \Delta_{n-1}$ with $A\hat{x} \in \text{conv}(A \setminus F)$. Since $A\hat{x} \in \text{conv}(A \setminus F)$, without loss of generality we may assume that $\hat{x} \in \Delta_{n-1}$ is chosen so that $i \in I(\hat{x}) \Rightarrow a_i \notin F$. Since F is a face, it follows that $Az \in F \Rightarrow I(z) \cap I(\hat{x}) = \emptyset$. In particular $\text{dist}(\hat{x}, Z(\hat{u})) = 2$ and thus

$$\min_{\substack{u \in \text{conv}(A) \\ x \in \Delta_{n-1} \setminus Z(u)}} \frac{2\|Ax - u\|}{\text{dist}(x, Z(u))} \leq \frac{2\|A\hat{x} - \hat{u}\|}{\text{dist}(\hat{x}, Z(\hat{u}))} = \|A\hat{x} - \hat{u}\| = \Phi(A).$$

□

Proof of Proposition 2. For (4) observe that

$$\max_{\substack{u \in \text{conv}(A) \\ x \in \Delta_{n-1} \setminus Z(u)}} \frac{\|Ax - v\|}{\text{dist}(x, Z(u))} = \max_{\substack{x, y \in \Delta_{n-1} \\ x \neq y}} \frac{\|A(x - y)\|}{\|x - y\|} = \max_{\substack{w \in \mathbb{R}^n \setminus \{0\} \\ \mathbf{1}^\top w = 0}} \frac{\|Aw\|}{\|w\|}.$$

For (5), observe that for all $x, y \in \Delta_{n-1}$ we have $x - y = \sum_{i=1}^n \langle e_i, x - y \rangle e_i$ and hence

$$\|x - y\| \leq \sum_{i=1}^n |\langle e_i, x - y \rangle| \|e_i\| \leq \max_{i=1, \dots, n} \|e_i\| \cdot \|x - y\|_1.$$

Thus Proposition 1 yields

$$\min_{\substack{u \in \text{conv}(A) \\ x \in \Delta_{n-1} \setminus Z(u)}} \frac{2\|Ax - u\|}{\text{dist}(x, Z(u))} \geq \min_{\substack{u \in \text{conv}(A) \\ x \in \Delta_{n-1} \setminus Z(u)}} \frac{2\|Ax - u\|}{\max_{i=1, \dots, n} \|e_i\| \cdot \text{dist}_1(x, Z(u))} = \frac{\Phi(A)}{\max_{i=1, \dots, n} \|e_i\|}.$$

□

3 Linear convergence of first-order methods

This section discusses linear convergence results for two first-order algorithms for the problem (1) namely the Frank-Wolfe with away steps (Algorithm 1) and the projected gradient method (Algorithm 2). Linear convergence results for both algorithms have been previously established in [1, 2, 6, 11, 12] under suitable assumptions. The goal of this section is to illustrate the role of the relative condition number $\frac{L_{f,A}}{\mu_{f,A}^*}$ in these linear convergence results. The role of the relative condition number $\frac{L_{f,A}}{\mu_{f,A}^*}$ is akin to the role of the usual condition number $\frac{L_f}{\mu_f}$ in the linear convergence of the gradient descent algorithm for unconstrained convex minimization.

Both proofs of linear convergence rely on the following elementary observation. If $a \leq 0$, $b > 0$, and $\alpha_{\max} > 0$ then

$$\min_{\alpha \in [0, \alpha_{\max}]} a\alpha + \frac{b}{2}\alpha^2 = \begin{cases} -\frac{a^2}{2b} & \text{if } \alpha_{\max} > -\frac{a}{b} \\ \alpha_{\max} \left(a + \frac{b}{2}\alpha_{\max} \right) \leq \frac{a}{2}\alpha_{\max} & \text{if } \alpha_{\max} \leq -\frac{a}{b}. \end{cases} \quad (20)$$

3.1 Frank-Wolfe algorithm with away steps

Algorithm 1 gives a description of the Frank-Wolfe algorithm with away steps for (1). This version of the algorithm has been previously discussed in [1, 6, 12]. Algorithm 1 relies on the following notation. Given $u = Ax \in \text{conv}(A)$, let $I(x) := \{i \in \{1, \dots, n\} : \langle e_i, x \rangle > 0\}$. The set $I(x)$ describes the *support* of $u = Ax$, that is, the indices of the columns of A that appear with positive weight in the convex combination $u = Ax$.

A critical detail in Algorithm 1 is the choice of step size α_k in Step 9. The construction of $L_{f,A}$ implies that for $\alpha \in [0, \alpha_{\max}]$

$$f(u_k + \alpha v) \leq f(u_k) + \alpha \langle \nabla f(u_k), v \rangle + \frac{L_{f,A}\alpha^2}{2} \|w\|_1^2 \leq f(u_k) + \alpha \langle \nabla f(u_k), v \rangle + 2L_{f,A}\alpha^2.$$

Algorithm 1 Frank-Wolfe algorithm with away steps

- 1: Pick $x_0 \in \Delta_{n-1}$; put $u_0 := Ax_0$; $k := 0$
 - 2: **for** $k = 0, 1, 2, \dots$ **do**
 - 3: $j := \operatorname{argmin}_{i=1, \dots, n} \langle \nabla f(u_k), a_i \rangle$; $\ell := \operatorname{argmax}_{i \in I(x_k)} \langle \nabla f(u_k), a_i \rangle$
 - 4: **if** $\langle \nabla f(u_k), a_j - u_k \rangle < \langle \nabla f(u_k), u_k - a_\ell \rangle$ or $|I(x_k)| = 1$ **then**
 - 5: $v := a_j - u_k$; $w := e_j - x_k$; $\alpha_{\max} := 1$ (regular step)
 - 6: **else**
 - 7: $v := u_k - a_\ell$; $w := x_k - e_\ell$; $\alpha_{\max} := \frac{\langle e_\ell, x_k \rangle}{1 - \langle e_\ell, x_k \rangle}$ (away step)
 - 8: **end if**
 - 9: choose $\alpha_k \in [0, \alpha_{\max}]$
 - 10: $x_{k+1} := x_k + \alpha_k w$; $u_{k+1} := u_k + \alpha_k v = Ax_{k+1}$
 - 11: **end for**
-

To simplify our analysis of linear convergence we will assume that α_k in Step 9 is chosen via

$$\alpha_k := \operatorname{argmin}_{\alpha \in [0, \alpha_{\max}]} \left\{ f(u_k) + \alpha \langle \nabla f(u_k), v \rangle + 2L_{f,A} \alpha^2 \right\} = \min \left\{ \alpha_{\max}, -\frac{\langle \nabla f(u_k), v \rangle}{4L_{f,A}} \right\}. \quad (21)$$

This is evidently possible only in the ideal case when $L_{f,A}$ is known. In the more realistic case when $L_{f,A}$ is not known, a standard backtracking procedure can be used to choose a constant $L > 0$ bounded above by a constant multiple of $L_{f,A}$ and such that the step size

$$\alpha_k := \operatorname{argmin}_{\alpha \in [0, \alpha_{\max}]} \left\{ f(u_k) + \alpha \langle \nabla f(u_k), v \rangle + 2L \alpha^2 \right\} = \min \left\{ \alpha_{\max}, -\frac{\langle \nabla f(u_k), v \rangle}{4L} \right\}$$

satisfies

$$f(u_k + \alpha_k v) \leq f(u_k) + \alpha_k \langle \nabla f(u_k), v \rangle + 2L \alpha_k^2.$$

Our ensuing analysis would then apply with $L_{f,A}$ replaced by a constant multiple of it. For the remainder of this subsection we will assume that α_k is indeed chosen via (21). Combining this assumption and (20) applied to $a = \langle \nabla f(u_k), v \rangle$ and $b = 4L_{f,A}$ we obtain

$$f(u_{k+1}) - f(u_k) \leq \begin{cases} -\frac{\langle \nabla f(u_k), v \rangle^2}{8L_{f,A}} & \text{if } \alpha_k < \alpha_{\max} \\ \frac{\langle \nabla f(u_k), v \rangle}{2} \alpha_{\max} & \text{if } \alpha_k = \alpha_{\max}. \end{cases} \quad (22)$$

The following result provides the crux of the linear convergence of Algorithm 1.

Lemma 3. *Suppose \mathbb{R}^n is endowed with the ℓ_1 norm $\|\cdot\|_1$, and $A \in \mathbb{R}^{m \times n}$ and $f : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$ are such that $\mu_{f,A}^* > 0$. Then the direction v chosen in Step 5 or Step 7 of Algorithm 1 satisfies*

$$\langle \nabla f(u_k), v \rangle^2 \geq \frac{\mu_{f,A}^*}{2} (f(u_k) - f^*) \quad (23)$$

and

$$\langle \nabla f(u_k), v \rangle \leq f^* - f(u_k). \quad (24)$$

Proof. The choice of regular versus away steps in Step 5 and Step 7 imply that if $u_k = Ax_k \in \text{conv}(A)$ then

$$2 \langle \nabla f(u_k), v \rangle \leq \min_{i=1, \dots, n} \langle \nabla f(u_k), a_i \rangle - \max_{i \in I(x_k)} \langle \nabla f(u_k), a_i \rangle \leq 0.$$

Let $z^* \in Z^*$ be such that $\|x_k - z^*\|_1 = \text{dist}(x_k, Z^*)$ and $u^* = Az^*$. Observe that $x_k - z^* = \delta(z - y)$ where $z, y \in \Delta_{n-1}$, $I(z) \subseteq I(x_k)$, and $\delta := \frac{\|x_k - z^*\|_1}{2} \leq 1$. Thus

$$\begin{aligned} \langle \nabla f(u_k), u_k - u^* \rangle &= \langle \nabla f(u_k), A(x_k - z^*) \rangle = \delta \langle \nabla f(u_k), A(z - y) \rangle \\ &\leq \delta \left(\max_{\substack{z \in \Delta_{n-1} \\ I(z) \subseteq I(x_k)}} \langle \nabla f(u_k), Az \rangle - \min_{y \in \Delta_{n-1}} \langle \nabla f(u_k), Ay \rangle \right) \\ &= \delta \left(\max_{i \in I(x_k)} \langle \nabla f(u_k), a_i \rangle - \min_{i=1, \dots, n} \langle \nabla f(u_k), a_i \rangle \right) \\ &\leq 2\delta |\langle \nabla f(u_k), v \rangle|. \end{aligned}$$

The construction of $\mu_{f,A}^*$, convexity of f , and the latter inequality yield

$$0 \leq \mu_{f,A}^* \leq \frac{2(f(u_k) - f^*)}{\text{dist}(z, Z^*)^2} = \frac{f(u_k) - f^*}{2\delta^2} \leq \frac{\langle \nabla f(u_k), u_k - u^* \rangle}{2\delta^2} \leq \frac{|\langle \nabla f(u_k), v \rangle|}{\delta}.$$

Therefore (23) follows. On the other hand, the choice of v and convexity of f yields

$$\langle \nabla f(u_k), v \rangle \leq \min_{u \in \text{conv}(A)} \langle \nabla f(u_k), u - u_k \rangle \leq \langle \nabla f(u_k), u^* - u_k \rangle \leq f^* - f(u_k)$$

and thus (24) follows as well. \square

Once we are equipped with Lemma 3, the following linear convergence result readily follows via a clever counting argument introduced in [6] and subsequently used in [1, 12]. To provide a full picture of this linear convergence result, the proof below briefly replicates the necessary material from [1, 6, 12].

Proposition 3. *Suppose \mathbb{R}^n is endowed with the ℓ_1 norm $\|\cdot\|_1$, and $A \in \mathbb{R}^{m \times n}$ and $f : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$ are such that $\frac{L_{f,A}}{\mu_{f,A}^*} < \infty$. If x_0 is a vertex of Δ_{n-1} then the iterates generated by Algorithm 1 satisfy*

$$f(u_k) - f^* \leq \left(1 - \min \left\{ \frac{\mu_{f,A}^*}{16L_{f,A}}, \frac{1}{2} \right\} \right)^{k/2} (f(u_0) - f^*).$$

Proof. Consider separately the three possible cases that can occur at iteration k , namely $\alpha_k < \alpha_{\max}$, $\alpha_k = \alpha_{\max} \geq 1$, and $\alpha_k = \alpha_{\max} < 1$.

Case 1: $\alpha_k < \alpha_{\max}$. In this case $|I(x_{k+1})| \leq |I(x_k)| + 1$. Furthermore, (22) and (23) imply that

$$f(u_{k+1}) - f(u_k) \leq -\frac{\mu_{f,A}^*}{16L_{f,A}} (f(u_k) - f^*).$$

Case 2: $\alpha_k = \alpha_{\max} \geq 1$. In this case $|I(x_{k+1})| \leq |I(x_k)|$. Furthermore, (22) and (24) imply that

$$f(u_{k+1}) - f(u_k) \leq -\frac{1}{2}(f(u_k) - f^*).$$

Case 3: $\alpha_k = \alpha_{\max} < 1$. In this case $|I(x_{k+1})| \leq |I(x_k)| - 1$. Furthermore, (22) implies that

$$f(u_{k+1}) - f(u_k) \leq 0.$$

Therefore to finish it suffices to show that in the first k iterations Case 3 can occur at most $k/2$ times. Since $|I(x_0)| = 1$ and $|I(x_i)| \geq 1$ for $i = 1, 2, \dots$, it follows that for each iteration when Case 3 occurred there must have been at least one previous iteration when Case 1 occurred. Hence in the first k iterations Case 3 could occur at most $k/2$ times. \square

3.2 Projected gradient

Algorithm 2 gives a description of the projected gradient algorithm for (1). This version can be seen as a particular case of more general gradient schemes like those discussed in [2, 11].

Algorithm 2 Projected gradient algorithm

- 1: Pick $x_0 \in \Delta_{n-1}$;
 - 2: **for** $k = 0, 1, 2, \dots$ **do**
 - 3: choose $L > 0$
 - 4: $x_{k+1} = \operatorname{argmin}_{x \in \Delta_{n-1}} \left\{ f(Ax_k) + \langle \nabla f(Ax_k), A(x - x_k) \rangle + \frac{L}{2} \|x - x_k\|_2^2 \right\}$
 - 5: **end for**
-

Like the choice of step size α_k in Step 9 of Algorithm 1, the choice of L in Step 3 is a critical detail in Algorithm 2. The construction of $L_{f,A}$ implies that

$$f(Ax) \leq f(Ax_k) + \langle \nabla f(Ax_k), A(x - x_k) \rangle + \frac{L_{f,A}}{2} \|x - x_k\|_2^2.$$

To simplify our analysis of linear convergence we will assume that Step 3 chooses $L = L_{f,A}$. This is possible only if $L_{f,A}$ is known. In the more realistic case when $L_{f,A}$ is not known, a standard backtracking procedure can be used to choose L bounded above by a constant multiple of $L_{f,A}$ and such that the next iterate $x_{k+1} \in \Delta_{n-1}$ chosen at Step 4 satisfies

$$f(Ax_{k+1}) \leq f(Ax_k) + \langle \nabla f(Ax_k), A(x_{k+1} - x_k) \rangle + \frac{L}{2} \|x_{k+1} - x_k\|_2^2.$$

Our ensuing analysis would then apply with $L_{f,A}$ replaced by a constant multiple of it. The assumption that $L = L_{f,A}$ in Step 3 of Algorithm 2 combined with (20) readily imply that for all $z \in \Delta_{n-1}$ such that $\langle \nabla f(Ax_k), A(z - x_k) \rangle \leq 0$

$$\begin{aligned} f(Ax_{k+1}) - f(Ax_k) &\leq \min_{\alpha \in [0,1]} \left(\langle \nabla f(Ax_k), A(z - x_k) \rangle \alpha + \frac{L_{f,A}}{2} \|z - x_k\|_2^2 \alpha^2 \right) \\ &\leq \begin{cases} -\frac{\langle \nabla f(Ax_k), A(z - x_k) \rangle^2}{2L_{f,A} \|z - x_k\|_2^2} & \text{if } -\frac{\langle \nabla f(Ax_k), A(z - x_k) \rangle}{L_{f,A} \|z - x_k\|_2^2} < 1 \\ \frac{\langle \nabla f(Ax_k), A(z - x_k) \rangle}{2} & \text{if } -\frac{\langle \nabla f(Ax_k), A(z - x_k) \rangle}{L_{f,A} \|z - x_k\|_2^2} \geq 1. \end{cases} \end{aligned} \quad (25)$$

Proposition 4. Suppose \mathbb{R}^n is endowed with the ℓ_2 norm $\|\cdot\|_2$, and $A \in \mathbb{R}^{m \times n}$ and $f : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$ are such that $\frac{L_{f,A}}{\mu_{f,A}^*} < \infty$. Then the sequence of iterates $\{x_k, k = 0, 1, \dots\}$ generated by Algorithm 2 satisfy

$$f(Ax_k) - f^* \leq \left(1 - \min \left\{ \frac{\mu_{f,A}^*}{4L_{f,A}}, \frac{1}{2} \right\}\right)^k (f(Ax_0) - f^*). \quad (26)$$

Proof. At iteration k let $z \in Z^*$ be such that $\|z - x_k\|_2^2 \leq \frac{2(f(Ax_k) - f^*)}{\mu_{f,A}^*}$. The convexity of f yields

$$\langle \nabla f(Ax_k), A(z - x_k) \rangle \leq f^* - f(Ax_k) \leq 0$$

and

$$\frac{\langle \nabla f(Ax_k), A(z - x_k) \rangle^2}{\|z - x_k\|_2^2} \geq \frac{\mu_{f,A}^*}{2} (f(Ax_k) - f^*)$$

Combining these inequalities and (25) it follows that

$$f(Ax_{k+1}) - f(Ax_k) \leq - \min \left\{ \frac{\mu_{f,A}^*}{4L_{f,A}}, \frac{1}{2} \right\} (f(Ax_k) - f^*).$$

Therefore (26) follows by induction. □

Acknowledgements

This research has been funded by NSF grant CMMI-1534850.

References

- [1] A. Beck and S. Shtern. Linearly convergent away-step conditional gradient for non-strongly convex functions. *Mathematical Programming*, 164:1–27, 2017.
- [2] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal of Imaging Services*, 2009.
- [3] S. Bubeck, Y. Lee, and M. Singh. A geometric alternative to Nesterov’s accelerated gradient descent. *arXiv preprint arXiv:1506.08187*, 2015.
- [4] D. Drusvyatskiy, M. Fazel, and S. Roy. An optimal first order method based on optimal quadratic averaging. *To Appear in SIAM Journal on Optimization*, 2018.
- [5] S. Karimi and S. Vavasis. A single potential governing convergence of conjugate gradient, accelerated gradient and geometric descent. *arXiv preprint arXiv:1712.09498*, 2017.
- [6] S. Lacoste-Julien and M. Jaggi. On the global linear convergence of Frank-Wolfe optimization variants. In *Advances in Neural Information Processing Systems (NIPS)*, 2015.

- [7] H. Lu, R. Freund, and Y. Nesterov. Relatively-smooth convex optimization by first-order methods, and applications. *To Appear in SIAM Journal on Optimization*, 2018.
- [8] C. Ma, N. Gudapati, M. Jahani, R. Tappenden, and M. Takáč. Underestimate sequences via quadratic averaging. *arXiv preprint arXiv:1710.03695*, 2017.
- [9] I. Necoara, Y. Nesterov, and F. Glineur. Linear convergence of first order methods for non-strongly convex optimization. *To Appear in Mathematical Programming*, 2018.
- [10] Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Applied Optimization. Kluwer Academic Publishers, 2004.
- [11] Y. Nesterov. Gradient methods for minimizing composite functions. *Mathematical Programming*, 140(1):125–161, 2013.
- [12] J. Peña and D. Rodríguez. Polytope conditioning and linear convergence of the Frank-Wolfe algorithm. *To Appear in Mathematics of Operations Research*, 2018.