

“EFFICIENT” SUBGRADIENT METHODS FOR GENERAL CONVEX OPTIMIZATION

JAMES RENEGAR

ABSTRACT. A subgradient method is presented for solving general convex optimization problems, the main requirement being that a strictly-feasible point is known. A feasible sequence of iterates is generated, which converges to within user-specified error of optimality. Feasibility is maintained with a line-search at each iteration, avoiding the need for orthogonal projections onto the feasible region (an operation that limits practicality of traditional subgradient methods). Lipschitz continuity is not required, yet the algorithm is shown to possess a convergence rate analogous to rates for traditional methods, albeit with error measured relatively, whereas traditionally error has been absolute. The algorithm is derived using an elementary framework that can be utilized to design other such algorithms.

1. Introduction

Given a convex optimization problem possessing little more structure than having a known strictly-feasible point, we provide a simple transformation to an equivalent convex optimization problem which has only linear equations as constraints, and has Lipschitz-continuous objective function defined on the whole space. Virtually any subgradient method can be applied to solve the equivalent problem. Relying only on a line-search during an iteration, the resulting iterate is made to be a feasible point for the original problem. Moreover, this sequence of feasible points converges to within user-specified error of optimality.

The algorithm can be implemented directly in terms of the original problem, leaving the equivalent problem hidden from view. In this introductory section we present two examples of algorithms and complexity bounds that ensue when standard subgradient methods are applied to the equivalent problem. Not until subsequent sections is the equivalent problem revealed.

Let \mathcal{E} denote a finite-dimensional real vector space with inner product $\langle \cdot, \cdot \rangle$ and associated norm $\| \cdot \|$.

Consider an optimization problem

$$\begin{aligned} \min \quad & f(x) \\ \text{s.t.} \quad & x \in \text{Feas} . \end{aligned} \tag{1.1}$$

Assume $f : \mathcal{E} \rightarrow (-\infty, \infty]$ is an extended-valued, lower-semicontinuous convex function. Equivalently, assume $f : \mathcal{E} \rightarrow \mathbb{R}$ has closed and convex epigraph, $\text{epi}(f) := \{(x, t) \mid x \in \text{dom}(f), t \geq f(x)\}$, (where $\text{dom}(f)$ is the (effective) domain of f (i.e., the set on which f is finite)).

Let

$$\text{Feas} = \{x \in S \mid Ax = b\} ,$$

where S is a closed, convex set with nonempty interior.

Assume \bar{e} is a known feasible point contained in $\text{int}(S \cap \text{dom}(f))$.

Research supported in part by NSF CCF Award #1552518.

Let \bar{D} denote the diameter of the sublevel set $\{x \in \text{Feas} : f(x) \leq f(\bar{e})\}$ (i.e., the supremum of distances between pairs of points in the set). We assume \bar{D} is finite, which together with previous assumptions implies the optimal value of (1.1) is attained at some feasible point. Let f^* denote the optimal value.

Let \hat{f} be a user-chosen constant satisfying $\hat{f} > f(\bar{e})$ (thus, $(\bar{e}, \hat{f}) \in \text{int}(\text{epi}(f))$). Define

$$\hat{r} := \sup\{r \mid \|x - \bar{e}\| \leq r \text{ and } Ax = b \Rightarrow x \in S \text{ and } f(x) \leq \hat{f}\}.$$

Note $\hat{r} > 0$, because $\bar{e} \in \text{int}(S \cap \text{dom}(f))$.

The scalars \bar{D} and \hat{r} appear in the complexity bounds for the representative algorithms presented below, but their values are not assumed to be known. We now introduce notation to be used in specifying the algorithms.

For fixed $x \in \mathcal{E}$ and $t \in \mathbb{R}$, and for scalars α , let

$$x(\alpha) := \bar{e} + \alpha \cdot (x - \bar{e}) \quad \text{and} \quad t(\alpha) := \hat{f} + \alpha \cdot (t - \hat{f}).$$

Define

$$\alpha_1(x) := \sup\{\alpha \mid x(\alpha) \in S\},$$

and

$$\alpha_2(x, t) := \sup\{\alpha \mid f(x(\alpha)) \leq t(\alpha)\}.$$

Note that $\alpha_1(x), \alpha_2(x, t) > 0$, due to our requirements for \bar{e} and \hat{f} .

We are concerned only with pairs (x, t) for which $Ax = b$ and $t < \hat{f}$. (Observe then, $Ax(\alpha) = b$ and $t(\alpha) < \hat{f}$ for all $\alpha \geq 0$.) We claim in this setting, at least one of the values $\alpha_1(x), \alpha_2(x, t)$ is finite. To establish the claim, assume $\alpha_2(x, t) = \infty$ – we show it easily follows that $\alpha_1(x)$ is finite.

Since $\alpha_2(x, t) = \infty$, the convex univariate function

$$\phi(\alpha) := f(x(\alpha)) - t(\alpha) \tag{1.2}$$

has negative value for all $\alpha \geq 0$. Hence, since $t < \hat{f}$ (by assumption) – and thus $t(\alpha) \rightarrow -\infty$ as $\alpha \rightarrow \infty$ – the function $\alpha \mapsto f(x(\alpha))$ is unbounded below on the interval $[0, \infty)$. In particular, $f(x(\alpha)) < f^*$ for some $\alpha > 0$, implying for this value of α , $x(\alpha) \notin \text{Feas}$, that is, $\alpha_1(x) < \alpha$, establishing the claim.

Let

$$\alpha(x, t) := \min\{\alpha_1(x), \alpha_2(x, t)\},$$

a value we now know to be positive and finite (assuming $Ax = b$ and $t < \hat{f}$). For given pairs (x, t) , the algorithms require that the scalar $\alpha(x, t)$ be computed. (This is the line search referred to above.) In general, of course, $\alpha(x, t)$ cannot be computed exactly. However, there is hope for quickly gaining a “good enough” approximation to $\alpha(x, t)$.

To understand, first observe for many optimization problems, for each x the value $\alpha_1(x)$ can readily be accurately approximated, or determined to be ∞ . All that is required here is finding whether a given half-line, with endpoint in $\text{int}(S)$, intersects the boundary of S , and if so, accurately approximating where the intersection occurs. For most feasible regions, this is far more easily accomplished than one of the main operations underlying traditional subgradient methods, where an (arbitrary) point outside Feas must be orthogonally projected onto Feas .

Assume for the moment that $\alpha_1(x)$ has been accurately approximated or determined to equal ∞ . For definiteness, assume $\alpha_1(x) = \infty$, in which case we know $\alpha_2(x, t)$ is finite. Approximating $\alpha_2(x, t)$ is the same as approximating the solution to $\phi(\alpha) = 0$, where ϕ is the function (1.2). As ϕ is a convex function satisfying $\phi(0) < 0$, approximating the root can be accomplished by first finding a value α sufficiently large so as to satisfy $\phi(\alpha) > 0$, and then proceeding to isolate the root using bisection.

On the other hand, if $\alpha_1(x)$ is finite, then to determine whether $\alpha_2(x, t)$ even needs to be computed, check whether $\phi(\alpha_1(x)) \leq 0$. If so, there is no need to compute $\alpha_2(x, t)$, because $\alpha(x, t) = \alpha_1(x)$. If instead, $\phi(\alpha_1(x)) > 0$, then to compute $\alpha_2(x, t)$ proceed by bisection, with initial interval $[0, \alpha_1(x)]$.

That bisection can be applied gives reason to hope a good approximation to $\alpha(x, t)$ can be quickly computed. Nonetheless, questions abound regarding what constitutes a “good enough” approximation. In what follows, we duck the issue, simply assuming $\alpha(x, t)$ can be computed exactly.

For x, t satisfying $Ax = b$ and $t < \hat{f}$, let

$$\pi'(x, t) := (\bar{e}, \hat{f}) + \alpha(x, t) \cdot ((x, t) - (\bar{e}, \hat{f})) . \quad (1.3)$$

Thus, in traveling from (\bar{e}, \hat{f}) in the direction $(x, t) - (\bar{e}, \hat{f})$, $\pi'(x, t)$ is the first point (x', t') encountered for which either $x' \in \text{bdy}(S)$ (boundary) or $(x', t') \in \text{bdy}(\text{epi}(f))$. Note in particular that if $(x', t') = \pi'(x, t)$, then x' is both feasible and lies in the domain of f .

For $x' \in \text{bdy}(S)$, let

$$G_1(x') := \left\{ \frac{-1}{\langle v, \bar{e} - x' \rangle} v \mid \vec{0} \neq v \in N_S(x') \right\} ,$$

where $N_S(x')$ is the normal cone to S at x' ($\{v \mid \forall x \in S, \langle v, x - x' \rangle \leq 0\}$). The denominator is negative, because $\bar{e} \in \text{int}(S)$.

For $(x', t') \in \text{bdy}(\text{epi}(f))$, let

$$G_2(x', t') := \left\{ \frac{-1}{\langle v, \bar{e} - x' \rangle + (\hat{f} - t')\delta} v \mid (\vec{0}, 0) \neq (v, \delta) \in N_{\text{epi}(f)}(x', t') \right\} ,$$

where here, normality is with respect to the inner product that assigns pairs (x_1, t_1) , (x_2, t_2) the value $\langle x_1, x_2 \rangle + t_1 t_2$.

If $(x', t') \in \text{bdy}(\text{dom}(f))$ and $x' \in \text{int}(\text{dom}(f))$ – hence $t' = f(x')$ – a more concrete description applies:

$$\begin{aligned} & x' \in \text{int}(\text{dom}(f)) \\ \Rightarrow G_2(x', f(x')) &= \left\{ \frac{1}{\hat{f} - (f(x') + \langle v, \bar{e} - x' \rangle)} v \mid v \in \partial f(x') \right\} , \end{aligned}$$

where $\partial f(x')$ is the subdifferential of f at x' (set of subgradients). We rely only on sets $G_2(x', t')$ for which $x' \in \text{Feas}$ – thus, if $\text{Feas} \cap \text{dom}(f) \subseteq \text{int}(\text{dom}(f))$, the more concrete description applies. (In the traditional literature on subgradient methods, the stronger condition $\text{Feas} \subset \text{int}(\text{dom}(f))$ is generally assumed, and hence the more concrete description certainly applies.)

Finally, letting $\text{bdy}_1 := \{(x', t') \mid x' \in \text{bdy}(S)\}$ and $\text{bdy}_2 := \text{bdy}(\text{epi}(f))$, define for $(x', t') \in \text{bdy}_1 \cup \text{bdy}_2$,

$$G(x', t') := \begin{cases} G_1(x') & \text{if } (x', t') \in \text{bdy}_1 \setminus \text{bdy}_2, \\ G_2(x', t') & \text{if } (x', t') \in \text{bdy}_2 \setminus \text{bdy}_1, \\ \text{hull}(G_1(x') \cup G_2(x', t')) & \text{if } (x', t') \in \text{bdy}_1 \cap \text{bdy}_2, \end{cases}$$

where “hull” denotes the convex hull.

Let P denote orthogonal projection onto the kernel of A .

Following is a representative algorithm that arises from the framework developed in subsequent sections.

Algorithm A

- (0) Input: $0 < \epsilon < 1$,
 $\bar{e} \in \text{int}(S \cap \text{dom}(f))$ satisfying $A\bar{e} = b$, and
 \hat{f} , a scalar satisfying $\hat{f} > f(\bar{e})$.
Initialize: $(x_0, t_0) := (\bar{e}, f(\bar{e}))$ and $(x'_0, t'_0) := (\bar{e}, f(\bar{e}))$.
(1) Iterate: Compute $\tilde{x}_{k+1} := x_k - \frac{\epsilon}{2\|P(g')\|^2} P(g')$, where $g' \in G(x'_k, t'_k)$.
Let $\alpha_{k+1} := \alpha(\tilde{x}_{k+1}, t_k)$ and $(x'_{k+1}, t'_{k+1}) := \pi'(\tilde{x}_{k+1}, t_k)$.
If $\alpha_{k+1} \geq 4/3$, then let $(x_{k+1}, t_{k+1}) := (x'_{k+1}, t'_{k+1})$,
else let $(x_{k+1}, t_{k+1}) := (\tilde{x}_{k+1}, t_k)$.
-

Critical to understanding the algorithm is the identity $(x'_k, t'_k) = \pi'(x_k, t_k)$. Iterates in the sequence $\{x_k\}$ need not be feasible nor in the domain of f , whereas x'_k is feasible and in the domain.

Unlike traditional subgradient methods, no orthogonal projections onto Feas are required. The only projections are onto the kernel of A , projections which are computed efficiently if the range space is low dimensional, and if as a preprocessing step the linear operator $(AA^*)^{-1}$ is computed and stored in memory (where A^* is the adjoint of A , and where we are assuming, of course, that A is surjective, for the inverse to exist).

Theorem 1.1. *For the feasible sequence $\{x'_k\}$ generated by Algorithm A,*

$$\ell \geq 8 \left(\frac{\bar{D}}{\hat{r}} \right)^2 \left(\frac{1}{\epsilon^2} + \frac{1}{\epsilon} \log_{4/3} \left(1 + \frac{\bar{D}}{\hat{r}} \right) \right) \quad (1.4)$$

$$\Rightarrow \min_{k \leq \ell} \frac{f(x'_k) - f^*}{\hat{f} - f^*} \leq \epsilon. \quad (1.5)$$

This result differs markedly from the traditional literature, in that there appears no Lipschitz constant for f . The theorem applies even when for no neighborhood of an optimal solution does a (local) Lipschitz constant exist, such as occurs, for example, when $\text{Feas} = \mathbb{R}^2$ and

$$f(x_1, x_2) = \begin{cases} 0 & \text{if } x = \vec{0}, \\ x_1^2 + x_2^2/x_1 & \text{if } x_1 > 0, \\ \infty & \text{otherwise.} \end{cases}$$

(For every $r > 0$, the function fails to be Lipschitz continuous on the relatively-open set $\{x \in \text{dom}(f) \mid \|x\| < r\}$.)

The equivalent problem is hidden from view, the problem to which a standard subgradient method is applied and then translated as Algorithm A. The objective function for the equivalent problem has Lipschitz constant bounded above by $1/\hat{r}$. Thus, a Lipschitz constant does in fact appear in (1.4), even though f need not be Lipschitz continuous.

A tradeoff to gaining independence from Lipschitz continuity is that the error (1.5) is measured relatively rather than absolutely. Dependence on $1/\epsilon^2$ occurs both in (1.4) and in the traditional literature.

The criterion “ $\alpha_{k+1} \geq 4/3$ ” is satisfied at most $\log_{4/3}(1 + \bar{D}/\hat{r})$ times, leading to the logarithmic term in (1.4). The primary importance of using the criterion to create two cases is due to the accuracy needed in solving the equivalent problem being dependent not only on ϵ , but also on f^* . Use of the criterion provides a means to compensate for not knowing f^* .

When f^* is known, a streamlined algorithm can be devised. This algorithm does not require ϵ as input.

Algorithm B

- (0) Input: f^* , the optimal value,
 $\bar{e} \in \text{int}(S \cap \text{dom}(f))$ satisfying $A\bar{e} = b$, and
 \hat{f} , a scalar satisfying $\hat{f} > f(\bar{e})$.
 Initialize: $x_0 := \bar{e}$, $(x'_0, t'_0) := (\bar{e}, f(\bar{e}))$ and $\alpha_0 = \alpha(x_0, f^*)$
- (1) Iterate: Compute $x_{k+1} := x_k + \frac{\alpha_k - 1}{\alpha_k \|P(g')\|^2} P(g')$, where $g' \in G(x'_k, t'_k)$.
 Let $(x'_{k+1}, t'_{k+1}) := \pi'(x_{k+1}, f^*)$ and $\alpha_{k+1} := \alpha(x_{k+1}, f^*)$.
-

Theorem 1.2. *For the feasible sequence $\{x'_k\}$ generated by Algorithm B, and for $0 < \epsilon < 1$,*

$$\begin{aligned} \ell \geq 4 \left(\frac{\bar{D}}{\hat{r}} \right)^2 \cdot \left(\frac{4}{3} \left(\frac{1-\epsilon}{\epsilon} \right)^2 + 4 \left(\frac{1-\epsilon}{\epsilon} \right) \right. \\ \left. + \log_2 \left(\frac{1-\epsilon}{\epsilon} \right) + \log_2 \left(\frac{\bar{D}}{\hat{r}} \right) + 1 \right) \\ \Rightarrow \min_{k \leq \ell} \frac{f(x'_k) - f^*}{\hat{f} - f^*} \leq \epsilon, \end{aligned} \tag{1.6}$$

When S is polyhedral and f is piecewise linear, then in addition to the implication (1.6), for some constants C_1 and C_2 there holds

$$\ell \geq C_1 \log(1/\epsilon) + C_2 \quad \Rightarrow \quad \min_{k \leq \ell} \frac{f(x'_k) - f^*}{\hat{f} - f^*} \leq \epsilon, \tag{1.7}$$

i.e., linear convergence. The constants, however, are highly dependent on Feas and f , making the linear convergence mostly a curiosity.

In the following four sections, focus is exclusively on convex conic linear optimization problems. In the first of these sections (§2), the equivalent problem is

explained and the key theory is developed. Perhaps most surprising is that the transformation to an equivalent problem is simple and the theory is elementary.

Representative algorithms for the conic setting are developed and analyzed in Sections 4 and 5, first in the ideal case where the optimal value is known, and then generally. (Groundwork for the development and analysis is laid in §3)

In Section 6, the general convex optimization problem (1.1) is recast into conic form, to which algorithms and analyses from earlier sections can be applied. The consequent complexity results are stated directly in terms of the original optimization problem (1.1), but the algorithms remain partially abstract in that some key computations are not expressed directly in terms of the original problem.

Finally, in Section 7, the remaining ties to the original optimization problem are established, allowing the algorithms applied to the conic recasting to instead be entirely expressed in terms of the original problem, resulting in Algorithms A and B above. The paper closes with the proofs of Theorems 1.1 and 1.2, and the proof of the claim regarding linear convergence (i.e., (1.7)).

We do not discuss how to compute appropriate input \bar{e} when such a point exists but is unknown, because at present we do not see with high generality how to *cleanly* apply the framework to accomplish this, let alone see how to apply the framework to compute a “good” choice for \bar{e} from among the points in $\text{int}(S \cap \text{dom}(f))$. (Still, compared to the traditional literature where it is assumed $\text{Feas} \subset \text{int}(\text{dom}(f))$ and (arbitrary) points can readily be orthogonally projected onto Feas , assuming \bar{e} is known and line-searches can be done seems to us considerably less restrictive.)

The core of this paper, §§2–6, is virtually identical with our arXiv posting [5], a paper focused on convex conic linear optimization problems rather than general convex optimization problems, and having only at the end a brief discussion of the relevance to general convex optimization. (Motivated by reviewers’ comments, we decided in revising the work for publication to give primary focus to general convex optimization, while still utilizing the conic setting as the natural venue for developing the equivalent problem to which traditional subgradient methods are applied.) Closely related to [5] is recent work of Freund and Lu [1], who develop first-order methods for convex optimization problems in which the objective function has a known lower bound and is assumed to satisfy a Lipschitz condition. Their perspective provides an interesting juxtaposition to the conic-oriented presentation in [5]. Also related is [4], in which accelerated methods for hyperbolic programming are presented.

We close the introduction by emphasizing the intent is to develop a framework, not to prescribe specific algorithms. Different subgradient methods applied in the framework yield different algorithms. Changing how a general optimization problem is recast into conic form also results in a different algorithm.

2. Key Theory

The key theory is elementary and yet has been overlooked in the literature, a blind spot.

Continue to let \mathcal{E} denote a finite-dimensional Euclidean space.

Let $\mathcal{K} \subset \mathcal{E}$ be a proper, closed, convex cone with nonempty interior.

Fix a vector $e \in \text{int}(\mathcal{K})$. We refer to e as the “distinguished direction.” For each $x \in \mathcal{E}$, let

$$\lambda_{\min}(x) := \inf\{\lambda \mid x - \lambda e \notin \mathcal{K}\},$$

that is, the scalar λ for which $x - \lambda e$ lies in the boundary of \mathcal{K} . (Existence and uniqueness of $\lambda_{\min}(x)$ follows from $e \in \text{int}(\mathcal{K}) \neq \mathcal{E}$ and the assumption that \mathcal{K} is a closed, convex cone.)

If, for example, $\mathcal{E} = \mathbb{S}^n$ ($n \times n$ symmetric matrices), $\mathcal{K} = \mathbb{S}_+^n$ (cone of positive semidefinite matrices), and $e = I$ (the identity), then $\lambda_{\min}(X)$ is the minimum eigenvalue of X .

On the other hand, if $\mathcal{K} = \mathbb{R}_+^n$ (non-negative orthant) and e is a vector with all positive coordinates, then $\lambda_{\min}(x) = \min_j x_j/e_j$ for $x \in \mathbb{R}^n$. Clearly, the value of $\lambda_{\min}(x)$ depends on the distinguished direction e (a fact the reader should keep in mind since the notation does not reflect the dependence).

Obviously, $\mathcal{K} = \{x \mid \lambda_{\min}(x) \geq 0\}$ and $\text{int}(\mathcal{K}) = \{x \mid \lambda_{\min}(x) > 0\}$. Also,

$$\lambda_{\min}(sx + te) = s\lambda_{\min}(x) + t \quad \text{for all } x \in \mathcal{E} \text{ and scalars } s \geq 0, t. \quad (2.1)$$

Let

$$\bar{\mathcal{B}} := \{v \in \mathcal{E} \mid e + v \in \mathcal{K} \text{ and } e - v \in \mathcal{K}\},$$

a closed, centrally-symmetric, convex set with nonempty interior. Define the gauge ([6, §15]) on \mathcal{E} by

$$\|u\|_{\infty} := \inf\{t \geq 0 \mid u = tv \text{ for some } v \in \bar{\mathcal{B}}\},$$

which is easily shown to be a norm if and only if \mathcal{K} is a pointed cone. Let $\bar{B}_{\infty}(x, r)$ denote the closed “ball” centered at x and of radius r . Clearly, $\bar{B}_{\infty}(0, 1) = \bar{\mathcal{B}}$, and $\bar{B}_{\infty}(e, 1)$ is the largest subset of \mathcal{K} that has symmetry point e , i.e., for each v , either both points $e + v$ and $e - v$ are in the set, or neither point is in the set.

Proposition 2.1. *The function $x \mapsto \lambda_{\min}(x)$ is concave and Lipschitz continuous:*

$$|\lambda_{\min}(x) - \lambda_{\min}(y)| \leq \|x - y\|_{\infty} \quad \text{for all } x, y \in \mathcal{E}.$$

Proof: Concavity follows easily from the convexity of \mathcal{K} , so we focus on establishing Lipschitz continuity.

Let $x, y \in \mathcal{E}$. According to (2.1), the difference $\lambda_{\min}(x + te) - \lambda_{\min}(y + te)$ is independent of t , and of course so is the quantity $\|(x + te) - (y + te)\|_{\infty}$. Consequently, in proving the Lipschitz continuity, we may assume x lies in the boundary of \mathcal{K} , that is, we may assume $\lambda_{\min}(x) = 0$. The goal, then, is to prove

$$|\lambda_{\min}(x + v)| \leq \|v\|_{\infty} \quad \text{for all } v \in \mathcal{E}. \quad (2.2)$$

We consider two cases. First assume $x + v$ does not lie in the interior of \mathcal{K} , that is, assume $\lambda_{\min}(x + v) \leq 0$. Then, to establish (2.2), it suffices to show $\lambda_{\min}(x + v) \geq -\|v\|_{\infty}$, that is, to show

$$x + v + \|v\|_{\infty} e \in \mathcal{K}. \quad (2.3)$$

However,

$$v + \|v\|_{\infty} e \in \bar{B}_{\infty}(\|v\|_{\infty} e, \|v\|_{\infty}) \subseteq \mathcal{K}, \quad (2.4)$$

the set containment due to \mathcal{K} being a cone and, by construction, $\bar{B}_{\infty}(e, 1) \subseteq \mathcal{K}$. Since $x \in \mathcal{K}$ (indeed, x is in the boundary of \mathcal{K}), (2.3) follows.

Now consider the case $x + v \in \mathcal{K}$, i.e., $\lambda_{\min}(x + v) \geq 0$. To establish (2.2), it suffices to show $\lambda_{\min}(x + v) \leq \|v\|_{\infty}$, that is, to show

$$x + v - \|v\|_{\infty} e \notin \text{int}(\mathcal{K}).$$

Assume otherwise, that is, assume

$$x = w + \|v\|_{\infty} e - v \quad \text{for some } w \in \text{int}(\mathcal{K}).$$

Since $\|v\|_{\infty} e - v \in \mathcal{K}$ (by the set containment on the right of (2.4)), it then follows that $x \in \text{int}(\mathcal{K})$, a contradiction to x lying in the boundary of \mathcal{K} . \square

In this section, the inner product is used primarily for expressing a conic optimization problem. To allow a distinction between inner products, we denote evaluation for the “modeling inner product” on a pair $u, v \in \mathcal{E}$ by $u \cdot v$, whereas in later sections we denote – as was done implicitly in the introduction – evaluation for the “computational inner product” by $\langle u, v \rangle$.

Let $\text{Affine} \subseteq \mathcal{E}$ be an affine space, i.e., the translate of a subspace. For fixed $c \in \mathcal{E}$, consider the conic program

$$\left. \begin{array}{l} \inf \quad c \cdot x \\ \text{s.t.} \quad x \in \text{Affine} \\ \quad \quad x \in \mathcal{K}. \end{array} \right\} \text{CP}$$

Assume $\text{Affine} \cap \text{int}(\mathcal{K})$ – the set of strictly feasible points – is nonempty. Let z^* denote the optimal value.

Assume c is not orthogonal to the subspace of which Affine is a translate, since otherwise all feasible points are optimal. This assumption implies that all optimal solutions for CP lie in the boundary of \mathcal{K} .

Fix a strictly feasible point, e . The point e serves as the distinguished direction.

For scalars $z \in \mathbb{R}$, let

$$\text{Affine}_z := \{x \in \text{Affine} \mid c \cdot x = z\}, \text{ and}$$

let \mathcal{L} denote the subspace of which these affine spaces are translates.

Presently we show that for any choice of z satisfying $z < c \cdot e$, CP can be easily transformed into an equivalent optimization problem in which the only constraint is $x \in \text{Affine}_z$. We make a simple observation.

Lemma 2.2. *Assume CP has bounded optimal value.*

If $x \in \text{Affine}$ satisfies $c \cdot x < c \cdot e$, then $\lambda_{\min}(x) < 1$.

Proof: It follows from (2.1) that if $\lambda_{\min}(x) \geq 1$, then $e + t(x - e)$ is feasible for all $t \geq 0$. As the function $t \mapsto c \cdot (e + t(x - e))$ is strictly decreasing (because $c \cdot x < c \cdot e$), this implies CP has unbounded optimal value, contrary to assumption. \square

For $x \in \mathcal{E}$ satisfying $\lambda_{\min}(x) < 1$, let $\pi(x)$ denote the point where the half-line beginning at e in direction $x - e$ intersects the boundary of \mathcal{K} :

$$\pi(x) := e + \frac{1}{1 - \lambda_{\min}(x)}(x - e) \tag{2.5}$$

(to verify correctness of the expression, observe (2.1) implies $\lambda_{\min}(\pi(x)) = 0$). We refer to $\pi(x)$ as the “radial projection” of x .

The proof of the following result is straightforward, but because it is central to our development, we label the result as a theorem.

Theorem 2.3. *Let z be any value satisfying $z < c \cdot e$. If x^* solves*

$$\begin{aligned} & \sup \quad \lambda_{\min}(x) \\ & \text{s.t.} \quad x \in \text{Affine}_z, \end{aligned} \tag{2.6}$$

then $\pi(x^)$ is optimal for CP. Conversely, if π^* is optimal for CP, then $x^* := e + \frac{c \cdot e - z}{c \cdot e - z^*}(\pi^* - e)$ is optimal for (2.6), and $\pi^* = \pi(x^*)$.*

Proof: Fix a value z satisfying $z < c \cdot e$. It is easily proven from the convexity of \mathcal{K} that $x \mapsto \pi(x)$ gives a one-to-one map from Affine_z onto

$$\{\pi \in \text{Affine} \cap \text{bdy}(\mathcal{K}) \mid c \cdot \pi < c \cdot e\}, \tag{2.7}$$

where $\text{bdy}(\mathcal{K})$ denotes the boundary of \mathcal{K} .

For $x \in \text{Affine}_z$, the CP objective value of $\pi(x)$ is

$$\begin{aligned} c \cdot \pi(x) &= c \cdot \left(e + \frac{1}{1 - \lambda_{\min}(x)}(x - e) \right) \\ &= c \cdot e + \frac{1}{1 - \lambda_{\min}(x)}(z - c \cdot e), \end{aligned} \tag{2.8}$$

a strictly-decreasing function of $\lambda_{\min}(x)$. Since the map $x \mapsto \pi(x)$ is a bijection between Affine_z and the set (2.7), the theorem readily follows. \square

CP has been transformed into an equivalent linearly-constrained maximization problem with concave, Lipschitz-continuous objective function. Virtually any subgradient method – rather, *supgradient* method – can be applied to this problem, the main cost per iteration being in computing a supgradient and projecting it onto the subspace \mathcal{L} .

For illustration, we digress to interpret the implications of the development thus far for the linear program


$$\left. \begin{aligned} & \min_{x \in \mathbb{R}^n} \quad c^T x \\ & \text{s.t.} \quad Ax = b \\ & \quad \quad x \geq 0, \end{aligned} \right\} \text{LP}$$

assuming $e = \mathbf{1}$ (the vector of all ones), in which case $\lambda_{\min}(x) = \min_j x_j$, and $\|\cdot\|_\infty$ is the ℓ_∞ norm, i.e., $\|v\|_\infty = \max_j |v_j|$. Let the number of rows of A be $m \geq 1$, assume the rows are linearly independent, and assume c is not a linear combination of the rows (otherwise all feasible points are optimal).

For any scalar $z < c^T \mathbf{1}$, Theorem 2.3 asserts that LP is equivalent to

$$\begin{aligned} & \max_x \quad \min_j x_j \\ & \text{s.t.} \quad Ax = b \\ & \quad \quad c^T x = z, \end{aligned} \tag{2.9}$$

in that when x is feasible for (2.9), x is optimal if and only if the projection $\pi(x) = \mathbf{1} + \frac{1}{1 - \min_j x_j}(x - \mathbf{1})$ is optimal for LP. The setup is shown schematically in the following figure:



Proposition 2.1 asserts that, as is obviously true, $x \mapsto \min_j x_j$ is ℓ_∞ -Lipschitz continuous with constant 1. Consequently, the function also is ℓ_2 -Lipschitz continuous with constant 1, as is relevant if supgradient methods rely on the standard inner product in computing supgradients.

With respect to the standard inner product, the supgradients of $x \mapsto \min_j x_j$ at x are the convex combinations of the standard basis vectors $e(k)$ for which $x_k = \min_j x_j$. Consequently, the projected supgradients at x are the convex combinations of the vectors \bar{P}_k for which $x_k = \min_j x_j$, where \bar{P}_k is the k^{th} column of the matrix projecting \mathbb{R}^n onto the nullspace of $\bar{A} = \begin{bmatrix} A \\ c^T \end{bmatrix}$, that is

$$\bar{P} := I - \bar{A}^T (\bar{A} \bar{A}^T)^{-1} \bar{A}.$$

If $m \ll n$, then \bar{P} is not computed in its entirety, but instead the matrix $\bar{M} = (\bar{A} \bar{A}^T)^{-1}$ is formed as a preprocessing step, at cost $O(m^2 n)$ (the inverse exists because the set consisting of c and the rows of A has been assumed linearly independent). Then, for any iterate x and an index k satisfying $x_k = \min_j x_j$, the projected supgradient \bar{P}_k is computed according to

$$u = \bar{M} \bar{A}_k \quad \rightarrow \quad v = \bar{A}^T u \quad \rightarrow \quad \bar{P}_k = e(k) - v,$$

for a cost of $O(m^2 + \#\text{non_zero_entries_in_}A)$ per iteration.

Before returning to the general theory, we note that if the choices are $\mathcal{E} = \mathbb{S}^n$, $\mathcal{K} = \mathbb{S}_+^n$ and $e = I$ (and thus $\lambda_{\min}(X)$ is the minimum eigenvalue of X), then with respect to the trace inner product, the supgradients at X for the function $X \mapsto \lambda_{\min}(X)$ are the convex combinations of the matrices vv^T , where $Xv = \lambda_{\min}(X)v$ and $\|v\|_2 = 1$.

Assume, henceforth, that CP has at least one optimal solution, and that z is a fixed scalar satisfying $z < c \cdot e$. Then the equivalent problem (2.6) has at least one optimal solution. Let x_z^* denote any of the optimal solutions for the equivalent problem, and recall z^* denotes the optimal value of CP. A useful characterization of the optimal value for the equivalent problem is easily provided.

Lemma 2.4.

$$\lambda_{\min}(x_z^*) = \frac{z - z^*}{c \cdot e - z^*}$$

Proof: By Theorem 2.3, $\pi(x_z^*)$ is optimal for CP – in particular, $c \cdot \pi(x_z^*) = z^*$. Thus, according to (2.8),

$$z^* = c \cdot e + \frac{1}{1 - \lambda_{\min}(x_z^*)} (z - c \cdot e).$$

Rearrangement completes the proof. \square

We focus on the goal of computing a point π which is feasible for CP and has better objective value than e in that

$$\frac{c \cdot \pi - z^*}{c \cdot e - z^*} \leq \epsilon, \quad (2.10)$$

where $0 < \epsilon < 1$. Thus, for the problem of primary interest, CP, the focus is on relative improvement in the objective value.

The following proposition provides a useful characterization of the accuracy needed in approximately solving the CP-equivalent problem (2.6) so as to ensure that for the computed point x , the projection $\pi = \pi(x)$ satisfies (2.10).

Proposition 2.5. *If $x \in \text{Affine}_z$ and $0 < \epsilon < 1$, then*

$$\frac{c \cdot \pi(x) - z^*}{c \cdot e - z^*} \leq \epsilon$$

if and only if

$$\lambda_{\min}(x_z^*) - \lambda_{\min}(x) \leq \frac{\epsilon}{1 - \epsilon} \frac{c \cdot e - z}{c \cdot e - z^*}.$$

Proof: Assume $x \in \text{Affine}_z$. For $y = x$ and $y = x_z^*$, we have the equality (2.8), that is,

$$c \cdot \pi(y) = c \cdot e + \frac{1}{1 - \lambda_{\min}(y)}(z - c \cdot e).$$

Thus,

$$\begin{aligned} \frac{c \cdot \pi(x) - z^*}{c \cdot e - z^*} &= \frac{c \cdot \pi(x) - c \cdot \pi(x_z^*)}{c \cdot e - c \cdot \pi(x_z^*)} \\ &= \frac{\frac{1}{1 - \lambda_{\min}(x)} - \frac{1}{1 - \lambda_{\min}(x_z^*)}}{-\frac{1}{1 - \lambda_{\min}(x_z^*)}} \\ &= \frac{\lambda_{\min}(x_z^*) - \lambda_{\min}(x)}{1 - \lambda_{\min}(x)}. \end{aligned}$$

Hence,

$$\begin{aligned} \frac{c \cdot \pi(x) - z^*}{c \cdot e - z^*} &\leq \epsilon \\ \Leftrightarrow \lambda_{\min}(x_z^*) - \lambda_{\min}(x) &\leq \epsilon(1 - \lambda_{\min}(x)) \\ \Leftrightarrow (1 - \epsilon)(\lambda_{\min}(x_z^*) - \lambda_{\min}(x)) &\leq \epsilon(1 - \lambda_{\min}(x_z^*)) \\ \Leftrightarrow \lambda_{\min}(x_z^*) - \lambda_{\min}(x) &\leq \frac{\epsilon}{1 - \epsilon}(1 - \lambda_{\min}(x_z^*)). \end{aligned}$$

Using Lemma 2.4 to substitute for the rightmost occurrence of $\lambda_{\min}(x_z^*)$ completes the proof. \square

3. Groundwork for Algorithm Design and Analysis

In Sections 4 and 5, we show how the key theory leads to algorithms and complexity results regarding the solution of the conic program CP. In this section, groundwork is laid for the development in those and subsequent sections.

Continue to assume CP has an optimal solution, denote the optimal value by z^* , and let e be a strictly feasible point, the distinguished direction. Given $\epsilon > 0$ and a value z satisfying $z < c \cdot e$, the approach is to apply supgradient methods to approximately solve

$$\begin{aligned} \max \quad & \lambda_{\min}(x) \\ \text{s.t.} \quad & x \in \text{Affine}_z, \end{aligned} \tag{3.1}$$

where by ‘‘approximately solve’’ we mean that $x \in \text{Affine}_z$ is computed for which

$$\lambda_{\min}(x_z^*) - \lambda_{\min}(x) \leq \frac{\epsilon}{1 - \epsilon} \frac{c \cdot e - z}{c \cdot e - z^*}.$$

Indeed, according to Proposition 2.5, the projection $\pi = \pi(x)$ will then satisfy

$$\frac{c \cdot \pi - z^*}{c \cdot e - z^*} \leq \epsilon.$$

Not until results regarding CP are applied to general convex optimization – Sections 6 and 7 – do we require a characterization of supgradients of the function $x \mapsto \lambda_{\min}(x)$, but as the characterization is interesting in itself, we present it as the first piece of groundwork laid in this section.

Let $P_{\mathcal{L}}$ denote orthogonal projection onto \mathcal{L} with respect to $\langle \cdot, \cdot \rangle$ (the computational inner product), and let $\hat{\partial}\lambda_{\min}(x)$ denote the supdifferential at x . The following proposition relates supdifferentials to normal cones (normal with respect to $\langle \cdot, \cdot \rangle$).

Proposition 3.1. *For all $x \in \mathcal{E}$,*

$$\hat{\partial}\lambda_{\min}(x) = \{v \mid -v \in N_{\mathcal{K}}(x - \lambda_{\min}(x)e) \text{ and } \langle e, v \rangle = 1\}.$$

Proof: We know

$$\lambda_{\min}(y + re) = \lambda_{\min}(y) + t \quad \text{for all } y \in \mathcal{E}, t \in \mathbb{R}, \tag{3.2}$$

from which follows for every x and t ,

$$\hat{\partial}\lambda_{\min}(x + te) = \hat{\partial}\lambda_{\min}(x) \tag{3.3}$$

and

$$v \in \hat{\partial}\lambda_{\min}(x) \quad \Rightarrow \quad \langle e, v \rangle = 1. \tag{3.4}$$

Due to (3.3), in proving the lemma we may assume $\lambda_{\min}(x) = 0$, i.e., $x \in \text{bdy}(\mathcal{K})$.

Since $\lambda_{\min}(x) = 0$, and the value of λ_{\min} is non-negative on \mathcal{K} ,

$$g \in \hat{\partial}\lambda_{\min}(x) \quad \Rightarrow \quad \forall y \in \mathcal{K}, \langle g, y - x \rangle \geq 0.$$

Thus, with (3.4), we have

$$\hat{\partial}\lambda_{\min}(x) \subseteq \{v \mid \langle e, v \rangle = 1 \text{ and } -v \in N_{\mathcal{K}}(x)\}. \tag{3.5}$$

On the other hand, if $-v \in N_{\mathcal{K}}(x)$, then

$$y \in \text{bdy}(\mathcal{K}) \quad \Rightarrow \quad 0 = \lambda_{\min}(y) \leq \lambda_{\min}(x) + \langle v, y - x \rangle$$

(using $\lambda_{\min}(x) = 0$). Thus, if in addition, $\langle e, v \rangle = 1$,

$$t \in \mathbb{R} \text{ and } y \in \text{bdy}(\mathcal{K}) \quad \Rightarrow \quad \lambda_{\min}(y + te) \leq \lambda_{\min}(x) + \langle v, (y + te) - x \rangle$$

(using 3.2). Since $\mathcal{E} = \{y + te \mid y \in \text{bdy}(\mathcal{K}) \text{ and } t \in \mathbb{R}\}$, the reverse inclusion to (3.5) thus holds, establishing the proposition. \square

Let $\|\cdot\|$ be the norm associated with $\langle \cdot, \cdot \rangle$, and for $z \in \mathbb{R}$, let

$$M_z := \sup \left\{ \frac{|\lambda_{\min}(x) - \lambda_{\min}(y)|}{\|x - y\|} \mid x, y \in \text{Affine}_z \text{ and } x \neq y \right\},$$

the Lipschitz constant for the map $x \mapsto \lambda_{\min}(x)$ restricted to Affine_z . Proposition 2.1 implies M_z is well-defined (finite), although unlike the Lipschitz constant for the norm appearing there (i.e., $\|\cdot\|_\infty$), M_z might exceed 1, depending on $\|\cdot\|$. The next piece of groundwork to be laid is a geometric characterization of an upper bound on M_z .

We claim the values M_z are identical for all z . To see why, consider that for $z_1, z_2 < c \cdot e$, a bijection from Affine_{z_1} onto Affine_{z_2} is provided by the map

$$x \mapsto y(x) := \frac{z_2 - z_1}{c \cdot e - z_1} e + \frac{c \cdot e - z_2}{c \cdot e - z_1} x.$$

Observe, using (2.1),

$$\lambda_{\min}(y(x)) = \frac{c \cdot e - z_2}{c \cdot e - z_1} \lambda_{\min}(x) + \frac{z_2 - z_1}{c \cdot e - z_1},$$

and thus

$$\lambda_{\min}(y(x)) - \lambda_{\min}(y(\bar{x})) = \frac{c \cdot e - z_2}{c \cdot e - z_1} (\lambda_{\min}(x) - \lambda_{\min}(\bar{x})) \quad \text{for } x, \bar{x} \in \text{Affine}_{z_1}.$$

Since, additionally, $\|y(x) - y(\bar{x})\| = \frac{c \cdot e - z_2}{c \cdot e - z_1} \|x - \bar{x}\|$, it is immediate that the values M_z are identical for all $z < c \cdot e$. A simple continuity argument then implies this value is equal to $M_{c \cdot e}$. Analogous reasoning shows $M_z = M_{c \cdot e}$ for all $z > c \cdot e$. In all, M_z is independent of z , as claimed.

Let M denote the common value, i.e., $M = M_z$ for all z .

The following proposition can be useful in modeling and in choosing the computational inner product. This result plays a central role in our complexity bounds for general convex optimization.

Let $\bar{B}(e, r) := \{x \mid \|x - e\| \leq r\}$.

Proposition 3.2. $M \leq 1/r_e$, where $r_e := \max\{r \mid \bar{B}(e, r) \cap \text{Affine}_{c \cdot e} \subseteq \mathcal{K}\}$

Proof: According to Proposition 2.1,

$$|\lambda_{\min}(x) - \lambda_{\min}(y)| \leq \|x - y\|_\infty \quad \text{for all } x, y.$$

Consequently, it suffices to show $\|x - y\| \geq r_e \|x - y\|_\infty$ for all $x, y \in \text{Affine}_{c \cdot e}$, i.e., it suffices to show for all $v \in \mathcal{L}$ that $\|v\| \geq r_e \|v\|_\infty$.

However, according to the discussion just prior to Proposition 2.1, $\bar{B}_\infty(e, 1)$ is the largest set which both is contained in \mathcal{K} and has symmetry point e , from which follows that $\bar{B}_\infty(e, 1) \cap \text{Affine}_{c \cdot e}$ is the largest set which is both contained in $\mathcal{K} \cap \text{Affine}_{c \cdot e}$ and has symmetry point e . Hence

$$\bar{B}(e, r_e) \cap \text{Affine}_{c \cdot e} \subseteq \bar{B}_\infty(e, 1) \cap \text{Affine}_{c \cdot e},$$

implying $\|v\| \geq r_e \|v\|_\infty$ for all $v \in \mathcal{L}$. □

Towards considering specific supgradient methods, we recall the following standard and elementary result, rephrased for our setting:

Lemma 3.3. Assume $z \in \mathbb{R}$, $x, y \in \text{Affine}_z$ and $g \in \hat{\partial} \lambda_{\min}(x)$.

For all scalars α ,

$$\|(x + \alpha P_{\mathcal{L}}(g)) - y\|^2 \leq \|x - y\|^2 - 2\alpha (\lambda_{\min}(y) - \lambda_{\min}(x)) + \alpha^2 \|P_{\mathcal{L}}(g)\|^2.$$

Proof: Letting $\tilde{g} := P_{\mathcal{L}}(g)$, simply observe

$$\begin{aligned} \|(x + \alpha\tilde{g}) - y\|^2 &= \|x - y\|^2 + 2\alpha\langle\tilde{g}, x - y\rangle + \alpha^2\|\tilde{g}\|^2 \\ &= \|x - y\|^2 - 2\alpha\langle g, y - x\rangle + \alpha^2\|\tilde{g}\|^2 \quad (\text{by } x - y \in \mathcal{L}) \\ &\leq \|x - y\|^2 - 2\alpha(\lambda_{\min}(y) - \lambda_{\min}(x)) + \alpha^2\|\tilde{g}\|^2, \end{aligned}$$

the inequality due to $g \in \hat{\partial}\lambda_{\min}(x)$. \square

4. Algorithm 1: When the Optimal Value is Known

Knowing z^* is not an entirely implausible situation. For example, if strict feasibility holds for a primal conic program and for its dual

$$\begin{array}{ll} \min & \bar{c}^T x \\ \text{s.t.} & \bar{A}x = \bar{b} \\ & x \in \bar{\mathcal{K}} \end{array} \quad \begin{array}{ll} \max & \bar{b}^T y \\ \text{s.t.} & \bar{A}^T y + s = \bar{c} \\ & s \in \bar{\mathcal{K}}^*, \end{array}$$

then the combined primal-dual conic program is known to have optimal value equal to zero:

$$\begin{array}{ll} \min & \bar{c}^T x - \bar{b}^T y \\ \text{s.t.} & \bar{A}x = \bar{b} \\ & \bar{A}^T y + s = \bar{c} \\ & (x, s) \in \bar{\mathcal{K}} \times \bar{\mathcal{K}}^*. \end{array}$$

Algorithm 1

- (0) Input: z^* , the optimal value of CP,
 e , a strictly feasible point for CP, and
 $\bar{x} \in \text{Affine}$ satisfying $c \cdot \bar{x} < c \cdot e$.
- Initialize: Let $x_0 = e + \frac{c \cdot e - z^*}{c \cdot e - c \cdot \bar{x}}(\bar{x} - e)$ (thus, $c \cdot x_0 = z^*$),
and let $\pi_0 = \pi(x_0)$ ($= \pi(\bar{x})$).
- (1) Iterate: Compute $x_{k+1} = x_k - \frac{\lambda_{\min}(x_k)}{\|P_{\mathcal{L}}(g_k)\|^2} P_{\mathcal{L}}(g_k)$, where $g_k \in \hat{\partial}\lambda_{\min}(x_k)$.
Let $\pi_{k+1} = \pi(x_{k+1})$.
-

All of the iterates x_k lie in Affine_{z^*} , and hence, $\lambda_{\min}(x_k) \leq 0$, with equality if and only if x_k is feasible (and optimal) for CP.

For all scalars $z < c \cdot e$ and for $x \in \text{Affine}_z$, define

$$\text{dist}_z(x) := \min\{\|x - x_z^*\| \mid x_z^* \text{ is optimal for (3.1)}\}.$$

Proposition 4.1. *The iterates for Algorithm 1 satisfy*

$$\max\{\lambda_{\min}(x_k) \mid k = \ell, \dots, \ell + m\} \geq -M \text{dist}_{z^*}(x_\ell) / \sqrt{m + 1}.$$

Proof: Letting $\tilde{g}_k := P_{\mathcal{L}}(g_k)$, Lemma 3.3 implies

$$\begin{aligned} &\text{dist}_{z^*}(x_{k+1})^2 \\ &\leq \text{dist}_{z^*}(x_k)^2 - 2(-\lambda_{\min}(x_k) / \|\tilde{g}_k\|^2)(0 - \lambda_{\min}(x_k)) + (\lambda_{\min}(x_k) / \|\tilde{g}_k\|)^2 \\ &= \text{dist}_{z^*}(x_k)^2 - (\lambda_{\min}(x_k) / \|\tilde{g}_k\|)^2, \end{aligned}$$

and thus by induction (and using $\|\tilde{g}_k\| \leq M$),

$$\begin{aligned} \text{dist}_{z^*}(x_{\ell+m+1})^2 &\leq \text{dist}_{z^*}(x_\ell)^2 - \sum_{k=\ell}^{\ell+m} (\lambda_{\min}(x_k)/M)^2 \\ &\leq \text{dist}_{z^*}(x_\ell)^2 - \frac{m+1}{M^2} \min\{\lambda_{\min}(x_k)^2 \mid k = \ell, \dots, \ell + m\}, \end{aligned}$$

implying the proposition (keeping in mind $\lambda_{\min}(x_k) \leq 0$). \square

We briefly digress to consider the case of \mathcal{K} being polyhedral, where already an interesting result is easily proven. The following corollary is offered only as a curiosity, as the constants typically are so large as to render the bound on ℓ meaningless except for minuscule ϵ .

Corollary 4.2. *Assume \mathcal{K} is polyhedral. There exist constants C_1 and C_2 (dependent on CP, e , \bar{x} and the computational inner product), such that for all $0 < \epsilon < 1$,*

$$\ell \geq C_1 + C_2 \log(1/\epsilon) \quad \Rightarrow \quad \min_{k \leq \ell} \frac{c \cdot \pi_k - z^*}{c \cdot e - z^*} \leq \epsilon.$$

For first-order methods, such a logarithmic bound in ϵ was initially established by Gilpin, Peña and Sandholm [2]. They did not assume an initial feasible point e was known, but neither did they require the computed approximate solution to be feasible (instead, constraint residuals were required to be small). They relied on an accelerated gradient method, along with the smoothing technique of Nesterov [3]. As is the case for the above result, they assumed the optimal value of CP to be known apriori, and they restricted \mathcal{K} to be polyhedral.

The proof of the corollary depends on the following simple lemma.

Lemma 4.3. *For Algorithm 1, the iterates satisfy*

$$\frac{c \cdot \pi_k - z^*}{c \cdot e - z^*} = \frac{-\lambda_{\min}(x_k)}{1 - \lambda_{\min}(x_k)}.$$

Proof: Immediate from $\pi_k = e + \frac{1}{1 - \lambda_{\min}(x_k)}(x_k - e)$ and $c \cdot x_k = z^*$. \square

Proof of Corollary 4.2: With \mathcal{K} being polyhedral, the concave function $x \mapsto \lambda_{\min}(x)$ is piecewise linear, and thus there exists a positive constant C such that

$$\text{dist}_{z^*}(x) \leq -C \lambda_{\min}(x) \quad \text{for all } x \in \text{Affine}_{z^*}.$$

Then Proposition 4.1 gives

$$\max\{\lambda_{\min}(x_k) \mid k = \ell, \dots, \ell + m\} \geq C M \lambda_{\min}(x_\ell) / \sqrt{m+1},$$

from which follows

$$\max\{\lambda_{\min}(x_k) \mid k = \ell, \dots, \ell + \lceil (2CM)^2 \rceil\} \geq \frac{1}{2} \lambda_{\min}(x_\ell),$$

i.e., $\lambda_{\min}(x_\ell)$ is “halved” within $\lceil (2CM)^2 \rceil$ iterations. The proof is easily completed using Lemma 4.3. \square

We now return to considering general convex cones \mathcal{K} .

The iteration bound provided by Proposition 4.1 bears little obvious connection to the geometry of the conic program CP, except in that the constant M is related to the geometry by Proposition 3.2. The other constant $-\text{dist}_{z^*}(x_k)$ – does not

at present have such a clear geometrical connection to CP. We next observe a meaningful connection.

The *level sets* for CP are the sets

$$\text{Level}_z = \text{Affine}_z \cap \mathcal{K} ,$$

that is, the largest feasible sets for CP on which the objective function is constant¹. If $z < z^*$, then $\text{Level}_z = \emptyset$.

If some level set is unbounded, then either CP has unbounded optimal value or can be made to have unbounded value with an arbitrarily small perturbation of c . Thus, in developing numerical optimization methods, it is natural to focus on the case that level sets for CP are bounded.

For scalars z , define the diameter of Level_z by

$$\text{diam}_z := \sup\{\|x - y\| \mid x, y \in \text{Level}_z\} ,$$

the diameter of Level_z . If $\text{Level}_z = \emptyset$, let $\text{diam}_z := -\infty$.

Lemma 4.4. *Assume $x \in \text{Affine}_{z^*}$, and let $\pi = \pi(x)$. Then*

$$\text{dist}_{z^*}(x) = (1 - \lambda_{\min}(x)) \text{dist}_{c \cdot \pi}(\pi) = \frac{\text{dist}_{c \cdot \pi}(\pi)}{1 - \frac{c \cdot \pi - z^*}{c \cdot e - z^*}} \leq \frac{\text{diam}_{c \cdot \pi}}{1 - \frac{c \cdot \pi - z^*}{c \cdot e - z^*}} .$$

Proof: Since

$$\pi = e + \frac{1}{1 - \lambda_{\min}(x)}(x - e) , \quad (4.1)$$

Theorem 2.3 implies that the maximizers of the map $y \mapsto \lambda_{\min}(y)$ over $\text{Affine}_{c \cdot \pi}$ are precisely the points of the form

$$x_{c \cdot \pi}^* = e + \frac{1}{1 - \lambda_{\min}(x)}(x_{z^*}^* - e) ,$$

where $x_{z^*}^*$ is a maximizer of the map when restricted to Affine_{z^*} (i.e., is an optimal solution of CP). Observing

$$\pi - x_{c \cdot \pi}^* = \frac{1}{1 - \lambda_{\min}(x)}(x - x_{z^*}^*) ,$$

it follows that

$$\text{dist}_{c \cdot \pi}(\pi) = \frac{1}{1 - \lambda_{\min}(x)} \text{dist}_{z^*}(x) ,$$

establishing the first equality in the statement of the lemma. The second equality then follows easily from (4.1) and $c \cdot x = z^*$. The inequality is due simply to $\pi, x_{c \cdot \pi}^* \in \text{Level}_{c \cdot \pi}$, for all optimal solutions $x_{c \cdot \pi}^*$ of the CP-equivalent problem (3.1) (with $z = c \cdot \pi$). \square

For scalars z , define

$$\text{Diam}_z := \max\{\text{diam}_{z'} \mid z' \leq z\} ,$$

the ‘‘horizontal diameter’’ of the sublevel set consisting of points x that are feasible for CP and satisfy $c \cdot x \leq z$. For $z^* < z < c \cdot e$, the value Diam_z can be thought of as a kind of condition number for CP, because Diam_z being large is an indication that

¹There is possibility of confusion here, as in the optimization literature, the terminology ‘‘level set’’ is often used for the portion of the feasible region on which the (convex) objective function does not exceed a specified value rather than – as for us – exactly equals the value. Our terminology is consistent with the general mathematical literature, where the region on which a function does not exceed a specified value is referred to as a sublevel set, not a level set.

the optimal value for CP is relatively sensitive to perturbations in the objective vector c .

For $z^* \leq z < c \cdot e$, define

$$\text{Dist}_z := \sup\{\text{dist}_{z'}(x) \mid z' \leq z \text{ and } x \in \text{Level}_{z'}\}.$$

Clearly, there holds the relation

$$\text{Dist}_z \leq \text{Diam}_z,$$

and hence if the “condition number” Diam_z is only of modest size, so is the value Dist_z .

Following is our main result for Algorithm 1. By substituting $\text{Diam}_{c \cdot \pi_0}$ for $\text{Dist}_{c \cdot \pi_0}$, and $1/r_e$ for M (where r_e is as in Proposition 3.2), the statement of the theorem becomes phrased in terms clearly reflecting the geometry of CP.

Theorem 4.5. *Assume $0 < \epsilon < \frac{c \cdot \pi_0 - z^*}{c \cdot e - z^*}$, where $\pi_0 = \pi(x_0)$ is the initial CP-feasible point for Algorithm 1 (i.e., assume π_0 does not itself satisfy the desired accuracy). Then*

$$\begin{aligned} \ell &\geq (2M \text{Dist}_{c \cdot \pi_0})^2 \left(\frac{4}{3} \left(\frac{1 - \epsilon}{\epsilon} \right)^2 + 4 \left(\frac{1 - \epsilon}{\epsilon} \right) \right. \\ &\quad \left. + \log_2 \left(\frac{c \cdot \pi_0 - z^*}{c \cdot e - z^*} \right) + \log_2 \left(\frac{1 - \epsilon}{1 - \frac{c \cdot \pi_0 - z^*}{c \cdot e - z^*}} \right) + 1 \right) \\ &\Rightarrow \min_{k \leq \ell} \frac{c \cdot \pi_k - z^*}{c \cdot e - z^*} \leq \epsilon. \end{aligned}$$

Proof: To ease notation, let $\lambda_k := \lambda_{\min}(x_k)$.

Let $k_0 = 0$ and recursively define k_{i+1} to be the first index for which $\lambda_{k_{i+1}} \geq \lambda_{k_i}/2$ (keeping in mind $\lambda_k \leq 0$ for all k). Proposition 4.1 implies

$$\begin{aligned} k_{i+1} - k_i + 1 &\leq \left(\frac{2M \text{dist}_{z^*}(x_{k_i})}{\lambda_{k_i}} \right)^2 \\ &= \left(2M \text{dist}_{c \cdot \pi_{k_i}}(\pi_{k_i}) \frac{1 - \lambda_{k_i}}{\lambda_{k_i}} \right)^2 \quad (\text{by Lemma 4.4}) \\ &\leq \left(2M \text{Dist}_{c \cdot \pi_0} \frac{1 - \lambda_{k_i}}{\lambda_{k_i}} \right)^2, \end{aligned} \tag{4.2}$$

where the final inequality is due to $c \cdot \pi_{k_i}$ ($i = 0, 1, \dots$) being a decreasing sequence (using Lemma 4.3).

Let i' be the first sub-index for which $\lambda_{k_{i'}} \geq -\epsilon/(1 - \epsilon)$. Lemma 4.3 implies

$$\frac{c \cdot \pi_{k_{i'}} - z^*}{c \cdot e - z^*} \leq \epsilon.$$

Thus, to prove the theorem, it suffices to show $\ell = k_{i'}$ satisfies the inequality in the statement of the theorem.

Note $i' > 0$ (because, by assumption, $\epsilon < \frac{c \cdot \pi_0 - z^*}{c \cdot e - z^*}$). Observe, then,

$$\begin{aligned} i' &< 1 + \log_2 \left(\frac{\lambda_0}{-\epsilon / (1 - \epsilon)} \right) \\ &= 1 + \log_2 \left(\frac{\frac{c \cdot \pi_0 - z^*}{c \cdot e - z^*}}{\epsilon} \right) + \log_2 \left(\frac{1 - \epsilon}{1 - \frac{c \cdot \pi_0 - z^*}{c \cdot e - z^*}} \right) \end{aligned} \quad (4.3)$$

(again using Lemma 4.3).

Additionally,

$$\begin{aligned} k_{i'} &= \sum_{i=0}^{i'-1} k_{i+1} - k_i \\ &\leq (2M \text{Dist}_{c \cdot \pi_0})^2 \sum_{i=0}^{i'-1} \left(\frac{1 - \lambda_{k_i}}{\lambda_{k_i}} \right)^2 \quad (\text{by (4.2)}) \\ &\leq (2M \text{Dist}_{c \cdot \pi_0})^2 \sum_{i=0}^{i'-1} \left(\frac{1 - 2^i \lambda_{k_{i'-1}}}{2^i \lambda_{k_{i'-1}}} \right)^2 \\ &\leq (2M \text{Dist}_{c \cdot \pi_0})^2 \sum_{i=0}^{i'-1} \left(\frac{1 + 2^i \epsilon / (1 - \epsilon)}{2^i \epsilon / (1 - \epsilon)} \right)^2 \\ &= (2M \text{Dist}_{c \cdot \pi_0})^2 \sum_{i=0}^{i'-1} \left(1 + \frac{1}{2^i} \frac{1 - \epsilon}{\epsilon} \right)^2 \\ &\leq (2M \text{Dist}_{c \cdot \pi_0})^2 \left(i' + 4 \frac{1 - \epsilon}{\epsilon} + \frac{4}{3} \left(\frac{1 - \epsilon}{\epsilon} \right)^2 \right). \end{aligned}$$

Using (4.3) to substitute for i' completes the proof. \square

5. Algorithm 2: When The Optimal Value Is Unknown

For the second algorithm, we discard the requirement of knowing z^* . Now ϵ (the desired relative-accuracy) is required as input.

Algorithm 2

- (0) Input: $0 < \epsilon < 1$,
 e , a strictly feasible point for CP, and
 $\bar{x} \in \text{Affine}$ satisfying $c \cdot \bar{x} < c \cdot e$.
- Initialize: $x_0 = \pi_0 = \pi(\bar{x})$
- (1) Iterate: Compute $\tilde{x}_{k+1} := x_k + \frac{\epsilon}{2 \|P_{\mathcal{L}} g_k\|^2} P_{\mathcal{L}} g_k$, where $g_k \in \hat{\partial} \lambda_{\min}(x_k)$.
Let $\pi_{k+1} := \pi(\tilde{x}_{k+1})$.
If $c \cdot (e - \pi_{k+1}) \geq \frac{4}{3} c \cdot (e - \tilde{x}_{k+1})$, let $x_{k+1} = \pi_{k+1}$;
else, let $x_{k+1} = \tilde{x}_{k+1}$.
-

Unsurprisingly, the iteration bound we obtain for Algorithm 2 is worse than the result for Algorithm 1, but perhaps surprisingly, the bound is not excessively worse, in that the factor for $1/\epsilon^2$ is essentially unchanged (it's the factor for $1/\epsilon$ that increases, although typically not by a large amount).

Theorem 5.1. Assume $0 < \epsilon < \frac{c \cdot \pi_0 - z^*}{c \cdot e - z^*}$. For the iterates of Algorithm 2,

$$\begin{aligned} \ell &\geq 8 (M \text{Dist}_{c \cdot \pi_0})^2 \left(\frac{1}{\epsilon^2} + \frac{1}{\epsilon} \log_{4/3} \left(\frac{1}{1 - \frac{c \cdot \pi_0 - z^*}{c \cdot e - z^*}} \right) \right) \\ &\Rightarrow \min_{k \leq \ell} \frac{c \cdot \pi_k - z^*}{c \cdot e - z^*} \leq \epsilon. \end{aligned}$$

Proof: In order to distinguish the iterates obtained by projecting to the boundary, we record a notationally-embellished rendition of the algorithm which introduces a distinction between “inner iterations” and “outer iterations”:

Algorithm 2 (notationally-embellished version):

- (0) **Input:** $0 < \epsilon < 1$, e and \bar{x} .
Initialize: $y_{1,0} = \pi(\bar{x})$,
 $i = 1$ (outer iteration counter),
 $j = 0$ (inner iteration counter).
 - (1) Compute $y_{i,j+1} = y_{i,j} + \frac{\epsilon}{2\|\tilde{g}_{i,j}\|^2} \tilde{g}_{i,j}$,
where $\tilde{g}_{i,j} = P_{\mathcal{L}} g_{i,j}$ and $g_{i,j} \in \hat{\partial} \lambda_{\min}(y_{i,j})$.
 - (2) If $c \cdot (e - \pi(y_{i,j+1})) \geq \frac{4}{3} c \cdot (e - y_{i,j+1})$,
then let $y_{i+1,0} = \pi(y_{i,j+1})$, $i \leftarrow i + 1$ and $j \leftarrow 0$;
else, let $j \leftarrow j + 1$.
 - (3) Go to step 1.
-

For each outer iteration i , all of the iterates $y_{i,j}$ have the same objective value. Denote the value by z_i . Obviously, z_1 is equal to the value $c \cdot \pi_0$ appearing in the statement of the theorem. Let

$$\text{Dist} := \text{Dist}_{c \cdot \pi_0} = \text{Dist}_{z_1}.$$

Step 2 ensures

$$c \cdot e - z_{i+1} \geq \frac{4}{3} (c \cdot e - z_i). \quad (5.1)$$

Thus, z_1, z_2, \dots is a strictly decreasing sequence. Consequently, as $y_{i,0} \in \text{Level}_{z_i}$, we have $\text{dist}_{z_i}(y_{i,0}) \leq \text{Dist}$ for all i .

From (5.1) we find for scalars $\delta > 0$ that

$$\frac{c \cdot e - z_{i+1}}{c \cdot e - z^*} < \delta \quad \Rightarrow \quad i < \log_{4/3} \left(\frac{\delta}{\frac{c \cdot e - z_1}{c \cdot e - z^*}} \right) = \log_{4/3} \left(\frac{\delta}{1 - \frac{z_1 - z^*}{c \cdot e - z^*}} \right),$$

and thus, for $\epsilon < 1$,

$$\frac{z_i - z^*}{c \cdot e - z^*} > \epsilon \quad \Rightarrow \quad i < 1 + \log_{4/3} \left(\frac{1 - \epsilon}{1 - \frac{z_1 - z^*}{c \cdot e - z^*}} \right). \quad (5.2)$$

Hence, if an outer iteration i fails to satisfy the inequality on the right, the initial inner iterate $y_{i,0}$ fulfills the goal of finding a CP-feasible point π satisfying $\frac{c \cdot \pi - z^*}{c \cdot e - z^*} \leq \epsilon$ (i.e., the algorithm has been successful no later than the start of outer iteration i). Also observe that (5.2) provides (letting $\epsilon \downarrow 0$) an upper bound on I , the total number of outer iterations:

$$I \leq 1 + \log_{4/3} \left(\frac{1}{1 - \frac{z_1 - z^*}{c \cdot e - z^*}} \right). \quad (5.3)$$

For $i = 1, \dots, I$, let J_i denote the number of inner iterates computed during outer iteration i , that is, J_i is the largest value j for which $y_{i,j}$ is computed. Clearly, $J_I = \infty$, whereas J_1, \dots, J_{I-1} are finite.

To ease notation, let $\lambda_{i,j} := \lambda_{\min}(y_{i,j})$, and let $\lambda_i^* := \lambda_{\min}(x_{z_i}^*)$, the optimal value of

$$\begin{aligned} \max \quad & \lambda_{\min}(x) \\ \text{s.t.} \quad & x \in \text{Affine}_{z_i} . \end{aligned}$$

According to Lemma 2.4,

$$\lambda_i^* = \frac{z_i - z^*}{c \cdot e - z^*} . \quad (5.4)$$

It is thus valid, for example, to substitute λ_i^* for $\frac{z_i - z^*}{c \cdot e - z^*}$ in (5.2). Additionally, (5.4) implies (5.1) to be equivalent to

$$1 - \lambda_{i+1}^* \geq \frac{4}{3}(1 - \lambda_i^*) . \quad (5.5)$$

For any point y , we have $\pi(y) = e + \frac{1}{1 - \lambda_{\min}(y)}(y - e)$, and thus,

$$\frac{c \cdot e - c \cdot \pi(y)}{c \cdot e - c \cdot y} = \frac{1}{1 - \lambda_{\min}(y)} .$$

Hence,

$$\frac{c \cdot e - c \cdot \pi(y)}{c \cdot e - c \cdot y} \geq \frac{4}{3} \quad \Leftrightarrow \quad \lambda_{\min}(y) \geq 1/4 .$$

Consequently,

$$\lambda_{i,j} < 1/4 \quad \text{for } j < J_i . \quad (5.6)$$

We use the following relation implied by Lemma 3.3:

$$\text{dist}_{z_i}(y_{i,j+1})^2 \leq \text{dist}_{z_i}(y_{i,j})^2 - \frac{\epsilon}{\|\bar{g}_{i,j}\|^2}(\lambda_i^* - \lambda_{i,j}) + \left(\frac{\epsilon}{2\|\bar{g}_{i,j}\|}\right)^2 . \quad (5.7)$$

We begin bounding the number of inner iterations by showing

$$\lambda_i^* \geq \max\{\frac{1}{2}, \epsilon\} \quad \Rightarrow \quad J_i \leq \frac{8(M \text{Dist})^2}{\epsilon} . \quad (5.8)$$

Indeed, for $j < J_i$,

$$\begin{aligned} & -\epsilon(\lambda_i^* - \lambda_{i,j}) + \frac{1}{4}\epsilon^2 \\ & < -\epsilon(\max\{\frac{1}{2}, \epsilon\} - \frac{1}{4}) + \frac{1}{4}\epsilon^2 \quad (\text{using (5.6)}) \\ & = \min\{\frac{1}{4}(\epsilon^2 - \epsilon), \frac{1}{4}\epsilon - \frac{3}{4}\epsilon^2\} \\ & \leq \frac{3}{4}\frac{1}{4}(\epsilon^2 - \epsilon) + \frac{1}{4}(\frac{1}{4}\epsilon - \frac{3}{4}\epsilon^2) \\ & = -\frac{1}{8}\epsilon . \end{aligned}$$

Thus, according to (5.7), for $j < J_i$,

$$\text{dist}_{z_i}(y_{i,j+1})^2 \leq \text{dist}_{z_i}(y_{i,j})^2 - \frac{\epsilon}{8M^2} ,$$

inductively giving

$$\begin{aligned} \text{dist}_{z_i}(y_{i,j+1})^2 & \leq \text{dist}_{z_i}(y_{i,0})^2 - \frac{(j+1)\epsilon}{8M^2} \\ & \leq \text{Dist}^2 - \frac{(j+1)\epsilon}{8M^2} . \end{aligned}$$

The implication (5.8) immediately follows.

The theorem is now readily established in the case $\epsilon \geq 1/2$. Indeed, because of the identity (5.4), the quantity on the right of (5.2) provides an upper bound on the number of outer iterations i for which $\lambda_i^* > \epsilon$, whereas the quantity on the right of (5.8) gives, assuming $\epsilon \geq 1/2$, an upper bound on the number of inner iterations for each of these outer iterations. However, for the first outer iteration satisfying $\lambda_i^* \leq \epsilon$, the initial iterate $y_{i,0}$ ($= \pi(y_{i,0})$) itself achieves the desired accuracy $\frac{c \cdot \pi - z^*}{c \cdot e - z^*} \leq \epsilon$. Thus, the total number of inner iterations made before the algorithm is successful is at most the product of the two quantities which is seen not to exceed the iteration bound in the statement of the theorem (using $\log_{4/3}(1 - \epsilon) \leq \log_{4/3}(1/2) < -1$).

Before considering the remaining case, $\epsilon < 1/2$, we establish a relation applying for all ϵ . For any outer iteration i for which $\lambda_i^* < 3/4$, and for any $0 < \epsilon < 1$, let

$$\widehat{J}_i := \left\lceil \frac{1}{\frac{3}{4} - \lambda_i^*} \left(\frac{M \text{Dist}}{\epsilon} \right)^2 - 1 \right\rceil .$$

We claim that either

$$J_i \leq \widehat{J}_i \quad \text{or} \quad \min \left\{ \frac{c \cdot \pi(y_{i,j}) - z^*}{c \cdot e - z^*} \mid j = 0, \dots, \widehat{J}_i \right\} \leq \epsilon . \quad (5.9)$$

Consequently, if $J_i > \widehat{J}_i$, the algorithm will achieve the goal of computing a point y satisfying $\frac{c \cdot \pi(y) - z^*}{c \cdot e - z^*} \leq \epsilon$ within \widehat{J}_i inner iterations during outer iteration i .

To establish (5.9), assume $\widehat{J}_i < J_i$ and yet the inequality on the right of (5.9) does not hold. (We obtain a contradiction.) For every $j \leq \widehat{J}_i$, Proposition 2.5 then implies

$$\begin{aligned} \lambda_i^* - \lambda_{i,j} &> \epsilon \frac{c \cdot e - z_i}{c \cdot e - z^*} \\ &= (1 - \lambda_i^*) \epsilon \quad (\text{by (5.4)}) , \end{aligned}$$

and hence, using (5.7),

$$\text{dist}_{z_i}(y_{i,j+1})^2 < \text{dist}_{z_i}(y_{i,j})^2 - \left(\frac{3}{4} - \lambda_i^*\right) (\epsilon/M)^2 ,$$

from which inductively follows

$$\begin{aligned} \text{dist}_{z_i}(y_{i,\widehat{J}_i+1})^2 &< \text{dist}_{z_i}(y_{i,0})^2 - (\widehat{J}_i + 1) \left(\frac{3}{4} - \lambda_i^*\right) (\epsilon/M)^2 \\ &\leq \text{Dist}^2 - (\widehat{J}_i + 1) \left(\frac{3}{4} - \lambda_i^*\right) (\epsilon/M)^2 \\ &\leq 0 , \end{aligned}$$

a contradiction. The claim is established.

Assume $\epsilon < 1/2$, the case remaining to be considered.

As each outer iteration i satisfying $\lambda_i^* \geq 1/2$ has only finitely many inner iterations, there must be at least one outer iteration i satisfying $\lambda_i^* < 1/2$. Let i be the first outer iteration for which $\lambda_i^* < 1/2$. From (5.3) and (5.8), the total number of inner iterations made before reaching outer iteration i is at most

$$\frac{8(M \text{Dist})^2}{\epsilon} \log_{4/3} \left(\frac{1}{1 - \frac{z_1 - z^*}{c \cdot e - z^*}} \right) . \quad (5.10)$$

According to (5.9), during outer iteration i , the algorithm either achieves its goal within \widehat{J}_i inner iterations, or the algorithm makes no more than \widehat{J}_i inner iterations before starting a new outer iteration. Assume the latter case. Then, for outer

iteration $i + 1$, the algorithm either achieves its goal within \widehat{J}_{i+1} inner iterations, or the algorithm makes no more than \widehat{J}_{i+1} inner iterations before starting a new outer iteration. Assume the latter case. In iteration $i + 2$, the algorithm definitely achieves its goal within \widehat{J}_{i+2} inner iterations, because there cannot be a subsequent outer iteration due, by (5.5), to

$$\frac{4}{3}(1 - \lambda_{i+2}^*) \geq \left(\frac{4}{3}\right)^3 (1 - \lambda_i^*) > \left(\frac{4}{3}\right)^3 \frac{1}{2} > 1.$$

The total number of inner iterations made before the algorithm achieves its goal is thus bounded by the sum of the quantity (5.10) and

$$\begin{aligned} & \widehat{J}_i + \widehat{J}_{i+1} + \widehat{J}_{i+2} \\ & < \left(\frac{1}{(1 - \frac{1}{2}) - \frac{1}{4}} + \frac{1}{\frac{4}{3}(1 - \frac{1}{2}) - \frac{1}{4}} + \frac{1}{\frac{4}{3} \frac{4}{3}(1 - \frac{1}{2}) - \frac{1}{4}} \right) \left(\frac{M \text{Dist}}{\epsilon} \right)^2 \\ & \quad \text{(using } \frac{3}{4} - \lambda_j^* = (1 - \lambda_j^*) - \frac{1}{4} \text{)} \\ & < 8 \left(\frac{M \text{Dist}}{\epsilon} \right)^2, \end{aligned}$$

completing the proof of the theorem. \square

6. Application to General Convex Optimization

We now return to the setting of Section 1, considering optimization problems of the form

$$\begin{aligned} \min \quad & f(x) \\ \text{s.t.} \quad & x \in \text{Feas}, \end{aligned} \tag{6.1}$$

where $f : \mathcal{E} \rightarrow (-\infty, \infty]$ is an extended-valued and lower-semicontinuous convex function, where $\text{Feas} = \{x \in S \mid Ax = b\}$, with S being a closed convex set, and where there is known a point $\bar{e} \in \text{int}(S \cap \text{dom}(f))$.

In this section we recast (6.1) into conic form, then interpret the complexity results obtained by applying Algorithms 1 and 2 to the conic reformulation. Before proceeding, we recall notation from §1.

Let $\langle \cdot, \cdot \rangle$ be the (computational) inner product on \mathcal{E} , and $\|\cdot\|$ the associated norm. Let P be the linear operator orthogonally projecting \mathcal{E} onto the kernel of A .

Recall \bar{D} denotes the diameter of the sublevel set $\{x \in \text{Feas} \mid f(x) \leq f(\bar{e})\}$. The diameter is assumed to be finite, implying the optimal value f^* is attained at some feasible point.

Recall \hat{f} is a user-chosen scalar satisfying $\hat{f} > f(\bar{e})$ (hence, $(\bar{e}, \hat{f}) \in \text{int}(\text{epi}(f))$), and recall

$$\hat{r} := \sup\{r \mid \|x - \bar{e}\| \leq r \text{ and } Ax = b \Rightarrow x \in S \text{ and } f(x) \leq \hat{f}\}, \tag{6.2}$$

a positive scalar (because $\bar{e} \in \text{int}(S \cap \text{dom}(f))$).

The values \bar{D} and \hat{r} are not assumed to be known, but do appear in complexity results.

For later reference, observe the convexity of f implies

$$f(\bar{e}) \leq \frac{\hat{r}}{\bar{D} + \hat{r}} f^* + \frac{\bar{D}}{\bar{D} + \hat{r}} \hat{f},$$

which in turn implies

$$\frac{1}{1 - \frac{f(\bar{e}) - f^*}{\hat{f} - f^*}} \leq 1 + \bar{D}/\hat{r} . \quad (6.3)$$

As S and $\text{epi}(f)$ are closed and convex, there exist closed, convex cones $\mathcal{K}_1, \mathcal{K}_2 \subseteq \mathcal{E} \times \mathbb{R} \times \mathbb{R}$ for which

$$S \times \mathbb{R} = \{(x, t) \mid (x, 1, t) \in \mathcal{K}_1\} \quad \text{and} \quad \text{epi}(f) = \{(x, t) \mid (x, 1, t) \in \mathcal{K}_2\} .$$

Letting $\mathcal{K} := \mathcal{K}_1 \cap \mathcal{K}_2$, clearly the optimization problem (4.1) is equivalent to

$$\begin{aligned} \min_{x,s,t} \quad & t \\ \text{s.t.} \quad & Ax = b \\ & s = 1 \\ & (x, s, t) \in \mathcal{K} , \end{aligned} \quad (6.4)$$

and has the same optimal value, f^* . The conic program (6.4) is of the same form as CP, the focus of preceding sections. Clearly,

$$\text{Affine} = \{(x, 1, t) \mid Ax = b\} , \quad \text{and for scalars } z, \text{ Affine}_z = \{(x, 1, z) \mid Ax = b\} .$$

For distinguished direction, choose $e = (\bar{e}, 1, \hat{f})$, which clearly lies in the interiors of \mathcal{K}_1 and \mathcal{K}_2 , and thus lies in $\text{int}(\mathcal{K})$. This distinguished direction, along with the cone \mathcal{K} , determines the map $(x, s, t) \mapsto \lambda_{\min}(x, s, t)$ on $\mathcal{E} \times \mathbb{R} \times \mathbb{R}$.

For $z < f^*$, the conic problem 6.4 – and hence the problem 6.1 – is, by Theorem 2.3, equivalent to

$$\begin{aligned} \max_{x,s,t} \quad & \lambda_{\min}(x, s, t) \\ \text{s.t.} \quad & (x, s, t) \in \text{Affine}_z , \end{aligned} \quad (6.5)$$

It is to this problem that supgradient methods are applied.

Choose the computational inner product on $\mathcal{E} \times \mathbb{R} \times \mathbb{R}$ to be

$$\langle (x_1, s_1, t_1), (x_2, s_2, t_2) \rangle := \langle x_1, x_2 \rangle + s_1 s_2 + t_1 t_2 ,$$

in which case

$$P_{\mathcal{L}}(x, s, t) = (P(x), 0, 0) .$$

Observe for all scalars z ,

$$\text{Level}_z = \{(x, 1, z) \mid x \in \text{Feas and } f(x) \leq z\} . \quad (6.6)$$

Recall that the Lipschitz constant M for the map $(x, s, t) \mapsto \lambda_{\min}(x, s, t)$ restricted to Affine_z is independent of z . Since by (6.2),

$$\text{Affine}_{\hat{f}} \cap \bar{B}(e, \hat{r}) \subseteq \mathcal{K} ,$$

Proposition 3.2 implies

$$M \leq 1/\hat{r} . \quad (6.7)$$

Choose the input \bar{x} to Algorithms 1 and 2 as $\bar{x} = (\bar{e}, 1, f(\bar{e}))$, which clearly is feasible for the conic program (6.4). Note that (6.6) then implies the horizontal diameter of the relevant sublevel set for the conic program satisfies

$$\text{Diam}_{f(\bar{e})} = \bar{D} . \quad (6.8)$$

Algorithm 2 requires input $0 < \epsilon < 1$, but not f^* . Applying Algorithm 2 results in a sequence of iterates $(x_k, 1, t_k)$ for which the projections $(x'_k, 1, t'_k) := \pi(x_k, 1, t_k)$

satisfy $x'_k \in \text{Feas}$ and $f(x'_k) \leq t'_k$ (simply because $(x'_k, 1, t'_k)$ is feasible for the conic program (6.4)).

Since $(x_0, 1, t_0) = (\bar{e}, 1, f(\bar{e})) \in \text{bdy}(\mathcal{K})$, we have $\pi(x_0, 1, t_0) = (x_0, 1, t_0)$ – in particular, the objective value of $\pi(x_0, 1, t_0)$ is $f(\bar{e})$. Consequently, the sequence of points x'_k not only lie in Feas, but by Theorem 5.1 satisfies

$$\ell \geq 8 \left(\frac{\bar{D}}{\hat{r}} \right)^2 \left(\frac{1}{\epsilon^2} + \frac{1}{\epsilon} \log_{4/3} \left(1 + \frac{\bar{D}}{\hat{r}} \right) \right) \quad (6.9)$$

$$\Rightarrow \min_{k \leq \ell} \frac{f(x'_k) - f^*}{\hat{f} - f^*} \leq \epsilon, \quad (6.10)$$

where for (6.9) we have used (6.3), (6.7), (6.8), and for (6.10) have used $f(x'_k) \leq t'_k$.

For Algorithm 1, which requires input f^* but not ϵ , points $(x'_k, 1, t'_k)$ are generated for which the sequence $\{x'_k\}$ is feasible and, by Theorem 4.5, satisfies

$$\ell \geq 4 \left(\frac{\bar{D}}{\hat{r}} \right)^2 \cdot \left(\frac{4}{3} \left(\frac{1-\epsilon}{\epsilon} \right)^2 + 4 \left(\frac{1-\epsilon}{\epsilon} \right) \right) \quad (6.11)$$

$$+ \log_2 \left(\frac{1-\epsilon}{\epsilon} \right) + \log_2 \left(\frac{\bar{D}}{\hat{r}} \right) + 1 \quad (6.12)$$

$$\Rightarrow \min_{k \leq \ell} \frac{f(x'_k) - f^*}{\hat{f} - f^*} \leq \epsilon. \quad (6.13)$$

We prove Theorem 1.1 by showing Algorithm A is “equivalent” to applying Algorithm 2 as above. Likewise, we prove Theorem 1.2 by showing algorithm B is equivalent to Algorithm 1. To be in position to establish these equivalences, however, we need a characterization of the supdifferentials $\hat{\partial}\lambda_{\min}(x, s, t)$, a characterization in terms of the original problem (6.1).

7. A Practical Characterization of Supdifferentials, and the Proofs of Theorems 1.1 and 1.2

We continue with the setting and notation developed in Section 6.

We provide a practical characterization of the supgradients for the “ λ_{\min} function,” a characterization expressed in terms of the original problem (6.1). The characterization provides the means for other algorithms designed using the framework to be implemented directly in terms of (6.1), avoiding computation of the conic reformulation.

For the reader’s convenience, we recall key sets appearing in Section 1 that also appear in the characterization below.

For $x' \in \text{bdy}(S)$, let

$$G_1(x') := \left\{ \frac{-1}{\langle v, \bar{e} - x' \rangle} v \mid \vec{0} \neq v \in N_S(x') \right\}.$$

For $(x', t') \in \text{bdy}(\text{epi}(f))$, let

$$G_2(x', t') := \left\{ \frac{-1}{\langle v, \bar{e} - x' \rangle + (\hat{f} - t')\delta} v \mid (\vec{0}, 0) \neq (v, \delta) \in N_{\text{epi}(f)}(x', t') \right\},$$

where here, normality is with respect to the inner product that assigns pairs (x_1, t_1) , (x_2, t_2) the value $\langle x_1, x_2 \rangle + t_1 t_2$.

Finally, letting $\text{bdy}_1 := \{(x', t') \mid x' \in \text{bdy}(S)\}$ and $\text{bdy}_2 := \text{bdy}(\text{epi}(f))$, define for $(x', t') \in \text{bdy}_1 \cup \text{bdy}_2$,

$$G(x', t') := \begin{cases} G_1(x') & \text{if } (x', t') \in \text{bdy}_1 \setminus \text{bdy}_2, \\ G_2(x', t') & \text{if } (x', t') \in \text{bdy}_2 \setminus \text{bdy}_1, \\ \text{hull}(G_1(x') \cup G_2(x', t')) & \text{if } (x', t') \in \text{bdy}_1 \cap \text{bdy}_2. \end{cases}$$

Proposition 7.1. *Assume $(x, 1, t) \in \text{Affine}$, and assume $t < \hat{f}$. If $(x', 1, t') = \pi(x, 1, t)$, then*

$$P_{\mathcal{L}}(\hat{\partial}\lambda_{\min}(x, 1, t)) = \{(P(v), 0, 0) \mid -v \in G(x', t')\}.$$

Proof: Proposition 3.1 implies for all $y \in \mathcal{E} \times \mathbb{R} \times \mathbb{R}$,

$$\begin{aligned} \hat{\partial}\lambda_{\min}(y) &= \{v \in -N_{\mathcal{K}}(y - \lambda_{\min}(y)e) \mid \langle e, v \rangle = 1\} \\ &= -\{v \in N_{\mathcal{K}}(y - \lambda_{\min}(y)e) \mid \langle e, v \rangle = -1\}. \end{aligned} \quad (7.1)$$

However,

$$\text{if } y = (x, 1, t) \in \text{Affine} \text{ and } t < \hat{f}, \quad (7.2)$$

then $\pi(y)$ is a positive multiple of $y - \lambda_{\min}(y)e$ (using $\lambda_{\min}(y) < 1$ (Lemma 2.2), and $\pi(y) = e - \frac{1}{1-\lambda_{\min}(y)}(y - e)$), and thus has the same normal cone to \mathcal{K} . Hence, in this case, (7.1) is equivalent to

$$\hat{\partial}\lambda_{\min}(y) = -\{v \in N_{\mathcal{K}}(\pi(y)) \mid \langle e, v \rangle = -1\}. \quad (7.3)$$

As $\text{int}(\mathcal{K}) \neq \emptyset$ and $\mathcal{K} = \mathcal{K}_1 \cap \mathcal{K}_2$, if $y \in \mathcal{K}$ then

$$N_{\mathcal{K}}(y) = \text{hull}(N_{\mathcal{K}_1}(y) \cup N_{\mathcal{K}_2}(y)).$$

Thus, for y as in (7.2), using (7.3) we have

$$\hat{\partial}\lambda_{\min}(y) = -\{v \in \text{hull}(N_{\mathcal{K}_1}(\pi(y)) \cup N_{\mathcal{K}_2}(\pi(y))) \mid \langle e, v \rangle = -1\}. \quad (7.4)$$

Since $e \in \text{int}(\mathcal{K}_i)$ ($i = 1, 2$),

$$y \in \mathcal{K}_i \wedge v \in N_{\mathcal{K}_i}(y) \wedge \langle e, v \rangle = 0 \quad \Rightarrow \quad v = 0.$$

Consequently, it follows from (7.4) that for y as in (7.2),

$$\hat{\partial}\lambda_{\min}(y) = -\text{hull}(\mathcal{N}_1(\pi(y)) \cup \mathcal{N}_2(\pi(y))),$$

where for $i = 1, 2$ and $y' \in \mathcal{K}_i$,

$$\mathcal{N}_i(y') := \begin{cases} \{v \in N_{\mathcal{K}_i}(y') \mid \langle e, v \rangle = -1\} & \text{if } y' \in \text{bdy}(\mathcal{K}_i), \\ \{\vec{0}\} & \text{if } y' \in \text{int}(\mathcal{K}_i). \end{cases}$$

Clearly, then,

$$P_{\mathcal{L}}(\hat{\partial}\lambda_{\min}(y)) = -\text{hull}(P_{\mathcal{L}}(\mathcal{N}_1(\pi(y))) \cup P_{\mathcal{L}}(\mathcal{N}_2(\pi(y)))) . \quad (7.5)$$

Since $S \times \mathbb{R} = \{(x, t) \mid (x, 1, t) \in \mathcal{K}_1\}$, if $y' = (x', 1, t') \in \text{bdy}(\mathcal{K}_1)$ then

$$N_{\mathcal{K}_1}(y') = \{(v, \gamma, 0) \mid v \in N_S(x') \wedge \langle x', v \rangle = 0\},$$

the equation to ensure perpendicularity to the ray through y' . Thus, a vector (v, γ, δ) is an element of $\mathcal{N}_1(y')$ if and only if $v \in N_S(x')$, $\delta = 0$ and

$$\begin{aligned}\langle x', v \rangle + \gamma &= 0 \\ \langle \bar{e}, v \rangle + \gamma &= -1\end{aligned}$$

(note v cannot be $\vec{0}$). However, for $\vec{0} \neq v \in N_S(y')$, there is a unique scaling of v for which there exist γ and δ satisfying the three equations, namely,

$$v \mapsto \frac{-1}{\langle v, \bar{e} - x' \rangle} (v, -\langle x', v \rangle, 0).$$

The denominator in the scaling is negative, because $\bar{e} \in \text{int}(S)$. It follows that

$$P_{\mathcal{L}}(\mathcal{N}_1) = \left\{ \frac{-1}{\langle v, \bar{e} - x' \rangle} (P(v), 0, 0) \mid \vec{0} \neq v \in N_S(x') \right\}. \quad (7.6)$$

Similarly, for $y' = (x', 1, t') \in \text{bdy}(\mathcal{K}_2)$, a vector (v, γ, δ) is in $\mathcal{N}_2(y')$ if and only $(v, \delta) \in N_{\text{epi}(f)}(x', t')$ and

$$\begin{aligned}\langle x', v \rangle + \gamma + \delta t' &= 0 \\ \langle \bar{e}, v \rangle + \gamma + \delta \hat{f} &= -1\end{aligned}$$

(note (v, δ) cannot be $(\vec{0}, 0)$). For $(\vec{0}, 0) \neq (v, \delta) \in N_{\text{epi}(f)}(x', t')$, the unique scaling for which there exists γ satisfying the equations is

$$(v, \delta) \mapsto \frac{-1}{\langle v, \bar{e} - x' \rangle + (\hat{f} - t')\delta} (v, -\langle x', v \rangle - t'\delta, \delta).$$

The denominator is negative, because $(\bar{e}, \hat{f}) \in \text{int}(\text{epi}(f))$. It follows that

$$P_{\mathcal{L}}(\mathcal{N}_2) = \left\{ \frac{-1}{\langle v, \bar{e} - x' \rangle + (\hat{f} - t')\delta} (P(v), 0, 0) \mid (\vec{0}, 0) \neq (v, \delta) \in N_{\text{epi}(f)}(x', t') \right\}. \quad (7.7)$$

Together, (7.5), (7.6) and (7.7) establish the proposition. \square

Proofs of Theorems 1.1 and 1.2, and of (1.7): We first establish Theorem 1.1, and then remark on the few changes required to establish Theorem 1.2 by the same approach.

By ‘‘Algorithm 2,’’ we mean the application of Algorithm 2 to the conic optimization problem (6.4), as presented in the preceding section.

Due to complexity bound (6.9), (6.10) established for Algorithm 2, it suffices to show the two algorithms are equivalent in that the iterates generated by Algorithm 2 are ‘‘identical’’ to those generated by Algorithm A.

The initial iterates for Algorithm 2 are $(x_0, 1, t_0) = (x'_0, 1, t'_0) = (\bar{e}, 1, f(\bar{e}))$, and for Algorithm A are $(x_0, t_0) = (x'_0, t'_0) = (\bar{e}, f(\bar{e}))$. For beginning an inductive proof showing the two algorithms are equivalent, observe that the initial iterate (x_0, t_0) (resp., (x'_0, t'_0)) for Algorithm A is obtained simply by eliminating the ‘‘1’’ from the initial iterate $(x_0, 1, t_0)$ (resp., $(x'_0, 1, t'_0)$) for Algorithm 2.

For the inductive step, assume eliminating ‘‘1’’ from the iterate $(x_k, 1, t_k)$ (resp., $(x'_k, 1, t'_k)$) computed by Algorithm 2 results in the iterate (x_k, t_k) (resp., (x'_k, t'_k)) computed by Algorithm A.

The first iterate computed by Algorithm 2 in Step 1 is

$$(\tilde{x}_{k+1}, 1, t_k) = (x_k, 1, t_k) + \frac{\epsilon}{2\|P_{\mathcal{L}}\bar{g}\|^2} P_{\mathcal{L}}\bar{g} \quad \text{where } \bar{g} \in \hat{\partial}\lambda_{\min}(x'_k, 1, t'_k),$$

whereas Algorithm A computes

$$\tilde{x}_{k+1} = x_k - \frac{\epsilon}{2\|Pg\|^2} Pg \quad \text{where } g \in G(x'_k, t'_k)$$

(and thereafter relies on the pair (\tilde{x}_{k+1}, t_k) in the same manner that Algorithm 2 relies on the triple $(\tilde{x}_{k+1}, 1, t_k)$). Thus, due to the identity provided by Proposition 7.1, the algorithms can be considered equivalent in this computation.

Algorithm 2 next computes $(x'_{k+1}, 1, t'_{k+1}) := \pi(\tilde{x}_{k+1}, 1, t_k)$, the point in the boundary of \mathcal{K} encountered when moving from $(\bar{e}, 1, \hat{f})$ in direction $(\tilde{x}_{k+1}, 1, t_k) - (\bar{e}, 1, \hat{f})$. Due to the definition of \mathcal{K} , however, this is the (only) feasible point $(x, 1, t)$ for which there exists $\alpha \geq 0$ such that

$$x = x(\alpha) := \bar{e} + \alpha \cdot (\tilde{x}_{k+1} - \bar{e}) \quad \text{and} \quad t = t(\alpha) := \hat{f} + \alpha \cdot (t_k - \hat{f})$$

and either $x \in \text{bdy}(S)$ or $(x, t) \in \text{bdy}(\text{epi}(f))$. Thus, for the iterate $(x'_{k+1}, 1, t'_{k+1})$ computed by Algorithm 2, we have $(x'_{k+1}, t'_{k+1}) = \pi'(\tilde{x}_{k+1}, t_k)$, where π' is defined by (1.3). Thus, eliminating “1” from the iterate $(x'_{k+1}, 1, t'_{k+1})$ computed by Algorithm 2 gives the iterate (x'_{k+1}, t'_{k+1}) computed by Algorithm A.

Observe, moreover, for Algorithm A, $t'_{k+1} = \hat{f} + \alpha_{k+1} \cdot (t_k - \hat{f})$, and thus,

$$\alpha_{k+1} = \frac{\hat{f} - t'_{k+1}}{\hat{f} - t_k}. \tag{7.8}$$

Algorithm A decides how to define its iterate (x_{k+1}, t_{k+1}) based on whether $\alpha_{k+1} \geq 4/3$, whereas Algorithm 2 decides how to define its iterate $(x_{k+1}, 1, t_{k+1})$ based on whether $\hat{f} - t'_{k+1} \geq \frac{4}{3}(\hat{f} - t_k)$. Thus, due to (7.8) and the equivalence of the two algorithms up until now, it is simple to verify that removing “1” from the iterate $(x_{k+1}, 1, t_{k+1})$ for Algorithm 2 gives the iterate (x_{k+1}, t_{k+1}) for Algorithm A.

Similarly, the proof of Theorem 1.2 is accomplished by showing Algorithm B is “equivalent” to Algorithm 1, and making use of the complexity result (6.11), (6.12), (6.13) for Algorithm 1. The proof of equivalence is virtually identical to the one above, the main difference being that the scalar α_k used by Algorithm B has to be related to the analogous scalar used by Algorithm 1 – in particular, for equivalence, the identity $\frac{\alpha(x_k, f^*) - 1}{\alpha(x_k, f^*)} = \lambda_{\min}(x_k, 1, f^*)$ is needed.

The identity, however, is easily established. Indeed, $\alpha(x_k, f^*)$ is the smallest scalar α for which the point $(x, t) = (\bar{e}, \hat{f}) + \alpha \cdot ((x_k, f^*) - (\bar{e}, \hat{f}))$ satisfies either $x \in \text{bdy}(S)$ or $(x, t) \in \text{bdy}(\text{epi}(f))$. Due to the definition of \mathcal{K} , however, this is readily seen to equal the smallest scalar α for which $e + \alpha \cdot ((x_k, 1, f^*) - e)$ lies in the boundary of \mathcal{K} . Hence, from (2.5), $\alpha(x_k, f^*) = \frac{1}{1 - \lambda_{\min}(x_k, 1, f^*)}$. Rearrangement establishes the desired identity.

Finally, immediately after the statement of Theorem 1.2, it is claimed in (1.7) that Algorithm B converges linearly when f is a piecewise linear function and Feas is polyhedral. This is immediate from the fact that Algorithm 1 converges linearly when \mathcal{K} is a polyhedral cone (Corollary 4.2). \square

Acknowledgements: The author expresses gratitude to the reviewers and the associate editor, for careful attention to detail and for suggestions on how to position the paper. The author thanks Yurii Nesterov for encouragement at an early, critical stage of the research, Rob Freund for discussions and feedback influencing the overall manner of presentation, and Wei Qian for conversations affecting the deduction and presentation of results for general convex optimization.

REFERENCES

- [1] Robert M Freund and Haihao Lu. New computational guarantees for solving convex optimization problems with first order methods, via a function growth condition measure. *arXiv preprint arXiv:1511.02974*, 2015.
- [2] A. Gilpin, J. Peña, and T. Sandholm. First-order algorithm with $O(\ln(1/\epsilon))$ convergence for ϵ -equilibrium in two-person zero-sum games. *Mathematical Programming*, 133(1-2):279–298, 2012.
- [3] Yu. Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, 2005.
- [4] James Renegar. Accelerated first-order methods for hyperbolic programming. *arXiv preprint arXiv:1512.07569*, 2015.
- [5] James Renegar. A framework for applying subgradient methods to conic optimization problems. *arXiv preprint arXiv:1503.02611*, 2015.
- [6] R Tyrrell Rockafellar. Convex analysis. *Princeton University Press*, 46:49, 1970.

SCHOOL OF OPERATIONS RESEARCH AND INFORMATION ENGINEERING, CORNELL UNIVERSITY,
ITHACA, NY, U.S.