

ACCELERATED FIRST-ORDER METHODS FOR HYPERBOLIC PROGRAMMING

JAMES RENEGAR

ABSTRACT. A framework is developed for applying accelerated methods to general hyperbolic programming, including linear, second-order cone, and semidefinite programming as special cases. The approach replaces a hyperbolic program with a convex optimization problem whose smooth objective function is explicit, and for which the only constraints are linear equations (one more linear equation than for the original problem). Virtually any first-order method can be applied. Iteration bounds for a representative accelerated method are derived.

1. Introduction

In [26], we showed that given a general conic optimization problem for which a strictly feasible point is known, there is a simple transformation to an equivalent convex optimization problem which has the same number of variables, has only linear equations as constraints (one more linear equation than the original problem), and has Lipschitz-continuous objective function defined on the whole space. Virtually any subgradient method can be applied to solve the equivalent problem, the cost per iteration dominated by computation of a subgradient and its orthogonal projection onto a subspace (the same subspace at every iteration, a situation for which preprocessing is effective). Iteration bounds for a representative subgradient method were established.

The iteration bounds, however, are of order $1/\epsilon^2$, as is to be expected for subgradient methods when the goal is to compute a point whose objective value is within ϵ of optimal ([17]). By contrast, a decade ago, Nesterov [20] obtained $O(1/\epsilon)$ iteration bounds for an accelerated first-order method applied to various minimization problems in which the objective is the sum of a smooth convex function and a nonsmooth simple convex function, where by “simple” is meant that the function is well-approximated by an implicitly-defined smooth function (a Moreau proximal smoothing) for which the gradient at a point can be computed by solving a simple optimization problem – an optimization problem so simple that the gradient can be expressed in closed form, that is, made explicit.

This work of Nesterov motivated many researchers, leading to a great number of related publications. An elegant formalization and insightful overview was provided by Beck and Teboulle [2], indicating the generality of the approach, at least in theory. The limitation of the approach is that for the implicitly-defined smooth function, the subproblem for computing the gradient at a point indeed has to be simple, else the cost of (approximately) computing the gradient can greatly exceed all other costs in an iteration. This requirement of simplicity significantly restricts the scope of optimization problems to which the approach can be applied in practice, although the approach has been very consequential for some particularly important classes of high-profile optimization problems, such as in compressive sensing (c.f., [3, 4, 6]).

Here we extend our investigation in [26], now applying accelerated methods rather than subgradient methods. However, whereas [26] considered *all* convex conic optimization problems, now attention is restricted to hyperbolic programming, still an extremely-general class of conic optimization problems, albeit a class for which the algebraic structure of the underlying cones naturally provides an explicit “smoothing” of the objective function for the equivalent optimization problem. Once the objective function is smoothed, as will become clear, virtually any first-order method can be applied.

Research supported in part by NSF CCF Award #1552518.

In §2.1, we provide a brief overview of the framework and “basic theory” introduced in [26, §2], all of which is elementary. In §2.2, we give a concise introduction to hyperbolic programming, a task easily accomplished by relying on the reader’s grasp of the basics for semidefinite programming.

Results regarding “smoothing” are presented in §3 (lengthy proofs are deferred to appendices). Besides presenting desirable theoretical properties of the explicit smoothing, we provide computational expressions for gradients, making evident the practicality of the smoothing especially when a hyperbolic program has “hyperbolic polynomial” which factors as a product of polynomials of low degree. The results are demonstrated by showing for a general second-order cone program, computing gradients for the smoothed function is easy.

In §4 we consider a representative accelerated method, due to Nesterov [22], establishing an $O(1/\epsilon)$ -iteration bound when the method is applied to the smoothing of a general hyperbolic program. Additionally, we give considerable attention to the constant hidden in the big- O , showing that if the accelerated method is instead applied to a carefully-chosen *sequence* of smoothed versions of the hyperbolic program, significantly greater efficiency can be obtained, that is, far fewer gradient computations can suffice for solving the hyperbolic program to within desired accuracy.

These results significantly extend and strengthen ones announced in [25] for the special case of semidefinite programming. The present paper and [26] should together be considered as the final, refined version of [25].

An especially-notable earlier work with a similar $O(1/\epsilon)$ -iteration bound for the special case of semidefinite programming is due to Lu, Nemirovski and Monteiro [16] (rather than a framework in which various accelerated methods can be applied, their focus is on a particular accelerated method aimed at solving large-scale problems). Also deserving of special mention is work of Lan, Lu and Monteiro [14], who obtain $O(1/\epsilon)$ -iteration bounds for primal-dual convex conic optimization problems (although for the cost of individual iterations is to be low, projection onto the underlying cones has to be inexpensive).

Closely related to the present paper and [26] is recent significant work of Freund and Lu [7], who develop first-order methods for convex optimization problems in which the objective function has a known lower bound. They consider general non-smooth, smooth and smoothed settings (although for the latter, they assume the smoothing is given, whereas the relevance of the present paper rests largely in contributing a particular smoothing). The perspective they develop provides an important juxtaposition to our conic-oriented presentation.

2. Preliminaries

2.1. Basic Theory. Here we provide a brief overview of the general framework introduced in [26, §2] and record relevant results proven there. All of this is elementary.

Let \mathcal{E} be a finite-dimensional real Euclidean space. Let $\mathcal{K} \subset \mathcal{E}$ be a proper, closed, convex cone with non-empty interior.

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

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

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 convexity of \mathcal{K} .)

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

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

a closed, centrally-symmetric, convex set with nonempty interior. Define a seminorm¹ on \mathcal{E} according to

$$\|u\|_{\infty} := \min\{t \geq 0 : u = tv \text{ for some } v \in \bar{\mathcal{B}}\}. \quad (2.2)$$

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.

It is straightforward to show $\|\cdot\|_{\infty}$ is a norm if and only if \mathcal{K} is pointed (i.e., contains no subspace other than $\{\vec{0}\}$).

Proposition 2.1. ([26, Prop 2.1]) *The function $x \mapsto \lambda_{\min}(x)$ is concave and Lipschitz continuous:*

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

Assume the Euclidean space \mathcal{E} is endowed with inner product written $u \cdot v$. 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

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

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} .

Assume $\text{Affine} \cap \text{int}(\mathcal{K})$ – the set of strictly feasible points – is nonempty. Fix a strictly feasible point, e . The point e serves as the distinguished direction.

For scalars $z \in \mathbb{R}$, we introduce the affine space

$$\text{Affine}_z := \{x \in \text{Affine} : c \cdot x = z\}.$$

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 need 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: If $\lambda_{\min}(x) \geq 1$, then by (2.1), $\lambda_{\min}(e + t(x - e)) \geq 1$ for all $t \geq 0$, and hence $x + t(x - e)$ is feasible for all $t \geq 0$. The lemma easily follows. \square

¹Recall that a seminorm $\|\cdot\|$ satisfies $\|tv\| = |t|\|v\|$ and $\|u + v\| \leq \|u\| + \|v\|$, but unlike a norm, is allowed to satisfy $\|v\| = 0$ for $v \neq 0$.

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} :

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

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

The following theorem has a simple proof, and yet the theorem is the foundation of both [26] and the present paper.

Theorem 2.3. ([26, Thm 2.3]) *Let z be any value satisfying $z < c \cdot e$. If x^* solves*

$$\begin{aligned} \sup \quad & \lambda_{\min}(x) \\ \text{s.t.} \quad & x \in \text{Affine}_z, \end{aligned} \tag{2.3}$$

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 (2.3), and $\pi^* = \pi(x^*)$.*

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} of which the affine space Affine_z is a translate. This approach was investigated in [26]. In the present paper, where the focus is on accelerated gradient methods, the objective function of the equivalent problem (2.3) is smoothed. Before discussing smoothing, however, we complete preliminaries.

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 (2.3) has at least one optimal solution. Let $\lambda_{\min,z}^$ denote the optimal value for the equivalent problem, and recall z^* denotes the optimal value of CP. A useful characterization of $\lambda_{\min,z}^*$ is easily provided.*

Lemma 2.4. ([26, Lemma 2.4])

$$\lambda_{\min,z}^* = \frac{z - z^*}{c \cdot e - z^*}.$$

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

$$\frac{c \cdot \pi - z^*}{c \cdot e - z^*} \leq \epsilon, \tag{2.4}$$

where $0 < \epsilon < 1$ is user-chosen. 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 (2.3) so as to ensure that for the computed point x , the projection $\pi = \pi(x)$ satisfies (2.4).

Proposition 2.5. ([26, Prop 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

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

In [26], we concluded developing the basic theory by observing the theory holds for convex conic optimization problems generally, such as problems of the form

$$\begin{aligned} \min_{x \in \mathcal{E}} \quad & c \cdot x \\ \text{s.t.} \quad & x \in \text{Affine} \\ & Ax + b \in \mathcal{K}' \end{aligned}$$

(here, A is a linear operator from \mathcal{E} to a Euclidean space \mathcal{E}' , $b \in \mathcal{E}'$ and \mathcal{K}' is a proper, closed, convex cone in \mathcal{E}' with nonempty interior). In the present paper as in [26], we focus on problems of the form CP because notationally the form is least cumbersome. However, for every result derived in the following sections, an essentially identical result holds for any hyperbolic program regardless of the form, even identical in the specific constants.

2.2. Hyperbolic Programming. Here we give a concise introduction to hyperbolic programming, easily accomplished by relying on the reader’s grasp of the basics for semidefinite programming.

A hyperbolic program is a convex conic optimization problem for which the underlying cone is a so-called “hyperbolicity cone.” The quintessential hyperbolicity cone is \mathbb{S}_+^n . The key role in qualifying \mathbb{S}_+^n as a hyperbolicity cone is played by the function $X \mapsto \det(X)$, which if the coefficients of X are viewed as variables, is a real polynomial of degree n on the Euclidean space \mathbb{S}^n . The critical property possessed by this polynomial is that there exists $E \in \text{int}(\mathbb{S}_+^n)$ – say, $E = I$ – such that for every $X \in \mathbb{S}^n$, the univariate polynomial $X \mapsto \det(X - \lambda I)$ has only real roots, the eigenvalues of X .

More generally, a closed convex cone $\mathcal{K} \subset \mathcal{E}$ with nonempty interior is said to be a *hyperbolicity cone*² if there exists a real homogeneous polynomial $p : \mathcal{E} \rightarrow \mathbb{R}$ for which

- (1) $\text{int}(\mathcal{K})$ is a connected component of the set $\{x : p(x) \neq 0\}$ (i.e., p is identically zero on the boundary of \mathcal{K} and nowhere zero in the interior of \mathcal{K}), and
- (2) for some $e \in \text{int}(\mathcal{K})$ and all $x \in \mathcal{E}$ – equivalently, by [8], for every $e \in \text{int}(\mathcal{K})$ and all $x \in \mathcal{E}$ – the univariate polynomial $\lambda \mapsto p(x - \lambda e)$ has only real roots.

We say “ \mathcal{K} has hyperbolic polynomial p ” (but it should be noted p is not unique – for example, p^2 also is a hyperbolic polynomial for \mathcal{K}).

It is easy to see that the non-negative orthant, \mathbb{R}_+^n , is a hyperbolicity cone with polynomial $p(x) = x_1 \cdots x_n$. Similarly, it is easy to see the second-order cone $\{(x, r) \in \mathbb{R}^n \times \mathbb{R} : \|x\|_2 \leq r\}$ has hyperbolic polynomial $(x, r) \mapsto r^2 - \|x\|_2^2$.

In requirement (2), the characterization “for some” is useful for verifying hyperbolicity of a given cone, whereas the characterization “for every” is critical to easily establishing desirable general properties of hyperbolicity cones. Gårding ([8]) proved the equivalence in the late 1950’s, initiating the study of hyperbolicity cones, but in the context of hyperbolic partial differential equations, not optimization.

Using the characterization “for every,” for example, it is simple to verify that if \mathcal{K}_1 and \mathcal{K}_2 are hyperbolicity cones whose interiors contain a common point, then $\mathcal{K}_1 \cap \mathcal{K}_2$ is a hyperbolicity cone, for which a hyperbolic polynomial is the product $p = p_1 p_2$. Similarly, it is easy to verify that if $T : \mathcal{E} \rightarrow \mathcal{E}$ is a linear transformation whose image intersects the interior of a hyperbolicity cone $\bar{\mathcal{K}}$, then the pullback $\mathcal{K} := \{x \in \mathcal{E} : T(x) \in \bar{\mathcal{K}}\}$ is a hyperbolicity cone, with hyperbolic polynomial $p(x) = \bar{p}(Tx)$.

²In the literature, the term “hyperbolicity cone” generally refers to the interior of \mathcal{K} , but for us it is expeditious to use the term for the closed cone \mathcal{K} .

Other examples of hyperbolicity cones include the so-called “derivative cones,” which form a nested sequence of relaxations of an initial hyperbolicity cone ([24]).

A *hyperbolic program* is a convex conic optimization problem for which the cone is a hyperbolicity cone. Since we are focusing – for notational expediency – on conic optimization problems having the form of CP, for us a hyperbolic program is an optimization problem

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

where \mathcal{K} is a hyperbolicity cone.

Due to the calculus of hyperbolicity cones (intersections, Cartesian products, pullbacks, etc.), a broad range of optimization problems can be cast as hyperbolic programs.

It was Güler [9] who initiated the study of hyperbolic programming, upon recognizing that the function $x \mapsto -\ln p(x)$ is a self-concordant barrier for \mathcal{K} (where p is a hyperbolic polynomial for \mathcal{K}), and consequently, the general interior-point method theory of Nesterov and Nemirovski [23] applies. We recommend [1] as a particularly elegant work examining hyperbolicity cones from the viewpoint of optimization.

Letting \mathcal{K} be a hyperbolicity cone with polynomial $p : \mathcal{E} \rightarrow \mathbb{R}$, and fixing $e \in \text{int}(\mathcal{K})$, for $x \in \mathcal{E}$ we refer to the roots of $\lambda \mapsto p(x - \lambda e)$ as the “eigenvalues of x ,” and denote them $\lambda_j(x)$ for $j = 1, \dots, n$, where $n := \deg(p)$. Thus, we list each eigenvalue according to its multiplicity (but we do not assume the eigenvalues are ordered).

The terminology “eigenvalues” is a mnemonic device for emphasizing the parallels of the general setting with the special case in which $\mathcal{K} = \mathbb{S}_+^n$ and $p(X) = \det(X)$.

The eigenvalues of x depend on the choice of e , but we do not introduce notation making explicit the dependence, because for us e will be fixed – in particular, e will be the “distinguished direction” of the basic theory (§2.1). Observe, then,

$$\min_j \lambda_j(x) = \lambda_{\min}(x) ,$$

where $\lambda_{\min}(x)$ is as in the basic theory. Also observe, by homogeneity of p ,

$$\lambda_{\max}(x) := \max_j \lambda_j(x) = -\lambda_{\min}(-x) . \quad (2.5)$$

Recall the seminorm $\|\cdot\|_\infty$, defined by (2.2) so as to have the property that $\bar{B}_\infty(e, 1)$ – the closed unit ball centered at e – is the largest set which both is contained in \mathcal{K} and has symmetry point e . Since for a vector v ,

$$e + v \in \mathcal{K} \text{ and } e - v \in \mathcal{K} \quad \Leftrightarrow \quad 1 + \lambda_{\min}(v) \geq 0 \text{ and } 1 + \lambda_{\min}(-v) \geq 0 ,$$

and since $\lambda_{\min}(-v) = -\lambda_{\max}(v)$, it follows for all $u \in \mathcal{E}$,

$$\|u\|_\infty = \max_j |\lambda_j(u)| . \quad (2.6)$$

3. Smoothing

The basic theory (§2.1) reveals that a convex conic optimization problem is equivalent to a convex optimization problem whose only constraints are linear equations:

$$\begin{array}{l} \min \quad c \cdot x \\ \text{s.t.} \quad x \in \text{Affine} \\ \quad \quad x \in \mathcal{K} \end{array} \quad \equiv \quad \begin{array}{l} \max \quad \lambda_{\min}(x) \\ \text{s.t.} \quad x \in \text{Affine}_z , \end{array} \quad (3.1)$$

where $\text{Affine}_z := \{x \in \text{Affine} : c \cdot x = z\}$ and z is assumed to satisfy $c \cdot z < c \cdot e$, with e being the distinguished direction.

In this section it is shown that if \mathcal{K} is a hyperbolicity cone, then the function $x \mapsto \lambda_{\min}(x)$ can be approximated by an explicit, concave, smooth function. If, for the problem on the right in (3.1), the objective function is replaced by the smooth function, we have an optimization problem to which accelerated first-order methods can be applied, so as to approximately solve the problem on the left, the conic problem of interest.

3.1. The smooth function and its gradients. To conform to the emphasis on convex (rather than concave) functions in the literature, in this subsection we focus on smoothing the convex function $x \mapsto \lambda_{\max}(x)$. Due to the relation (2.5), the results easily translate to smoothing the concave function $x \mapsto \lambda_{\min}(x)$.

Assume \mathcal{K} is a hyperbolicity cone, with hyperbolic polynomial p of degree n , and let $e \in \text{int}(\mathcal{K})$ be the distinguished direction, giving rise to eigenvalues $\lambda_j(x)$ ($j = 1, \dots, n$) for every $x \in \mathcal{E}$.

Our approach is entirely motivated by Nesterov [21], who focused on the case $\mathcal{K} = \mathbb{S}_+^n$.

For $\mu > 0$, define the function $f_\mu : \mathcal{E} \rightarrow \mathbb{R}$ by

$$f_\mu(x) := \mu \ln \sum_j \exp(\lambda_j(x)/\mu) .$$

It is easily verified for all $x \in \mathcal{E}$ that

$$\lambda_{\max}(x) \leq f_\mu(x) \leq \lambda_{\max}(x) + \mu \ln n , \quad (3.2)$$

where n is the degree of the hyperbolic polynomial p (and hence is the number of eigenvalues, counting multiplicities).

Dual to the seminorm $\| \cdot \|_\infty$ is the extended-real-valued convex function

$$\|g\|_\infty^* := \sup_{\|v\|_\infty \leq 1} \langle g, v \rangle$$

(this is a norm when $\| \cdot \|_\infty$ is a norm, i.e., when \mathcal{K} is a regular cone).

Proposition 3.1. *The function f_μ is convex and infinitely Fréchet-differentiable. Moreover,*

$$\|\nabla f_\mu(x) - \nabla f_\mu(y)\|_\infty^* \leq \frac{1}{\mu} \|x - y\|_\infty \quad \text{for all } x, y \in \mathcal{E} .$$

The proposition is mostly a corollary of Nesterov [21] and a deep theorem regarding hyperbolicity cones, the Helton-Vinnikov Theorem. The proof is deferred to Appendix A.

For f_μ to be relevant to practice, its gradients must be readily computable. We now discuss computation of the gradients.

A key observation is that for many problems, the underlying hyperbolicity cone \mathcal{K} is either a Cartesian product of hyperbolicity cones \mathcal{K}_i , or is the intersection of hyperbolicity cones \mathcal{K}_i . In such cases, computing gradients mostly reduces to consideration of the individual cones \mathcal{K}_i .

Proposition 3.2. *Assume $\mathcal{K}_1, \dots, \mathcal{K}_\ell$ are hyperbolicity cones in a Euclidean space \mathcal{E} , with hyperbolic polynomials p_1, \dots, p_ℓ . Assume $e \in \text{int}(\mathcal{K}_i)$ for all i , and thus e can serve as distinguished direction for every i , thereby defining for each $x \in \mathcal{E}$ a set of eigenvalues $\{\lambda_{ij}(x) : j = 1, \dots, n_i\}$ (the roots of $\lambda \mapsto p_i(x - \lambda e)$), where n_i is the degree of p_i). Let*

$$f_{i,\mu}(x) := \mu \ln \sum_{j=1}^{n_i} \exp(\lambda_{ij}(x)/\mu) .$$

Let $\mathcal{K} := \cap_i \mathcal{K}_i$, for which a hyperbolic polynomial is the product $p = p_1 \cdots p_\ell$. For every $x \in \mathcal{E}$, the corresponding set of eigenvalues – the roots of $\lambda \mapsto p(x - \lambda e)$ – is $\cup_{i=1}^\ell \{\lambda_{ij}(x) : j = 1, \dots, n_i\}$, and the corresponding smooth function f_μ satisfies

$$\nabla f_\mu(x) = \frac{1}{\sum_{i=1}^\ell (\sum_{j=1}^{n_i} \exp(\lambda_{ij}(x)/\mu))} \sum_{i=1}^\ell \left(\sum_{j=1}^{n_i} \exp(\lambda_{ij}(x)/\mu) \right) \nabla f_{i,\mu}(x),$$

a convex combination of the gradients $\nabla f_{i,\mu}(x)$.

Proof: The only claim that perhaps is not obvious is the identity for $\nabla f_\mu(x)$. However, the reasoning for the identity becomes simple upon observing the functions $f_{i,\mu}$ are smooth (Proposition 3.1), and hence so are the functions $x \mapsto \sum_{j=1}^{n_i} \exp(\lambda_{ij}(x)/\mu)$. Consequently,

$$\begin{aligned} \nabla f_\mu(x) &= \frac{\mu}{\sum_{i=1}^\ell \sum_{j=1}^{n_i} \exp(\lambda_{ij}(x)/\mu)} \sum_{i=1}^\ell \nabla \left(\sum_{j=1}^{n_i} \exp(\lambda_{ij}(x)/\mu) \right) \\ &= \frac{1}{\sum_{i=1}^\ell \sum_{j=1}^{n_i} \exp(\lambda_{ij}(x)/\mu)} \sum_{i=1}^\ell \left(\sum_{j=1}^{n_i} \exp(\lambda_{ij}(x)/\mu) \right) \nabla f_{i,\mu}(x), \end{aligned}$$

establishing the identity. \square

Similarly, if \mathcal{K} is the Cartesian-product of cones $\mathcal{K}_i \subseteq \mathcal{E}_i$ with hyperbolic polynomials p_i ($i = 1, \dots, \ell$), then $p(x_1, \dots, x_m) := p_1(x_1) \cdots p_\ell(x_\ell)$ is a hyperbolic polynomial for \mathcal{K} . Moreover, if $e = (e_1, \dots, e_\ell) \in \text{int}(\mathcal{K})$ is the distinguished direction defining eigenvalues for points $x = (x_1, \dots, x_\ell)$ in the Euclidean space $\mathcal{E} = \mathcal{E}_1 \times \cdots \times \mathcal{E}_\ell$, then the smooth function f_μ satisfies

$$\nabla f_\mu(x) = (g_1, \dots, g_\ell) \quad \text{where } g_i = \frac{\sum_{j=1}^{n_i} \exp(\lambda_{ij}(x_i)/\mu)}{\sum_{k=1}^\ell \sum_{j=1}^{n_k} \exp(\lambda_{kj}(x_k)/\mu)} \nabla f_{i,\mu}(x_i),$$

and where $\{\lambda_{ij}(x_i) : j = 1, \dots, n_i\}$ is the set of roots of $\lambda \mapsto p_i(x_i - \lambda e_i)$.

Now we turn to establishing a formula for computing $\nabla f_\mu(x)$ once the eigenvalues $\lambda_j(x)$ are known (to high precision). In the context of Proposition 3.2, the identity is pertinent to computing the gradients $\nabla f_{i,\mu}(x)$.

The identity is easily motivated in the case that each eigenvalue $\lambda_j(x)$ is simple (i.e., is a root of multiplicity one for the polynomial $\lambda \mapsto p(x - \lambda e)$).

If $\lambda_j(x)$ is a simple eigenvalue, then $y \mapsto \lambda_j(y)$ is analytic in a neighborhood of x (because simple roots of a univariate polynomial are analytic in small perturbations of the coefficients, and because the coefficients of $\lambda \mapsto p(x - \lambda e)$ are analytic in x). In particular, when each eigenvalue $\lambda_j(x)$ is simple, we have

$$\nabla f_\mu(x) = \frac{1}{\sum_j \exp(\lambda_j(x)/\mu)} \sum_j \exp(\lambda_j(x)/\mu) \nabla \lambda_j(x). \quad (3.3)$$

On the other hand, differentiating both sides of the equation $p(y - \lambda_j(y)e) = 0$ and evaluating at x leads to

$$\nabla \lambda_j(x) = \frac{1}{p^{(1)}(x - \lambda_j(x)e)} \nabla p(x - \lambda_j(x)e), \quad (3.4)$$

where we define for all $y \in \mathcal{E}$,

$$p^{(1)}(y) := \left. \frac{d}{dt} p(y + te) \right|_{t=0},$$

a directional derivative. (The denominator in (3.4) is nonzero because, by assumption, $\lambda_j(x)$ is a simple eigenvalue – a simple root of $\lambda \mapsto p(x - \lambda e)$ – and hence $t = -\lambda_j(x)$ is a simple root of $t \mapsto p(x + te)$.)

Substituting (3.4) into (3.3) gives an expression for $\nabla f_\mu(x)$ that can be evaluated once the eigenvalues $\lambda_j(x)$ have been computed (to high precision):

$$\nabla f_\mu(x) = \frac{1}{\sum_j \exp(\lambda_j(x)/\mu)} \sum_j \frac{\exp(\lambda_j(x)/\mu)}{p^{(1)}(x-\lambda_j(x)e)} \nabla p(x-\lambda_j(x)e), \quad (3.5)$$

a convex combination of the vectors $\frac{1}{p^{(1)}(x-\lambda_j(x)e)} \nabla p(x-\lambda_j(x)e)$ (for $j = 1, \dots, n$, where n is the degree of p), vectors independent of μ .

In order to provide an identity similar to (3.5) which holds even when eigenvalues are not simple, we introduce the (directional-)derivative polynomials

$$p^{(m)}(y) := \frac{d^m}{dt^m} p(y+te)|_{t=0}, \quad \text{for } m = 0, 1, \dots, n,$$

where n is the degree of p . (Like p , these are hyperbolic polynomials for which e can serve as distinguished direction. Moreover, using the (inverse) Discrete Fourier Transform to perform interpolation, these polynomials and their gradients can be evaluated efficiently if p and its gradients can be evaluated efficiently (see [24]).)

Proposition 3.3. *For $x \in \mathcal{E}$, assume $\{\lambda_j(x)\}$ is the set of distinct eigenvalues of x , and let m_j denote the multiplicity of $\lambda_j(x)$ (that is, the multiplicity of $\lambda_j(x)$ as a root of $\lambda \mapsto p(x-\lambda e)$). Then*

$$\nabla f_\mu(x) = \frac{1}{\sum_j m_j \exp(\lambda_j(x)/\mu)} \sum_j \frac{m_j \exp(\lambda_j(x)/\mu)}{p^{(m_j)}(x-\lambda_j(x)e)} \nabla p^{(m_j-1)}(x-\lambda_j(x)e),$$

a convex combination of the vectors $\frac{1}{p^{(m_j)}(x-\lambda_j(x)e)} \nabla p^{(m_j-1)}(x-\lambda_j(x)e)$ (vectors which are independent of μ).

The proof of Proposition 3.3 is deferred to Appendix B.

As a simple example, consider a second-order constraint $\|Ax+b\|_2 \leq \alpha^T x + \beta$, where A is an $m \times n$ matrix and $b \in \mathbb{R}^m$, $\alpha \in \mathbb{R}^n$, $\beta \in \mathbb{R}$. Assume $x = \bar{e}$ satisfies the inequality strictly.

Introducing an additional variable r , the constraint can be put into conic form via a rotated Lorentz cone:

$$\begin{aligned} r &= 1 \\ (x, r) \in \mathcal{K} &:= \{(x, r) : r \geq 0 \text{ and } \|Ax+rb\|_2^2 \leq r(\alpha^T x + \beta r)\}. \end{aligned}$$

A hyperbolic polynomial for \mathcal{K} is given by

$$\begin{aligned} p(x, r) &:= \frac{1}{2} (\|Ax+rb\|_2^2 - r(\alpha^T x + \beta r)) \\ &= \frac{1}{2} \begin{bmatrix} x^T & r \end{bmatrix} \begin{bmatrix} A^T A & A^T b - \frac{1}{2} \alpha \\ b^T A - \frac{1}{2} \alpha^T & \|b\|_2^2 - \beta \end{bmatrix} \begin{bmatrix} x \\ r \end{bmatrix}, \end{aligned} \quad (3.6)$$

and $e := (\bar{e}, 1)$ can serve as distinguished direction.

Then for $y = \begin{bmatrix} x \\ r \end{bmatrix}$, the eigenvalues are simply the roots $\lambda^+(y)$ and $\lambda^-(y)$ for the quadratic polynomial $\lambda \mapsto \frac{1}{2}(y-\lambda e)^T B(y-\lambda e)$, where B is the matrix in (3.6). Moreover, we have the following easily-verified identities: $\nabla p(y) = By$, $p^{(1)}(y) = e^T By$, $\nabla p^{(1)}(y) = Be$ and $p^{(2)}(y) = e^T Be$.

Fixing $y \in \mathbb{R}^{n+1}$, let $\lambda^+ := \lambda^+(y)$, $y^+ := y - \lambda^+ e$, and define λ^- , y^- similarly. If $\lambda^+ \neq \lambda^-$, then

$$\nabla f_\mu(y) = \frac{1}{\exp(\lambda^+/\mu) + \exp(\lambda^-/\mu)} \left(\frac{\exp(\lambda^+/\mu)}{e^T B y^+} B y^+ + \frac{\exp(\lambda^-/\mu)}{e^T B y^-} B y^- \right),$$

whereas if $\lambda^+ = \lambda^-$ (i.e., if y lies in the subspace spanned by e and the null space of B), then

$$\nabla f_\mu(y) = \frac{1}{e^T B e} B e,$$

independent of μ .

More generally, consider a collection of constraints $\|A_i x + b_i\|_2 \leq \alpha_i^T x + \beta_i$ ($i = 1, \dots, \ell$), and assume $x = \bar{e}$ strictly satisfies each inequality. The constraints can be expressed in conic form as

$$\begin{aligned} r &= 1 \\ (x, r) &\in \mathcal{K} := \bigcap_i \mathcal{K}_i, \end{aligned}$$

where $\mathcal{K}_i := \{(x, r) : r \geq 0 \text{ and } \|A_i x + r b_i\|_2^2 \leq r(\alpha_i^T x + \beta_i r)\}$. A hyperbolic polynomial for \mathcal{K} is given by the product $p = p_1 \cdots p_\ell$, where $p_i(y) = \frac{1}{2} y^T B_i y$ and B_i is the matrix in (3.6) but with entries subscripted. Again $e = (\bar{e}, 1)$ can serve as distinguished direction.

Now f_μ is a smoothing of the function $y \mapsto \max\{\lambda_i^+(y), \lambda_i^-(y) : i = 1, \dots, \ell\}$. To compute $\nabla f_\mu(y)$, one can first perform the simple computations above to compute $\nabla f_{i,\mu}(y)$ for each i , then compute the convex combination of these gradients using the weights prescribed by Proposition 3.2.

3.2. The setup for applying accelerated methods. For $\mu > 0$, let $\hat{f}_\mu : \mathcal{E} \rightarrow \mathbb{R}$ be the smooth concave function

$$\hat{f}_\mu(x) := -f_\mu(-x) = -\mu \ln \sum_j \exp(-\lambda_j(x)/\mu),$$

an approximation to the function $x \mapsto \lambda_{\min}(x)$ and which satisfies

$$\lambda_{\min}(x) - \mu \ln n \leq \hat{f}_\mu(x) \leq \lambda_{\min}(x). \quad (3.7)$$

(The “hat” on \hat{f}_μ is the graph of a concave function, and is meant to remind the reader that \hat{f}_μ is concave.)

We have three optimization problems,

$$\text{HP} \begin{cases} \min & c \cdot x \\ \text{s.t.} & Ax = b \\ & x \in \mathcal{K} \end{cases} \equiv \max_{\text{s.t. } x \in \text{Affine}_z} \lambda_{\min}(x) \approx \max_{\text{s.t. } x \in \text{Affine}_z} \hat{f}_\mu(x), \quad (3.8)$$

the first two being equivalent (assuming $z < c \cdot e$), and the third being a smoothing of the second. To approximately solve HP, we apply an accelerated method to the third problem, with μ chosen appropriately.

In applying an accelerated method, we adopt a Euclidean perspective, endowing \mathcal{E} with an inner product $\langle \cdot, \cdot \rangle$ to be used for computation, so as to make it easy to retain feasibility for the smoothed problem, simply by orthogonally projecting gradients onto \mathcal{L} , the subspace of which Affine_z is a translate. Let $P_{\mathcal{L}}$ denote the projection operator.

The inner product $\langle \cdot, \cdot \rangle$ – the “computational inner product” – can differ from the inner product used in expressing HP, i.e., the inner product $(u, v) \mapsto u \cdot v$. We impose no conditions on the computational inner product.

Let $\|\cdot\|$ denote the corresponding Euclidean norm, $\|v\| = \langle v, v \rangle^{1/2}$.

The computational inner product provides an identification of \mathcal{E} with its dual space. Henceforth we use $\nabla \hat{f}_\mu(x)$ to denote the gradient viewed as an element of \mathcal{E} .³

In this setup, the relevant Lipschitz constant for accelerated methods is a value $L_{\mu,z}$ satisfying

$$\|P_{\mathcal{L}} \nabla \hat{f}_\mu(x) - P_{\mathcal{L}} \nabla \hat{f}_\mu(y)\| \leq L_{\mu,z} \|x - y\| \quad \text{for all } x, y \in \text{Affine}_z.$$

³That is, $g = \nabla \hat{f}_\mu(x)$ is the vector in \mathcal{E} satisfying $0 = \lim_{\|\Delta x\| \rightarrow 0} \frac{\hat{f}_\mu(x + \Delta x) - \hat{f}_\mu(x) - \langle g, \Delta x \rangle}{\|\Delta x\|}$.

Generally, the best such constant depends on z , but the dependence is mild. Considering the special case $\mu = 1$, we are thus led to define a Lipschitz constant applying to all translates of \mathcal{L} :

$$L_1 = \sup_{u \in \mathcal{E}} \sup \left\{ \frac{\|P_{\mathcal{L}} \nabla \hat{f}_1(u+v) - P_{\mathcal{L}} \nabla \hat{f}_1(u+w)\|}{\|v-w\|} : v, w \in \mathcal{L} \text{ and } v \neq w \right\}. \quad (3.9)$$

Depending on the subspace \mathcal{L} , the value L_1 can be far smaller than the Lipschitz constant for the unrestricted gradient map, that is, for the case $\mathcal{L} = \mathcal{E}$.

Observe that $\hat{f}_\mu(x) = \mu \hat{f}_1(\frac{1}{\mu}x)$ (due to positivity of μ , the definition of \hat{f}_μ , and the homogeneity $\lambda_j(\frac{1}{\mu}x) = \lambda_j(x)/\mu$). Consequently, if \hat{f}_μ is substituted for \hat{f}_1 in (3.9), the supremum is equal to L_1/μ .

The value L_1 both reflects properties of the specific subspace \mathcal{L} and immediately gives the Lipschitz constant L_1/μ for the gradient map $x \mapsto P_{\mathcal{L}} \nabla \hat{f}_\mu(x)$ restricted to any of the affine spaces Affine_z (thus eliminating the need for a Lipschitz constant $L_{\mu,z}$ subscripted with parameters μ and z).

The following corollary to Proposition 3.1 provides geometric insight into how the constant L_1 reflects the particular subspace \mathcal{L} in relation to the hyperbolicity cone \mathcal{K} , as well as reflects the choices for the distinguished direction e and the computational inner product $\langle \cdot, \cdot \rangle$ (for which, we recall, the norm is written $\| \cdot \|$).

Define

$$r_e = \sup \{ r : B(e, r) \cap \text{Affine}_{c \cdot e} \subseteq \mathcal{K} \}, \quad (3.10)$$

where $B(e, r)$ is the $\| \cdot \|$ -ball of radius r centered at e .

Corollary 3.4. $L_1 \leq 1/r_e^2$.

Proof: As $\bar{B}_\infty(e, 1)$ is the largest set both contained in \mathcal{K} and having symmetry point e , the intersection $\bar{B}_\infty(e, 1) \cap \text{Affine}_{c \cdot e}$ is the largest set both contained in $\mathcal{K} \cap \text{Affine}_{c \cdot e}$ and having symmetry point e . Thus, the Euclidean norm satisfies

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

from which follows

$$\|\bar{v}\|_\infty \leq \|\bar{v}\|/r_e \quad \text{for all } \bar{v} \in \mathcal{L}.$$

Consequently, for $u \in \mathcal{E}$ and $v, w \in \mathcal{L}$,

$$\begin{aligned} & \|P_{\mathcal{L}} \nabla \hat{f}_1(u+v) - P_{\mathcal{L}} \nabla \hat{f}_1(u+w)\| \\ &= \max \{ \langle \bar{v}, \nabla \hat{f}_1(u+v) - \nabla \hat{f}_1(u+w) \rangle : \bar{v} \in \mathcal{L} \text{ and } \|\bar{v}\| \leq 1 \} \\ &\leq \max \{ \langle \bar{v}, \nabla \hat{f}_1(u+v) - \nabla \hat{f}_1(u+w) \rangle : \bar{v} \in \mathcal{E} \text{ and } \|\bar{v}\|_\infty \leq 1/r_e \} \\ &= \|\nabla \hat{f}_1(u+v) - \nabla \hat{f}_1(u+w)\|_\infty^* / r_e \\ &\leq \|v-w\|_\infty / r_e \quad (\text{by Proposition 3.1}) \\ &\leq \|v-w\| / r_e^2, \end{aligned}$$

completing the proof (due to the definition of L_1). \square

We close the section with a lemma that ties accuracy in solving the smoothed problem to the accuracy in solving HP, the problem of primary interest.

Recall z^* denotes the optimal value of HP, the optimization problem of primary interest – the first problem in (3.8). Recall $\lambda_{\min, z}^*$ is the optimal value of the HP-equivalent problem – the second problem in (3.8) – and x_z^* is any optimal solution. Let $f_{\mu, z}^*$ denote the optimal value of the smoothed problem – the third problem – and let $x_{\mu, z}^*$ be an optimal solution.

Lemma 3.5. *If $x \in \text{Affine}_z$ where $z < c \cdot e$, then the projection $\pi = \pi(x)$ satisfies*

$$\frac{c \cdot \pi - z^*}{c \cdot e - z^*} \leq \left(\hat{f}_{\mu,z}^* - \hat{f}_\mu(x) + \mu \ln n \right) \frac{c \cdot e - z^*}{c \cdot e - z} .$$

Proof: Define $\epsilon := \frac{c \cdot \pi - z^*}{c \cdot e - z^*}$. According to Proposition 2.5,

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

Thus, to prove the lemma it suffices to show

$$\lambda_{\min,z}^* - \lambda_{\min}(x) \leq \hat{f}_{\mu,z}^* - \hat{f}_\mu(x) + \mu \ln n .$$

However, using (3.7),

$$\begin{aligned} \lambda_{\min,z}^* - \lambda_{\min}(x) &= (\lambda_{\min,z}^* - \hat{f}_\mu(x_z^*)) + (\hat{f}_\mu(x_z^*) - \hat{f}_{\mu,z}^*) \\ &\quad + (\hat{f}_{\mu,z}^* - \hat{f}_\mu(x)) + (\hat{f}_\mu(x) - \lambda_{\min}(x)) \\ &\leq \mu \ln n + 0 + (\hat{f}_{\mu,z}^* - \hat{f}_\mu(x)) + 0 , \end{aligned}$$

completing the proof. \square

4. Applying Accelerated First-Order Methods

Virtually any (accelerated) first-order method can be applied to solving

$$\begin{aligned} \max \quad & \hat{f}_\mu(x) \\ \text{s.t.} \quad & x \in \text{Affine}_z , \end{aligned} \tag{4.1}$$

as feasibility can be retained using only orthogonal projection of gradients onto a subspace, the same subspace at every iteration. We develop results for an accelerated method due to Nesterov [22]. Among the accelerated methods not requiring (an upper bound on) the Lipschitz constant as input, this algorithm streamlines most nicely when specialized to the setting where the only projections are onto a subspace. However, similar complexity results could be obtained if instead we employed the universal method of Lan [13], or of Nesterov [18], and would be obtained if we employed any of numerous accelerated methods that *do* require a Lipschitz constant as input.

Following is the accelerated method [22, (4.9)] specialized to our setting (see Appendix C for discussion). To simplify specification of the algorithm, the distinguished direction e is viewed as fixed, not input.

AGM (Accelerated Gradient Method)

(0) Input: $L_1^\bullet > 0$, an estimate of the value L_1 defined by (3.9),

$\mu > 0$, and

$x_0 \in \text{Affine}$ satisfying $c \cdot x_0 < c \cdot e$.

Initialize: $L = L_1^\bullet / \mu$, $v = x_0$, $A = 0$ and $k = 0$.

(1) Compute $a = \frac{1 + \sqrt{1 + 2AL}}{L}$, $y = \frac{A}{A+a}x_k + \frac{a}{A+a}v$ and $x = y + \frac{1}{L}P_{\mathcal{L}}\nabla\hat{f}_\mu(y)$.

(2) If $\langle \nabla\hat{f}_\mu(x), y - x \rangle > \frac{1}{L} \|P_{\mathcal{L}}\nabla\hat{f}_\mu(x)\|^2$, let $L \leftarrow 2L$ and return to Step 1.

(3) Define $x_{k+1} = x$ and let $v \leftarrow v + aP_{\mathcal{L}}\nabla\hat{f}_\mu(x)$, $L \leftarrow L/2$, $A \leftarrow A + a$,
 $k \leftarrow k + 1$, then go to Step 1.

An important aspect of AGM is that in moving from iterate x_k to x_{k+1} , numerous gradient steps $x = y + \frac{1}{L}P_{\mathcal{L}}\nabla\hat{f}_\mu(y)$ might have to be tried before arriving at a step which is not rejected

in (2). We are thus careful to distinguish between the number of iterates x_k in a sequence and the number of gradients computed when constructing the sequence.

It is easily seen that all iterates x_k lie in Affine_z for $z = c \cdot x_0$.

Recall $\hat{f}_{\mu,z}^*$ denotes the optimal value of (4.1), and $x_{\mu,z}^*$ is any optimal solution.

Theorem 4.1. (Nesterov) *Let $z = c \cdot x_0$ and $k^\bullet := \max\{0, \lceil \log_2(L_1^\bullet/L_1) \rceil\}$. Then*

$$k > k^\bullet \quad \Rightarrow \quad \hat{f}_{\mu,z}^* - \hat{f}_\mu(x_k) \leq \frac{2L_1}{\mu} \left(\frac{\|x_0 - x_{\mu,z}^*\|}{k - k^\bullet} \right)^2.$$

Moreover, for $k > k^\bullet$, the total number of gradients computed before x_k is determined does not exceed

$$4(k - k^\bullet) + 2 \lceil \log_2(L_1^\bullet/L_1) \rceil + 6. \quad (4.2)$$

Results from [22] can readily be assembled into a proof of Theorem 4.1, as is demonstrated in Appendix C.

For recasting Theorem 4.1 into a form meaningful to HP (the optimization problem of primary interest), we follow the approach introduced in [26], to which the present paper is a sequel.

The *level sets* for HP are the sets

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

that is, the largest feasible sets for HP on which the objective function is constant⁴. If $z < z^*$, then $\text{Level}_z = \emptyset$.

If some level set is unbounded, then either HP has unbounded optimal value or can be made to have unbounded optimal value with an arbitrarily small perturbation of c . In the same vein, a level set with large diameter is an indication that perturbations can result in relatively large changes in optimal value.

For scalars z , let

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

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

Recall $\lambda_{\min,z}^*$ denotes the optimal value of the HP-equivalent problem (the problem to which (4.1) is a smoothing).

Corollary 4.2. *Assume input x_0 is strictly feasible for HP, and let $z := c \cdot x_0$. Assume ϵ' is a scalar satisfying $0 < \epsilon' \leq 2\lambda_{\min}(x_0)$. If input μ is chosen as $\mu = \frac{1}{2}\epsilon'/\ln(n)$, then*

$$k > k^\bullet + \frac{2 \text{diam}_z \sqrt{2L_1 \ln n}}{\epsilon'} \quad \Rightarrow \quad \lambda_{\min,z}^* - \lambda_{\min}(x_k) \leq \epsilon',$$

where $k^\bullet := \max\{0, \lceil \log_2(L_1^\bullet/L_1) \rceil\}$.

Moreover, for $k > k^\bullet$, the total number of gradients computed before x_k is determined does not exceed $4(k - k^\bullet) + 2 \lceil \log_2(L_1^\bullet/L_1) \rceil + 6$.

Proof: Fix $\mu = \frac{1}{2}\epsilon'/\ln(n)$. It is readily seen from (3.7) that for $x \in \text{Affine}_z$,

$$\hat{f}_\mu^* - \hat{f}_\mu(x) \leq \frac{1}{2}\epsilon' \quad \Rightarrow \quad \lambda_{\min,z}^* - \lambda_{\min}(x) \leq \epsilon'.$$

⁴There 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 choice of 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.

Thus, Theorem 4.1 is easily seen to imply

$$k > k^\bullet + \frac{2\|x_0 - x_{\mu,z}^*\| \sqrt{2L_1 \ln n}}{\epsilon'} \Rightarrow \lambda_{\min,z}^* - \lambda_{\min}(x) \leq \epsilon'.$$

To complete the proof, it suffices to show $x_{\mu,z}^*$ is feasible for HP, as then the inequality $\|x_0 - x_{\mu,z}^*\| \leq \text{diam}_z$ follows from the assumed feasibility of x_0 .

However, since $0 < \epsilon' \leq 2\lambda_{\min}(x_0)$, we have

$$\mu = \frac{1}{2}\epsilon' / \ln(n) \leq \lambda_{\min}(x_0) / \ln(n),$$

and thus (3.7) implies $\hat{f}_\mu(x_0) \geq 0$. Hence, $\hat{f}_\mu(x_{\mu,z}^*) \geq 0$, which by (3.7) implies $\lambda_{\min}(x_{\mu,z}^*) \geq 0$, that is, implies $x_{\mu,z}^*$ is indeed feasible for HP. \square

For the iterates x_k of the accelerated method, define $\pi_k = \pi(x_k)$. The main algorithmic goal is achieved when π_k is computed for which $\frac{c \cdot \pi_k - z^*}{c \cdot e - z^*} \leq \epsilon$. Proposition 2.5 shows the goal is accomplished precisely when iterate x_k satisfies

$$\lambda_{\min,z}^* - \lambda_{\min}(x_k) \leq \frac{\epsilon}{1 - \epsilon} \frac{c \cdot e - z}{c \cdot e - z^*}. \quad (4.3)$$

It is thus tempting to try making use of Corollary 4.2 with ϵ' chosen to be the value on the right of (4.3). However, this is an undesirable approach for multiple reasons, one being that the quantity on the right of (4.3) depends on the optimal value z^* (which we want to avoid assuming known), and another being that the resulting bound on the total number of gradients computed would be of the form

$$\mathcal{O} \left(\frac{\text{diam}_z \sqrt{L_1 \ln n}}{\epsilon} \cdot \frac{c \cdot e - z^*}{c \cdot e - z} + \left| \log \frac{L_1^\bullet}{L_1} \right| \right), \quad (4.4)$$

If z – the objective value for the known feasible point x_0 – is only slightly smaller than $c \cdot e$, then the ratio $\frac{c \cdot e - z^*}{c \cdot e - z}$ can be large compared to $1/\epsilon$, in which case the bound (4.4) is entirely uninteresting.

To arrive at an algorithm which does not require z^* to be known, and for which the disconcerting product $\frac{1}{\epsilon} \frac{c \cdot e - z^*}{c \cdot e - z}$ does not appear as it does in (4.4), we must go beyond simply applying an accelerated method to the smoothed problem (4.1). Towards presenting a strategy which accomplishes the goals, we first record a result highlighting a case where the ratio $\frac{c \cdot e - z^*}{c \cdot e - z}$ is small, leading to a bound of the form (4.4) which is nice, not disconcerting.

Corollary 4.3. *Assume the input x_0 is strictly feasible for HP and satisfies*

$$\lambda_{\min}(x_0) \geq 1/6 \quad \text{and} \quad \frac{c \cdot e - z^*}{c \cdot e - z} \leq 3, \quad (4.5)$$

where $z := c \cdot x_0$. Let $0 < \epsilon < 1$.

If input μ is chosen as $\mu = \frac{1}{6}\epsilon / \ln(n)$, then for the projected iterates $\pi_k := \pi(x_k)$,

$$k > k^\bullet + \frac{6 \text{diam}_z \sqrt{2L_1 \ln n}}{\epsilon} \Rightarrow \frac{c \cdot \pi_k - z^*}{c \cdot e - z^*} \leq \epsilon, \quad (4.6)$$

where $k^\bullet := \max\{0, \lceil \log_2(L_1^\bullet/L_1) \rceil\}$.

Moreover, for $k > k^\bullet$, the total number of gradients computed before x_k is determined does not exceed $4(k - k^\bullet) + 2 \lceil \log_2(L_1^\bullet/L_1) \rceil + 6$.

Proof: Let $\epsilon' := \epsilon/3 \leq 1/3$. Then, by assumption, $\epsilon' \leq 2\lambda_{\min}(x_0)$, and hence Corollary 4.2 can be applied with $\mu = \frac{1}{2}\epsilon'/\ln(n) = \frac{1}{6}\epsilon/\ln(n)$, giving

$$\begin{aligned} k > k^\bullet + \frac{6 \operatorname{diam}_z \sqrt{2L_1 \ln n}}{\epsilon} &\Rightarrow \lambda_{\min,z}^* - \lambda_{\min}(x_k) \leq \epsilon/3 \\ &\Rightarrow \lambda_{\min,z}^* - \lambda_{\min}(x_k) \leq \frac{\epsilon}{1-\epsilon} \frac{c \cdot e - z}{c \cdot e - z^*}, \end{aligned}$$

where the last implication is due to the rightmost inequality assumed in (4.5). Invoking Proposition 2.5 completes the proof. \square

The threshold on k given by (4.6) can be given geometric meaning by recalling the value r_e defined by (3.10), a value reflecting the extent to which the computational inner product conforms to the geometry of the set obtained by intersecting the cone \mathcal{K} with $\text{Affine}_{c \cdot e}$. Substituting $L_1 \leq 1/r_e^2$ (by Corollary 3.4) gives

$$k > k^\bullet + \frac{6 \sqrt{\ln n}}{\epsilon} \frac{\operatorname{diam}_z}{r_e} \Rightarrow \frac{c \cdot \pi_k - z^*}{c \cdot e - z^*} \leq \epsilon.$$

Here, the strongest result is obtained, of course, if the computational inner product is an inner product which minimizes the ratio $\operatorname{diam}_z/r_e$.

Corollary 4.3 is limited in its scope, as the second assumption in (4.5) is strong. Nonetheless, to obtain an efficient algorithm, we apply the accelerated method sequentially in such a manner that on the final application, the assumptions of Corollary 4.3 are satisfied.

Our main algorithm, given below, has both outer iterations and inner iterations. At the beginning of an outer iteration, two different inputs to AGM are specified, on which the method is run separately, in parallel. To simplify the presentation, it is useful to introduce notation which reflects the relevant choices of input to AGM – here, the approximation L_1^\bullet of the Lipschitz constant L_1 should be viewed as fixed (along with e , the distinguished direction):

$$\text{AGM}(\mu, v) \quad - \quad \text{AGM run with input } e, L_1^\bullet, \mu \text{ and } x_0 = v.$$

Likewise, to simplify specification of the main algorithm, L_1^\bullet (and e) are viewed as fixed, not as input.

MainAlgo

Input: $0 < \epsilon < 1$ (the accuracy desired in solving HP), and

$$v_1 \in \text{Affine satisfying } c \cdot v_1 < c \cdot e \text{ and }^5 \lambda_{\min}(v_1) = 1/4.$$

Fix: $\mu_1 = 1/(12 \ln n)$, and

$$\mu_2 = \epsilon/(6 \ln n).$$

Initiate: $\ell = 1$ (counter for outer iterations).

Procedure:

(1) Start, in parallel, $\text{AGM}_1 := \text{AGM}(\mu_1, v_\ell)$ and $\text{AGM}_2 := \text{AGM}(\mu_2, v_\ell)$, letting the implementations run indefinitely unless an iterate w_j for AGM_1 is obtained for which $\lambda_{\min}(w_j) \geq 1/2$, in which case go to Step 2.

(2) For the iterate w_j satisfying $\lambda_{\min}(w_j) \geq 1/2$, let $v_{\ell+1} := e + \frac{3}{4} \frac{1}{1-\lambda_{\min}(w_j)}(w_j - e)$, an “outer iterate.” Let $\ell \leftarrow \ell + 1$, then go to Step 1, restarting AGM_1 and AGM_2 with the new v_ℓ (but leaving μ_1 and μ_2 unchanged).

⁵For any vector $v \in \text{Affine}$ satisfying $c \cdot v < c \cdot e$, the vector $v_1 = e + \frac{3}{4} \frac{1}{1-\lambda_{\min}(v)}(v - e)$ is appropriate as input.

In practice, a smarter way to proceed might be to update the estimate L_1^\bullet based on the last values for L just before AGM_1 and AGM_2 are restarted.

Denote the sequence of iterates for AGM_1 by w_1, w_2, \dots , where the indexing is *not* restarted when AGM_1 is restarted. (Thus, if w_j satisfies the property which causes Step 2 to be invoked and the two implementations to be restarted, the next iterate produced by AGM_1 is w_{j+1} .) Likewise, denote the sequence of iterates for AGM_2 by x_1, x_2, \dots , where the indexing is *not* restarted when AGM_2 is restarted.

To establish a bound on the total number of gradients computed, we assume AGM_1 and AGM_2 compute gradients simultaneously, so that at any given time, both implementations have computed the same number of gradients.

Let $z_1 := c \cdot v_1$, where v_1 is input to `MainAlgo`, and define

$$\text{Diam}_{z_1} = \max\{\text{diam}_z : z \leq z_1\}.$$

Theorem 4.4. *MainAlgo produces an iterate x_k for which $\pi_k := \pi(x_k)$ satisfies $\frac{c \cdot \pi_k - z^*}{c \cdot e - z^*} \leq \epsilon$, and does so within computing a total number of gradients not exceeding*

$$\mathcal{O} \left(\frac{\text{Diam}_{z_1} \sqrt{L_1 \ln n}}{\epsilon} + \left(1 + \log \frac{c \cdot e - z^*}{c \cdot e - z_1} \right) \cdot \left(1 + \text{Diam}_{z_1} \sqrt{L_1 \ln n} + \left| \log \frac{L_1^\bullet}{L_1} \right| \right) \right).$$

Remark 1: Unlike the bound (4.4), here the ratio $\frac{c \cdot e - z^*}{c \cdot e - z_1}$ has been entirely separated from terms involving ϵ . Moreover, now only the logarithm of the ratio appears.

Remark 2: Similar to the observation immediately following Corollary 4.3, if $1/r_e^2$ is substituted for L_1 in $\sqrt{L_1 \ln n}$ (Corollary 3.4), the result is strongest when the computational inner product is chosen to minimize the ratio Diam_{z_1}/r_e .

Proof of Theorem 4.4: As noted just before Theorem 4.1, all iterates of AGM have the same objective value as the initial iterate. Consequently, for sequences of iterates w_j and x_k generated by `MainAlgo`, all iterates have the same objective value so long as Step 2 is not invoked while the sequences are being generated. Let us examine the change in objective value when Step 2 is invoked.

Step 2 is called when an iterate w_j is obtained for which $\lambda_{\min}(w_j) \geq 1/2$. Fixing such an iterate for now, let $z := c \cdot w_j$. AGM_1 and AGM_2 are started anew with input $v_+ = e + \frac{3}{4} \frac{1}{1 - \lambda_{\min}(w_j)} (w_j - e)$. Then, until Step 2 is again invoked, all iterates have objective value

$$z_+ := c \cdot v_+ = c \cdot e + \frac{3}{4} \frac{1}{1 - \lambda_{\min}(w_j)} (z - c \cdot e),$$

which satisfies

$$\frac{c \cdot e - z_+}{c \cdot e - z} \geq \frac{3}{4} \frac{1}{1 - \frac{1}{2}} = \frac{3}{2}. \quad (4.7)$$

We observe, too, that for Step 2 to have been invoked implies

$$\frac{c \cdot e - z}{c \cdot e - z^*} \leq \frac{1}{2}. \quad (4.8)$$

Indeed,

$$\begin{aligned} \frac{c \cdot e - z}{c \cdot e - z^*} &= 1 - \frac{z - z^*}{c \cdot e - z^*} \\ &= 1 - \lambda_{\min, z}^* \quad (\text{by Lemma 2.4}) \\ &\leq 1 - \lambda_{\min}(w_k) \\ &\leq 1/2. \end{aligned}$$

Letting $v_1, v_2, \dots, v_{\ell'}$ be the first ℓ' outer iterates computed by MainAlgo, and letting $z_\ell := c \cdot v_\ell$, from (4.7) follows

$$(c \cdot e - z_\ell) \geq \left(\frac{3}{2}\right)^{\ell-1} (c \cdot e - z_1) \quad (\ell = 1, \dots, \ell'), \quad (4.9)$$

whereas (4.8) implies

$$\frac{1}{2}(c \cdot e - z^*) \geq c \cdot e - z_{\ell'-1}. \quad (4.10)$$

Together, (4.9) and (4.10) give $\frac{1}{2}(c \cdot e - z_1) \geq \left(\frac{3}{2}\right)^{\ell'-2} (c \cdot e - z_1)$, and hence,

$$\ell' < 1 + \log_{3/2} \frac{c \cdot e - z^*}{c \cdot e - z_1} \quad (4.11)$$

(where we have used $\log_{3/2}(1/2) < -1$). Trivially, the number of outer iterates is thus finite. Henceforth we use ℓ' to denote the final index of the outer iterations, that is, $v_{\ell'}$ is the last outer iterate computed by MainAlgo.

Next, we claim for the decreasing sequence $z_1, z_2, \dots, z_{\ell'}$,

$$\ell \leq \ell' - 2 \quad \Rightarrow \quad \frac{c \cdot e - z_\ell}{c \cdot e - z^*} \leq 1/3. \quad (4.12)$$

To understand, assume $\ell < \ell'$ and z_ℓ does not satisfy the inequality on the right of (4.12). Then (4.7) implies

$$c \cdot e - z_{\ell+1} \geq \frac{3}{2}(c \cdot e - z_\ell) > \frac{1}{2}(c \cdot e - z^*),$$

and thus, by Lemma 2.4,

$$\lambda_{\min, z_{\ell+1}}^* = 1 - \frac{c \cdot e - z_{\ell+1}}{c \cdot e - z^*} < 1/2.$$

Consequently, when AGM_1 and AGM_2 are started anew with input $v_{\ell+1}$, each iterate w_j must satisfy $\lambda_{\min}(w_j) < 1/2$, implying Step 2 will never again be invoked, that is, implying $\ell+1 = \ell'$. The implication (4.12) is thus established.

Assuming for index ℓ that the inequality on the right of (4.12) is satisfied, we now provide an upper bound on the number of gradients computed by AGM_1 in the ℓ^{th} outer iteration, that is, the number of gradients computed between when AGM_1 and AGM_2 were started with input v_ℓ and when an iterate w_j is encountered for which $\lambda_{\min}(w_j) \geq 1/2$. Here we rely on the equality $\lambda_{\min}(v_\ell) = 1/4$ (MainAlgo explicitly requires this equality for input v_1 , and it holds for subsequent outer-iterates due to (2.1) and to the definition of $v_{\ell+1}$ in Step 2 of MainAlgo).

Since $\frac{c \cdot e - z_\ell}{c \cdot e - z^*} \leq \frac{1}{3}$ (by assumption), Lemma 2.4 implies

$$\lambda_{\min, z_\ell}^* = 1 - \frac{c \cdot e - z_\ell}{c \cdot e - z^*} \geq 2/3.$$

Consequently, applying Corollary 4.2 with the choices $x_0 = v_\ell$ and $\epsilon' = 1/6$, we find that in the ℓ^{th} outer iteration, AGM_1 determines an iterate satisfying $\lambda_{\min}(w_j) \geq 1/2$ within a number of gradient computations not exceeding

$$48 \text{diam}_{z_\ell} \sqrt{2L_1 \ln n} + 2 \lceil \log_2(L_1^\bullet/L_1) \rceil + 6. \quad (4.13)$$

(Upon determining w_j , of course, the $(\ell+1)^{\text{th}}$ outer iteration commences.) As AGM_1 and AGM_2 are assumed to compute gradients simultaneously, the same bound holds for AGM_2 .

In addition to having a bound on the number of gradients computed during outer iteration ℓ when $\frac{c \cdot e - z_\ell}{c \cdot e - z^*} \leq \frac{1}{3}$, we now know for $\ell = \ell'$ (the final outer iteration), $\frac{c \cdot e - z_\ell}{c \cdot e - z^*} > \frac{1}{3}$, simply because Main Algorithm never terminates. Moreover, according to the implication (4.12), outer iteration $\ell = \ell' - 1$ is the only other outer iteration that might satisfy $\frac{c \cdot e - z_\ell}{c \cdot e - z^*} > \frac{1}{3}$.

Consider the case that $\ell = \ell'$ is the only outer iteration for which $\frac{c \cdot e - z_\ell}{c \cdot e - z^*} > \frac{1}{3}$. Then, from above, the total number of gradients computed by AGM_1 before reaching outer iteration ℓ' does not exceed

$$(\ell' - 1) \left(48 \text{Diam}_{z_1} \sqrt{2L_1 \ln n} + 2|\log_2(L_1^\bullet/L_1)| + 6 \right). \quad (4.14)$$

On the other hand, since $\frac{c \cdot e - z_{\ell'}}{c \cdot e - z^*} > 1/3$ and $\lambda_{\min}(v_{\ell'}) = 1/4 > 1/6$, Corollary 4.3 implies that at some time during outer iteration ℓ' , an iterate x_k will be determined by AGM_2 for which

$$\frac{c \cdot \pi(x_k) - z^*}{c \cdot e - z^*} \leq \epsilon. \quad (4.15)$$

Furthermore, the corollary implies that during outer iteration ℓ' , the number of gradients computed by AGM_2 before that time is reached is at most

$$24 \text{Diam}_{z_1} \sqrt{2L_1 \ln n} / \epsilon + 2|\log_2(L_1^\bullet/L_1)| + 6. \quad (4.16)$$

Since AGM_1 and AGM_2 are assumed to compute gradients simultaneously, the same bound holds for AGM_1 .

Clearly, then, for the case that $\ell = \ell'$ is the only outer iteration for which $\frac{c \cdot e - z_\ell}{c \cdot e - z^*} > \frac{1}{3}$, the total number of gradients computed by AGM_1 before (4.15) is satisfied does not exceed the sum of (4.14) and (4.16). Likewise for AGM_2 . Using (4.11) to substitute for ℓ' completes the proof of the theorem in this case.

Finally, consider the remaining case, where $\frac{c \cdot e - z_\ell}{c \cdot e - z^*} > \frac{1}{3}$ for $\ell = \ell' - 1$ as well as for $\ell = \ell'$ (although $\frac{c \cdot e - z_\ell}{c \cdot e - z^*} \leq \frac{1}{3}$ for $\ell < \ell' - 1$). Clearly, replacing $\ell' - 1$ in (4.14) by $\ell' - 2$ results in an upper bound on the number gradients computed by AGM_1 before reaching outer iteration $\ell' - 1$.

If in outer iteration $\ell = \ell' - 1$ or $\ell = \ell'$, the number of gradients computed by AGM_2 during the iteration reaches the quantity (4.16), then by the same reasoning as above, an iterate x_k satisfying (4.15) will have been obtained. Consequently, the combined number of gradients computed by AGM_2 in outer iterations $\ell' - 1$ and ℓ' before reaching an iterate x_k satisfying (4.15) is at most twice the value (4.16). Adding this quantity to (4.14) bounds the total number of gradients computed by AGM_2 from the time MainAlgo is initiated until x_k satisfying (4.15) is reached. The same bound holds for AGM_1 . \square

APPENDIX A. Proof of Proposition 3.1

For the proposition, \mathcal{K} is a hyperbolicity cone in \mathcal{E} , with hyperbolic polynomial p , of degree n . Fix $e \in \text{int}(\mathcal{K})$, the distinguished direction for defining eigenvalues $\lambda_j(x)$ for every $x \in \mathcal{E}$, these being the roots of the univariate polynomial $\lambda \mapsto p(x - \lambda e)$.

It suffices to establish the proposition for the case $\mu = 1$, because for any $\mu > 0$ we have $f_\mu(x) = \mu f_1(\frac{1}{\mu}x)$ (due to the homogeneity $\lambda_j(\frac{1}{\mu}x) = \lambda_j(x)/\mu$). Thus, we aim to show the function

$$f(x) := \ln \sum_{j=1}^n \exp(\lambda_j(x)) \quad (\text{A.1})$$

is convex and infinitely Fréchet-differentiable, and satisfies

$$\|\nabla f(x) - \nabla f(y)\|_\infty^* \leq \|x - y\|_\infty \quad \text{for all } x, y \in \mathcal{E}. \quad (\text{A.2})$$

Nesterov [21, (10) and Thm 3] established these facts for the special case $\mathcal{K} = \mathbb{S}^n$, with hyperbolic polynomial $X \mapsto \det(X)$ and distinguished direction I , in which case, of course, $\lambda_j(X)$ are the usual eigenvalues of matrix X . As we rely heavily on Nesterov's results, we capitalize the function in the special case:

$$F(X) = \ln \sum_j \exp(\lambda_j(X)).$$

Our goals are almost entirely accomplished by a straightforward leveraging of Nesterov's results using the following deep theorem.

Helton-Vinnikov Theorem. *If \mathcal{L} is a 3-dimensional subspace of \mathcal{E} , and $e \in \mathcal{L} \cap \text{int}(K)$, then there exists a linear transformation $T : \mathcal{L} \rightarrow \mathbb{S}^n$ satisfying*

$$T(e) = I \quad \text{and} \quad p(x) = p(e) \det(T(x)) \quad \text{for all } x \in \mathcal{L}.$$

This is a ‘‘homogeneous version’’ of the theorem proven by Helton and Vinnikov [11] (also see [28]). The homogenous version was first recorded by Lewis, Parrilo and Ramana in [15], and observed there to settle affirmatively the Lax conjecture (i.e., Lax had conjectured that the statement of the above theorem is true).

Recently, Hanselka [10] gave a proof of the H-V (Helton-Vinnikov) Theorem very different than the original proof, his approach being largely algebraic rather than geometric. Both proofs rely on deep mathematics.

A simple observation is that convexity is a 1-dimensional property – a function with domain \mathcal{E} is convex if and only if it is convex when restricted to lines in \mathcal{E} . Thus, to prove convexity of f , it suffices to consider general pairs $x, y \in \mathcal{E}$ and show f is convex when restricted to the line containing x and y , which certainly is the case if f is convex on the subspace \mathcal{L} spanned by x, y and e . However, letting T be as in the H-V Theorem, it is easy to see $\{\lambda_j(x)\} = \{\lambda_j(T(x))\}$ for every $x \in \mathcal{L}$, and thus, $f(x) = F(T(x))$. Convexity of f on \mathcal{L} is then immediate from convexity of F .⁶

Likewise, for any convex function $f \in \mathcal{C}^1$, Lipschitz continuity of the gradient map is a 1-dimensional property, in that, for example, for any norm $\|\cdot\|_\Delta$ and its dual $\|g\|_\Delta^* := \max\{u : \langle g, u \rangle \leq 1\}$,

$$\|\nabla f(x) - \nabla f(y)\|_\Delta^* \leq L \|x - y\|_\Delta \quad \text{for all } x, y \in \mathcal{E}$$

$$\Leftrightarrow$$

$$f(y) \leq f(x) + \frac{d}{dt} f(x + t(y - x))|_{t=0} + \frac{L}{2} \|y - x\|_\Delta^2 \quad \text{for all } x, y \in \mathcal{E}. \quad (\text{A.3})$$

(The implication ‘‘ \Rightarrow ’’ is standard, but the implication ‘‘ \Leftarrow ’’ holds as well, as is shown in [19] for Euclidean norms, and is remarked in [21, Appendix] to hold generally⁷.)

Thus, to establish the desired Lipschitz continuity (A.2) for the gradient map of the function f defined by (A.1), it suffices to show $f \in \mathcal{C}^1$ and to establish the inequality (A.3) for a general pair $x, y \in \mathcal{E}$, with $L = 1$ and $\|\cdot\|_\Delta = \|\cdot\|_\infty$.

⁶Invoking the Helton-Vinnikov Theorem to establish convexity of f is overkill. For a more insightful understanding of the convexity of f (and the convexity of a vast range of other ‘‘eigenvalue’’ functions), see [1].

⁷Indeed, in the proof of [19, Thm. 2.1.5], for the function $\phi(y) := f(y) - \langle \nabla f(x), y - x \rangle$, simply replace the point $y - \frac{1}{L} \nabla \phi(y)$ with $y - \frac{\|\nabla \phi(y)\|_\Delta^*}{L} u$ where $\|u\|_\Delta = 1$ and $\langle \nabla \phi(y), u \rangle = \|\nabla \phi(y)\|_\Delta^*$.

However, assuming for the moment that f is continuously differentiable, letting \mathcal{L} be the subspace spanned by x, y and e , and letting T be as in the H-V Theorem, for $X := T(x)$ and $Y := T(y)$ we have

$$\begin{aligned} f(y) &= F(Y) \\ &\leq F(X) + \frac{d}{dt}F(X + t(Y - X))|_{t=0} + \frac{1}{2}\|Y - X\|_\infty^2 \\ &= f(x) + \frac{d}{dt}f(x + t(y - x))|_{t=0} + \frac{1}{2}\|y - x\|_\infty^2, \end{aligned}$$

where the inequality is due to Nesterov [21, (10) and Thm 3] already having established the desired Lipschitz continuity for the gradient map of F . Hence, to complete the proof of the desired Lipschitz continuity for gradient map of f , it only remains to show $f \in \mathcal{C}^1$.

Showing that f is continuously Fréchet-differentiable, however, seems out-of-reach by using a simple application of the H-V Theorem, as already two of the three dimensions of \mathcal{L} are taken by e and the point x at which continuous differentiability is to be established. Consequently, for establishing $f \in \mathcal{C}^1$, we use a route that bypasses the H-V Theorem. With only slight addition to the effort, we show $f \in \mathcal{C}^\infty$.

Endow \mathcal{E} with a coordinate system coming from a basis orthonormal with respect to $\langle \cdot, \cdot \rangle$ (any inner product suffices), allowing us to write $x = (x_1, \dots, x_m)$, where $m = \dim(\mathcal{E})$. For a tuple of non-negative integers $\alpha = (i_1, \dots, i_m)$, write $\frac{\partial^\alpha}{\partial x^\alpha}$ for the differential operator $\frac{\partial^{i_1}}{\partial x_1^{i_1}} \frac{\partial^{i_2}}{\partial x_2^{i_2}} \dots \frac{\partial^{i_m}}{\partial x_m^{i_m}}$.

To prove $f \in \mathcal{C}^\infty$, it suffices to show the function $x \mapsto \sum_j \exp(\lambda_j(x))$ is infinitely Fréchet-differentiable. Since $\exp(t) = \sum_{k=0}^\infty \frac{1}{k!} t^k$, it thus suffices to show the functions

$$S_k(x) = \sum_j \lambda_j(x)^k \quad (k = 0, 1, \dots)$$

are infinitely differentiable and have the property that for every tuple of non-negative integers $\alpha = (i_1, \dots, i_m)$, the sequence of functions

$$x \mapsto \sum_{k=0}^\ell \frac{1}{k!} \frac{\partial^\alpha}{\partial x^\alpha} S_k(x) \quad \ell = 0, 1, \dots$$

converges uniformly on bounded subsets of \mathcal{E} . In turn, to establish the uniform convergence, it suffices to show for every $R > 0$ and every α ,

$$\lim_{\ell \rightarrow \infty} \sum_{k=\ell+1}^\infty \frac{1}{k!} \max_{\|x\| \leq R} \left| \frac{\partial^\alpha}{\partial x^\alpha} S_k(x) \right| = 0. \quad (\text{A.4})$$

It happens that S_k is not only infinitely differentiable, but is a homogeneous polynomial of degree k . Indeed, it is a standard fact in the literature on hyperbolicity cones that for any homogeneous polynomial $Q : \mathbb{R}^n \rightarrow \mathbb{R}$ which is symmetric (i.e., invariant under permutations of its arguments), the function $x \mapsto Q(\lambda_1(x), \dots, \lambda_n(x))$ is a homogeneous polynomial of the same degree of Q (this result follows from the fundamental theorem of symmetric polynomials and the fact that the result is easily shown true when Q is any elementary symmetric polynomial (c.f., [24, Prop. 18])).

It only remains to establish (A.4), which we show is a simple corollary to the following theorem published in 1928. (We are indebted to Greg Blekherman [5] for making us aware of this theorem.)

Theorem A.2. (Kellogg [12, Thm IV]) *If $P : \mathbb{R}^m \rightarrow \mathbb{R}$ is a homogeneous polynomial of degree k , then*

$$\max_{\|x\| \leq 1} \|\nabla P(x)\| \leq k \max_{\|x\| \leq 1} |P(x)|.$$

To conclude the proof of the proposition, let

$$M := \max_{\|x\| \leq 1} \sum_j |\lambda_j(x)| .$$

(the maximum is attained and finite, because the roots of a univariate polynomial vary continuously in the coefficients (so long as the leading coefficient does not vanish), and thus the roots of $\lambda \mapsto p(\lambda e - x)$ vary continuously in x). Clearly, for every $k \geq 1$,

$$\max_{\|x\| \leq 1} |S_k(x)| \leq M^k .$$

Consequently, as S_k is a homogeneous polynomial of degree k , induction using Theorem A.2 shows for every tuple of non-negative integers $\alpha = (i_1, \dots, i_m)$ satisfying $I := \sum_j i_j \leq k$,

$$\left| \frac{\partial^\alpha}{\partial x^\alpha} S^k(x) \right| \leq \frac{k!}{(k-I)!} M^k \|x\|^{k-I} .$$

Thus, for $\ell \geq I$,

$$\sum_{k=\ell+1}^{\infty} \frac{1}{k!} \max_{\|x\| \leq R} \left| \frac{\partial^\alpha}{\partial x^\alpha} S_k(x) \right| \leq M^I \sum_{k=\ell+1}^{\infty} \frac{1}{(k-I)!} (MR)^{k-I} ,$$

from which (A.4) is immediate.

APPENDIX B. Proof of Proposition 3.3

For ease of reference, we restate the proposition.

Proposition B.1. *For $x \in \mathcal{E}$, assume $\{\lambda_j(x)\}$ is the set of distinct eigenvalues of x , and let m_j denote the multiplicity of $\lambda_j(x)$ (that is, the multiplicity of $\lambda_j(x)$ as a root of $\lambda \mapsto p(x - \lambda e)$). Then*

$$\nabla f_\mu(x) = \frac{1}{\sum_j m_j \exp(\lambda_j(x)/\mu)} \sum_j \frac{m_j \exp(\lambda_j(x)/\mu)}{p^{(m_j)}(x - \lambda_j(x)e)} \nabla p^{(m_j-1)}(x - \lambda_j(x)e) , \quad (\text{B.1})$$

a convex combination of the vectors $\frac{1}{p^{(m_j)}(x - \lambda_j(x)e)} \nabla p^{(m_j-1)}(x - \lambda_j(x)e)$ (vectors which are independent of μ).

Proof: We first establish that the statement is correct for the hyperbolic polynomial $P(X) = \det(X)$ ($X \in \mathbb{S}^n$) and distinguished direction I , and then rely on the Helton-Vinnikov Theorem (Appendix A) to establish the theorem generally.

For $P(X) = \det(X)$, it is readily verified by induction that the value $P^{(i)}(X)$ is $i!$ (i factorial) times the sum of the $(n - i)^{\text{th}}$ order principal minors of X . Thus, assuming D is a diagonal matrix having zero for its first m diagonal entries, and having nonzero eigenvalues $\lambda_{m+1}, \dots, \lambda_n$, we have $P^{(m)}(D) = m! \prod_{i=m+1}^n \lambda_i$. Moreover, for gradients computed with respect to the trace inner product, it is straightforward (even if notationally cumbersome) to show $\nabla P^{(m-1)}(D)$ is the diagonal matrix whose first m diagonal entries are $(m - 1)! \prod_{i=m+1}^n \lambda_i$ and whose other entries are all zero. Consequently, $\frac{m}{P^{(m)}(D)} \nabla P^{(m-1)}(D)$ is the diagonal matrix whose first m diagonal entries are 1, and whose other entries are 0.

For any $n \times n$ orthogonal matrix Q , it follows that

$$\frac{m}{P^{(m)}(QDQ^T)} \nabla P^{(m-1)}(QDQ^T) = Q_{\cdot, 1:m} (Q_{\cdot, 1:m})^T ,$$

where $Q_{\cdot, 1:m}$ is the matrix formed from the first m columns of Q . Thus,

$\frac{m}{P^{(m)}(QDQ^T)} \nabla P^{(m-1)}(QDQ^T)$ is the matrix projecting \mathbb{R}^n orthogonally onto the null space of QDQ^T .

Consequently, if for $X \in \mathbb{S}^n$, $\lambda_j(X)$ is an eigenvalue of multiplicity m_j , then the matrix $\frac{m_j}{P^{(m_j)}(X - \lambda_j(X)I)} \nabla P^{(m_j-1)}(X - \lambda_j(X)I)$ projects \mathbb{R}^n orthogonally onto the eigenspace corresponding to $\lambda_j(X)$.

Since for the function $F_\mu(X) = \mu \ln \sum_j \exp(\lambda_j(X)/\mu)$, it already is known that

$$\nabla F_\mu(X) = \frac{1}{\sum_j \exp(\lambda_j(X)/\mu)} Q \begin{bmatrix} \exp(\lambda_1(X)/\mu) & & \\ & \ddots & \\ & & \exp(\lambda_n(X)/\mu) \end{bmatrix} Q^T$$

(where $Q \begin{bmatrix} \lambda_1(X) & & \\ & \ddots & \\ & & \lambda_n(X) \end{bmatrix} Q^T$ is an eigendecomposition of X), the proposition immediately follows for the special case of the hyperbolic polynomial $P(X) = \det(X)$ and distinguished direction I .

More generally, let \mathcal{K} be a hyperbolicity cone in a Euclidean space \mathcal{E} , and assume p is a hyperbolic polynomial for \mathcal{K} . Let n be the degree of p , and assume $e \in \text{int}(\mathcal{K})$, the distinguished direction for defining eigenvalues.

We now briefly argue that in proving the proposition, we may assume \mathcal{K} is a regular cone, that is, contains no subspace other than $\{\vec{0}\}$ and has nonempty interior (the latter condition being satisfied simply because $e \in \text{int}(\mathcal{K})$, by assumption).

Let \mathcal{L} denote the lineality space of \mathcal{K} , the largest subspace contained in \mathcal{K} . Then for every $e' \in e + \mathcal{L}$ (Minkowski sum), and for every $x \in \mathcal{E}$, the eigenvalues for x as defined by e' are the same as those defined by e . Moreover, for every $x' \in x + \mathcal{L}$, the eigenvalues of x' are the same as those of x . (For simple justification of these claims, see [27, §8].)

Since the eigenvalues defined by $e' \in e + \mathcal{L}$ are the same as those defined by e , also the same are the values of the function f_μ and polynomials $p^{(i)}$. Consequently, we may assume $e \in \mathcal{L}^\perp$ (where orthogonality is with respect any fixed inner product).

Since the eigenvalues of $x' \in x + \mathcal{L}$ are the same as those of x , and since this holds for all $x \in \mathcal{E}$, we have $\nabla f_\mu(x') = \nabla f_\mu(x)$ and $\nabla p^{(i)}(x') = \nabla p^{(i)}(x)$ for all $x' \in x + \mathcal{L}$, and we have that the gradients lie in \mathcal{L}^\perp .

In light of these observations, in proving the proposition we may restrict attention to x lying in the cone $\mathcal{K} \cap \mathcal{L}^\perp$, a regular cone in the Euclidean space $\mathcal{E} \cap \mathcal{L}^\perp$, and which has hyperbolic polynomial $p|_{\mathcal{E} \cap \mathcal{L}^\perp}$. Thus, as claimed, we may assume \mathcal{K} is a regular cone.

Although gradients depend on the particular inner product, the desired identity (B.1) does not. This allows us to proceed with an inner product of our choosing. We choose the inner product $\langle u, v \rangle_e := \langle u, H(e)v \rangle$, where $H(e)$ is the Hessian at e for the self-concordant logarithmic barrier function $x \mapsto -\ln p(x)$. (Regularity of \mathcal{K} implies invertibility of $H(e)$, and thus guarantees $\langle \cdot, \cdot \rangle_e$ is indeed an inner product – elementary proofs of these well-known facts can be found in [27, §8].)

Fix $x \in \mathcal{E}$, and let

$$g_1 := \nabla f_\mu(x), \quad g_2 := \frac{1}{\sum_j m_j \exp(\lambda_j(x)/\mu)} \sum_j \frac{m_j \exp(\lambda_j(x)/\mu)}{p^{(m_j)}(x - \lambda_j(x)e)} \nabla p^{(m_j-1)}(x - \lambda_j(x)e),$$

where $\{\lambda_j(x)\}$ are the distinct eigenvalues of x , and where m_j is the multiplicity of $\lambda_j(x)$. Our goal is to show $g_1 = g_2$.

Consider a subspace \mathcal{L} whose dimension is at most 3, and which satisfies $e, x \in \mathcal{L} \subseteq \mathcal{E}$. Let $T : \mathcal{L} \rightarrow \mathbb{R}$ be a linear transformation as in the Helton-Vinnikov Theorem:

$$p(e) = I \quad \text{and} \quad p(y) = p(e) \det(T(y)) \quad \text{for all } y \in \mathcal{L}. \quad (\text{B.2})$$

From (B.2) and the definition of $\langle \cdot, \cdot \rangle_e$ follows

$$\langle u, v \rangle_e = \text{trace}(T(u)T(v)) \text{ for all } u, v \in \mathcal{L},$$

that is, T is an isometry between \mathcal{L} (endowed with the inner product $\langle \cdot, \cdot \rangle_e$) and the subspace $T(\mathcal{L})$ (endowed with the trace inner product). Consequently, if \mathcal{C}^1 functions $h : \mathcal{E} \rightarrow \mathbb{R}$ and $H : \mathbb{S}^n \rightarrow \mathbb{R}$ satisfy $h(y) = H(T(y))$ for all $y \in \mathcal{L}$, then for all $y \in \mathcal{L}$,

$$T(\text{proj}_{\mathcal{L}} \nabla h(y)) = \text{Proj}_{T(\mathcal{L})} \nabla H(T(y)), \quad (\text{B.3})$$

where $\text{proj}_{\mathcal{L}}$ (resp., $\text{Proj}_{T(\mathcal{L})}$) is the operator projecting \mathcal{E} orthogonally onto \mathcal{L} (resp., projecting \mathbb{S}^n orthogonally onto $T(\mathcal{L})$).

On the other hand, clear from (B.2) is that for every i ,

$$p^{(i)}(y) = p(e)P^{(i)}(T(y)) \text{ for all } y \in \mathcal{L}. \quad (\text{B.4})$$

Also evident is

$$f_{\mu}(y) = F_{\mu}(T(y)) \text{ for all } y \in \mathcal{L}. \quad (\text{B.5})$$

Letting $X := T(x)$ and $G_1 := \nabla F_{\mu}(X)$, from (B.3) and (B.5) follows

$$T(\text{proj}_{\mathcal{L}} g_1) = \text{Proj}_{T(\mathcal{L})} G_1.$$

Similarly, from (B.3) and (B.4) follows

$$T(\text{proj}_{\mathcal{L}} g_2) = \text{Proj}_{T(\mathcal{L})} G_2,$$

where

$$G_2 := \frac{1}{\sum_j m_j \exp(\lambda_j(X)/\mu)} \sum_j \frac{m_j \exp(\lambda_j(X)/\mu)}{P^{(m_j)}(X - \lambda_j(X)I)} \nabla P^{(m_j-1)}(X - \lambda_j(X)I).$$

However, we have already seen $G_1 = G_2$. Thus, for any 3-dimensional subspace \mathcal{L} containing e and x ,

$$\text{proj}_{\mathcal{L}} g_1 = \text{proj}_{\mathcal{L}} g_2. \quad (\text{B.6})$$

Letting \mathcal{L} be the subspace spanned by e , x and g_1 , it follows from (B.6) that g_1 is the orthogonal projection of g_2 onto the line $\{tg_1 : t \in \mathbb{R}\}$. Similarly, letting \mathcal{L} be the subspace spanned by e , x and g_2 , we find g_2 is the orthogonal projection of g_1 onto the line $\{tg_2 : t \in \mathbb{R}\}$. We conclude $g_1 = g_2$. \square

APPENDIX C. Proof of Theorem 4.1

Theorem 4.1, due to Nesterov, is not quite to be found in [22] although he develops all of the key ingredients. He assumes, for example, that the estimated value for the Lipschitz constant is a lower bound, but does not record the (easier) case that the estimated value is an upper bound.

Additionally, the development in [22] is done in the context of general prox functions, whereas we need only the prox function $x \mapsto \frac{1}{2} \|x\|^2$, as our problems have been recast to have only linear equations as constraints and hence nothing more than orthogonal projection of gradients onto a subspace is needed to ensure iterates remain feasible.

Because Theorem 4.1 is not quite to be found in [22], and because the development there can be streamlined in our setting, in this appendix we provide a proof of the theorem, relying solely on Nesterov's ideas.

Let \mathcal{E} be a Euclidean space, with inner product $\langle \cdot, \cdot \rangle$ and associated norm $\|\cdot\|$. Let $f : \mathcal{E} \rightarrow \mathbb{R}$ be a convex function with Lipschitz-continuous gradient. Let L_f be the Lipschitz constant (possibly unknown to the user). The goal is simply to minimize f with no constraints (linear

equations are easily handled in the theory by taking projected gradient steps rather than gradient steps).

Following is the accelerated method [22, (4.9)], streamlined for this simple setting.

Accelerated Gradient Method

- (0) Input: $x_0 \in \mathcal{E}$ and $L' > 0$, an estimate of L_f .
 Initialize: $v = x_0$, $L = L'$, $A = 0$ and $k = 0$.
- (1) Compute $a = \frac{1+\sqrt{1+2AL}}{L}$, $y = \frac{A}{A+a}x_k + \frac{a}{A+a}v$ and $x = y - \frac{1}{L}\nabla f(y)$.
- (2) If $\langle \nabla f(x), y - x \rangle < \frac{1}{L}\|\nabla f(x)\|^2$ then let $L \leftarrow 2L$ and return to Step 1.
- (3) Define $x_{k+1} = x$ and let $v \leftarrow v - a\nabla f(x)$, $L \leftarrow L/2$, $A \leftarrow A + a$,
 $k \leftarrow k + 1$, and go to Step 1.
-

Remark: Whereas in Step 2 we have $L \leftarrow 2L$, Nesterov has $L \leftarrow \gamma_u L$, where $\gamma_u > 1$ is user-chosen. Similarly, in Step 3 he has $L \leftarrow L/\gamma_d$. He assumes $\gamma_d \geq \gamma_u$.

Theorem C.1. (Nesterov [22]) *Let $k' := \max\{0, \lceil \log_2(L'/L_f) \rceil\}$. Then*

$$k > k' \quad \Rightarrow \quad f(x_k) - f^* \leq 2L_f \left(\frac{\|x_0 - x^*\|}{k - k'} \right)^2. \quad (\text{C.1})$$

Moreover, for $k > k'$, the total number of gradients computed before x_k is determined does not exceed

$$4(k - k') + 2 \lceil \log_2(L'/L_f) \rceil + 6. \quad (\text{C.2})$$

In establishing the theorem, we rely on Nesterov's notation, as given in the following rendition of the algorithm.

(Notationally-embellished version of the accelerated method)

- (0) Input: $x_0 \in \mathcal{E}$ and $L' > 0$, a guessed estimate of L_f .
 Initialize: $v_0 = x_0$, $L_0 = L'$, $L = L_0$, $A_0 = 0$, $k = 0$.
- (1) Compute $a = \frac{1+\sqrt{1+2A_k L}}{L}$, $y = \frac{A_k}{A_k+a}x_k + \frac{a}{A_k+a}v_k$ and $x = y - \frac{1}{L}\nabla f(y)$.
- (2) If $\langle \nabla f(x), y - x \rangle < \frac{1}{L}\|\nabla f(x)\|^2$ then let $L \leftarrow 2L$ and return to Step 1.
- (3) Define $y_k = y$, $M_k = L$,
 $x_{k+1} = x$, $v_{k+1} = v_k - a\nabla f(x)$, $L_{k+1} = L/2$, $a_{k+1} = a$, $A_{k+1} = A_k + a$,
 and let $L \leftarrow L/2$, $k \leftarrow k + 1$, then go to Step 1.
-

Inductively define the convex functions ψ_k by $\psi_0(x) := \frac{1}{2}\|x - x_0\|^2$ and

$$\psi_{k+1}(x) := \psi_k(x) + a_{k+1}(f(x_{k+1}) + \langle \nabla f(x_{k+1}), x - x_{k+1} \rangle).$$

It is easy to verify inductively that v_k is the minimizer of ψ_k .

Note that $x \mapsto \frac{1}{A_{k+1}}\psi_{k+1}(x)$ can be thought of as an approximation to f , in that

$$\frac{1}{A_{k+1}}\psi_{k+1}(x) = \frac{1}{2A_{k+1}}\|x - x_0\|^2 + \frac{1}{\sum_{j=1}^{k+1} a_j} \sum_{j=1}^{k+1} a_j (f(x_j) + \langle \nabla f(x_j), x - x_j \rangle),$$

a quadratic function plus a convex combination of linear approximations to f (the weight $1/A_{k+1}$ on the quadratic term, as will be seen, goes to zero as $k \rightarrow \infty$).

The proof of the implication (C.1) is accomplished by inductively establishing the following inequalities and implication:

$$\psi_k(x) \leq A_k f(x) + \frac{1}{2} \|x - x_0\|^2 \quad \text{for all } x \in \mathcal{E}, \quad (\text{C.3})$$

$$A_k f(x_k) \leq \psi_k^*, \quad (\text{C.4})$$

$$k > k' \quad \Rightarrow \quad A_k \geq \frac{1}{4}(k - k')^2 / L_f, \quad (\text{C.5})$$

where ψ_k^* is the optimal value of ψ_k , and where k' as in the theorem, that is, $k' = \max\{0, \lceil \log_2(L'/L_f) \rceil\}$. Clearly, (C.3) and (C.4) imply $A_k f(x_k) \leq A_k f^* + \frac{1}{2} \|x^* - x_0\|^2$, which combined with (C.5) gives (C.1).

Proposition C.2. *Both (C.3) and (C.4) are valid.*

Proof: To establish (C.3), first note the inequality trivially holds for $k = 0$. Assuming the inequality holds for k , we have for all x ,

$$\begin{aligned} \psi_{k+1}(x) &= \psi_k(x) + a_{k+1}(f(x_{k+1}) + \langle \nabla f(x_{k+1}), x - x_{k+1} \rangle) \\ &\leq \psi_k(x) + a_{k+1}f(x) \quad (\text{convexity}) \\ &\leq (A_k + a_{k+1})f(x) + \frac{1}{2} \|x - x_0\|^2 \\ &= A_{k+1}f(x) + \frac{1}{2} \|x - x_0\|^2. \end{aligned}$$

Now we establish the inequality (C.4), which trivially holds as equality when $k = 0$. Assume the inequality holds for k . To ease notation, remove the subscript “ k ” and use “ $+$ ” in place of “ $k + 1$.” (Thus, for example, ψ_k becomes ψ , and ψ_{k+1} becomes ψ_+).

Since v_+ minimizes ψ_+ ,

$$\psi_+^* = \psi(v_+) + a_+(f(x_+) + \langle \nabla f(x_+), v_+ - x_+ \rangle). \quad (\text{C.6})$$

However, since ψ is the sum of a linear function and the quadratic $x \mapsto \frac{1}{2} \|x - x_0\|^2$, and since v minimizes ψ ,

$$\begin{aligned} \psi(v_+) &= \psi^* + \frac{1}{2} \|v_+ - v\|^2 \\ &\geq A f(x) + \frac{1}{2} \|v_+ - v\|^2 \quad (\text{inductive hypothesis}) \\ &\geq A(f(x_+) + \langle \nabla f(x_+), x - x_+ \rangle) + \frac{1}{2} \|v_+ - v\|^2 \quad (\text{convexity}), \end{aligned}$$

and hence (C.6) implies (using $A + a_+ = A_+$)

$$\begin{aligned} \psi_+^* &\geq A_+ f(x_+) + \frac{1}{2} \|v_+ - v\|^2 \\ &\quad + \langle \nabla f(x_+), A(x - x_+) + a_+(v_+ - x_+) \rangle \\ &= A_+ f(x_+) \\ &\quad + A_+ \langle \nabla f(x_+), y - x_+ \rangle - \frac{1}{2} a_+^2 \|\nabla f(x_+)\|^2 \\ &\geq A_+ f(x_+), \end{aligned}$$

where the equality is obtained by substituting $v_+ = v - a_+ \nabla f(x_+)$ and then

$$A(x - x_+) + a_+(v - x_+) = Ax + a_+v - A_+x_+ = A_+(y - x_+),$$

and where the final inequality comes from the condition in Step 2 allowing passage to Step 3:

$$\langle \nabla f(x_+), y - x_+ \rangle \geq \frac{1}{M_+} \|\nabla f(x_+)\|^2 = \frac{a_+^2}{2A_+} \|\nabla f(x_+)\|^2.$$

The proof is complete. \square

To establish the implication (C.1), it remains to establish (C.5). This is done using [22, Lemma 5], which shows that if $L \geq L_f$, then the algorithm moves immediately from Step 2 to Step 3. Here is a proof of the lemma in our simple setting.

Lemma C.3. *Let $L > 0$, $y \in \mathcal{E}$ and $x = y - \frac{1}{L}\nabla f(y)$. Then*

$$L \geq L_f \quad \Rightarrow \quad \langle \nabla f(x), y - x \rangle \geq \frac{1}{L} \|\nabla f(x)\|^2 .$$

Proof: The proof makes use of the fact that for all $x, y \in \mathcal{E}$ and $L \geq L_f$,

$$\langle \nabla f(y) - \nabla f(x), y - x \rangle \geq \frac{1}{L} \|\nabla f(x) - \nabla f(y)\|^2$$

(see Theorem 2.1.5 in Nesterov's introductory book [19]). Substituting $y - x = \frac{1}{L}\nabla f(y)$, multiplying both sides by L and then rearranging, we have

$$\langle \nabla f(x), \nabla f(y) \rangle \geq \|\nabla f(x)\|^2 .$$

Substituting $\nabla f(y) = L(y - x)$ completes the proof. \square

Recall $k' := \max\{0, \lceil \log_2(L'/L_f) \rceil\} = \max\{0, \lceil \log_2(L_0/L_f) \rceil\}$.

Lemma C.4. *The first index k satisfying $L_k \leq L_f$ is $k = k'$. Moreover, for all $k > k'$, both $L_k \leq L_f$ and $M_k \leq 2L_f$.*

Proof: Lemma C.3 implies that for $L_k \geq L_f$, the algorithm goes immediately from Step 2 to Step 3, that is, L is not replaced by $2L$ but instead L_{k+1} is defined as $L_k/2$. It follows that k' is the first index k for which $L_k \leq L_f$.

Similar simple (inductive) arguments using Lemma C.3 complete the proof. \square

Now we can establish (C.5), and thus complete the proof of the implication (C.1) in Theorem C.1.

Lemma C.5. *The implication (C.5) holds, that is,*

$$k > k' \quad \Rightarrow \quad A_k \geq \frac{(k - k')^2}{4L_f} .$$

Proof: In the proof of [22, Lemma 8], Nesterov easily shows that regardless of the initial value L_0 ,

$$1 \leq 2M_k \left(A_{k+1}^{1/2} - A_k^{1/2} \right)^2 .$$

Since $M_k \leq 2L_f$ whenever $k > k'$ (Lemma C.4), and since $A_{k+1} > A_k$, we thus have

$$k > k' \quad \Rightarrow \quad A_{k+1}^{1/2} \geq A_k^{1/2} + \frac{1}{2\sqrt{L_f}} ,$$

from which the lemma follows by induction. \square

To complete the proof of Theorem C.1, it remains to prove the bound (C.2) on the number of gradient evaluations. Following Nesterov's choice of notation, let N_k denote the total number of gradients computed before x_{k+1} becomes defined. The following lemma establishes the desired bound.

Lemma C.6.

$$k \geq k' \quad \Rightarrow \quad N_k < 4(k - k') + 2|\log_2(L_0/L_f)| + 6$$

Proof: For the case $L_0 \leq L_f$, the proof of [22, Lemma 6] establishes a slightly stronger result (choose $\gamma_u = \gamma_d = 2$ as the parameter values there). Thus, it remains to establish the lemma when $L_0 > L_f$.

In the proof of [22, Lemma 6], Nesterov shows for any initial value $L_0 > 0$,

$$N_k \leq 4(k+1) + 2 \log_2(L_{k+1}/L_0). \quad (\text{C.7})$$

Assume $L_0 > L_f$. Since k' is the first index k for which $L_k \leq L_f$ (Lemma C.4), we know for $k < k'$ that $L_{k+1} = \frac{1}{2}L_k$, and hence,

$$\frac{1}{2}L_f < L_{k'} = \left(\frac{1}{2}\right)^{k'} L_0 \leq L_f. \quad (\text{C.8})$$

Since $L_k/L_{k'}$ is a power of 2 for all k , and since $L_k \leq L_f$ for all $k > k'$ (Lemma C.4), it follows from (C.8) that

$$k \geq k' \quad \Rightarrow \quad L_k \leq L_{k'}. \quad (\text{C.9})$$

From (C.7), (C.8) and (C.9) we find

$$\begin{aligned} k \geq k' \quad \Rightarrow \quad N_k &\leq 4(k+1) + 2 \log_2(L_{k+1}/L_{k'}) + 2 \log_2(L_{k'}/L_0) \\ &\leq 4(k+1) + 0 - 2k' \\ &= 4(k - k' + 1) + 2k' \\ &= 4(k - k' + 1) + 2 \lceil \log_2(L_0/L_f) \rceil, \end{aligned}$$

completing the proof. \square

REFERENCES

- [1] Heinz H Bauschke, Osman Guler, Adrian S Lewis, and Hristo S Sendov. Hyperbolic polynomials and convex analysis. *Canadian Journal of Mathematics*, 53(3):470–488, 2001.
- [2] Amir Beck and Marc Teboulle. Smoothing and first order methods: A unified framework. *SIAM Journal on Optimization*, 22(2):557–580, 2012.
- [3] Stephen Becker, Jérôme Bobin, and Emmanuel J Candès. NESTA: a fast and accurate first-order method for sparse recovery. *SIAM Journal on Imaging Sciences*, 4(1):1–39, 2011.
- [4] Stephen R Becker, Emmanuel J Candès, and Michael C Grant. Templates for convex cone problems with applications to sparse signal recovery. *Mathematical Programming Computation*, 3(3):165–218, 2011.
- [5] Grigoriy Blekherman. There are significantly more nonnegative polynomials than sums of squares. *Israel Journal of Mathematics*, 153(1):355–380, 2006.
- [6] Simon Foucart and Holger Rauhut. *A mathematical introduction to compressive sensing*. Springer, 2013.
- [7] 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.
- [8] Lars Garding. An inequality for hyperbolic polynomials. *Journal of Mathematics and Mechanics*, 8(6):957–965, 1959.
- [9] Osman Güler. Hyperbolic polynomials and interior point methods for convex programming. *Mathematics of Operations Research*, 22(2):350–377, 1997.
- [10] Christoph Hanselka. Definite determinantal representations of ternary hyperbolic forms. *arXiv preprint arXiv:1411.1661*, 2014.
- [11] J William Helton and Victor Vinnikov. Linear matrix inequality representation of sets. *Communications on Pure and Applied Mathematics*, 60(5):654–674, 2007.
- [12] Oliver Dimon Kellogg. On bounded polynomials in several variables. *Mathematische Zeitschrift*, 27(1):55–64, 1928.
- [13] Guanghui Lan. Bundle-level type methods uniformly optimal for smooth and nonsmooth convex optimization. *Mathematical Programming*, 149(1-2):1–45, 2015.
- [14] Guanghui Lan, Zhaosong Lu, and Renato DC Monteiro. Primal-dual first-order methods with $O(1/\epsilon)$ iteration-complexity for cone programming. *Mathematical Programming*, 126(1):1–29, 2011.
- [15] Adrian Lewis, Pablo Parrilo, and Motakuri Ramana. The Lax conjecture is true. *Proceedings of the American Mathematical Society*, 133(9):2495–2499, 2005.
- [16] Zhaosong Lu, Arkadi Nemirovski, and Renato DC Monteiro. Large-scale semidefinite programming via a saddle point mirror-prox algorithm. *Mathematical Programming*, 109(2-3):211–237, 2007.
- [17] Arkadi Nemirovski and David Yudin. Problem complexity and method efficiency in optimization. 1983.
- [18] Yu Nesterov. Universal gradient methods for convex optimization problems. *Mathematical Programming*, pages 1–24, 2014.
- [19] Yurii Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Springer, 2004.

- [20] Yurii Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, 2005.
- [21] Yurii Nesterov. Smoothing technique and its applications in semidefinite optimization. *Mathematical Programming*, 110(2):245–259, 2007.
- [22] Yurii Nesterov. Gradient methods for minimizing composite functions. *Mathematical Programming*, 140(1):125–161, 2013.
- [23] Yurii Nesterov and Arkadi Nemirovski. *Interior-point polynomial algorithms in convex programming*. SIAM, 1994.
- [24] James Renegar. Hyperbolic programs, and their derivative relaxations. *Foundations of Computational Mathematics*, 6(1):59–79, 2006.
- [25] James Renegar. Efficient first-order methods for linear programming and semidefinite programming. *arXiv preprint arXiv:1409.5832*, 2014.
- [26] James Renegar. A framework for applying subgradient methods to conic optimization problems (version 2). *arXiv preprint arXiv:1503.02611*, 2015.
- [27] James Renegar and Mutiara Sondjaja. A polynomial-time affine-scaling method for semidefinite and hyperbolic programming. *arXiv preprint arXiv:1410.6734*, 2014.
- [28] Victor Vinnikov. LMI representations of convex semialgebraic sets and determinantal representations of algebraic hypersurfaces: past, present, and future. In *Mathematical methods in systems, optimization, and control*, pages 325–349. Springer, 2012.

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