

Generalized self-concordant Hessian-barrier algorithms

Pavel Dvurechensky¹, Mathias Staudigl², Cesar A. Uribe³

submitted: February 19, 2020

¹ Weierstrass Institute
Mohrenstr. 39
10117 Berlin
Germany

E-Mail: pavel.dvurechensky@wias-berlin.de

² Maastricht University
Department of Quantitative Economics
P.O. Box 616,
NL-6200 MD Maastricht, The Netherlands
E-Mail: m.staudigl@maastrichtuniversity.nl

³ Laboratory for Information and Decision Systems (LIDS)
Institute for Data, Systems, and Society (IDSS)
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139
USA
E-Mail: cauribe@mit.edu

No. 2693

Berlin 2020



Edited by
Weierstraß-Institut für Angewandte Analysis und Stochastik (WIAS)
Leibniz-Institut im Forschungsverbund Berlin e. V.
Mohrenstraße 39
10117 Berlin
Germany

Fax: +49 30 20372-303
E-Mail: preprint@wias-berlin.de
World Wide Web: <http://www.wias-berlin.de/>

Generalized self-concordant Hessian-barrier algorithms

Pavel Dvurechensky, Mathias Staudigl, Cesar A. Uribe

Abstract

Many problems in statistical learning, imaging, and computer vision involve the optimization of a non-convex objective function with singularities at the boundary of the feasible set. For such challenging instances, we develop a new interior-point technique building on the Hessian-barrier algorithm recently introduced in Bomze, Mertikopoulos, Schachinger and Staudigl, [SIAM J. Opt. 2019 29(3), pp. 2100-2127], where the Riemannian metric is induced by a generalized self-concordant function. This class of functions is sufficiently general to include most of the commonly used barrier functions in the literature of interior point methods. We prove global convergence to an approximate stationary point of the method, and in cases where the feasible set admits an easily computable self-concordant barrier, we verify worst-case optimal iteration complexity of the method. Applications in non-convex statistical estimation and L^p -minimization are discussed to given the efficiency of the method.

1 Introduction

In this paper, we consider the following constrained minimization problem, which has plenty of applications in diverse disciplines, including machine learning, signal processing, statistics, and operations research

$$f^* = \min\{f(x) : x \in \bar{C}, Ax = b\} \quad (\text{P})$$

Here $\bar{C} \subset \mathbb{R}^n$ is a nonempty, closed and convex set, and f is a (possibly) non-smooth, non-convex function from $\mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$. A special case of (P) are regularized statistical estimation problems where the aim is to find a parameter vector $x \in \mathbb{R}^n$ in order to minimize a composite objective of the form

$$f(x) := f_0(x) + f_1(c(x)). \quad (1)$$

In such applications the function $f_0 : \mathbb{R}^n \rightarrow \mathbb{R}$ is a continuous data fidelity term, $f_1 : \mathbb{R}^d \rightarrow (-\infty, \infty]$ is a regularizer and $c : \mathbb{R}^n \rightarrow \mathbb{R}^d$ is some link function mapping the parameters to a usually lower dimensional subspace. Typical formulations of such problems can be given in the form of

$$f(x) = \frac{1}{2} \|s - Wx\|_2^2 + \sum_{i=1}^m \varphi_i(\|D_i x\|^p) \quad (2)$$

in which $f_0(x) = \frac{1}{2} \|s - Wx\|_2^2$, $f_1(y) = \sum_{i=1}^m \varphi_i(y_i) : \mathbb{R}^q \rightarrow [-\infty, \infty)$, $\varphi_i : \mathbb{R}^{d_i} \rightarrow \mathbb{R}$ is continuously differentiable, and $c(x) = (\|D_1 x\|^p, \dots, \|D_m x\|^p)^\top$ for some $p \in (0, 1)$, $D_i \in \mathbb{R}^{d_i \times n}$, $d = \sum_{i=1}^m d_i$. In fact, the formulation (2) includes many well-known problems in statistics: fused lasso [48], grouping pursuit [45], etc. From an optimization perspectives, these regularized estimation problems are challenging since they are non-convex and not globally Lipschitz continuous and thus belong to the class of NP-hard problems. Even worse, they may even fail to be differentiable. As an example, the L^p regularization problem with link function $c(x) = (|x_1|^p, \dots, |x_n|^p)$ for $0 < p < 1$ and $f_1(y) = \sum_{i=1}^n y_i$ fails to be even directionally differentiable when $x_i = 0$ for some $i = 1, 2, \dots, n$.

A common tenet of all recent applications of (P) is that the problem involves a huge number of variables. This makes the application of classical interior-point solvers infeasible. Instead, first-order methods (FOMs) with cheap per-iteration implementation costs are the method of choice [47]. The most impactful success stories of FOMs have been achieved under the quite demanding assumption that the objective function is convex and smooth, and the feasible set admits a proximal-friendly formulation. Indeed, if proximal-based projection operators onto the feasible set $\mathcal{X} := \{x \in \bar{\mathcal{C}} : Ax = b\}$ are easy to evaluate, black-box based FOMs such as mirror descent, projected subgradient, and conditional gradient methods can be tuned to successfully solve (P) up to ε -accuracy. How to handle non-convex objective functions with first-order methods is still a challenging problem receiving a lot of interest from various different perspectives, in particular in statistical and deep learning. Beside this implicit assumption in all projection-based FOMs, another fundamental assumption in all these methods is the availability of a *descent lemma* [39, Lem.1.2.3]. Sufficient for such an a-priori estimate is that the objective function's gradient is a Lipschitz continuous function. Already the above mentioned application to statistical estimation shows that this Lipschitz-smoothness assumption is too demanding to cover such important applications. Only recently, the path-breaking paper [4] resolved this problem by introducing the concept of *relative smoothness* as a surrogate for the demanding Lipschitz gradient assumption (see also [35] for elaborations and applications). Based on relative-smoothness, they derive a new descent lemma where the usual quadratic approximation is replaced by a non-Euclidean proximity measure, which captures the objective function and the geometry of the underlying domain all at once. The corresponding proximal-based subgradient algorithm comes with global convergence guarantees and complexity estimates for *convex composite* models. This beautiful, and practically relevant, approach has been recently extended to a non-convex composite model in [8], where new complexity estimates are derived as well.

1.1 Our Approach

This work is concerned with a different approach to tackle non-convex optimization problems avoiding knowledge of a global Lipschitz constant. Our work is inspired by the recent Riemannian gradient methods developed in [11], where a rather large class of *Hessian barrier algorithms* (HBA) has been constructed as numerical schemes for solving Lipschitz-smooth, non-convex optimization problems over the polyhedron $\{x \in \mathbb{R}^n : x \geq 0, Ax = b\}$. The construction of HBA is motivated by looking at an *explicit* numerical discretization of a continuous-time dynamical system introduced by [2, 9] as a theoretical method for solving *convex* linearly constrained smooth optimization problems. However, these schemes remained at a conceptual level, and the usefulness of these dynamical systems for effectively solving constrained optimization of the form (P) remained completely unanswered. HBA laid the foundations to an algorithmic analysis for these dynamical systems, and investigated their efficiency when solving linearly constrained and smooth non-convex optimization problems. HBA first identifies the feasible set as a Riemannian manifold with a metric induced by the Hessian matrix of a \mathcal{C}^2 barrier-like function h (a *barrier-generating kernel*). Once the geometry has been defined, a step-size strategy is designed ensuring feasibility and a sufficient decrease of the objective. The analysis in [11] relied heavily on the classical Lipschitz-descent lemma [39], and involved an Armijo line-search procedure. It has been shown that this approach generalizes many classical interior point techniques like affine scaling [5], and also contains Lotka-volterra systems as a special case [51]. A complexity analysis was achieved in the case where the objective function is quadratic, and it has been shown that a proper choice of the Riemannian metric affects the complexity of the method [11, Thm 5.1]. In this paper we significantly extend the results obtained in [11] by constructing a new first-order interior point method for solving problem (P) under very mild assumptions on the data. In order to explain

the approach described in this paper, let us go back to the classical way of solving problem (P) when f is *convex*. The most famous algorithm for solving such problems are interior point methods, which solve conic optimization problems in polynomial time [43]. The key structure exploited in conventional IPMs is the existence of a *self-concordant barrier* (SCB) for the set constraint $\bar{\mathcal{C}}$. In such cases one considers the potential function

$$F_\mu(x) = f(x) + \mu h(x),$$

where $\mu > 0$ is a penalty parameter and h is a barrier function over the set $\bar{\mathcal{C}}$. By fixing a sequence of barrier parameters $(\mu_k)_{k \geq 0}$ with $\mu_k \downarrow 0$ and solving the sequence of minimization problems $\min_x F_{\mu_k}(x)$ along this sequence generates the *analytic central path* $\{x_\mu^* : \mu > 0\}$ as it converges to the solution of the actual problem of interest (P). For proving convergence of the analytic central path, SCBs are the key tool to prove polynomial-solvability of the barrier problem by sequentially using Newton's method.¹ The new Hessian-barrier method we propose in this paper follows similar ideas. We first embed the original optimization problem (P) into a potential-reduction scheme involving the potential function F_μ . However, instead of classical self-concordance theory our analysis works easily on a much broader class of *generalized self-concordant functions* (GSC), which have recently been introduced in [46]. As we show in this paper, GSC functions provide a very attractive class of barrier-generating kernels as their Hessian matrix induce a Riemannian metric under which a full-fledged complexity analysis can be performed.

1.2 Available complexity results

We now review results on complexity analysis of non-convex optimization problems. Since the Hessian-barrier method uses only first-order information about f , we restrict our review to first-order methods as well. In unconstrained non-convex optimization the usual criticality measure is the norm of the gradient of the objective function. Hence, oracle complexity of a given algorithm refers to the number of oracle queries until $\|\nabla f(x^k)\| \leq \varepsilon$ for a targeted tolerance level $\varepsilon > 0$. An in-depth survey of known complexity bounds can be found in [16, 17]. We can give no justice to the huge literature on complexity results for first-order methods but provide below a partial survey of known complexity estimates in order to put our results into perspective.

Smooth, non-convex For quadratic programming problems with linear constraints, the authors in [55] proved that an ε -KKT point is computed in $\mathcal{O}(\varepsilon^{-1} \log(\varepsilon^{-1}))$ iterations. A recent manifestation of the effect of Riemannian geometry on the convergence to ε -KKT points can be found in [11]. For general unconstrained nonconvex optimization, it was shown in [39] that a steepest descent with line search method can find an ε -stationary point in $\mathcal{O}(\varepsilon^{-2})$ iterations. An accelerated method with the same guarantee can be found in [27]. The same worst case complexity result holds for trust-region methods [26]. The results for gradient method were generalized to the case of simple projection-friendly constraints and Hölder derivatives in [25]. Accelerated methods with complexity $\tilde{\mathcal{O}}(\varepsilon^{-7/4})$ under additional assumption of Lipschitz second derivative are proposed in [1, 14, 15, 29].

Lipschitz continuous, nonconvex Cartis, Gould and Toint [18] estimated the worst-case complexity of a first-order trust-region or quadratic regularization method for solving unconstrained, non-convex minimization problems of the form (1), where $f_1 : \mathbb{R}^m \rightarrow \mathbb{R}$ is convex but may be nonsmooth and

¹Recently, the path-following method was extended for self-concordant functions which are not self-concordant barriers in [21].

$c : \mathbb{R}^n \rightarrow \mathbb{R}^d$ is continuously differentiable. Their method takes at most $\mathcal{O}(\varepsilon^{-2})$ iterations to reduce the size of a suitably defined first-order criticality measure below ε .

Non-Lipschitz, nonconvex The authors in [24] extended the complexity result of [55] to the L^p -minimization problem over a polytop. They showed that finding an ε -scaled stationary point or global minimizer requires at most $\mathcal{O}(\varepsilon^{-1} \log(\varepsilon^{-1}))$ iterations. For general linear constrained non-convex minimization problems [28] obtained an iteration complexity of $\mathcal{O}(\varepsilon^{-2})$ to reach an ε -KKT point for optimization problems whose feasible set is defined by linear equality and non-negativity constraints. In the case of L^p minimization for $p \in (0, 1)$ over box constraints [6] develop a first-order interior point method yielding $\mathcal{O}(\varepsilon^{-2})$ worst-case iteration complexity in order to return an ε -scaled stationary point.

1.3 Our Contribution

In relation to the above summarized literature, we provide here an easy-to-implement first-order method for non-convex non-smooth optimization problems, without requiring knowledge on the Lipschitz constant of the objective function. Specifically, the main contributions of this paper are summarized as follows:

- 1 We provide a new first-order interior point method based on the HBA method for non-convex and non-smooth optimization problems (P) without Lipschitz smoothness conditions.
- 2 We are the first who provide a first-order interior point analysis based on GSC functions.
- 3 We show how some model parameters can be made adaptive, making the method even more efficient in practice.
- 4 Our method comes with an explicit construction of optimal step-size policies and convergence guarantees.
- 5 We demonstrate optimal iteration complexity estimates on the order of $\mathcal{O}(\varepsilon^{-2})$ to reach a generalized stationary point, and connect this to classical ε -KKT conditions in case where the barrier-generating kernel is a SC-B. This answers an open question raised in Remark 4.1. in [9], since HBA can be seen as a descendent of the *A-driven descent methods* defined in that paper. Also, in view of the partial literature survey given above, this rate is optimal.

The rest of this paper is organized as follows: Section 2 introduces the standing assumptions used in this paper, and introduces the class of generalized self-concordant functions. Section 3 defines conceptually the HBA method and introduces the optimal step-size policy associated with it. An adaptive variant of this base scheme is discussed in Section 4.2. Section 5 includes the main results in terms of convergence and complexity of the method. Section 6 contains numerical examples.

Notation Given a k -times continuously differentiable function $f : \mathcal{C} \rightarrow \mathbb{R}$ and vectors $v_1, \dots, v_k \in \mathbb{R}^n$. For $x \in \text{ri}(\mathcal{C})$ and $1 \leq j \leq k$, we define recursively

$$D^j f(x)[v_1, \dots, v_j] := \lim_{\varepsilon \rightarrow 0} \frac{D^{j-1} f(x + \varepsilon v_j)[v_1, \dots, v_{j-1}] - D^{j-1} f(x)[v_1, \dots, v_{j-1}]}{\varepsilon}.$$

As a convention $D^0 f(x) = f(x)$ and for $k = 1$ we recover the directional derivative $f'(x; v)$. Given a positive semi-definite matrix $H \in \mathbb{R}^{n \times n}$, we define the norm $\|a\|_H := \sqrt{\langle Ha, a \rangle}$, and the dual norm $\|a\|_H^* = \sup\{\langle a, d \rangle : \|d\|_H = 1\}$. If H is invertible, the dual norm is given by $\|a\|_H^* = \sqrt{\langle H^{-1}a, a \rangle}$. For a given $n \times n$ matrix A , let us define the operator norm $|A| := \sup\{\|Ax\| : \|x\| = 1\}$. Let $\mathcal{C} \subset \mathbb{R}^n$ be a convex set with closure $\bar{\mathcal{C}}$. Define the tangent cone $\text{TC}_{\mathcal{C}}(x) := \text{cl}[\mathbb{R}_+(\mathcal{C} - x)]$, and the corresponding polar cone $\text{NC}_{\bar{\mathcal{C}}}(x) = \{p \in \mathbb{R}^n : \langle p, y - x \rangle \leq 0, \forall y \in \bar{\mathcal{C}}\}$, which is called the normal cone.

2 Setup and preliminaries

2.1 Elements of Riemannian geometry

A key notion in our considerations is that of a *Riemannian metric*, i.e. a position-dependent variant of the ordinary (Euclidean) scalar product between vectors [30]. To define it, recall first that a *scalar product* on \mathbb{R}^n is a symmetric, positive-definite bilinear form $\langle \cdot, \cdot \rangle : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$. This scalar product defines a norm in the usual way and it can be represented equivalently via its *metric tensor*, that is, a symmetric, positive-definite matrix $H \in \mathbb{R}^{n \times n}$ with components

$$H_{ij} = \langle e_i, e_j \rangle \quad (3)$$

in the standard basis $\{e_i\}_{i=1}^n$ of \mathbb{R}^n . A *Riemannian metric* on a nonempty open set $\mathcal{C} \subseteq \mathbb{R}^n$ is then defined to be a smooth assignment of scalar products $\langle \cdot, \cdot \rangle_x$ to each $x \in \mathcal{C}$ or, equivalently, a smooth field $H(x)$ of symmetric positive-definite matrices on \mathcal{C} .

Given a Riemannian metric on \mathcal{C} , the *Riemannian gradient* of a smooth function $\phi : \mathcal{C} \rightarrow \mathbb{R}$ at $x \in \mathcal{C}$ is defined via the characterization

$$\langle \text{grad } \phi(x), z \rangle_x = \phi'(x; z) \quad \text{for all } z \in \mathbb{R}^n. \quad (4)$$

More concretely, by expressing everything in components, it is easy to see that $\text{grad } \phi(x)$ is given by the explicit expression

$$\text{grad } \phi(x) = [H(x)]^{-1} \nabla \phi(x). \quad (5)$$

Bringing the above closer to our setting, let $\mathcal{A}_0 \subseteq \mathbb{R}^n$ be a subspace of \mathbb{R}^n and let \mathcal{A} be an affine translate of \mathcal{A}_0 such that $\mathcal{X}^\circ \equiv \mathcal{C} \cap \mathcal{A}$ is nonempty. Then, viewing \mathcal{X}° as an open subset of \mathcal{A} , the *gradient of ϕ restricted to \mathcal{X}°* is defined as the unique vector $\text{grad}_{\mathcal{X}^\circ} \phi(x) \equiv \text{grad } \phi|_{\mathcal{X}^\circ}(x) \in \mathcal{A}_0$ such that

$$\text{grad}_{\mathcal{X}^\circ} \phi(x) = \phi'(x, d) \quad \text{for all } d \in \mathcal{A}_0. \quad (6)$$

Hence, specializing all this to the problem at hand, let $H(x)$ be a Riemannian metric on the open set $\mathcal{C} \subset \mathbb{R}^n$ and set

$$\mathcal{A}_0 := \ker(A) := \{d \in \mathbb{R}^n : Ad = 0\}, \quad \mathcal{A} := \{x \in \mathbb{R}^n : Ax = b\}.$$

Then, a straightforward exercise in matrix algebra shows that the gradient of f restricted to $\mathcal{X}^\circ = \mathcal{C} \cap \mathcal{A}_0$ can be written in closed form as

$$\text{grad}_{\mathcal{X}^\circ} f(x) = P_x [H(x)]^{-1} \nabla f(x) \quad (7)$$

with

$$P_x := \text{Id} - [H(x)]^{-1} A^\top (A [H(x)]^{-1} A^\top)^{-1} A. \quad (8)$$

2.2 Generalized self-concordant functions

The Hessian-barrier method's main assumption is that the set constraint $\bar{\mathcal{C}}$ admits an easy-to-compute *generalized self-concordant function*. In the next sections we describe the subclass of admissible functions on which the subsequent constructions build on. We begin with the notion of *kernel generating distance*, by following the very general setup introduced in [3].

Definition 2.1 (Kernel generating distance). Let \mathcal{C} be a nonempty convex and open subset of \mathbb{R}^n . Associated with \mathcal{C} , a function $h : \mathcal{C} \rightarrow (-\infty, \infty]$ is called a *kernel generating distance* if

- (a) h is proper, lower semi-continuous and convex, with $\text{dom } h \subset \bar{\mathcal{C}}$ and $\text{dom } \partial h = \text{ri dom } h = \mathcal{C}$.
- (b) $h \in \mathbf{C}^1(\mathcal{C})$.

Denote the class of kernel generating distances by $\mathcal{G}(\mathcal{C})$.

We now add the additional structure on the kernel generating distance $h \in \mathcal{G}(\mathcal{C})$ we use in our algorithmic design.

Definition 2.2. [46] Let $\phi \in \mathbf{C}^3(\text{dom } \phi; \mathbb{R})$ be a closed convex function with $\text{dom } \phi$ open. Given $\nu > 0$ and $M > 0$ some constants, we say that ϕ is (M, ν) *generalized self-concordant* (GSC) if

$$|\phi'''(t)| \leq M\phi''(t)^{\frac{\nu}{2}} \quad \forall t \in \text{dom } \phi. \quad (9)$$

This definition generalizes to multivariate functions by requiring GSC along every straight line. Specifically, let $h : \mathbb{R}^n \rightarrow \mathbb{R}$ be a closed convex, lower semi-continuous function with open and convex effective domain $\text{dom } h = \mathcal{C} \subset \mathbb{R}^n$. For $x \in \mathcal{C}$ and $u, v \in \mathbb{R}^n$, define the real-valued function $\phi(t) := \langle \nabla^2 h(x + tv)u, u \rangle$. For $t \in \text{dom } \phi$, one sees that

$$\phi'(t) = \langle D^3 h(x + tv)[v]u, u \rangle,$$

so that we can define generalized self-concordance of a function by formulating conditions on the behavior of $\phi'(0)$.

Definition 2.3. [46] A closed convex function $h \in \mathbf{C}^3(\text{dom } h)$, with $\text{dom } h$ open, is called (M, ν) *generalized self-concordant of the order* $\nu > 0$ and constant $M \geq 0$ if for all $x \in \text{dom } h$

$$|\langle D^3 h(x)[v]u, u \rangle| \leq M \|u\|_x^2 \|v\|_x^{\nu-2} \|v\|_2^{3-\nu} \quad \forall u, v \in \mathbb{R}^n. \quad (10)$$

We denote this class of functions as $\mathcal{H}_{M,\nu}(\text{dom } h)$.

As in the theory of standard self-concordant functions, the precise value of the scale parameter $M > 0$ is not of big importance for theoretical considerations. In fact, it is easy to see that we can always rescale the function so that the definition of a GSC function holds for $M = 2$.

Lemma 2.4. *If $\phi \in \mathbf{C}^3(\text{dom } \phi; \mathbb{R})$ belongs to the class $\mathcal{H}_{M,\nu}(\text{dom } h)$, then $(\frac{M}{2})^{\frac{2}{\nu-2}} \phi \in \mathcal{H}_{2,\nu}(\text{dom } h)$.*

Proof. Let $\psi(t) := (\frac{M}{2})^{\frac{2}{\nu-2}} \phi(t)$. Then, for all $t \in \text{dom } \phi$, we have

$$|\psi'''(t)| = \left(\frac{M}{2}\right)^{\frac{2}{\nu-2}} |\phi'''(t)| \leq M \left(\frac{M}{2}\right)^{\frac{2}{\nu-2}} \phi''(t)^{\nu/2} = 2\psi''(t)^{\nu/2}.$$

■

The function $h \in \mathcal{H}_{M,\nu}(\text{dom } h)$ defines a semi-norm

$$\|d\|_x := \sqrt{\langle \nabla^2 h(x) d, d \rangle}, \quad (11)$$

with dual norm

$$\|d\|_x^* := \sup_{d \in \mathbb{R}^n} \{2\langle d, a \rangle - \|d\|_x^2\}. \quad (12)$$

Note that if $H(x) \equiv \nabla^2 h(x) \succ 0$ then $\|\cdot\|_x$ is a real norm (necessarily equivalent to the euclidean norm), and $\|d\|_x^* = \sqrt{\langle [H(x)]^{-1} d, d \rangle}$. The barrier-character of functions $h \in \mathcal{H}_{M,\nu}(\text{dom } h)$ is made clear in the next Lemma.

Lemma 2.5 ([39], Thm. 4.1.4). *For every sequence $(x^k)_{k \geq 0}$ such that $(x^k)_{k \geq 0} \subset \text{dom } h$ and $x^k \rightarrow x \in \text{bd}(\text{dom } h)$, we have $\lim_{k \rightarrow \infty} h(x^k) = \infty$.*

Proof. For all $k \geq 0$, we have

$$h(x^k) \geq h(x^0) + \langle \nabla h(x^0), x^k - x^0 \rangle.$$

If the sequence $(h(x^k))_{k \geq 0}$ is bounded from above, we can descent to a subsequence along which $h(x^k) \rightarrow \bar{h}$ (we omit the relabeling). Then, for all $k \geq 0$, $z^k = (x^k, h(x^k)) \in \text{epi}(h)$, and $z^k \rightarrow z = (x, \bar{h}) \in \text{epi}(h)$, since the function is closed. Hence, $x \in \text{dom } h$. A contradiction. ■

Given $\nu \in (2, 4]$ and $h \in \mathcal{H}_{M,\nu}(\mathcal{C})$, we define the distance function

$$d_\nu(x, y) := \begin{cases} M \|y - x\|_2 & \text{if } \nu = 2, \\ \frac{\nu-2}{2} M \|y - x\|_2^{3-\nu} \cdot \|y - x\|_x^{\nu-2} & \text{if } \nu > 2. \end{cases} \quad (13)$$

The *Dikin Ellipsoid* with respect to the distance function d_ν is defined as

$$\mathcal{W}(x; r) := \{y \in \mathbb{R}^n : d_\nu(x, y) < r\} \quad \forall (x, r) \in \text{dom } h \times \mathbb{R}. \quad (14)$$

Lemma 2.6 ([46], Prop. 7). *Let $h \in \mathcal{H}_{M,\nu}(\mathcal{C})$ be a barrier-generating kernel of order $\nu \in (2, 4]$. We have $\mathcal{W}(x; 1) \subseteq \text{dom } h$ for all $x \in \text{dom } h$.*

Remark 2.1. The familiar inclusion $\mathcal{W}(x; 1) \subset \text{dom } h$ is only true if $\nu > 2$. This is intuitive, since for $\nu = 2$, the local norm effectively boils down to the euclidean norm, and thus is not adaptive to the local geometry. As a consequence, our algorithmic scheme will take as inputs functions $h \in \mathcal{H}_{M,\nu}(\mathcal{C})$ with $\nu > 2$ only. This covers the important case of standard self-concordant functions, as well as many entropy-based barrier functions familiar from the literature on Bregman proximal gradient methods. However, our method also works well for generalized self-concordant function of order $\nu \in (3, 4]$. This range cannot be analyzed by proximal based algorithms studied in recent work [46, 49] and [50].

We define the *Bregman divergence* associated to $h \in \mathcal{H}_{M,\nu}(\mathcal{C})$ as

$$D_h(x, y) := h(x) - h(y) - \langle \nabla h(y), x - y \rangle \quad \text{for } x \in \mathcal{C}, y \in \mathcal{C}. \quad (15)$$

Since this divergence function will be a crucial quantity of interest in measuring the per-iteration progress of our method, it is instrumental to have universal bounds on the function values. For the class of self-concordant functions, such bounds are classical to the field (see e.g. [43]). For the Bregman divergence induced by the class of generalized self-concordant functions, a similar universal bound can be reported.

Lemma 2.7 ([46], Prop. 10). *Let $x \in \text{dom } h$ for $h \in \mathcal{H}_{M,\nu}(\mathcal{C})$ and $\nu \in (2, 4]$. Then*

$$\omega_\nu(-d_\nu(x, y)) \|y - x\|_x^2 \leq D_h(y, x) \leq \omega_\nu(d_\nu(x, y)) \|y - x\|_x^2, \quad (16)$$

for all $y \in \mathcal{W}(x; 1)$, where

$$\omega_\nu(t) := \begin{cases} \frac{-t - \ln(1-t)}{t^2} & \text{if } \nu = 3, \\ \frac{(1-t)\ln(1-t) + t}{t^2} & \text{if } \nu = 4, \\ \left(\frac{\nu-2}{4-\nu}\right) \frac{1}{t} \left[\frac{\nu-2}{2(3-\nu)t} ((1-t)^{\frac{2(3-\nu)}{2-\nu}} - 1) - 1 \right] & \text{otherwise.} \end{cases} \quad (17)$$

Let $h \in \mathcal{H}_{M,\nu}(\mathcal{C})$ and $x_j \rightarrow \text{bd}(\mathcal{C}) = \bar{\mathcal{C}} \setminus \mathcal{C}$. Then, by Lemma 2.5, $h(x_j) \rightarrow \infty$. We claim that this implies $\|\nabla h(x_j)\|_2 \rightarrow \infty$. Indeed, by convexity, for all $y \in \mathcal{C}$, we have

$$h(y) \geq h(x_j) + \langle \nabla h(x_j), y - x_j \rangle.$$

Therefore, by Cauchy-Schwarz

$$h(y) \geq h(x_j) - \|y - x_j\| \cdot \|\nabla h(x_j)\|_2$$

If $\|\nabla h(x_j)\|_2$ would be bounded, the right-hand side diverges to ∞ , whereas the left hand side is bounded for $y \in \text{dom}(h)$. This gives a contradiction. We conclude that h is *Legendre*:

$$x_j \rightarrow x^* \in \text{bd}(\mathcal{C}) \Rightarrow \|\nabla h(x_j)\|_2 \rightarrow \infty. \quad (18)$$

Lemma 2.8. *If $h \in \mathcal{H}_{M,\nu}(\text{dom } h)$ and $\text{dom } h \subset \mathbb{R}^n$ contains no lines, then $H(x) \succ 0$ for all $x \in \text{dom } h$.*

Proof. Define the recessive subspace $E_h := \{d \in \mathbb{R}^n : \|d\|_x = 0 \text{ for some } x \in \text{dom } h\}$. From [46, Prop. 8] we deduce that for all $r \in (0, 1)$ and $y \in \mathbb{R}^n$ with $r = d_\nu(x, y)$, we have

$$(1-r)^{\frac{2}{\nu-2}} \nabla^2 h(x) \leq \nabla^2 h(y) \leq (1-r)^{\frac{-2}{\nu-2}} \nabla^2 h(x). \quad (19)$$

Let $Z_d := \{x \in \text{dom } h : \|d\|_x = 0\}$. (19) implies that $x \in Z_d \Rightarrow y \in Z_d$ for all $y \in \mathcal{W}(x; r)$, and therefore Z_d is open. Since $h \in \mathbf{C}^3(\text{dom } h)$, it is closed as well. Therefore Z_d is either empty or the entire set \mathbb{R}^n . This implies that E_h is either empty or \mathbb{R}^n . From here the result follows from [39, Thm. 4.1.3]. \blacksquare

In order to fully understand the behavior of Newton methods involving standard self-concordant function $h \in \mathcal{H}_{2,3}(\text{dom } h)$, the general theory laid out by [43] introduced the concept of a *self-concordant barrier* (SC-B).

Definition 2.9. For some scalar $\theta \geq 0$, a function $h \in \mathcal{H}_{2,3}(\mathcal{C})$ is a self-concordant barrier of order $\theta > 0$ (θ -SCB) if

$$\|\nabla h(x)\|_x^* \leq \sqrt{\theta} \quad \forall x \in \mathcal{C}. \quad (20)$$

It is very remarkable that every convex body, i.e. every compact convex set with nonempty interior, admits a θ -SCB with θ a dimension-dependent constant. This class of functions is the main driver in standard interior-point solvers, and conceptually it reveals the two main problems IPMs face when confronted with large-scale problems: (i) Even if a suitable self-concordant barrier can be computed, the order parameter θ is dimension dependent, and thus iteration complexity cannot be dimension-free. However, Bubeck and Eldan [13] have recently shown that the Nesterov-Nemirovski universal barrier (see below) is a θ -SCB on \mathcal{C} with $\theta = (1 + \varepsilon_n)n$ and $\varepsilon_n \leq 100\sqrt{\log(n)}/n$. (ii) The construction of the *universal barrier* for a convex body $\bar{\mathcal{C}} \subset \mathbb{R}^n$ due to Nesterov and Nemirovski [43] is based on the log-Laplace transform $h^*(w) = \log\left(\int_{\bar{\mathcal{C}}} \exp(\langle x, w \rangle) dx\right)$ and its convex conjugate $h(x) = \sup_{w \in \mathbb{R}^n} \{\langle w, x \rangle - h^*(w)\}$. Unless the set \mathcal{C} is special, computing a universal barrier is infeasible.

2.3 Barrier generating kernels

The class of metric generating functions of interest in this paper is defined as a subset $\mathcal{F}_{M,\nu}(\mathcal{C})$ contained in $\mathcal{G}(\mathcal{C}) \cap \mathcal{H}_{M,\nu}(\mathcal{C})$, whose exact definition depends on the order parameter ν . If $\nu \in (2, 3]$, we take $\mathcal{F}_{M,\nu}(\mathcal{C}) = \mathcal{G}(\mathcal{C}) \cap \mathcal{H}_{M,\nu}(\mathcal{C})$. If $\nu \in (3, 4]$, we additionally assume *coercivity* of the kernel-generating distance h , i.e that $h(x) \rightarrow \infty$ whenever $\|x\| \rightarrow \infty$.

Definition 2.10. The class of *barrier-generating kernels* is defined as

$$\mathcal{F}_{M,\nu}(\mathcal{C}) := \begin{cases} \mathcal{G}_{M,\nu}(\mathcal{C}) \cap \mathcal{H}_{M,\nu}(\mathcal{C}) & \text{if } \nu \in (2, 3], \\ \{h \in \mathcal{G}_{M,\nu}(\mathcal{C}) \cap \mathcal{H}_{M,\nu}(\mathcal{C}) : h \text{ is coercive}\} & \text{if } \nu \in (3, 4]. \end{cases}$$

2.4 Examples

In order to illustrate the flexibility of the framework of sets endowed with barrier-generating kernels, we collect below some representative examples taken from the literature. For many more examples, we refer the reader to [43].

The first set of examples are tailored to product domains $\mathcal{C} = \mathcal{C}_1 \times \cdots \times \mathcal{C}_n$, where each \mathcal{C}_i is an open convex subset of the real-line. For such domains, *decomposable* barrier-generating kernels are an attractive choice since their Hessian matrix is diagonal. Specifically, we consider functions $h \in \mathcal{F}_{M,\nu}(\mathcal{C})$ of the form

$$h(x_1, \dots, x_n) = \sum_{i=1}^n \phi_i(x_i),$$

where each function $\phi_i \in \mathcal{F}_{M_{\phi_i},\nu}(\mathcal{C}_i)$. By [46, Prop. 1], h is generalized self-concordant with $\text{dom } h = \bigcap_{i=1}^n \text{dom } \phi_i$ and constant $M := \max\{M_{\phi_1}, \dots, M_{\phi_n}\}$. For different structures of the sets \mathcal{C}_i , we can propose different barrier-generating kernels ϕ_i . Here are some illustrative examples.

- 1 *Burg entropy*: $\phi(t) = -\log(t)$ for $t > 0$ is an element of $\mathcal{F}_{2,3}(\mathbb{R}_{++})$;
- 2 *Entropy-Barrier*: $\phi(t) = t \log(t) - \log(t)$ for $t > 0$ is an element of $\mathcal{F}_{2,3}(\mathbb{R}_{++})$;
- 3 Consider the function $\phi(t) = (1 - t/\kappa)^{-\kappa}$ for $\kappa > 0$ and $t \in (-\infty, \kappa)$. Then $\text{dom } \phi = (-\infty, \kappa)$ and one can check that $\phi \in \mathcal{F}_{M,\nu}((-\infty, \kappa))$, where $M = \frac{2+\kappa}{\kappa} \left(\frac{\kappa}{1+\kappa}\right)^{\frac{1}{2+\kappa}}$ and $\nu = \frac{2(3+\kappa)}{2+\kappa} \in (2, 3)$ for $\kappa > 0$.
- 4 The function $\phi(t) = \frac{1}{\sqrt{1-t^2}}$ defines an element in $\mathcal{F}_{M,\nu}((-1, 1))$ for M a constant smaller than 3.25 and $\nu = 14/5$.

Remark 2.2. Let $\bar{\mathcal{C}} = \prod_{i=1}^n [a_i, b_i]$ be a high-dimensional box of dimension $n \gg 1$ where $+\infty > b_i \geq a_i > -\infty$. According to [20], this geometry is computational challenging for standard proximal methods.² This geometry can be easily endowed with a barrier-generating kernel $h \in \mathcal{F}_{M,\nu}(\mathcal{C})$ given by a sum of Burg entropies or Entropy barriers, for instance, leading to a simple Riemannian metric on the interior of this box.

²To be clear, the challenge is not to find a good distance generating function, but rather the scalability of the mirror descent algorithm. We refer the reader to in-depth discussion in [20] for details.

The above examples provide a snapshot of common barrier-generating kernels used in practice. However, it is possible to combine these functions to obtain mixture functions that preserve the properties imposed on an element $h \in \mathcal{F}_{M,\nu}(\mathcal{C})$. In particular, it is easy to find barrier generating kernels for geometries which are given as intersections of open convex sets $\mathcal{C}_1, \dots, \mathcal{C}_J$, each admitting a generalized self-concordant function $h_j \in \mathcal{F}_{M_j,\nu}(\mathcal{C}_j)$, $1 \leq j \leq J$. This appears in second-order cone programming problems, which have as a special case optimization problems with quadratic constraints (see [33] for a survey). The typical sets appearing in such optimization problems are the following:

- $\mathcal{C} = \mathbb{R}_{++}^n$, $h(x) = -\sum_i \log(x_i)$ is a n -SCB and an element of $h \in \mathcal{F}_{2,3}(\mathcal{C})$;
- Consider the second-order cone $\bar{\mathcal{C}} \equiv \mathcal{L}^n := \text{cl}(\{x = (t, w) \in \mathbb{R} \times \mathbb{R}^{n-1} : t > \|w\|_2\})$. The function

$$h(x) = -\log(t^2 - \|w\|_2^2) \quad x := (t, w) \in \mathcal{C},$$

is a barrier-generating kernel belonging to $\mathcal{F}_{2,3}(\mathcal{C})$. It is also an 2-SCB.

- Consider the cone of positive definite symmetric $n \times n$ matrices with real entries $\Sigma_{++}^n := \{x \in \mathbb{R}^{n \times n} : x \succ 0, x^\top = x\}$, and set $\bar{\mathcal{C}} := \text{cl}(\Sigma_{++}^n)$. The function $h(x) = -\log \det(x)$ for $x \in \mathcal{C}$ is a barrier-generating kernel of class $\mathcal{F}_{2,3}(\mathcal{C})$. It is also an n -SCB.
- Let B be a $p \times n$ matrix with rows $b_1^\top, \dots, b_p^\top$, and d a given vector in \mathbb{R}^p . Consider the polyhedral set $\bar{\mathcal{C}} := \{x \in \mathbb{R}^n : Bx \leq d\}$. Assume that $\mathcal{C} = \{x \in \mathbb{R}^n : Bx < d\}$ is nonempty (Slater condition). Then, the function $h(x) = \sum_{j=1}^p -\log(d_j - b_j^\top x)$ is a barrier generating kernel belonging to the class $\mathcal{F}_{2,3}(\mathcal{C})$.
- Let $\mathcal{M}_{m,n}$ be the space of real $m \times n$ matrices with inner product $\langle A, B \rangle = \text{tr}(AB^\top)$. The standard operator norm is defined as $|Q| := \max\{\|Qw\| : \|w\| = 1\}$. Consider the set $\bar{\mathcal{C}} = \text{cl}(\{x = (t, Q) \in \mathbb{R} \times \mathcal{M}_{m,n} : t > |Q|\})$. This set admits a barrier-generating kernel $h \in \mathcal{F}_{2,3}(\mathcal{C})$ given by $h(x) = \log \det(t \text{Id} - \frac{1}{t} QQ^\top) - \log(t)$ for $x = (t, Q) \in \mathcal{C}$.

2.5 The minimization problem

We are given a matrix $A \in \mathbb{R}^{m \times n}$ of full row rank m and $b \in \text{im}(A)$. Define the sets $\mathcal{A} = \{x \in \mathbb{R}^n : Ax = b\}$, and $\mathcal{A}_0 = \ker(A)$, so that $\mathcal{A}_0^\perp = \text{im}(A^\top)$. Let \mathcal{C} be a nonempty open convex set in \mathbb{R}^n with closure $\bar{\mathcal{C}}$ that is not contained in any $(n-1)$ -dimensional affine subspace. Throughout the rest of this paper the following assumption is taken as a standing hypothesis.

Assumption 1. The set \mathcal{C} is nonempty, convex and contains no lines.

Combining this assumption with Lemma 2.8, we know that the Hessian matrix $H(x) = \nabla^2 h(x)$ is positive definite on \mathcal{C} . The matrix-valued function $H : \mathcal{C} \rightarrow \Sigma_{++}^n$ defines a Riemannian manifold $(\mathcal{C}, \|\cdot\|_x)$, with Riemannian metric given by (11). We are also given a lower semi-continuous $f : \mathbb{R}^n \rightarrow (-\infty, +\infty]$. The problem we aim to solve is the minimization problem

$$f^* := \inf\{f(x) : x \in \bar{\mathcal{C}}, Ax = b\} \tag{P}$$

The feasible set of (P) is denoted as $\mathcal{X} = \bar{\mathcal{C}} \cap \mathcal{A}$, and we shall denote by \mathcal{X}° the relative interior of \mathcal{X} , that is, $\mathcal{X}^\circ = \{x \in \mathbb{R}^n : x \in \mathcal{C}, Ax = b\}$. As a standing hypothesis, we shall impose the following Slater constraint qualification condition:

Assumption 2. $\mathcal{X}^\circ \neq \emptyset$.

Assumption 3. The level sets of the objective function are bounded: Given $x^0 \in \mathcal{X}^\circ$ there exists $R > 0$ such that $\sup\{\|x\|_\infty : f(x) \leq f(x^0)\} \leq R$.

For $\varepsilon > 0$, an ε global minimizer is defined as a feasible solution x_ε such that

$$f(x_\varepsilon) - \inf_{x \in \mathcal{X}} f(x) \leq \varepsilon. \quad (21)$$

It is well known that finding an ε -global minimizer is a strongly NP-hard problem (see e.g. [24]). Even worse, it is also well known that in general, finding a descent direction for a non-convex non-smooth function is NP-hard. As concrete illustration, even deciding whether the function

$$f(x) = (1 - 1/\gamma) \max_{1 \leq i \leq n} |x_i| - \min_{1 \leq i \leq n} |x_i| + |\langle c, x \rangle|,$$

where $x \in \mathbb{R}^n, c \in \mathbb{N}^n$ and $\gamma = \sum_{i=1}^n c_i$, admits a descent direction is NP-hard [41, Lem. 1]. Therefore, in this paper we restrict ourselves to objective functions of very special structure. Namely, we consider the problem of minimizing a real-valued function $f : \mathbb{R}^n \rightarrow (-\infty, \infty]$ which is continuously differentiable on an open convex set $\mathcal{C} \subset \mathbb{R}^n$, and possibly non-differentiable at the boundary $\text{bd } \mathcal{C} = \bar{\mathcal{C}} \setminus \mathcal{C}$.

Assumption 4. $f : \mathbb{R}^n \rightarrow (-\infty, \infty]$ is a proper and lower semi-continuous function with $f \in \mathbf{C}^1(\mathcal{C})$.

The smoothness condition formulated in Assumption 4 is silent about the behavior of the function at the boundary $\text{bd}(\mathcal{C})$. In case where the function f is twice continuously differentiable Cartis, Gould and Toint defined in [19] the following criticality measure at $x \in \mathcal{X}$

$$\chi_{CGT}(x) := \left| \min_{x+d \in \mathcal{X}, \|d\|_2 \leq 1} \langle \nabla f(x), d \rangle \right|. \quad (22)$$

They subsequently proved $\mathcal{O}(\varepsilon^{-2})$ iteration complexity for reaching a point with $\chi_{CGT}(x) \leq \varepsilon$. We propose a similar criticality measure here, but make use of the local norm. In particular, we consider the primal-dual stationarity measure at $(x, y) \in \mathcal{X}^\circ \times \mathbb{R}^m$ given by

$$\chi(x, y) := \|\nabla f(x) - A^\top y\|_x^*. \quad (23)$$

Definition 2.11. Given $\varepsilon > 0$, a pair $(x^*, y^*) \in \mathcal{X}^\circ \times \mathbb{R}^m$ is called ε -stationary if $\chi(x, y) \leq \varepsilon$.

In order to motivate this criticality measure, we first recall the classical Fenchel-Young inequality

$$|\langle u, v \rangle| \leq \|u\|_x \cdot \|v\|_x^* \quad \forall u, v \in \mathbb{R}^n, x \in \mathcal{X}^\circ. \quad (24)$$

Hence, for $v = \nabla f(x) - A^\top y$ and $u \in \mathcal{A}_0$ with $\|u\|_x = 1$, this inequality readily gives us

$$-\chi(x, y) \leq \langle \nabla f(x) - A^\top y, u \rangle = \langle \nabla f(x), u \rangle \leq \chi(x, y),$$

and in particular,

$$\left| \min_{u \in \mathcal{A}_0, \|u\|_x = 1} \langle \nabla f(x), u \rangle \right| \leq \chi(x, y).$$

Thus, the primal-dual criticality measure is an upper bound of a version of the Cartis-Gould-Toint criticality measure $\chi_{CGT}(x)$, and we note in passing that if $\chi(x, y) \leq \varepsilon$, then automatically $\langle \nabla f(x), u \rangle \in$

$(-\varepsilon, \varepsilon)$ for all $u \in \mathcal{A}_0$ satisfying $\|u\|_x = 1$. One potentially troublesome part in the definition of our proposed criticality measure is that it is formulated in terms of the local norm. Hence, we would need to evaluate the inverse matrix $[H(x)]^{-1}$ (provided it exists at x). However, for our algorithm this is not a problem since we will have very good control about the location of the iterates. Indeed, as will be seen in Section 5, the algorithm will take values on a compact set $\mathcal{S}_\mu(x^0)$ in \mathcal{X}° (this is similar to proximal based self-concordant algorithms and exploited in [36] in the convergence analysis). On this set, we have very good control on the eigenvalues of the Hessian matrix $H(x) = \nabla^2 h(x)$, and in fact, under assumption spelled out explicitly in the sections to follow, we can provide upper and lower bounds on the eigenvalues of $H(x)$ over the set $\mathcal{S}_\mu(x^0)$, denoted as $0 < \sigma_h < \tau_h < \infty$. Therefore, during the working phase of the algorithm, we produce a primal-dual sequence (x^k, y^k) along which the criticality measure is sandwiched as

$$\tau_h^{-1/2} \|\nabla f(x^k) - A^\top y^k\|_2 \leq \|\nabla f(x^k) - A^\top y^k\|_{x^k}^* \leq \sigma_h^{-1/2} \|\nabla f(x^k) - A^\top y^k\|_2. \quad (25)$$

Therefore, if the euclidean norm of the vector $\nabla f(x^k) - A^\top y^k$ falls below a cut-off $\varepsilon > 0$, we have reached an ε -stationary point in the sense of Definition 2.11.

Remark 2.3. This notion of stationarity is also motivated by the structure of the KKT conditions satisfied by a solution candidate for problem (P) taking values in the relative interior $\mathcal{X}^\circ = \mathcal{C} \cap \mathcal{A}$. As a concrete illustration, let us consider the set $\bar{\mathcal{C}} = \mathbb{R}_+^n$, so that we are in the setting of [11]. The complementary slackness condition for the resulting optimization problem (P) reads as

$$X(\nabla f(x) - A^\top y) = 0 \quad (26)$$

where $X = \text{diag}\{x_1, \dots, x_n\}$. Hence, a reasonable definition of an ε -KKT point under a Riemannian-Hessian structure induced by the Hessian of the function $h(x) = -\sum_{i=1}^n \ln x_i$ would read as

$$\|X(\nabla f(x) - A^\top y)\|_\infty \leq \|\nabla f(x) - A^\top y\|_x^* \leq \varepsilon.$$

This ε -KKT definition has also been used in [28].

3 The Hessian-barrier method

In this section we describe a conceptual version of the Hessian-barrier method. To this end, we are given an open nonempty set $\mathcal{C} \subset \mathbb{R}^n$ satisfying Assumption 1, admitting a computable barrier generating kernel $h \in \mathcal{F}_{M,\nu}(\mathcal{C})$.

3.1 Defining the search directions

For a pair $(x, g) \in \mathcal{C} \times \mathbb{R}^n$, define the functions

$$\psi(x, g) := \min_v \left\{ \langle g, v \rangle + \frac{1}{2} \|v\|_x^2 : Av = 0 \right\}, \text{ and}$$

$$V(x, g) := \operatorname{argmin}_v \left\{ \langle g, v \rangle + \frac{1}{2} \|v\|_x^2 : Av = 0 \right\}.$$

Computing the vector $V(x, g)$ means finding a pair $(V(x, g), y(x, g)) = (v, y) \in \mathbb{R}^n \times \mathbb{R}^m$ solving the Newton-type of system

$$\begin{bmatrix} H(x) & -A^\top \\ -A & 0 \end{bmatrix} \cdot \begin{bmatrix} v \\ y \end{bmatrix} = \begin{bmatrix} -g \\ 0 \end{bmatrix}. \quad (27)$$

In particular, the complexity of computing $V(x, g)$ is of the same order as finding a Newton direction, and the practical efficiency of the method depends heavily on the structure of the matrices $H(x)$ and A , respectively. In any case, given that $H(x) \in \Sigma_{++}^n$, we obtain a closed form expression for the vector $V(x, g)$ as

$$V(x, g) = -P_x[H(x)]^{-1}g \quad \forall (x, g) \in \mathcal{C} \times \mathbb{R}^n, \text{ and} \quad (28)$$

$$y(x, g) = (A[H(x)]^{-1}A^\top)^{-1}A[H(x)]^{-1}g \quad (29)$$

where the matrix valued function $P : \mathcal{C} \rightarrow \mathbb{R}^{n \times n}$ defined in (8). We just remark that, given the matrix A being of full rank, the function $x \mapsto A[H(x)]^{-1}A^\top$ is invertible [11]. Computational efficiency considerations will be made later. We close this section by establishing some general properties of the mapping ψ and V .

Proposition 3.1. *The following assertions are true:*

(a) *The mapping $\psi(x, \cdot) : \mathbb{R}^n \rightarrow \mathbb{R}$ is concave and continuously differentiable with*

$$\nabla_g \psi(x, g) = V(x, g) \quad \forall (x, u) \in \mathcal{C} \times \mathbb{R}^n.$$

(b) *If h is K -strongly convex under the ℓ_2 norm, then for every $x \in \mathcal{C}$, the mappings $V(x, \cdot)$ and $\psi(x, \cdot)$ are $\frac{1}{K}$ -Lipschitz,*

(c) *If $g = p + z \in \text{NC}_{\mathcal{X}}(x)$ with $p \in \text{NC}_{\bar{c}}(x)$, $z \in \mathcal{A}_0^\perp$, then*

$$V(x, p + z) = V(x, p)$$

Proof. Since $\psi(x, g)$ is the pointwise minimum of a linear function, it must be concave. The integrability condition on the vector field $g \mapsto V(x, g)$ is a straightforward computation. Parts (b) and (c) are standard, and follow from the general analysis of such projection schemes as in [40]. It is however instructive here to go over the computations. First, the K -strong convexity of the norm ensures that $u \mapsto \psi(x, u)$ is well defined and convex. In particular, $g \mapsto V(x, g)$ is uniquely defined by eq. (28). Therefore,

$$\psi(x, g) = \frac{1}{2} \|V(x, g)\|_x^2 \quad \forall (x, g) \in \mathcal{C} \times \mathbb{R}^n.$$

Let $g_1, g_2 \in \mathbb{R}^n$ be arbitrary and set $v_1 = V(x, g_1)$, $v_2 = V(x, g_2)$. The optimality conditions at a given point $x \in \mathcal{C}$ imply

$$(g_1 - H(x)v_1)^\top (v_2 - v_1) = 0 \text{ and}$$

$$(g_2 - H(x)v_2)^\top (v_1 - v_2) = 0.$$

Adding both, and using the K -strong convexity shows

$$(g_1 - g_2)^\top (v_1 - v_2) = (H(x)(v_1 - v_2))^\top (v_1 - v_2) \geq K \|v_1 - v_2\|_2^2.$$

Using the Cauchy-Schwarz inequality, we arrive at

$$\|V(x, g_1) - V(x, g_2)\|_2 \leq \frac{1}{K} \|g_1 - g_2\|.$$

■

Corollary 3.2. *For all $x \in \mathcal{X}^\circ$, we have*

$$V(x, g) = 0 \quad \forall g \in \text{NC}_{\mathcal{X}}(x) = \text{NC}_{\bar{c}}(x) + \mathcal{A}_0^\perp.$$

Proof. Just observe that for $x \in \mathcal{X}^\circ$ we have $\text{NC}_{\bar{c}}(x) = \{0\}$ and thus $g \in \mathcal{A}_0^\perp$. Hence, the claim follows from (c) of Proposition 3.1. ■

3.2 The Hessian-Barrier potential reduction algorithm

Based on the family of search directions $V(x, g)$, we now tailor the gradient input $g \in \mathbb{R}^n$ to derive a potential reduction algorithm solving problem (P). Throughout this paper we will work with a pair of functions (f, h) such that:

- (i) $h \in \mathcal{F}_{M,\nu}(\mathcal{C})$ with some parameters $M > 0$ and $\nu \in (2, 4]$;
- (ii) $\text{cl}(\text{dom } h) = \bar{\mathcal{C}}$;
- (iii) $f : \mathbb{R}^n \rightarrow (-\infty, \infty]$ obeys Assumption 4 and $\text{dom}(h) \subseteq \text{dom}(f)$.

The next definition, due to [4, 35] and [8], is fundamental to our analysis.

Definition 3.3. The pair of functions (f, h) is L -smooth if there exists a constant $L > 0$ such that

$$f(y) - f(x) - \langle \nabla f(x), y - x \rangle \leq LD_h(y, x) \quad \forall x, y \in \mathcal{C} \quad (\text{L})$$

It is easy to check that (f, h) being L -smooth is equivalent to $Lh - f$ being convex. Define the potential function

$$F_\mu(x) = f(x) + \mu h(x) \quad \forall x \in \text{dom } h. \quad (30)$$

If (f, h) is an L -smooth pair, then the function $(L + \mu)h - F_\mu$ must be convex. Therefore,

$$F_\mu(x) \leq F_\mu(y) + \langle \nabla F_\mu(y), x - y \rangle + (L + \mu)D_h(x, y) \quad \forall (x, y) \in \text{dom } h \times \text{dom } h. \quad (31)$$

This inequality is in fact a modified descent lemma, in the spirit of [4], for the non-convex, non-smooth composite function F_μ .

Remark 3.1. If $\mathcal{C} = \mathbb{R}^n$ and $h(x) = \frac{1}{2}\|x\|^2$, the pair (f, h) is L -smooth if and only if the classical descent inequality

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

holds for all $x, y \in \mathbb{R}^n$, i.e. the parameter L is a surrogate for the Lipschitz constant of the Euclidean gradient map $x \mapsto \nabla f(x)$.

Define the search direction

$$d_\mu(x) := V(x, \nabla F_\mu(x)) \quad \forall x \in \mathcal{X}^\circ. \quad (32)$$

From the first-order optimality condition of the search direction (27), we know that

$$\langle \nabla F_\mu(x), d_\mu(x) \rangle = -\|d_\mu(x)\|_x^2 \quad \forall x \in \mathcal{X}^\circ. \quad (33)$$

The associated dual variable is obtained by the evaluation of (29) as $y_\mu(x) := y(x, \nabla F_\mu(x))$.

Lemma 3.4. *The dual function $y_\mu : \mathcal{C} \rightarrow \mathbb{R}^m$ is continuous.*

Proof. By (29), the function $y_\mu(x)$ has the explicit expression $y_\mu(x) = (A[H(x)]^{-1}A^\top)^{-1}A[H(x)]^{-1}\nabla F_\mu(x)$. Since $h \in \mathbf{C}^3(\mathcal{C})$, the matrix-valued mapping $H(x) = \nabla^2 h(x)$ is continuous, and $x \mapsto \nabla F_\mu(x)$ is continuous as well on \mathcal{C} . The claim follows. \blacksquare

Define

$$\lambda_\mu(x) := \|d_\mu(x)\|_x, \text{ and } \beta_\mu(x) := \|d_\mu(x)\|_2, \quad (34)$$

and the transfer function

$$T_\mu(x, \alpha) := x + \alpha d_\mu(x) \quad x \in \mathcal{C}, \alpha > 0, \mu > 0. \quad (35)$$

This mapping $T_\mu : \mathcal{C} \times [0, \infty) \rightarrow \mathbb{R}^n$ will serve as the generator of the numerical algorithm. Thanks to generalized self-concordance, we can easily determine the step length values $\alpha > 0$ guaranteeing that $T_\mu(x, \alpha) \in \mathcal{X}^\circ$ for $x \in \mathcal{X}^\circ$. Indeed, by Lemma 2.6, a sufficient condition ensuring that we stay in the interior of the feasible set is to set $\alpha > 0$ such that $d_\nu(x, T_\mu(x, \alpha)) < 1$. This leads to the bound $\alpha M^{\frac{\nu-2}{2}} \lambda_\mu(x)^{\nu-2} \beta_\mu(x)^{3-\nu} < 1$. Let us define

$$\delta_\mu(x) := M^{\frac{\nu-2}{2}} \lambda_\mu(x)^{\nu-2} \beta_\mu(x)^{3-\nu}, \quad (36)$$

so that $d_\nu(x, T_\mu(x, \alpha)) = \alpha \delta_\mu(x)$ for all $x \in \mathcal{X}^\circ, \alpha \geq 0$. Furthermore, if $\delta_\mu(x) > 0$, we define $\bar{\alpha}_\mu(x) := 1/\delta_\mu(x)$. Therefore, any choice of step size $\alpha \in (0, \bar{\alpha}_\mu(x))$, delivers a feasible step. Furthermore, for all $\alpha \in (0, \bar{\alpha}_\mu(x))$, we can apply the general descent inequality (31) to the L -smooth pair (f, h) , so that (31), (33) and (36) give us the per-iteration estimate

$$\begin{aligned} F_\mu(T_\mu(x, \alpha)) &\leq F_\mu(x) + \alpha \langle \nabla F_\mu(x), d_\mu(x) \rangle + (L + \mu) D_h(T_\mu(x, \alpha), x) \\ &= F_\mu(x) - \alpha \lambda_\mu(x)^2 + (L + \mu) D_h(T_\mu(x, \alpha), x). \end{aligned}$$

Combining this with Lemma 2.7, we see that for all $x \in \mathcal{C}$ and $\alpha \in (0, \bar{\alpha}_\mu(x))$

$$\begin{aligned} F_\mu(T_\mu(x, \alpha)) &\leq F_\mu(x) - \alpha \lambda_\mu(x)^2 + (L + \mu) \omega_\nu(\alpha \delta_\mu(x)) \alpha^2 \lambda_\mu(x)^2 \\ &= F_\mu(x) - \eta_\mu(x, \alpha), \end{aligned} \quad (37)$$

where we have set

$$\eta_\mu(x, t) := t \lambda_\mu(x)^2 - (L + \mu) \omega_\nu(t \delta_\mu(x)) t^2 \lambda_\mu(x)^2. \quad (38)$$

Note that the barrier-generating kernel $h \in \mathcal{F}_{M, \nu}(\mathcal{C})$ only appears in this per-iteration bound via the local norm of the search direction $\lambda_\mu(x)$. As such, the above bound can be seen as worst-case bound on the potential function decrease. This worst-case point of view is however very useful in determining an explicit step-size policy, akin to the recently proposed prox-based algorithms for convex composite self-concordant minimization [46, 49].

Proposition 3.5. *For all $x \in \mathcal{X}^\circ, \mu, L > 0$ and $\alpha \in (0, \bar{\alpha}_\mu(x))$, we have $T_\mu(x, \alpha) \in \mathcal{X}^\circ$. The optimal step-size rule, in the analytical worst-case sense, is given by*

$$\alpha_\mu(x, L) := \begin{cases} \frac{1}{\delta_\mu(x)} \left[1 - \left(1 + \frac{\delta_\mu(x)}{L+\mu} \frac{4-\nu}{\nu-2} \right)^{-\frac{\nu-2}{4-\nu}} \right] & \text{if } \nu \in (2, 3) \cup (3, 4), \\ \frac{1}{\delta_\mu(x) + L + \mu} & \text{if } \nu = 3, \\ \frac{1}{\delta_\mu(x)} \left[1 - \exp\left(-\frac{\delta_\mu(x)}{L+\mu}\right) \right] & \text{if } \nu = 4. \end{cases} \quad (39)$$

The proof of this Proposition is a rather technical computation, and therefore delegated to Appendix A. It is however interesting to note that the self-concordance parameter ν plays a somewhat symmetric role around its values $\nu \in (2, 3) \cup (3, 4)$. Moreover, it is remarkable that the theoretical upper bound on the step size ensuring feasibility, $\bar{\alpha}_\mu(x)$, is independent of the constant L . It appears only when we




Figure 1: Step size $\alpha = s_{\nu,L}(t)$ for $L + \mu = 10$ and $\nu \in \{2.6, 3, 3.5, 4\}$.

compute the optimal step size $\alpha_{\mu}(x, L)$, which, in turn, is available in a closed-form expression. This functional form of the step size policy allows for a direct comparison in dependence of the generalized self-concordance parameter $\nu \in (2, 4]$. In Figure 1 we provide a numerical illustration on the ordering of the step sizes, realizing for fixed parameter pair (μ, L) , we can think of the function $\alpha_{\mu}(x, L)$ as the output of a function of the composition $s_{\mu,L,\nu} \circ \delta_{\mu}$, where $s_{\mu,L,\nu} : (0, \infty) \rightarrow (0, \infty)$ is given by

$$s_{\mu,L,\nu}(t) := \begin{cases} \frac{1}{t} \left[1 - \left(1 + \frac{t}{L+\mu} \frac{4-\nu}{\nu-2} \right)^{-\frac{\nu-2}{4-\nu}} \right] & \text{if } \nu \in (2, 3) \cup (3, 4), \\ \frac{1}{t+L+\mu} & \text{if } \nu = 3, \\ \frac{1}{t} \left[1 - \exp\left(-\frac{t}{L+\mu}\right) \right] & \text{if } \nu = 4. \end{cases}$$

We see that larger parameters ν lead to higher step sizes and thus to more aggressive schemes. This clearly indicates that the choice of the metric-inducing kernel matters in the design of the algorithm. Observe that the function $s_{\mu,L,\nu}$ is well defined at $t = 0$ and attains the same value $1/(L + \mu)$ for all $\nu \in (2, 4]$.

Both parameters μ and L are seen to have the same effect on the step size policy: Larger values imply smaller step sizes. Hence, for optimization purposes it is of utmost interest to pick these parameters in a way that prevents $\text{HBA}(\mu, L)$ making to small steps. However, the two parameters also play different roles in the design of the method. While μ is a *barrier parameter* guaranteeing that the algorithm is an interior-point method, the parameter L is dictated by the pair (f, h) in order to guarantee the descent property (L). Intuitively, we would like to run the algorithm with a numerically small value of μ . The descent property (L) tells us that for any $\tilde{L} \geq L$, we can guarantee a sufficient decrease in the potential function, so the design question becomes how small the the parameter L can be chosen with a guarantee to obtain a sufficient decrease. In order to answer these questions, we will develop an adaptive version of the base algorithm $\text{HBA}(\mu, L)$. Since the analysis of this adaptive version will rely on general results obtained for the base scheme $\text{HBA}(\mu, L)$, we start our mathematical analysis with the assumption that both parameters are fixed, and later make their choice dynamic.

4 The Hessian-barrier algorithm

Let $c > 0$ be a positive constant. For the construction of our algorithmic scheme it will be important to have access to a c -analytic center, i.e. a point $x^0 \in \mathcal{X}^\circ$ such that

$$h(x) \geq h(x^0) - c \quad \forall x \in \mathcal{X}. \quad (40)$$

Data: kernel generating distance $h \in \mathcal{F}_{M,\nu}(\mathcal{C})$ such that (f, h) is L -smooth;

Barrier parameter $\mu > 0$.

Result: Stationary point of F_μ .

Initial point: c-analytic center $x^0 \in \mathcal{X}^\circ$;

while $k = 0, 1, \dots, k_{\max}$ **do**

 obtain x^k and $\nabla F_\mu(x^k)$;

if *Stopping condition not satisfied* **then**

 Solve the linear system (27) for $x = x^k$ and $g = \nabla F_\mu(x^k)$;

 Denote by $(v_\mu^k, y_\mu^k)^\top$ the solution;

 Compute step size $\alpha_k = \alpha_\mu(x^k, L)$;

 Update $x^{k+1} = x^k + \alpha_k v_\mu^k$.

else

 Stop and report x^k as the solution

end

end

Algorithm 1: HBA(μ, L)

To obtain such a point x^0 , we need to be able to approximately solve the minimization problem

$$\min_{x \in \mathcal{X}} h(x).$$

This assumption is very common in potential reduction schemes [28, 54, 55]. In case where \mathcal{X} is bounded, existence of an exact analytic center is guaranteed and we can use any efficient solver for computing it. In the case where \mathcal{X} is unbounded, existence of an exact analytic center is guaranteed if the classical existence condition given by the Weierstrass' theorem are satisfied. If $h \in \mathcal{F}_{M,\nu}(\mathcal{C})$ for $\nu \in (3, 4]$ then by Definition 2.10 the function h is coercive, and therefore the program (40) always has a solution. In the remaining cases, where $\nu \in (2, 3]$ we either have to assume that the feasible set is bounded or that there exists a point $x \in \mathcal{X}$ under which the dual norm $\|\nabla h(x)\|_x^*$ is β_{small} . A precise meaning of this statement can be given by adapting the relevant arguments in [46] to the current geometry with linear equality constraints and leave this, somewhat off-topic, exercise to the reader.

4.1 HBA iterations for fixed μ

Given the general template described in Section 3, the first algorithmic scheme of interest in this paper is easy to describe. Assuming the standing hypothesis Assumptions 1-4 in place, we are given a pair of functions (f, h) such that $h \in \mathcal{F}_{M,\nu}(\mathcal{C})$, and L -smoothness holds for some $L \geq 0$. Given a c-approximate analytic center $x^0 \in \mathcal{X}^\circ$ as initial condition, we generate a sequence $(x^k)_{k \geq 0}$ recursively by

$$x^{k+1} = T_\mu(x^k, \alpha_k) = x^k + \alpha_k v_\mu^k \quad \forall k \geq 0,$$

where $T_\mu(x, \alpha)$ is the transfer function defined in (35), with step-size $\alpha_k = \alpha_\mu(x^k, L)$ and search direction $v_\mu^k = d_\mu(x^k)$. The pseudo-code corresponding to the conceptual implementation of HBA(μ, L) reads as Algorithm 1.

Let us make some remarks on the computational efficiency of HBA(μ, L). Since \mathcal{C} contains no lines and the sequence $(x^k)_{k \geq 0}$ stays in the relative interior for the optimization problem's feasible set, we guarantee that $H(x^k) \succ 0$ for every iteration of the algorithm. Hence, the main computational

Data: kernel generating distance $h \in \mathcal{F}_{M,\nu}(\mathcal{C})$ such that (f, h) is L -smooth;

Barrier parameter $\mu > 0$.

Result: Stationary point of F_μ .

Initial point: c-analytic center $x^0 \in \mathcal{X}^\circ$;

while $k = 0, 1, \dots, k_{\max}$ **do**

 obtain x^k ;

if *Stopping condition not satisfied* **then**

 Solve the linear system (27) for $x = x^k$ and $g = \nabla F_\mu(x^k)$;

 Denote by $(v_\mu^k, y_\mu^k)^\top$ the solution;

 Find the smallest $i_k \geq 0$ such that $z^k = x^k + \alpha_\mu(x^k, 2^{i_k-1} L_k) v_\mu^k$ satisfies

$$f(z^k) \leq f(x^k) + \langle \nabla f(x^k), x^{k+1} - x^k \rangle + 2^{i_k-1} L_k D_h(z^k, x^k). \quad (41)$$

 Set $L_{k+1} = 2^{i_k-1} L_k$;

 Update $x^{k+1} = z^k$.

else

 Stop and report x^k as the solution.

end

end

Algorithm 2: AHBA(μ)

step in HBA(μ, L) (27) is always well-posed while executing the algorithm, and delivers a unique solution. The complexity of (27) is the same as for a Newton method; It requires $\mathcal{O}(n^3)$ operations via either a Cholesky decomposition, or a well implemented conjugate gradient (CG) method. Hence, in terms of per-iteration complexity HBA(μ, L) is comparable with Newton methods. In many applications, however, the matrix $H(x^k)$ has a special structure which makes the application of heavy analytic machinery unnecessary. Indeed, most barrier functions h used in the literature are additively separable and the resulting Hessian matrix H is therefore diagonal. In this case, the computational complexity is essentially determined by the density of the matrix A , and in many applications (e.g. resource allocations problems where A embodies network flow constraints) we will be able to implement a closed-form expression for the search direction v_μ^k . Hence, in such favorable instances, the per-iteration computational overhead of implementing HBA(μ, L) is rather small.

4.2 Adaptive HBA

The basic algorithmic scheme HBA(μ, L) utilizes knowledge of the barrier parameter μ and the L -smoothness parameter $L > 0$ in order to determine the step size $\alpha_\mu(x, L)$. Knowing a-priori the parameter μ is not a very demanding, since it is chosen by the user at the beginning of the implementation. However, running the basic Hessian-barrier algorithm scheme with a stiff parameter L might be inefficient since it forces us to rescale the step size with the same constant factor $\mu + L$ globally. This might lead to unnecessary small steps, resulting in long run times of the method. To overcome this drawback we present in this section a new adaptive method of Hessian-barrier algorithm, where all necessary information about L can be accumulated by an appropriate "line-search" strategy. The thus resulting Adaptive-Hessian-barrier algorithm method closely resembles ideas spelled out in "universal gradient methods" defined by Nesterov in [42].

5 Complexity analysis of HBA

We organize our discussion on the long-run properties of $\text{HBA}(\mu, L)$ and $\text{AHBA}(\mu)$ in two parts. The first part is concerned with the asymptotic convergence properties of the two methods. In the second part we will discuss the non-asymptotic complexity properties of the method. Throughout this section we assume that Assumptions 1-4 are satisfied.

5.1 Asymptotic convergence

Let $(x^k)_{k \geq 0}$ be a sequence generated by $\text{HBA}(\mu, L)$, with search direction $v_\mu^k = d_\mu(x^k)$ and step-size policy $\alpha_k = \alpha_\mu(x^k, L)$. Let us introduce the associated sequences $(\lambda_k)_{k \geq 0}, (\beta_k)_{k \geq 0}, (\delta_k)_{k \geq 0} \subset [0, \infty)$ by

$$\lambda_k := \|v_\mu(x^k)\|_{x^k}, \quad \beta_k := \|v_\mu(x^k)\|_2, \quad \delta_k := M \left(\frac{\nu - 2}{2} \right) \lambda_k^{\nu-2} \beta_k^{3-\nu}. \quad (42)$$

Let us define the per-iteration progress along the thus produced sequence, quantified in (37), as

$$\Delta_k := \eta_\mu(x^k, \alpha_k) \quad \forall k \geq 0. \quad (43)$$

From Proposition 3.5, we immediately deduce an online version of the step-sizes together with a descent inequality for the potential function, summarized in the next Proposition.

Proposition 5.1. *Let $(x^k)_{k \geq 0}$ be the sequence generated by $\text{HBA}(\mu, L)$ with the step size policy*

$$\alpha_k := \begin{cases} \frac{1}{\delta_k} \left[1 - \left(1 + \frac{\delta_k}{L+\mu} \frac{4-\nu}{\nu-2} \right)^{-\frac{\nu-2}{4-\nu}} \right] & \text{if } \nu \in (2, 3) \cup (3, 4), \\ \frac{1}{\delta_k + L + \mu} & \text{if } \nu = 3, \\ \frac{1}{\delta_k} \left[1 - \exp\left(-\frac{\delta_k}{L+\mu}\right) \right] & \text{if } \nu = 4. \end{cases} \quad (44)$$

Then, for all $k \geq 0$ we have

$$F_\mu(x^{k+1}) \leq F_\mu(x^k) - \Delta_k, \quad (45)$$

where Δ_k is defined in (43). Moreover, this step size rule is optimal in the worst-case analytic sense.

We next provide some general properties of $\text{HBA}(\mu, L)$.

Proposition 5.2. *Let $(x^k)_{k \geq 0}$ be generated by $\text{HBA}(\mu, L)$, and set $f^* := \inf_{x \in \mathcal{X}} f(x)$. Then, the following assertions hold:*

- (a) $(F_\mu(x^k))_{k \geq 0}$ is non-increasing;
- (b) $\sum_{k \geq 0} \Delta_k < \infty$, and hence the sequence $(\Delta_k)_{k \geq 0}$ converges to 0;
- (c) $\min_{0 \leq k < K} \Delta_k \leq \frac{1}{K} [f(x^0) - f^* + \mu c]$.

Proof. Unraveling the expressions in eq. (31), we get for all $k \geq 0$,

$$f(x^{k+1}) - f(x^k) \leq -\Delta_k + \mu [h(x^k) - h(x^{k+1})].$$

Telescoping this expression shows that for all $K \geq 1$,

$$f(x^K) - f(x^0) \leq - \sum_{k=0}^{K-1} \Delta_k + \mu[h(x^0) - h(x^K)].$$

Since x^0 is a c-analytic center, the left-hand side in the above display can be majorized to obtain the bound

$$f(x^K) - f(x^0) \leq - \sum_{k=0}^{K-1} \Delta_k + \mu c.$$

Since $\Delta_k > 0$, the sequence $(F_\mu(x^k))_{k \geq 0}$ is monotonically decreasing. Since $h(x) - h(x^0) \geq -c$, and f is bounded from below, the potential function F_μ is bounded from below as well. Therefore $\liminf_{k \rightarrow \infty} F_\mu(x^k) = \lim_{k \rightarrow \infty} F_\mu(x^k)$ exists and equals a number $F_\mu^* \in (-\infty, \infty)$. It follows that $(x^k)_{k \geq 0} \subset \text{dom } h = \mathcal{C}$, and therefore $\lim_{k \rightarrow \infty} f(x^k)$ exists as well. Calling $f^* := \inf\{f(x) : x \in \mathcal{X}\} > -\infty$, we conclude that for all $K \geq 1$,

$$\sum_{k=0}^{K-1} \Delta_k \leq f(x^0) - f(x^K) + \mu c \leq f(x^0) - f^* + \mu c, \quad (46)$$

and

$$\min_{1 \leq k \leq K} \Delta_k \leq \frac{1}{K} [f(x^0) - f^* + \mu c]. \quad (47)$$

Hence, $\lim_{k \rightarrow \infty} \Delta_k = 0$. ■

We turn now to the convergence properties of HBA(μ, L). Our aim is to show that accumulation points of the sequence $(x^k)_{k \geq 0}$ generated by the algorithm are stationary points of the potential function F_μ . We start by proving some auxiliary results.

Lemma 5.3. *Let $(x^k)_{k \geq 0}$ be generated by HBA(μ, L). Then, $(x^k)_{k \geq 0}$ is bounded.*

Proof. Since $(F_\mu(x^k))_{k \geq 0}$ is monotonically decreasing, we have

$$f(x^{k+1}) - f(x^k) \leq -\Delta_k - \mu[h(x^{k+1}) - h(x^k)].$$

Hence, for all $K \geq 1$, using the c-analytic center property of the initial condition x^0 , we get

$$f(x^K) \leq f(x^0) + \mu c.$$

Hence, $x^k \in \text{lev}_f(f(x^0) + \mu c)$. Since f has bounded level sets (Assumption 3), the entire sequence $(x^k)_{k \geq 0}$ is bounded. ■

Define the limit set

$$\omega(x^0) := \{p \in \mathcal{X} : \exists (k_q)_{q \in \mathbb{N}} \uparrow \infty, \lim_{k_q \rightarrow \infty} x^{k_q} = p\}. \quad (48)$$

Thanks to Lemma 5.3, standard results imply that $\omega(x^0)$ is nonempty, connected and compact (see e.g. [10, Lem.5]). Furthermore, $\lim_{k \rightarrow \infty} \text{dist}(x^k, \omega(x^0)) = 0$. For $x \in \mathcal{X}^\circ$, define

$$\mathcal{S}_\mu(x) := \{y \in \mathcal{X} : F_\mu(y) \leq F_\mu(x)\} = \text{lev}_{F_\mu}(F_\mu(x)) \cap \mathcal{A}. \quad (49)$$

Since HBA(μ, L) is a descent method for the potential function F_μ , we immediately conclude that $(x^k)_{k \geq 0} \subseteq \mathcal{S}_\mu(x^0)$.

Lemma 5.4. *Let x^0 be a c -analytic center. Then $\mathcal{S}_\mu(x^0)$ is a compact subset in $\mathcal{X}^\circ = \mathcal{C} \cap \mathcal{A}$.*

Proof. Note that

$$\begin{aligned}\mathcal{S}_\mu(x^0) &= \{x \in \mathcal{X} : f(x) \leq f(x^0) + \mu[h(x^0) - h(x)]\} \\ &\subseteq \{x \in \mathcal{X} : f(x) \leq f(x^0) + \mu c\} \\ &= \text{lev}_f(f(x^0) + \mu c) \cap \mathcal{A}.\end{aligned}$$

Since f has bounded level sets (Assumption 3), the set $\mathcal{S}_\mu(x^0)$ is bounded as well. It remains to prove that the set $\mathcal{S}_\mu(x^0)$ is closed. To that end, let $(x_j)_{j \geq 1}$ be a converging sequence with $\lim_{j \rightarrow \infty} x_j = \bar{x}$ and $x_j \in \mathcal{S}_\mu(x^0)$ for all $j \geq 1$. Then, $f(x^j) + \mu h(x^j) \leq r \equiv f(x^0) + \mu h(x^0)$ for all $j \geq 1$. If $\bar{x} \in \text{bd}(\mathcal{C})$, then $h(x_j) \rightarrow \infty$ and we immediately obtain a contradiction. Hence, $x \in \mathcal{C} \cap \mathcal{A}$, and the restriction of the composite function $f + \mu h$ on this domain is continuous. We conclude $\bar{x} \in \mathcal{S}_\mu(x^0)$. ■

Corollary 5.5. $\omega(x^0) \subset \mathcal{X}^\circ$.

Let $\sigma_{\min}(x)$ denote the smallest and $\sigma_{\max}(x)$ the largest eigenvalue of the Hessian $H(x) = \nabla^2 h(x)$. Since $H(x) \succ 0$ for all $x \in \mathcal{X}^\circ$, we conclude that $\sigma_{\min}(x) > 0$. Moreover, the compactness of the set $\mathcal{S}_\mu(x^0)$ allows us to define the positive constant

$$\sigma_h := \min_{x \in \mathcal{S}_\mu(x^0)} \sigma_{\min}(x) \quad (50)$$

Hence, along the iterates of $\text{HBA}(\mu, L)$, we have

$$\lambda_k \geq \sqrt{\sigma_h} \beta_k \quad \forall k \geq 0. \quad (51)$$

Since

$$\alpha_k \delta_k = \mathbf{d}_\nu(x^k, x^{k+1}) = M \left(\frac{\nu}{2} - 1 \right) \lambda_k^{\nu-2} \beta_k^{3-\nu},$$

the following lower and upper bounds can be established for $\nu \in (2, 3]$:

$$M \left(\frac{\nu}{2} - 1 \right) \sigma_h^{\frac{\nu-2}{2}} \alpha_k \beta_k \leq \alpha_k \delta_k \leq M \left(\frac{\nu}{2} - 1 \right) \sigma_h^{-\frac{3-\nu}{2}} \alpha_k \lambda_k. \quad (52)$$

This inequality will be key to prove convergence of the method to a stationary point of the potential function when $\nu \in (2, 3]$. For $\nu \in (3, 4]$, we will need to upper bound the local norm of the search direction, λ_k , as well. Let $\sigma_{\max}(x) \in (0, \infty]$ be the largest eigenvalue of the Hessian matrix $H(x)$. Since $\mathcal{S}_\mu(x^0)$ is a compact set in \mathcal{X}° , the quantity

$$\tau_h := \max_{x \in \mathcal{S}_\mu(x^0)} \sigma_{\max}(x) \quad (53)$$

is well-defined and finite. Given these bounds, we see that for all $x \in \mathcal{S}_\mu(x^0)$ we have

$$\sigma_h \text{Id} \leq H(x) \leq \tau_h \text{Id} \quad (54)$$

so that the function h is σ_h -smooth and τ_h -strongly convex on the compact set $\mathcal{S}_\mu(x^0)$. The quantity $\kappa_h = \frac{\tau_h}{\sigma_h} \geq 1$ is the condition number of h . Hence, along the sequence $(x^k)_{k \geq 0}$ generated by $\text{HBA}(\mu, L)$, we can upper bound the local norm of the search direction by

$$\lambda_k \leq \sqrt{\tau_h} \beta_k \quad \forall k \geq 0. \quad (55)$$

All these estimates together will be needed to prove the main result of this section, represented by the following Theorem.

Theorem 5.6. *Let $(x^k)_{k \geq 0}$ be the sequence generated by HBA(μ, L) with step-size policy $(\alpha_k)_{k \geq 0}$ described in (44). Then, $\omega(x^0) \subseteq \{x \in \mathcal{X} : (\exists y \in \mathbb{R}^m) : \nabla F_\mu(x) - A^\top y = 0\}$.*

Proof. See Appendix B. ■

As a consequence of this Theorem, it follows that the trajectory $(x^k)_{k \geq 0}$ exhibits a decaying energy in the metric-like function d_ν :

Corollary 5.7. $\lim_{k \rightarrow \infty} d_\nu(x^k, x^{k+1}) = 0$.

Proof. By definition, $d_\nu(x^k, x^{k+1}) = \alpha_k \delta_k$ for all $k \geq 0$ and $\nu \in (2, 4]$. In Appendix B we have shown that $\liminf_{k \rightarrow \infty} \alpha_k > 0$ and $\limsup_{k \rightarrow \infty} \delta_k = 0$. The claim follows. ■

5.2 Non-asymptotic bounds

In this section we provide complexity estimates for the non-adaptive base algorithm HBA(μ, L). To do so, we report first a useful technical corollary of the proof of Theorem 5.6.

Lemma 5.8. *Let $(x^k)_{k \geq 0}$ be generated by HBA(μ, L), with corresponding potential reduction sequence $(\Delta_k)_{k \geq 0}$ defined in (43). For each generalized self-concordance parameter $\nu \in (2, 4]$, there exists a strictly increasing function $\tilde{\omega}_\nu : (0, \infty) \rightarrow (0, \infty)$ satisfying*

$$\Delta_k \geq \tilde{\omega}_\nu(\lambda_k) \quad \forall k \geq 0. \quad (56)$$

In particular, this function is given by

$$\tilde{\omega}_\nu(t) := \begin{cases} \tilde{\gamma}_\nu t \min \left\{ \frac{2\sigma_h^{\frac{3-\nu}{2}}}{M(\nu-2)}, \frac{t}{-\mathfrak{b}(L+\mu)} \right\} & \text{if } \nu \in (2, 3), \\ \tilde{\gamma}_\nu t \min \left\{ \frac{2}{M(\nu-2)} \tau_h^{-\frac{3-\nu}{2}}, \frac{t}{-\mathfrak{b}(L+\mu)} \right\} & \text{if } \nu \in (3, 4), \\ \frac{2(1-\ln(2))t}{M(L+\mu)} \min \left\{ (L+\mu), \frac{M}{2}t \right\} & \text{if } \nu = 3, \\ t \exp(-1) \min \left\{ \frac{1}{\sqrt{\tau_h}M}, \frac{t}{L+\mu} \right\} & \text{if } \nu = 4, \end{cases}$$

where

$$\mathfrak{b} := \frac{2-\nu}{4-\nu} \text{ for } \nu \in (2, 4), \text{ and}$$

$$\tilde{\gamma}_\nu := 1 + \frac{4-\nu}{2(3-\nu)} \left(1 - 2^{\frac{2(3-\nu)}{4-\nu}} \right) \text{ for } \nu \in (2, 3) \cup (3, 4).$$

Proof. The proof follows from eqs. (76), (81), (84), and (86) in Appendix B. ■

A remarkable observation we can make from this Corollary is that the eigenvalue bounds number τ_h and σ_h only appear for the generalized self-concordance parameters $\nu \in (2, 4] \setminus \{3\}$.

Lemma 5.8 is key to prove the first iteration complexity bounds to estimate the number of steps needed to ensure that the local norm of the search direction is smaller than a user-defined tolerance. In the context of proximal algorithms for solving composite self-concordant minimization problems with convex data, a similar result has been established by [49]. We instead derive such a basic complexity estimate in the setting of Hessian-barrier methods for non-convex optimization problems without Lipschitz gradient assumptions and generalized self-concordant penalties.

Lemma 5.9. *Suppose Assumptions 1-4 hold. Let $(x^k)_{k \geq 0}$ be the sequence generated by $\text{HBA}(\mu, L)$. Define the stopping time*

$$\mathbf{N}(\varepsilon, x^0, \nu, L) := \min \{k \geq 0 : \lambda_k < \varepsilon\}. \quad (57)$$

Then,

$$\mathbf{N}(\varepsilon, x^0, \nu, L) \leq \left\lceil \frac{f(x^0) - f^* + \mu c}{\tilde{\omega}_\nu(\varepsilon)} \right\rceil. \quad (58)$$

Proof. By definition, for all $0 \leq k \leq \mathbf{N}(\varepsilon, x^0, \nu, L) - 1$, we have $\lambda_k \geq \varepsilon$ and, due to the strong monotonicity of the function $\tilde{\omega}_\nu$, that $\tilde{\omega}_\nu(\lambda_k) \geq \tilde{\omega}_\nu(\varepsilon)$. Therefore, using the per-iteration descent of the potential function given by

$$F_\mu(x^{k+1}) - F_\mu(x^k) \leq -\Delta_k \leq -\tilde{\omega}_\nu(\lambda_k) \leq -\tilde{\omega}_\nu(\varepsilon),$$

we readily conclude for $N > \mathbf{N}(\varepsilon, x^0, \nu, L)$,

$$f^* \leq f(x^N) \leq f(x^0) - N\tilde{\omega}_\nu(\varepsilon) + \mu c < f^*.$$

Solving for N gives the claimed bound. ■

Our second iteration complexity result gives a more precise estimate on the number of steps needed to make the local norm of the search direction as small as desired. In particular, the next estimate provides us with an easy-to-implement stopping criterion for $\text{HBA}(\mu, L)$, building on the insights gained from Lemma 5.9. Let $\varepsilon > 0$ be a target precision level, specified before the algorithm is started, and set $\mu = 4\varepsilon$. We elect to terminate $\text{HBA}(4\varepsilon, L)$ whenever $F_{4\varepsilon}(x^{K+1}) - F_{4\varepsilon}(x^K) \geq -\frac{\hat{\gamma}_\nu \varepsilon^2}{L+4\varepsilon}$ at iteration K , and report the iterate x^K . When this happens for the first time, we will show that $\lambda_K \leq \varepsilon$. If this stopping criterion is not satisfied, we continue with the execution of the protocol $\text{HBA}(4\varepsilon, L)$ until an upper bound on the number of iterations $K = O(\varepsilon^{-2})$ is reached. Implementing this stopping criterion, we therefore are guaranteed to reach a point x^K either satisfying $f(x^K) - \inf_{x \in \mathcal{X}} f(x) = f(x^K) - f^* \leq \varepsilon$, or else $\lambda_K \leq \varepsilon$. Together with this stopping criterion we see that $\text{HBA}(4\varepsilon, L)$ solves a constrained problem with potential non-differentiability at the boundary, with an iteration complexity of $O(\varepsilon^{-2})$. For this type of problem, such a rate is the best known in the literature [39]. It is also worth emphasizing that the transfer map T_μ admits a closed form expression and the complexity of computing this map is on the same order as standard Newton methods.

We are now ready to present the main complexity result for $\text{HBA}(4\varepsilon, L)$. Define the constant

$$\hat{\gamma}_\nu := \begin{cases} \tilde{\gamma}_\nu \frac{4-\nu}{\nu-2} & \text{if } \nu \in (2, 3) \cup (3, 4) \\ 1 & \text{if } \nu = 3, \\ \exp(-1) & \text{if } \nu = 4. \end{cases} \quad (59)$$

Consider the stopping time

$$\mathbf{K}_1(\varepsilon, x^0, \nu, L) := \min \left\{ k \geq 0 : \tilde{\omega}_\nu(\lambda_k) < \frac{\hat{\gamma}_\nu \varepsilon^2}{L + 4\varepsilon} \right\}. \quad (60)$$

Furthermore, let us define

$$\mathbf{K}_2(\varepsilon, x^0, \nu, L) := \left\lceil \frac{(4\varepsilon(c - 1/4) + f(x^0) - f^*)(L + 4\varepsilon)}{\hat{\gamma}_\nu \varepsilon^2} \right\rceil \quad (61)$$

and $\mathbf{K}_{\max}(\varepsilon, x^0, \nu, L) = \min\{\mathbf{K}_1(\varepsilon, x^0, \nu, L), \mathbf{K}_2(\varepsilon, x^0, \nu, L)\}$.

Theorem 5.10. *Let $(x^k)_{k \geq 0}$ be generated by $\text{HBA}(4\varepsilon, L)$. Either the algorithm stops at $k_{\max} = \mathbf{K}_2(\varepsilon, x^0, \nu, L)$, and reaches a point $x^{k_{\max}} \in \mathcal{X}^\circ$ satisfying*

$$f(x^{k_{\max}}) - f^* \leq \varepsilon,$$

or else, it stops after $k_{\max} = \mathbf{K}_1(\varepsilon, x^0, \nu, L)$ iterations, and we reach a point where $\lambda_{k_{\max}} \leq \varepsilon$.

Proof. If $k_{\max} = \mathbf{K}_2(\varepsilon, x^0, \nu, L) \leq \mathbf{K}_1(\varepsilon, x^0, \nu, L)$, then for all $0 \leq k \leq k_{\max} - 1$ we have by definition $\tilde{\omega}_\nu(\lambda_k) \geq \hat{\gamma}_\nu \frac{\varepsilon^2}{L+4\varepsilon}$. Hence,

$$f^* \leq f(x^{k_{\max}}) \leq f(x^0) - k_{\max} \hat{\gamma}_\nu \frac{\varepsilon^2}{L+4\varepsilon} + 4\varepsilon c \leq f^* + \varepsilon.$$

For the second claim, suppose that $k_{\max} = \mathbf{K}_1(\varepsilon, x^0, \nu, L) < \mathbf{K}_2(\varepsilon, x^0, \nu, L)$, i.e. the algorithm stops before the objective function value is within ε of the global minimal value. Since $\lambda_k \rightarrow 0$, for all $\sigma > 0$ there exists a k such that $\lambda_k < \sigma$. Hence, let us fix a sufficiently small tolerance level $\varepsilon \in (0, 1)$ so that the function $\tilde{\omega}_\nu$ is determined by terms including λ_k^2 . Specifically, the following computations can be made for each generalized self-concordant parameter ν : When $\nu \in (2, 3) \cup (3, 4)$ we see $\tilde{\omega}_\nu(\lambda_k) = \frac{\tilde{\gamma}_\nu}{-b(L+4\varepsilon)} \lambda_k^2$ for k large enough, which smaller than $\frac{\tilde{\gamma}_\nu}{-b} \frac{\varepsilon^2}{L+4\varepsilon}$ exactly if $\lambda_k < \varepsilon$. For $\nu \in \{3, 4\}$ the same reasoning applies, proving the claim. ■

Remark 5.1. Evaluating the stopping criterion \mathbf{K}_1 appears to be expensive, since we have to keep track of the local norm of the search direction λ_k . However, since $\tilde{\omega}_\nu$ is monotone, we can replace λ_k with the more conservative figure $\sqrt{\tau_h} \beta_k$. Hence, if a bound on τ_h is available, we only have to monitor the evolution of the Euclidean length of the search direction.

While the above result is formulated in terms of convergence to stationary points of the potential function, our aim is actually to approximately solve the optimization problem (P). In order to connect these two conditions, we rely on our characterization of ε -KKT points. Using the relation (51), we see

$$\|\nabla f(x^k) - A^\top y^k\|_{x^k}^* = \|\sqrt{H(x^k)^{-1}}(\nabla f(x^k) - A^\top y^k)\|_2 \geq \frac{1}{\sqrt{\tau_h}} \|\nabla f(x^k) - A^\top y^k\|_2.$$

Therefore,

$$\begin{aligned} \|\nabla f(x^k) - A^\top y^k\|_2 &\leq \sqrt{\tau_h} \|\nabla f(x^k) - A^\top y^k\|_{x^k}^* \\ &\leq \sqrt{\tau_h} (\|\nabla F_\mu(x^k) - A^\top y^k\|_{x^k}^* + \mu \|\nabla h(x)\|_{x^k}^*) \end{aligned}$$

Recall that $\|\nabla F_\mu(x^k) - A^\top y^k\|_{x^k}^* = \lambda_k$. Furthermore, we know that $(x^k)_{k \geq 0} \subseteq \mathcal{S}_\mu(x^0)$, a compact set in \mathcal{X}° . Since $h \in \mathbf{C}^3(\mathcal{C})$ and \mathcal{C} contains no lines, the mapping, the norm $x \mapsto \|\cdot\|_x^*$ is a continuous function on compact subsets of \mathcal{C} . Hence, the quantity

$$M_\mu(x^0) := \max_{x \in \mathcal{S}_\mu(x^0)} \|\nabla h(x)\|_x^*,$$

is well-defined and finite. In terms of this quantity we see that

$$\|\nabla f(x^k) - A^\top y^k\|_2 \leq \sqrt{\tau_h} (\lambda_k + \mu M_\mu(x^0)), \quad (62)$$

so that for $k \geq \mathbf{N}(\varepsilon/\sqrt{\tau_h}, x^0, \nu, L)$, and $\mu = \varepsilon/\sqrt{\tau_h}$ we get

$$\|\nabla f(x^k) - A^\top y^k\|_2 = \mathcal{O}(\varepsilon).$$

Combined with the inequalities (25) we therefore conclude that $\chi(x^k, y^k) = \mathcal{O}(\varepsilon)$. i.e. we get an ε -stationary point in the sense of Definition 2.11. Note that $M_\mu(x^0)$ is an algorithm independent constant, which can be computed before the method is started. Still it requires the solution of an optimization problem which can be fairly complicated in concrete instances, so it is definitely worthwhile searching for settings where this bound can be improved. Additionally, the complexity of the algorithm now explicitly depends on the eigenvalue bound τ_h of the barrier-generating kernel h , which means that if this number is big, the run time could become quite large.³ Motivated by these observations, we next provide a refinement of this complexity result under the additional assumption that $h \in \mathcal{F}_{2,3}(\mathcal{C})$ is a θ -self-concordant barrier in the sense of (20).

Corollary 5.11. *Let $\varepsilon > 0$ be a given tolerance level. If $h \in \mathcal{F}_{2,3}(\mathcal{C})$ is a θ -SCB, then running $\text{HBA}(\varepsilon/\sqrt{\theta}, L)$ yields either an $2\sqrt{\tau_h}\varepsilon$ -stationary point, or an ε global minimum.*

Proof. For a θ -SCB $h \in \mathcal{F}_{2,3}(\mathcal{C})$, the complexity estimate in Theorem 5.10 yields the following estimates: If

$$k_{\max} = \mathbf{K}_1(\varepsilon, x^0, 3, L) = \min \left\{ k \geq 0 : \min \left\{ \lambda_k, \frac{\lambda_k^2}{L + \varepsilon/\sqrt{\theta}} \right\} \leq \frac{\varepsilon^2}{L + \varepsilon/\sqrt{\theta}} \right\},$$

then we know that the local norm of the gradient of the potential function is small, $\lambda_{k_{\max}} \leq \varepsilon$. Then (62) gives us

$$\|\nabla f(x^{k_{\max}}) - A^\top y^{k_{\max}}\|_2 \leq \sqrt{\tau_h} \left(\lambda_{k_{\max}} + \mu\sqrt{\theta} \right).$$

Choosing $\mu = \varepsilon/\sqrt{\theta}$, and using again the relation (25), the point $x^{k_{\max}}$ is seen to be a $2\sqrt{\tau_h}\varepsilon$ stationary point, in the sense of Definition 2.11. If instead $k_{\max} = \mathbf{K}_2(\varepsilon, x^0, 3, L)$, we know we are ε -close to the global minimum. ■

5.3 Analysis of AHBA(μ)

The analysis of the adaptive version of our method follows similar lines as for the mother scheme $\text{HBA}(\mu, L)$. The key innovation of the adaptive method is that it produces four recursive sequences $(x^k)_{k \geq 0}$, $(y^k)_{k \geq 0}$, $(\alpha_k)_{k \geq 0}$, and $(L_k)_{k \geq 0}$, where $\alpha_k = \alpha_\mu(x^k, L_k)$. We first show finite termination of the line search subroutine at each iteration, and establish a bound on the total number of function evaluations needed for its execution. The result is a generalization of the arguments in [7, 44] for the case of relative smoothness in the non-convex case.

Lemma 5.12. *Suppose that we run AHBA(μ) for $N \geq 1$ rounds. Then, the total number of function evaluations \mathcal{E}_N , needed to satisfy (41) in each of these $k = 1, 2, \dots, N$ rounds, is at most*

$$\mathcal{E}_N \leq 2N + \log_2 \left(\frac{2L}{L_0} \right). \quad (63)$$

Proof. Let $k = 1, 2, \dots, N$ be an arbitrary iteration count. It is quite easy to see that the search cycle for i_k is finite. Indeed since, by Definition 3.3, there exists such L that for any $x, y \in \mathcal{C}$

$$f(y) \leq f(x) + \langle \nabla f(x), y - x \rangle + LD_h(y, x),$$

³However, both these remarks hold also for mirror descent type of methods, where the prox-function should be appropriately chosen since its properties affect the complexity bound

the search cycle for i_k terminates no later than $L_{k+1} = 2^{i_k-1}L_k \geq L$. At the same time, since $i_k \geq 0$ is the smallest integer for which (41) holds, we have for $L_{k+1}/2 = 2^{i_k-2}L_k$ the inequality

$$f(z^k) > f(x^k) + \langle \nabla f(x^k), z^k - x^k \rangle + 2^{i_k-2}L_k D_h(z^k, x^k).$$

Hence, $2^{i_k-2}L_k < L$, or $L_{k+1} = 2^{i_k-1}L_k \leq 2L$. Let us estimate the total number of function evaluations needed to ensure (41). On each iteration k , the number of function calls is $i_k + 1 = 2 + \log_2 \frac{L_{k+1}}{L_k}$. Thus, the total number of function evaluations for N rounds of execution of AHBA(μ, L) is thus

$$\mathcal{E}_N = \sum_{k=1}^N (i_k + 1) = \sum_{k=1}^N \left(2 + \log_2 \frac{L_{k+1}}{L_k} \right) \leq 2N + \log_2 \frac{2L}{L_0},$$

where we used the bound $L_k \leq 2L$. ■

This shows that $\mathcal{E}_N = \mathcal{O}(N)$, meaning that on average only a single function call is needed to satisfy the line search criterion (41). Thus, the performance of AHBA(μ) is well described by the estimates for the overall iteration complexity of the method.

From the analysis of the base scheme HBA(μ, L), we immediately deduce that the sequence $(x^k)_{k \geq 0}$ generated by AHBA(μ) satisfy the per-iteration descent

$$F_\mu(x^{k+1}) \leq F_\mu(x^k) - \eta_k(x^k, \alpha_k) \equiv F_\mu(x^k) - \Delta_k,$$

with the only difference that now the step size α_k is adaptively adjusted by evaluating the expression $\alpha_\mu(x^k, 2^{i_k-1}L_k)$. From Lemma 5.8, we see that $\tilde{\omega}_\nu$ is a decreasing function of L . At the same time, $2^{i_k-1}L_k \leq 2L$ as it was shown above. This means that the adaptive versions of Lemma 5.9, Theorem 5.10, and Corollary 5.11 are obtained by the change $L \rightarrow 2L$. We see that the number of oracle calls increases for the adaptive version in comparison to non-adaptive. Nevertheless, the adaptive algorithm can use smaller values of L and, hence, make longer steps, leading to faster convergence in practice.

6 Numerical Results

6.1 Statistical learning with non-convex regularization

We consider the non-convex statistical learning problem

$$\min_{\beta \in \mathbb{R}^d} \frac{1}{2} \|y - W\beta\|_2^2 + \sum_{i=1}^d p_\zeta(|\beta_i|) \quad (64)$$

where $\ell(\beta) := \frac{1}{2} \|y - W\beta\|_2^2$ is the quadratic data fitting term and $p_\zeta : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a *folded concave penalty* [31, 32, 34, 53], meaning that for given $a > 2, \zeta > 0$:

- (i) $t \mapsto p_\zeta(t)$ is non-decreasing and concave with $p_\zeta(0) = 0$ and $p_\zeta(t) > 0$ for $t > 0$;
- (ii) $t \mapsto p_\zeta(t)$ is differentiable on $[0, \infty)$;
- (iii) $p'_\zeta(t) = 0$ for all $t \geq a\zeta$ and $0 \leq p'_\zeta(t)$ for any $t \geq 0$.

A specific example would be *smoothly clipped absolute deviation* (SCAD) penalty [22] given by

$$p_\zeta(t) = \begin{cases} \zeta t & \text{if } 0 \leq t \leq \zeta, \\ \frac{1}{a-1} \left(-\frac{\zeta^2}{2} + a\zeta t - \frac{t^2}{2} \right) & \text{if } \zeta \leq t \leq a\zeta, \\ \frac{a+1}{2} \zeta^2 & \text{if } t > a\zeta. \end{cases}$$

Note that the composite function $t \mapsto (p_\zeta \circ |\cdot|)(t)$ is continuous, but not differentiable at $t = 0$. Hence, the objective function (64) is not smooth and non-convex. Doing some simple variable transformations, the regularized least-squares problem (64) can be put into an optimization problem fitting the structure of this paper. Let us introduce new variables $\beta_i^+ := \max\{\beta_i, 0\}$ and $\beta_i^- := \max\{-\beta_i, 0\}$, so that $\beta_i^+ + \beta_i^- = |\beta_i|$. We additionally allow the inclusion of a-priori upper bounds on the parameter vector. This gives rise to a box-constrained reformulation of (64) of the form

$$\begin{aligned} & \min_{\beta^+ \in \mathbb{R}^d, \beta^- \in \mathbb{R}^d} \ell(\beta^+ - \beta^-) + \sum_{i=1}^d p_\zeta(\beta_i^+ + \beta_i^-), \\ & \text{s.t. } 0 \leq \beta_i^- \leq u_i, \quad 0 \leq \beta_i^+ \leq u_i \quad 1 \leq i \leq d. \end{aligned}$$

To bring this problem into a formulation fitting this paper, we first relabel the pair $(\beta^-, \beta^+) \in \mathbb{R}^d \times \mathbb{R}^d$ into one long vector $x := (x_1, \dots, x_d, x_{d+1}, \dots, x_{2d})$, where the first d variables correspond to the positive part and the remaining d variables represent the negative part. Call $n := 2d$ we define the data fitting term to be $f_0(x) := \ell(Bx)$, where $B : \mathbb{R}^n \rightarrow \mathbb{R}^d$ is the linear operator $(Bx)_i := x_{d+i} - x_i$ for all $i \in \{1, 2, \dots, d\}$. The regularizer can be written as $f_1(Dx) := \sum_{i=1}^d p_\zeta(x_i + x_{d+i})$, corresponding to the composition of the function $\mathbb{R}^d \ni y \mapsto f_1(y) = \sum_{i=1}^d p_\zeta(y_i)$ with the linear operator $D : \mathbb{R}^n \rightarrow \mathbb{R}^d$ given by $(Dx)_i = x_i + x_{d+i}$ for all $i \in \{1, 2, \dots, d\}$. Define $\mathcal{X} = \bar{\mathcal{C}} := \prod_{i=1}^n [0, u_i]$, so that our non-convex minimization problem reads as

$$\min_{x \in \mathcal{X}} \{f(x) := f_0(x) + f_1(Dx)\}, \quad (65)$$

where $f_0(x) := \frac{1}{2}x^\top Qx + x^\top q$ is a convex quadratic function with Hessian $Q := B^\top W^\top W B$ and $q^\top := -B^\top W^\top y$. Note that $\bar{\mathcal{C}}$ admits a simple self-concordant function (e.g. the Burg entropy as described in Example 2.4), but is not prox-friendly (see Remark 2.2).

The quadratic loss function $f_0(x)$ is convex and has a Lipschitz continuous gradient with Lipschitz constant $\rho := |Q|$. Rescaling the data appropriately, we can assume without loss of generality that $\rho \geq 1$. Hence, for the data fidelity part, a standard Lipschitz-descent lemma [39] applies and gives

$$f_0(y) \leq f_0(x) + \langle \nabla f_0(x), y - x \rangle + \frac{\rho}{2} \|y - x\|_2^2. \quad (66)$$

For $\theta \in \mathbb{R}_+^d$, the regularizing term reads as $f_1(\theta) = \sum_{i=1}^d p_\zeta(\theta_i)$, and each summand in this expression is a concave and differentiable function on $(0, \infty)$. Hence, for all $s, t > 0$, we have

$$p_\zeta(s) \leq p_\zeta(t) + p'_\zeta(t)(s - t).$$

For any two vectors $\theta^{(1)}, \theta^{(2)} \in \mathbb{R}_{++}^d$ this implies that

$$f_1(\theta^{(2)}) \leq f_1(\theta^{(1)}) + \langle \nabla f_1(\theta^{(1)}), \theta^{(2)} - \theta^{(1)} \rangle.$$

Evaluating this expression at the vectors $\theta^{(1)} = Dx$ and $\theta^{(2)} = Dy$, we obtain

$$f_1(Dy) \leq f_1(Dx) + \langle D^\top \nabla f_1(Dx), y - x \rangle. \quad (67)$$

Adding (66) with (67), we see that

$$f(y) \leq f(x) + \langle \nabla f(x), y - x \rangle + \frac{\rho}{2} \|y - x\|_2^2 \quad \forall y, x \in \mathcal{X}^\circ. \quad (68)$$

For the rest of the analysis we assume that $u_i = \infty$, so that no external upper bounds on the parameter vectors are imposed. Thus, $\mathcal{X} = \bar{\mathcal{C}} = \mathbb{R}_+^n$, and the natural barrier-generating kernel for this set is the Burg entropy $h(x) = -\sum_{i=1}^n \ln(x_i)$, inducing the Riemannian metric $H(x) = \text{diag}\{x_1^{-2}, \dots, x_n^{-2}\}$, and Bregman divergence

$$D_h(y, x) = h(y) - h(x) - \langle \nabla h(x), y - x \rangle = \sum_{i=1}^n \ln\left(\frac{y_i}{x_i}\right) + \sum_{i=1}^n \frac{y_i}{x_i} - n.$$

In terms of the potential function $F_\mu(x) = f(x) + \mu h(x)$, the combined descent inequality (68) reads as

$$F_\mu(y) \leq F_\mu(x) + \langle \nabla F_\mu(x), y - x \rangle + \frac{\rho}{2} \|y - x\|_2^2 + \mu D_h(y, x).$$

Defining the regularized Burg entropy $\tilde{h}_{\mu, \rho} := -\sum_{i=1}^n \ln(x_i) + \frac{\rho}{2\mu} \|x\|_2^2$, we can write the descent inequality for the potential function in more concise terms as

$$F_\mu(y) \leq F_\mu(x) + \langle \nabla F_\mu(x), y - x \rangle + \mu D_{\tilde{h}_{\mu, \rho}}(y, x).$$

Note that $\tilde{h}_{\mu, \rho} \in \mathcal{F}_{2,3}(\mathbb{R}_{++}^n)$. This shows that the regularized statistical learning problem can be solved with HBA($\mu, 0$). We apply the model to the Prostate Cancer data set available at <https://web.stanford.edu/~hastie/ElemStatLearn/data.html>. This data set consists of a total of 97 samples with 8 dimensions each, from which 67 are used to train the model and 30 are used for validation. Thus, in this case we have a matrix $W \in \mathbb{R}^{67 \times 8}$ and $y \in \mathbb{R}^{67}$. Moreover, we have used the following set of parameter values $\zeta = 0.01$, $a = 10$, and $\mu = 1 \cdot 10^{-3}$. Once a model $\hat{\beta}$ is found, such value is used to predict a output $A_{\text{test}\hat{\beta}}$ for the test database A_{test} . Figure 2(a) shows the true output of the test database for each of the 30 data points colored in black, and the predicted output for the same points generated by the output of the Adaptive HBA algorithm. Moreover, Figure 2(b) shows the gradient norm value versus the number of iterations of the algorithm, and the test error in color red. AHBA(μ) reaches a test error of 0.363. This value improves upon the 0.4194 test error reported in [6], and the 0.479 test error reported in [23, Table 3.3].

6.2 L^p -minimization

Consider the optimization problem

$$\begin{aligned} \min \quad & f(x) = \sum_{i=1}^n x_i^p \\ \text{subject to} \quad & x \in \mathcal{X} = \mathbb{R}_+^n \cap \mathcal{A} \end{aligned} \quad (69)$$

where the problem inputs consist of $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$ and $p \in (0, 1]$. Sparse signal or solution reconstruction by solving problem (69), especially for the case where $p \in (0, 1)$, has recently received considerable attention; see e.g. [12]. In signal reconstruction, one typically has linear measurements $b = Ax$, where x is a sparse signal, i.e. the sparsest or smallest support cardinality solution of the linear system. This sparse signal is recovered by solving the inverse problem (69) with the non-smooth, non-convex objective function $\|x\|_0 = |\{i \in \{1, 2, \dots, n\} | x_i > 0\}|$. The L_0 -norm optimization problem is shown to be NP-hard. When $p = 1$, the problem is reduced to a linear program, and hence it can be solved in polynomial time. If $p > 1$, the problem (69) becomes




Figure 2: Fitted values and Test error of $HBA(\mu, 0)$ for the Prostate Cancer data with SCAD regularization. As parameters for the SCAD regularizer we have used $\zeta = 0.01$, $a = 10$. For the barrier parameter the value $\mu = 10^{-3}$ has been chosen.

a convex optimization problem, and thus is also efficiently solvable with fast interior point methods. Only recently, the challenging case where $p \in (0, 1)$ has been thoroughly investigated in [6, 24]. We aim to solve this NP-hard problem with $AHBA(\mu)$. Given the geometry, it is natural to look at the barrier-generating kernel $h(x) = -\sum_{j=1}^n \ln(x_j)$, so that the potential function F_μ becomes $F_\mu(x) = \|x\|_p^p - \mu \sum_{i=1}^n \ln(x_i)$. Note that the objective function $f(x) = \|x\|_p^p$ is twice continuously differentiable on \mathcal{X}° and concave. Hence, L -smoothness holds for any $L > 0$.

To test the performance of our method, we have set up numerical experiments and recorded the recovery rate of the true underlying signal for each level of sparsity. Specifically, we generate a binary signal of length 120, denoted as \hat{x} , and various sparsity patterns. The excellent recovery properties of $AHBA(\mu)$ with 5 non-zero entries is displayed in Figure 3(a) and with 10 non-zero entries in Figure 3(b), in which the original signal is marked as black circles \circ and the recovered one is marked as red crosses \times . In each case, we generated an observation matrix as an orthogonal positive sensing matrix A , and a set of 30 observations. Moreover we have used as parameter values $p = 0.5$ and $\mu = 1$. Figure 3(c) reveals the general pattern of the recovery rates of the true signal.

7 Conclusion

In this paper we introduced and studied a new class of interior point methods based on the Hessian-barrier technique originally developed in [11]. Using the theory of generalized self-concordant functions we can significantly extend the applicability of this method to cover general non-convex optimization problems on polyhedral domains with a set constraint admitting a generalized self-concordance set-up. Theoretical convergence and complexity results are proven, showing that the method achieves the optimal iteration complexity $O(\varepsilon^{-2})$. We have tested the method empirically and verified that the method performs also well in practice. There are many important directions for future investigations to be made. First, it is very important to relax the present algorithmic scheme to allow for inexact computations and to derive a path-following approach allowing the barrier parameter μ to vary over the run time of the algorithm. Allowing for numerical and random noise is of relevance when HBA methods are to be designed in distributed optimization settings [52, 56], something we plan to do in the future, and the path-following approach might allow us to strengthen the convergence properties

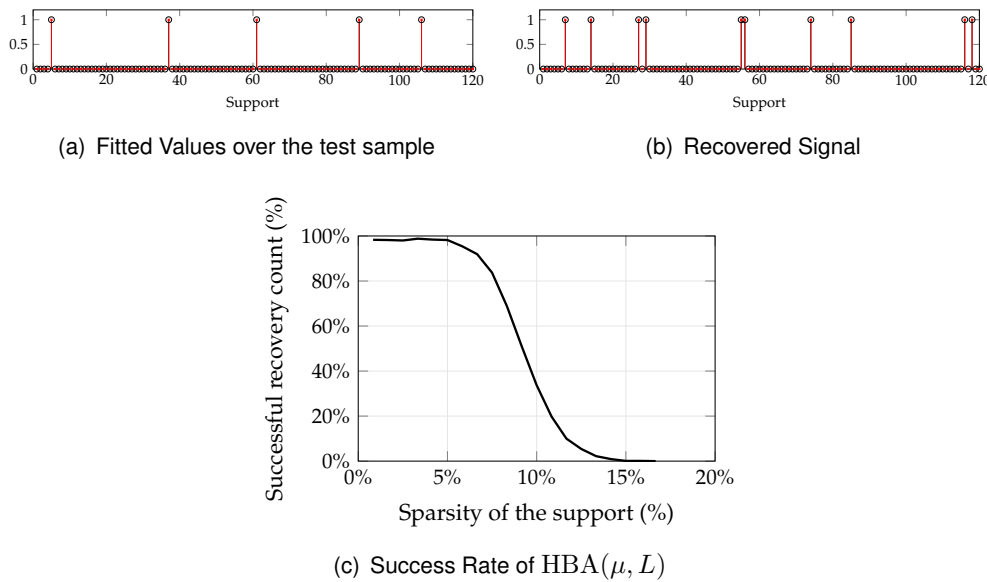


Figure 3: Recovery pattern and recovery rates for the L^p minimization problem solved by $AHBA(\mu)$. As parameter values we have chosen $p = 0.5$ and $\mu = 1$.

of the algorithm.

The most costly step of $HBA(\mu, L)$, and its adaptive version $AHBA(\mu)$, is the solution of the linear system (27). If the function f appears as a finite sum, a very important direction for future research is to either use preconditioning techniques or randomization and sketching ideas, to speed up the computation. Second, the method should be also a competitive first-order scheme for stochastic optimization. In fact, self-concordant functions have been already successfully used in online learning [38] and random sampling [37], and the class of generalized self-concordant functions may provide interesting extensions of these seminal contributions. Finally, it will be important to identify acceleration strategies for the basic HBA template to make it even more attractive for large-scale application in engineering and machine learning. This is another challenging line of research we plan to pursue in the near future.

Acknowledgments.

M. Staudigl would like to thank Panayotis Mertikopoulos for years of fruitful collaborations, and extensive feedback on this paper. The research of M. Staudigl has been supported by the COST Action CA16228 "European Network for Game Theory". The work of C.A. Uribe was partially supported by Yahoo! Research Faculty Engagement Program.

A Proof of Proposition 3.5

For each parameter $\nu \in (2, 4]$ we derive the corresponding optimal step-size policy via a simple optimization argument. This will prove the claimed optimality of the policy.

A.1 The case $\nu \in (2, 3)$

Using the definition

$$\omega_\nu(t) = \left(\frac{\nu - 2}{4 - \nu} \right) \frac{1}{t} \left[\frac{\nu - 2}{2t(3 - \nu)} \left((1 - t)^{\frac{2(3-\nu)}{2-\nu}} - 1 \right) - 1 \right],$$

we obtain

$$\begin{aligned} \eta_\mu(x, t) &= t\lambda_\mu^2(x) - t^2\lambda_\mu^2(x)(L + \mu) \frac{(\nu - 2)^2}{2t^2\delta_\mu(x)^2(3 - \nu)(4 - \nu)} \left((1 - t\delta_\mu(x))^{\frac{2(3-\nu)}{2-\nu}} - 1 \right) \\ &\quad + t \frac{\lambda_\mu^2(x)}{\delta_\mu(x)} (L + \mu) \frac{\nu - 2}{4 - \nu} \\ &= t \left(\lambda_\mu^2(x) + \lambda_\mu^2(x) \frac{L + \mu \nu - 2}{\delta_\mu(x) (4 - \nu)} \right) \\ &\quad - \left(\frac{\lambda_\mu(x)}{\delta_\mu(x)} \right)^2 (L + \mu) \frac{(\nu - 2)^2}{2(3 - \nu)(4 - \nu)} \left((1 - t\delta_\mu(x))^{\frac{2(3-\nu)}{2-\nu}} - 1 \right). \end{aligned}$$

For $t\delta_\mu(x) \in (0, 1)$, this gives

$$\begin{aligned} \frac{\partial}{\partial t} \eta_\mu(x, t) &= \lambda_\mu^2(x) \left(1 + \frac{L + \mu \nu - 2}{\delta_\mu(x) (4 - \nu)} \right) - \frac{\lambda_\mu^2(x) (L + \mu)(\nu - 2)}{\delta_\mu(x) (4 - \nu)} (1 - t\delta_\mu(x))^{-\frac{4-\nu}{\nu-2}}, \\ \frac{\partial^2}{\partial t \partial t} \eta_\mu(x, t) &= -\lambda_\mu^2(x) (L + \mu) (1 - t\delta_\mu(x))^{-\frac{2}{2-\nu}} < 0. \end{aligned}$$

Solving the stationarity condition $\frac{\partial}{\partial t} \eta_\mu(x, t) \big|_{t=\alpha_\mu(x, L)} = 0$, gives

$$\alpha_\mu(x, L) \delta_\mu(x) = 1 - \left(1 + \frac{\delta_\mu(x) (4 - \nu)}{L + \mu \nu - 2} \right)^{-\frac{\nu-2}{4-\nu}}. \quad (70)$$

Since $\frac{\nu-2}{4-\nu} \in (0, 1)$ for $\nu \in (2, 3)$, the Bernoulli inequality gives

$$\left(1 + \frac{\delta_\mu(x) (4 - \nu)}{L + \mu \nu - 2} \right)^{\frac{\nu-2}{4-\nu}} \leq 1 + \frac{\delta_\mu(x)}{L + \mu},$$

so that, for $L \geq 0$,

$$1 - \left(1 + \frac{\delta_\mu(x) (4 - \nu)}{L + \mu \nu - 2} \right)^{-\frac{\nu-2}{4-\nu}} \leq \frac{\delta_\mu(x)}{\delta_\mu(x) + (L + \mu)} < 1.$$

Hence, setting

$$\alpha_\mu(x, L) = \frac{1}{\delta_\mu(x)} \left[1 - \left(1 + \frac{\delta_\mu(x) (4 - \nu)}{L + \mu \nu - 2} \right)^{-\frac{\nu-2}{4-\nu}} \right] \quad (71)$$

gives $\eta_\mu(x, \alpha_\mu(x, L)) > \eta_\mu(x, 0) = 0$ and $\mathbf{d}_\nu(x, T_\mu(x, \alpha_\mu(x, L))) < 1$.

A.2 The case $\nu = 3$.

We have $\omega_3(t) = \frac{-1}{t^2}(t + \ln(1-t))$ for $t \in (-\infty, 1)$, and $d_3(x, y) = \frac{M}{2}\|y - x\|_x$. Hence, $\delta_\mu(x) = \frac{M}{2}\lambda_\mu(x)$, and

$$\eta_\mu(x, t) = t\lambda_\mu^2(x) + \left(\frac{\lambda_\mu(x)}{\delta_\mu(x)}\right)^2 (L + \mu) [t\delta_\mu(x) + \ln(1 - t\delta_\mu(x))].$$

Therefore,

$$\begin{aligned} \frac{\partial}{\partial t}\eta_\mu(x, t) &= \lambda_\mu(x)^2 + \left(\frac{\lambda_\mu(x)}{\delta_\mu(x)}\right)^2 (L + \mu) \left[\delta_\mu(x) - \frac{\delta_\mu(x)}{1 - t\delta_\mu(x)}\right] \\ \frac{\partial^2}{\partial t^2}\eta_\mu(x, t) &= -\frac{\lambda_\mu(x)^2}{\delta_\mu(x)}(L + \mu)(1 - t\delta_\mu(x))^{-2} < 0 \end{aligned}$$

Solving for the stationary condition $\frac{\partial}{\partial t}|_{t=\alpha_\mu(x)}\eta_\mu(x, t) = 0$, we get

$$1 - \alpha_\mu(x)\delta_\mu(x) = t(L + \mu).$$

Hence, setting

$$\alpha_\mu(x) = \frac{1}{\delta_\mu(x) + L + \mu}, \tag{72}$$

we observe that $\eta_\mu(x, \alpha_\mu(x, L)) > \eta_\mu(x, 0) = 0$, and $\alpha_\mu(x, L)\delta_\mu(x) < 1$. Therefore, $d_3(x, T_\mu(x, \alpha_\mu(x, L))) < 1$.

A.3 The case $\nu = 4$.

We have $\omega_4(t) = \frac{(1-t)\ln(1-t)+t}{t^2}$, and $d_4(x, y) = M\|y - x\|_2^{-1}\|y - x\|_x^2$. Hence, $\delta_\mu(x) = M\frac{\lambda_\mu^2(x)}{\beta_\mu(x)}$, and

$$\eta_\mu(x, t) = t\lambda_\mu(x)^2 - \left(\frac{\lambda_\mu(x)}{\delta_\mu(x)}\right)^2 (L + \mu)[t\delta_\mu(x) + (1 - t\delta_\mu(x))\ln(1 - t\delta_\mu(x))].$$

Therefore,

$$\begin{aligned} \frac{\partial}{\partial t}\eta_\mu(x, t) &= \lambda_\mu(x)^2 - \left(\frac{\lambda_\mu(x)}{\delta_\mu(x)}\right)^2 (L + \mu)\ln(1 - t\delta_\mu(x)), \\ \frac{\partial^2}{\partial t^2}\eta_\mu(x, t) &= -\frac{\lambda_\mu^2(x)(L + \mu)}{1 - t\delta_\mu(x)} < 0. \end{aligned}$$

Solving for stationarity $\frac{\partial}{\partial t}|_{t=\alpha_\mu(x, L)}\eta_\mu(x, t) = 0$, gives

$$\frac{-\delta_\mu(x)}{L + \mu} = \ln(1 - \alpha_\mu(x, L)\delta_\mu(x, L)),$$

so that

$$\alpha_\mu(x, L) = \frac{1}{\delta_\mu(x, L)} \left[1 - \exp\left(-\frac{\delta_\mu(x)}{L + \mu}\right)\right]. \tag{73}$$

It follows $d_4(x, T_\mu(x, \alpha_\mu(x, L))) = 1 - \exp(-\delta_\mu(x)/(L + \mu)) \in (0, 1)$.

A.4 The case $\nu \in (3, 4)$.

The basic computations for this range can be copied from the case $\nu \in (2, 3)$. Doing so, we immediately arrive at the step size policy

$$\alpha_\mu(x, L)\delta_\mu(x) = 1 - \left(1 + \frac{\delta_\mu(x) (4 - \nu)}{L + \mu \nu - 2}\right)^{-\frac{\nu-2}{4-\nu}}. \quad (74)$$

From here, we can continue all the computations as for the case $\nu \in (2, 3)$ to conclude that the step size $\alpha_\mu(x, L)$ is given by (71). Note that $\frac{2-\nu}{4-\nu} < 0$, so that $(1 + \frac{\delta_\mu(x) (4-\nu)}{L+\mu \nu-2})^{\frac{2-\nu}{4-\nu}} \in (0, 1)$, and therefore $\alpha_\nu(x)\delta_\mu(x) \in (0, 1)$. All other conclusions derived for $\nu \in (2, 3)$ apply to the present setting as well.

B Proof of Theorem 5.6

We denote by $\Delta_k \equiv \eta_\mu(x^k, \alpha_k)$, where x^k is the iterate of HBA(μ, L), and $\alpha_k \equiv \alpha_\mu(x^k, L)$ is the associated step size. Similarly, we define the sequence λ_k, β_k and δ_k as in (42).

B.1 The case $\nu \in (2, 3)$

An explicit calculation shows that

$$\begin{aligned} \Delta_k &= \frac{\lambda_k^2}{\delta_k} \left[1 - \frac{4 - \nu}{2(3 - \nu)} \left(1 + \frac{\delta_k (4 - \nu)}{L + \mu \nu - 2} \right)^{(2-\nu)/(4-\nu)} \right] \\ &\quad + \left(\frac{\lambda_k}{\delta_k} \right)^2 \frac{(\nu - 2)(L + \mu)}{2(3 - \nu)} \left[1 - \left(1 + \frac{\delta_k (4 - \nu)}{L + \mu \nu - 2} \right)^{(2-\nu)/(4-\nu)} \right]. \end{aligned}$$

To make the analysis of this expression more convenient, we introduce the quantities

$$\begin{aligned} t_k &:= 1 - \frac{1}{\mathbf{b}} \frac{\delta_k}{L + \mu} \in (1, +\infty), \text{ and} \\ \mathbf{a} &:= \frac{4 - \nu}{2(3 - \nu)} \in (1, +\infty), \quad \mathbf{b} := \frac{2 - \nu}{4 - \nu} \in (-1, 0). \end{aligned}$$

Then $\frac{(\nu-2)(L+\mu)}{\delta_k} = \frac{(4-\nu)}{t_k-1}$ and

$$\begin{aligned} \Delta_k &= \frac{\lambda_k^2}{\delta_k} \left(1 - \mathbf{a} t_k^{\mathbf{b}} + \frac{\mathbf{a}}{t_k - 1} (1 - t_k^{\mathbf{b}}) \right) \\ &= \frac{\lambda_k^2}{\delta_k} \left(1 + \frac{\mathbf{a}}{t_k - 1} - \mathbf{a} t_k^{\mathbf{b}} \left(1 + \frac{1}{t_k - 1} \right) \right) \\ &= \frac{\lambda_k^2}{\delta_k} \left(1 + \frac{\mathbf{a}}{t_k - 1} - \frac{\mathbf{a} t_k^{\mathbf{b}+1}}{t_k - 1} \right). \end{aligned}$$

Let us define a function $\gamma(t)$ such that $\Delta_k = \frac{\lambda_k^2}{\delta_k} \gamma(t_k)$. Our next goal is to show that, for $t \in [2, +\infty)$, $\gamma(t)$ is below bounded by some positive constant and, for $t \in (1, 2]$, $\gamma(t)$ is below bounded by some positive constant multiplied by $t - 1$.

1. $t \in [2, +\infty)$. We will show that $\gamma'(t) \geq 0$, whence $\gamma(t) \geq \gamma(2)$. Thus, we need to show that

$$0 \leq \gamma'(t) = -\frac{a}{(t-1)^2} \underbrace{\left(1 - (b+1)t^b + bt^{b+1}\right)}_{=: \psi(t)}.$$

Since $a > 1$, to show that $\gamma'(t) \geq 0$ it is enough to show that $\psi(t) \leq 0$. Since $b \in (-1, 0)$ and $t \geq 2$,

$$\psi'(t) = b(b+1)t^b - b(b+1)t^{b-1} = b(b+1)t^{b-1}(t-1) \leq 0.$$

Whence, $\psi(t) \leq \psi(2)$ for all $t \in [2, +\infty)$. It remains to show that $\psi(2) \leq 0$. Let us consider $\psi(2) = \varphi(b) := 1 - (b+1)2^b + b2^{b+1} = 1 + b2^b - 2^b$ as a function of $b \in (-1, 0)$. Clearly, $\varphi(-1) = \varphi(0) = 0$, and it is easy to check via the intermediate value theorem that $\varphi(b) < 0$ for all $b \in (-1, 0)$. We conclude that for $t \geq 2$ we get $\Delta_k \geq \frac{\lambda_k^2}{\delta_k} \gamma(2)$.

2. $t \in (1, 2]$. We will show that $\frac{d}{dt}(\gamma(t)/(t-1)) \leq 0$, whence $\gamma(t) \geq (t-1)\gamma(2)$. Thus, we need to show that

$$\begin{aligned} 0 &\geq \frac{d}{dt} \left(\frac{1}{t-1} + \frac{a}{(t-1)^2} - \frac{at^{b+1}}{(t-1)^2} \right) \\ &= \frac{1}{(t-1)^3} (-t+1-2a+a(b+1)t^b - a(b-1)t^{b+1}) \equiv \frac{1}{(t-1)^3} \psi(t). \end{aligned}$$

Therefore, our next step is to show that $\psi(t) \leq 0$. We have

$$\begin{aligned} \psi'(t) &= -1 + a(b+1)bt^{b-1} - a(b-1)(b+1)t^b, \\ \psi''(t) &= ab(b+1)(b-1)t^{b-2} - a(b-1)b(b+1)t^{b-1} \\ &= ab(b+1)(b-1)t^{b-2}(1-t). \end{aligned}$$

By definition, $a(b+1) = 1$. Hence, since $t > 1$ and $b \in (-1, 0)$, we observe that $\psi''(t) \leq 0$. Thus, $\psi'(t) \leq \psi'(1) = 0$, and consequently, $\psi(t) \leq \psi(1) = 0$, for all $t \in (1, 2]$. This proves the claim $\gamma(t)/(t-1) \geq \gamma(2)$ for $t \in (1, 2]$.

Combining both cases, we obtain that $\gamma(t) \geq \min\{\gamma(2), (t-1)\gamma(2)\}$, where $\gamma(2) = 1 - a + a2^{1/a}$, using the fact that $b+1 = 1/a$. Unraveling this expression by using the definition of the constant a , we see that $\gamma(2)$ depends only on the self-concordance parameter $\nu \in (2, 3)$. In light of this, let us introduce the constant

$$\tilde{\gamma}_\nu := 1 + \frac{4-\nu}{2(3-\nu)} \left(1 - 2^{2(3-\nu)/(4-\nu)}\right). \tag{75}$$

Observe that $\tilde{\gamma}_2 = 0$ and, by a simple application of l'Hôpital's rule, $\lim_{\nu \uparrow 3} \tilde{\gamma}_\nu = 1 - \log(2) \in (0, 1)$. Hence $\gamma(2) \equiv \tilde{\gamma}_\nu \in (0, 1)$ for all $\nu \in (2, 3)$. We conclude,

$$\Delta_k \geq \frac{\tilde{\gamma}_\nu \lambda_k^2}{\delta_k} \min \left\{ 1, \frac{-1}{b} \frac{\delta_k}{L + \mu} \right\} = \tilde{\gamma}_\nu \min \left\{ \frac{\lambda_k^2}{\delta_k}, \frac{\lambda_k^2}{L + \mu} \frac{-1}{b} \right\}.$$

Since $\lambda_k \geq \sqrt{\sigma_h} \beta_k$, $\delta_k = M(\frac{\nu}{2} - 1) \lambda_k^{\nu-2} \beta_k^{3-\nu}$, the following lower and upper bounds can be established for $\nu \in (2, 3)$:

$$M \left(\frac{\nu}{2} - 1 \right) \sigma_h^{\frac{\nu-2}{2}} \alpha_k \beta_k \leq \alpha_k \delta_k \leq M \left(\frac{\nu}{2} - 1 \right) \sigma_h^{-\frac{3-\nu}{2}} \alpha_k \lambda_k.$$

This estimate implies first that $\delta_k \leq M(\nu/2 - 1) \sigma_h^{-(3-\nu)/2} \lambda_k$, and second

$$\frac{\lambda_k^2}{\delta_k} \geq \frac{2\lambda_k}{M(\nu-2)} \sigma_h^{\frac{3-\nu}{2}}.$$

This yields the bound

$$\Delta_k \geq \tilde{\gamma}_\nu \lambda_k \min \left\{ \frac{2\sigma_h^{(3-\nu)/2}}{M(\nu-2)}, \frac{4-\nu}{(\nu-2)(L+\mu)} \lambda_k \right\} \quad \forall k \geq 0. \quad (76)$$

Recall from Proposition 5.2 that $\lim_{k \rightarrow \infty} \Delta_k = 0$ always holds. Consequently, combining (76) with (51), we immediately see $\lim_{k \rightarrow \infty} \lambda_k = \lim_{k \rightarrow \infty} \beta_k = 0$. Now observe that

$$\alpha_k = \frac{1}{\delta_k} \left[1 - \left(1 + \frac{\delta_k}{L+\mu} \frac{4-\nu}{\nu-2} \right)^{\frac{2-\nu}{4-\nu}} \right] =: \frac{Q(\delta_k)}{\delta_k},$$

and $\delta_k = M(\nu/2 - 1)\lambda_k^{\nu-2}\beta_k^{3-\nu}$. Thus,

$$\lim_{k \rightarrow \infty} \delta_k = 0, \quad (77)$$

and $\lim_{k \rightarrow \infty} Q(\delta_k) = 0$. By l'Hôpital rule

$$\lim_{k \rightarrow \infty} \alpha_k = \lim_{k \rightarrow \infty} \frac{Q(\delta_k)}{\delta_k} = \lim_{k \rightarrow \infty} Q'(\delta_k) = \frac{1}{L+\mu} > 0. \quad (78)$$

Finally, by definition of the search direction, there exists a sequence of dual variables $(y^k)_{k \geq 0} \subset \mathbb{R}^m$, explicitly defined by (29), for which

$$\|\nabla F_\mu(x^k) - A^\top y^k\|_{x^k}^* = \lambda_k \quad \forall k \geq 0.$$

We therefore observe first that $\lim_{k \rightarrow \infty} \|\nabla F_\mu(x^k) - A^\top y^k\|_{x^k}^* = 0$, and second

$$\|\nabla F_\mu(x^k) - A^\top y^k\|_2 \leq |H(x^k)|^{1/2} \lambda_k.$$

Since $(x^k)_{k \geq 0} \subset \mathcal{S}_\mu(x^0)$, $H(x^k) \succ 0$, and $h \in \mathbf{C}^3(\text{dom } h)$, using (53), we conclude that

$$\lim_{k \rightarrow \infty} \|\nabla F_\mu(x^k) - A^\top y^k\|_2 \leq \sqrt{\tau_h} \lim_{k \rightarrow \infty} \lambda_k = 0.$$

B.2 The case $\nu = 3$

A direct substitution for Δ_k gives us

$$\Delta_k = \frac{\lambda_k^2}{\frac{M}{2}\lambda_k + L + \mu} + \frac{4}{M^2}(L + \mu) \left[\frac{\frac{M}{2}\lambda_k}{\frac{M}{2}\lambda_k + L + \mu} + \ln \left(\frac{L + \mu}{\frac{M}{2}\lambda_k + L + \mu} \right) \right]. \quad (79)$$

Denote $t_k := (L + \mu)/(\frac{M}{2}\lambda_k)$, $\delta_k = \frac{M}{2}\lambda_k$. Then

$$\alpha_k = \frac{2}{M\lambda_k} \frac{1}{1+t_k} = \frac{1}{\delta_k + L + \mu},$$

so that

$$\frac{\alpha_k M \lambda_k}{2} = \frac{1}{1+t_k}, \quad \text{and } L + \mu = \frac{M}{2} \lambda_k t_k.$$

This implies that

$$\Delta_k = \frac{2\lambda_k}{M} \frac{1}{1+t_k} + \frac{2\lambda_k}{M} t_k \left[\frac{1}{1+t_k} + \ln \left(\frac{t_k}{1+t_k} \right) \right],$$

$$= \frac{2\lambda_k}{M} \left(1 + t_k \ln \left(\frac{t_k}{1+t_k} \right) \right). \quad (80)$$

Consider the function $\gamma : (0, \infty) \rightarrow (0, \infty)$, given by $\gamma(t) := 1 + t \ln \left(\frac{t}{1+t} \right)$. When $t \in (0, 1)$, since

$$\begin{aligned} \gamma'(t) &= \ln \left(\frac{t}{1+t} \right) + t \frac{1+t}{t} \left(\frac{1}{1+t} - \frac{t}{(1+t)^2} \right) \\ &= \ln \left(1 - \frac{1}{1+t} \right) + \frac{1}{1+t} < 0, \end{aligned}$$

we conclude that $\gamma(t)$ is decreasing for $t \in (0, 1)$. Hence, $\gamma(t) \geq \gamma(1) = 1 - \ln 2$, for all $t \in (0, 1)$. On the other hand, if $t \geq 1$,

$$\frac{d}{dt} \left(\frac{\gamma(t)}{1/t} \right) = \frac{d}{dt} (t\gamma(t)) = 1 + 2t \ln \left(\frac{t}{1+t} \right) + \frac{t}{1+t} \geq 0.$$

Hence, $t \mapsto \frac{\gamma(t)}{1/t}$ is an increasing function for $t \geq 1$, and thus $\gamma(t) \geq \frac{1-\ln 2}{t}$, for all $t \geq 1$. Summarizing these two cases we see $\Delta_k \geq \frac{2\lambda_k}{M} \min\{1, 1/t_k\}(1 - \ln(2))$, which after rearranging, can be stated as

$$\Delta_k \geq \frac{2(1 - \ln(2))\lambda_k}{M(L + \mu)} \min \left\{ L + \mu, \frac{M}{2} \lambda_k \right\} \quad \forall k \geq 0. \quad (81)$$

From Proposition 5.2 we know that $\lim_{k \rightarrow \infty} \Delta_k = 0$, and consequently,

$$\lim_{k \rightarrow \infty} \delta_k = \lim_{k \rightarrow \infty} \lambda_k = 0, \text{ as well as } \lim_{k \rightarrow \infty} \alpha_k = \frac{1}{L + \mu}. \quad (82)$$

Eq. (51) shows that $\lim_{k \rightarrow \infty} \beta_k = 0$. As in the case $\nu \in (2, 3)$, we arrive at the conclusion $\lim_{k \rightarrow \infty} \|\nabla F_\mu(x^k) - A^\top y^k\|_2 = 0$.

B.3 The case $\nu \in (3, 4)$

Similarly to the case $\nu \in (2, 3)$, denote $t_k = 1 + \frac{\delta_k}{L+\mu} \frac{4-\nu}{\nu-2} \in (1, +\infty)$, $\mathbf{a} = \frac{4-\nu}{2(3-\nu)} \in (-\infty, 0)$, $\mathbf{b} = \frac{2-\nu}{4-\nu} \in (-\infty, -1)$. Then the expression for the Δ_k is the same as for $\nu \in (2, 3)$:

$$\Delta_k = \frac{\lambda_k^2}{\delta_k} \left(1 + \frac{\mathbf{a}}{t_k - 1} - \frac{\mathbf{a}t_k^{\mathbf{b}+1}}{t_k - 1} \right).$$

Let us define a function $\gamma(t)$ such that $\Delta_k = \frac{\lambda_k^2}{\delta_k} \gamma(t_k)$. Our next goal is to show that, for $t \in [2, +\infty)$, $\gamma(t)$ is below bounded by some positive constant and, for $t \in (1, 2]$, $\gamma(t)$ is below bounded by some positive constant multiplied by $t - 1$.

1. $t \in [2, +\infty)$. We will show that $\gamma'(t) \geq 0$, whence $\gamma(t) \geq \gamma(2)$. Thus, we need to show that

$$0 \leq \frac{d}{dt} \left(1 + \frac{\mathbf{a}}{t-1} - \frac{\mathbf{a}t^{\mathbf{b}+1}}{t-1} \right) = -\frac{\mathbf{a}}{(t-1)^2} \underbrace{(1 - (\mathbf{b}+1)t^{\mathbf{b}} + \mathbf{b}t^{\mathbf{b}+1})}_{=:\psi(t)}.$$

Since $\mathbf{a} \leq 0$, to show that $\gamma'(t) \geq 0$ it is enough to show that $\psi(t) \geq 0$. Since $\mathbf{b} < -1$,

$$\psi'(t) = \mathbf{b}(\mathbf{b}+1)t^{\mathbf{b}} - \mathbf{b}(\mathbf{b}+1)t^{\mathbf{b}-1} = \mathbf{b}(\mathbf{b}+1)t^{\mathbf{b}-1}(t-1) \geq 0,$$

whence, $\psi(t) \geq \psi(2)$, $t \in [2, +\infty)$. It remains to show that $\psi(2) \geq 0$. Let us consider $\psi(2) = 1 - (b+1)2^b + b2^{b+1} = 1 + b2^b - 2^b$ as a function of b . For all possible values $b \in (-\infty, -1)$ one can check numerically that $\psi(2) \in (0, 1)$. Hence, $\psi(t) \geq 0$ for all $t \geq 2$.

2. $t \in (1, 2]$. We will show that $\frac{d}{dt}(\gamma(t)/(t-1)) \leq 0$, whence $\gamma(t) \geq (t-1)\gamma(2)$. Thus, we need to show that

$$\begin{aligned} 0 &\geq \frac{d}{dt} \left(\frac{1}{t-1} + \frac{a}{(t-1)^2} - \frac{at^{b+1}}{(t-1)^2} \right) \\ &= -\frac{1}{(t-1)^2} - \frac{2a}{(t-1)^3} - \frac{a(b+1)t^b}{(t-1)^2} + \frac{2at^{b+1}}{(t-1)^3} \\ &= \frac{1}{(t-1)^3} (-t+1-2a-a(b+1)t^{b+1}+a(b+1)t^b+2at^{b+1}) \\ &= \frac{1}{(t-1)^3} \underbrace{(-t+1-2a+a(b+1)t^b-a(b-1)t^{b+1})}_{=:\psi(t)}. \end{aligned}$$

Our next step is to show that $\psi(t) \leq 0$. We have

$$\begin{aligned} \psi'(t) &= -1 + a(b+1)bt^{a-1} - a(b-1)(b+1)t^b \\ \psi''(t) &= ab(b+1)(b-1)t^{b-2} - a(b-1)b(b+1)t^{b-1} \\ &= ab(b+1)(b-1)t^{b-2}(1-t). \end{aligned}$$

Using the definition of a, b , and the fact that $\nu \in (3, 4)$, we obtain that $a(b+1) = 1$. Hence, since $t > 1$, we obtain that $\psi''(t) \leq 0$. Thus, $\psi'(t) \leq \psi'(1) = 0$, $\psi(t) \leq \psi(1) = 0$, and $\gamma(t)/(t-1) \geq \gamma(2)$.

Combining both cases, we obtain that $\gamma(t) \geq \min\{\gamma(2), (t-1)\gamma(2)\}$. Note that

$$\begin{aligned} \gamma(2) &\equiv \tilde{\gamma}_\nu := 1 + \frac{4-\nu}{2(3-\nu)} - \frac{4-\nu}{2(3-\nu)} 2^{(2-\nu)/(4-\nu)+1} \\ &= 1 + \frac{1 - \exp\left(\frac{2(3-\nu)}{4-\nu} \ln(2)\right)}{\frac{2(3-\nu)}{4-\nu}} \end{aligned}$$

Via L'Hôpital's rule, one can check that $\lim_{\nu \downarrow 3} \hat{\gamma}(\nu) = 1 - \ln(2) \in (0, 1)$, and $\lim_{\nu \uparrow 4} \hat{\gamma}(\nu) = 1$, since $\lim_{\nu \uparrow 4} \frac{2(3-\nu)}{4-\nu} = -\infty$. Consequently,

$$\Delta_k \geq \tilde{\gamma}_\nu \min \left\{ \frac{\lambda_k^2}{\delta_k}, \frac{1}{-b} \frac{\lambda_k^2}{L + \mu} \right\}, \quad \tilde{\gamma}_\nu \in (1 - \ln(2), 1).$$

By Proposition 5.2, we know that $\lim_{k \rightarrow \infty} \Delta_k = 0$. Therefore, either $\lambda_k \rightarrow 0$, or $\frac{\lambda_k^2}{\delta_k} \rightarrow 0$. Suppose there exists $\varepsilon > 0$ such that $\lambda_k \geq \varepsilon$ for all $k \geq 0$. Then,

$$\frac{\lambda_k^2}{\delta_k} = \frac{2}{M(\nu-2)} \lambda_k^{4-\nu} \beta_k^{\nu-3} \geq \frac{2}{M(\nu-2)} \varepsilon^{4-\nu} \beta_k^{\nu-3}.$$

Hence, $\beta_k \rightarrow 0$ must hold. But then eq. (55) implies $\lambda_k \rightarrow 0$. A contradiction. It follows that $\lambda_k \rightarrow 0$, and therefore, by (51), $\beta_k \rightarrow 0$. Using that

$$\delta_k = M(\nu/2 - 1) \lambda_k^{\nu-2} \beta_k^{3-\nu} \stackrel{(55)}{\leq} M(\nu/2 - 1) \tau_h^{\frac{\nu-2}{2}} \beta_k,$$

we arrive at the string of inequalities

$$\frac{\lambda_k^2}{\delta_k} = \frac{2}{M(\nu-2)} \lambda_k^{4-\nu} \beta_k^{\nu-3} \stackrel{(55)}{\geq} \frac{2}{M(\nu-2)} \tau_h^{-\frac{\nu-3}{2}} \lambda_k$$

Hence,

$$\lim_{k \rightarrow \infty} \delta_k = 0, \text{ and } \lim_{k \rightarrow \infty} \alpha_k = \frac{1}{L + \mu}. \quad (83)$$

We conclude $\lim_{k \rightarrow \infty} \|\nabla F_\mu(x^k) - A^\top y^k\|_2 = 0$. Moreover, we obtain the explicit bound

$$\Delta_k \geq \lambda_k \tilde{\gamma}_\nu \min \left\{ \frac{2}{M(\nu-2)} \tau_h^{-\frac{\nu-3}{2}}, \frac{1}{-b} \frac{\lambda_k}{L + \mu} \right\}. \quad (84)$$

B.4 The case $\nu = 4$

We can compute the per-iteration potential reduction as

$$\begin{aligned} \Delta_k &= \frac{\lambda_k^2}{\delta_k} \left[1 - \exp\left(-\frac{\delta_k}{L + \mu}\right) \right] \\ &\quad - \left(\frac{\lambda_k}{\delta_k}\right)^2 (L + \mu) \left(\left[1 - \exp\left(-\frac{\delta_k}{L + \mu}\right) \right] - \frac{\delta_k}{L + \mu} \exp\left(-\frac{\delta_k}{L + \mu}\right) \right). \end{aligned}$$

To analyze this expression, denote by $t_k^{-1} := \frac{\delta_k}{L + \mu}$. Then

$$\Delta_k = \frac{\lambda_k^2}{\delta_k} \left(1 - t_k + t_k \exp\left(-\frac{1}{t_k}\right) \right) \geq 0.$$

Let us define a function $\gamma(t)$ such that $\Delta_k = \frac{\lambda_k^2}{\delta_k} \gamma(t_k)$. Our next goal is to show that, for $t \in (0, 1]$, $\gamma(t)$ is below bounded by some positive constant and, for $t \geq 1$, $\gamma(t)$ is below bounded by some positive constant divided by t .

1. $t \in (0, 1]$. We will show that $\gamma'(t) \leq 0$, whence $\gamma(t) \geq \gamma(1)$. Indeed, for $t \in (0, 1]$,

$$\begin{aligned} \gamma'(t) &= -1 + \exp(-1/t)(1 + 1/t) \\ &< -1 + 2 \exp(1/t) \\ &\leq -1 + 2 \exp(-1) < 0. \end{aligned}$$

Thus, we have $\gamma(t) \geq \gamma(1) = \exp(-1)$.

2. $t \in [1, +\infty)$. We will show that $\frac{d}{dt} \left(\frac{\gamma(t)}{1/t} \right) \geq 0$, whence $\gamma(t) \geq \frac{\gamma(1)}{t}$.

$$\frac{d}{dt} \left(t \left(1 - t + t \exp\left(-\frac{1}{t}\right) \right) \right) = \exp\left(-\frac{1}{t}\right) (2t + 1) + 1 - 2t. \quad (85)$$

Using the Taylor expansion for $\ln(1+x)$ and $\ln(1-x)$ for $x \in (0, 0.5]$, we have

$$\begin{aligned} \ln(1+x) - \ln(1-x) &= x - \frac{x^2}{2} + \frac{x^3}{3} + \sum_{k=4}^{\infty} \frac{(-1)^k x^k}{k} - \left(-x - \frac{x^2}{2} + \frac{x^3}{3} - \sum_{k=4}^{\infty} \frac{x^k}{k} \right) \\ &= 2x + \frac{2x^3}{3} + \sum_{k=2}^{\infty} \frac{2x^{2k+1}}{2k+1} \geq 2x. \end{aligned}$$

Setting $x = \frac{1}{2t}$ for $t \geq 1$, we obtain

$$\begin{aligned} & \ln(1 + 1/(2t)) - \ln(1 - 1/(2t)) \geq 1/t \\ \Leftrightarrow & \ln(2t(1 + 1/(2t))) - \ln(2t(1 - 1/(2t))) \geq 1/t \\ \Leftrightarrow & \ln(2t + 1) - \ln(2t - 1) \geq 1/t \\ \Leftrightarrow & -1/t + \ln(2t + 1) \geq \ln(2t - 1) \\ \Leftrightarrow & \exp\left(-\frac{1}{t}\right) (2t + 1) + 1 - 2t \geq 0. \end{aligned}$$

which, combined with (85) proves that $\frac{d}{dt} \left(\frac{\gamma(t)}{1/t} \right) \geq 0$ for $t \geq 1$. Thus, we have that, for $t \geq 1$, $\gamma(t) \geq \frac{\gamma(1)}{t} = \frac{\exp(-1)}{t}$.

Combining two cases $t_k \in (0, 1]$ and $t_k \in [1, +\infty)$, we obtain that $\gamma(t) \geq \min\{\gamma(1), \gamma(1)/t\}$ and, since $t_k^{-1} := \frac{\delta_k}{L+\mu}$

$$\Delta_k = \frac{\lambda_k^2}{\delta_k} \gamma(t_k) \geq \frac{\lambda_k^2}{\delta_k} \min\{\gamma(1), \gamma(1)/t_k\} = \exp(-1) \min\left\{ \frac{\lambda_k^2}{\delta_k}, \frac{\lambda_k^2}{L+\mu} \right\}.$$

By Proposition 5.2, we know that $\lim_{k \rightarrow \infty} \Delta_k = 0$. Thus, either $\lambda_k \rightarrow 0$, or $\frac{\lambda_k^2}{\delta_k} \rightarrow 0$. Suppose there exists $\varepsilon > 0$ such that $\lambda_k \geq \varepsilon > 0$ for all $k \geq 0$. Then it must be true $\frac{\lambda_k^2}{\delta_k} \rightarrow 0$. Then, $\frac{\lambda_k^2}{\delta_k} = \frac{\beta_k}{M}$, and therefore $\beta_k \rightarrow 0$ must be true. But then (55) yields the contradiction $\lambda_k \rightarrow 0$. We are therefore forced to conclude that $\lim_{k \rightarrow \infty} \lambda_k = 0$, and from (51) it then follows $\lim_{k \rightarrow \infty} \beta_k = 0$. Furthermore, using (55),

$$\delta_k = M \frac{\lambda_k^2}{\beta_k} \leq M \tau_h \beta_k,$$

so that $\lim_{k \rightarrow \infty} \delta_k = 0$ and $\frac{\lambda_k^2}{\delta_k} \geq \frac{\lambda_k}{\sqrt{\tau_h M}}$. This gives the final estimate

$$\Delta_k \geq \exp(-1) \lambda_k \min\left\{ \frac{1}{\sqrt{\tau_h M}}, \frac{\lambda_k}{L+\mu} \right\}. \quad (86)$$

A simple application of l'Hôpital's rule gives $\lim_{k \rightarrow \infty} \alpha_k = \frac{1}{L+\mu}$, and $\lim_{k \rightarrow \infty} \|\nabla F_\mu(x^k) - A^\top y^k\|_2 = 0$.

Proof of Theorem 5.6. Combining all the results just derived for each generalized self-concordant parameter $\nu \in (2, 4]$, we conclude that always $\lim_{k \rightarrow \infty} \|\nabla F_\mu(x^k) - A^\top y^k\|_2 = 0$. Corollary 5.5 shows that $(x^k)_{k \geq 0} \subseteq \mathcal{S}_\mu(x^0)$, which is a compact set by Lemma 5.4 contained in \mathcal{X}° . Since $\nabla F_\mu(x) = \nabla f(x) + \mu \nabla h(x)$ is a continuous function on \mathcal{X}° , we conclude that along every convergent subsequence $(x^{k_q})_{q \in \mathbb{N}}$ with limit $\bar{x} \in \mathcal{X}^\circ$, we have

$$\lim_{q \rightarrow \infty} A^\top y^{k_q} = \lim_{q \rightarrow \infty} \nabla F_\mu(x^{k_q}) = \nabla F_\mu(\bar{x}).$$

Recall that $y^{k_q} = y_\mu(x^{k_q})$ and the map $x \mapsto y_\mu(x)$ is continuous by Lemma 3.4. Denote by $\bar{y} \in \mathbb{R}^m$ the corresponding limit of the convergent subsequence $(y^{k_q})_{q \in \mathbb{N}}$, we conclude that $\nabla F_\mu(\bar{x}) = A^\top \bar{y}$. Since the convergent subsequence $(x^{k_q})_{q \in \mathbb{N}}$ has been chosen arbitrarily, the claim $\omega(x^0) \subseteq \{x \in \mathcal{X} \mid (\exists y \in \mathbb{R}^m) : \nabla F_\mu(x) - A^\top y = 0\}$ follows. ■

References

- [1] Naman Agarwal, Zeyuan Allen-Zhu, Brian Bullins, Elad Hazan, and Tengyu Ma. Finding approximate local minima faster than gradient descent. pages 1195–1199. ACM, 2017.
- [2] Felipe Alvarez, Jérôme Bolte, and Olivier Brahic. Hessian Riemannian gradient flows in convex programming. *SIAM Journal on Control and Optimization*, 43(2):477–501, 2004.
- [3] A. Auslender and M. Teboulle. Interior gradient and proximal methods for convex and conic optimization. *SIAM Journal on Optimization*, 16(3):697–725, 2016/06/22 2006.
- [4] Heinz H. Bauschke, Jérôme Bolte, and Marc Teboulle. A descent lemma beyond lipschitz gradient continuity: First-order methods revisited and applications. *Mathematics of Operations Research*, 42(2):330–348, 2017/09/12 2016.
- [5] D. A. Bayer and J. C. Lagarias. The nonlinear geometry of linear programming. i. affine and projective scaling trajectories. *Trans. Amer. Math. Soc*, pages 499–526, 1989.
- [6] Wei Bian, Xiaojun Chen, and Yinyu Ye. Complexity analysis of interior point algorithms for non-lipschitz and nonconvex minimization. *Mathematical Programming*, 149(1):301–327, 2015.
- [7] Lev Bogolubsky, Pavel Dvurechensky, Alexander Gasnikov, Gleb Gusev, Yurii Nesterov, Andrei M Raigorodskii, Aleksey Tikhonov, and Maksim Zhukovskii. Learning supervised pagerank with gradient-based and gradient-free optimization methods. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 4914–4922. Curran Associates, Inc., 2016. arXiv:1603.00717.
- [8] J. Bolte, S. Sabach, M. Teboulle, and Y. Vaisbourd. First order methods beyond convexity and lipschitz gradient continuity with applications to quadratic inverse problems. *SIAM Journal on Optimization*, 28(3):2131–2151, 2019/09/05 2018.
- [9] J. Bolte and M. Teboulle. Barrier operators and associated gradient-like dynamical systems for constrained minimization problems. *SIAM Journal on Control and Optimization*, 42(4):1266–1292, 2003.
- [10] Jérôme Bolte, Shoham Sabach, and Marc Teboulle. Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Mathematical Programming*, 146(1):459–494, 2014.
- [11] Immanuel M Bomze, Panayotis Mertikopoulos, Werner Schachinger, and Mathias Staudigl. Hessian barrier algorithms for linearly constrained optimization problems. *SIAM Journal on Optimization*, 29(3):2100–2127, 2019.
- [12] A. Bruckstein, D. Donoho, and M. Elad. From sparse solutions of systems of equations to sparse modeling of signals and images. *SIAM Review*, 51(1):34–81, 2019/10/05 2009.
- [13] Sébastien Bubeck and Ronen Eldan. The entropic barrier: Exponential families, log-concave geometry, and self-concordance. *Mathematics of Operations Research*, 44(1):264–276, 2019/10/02 2018.
- [14] Yair Carmon, John C Duchi, Oliver Hinder, and Aaron Sidford. Convex until proven guilty: Dimension-free acceleration of gradient descent on non-convex functions. pages 654–663. JMLR. org, 2017.
- [15] Yair Carmon, John C Duchi, Oliver Hinder, and Aaron Sidford. Accelerated methods for nonconvex optimization. *SIAM Journal on Optimization*, 28(2):1751–1772, 2018.
- [16] Yair Carmon, John C. Duchi, Oliver Hinder, and Aaron Sidford. Lower bounds for finding stationary points i. *Mathematical Programming*, 2019.
- [17] Yair Carmon, John C. Duchi, Oliver Hinder, and Aaron Sidford. Lower bounds for finding stationary points ii: first-order methods. *Mathematical Programming*, 2019.
- [18] C. Cartis, N. Gould, and P. Toint. On the evaluation complexity of composite function minimization with applications to nonconvex nonlinear programming. *SIAM Journal on Optimization*, 21(4):1721–1739, 2019/09/30 2011.
- [19] Coralia Cartis, Nicholas IM Gould, and Ph L Toint. An adaptive cubic regularization algorithm for nonconvex optimization with convex constraints and its function-evaluation complexity. *IMA Journal of Numerical Analysis*, 32(4):1662–1695, 2012.
- [20] Bruce Cox, Anatoli Juditsky, and Arkadi Nemirovski. Dual subgradient algorithms for large-scale nonsmooth learning problems. *Mathematical Programming*, 148(1):143–180, 2014.
- [21] Pavel Dvurechensky and Yurii Nesterov. Global performance guarantees of second-order methods for unconstrained convex minimization. CORE Discussion Paper 2018/32, December 2018.
- [22] Jianqing Fan and Runze Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96(456):1348–1360, 2001.
- [23] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. *The elements of statistical learning*, volume 1. Springer series in statistics Springer, Berlin, 2001.

- [24] Dongdong Ge, Xiaoye Jiang, and Yinyu Ye. A note on the complexity of l_p minimization. *Mathematical Programming*, 129(2):285–299, 2011.
- [25] Saeed Ghadimi, Guanghui Lan, and Hongchao Zhang. Generalized uniformly optimal methods for nonlinear programming. *Journal of Scientific Computing*, 79(3):1854–1881, 2019. arXiv:1508.07384.
- [26] S. Gratton, A. Sartenaer, and P. Toint. Recursive trust-region methods for multiscale nonlinear optimization. *SIAM Journal on Optimization*, 19(1):414–444, 2019/09/30 2008.
- [27] S. V. Guminov, Yu. E. Nesterov, P. E. Dvurechensky, and A. V. Gasnikov. Accelerated primal-dual gradient descent with linesearch for convex, nonconvex, and nonsmooth optimization problems. *Doklady Mathematics*, 99(2):125–128, Mar 2019.
- [28] Gabriel Haeser, Hongcheng Liu, and Yinyu Ye. Optimality condition and complexity analysis for linearly-constrained optimization without differentiability on the boundary. *Mathematical Programming*, 2018.
- [29] Chi Jin, Praneeth Netrapalli, and Michael I. Jordan. Accelerated gradient descent escapes saddle points faster than gradient descent. In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet, editors, *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 1042–1085. PMLR, 2018.
- [30] John M. Lee. *Riemannian manifolds : an introduction to curvature*. Springer-Verlag, 1997.
- [31] Hongcheng Liu, Tao Yao, and Runze Li. Global solutions to folded concave penalized nonconvex learning. *Ann. Statist.*, 44(2):629–659, 2016.
- [32] Hongcheng Liu, Tao Yao, Runze Li, and Yinyu Ye. Folded concave penalized sparse linear regression: sparsity, statistical performance, and algorithmic theory for local solutions. *Mathematical Programming*, 166(1):207–240, 2017.
- [33] Miguel Sousa Lobo, Lieven Vandenbergh, Stephen Boyd, and Hervé Lebret. Applications of second-order cone programming. *Linear Algebra and its Applications*, 284(1):193–228, 1998.
- [34] Po-Ling Loh and Martin J. Wainwright. Regularized m-estimators with nonconvexities: Statistical and algorithmic theory for local optima. *Journal of Machine Learning Research*, 16:559–616, 2015.
- [35] H. Lu, R. Freund, and Y. Nesterov. Relatively smooth convex optimization by first-order methods, and applications. *SIAM Journal on Optimization*, 28(1):333–354, 2019/03/21 2018.
- [36] Zhaosong Lu. Randomized block proximal damped newton method for composite self-concordant minimization. *SIAM Journal on Optimization*, 27(3):1910–1942, 2019/11/01 2017.
- [37] Hariharan Narayanan. Randomized interior point methods for sampling and optimization. *Ann. Appl. Probab.*, 26(1):597–641, 2016.
- [38] Hariharan Narayanan and Alexander Rakhlin. Random walk approach to regret minimization. In J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems 23*, pages 1777–1785. Curran Associates, Inc., 2010.
- [39] Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Kluwer, Dordrecht, 2004.
- [40] Yu. Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, 2005.
- [41] Yu. Nesterov. Gradient methods for minimizing composite functions. *Mathematical Programming*, 140(1):125–161, 2013.
- [42] Yu Nesterov. Universal gradient methods for convex optimization problems. *Mathematical Programming*, 152(1):381–404, 2015.
- [43] Yu. Nesterov and A. Nemirovski. *Interior Point Polynomial methods in Convex programming*. SIAM Publications, 1994.
- [44] Yurii Nesterov and Boris Polyak. Cubic regularization of newton method and its global performance. *Mathematical Programming*, 108(1):177–205, 2006.
- [45] Xiaotong Shen and Hsin-Cheng Huang. Grouping pursuit through a regularization solution surface. *Journal of the American Statistical Association*, 105(490):727–739, 2010.
- [46] Tianxiao Sun and Quoc Tran-Dinh. Generalized self-concordant functions: a recipe for newton-type methods. *Mathematical Programming*, 2018.
- [47] Marc Teboulle. A simplified view of first order methods for optimization. *Mathematical Programming*, 170(1):67–96, 2018.
- [48] Robert Tibshirani, Michael Saunders, Saharon Rosset, Ji Zhu, and Keith Knight. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(1):91–108, 2019/04/05 2005.

- [49] Quoc Tran-Dinh, Anastasios Kyrillidis, and Volkan Cevher. Composite self-concordant minimization. *The Journal of Machine Learning Research*, 16(1):371–416, 2015.
- [50] Quoc Tran-Dinh, Tianxiao Sun, and Shu Lu. Self-concordant inclusions: a unified framework for path-following generalized newton-type algorithms. *Mathematical Programming*, 177(1):173–223, 2019.
- [51] Paul Tseng, Immanuel M. Bomze, and Werner Schachinger. A first-order interior-point method for linearly constrained smooth optimization. *Mathematical Programming*, 127(2):399–424, 2011.
- [52] César A Uribe, Soomin Lee, Alexander Gasnikov, and Angelia Nedić. A dual approach for optimal algorithms in distributed optimization over networks. *arXiv preprint arXiv:1809.00710*, 2018.
- [53] Zhaoran Wang, Han Liu, and Tong Zhang. Optimal computational and statistical rates of convergence for sparse nonconvex learning problems. *Ann. Statist.*, 42(6):2164–2201, 2014.
- [54] Yinyu Ye. On affine scaling algorithms for nonconvex quadratic programming. *Mathematical Programming*, 56(1):285–300, 1992.
- [55] Yinyu Ye. On the complexity of approximating a KKT point of quadratic programming. *Mathematical Programming*, 80(2):195–211, 1998.
- [56] Yuchen Zhang and Xiao Lin. Disco: Distributed optimization for self-concordant empirical loss. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 362–370. PMLR, 06 2015.