

Decomposable Non-Smooth Convex Optimization with Nearly-Linear Gradient Oracle Complexity

Sally Dong
University of Washington, Seattle
sallyqd@uw.edu

Haotian Jiang
University of Washington, Seattle
jhtdavid@uw.edu

Yin Tat Lee[†]
University of Washington, Seattle
yintat@uw.edu

Swati Padmanabhan
University of Washington, Seattle
pswati@uw.edu

Guanghao Ye
Massachusetts Institute of Technology
ghye@mit.edu

August 9, 2022

Abstract

Many fundamental problems in machine learning can be formulated by the convex program

$$\min_{\theta \in \mathbb{R}^d} \sum_{i=1}^n f_i(\theta),$$

where each f_i is a convex, Lipschitz function supported on a subset of d_i coordinates of θ . One common approach to this problem, exemplified by stochastic gradient descent, involves sampling one f_i term at every iteration to make progress. This approach crucially relies on a notion of uniformity across the f_i 's, formally captured by their condition number. In this work, we give an algorithm that minimizes the above convex formulation to ϵ -accuracy in $\tilde{O}(\sum_{i=1}^n d_i \log(1/\epsilon))$ gradient computations, with no assumptions on the condition number. The previous best algorithm independent of the condition number is the standard cutting plane method, which requires $O(nd \log(1/\epsilon))$ gradient computations. As a corollary, we improve upon the evaluation oracle complexity for decomposable submodular minimization by [AKMSV21]. Our main technical contribution is an adaptive procedure to select an f_i term at every iteration via a novel combination of cutting-plane and interior-point methods.

*Author names are listed in alphabetical order.

[†]Supported by NSF awards CCF-1749609, DMS-1839116, DMS-2023166, CCF-2105772, a Microsoft Research Faculty Fellowship, a Sloan Research Fellowship, and a Packard Fellowship

1 Introduction

Many fundamental problems in machine learning are abstractly captured by the convex optimization formulation

$$\text{minimize}_{\theta \in \mathbb{R}^d} \sum_{i=1}^n f_i(\theta), \quad (1.1)$$

where each f_i is a convex, Lipschitz function. For example, in empirical risk minimization, each f_i measures the loss incurred by the i -th data point from the training set. In generalized linear models, each f_i represents a link function applied to a linear predictor evaluated at the i -th data point.

The ubiquity of (1.1) in the setting with smooth f_i 's has spurred the development of well-known variants of stochastic gradient methods [RM51; BC03; Zha04; Bot12] such as [RSB12; SZ13b; JZ13; MZJ13; DBL14; Mai15; AY16; HL16; SLB17]; almost universally, these algorithms leverage the “sum structure” of (1.1) by sampling, in each iteration, one f_i with which to make progress. These theoretical developments have in turn powered tremendous empirical success in machine learning through widely used software packages such as `libSVM` [CL11].

In many practical applications, (1.1) appears with non-smooth f_i 's, as well as the additional structure that each f_i depends only on a subset of the problem parameters θ . One notable example is decomposable submodular function minimization¹ (SFM), which has proven to be expressive in diverse contexts such as determinantal point processes [KT10], MAP inference in computer vision [KLT09; VKR09; FJPZ13], hypergraph cuts [VBK20], and covering functions [SK10]. Another application is found in generalized linear models when the data is high dimensional and sparse. In this setting, f_i depends on a restricted subset of the parameters θ that correspond to the features of the data point with non-zero value. Last but not least, the case with each f_i depending on a small subset of the parameters is also called sparse separable optimization and has applications in sparse SVM and matrix completion [RRWN11].

In this work, we initiate a systematic study of algorithms for (1.1) without the smoothness assumption². Motivated by the aforementioned applications, we introduce the additional structure that each f_i depends on a subset of the coordinates of θ . As is standard in the black-box model for studying first-order convex optimization methods, we allow sub-gradient oracle access to each f_i .

Problem 1.1. *Let $f_1, f_2, \dots, f_n : \mathbb{R}^d \mapsto \mathbb{R}$ be convex, L -Lipschitz, and possibly non-smooth functions, where each f_i depends on d_i coordinates of θ , and is accessible via a (sub-)gradient oracle. Define $m \stackrel{\text{def}}{=} \sum_{i=1}^n d_i$ to be the “total effective dimension” of the problem. Let $\theta^* \stackrel{\text{def}}{=} \arg \min_{\theta \in \mathbb{R}^d} \sum_{i=1}^n f_i(\theta)$ be a minimizer of (1.1), and let $\theta^{(0)}$ be an initial point such that $\|\theta^{(0)} - \theta^*\|_2 \leq R$. We want to compute a vector $\theta \in \mathbb{R}^d$ satisfying*

$$\sum_{i=1}^n f_i(\theta) \leq \epsilon LR + \sum_{i=1}^n f_i(\theta^*). \quad (1.2)$$

Prior works. We focus on the weakly-polynomial regime and therefore restrict ourselves to algorithms with $\text{polylog}(1/\epsilon)$ gradient oracle complexities. Table 1 summarizes the performance of all well-known algorithms applied to Problem 1.1. Note that the variants of gradient descent each require bounded condition number. The results of [Nes83; All17] and cutting plane methods are all complemented by matching lower bounds [WS16; Nes04].

Even with smooth f_i 's, first-order methods perform poorly when the condition number is large, or when there is a long chain of variable dependencies. These instances commonly arise in applications; an example from signal processing is

$$\text{minimize}_{\mathbf{x}} \left\{ (\mathbf{x}_1 - 1)^2 + \sum_{i=2}^{n-1} (\mathbf{x}_i - \mathbf{x}_{i+1})^2 + \mathbf{x}_n^2 \right\}, \quad (1.3)$$

whose variables form an $O(n)$ -length chain of dependencies, and has condition number $\kappa = \Theta(n^2)$ and $\bar{\kappa} = \Theta(n^3)$. Gradient descent algorithms such as [Nes83] and [All17] therefore require $\Omega(n^2)$ gradient queries, despite the problem's total effective dimension being only $O(n)$.

¹In decomposable submodular minimization, each f_i corresponds to the Lovász extension of the individual submodular function and is therefore generally non-smooth.

²A function f is said to be β -smooth if $f(y) \leq f(x) + \langle \nabla f(x), y - x \rangle + \beta/2 \|y - x\|_2^2$ for all x, y and α -strongly-convex if $f(y) \geq f(x) + \langle \nabla f(x), y - x \rangle + \alpha/2 \|y - x\|_2^2$ for all x, y . The condition number of f is defined to be $\kappa = \beta/\alpha$.

Authors	Algorithm Type	Gradient Queries	Non-smooth OK?
[Cau+47]	Gradient Descent (GD)	$O(n\kappa \log(1/\epsilon))$	
[Nes83]	Accelerated (Acc.) GD	$O(n\sqrt{\kappa} \log(1/\epsilon))$	
[RSB12; SZ13b; JZ13]	Stochastic (Stoch.) Variance-Reduced GD	$O((n + \bar{\kappa}) \log(1/\epsilon))$	
[SZ13a; LMH15; FGKS15; ZL15; AB15]	Acc. Stoch. Variance-Reduced GD	$O((n + \sqrt{n\bar{\kappa}}) \log(\bar{\kappa}) \log(1/\epsilon))$	
[All17]	Acc. Stoch. Variance-Reduced GD	$O((n + \sqrt{n\bar{\kappa}}) \log(1/\epsilon))$	
[KTE88; NN89; Vai89; BV02; LSW15; JLSW20]	Cutting-Plane Method (CPM)	$O(nd \log(1/\epsilon))$	✓
[LV21; DLY21]	Robust Interior-Point Method (IPM)	$O(\sum_{i=1}^n d_i^{3.5} \log(1/\epsilon))$	✓

Table 1: Gradient oracle complexities for solving (1.1) to ϵ -additive accuracy. κ denotes the condition number of $\sum_i f_i$, and $\bar{\kappa}$ is a variant of the condition number defined to be the sum of smoothness of the f_i 's divided by the strong convexity of $\sum_i f_i$.

On the other hand, cutting-plane methods (CPM) and robust interior-point methods (IPM) both trade off the dependency on condition number for worse dependencies on the problem dimension.

These significant gaps in the existing body of work motivate the following question:

Can we solve Problem 1.1 using a nearly-linear (in total effective dimension) number of subgradient oracle queries?

In this paper, we give an affirmative answer to this question.

1.1 Our results

We present an algorithm to solve Problem 1.1 with gradient oracle complexity nearly-linear in the total effective dimension.

Theorem 1.2 (Main Result (Informal); see Theorem 4.11 for formal statement). *We give an algorithm that provably solves Problem 1.1 using $O(m \log(m/\epsilon))$ subgradient oracle queries, where $m \stackrel{\text{def}}{=} \sum_{i=1}^n d_i$.*

Intuitively, the number of gradient queries for each f_i should be thought of as $\tilde{O}(d_i)$ in our algorithm, which nearly matches that of the standard cutting-plane method for minimizing the individual function f_i . The nearly-linear dependence on m overall is obtained by leveraging the additional structure on the f_i 's and stands in stark contrast to the $O(nd)$ query complexity of CPM, which is significantly worse in the case where each $d_i \ll d$. Furthermore, we improve over the current best gradient descent algorithms in the case where the f_i 's have a large condition number.

Based on the query complexity of the standard cutting-plane method, we have the following lower bound matching our algorithm's query complexity up to a log m -factor:

Theorem 1.3. *There exist functions $f_1, \dots, f_n : \mathbb{R}^d \mapsto \mathbb{R}$ for which a total of $\Omega(m \log(1/\epsilon))$ gradient queries are required to solve Problem 1.1.*

An immediate application of Theorem 1.2 is to decomposable submodular function minimization:

Theorem 1.4 (Decomposable SFM). *Let $V = \{1, 2, \dots, n\}$, and let $F : 2^V \mapsto [-1, 1]$ be given by $F(S) = \sum_{i=1}^n F_i(S \cap V_i)$, each $F_i : 2^{V_i} \mapsto \mathbb{R}$ a submodular function on $V_i \subseteq V$ with $|V_i| \leq k$. We can find an ϵ -additive approximate minimizer of F in $O(nk^2 \log(nk/\epsilon))$ evaluation oracle calls.*

Theorem 1.4 significantly improves over the evaluation oracle complexity of $\tilde{O}(nk^6 \log(1/\epsilon))$ given in [AKMSV21] when the dimension k of each function F_i is large. For non-decomposable SFM, i.e. $n = 1$ and $|V_1| = k$, the current best weakly-polynomial time SFM algorithm³ finds an ϵ -approximate minimizer in time $O(k^2 \log(k/\epsilon))$ [LSW15]. Therefore, our result in Theorem 1.4 can be viewed as a generalization of the evaluation oracle complexity for non-decomposable SFM in [LSW15], and the dependence on k in Theorem 1.4 might be the best possible. We defer the details of decomposable SFM to Appendix A.

³Here, we focus on the weakly-polynomial regime, where the runtime dependence on ϵ is $\log(1/\epsilon)$.

1.1.1 Limitations

Some limitations of our algorithm are as follows: When each f_i depends on the entire d -dimensional vector θ , as opposed to a subset of the coordinates of size $d_i \ll d$, our gradient complexity simply matches that of CPM. We would like to highlight, though, that our focus is in fact the regime $d_i \ll d$. When the f_i 's are strongly-convex and smooth, our gradient complexity improves over Table 1 only when κ is large compared to d_i . Finally, note that we consider only the gradient oracle complexity in our work; our algorithm's implementation requires sampling a Hessian matrix and a gradient vector at every iteration, which incur an additional polynomial factor in the overall running time.

1.2 Technical challenges in prior works

We now describe the key technical challenges that barred existing algorithms from solving Problem 1.1 in the desired nearly-linear gradient complexity.

Gradient descent and variants. As mentioned in Section 1, the family of gradient descent algorithms presented in Table 1 are not applicable to Problem 1.1 without the smoothness assumption. When the objective in Problem 1.1 is smooth but has a large condition number, even the optimal deterministic algorithm, Accelerated Gradient Descent (AGD) [Nes83] can perform poorly. For example, when applied to (1.3), AGD updates only one coordinate in each step (thereby requiring n steps), with each step performing n gradient queries (one on each term in the problem objective), yielding a total gradient complexity of $\Omega(n^2)$ [Nes83]. For a similar reason, the fastest randomized algorithm, Katyusha [All17] also incurs a gradient complexity of $\Omega(n^2)$ [WS16].

Cutting-plane methods (CPM). Given a convex function f with its set \mathcal{S} of minimizers, CPM minimizes f by maintaining a convex search set $\mathcal{E}^{(k)} \supseteq \mathcal{S}$ in the k^{th} iteration, and iteratively shrinking $\mathcal{E}^{(k)}$ using the subgradients of f . Specifically, this is achieved by noting that for any $\mathbf{x}^{(k)}$ chosen from $\mathcal{E}^{(k)}$, if the gradient oracle indicates $\nabla f(\mathbf{x}^{(k)}) \neq 0$, (i.e. $\mathbf{x}^{(k)} \notin \mathcal{S}$), then the convexity of f guarantees $\mathcal{S} \subseteq \mathcal{H}^{(k)} \stackrel{\text{def}}{=} \{\mathbf{y} : \langle \nabla f(\mathbf{x}^{(k)}), \mathbf{y} - \mathbf{x}^{(k)} \rangle \leq 0\}$, and hence $\mathcal{S} \subseteq \mathcal{H}^{(k)} \cap \mathcal{E}^{(k)}$. The algorithm continues by choosing $\mathcal{E}^{(k+1)} \supseteq \mathcal{E}^{(k)} \cap \mathcal{H}^{(k)}$, and different choices of $\mathbf{x}^{(k)}$ and $\mathcal{E}^{(k)}$ yield different rates of shrinkage of $\mathcal{E}^{(k)}$ until a point in \mathcal{S} is found.

Solving Problem 1.1 via the current fastest CPM [JLSW20] takes $\tilde{O}(d)$ iterations, each invoking the gradient oracle on every f_i to compute $\nabla f(\mathbf{x}^{(k)}) = \sum_{i=1}^n \nabla f_i(\mathbf{x}^{(k)})$. This results in $\tilde{O}(nd)$ gradient queries overall, which can be quadratic in n when $d = \Theta(n)$ even if each f_i depends on only $d_i = O(1)$ coordinates. Similar to gradient descent and its variants, the poor performance of CPM on Problem 1.1 may therefore be attributed to their inability to query the right f_i required to make progress.

Interior-point methods (IPM). IPM solves the convex program $\min_{\mathbf{u} \in \mathcal{S}} \langle \mathbf{c}, \mathbf{u} \rangle$ by solving a sequence of unconstrained problems $\min_{\mathbf{u}} \Psi_t(\mathbf{u}) \stackrel{\text{def}}{=} \{t \cdot \langle \mathbf{c}, \mathbf{u} \rangle + \psi_{\mathcal{S}}(\mathbf{u})\}$ parametrized by increasing t , where $\psi_{\mathcal{S}}$ is a *self-concordant barrier* function that enforces feasibility by becoming unbounded as it approaches the boundary of the feasible set \mathcal{S} . The algorithm starts at $t = 0$, for which an approximate minimizer \mathbf{x}_0^* of $\psi_{\mathcal{S}}$ is known, and alternates between increasing t and updating to an approximate minimizer \mathbf{x}_t^* of the new Ψ_t via Newton's method. For a sufficiently large t , the minimizer \mathbf{x}_t^* also approximately optimizes the original problem $\min_{\mathbf{u} \in \mathcal{S}} \langle \mathbf{c}, \mathbf{u} \rangle$ with sub-optimality gap $O(\nu/t)$, where ν is the self-concordance parameter of the barrier function used.

To apply IPM to Problem 1.1, we may first transform (1.1) to $\min_{(\mathbf{u}, \mathbf{z}) \in \mathcal{K}} \sum_i \mathbf{z}_i$, where $\mathcal{K} = \{(\mathbf{u}, \mathbf{z}) : (\mathbf{u}_i, \mathbf{z}_i) \in \mathcal{K}_i, \forall i \in [n]\}$ is the feasible set. Using the universal barrier ψ_i for each \mathcal{K}_i [NN94a], the number of iterations of IPM is $\tilde{O}(\sqrt{\sum_{i=1}^n d_i})$, each requiring the computation of the Hessian and gradient of ψ_i for all $i \in [n]$, leading to a total of $\tilde{O}(n^{1.5})$ sub-gradient queries to f_i 's even when all $d_i = O(1)$. Even when leveraging the recent framework of robust IPM for linear programs [LV21], the computation of each Hessian (by sampling the corresponding \mathcal{K}_i [JLLV21]) yields a total sub-gradient oracle complexity of $\tilde{O}(\sum_{i=1}^n d_i^{3.5})$, far from the complexity we seek.

1.3 Our algorithmic framework

We now give an overview of the techniques developed in this work to overcome the above barriers. First, we transform (1.1) into a convex program over structured convex sets:

$$\begin{aligned} & \text{minimize} && \langle \mathbf{c}, \mathbf{x} \rangle, \\ & \text{subject to} && \mathbf{x}_i \in \mathcal{K}_i \subseteq \mathbb{R}^{d_i+1} \quad \forall i \in [n] \\ & && \mathbf{A}\mathbf{x} = \mathbf{b}. \end{aligned} \tag{1.4}$$

where \mathbf{x} is the concatenation of the vectors $\mathbf{x}_1, \dots, \mathbf{x}_n$, and notably the convex sets \mathcal{K}_i are all disjoint. Under this transformation we do not have explicit, closed-form expressions for each \mathcal{K}_i ; however, the subgradient oracle for f_i can be transformed equivalently to a separation oracle \mathcal{K}_i . We define $\mathcal{K} \stackrel{\text{def}}{=} \mathcal{K}_1 \times \mathcal{K}_2 \times \dots \times \mathcal{K}_n$.

Main idea: combining CPM and IPM. Recall that CPM maintains a convex set which initially contains the feasible region and gradually shrinks around the minimizer, while IPM maintains a point inside the feasible region that moves toward the minimizer. Our novel idea is to combine both methods and maintain an inner convex set $\mathcal{K}_{\text{in},i}$ as well as an outer convex set $\mathcal{K}_{\text{out},i}$ for each $i \in [n]$, such that $\mathcal{K}_{\text{in},i} \subseteq \mathcal{K}_i \subseteq \mathcal{K}_{\text{out},i}$. We define \mathcal{K}_{in} and \mathcal{K}_{out} analogously to \mathcal{K} . When Inequality (3.4) and Inequality (3.3) are satisfied for all $i \in [n]$, we make IPM-style updates without needing to make any oracle calls. When Inequality (3.3) is violated for some $i \in [n]$, we query the separation oracle at the point $\mathbf{x}_{\text{out},i}^*$ defined as the centroid of $\mathcal{K}_{\text{out},i}$ (c.f. Proposition 3.2). Based on the oracle’s response, we iteratively either grow $\mathcal{K}_{\text{in},i}$ (and, thus, \mathcal{K}_{in}) outward or shrink $\mathcal{K}_{\text{out},i}$ (and, thus, \mathcal{K}_{out}) inward, until ultimately they approximate \mathcal{K} around the optimum point.

First benefit: large change in volume. If the point $\mathbf{x}_{\text{out},i}^*$ violates Inequality (3.3) for some $i \in [n]$, we query the separation oracle to see if $\mathbf{x}_{\text{out},i}^* \in \mathcal{K}_i$ or not. If $\mathbf{x}_{\text{out},i}^* \in \mathcal{K}_i$, then it is used to expand $\mathcal{K}_{\text{in},i}$, yielding in a large volume increase for $\mathcal{K}_{\text{in},i}$. On the other hand, if $\mathbf{x}_{\text{out},i}^* \notin \mathcal{K}_i$, the fact that it is the centroid of $\mathcal{K}_{\text{out},i}$ results in a large volume decrease for $\mathcal{K}_{\text{out},i}$ when it is intersected with a halfspace through $\mathbf{x}_{\text{out},i}^*$. Thus, our algorithm witnesses a large change in volume of one of $\mathcal{K}_{\text{in},i}$ and $\mathcal{K}_{\text{out},i}$, regardless of the answer from the oracle. Just like in standard CPM, this rapid change in volume is crucial to achieving the algorithm’s oracle complexity.

Second benefit: making a smart choice about querying f_i . Since the algorithm maintains both an inner and outer set approximating \mathcal{K} , by checking if $\mathcal{K}_{\text{in},i}$ and $\mathcal{K}_{\text{out},i}$ differ significantly (Inequality (3.3) essentially performs this function), we can determine if \mathcal{K}_i is poorly approximated, and if so, improve the inner and outer approximations of the true feasible set. Choosing the right \mathcal{K}_i translates to choosing the right f_i to make progress with at an iteration; thus, we address the central weakness of the gradient descent variants in solving (1.1).

2 Notation and preliminaries

We lay out the notation used in our paper as well as the definitions and prior known results that we rely on. We use lowercase boldface letters to denote (column) vectors and uppercase boldface letters to denote matrices. We use \mathbf{x}_i to denote the i^{th} block of coordinates in the vector \mathbf{x} (the ordering of these blocks is not important in our setup). We use \succeq and \preceq to denote the Loewner ordering of matrices.

We use $\langle \mathbf{x}, \mathbf{y} \rangle$ to mean the Euclidean inner product $\mathbf{x}^\top \mathbf{y}$. A subscript \mathbf{x} in the inner product notation means it is induced by the Hessian of some function (which is clear from context) at \mathbf{x} ; for example, $\langle \mathbf{u}, \mathbf{v} \rangle_{\mathbf{x}} = \mathbf{u}^\top \nabla_{ii}^2 \psi(\mathbf{x}) \mathbf{v}$ with ψ inferred from context. We define the local norm of \mathbf{v} at \mathbf{x} analogously: $\|\mathbf{v}\|_{\mathbf{x}} = \sqrt{\langle \mathbf{v}, \nabla^2 \psi(\mathbf{x}) \cdot \mathbf{v} \rangle}$. We also define the norm $\|\mathbf{v}\|_{\mathbf{x},1} \stackrel{\text{def}}{=} \sum_{i=1}^n \|\mathbf{v}\|_{\mathbf{x}_i}$.

We use ψ to represent barrier functions and Φ to represent potential functions, with appropriate subscripts and superscripts to qualify them as needed.

2.1 Facts from convex analysis

In this section, we present some definitions and properties from convex analysis that are useful in our paper. These results are standard and may be found in, for example, [Roc70; BBV04].

Definition 2.1. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$. Then the function $f^* : \mathbb{R}^n \rightarrow \mathbb{R}$ defined as

$$f^*(\mathbf{y}) = \sup_{\mathbf{x} \in \text{dom}(f)} [\langle \mathbf{x}, \mathbf{y} \rangle - f(\mathbf{x})]$$

is called the Fenchel conjugate of the function f . An immediate consequence of the definition (and by applying the appropriate convexity-preserving property) is that f^* is convex, regardless of the convexity of f . We use the superscript $*$ on functions to denote their conjugates.

Lemma 2.2 (Biconjugacy). *For a closed, convex function f , we have $f = f^{**}$.*

Lemma 2.3 ([Roc70]). *For a closed, convex differentiable function f , we have $\mathbf{y} = \nabla f(\mathbf{x}) \iff \mathbf{x} = \nabla f^*(\mathbf{y})$.*

Lemma 2.4 ([Roc70]). *For a strictly convex, twice-differentiable function f , we have $\nabla^2 f^*(\nabla f(\mathbf{x})) = (\nabla^2 f(\mathbf{x}))^{-1}$.*

Definition 2.5 (Polar of a Set [RW09]). Given a set $\mathcal{S} \subseteq \mathbb{R}^n$, its polar is defined as

$$\mathcal{S}^\circ \stackrel{\text{def}}{=} \{\mathbf{y} \in \mathbb{R}^n : \langle \mathbf{y}, \mathbf{x} \rangle \leq 1, \forall \mathbf{x} \in \mathcal{S}\}.$$

Lemma 2.6 ([Roc70]). *Let $\mathcal{S} \subseteq \mathbb{R}^n$ be a closed, compact, convex set, and let \mathbf{y} be a point. Then $(\text{conv}\{\mathcal{S}, \mathbf{y}\})^\circ \subseteq \mathcal{S}^\circ \cap \mathcal{H}$, where \mathcal{H} is the halfspace defined by $\mathcal{H} = \{\mathbf{z} \in \mathbb{R}^n : \langle \mathbf{z}, \mathbf{y} \rangle \leq 1\}$.*

2.2 Background on interior-point methods

Our work draws heavily upon geometric properties of self-concordant functions, which underpin the rich theory of interior-point methods. We list below the formal results needed for our analysis, and refer the reader to [NN94a; Ren01] for a detailed exposition of this function class. We begin with the definitions of self-concordant functions and self-concordant barriers:

Definition 2.7 (Self-concordance [NN94a]). A function $F : Q \mapsto \mathbb{R}$ is a self-concordant function on a convex set Q if for any $\mathbf{x} \in Q$ and any \mathbf{h} ,

$$|D^3 F(\mathbf{x})[\mathbf{h}, \mathbf{h}, \mathbf{h}]| \leq 2(D^2 F(\mathbf{x})[\mathbf{h}, \mathbf{h}])^{3/2},$$

where $D^k F(\mathbf{x})[\mathbf{h}_1, \dots, \mathbf{h}_k]$ denotes the k -th derivative of F at \mathbf{x} along the directions $\mathbf{h}_1, \dots, \mathbf{h}_k$. We say F is a ν -self-concordant barrier if F further satisfies $\nabla F(\mathbf{x})^\top (\nabla^2 F(\mathbf{x}))^{-1} \nabla F(\mathbf{x}) \leq \nu$ for any $\mathbf{x} \in Q$.

Theorem 2.8 (Theorem 2.3.3 from [Ren01]). *If f is a self-concordant barrier, then for all \mathbf{x} and $\mathbf{y} \in \text{dom}(f)$, we have $\langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle \leq \nu$, where ν is the self-concordance of f .*

Theorem 2.9 (Theorem 2.3.4 from [Ren01]). *If f is a ν -self-concordant barrier such that $\mathbf{x}, \mathbf{y} \in \text{dom}(f)$ satisfy $\langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle \geq 0$, then $\mathbf{y} \in \mathcal{B}_{\mathbf{x}}(\mathbf{x}, 4\nu + 1)$.*

We now state the following result from self-concordance calculus.

Theorem 2.10 (Theorem 3.3.1 of [Ren01]). *If f is a (strongly nondegenerate) self-concordant function, then so is its Fenchel conjugate f^* .*

The following result gives a bound on the quadratic approximation of a function, with the distance between two points measured in the local norm. The convergence of Newton's method can be essentially explained by this result.

Theorem 2.11 (Theorem 2.2.2 of [Ren01]). *If f is a self-concordant function, $\mathbf{x} \in \text{dom}(f)$, and $\mathbf{y} \in \mathcal{B}_{\mathbf{x}}(\mathbf{x}, 1)$, then*

$$f(\mathbf{y}) \leq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{1}{2} \|\mathbf{y} - \mathbf{x}\|_{\mathbf{x}}^2 + \frac{\|\mathbf{y} - \mathbf{x}\|_{\mathbf{x}}^3}{3(1 - \|\mathbf{y} - \mathbf{x}\|_{\mathbf{x}})},$$

where $\|\mathbf{y} - \mathbf{x}\|_{\mathbf{x}}^2 \stackrel{\text{def}}{=} \langle \mathbf{y} - \mathbf{x}, \nabla^2 f(\mathbf{x}) \cdot (\mathbf{y} - \mathbf{x}) \rangle$.

Finally, we need the following definitions of entropic barrier and universal barrier.

Definition 2.12 ([BE15; Che21]). Given a convex body $\mathcal{K} \subseteq \mathbb{R}^n$ and some fixed $\theta \in \mathbb{R}^n$, define the function $f(\theta) = \log \left[\int_{\mathbf{x} \in \mathcal{K}} \exp(\langle \mathbf{x}, \theta \rangle) d\mathbf{x} \right]$. Then the Fenchel conjugate $f^* : \text{int}(\mathcal{K}) \rightarrow \mathbb{R}$ is a self-concordant barrier termed the *entropic barrier*. The entropic barrier is n -self-concordant.

Definition 2.13 ([NN94b; LY21]). Given a convex body $\mathcal{K} \subseteq \mathbb{R}^n$, the *universal barrier* of \mathcal{K} is defined as $\psi : \text{int}(\mathcal{K}) \rightarrow \mathbb{R}$ by

$$\psi(\mathbf{x}) = \log \text{vol}((\mathcal{K} - \mathbf{x})^\circ)$$

where $(\mathcal{K} - \mathbf{x})^\circ = \{\mathbf{y} \in \mathbb{R}^n : \mathbf{y}^\top (\mathbf{z} - \mathbf{x}) \leq 1, \forall \mathbf{z} \in \mathcal{K}\}$ is the polar set of \mathcal{K} with respect to \mathbf{x} . The universal barrier is n -self-concordant.

2.3 Facts from convex geometry

Since our analysis is contingent on the change in the volume of convex bodies when points are added to them or when they are intersected with halfspaces, we invoke the classical result by Grünbaum several times. We therefore state its relevant variants next: [Theorem 2.14](#) applies to log-concave distributions, and [Corollary 2.16](#) is its specific case, since the indicator function of a convex set is a log-concave function [BBV04].

Theorem 2.14 ([Grü60; BKLLS20]). *Let f be a log-concave distribution on \mathbb{R}^d with centroid \mathbf{c}_f . Let $\mathcal{H} = \{\mathbf{u} \in \mathbb{R}^d : \langle \mathbf{u}, \mathbf{v} \rangle \geq q\}$ be a halfspace defined by a normal vector $\mathbf{v} \in \mathbb{R}^d$. Then, $\int_{\mathcal{H}} f(\mathbf{z}) d\mathbf{z} \geq \frac{1}{e} - t^+$, where $t = \frac{q - \langle \mathbf{c}_f, \mathbf{v} \rangle}{\sqrt{\mathbb{E}_{\mathbf{y} \sim f} \langle \mathbf{v}, \mathbf{y} - \mathbf{c}_f \rangle^2}}$ is the distance of the centroid to the halfspace scaled by the standard deviation along the normal vector \mathbf{v} and $t^+ \stackrel{\text{def}}{=} \max\{0, t\}$.*

Remark 2.15. A crucial special case of [Theorem 2.14](#) is that cutting a convex set through its centroid yields two parts, the smaller of which has volume at least $1/e$ times the original volume and the larger of which is at most $1 - 1/e$ times the original total volume.

Corollary 2.16 ([Grü60]). *Let \mathcal{K} be a convex set with centroid μ and covariance matrix Σ . Then, for any point \mathbf{x} satisfying $\|\mathbf{x} - \mu\|_{\Sigma^{-1}} \leq \eta$ and a halfspace \mathcal{H} such that $\mathbf{x} \in \mathcal{H}$, we have $\text{vol}(\mathcal{K} \cap \mathcal{H}) \geq \text{vol}(\mathcal{K}) \cdot (1/e - \eta)$.*

Finally, we need the following facts.

Fact 2.17 (Volumes of standard objects). *The volume of a q -dimensional Euclidean ball is given by $\text{vol}(\mathcal{B}_q(0, \bar{R})) = \frac{\pi^{q/2}}{\Gamma(1+q/2)} \bar{R}^q$, and the volume of a q -dimensional cone = $\frac{1}{q+1} \cdot \text{volume of base} \cdot \text{height}$.*

3 Our algorithm

We begin by reducing [Problem 1.1](#) to the following slightly stronger formulation (see [Theorem 4.11](#) for the detailed reduction):

$$\begin{aligned} & \text{minimize} && \langle \mathbf{c}, \mathbf{x} \rangle, \\ & \text{subject to} && \mathbf{x}_i \in \mathcal{K}_i \subseteq \mathbb{R}^{d_i+1} \quad \forall i \in [n] \\ & && \mathbf{A}\mathbf{x} = \mathbf{b}. \end{aligned} \tag{3.1}$$

where \mathbf{x} is a concatenation of vectors \mathbf{x}_i 's, and the \mathcal{K}_i 's are disjoint convex sets. This formulation decouples the overlapping support of the original f_i 's by introducing additional variables tied together through the linear system $\mathbf{A}\mathbf{x} = \mathbf{b}$. Each \mathcal{K}_i is constructed by applying a standard epigraph trick to the function f_i .

Note that we do not have a closed-form expression for \mathcal{K}_i . Instead, the subgradient oracle for f_i translates to a *separation oracle* for \mathcal{K}_i : on a point \mathbf{z}_i queried by the oracle, the oracle either asserts $\mathbf{z}_i \in \mathcal{K}_i$, or returns a separating hyperplane that separates \mathbf{z}_i from \mathcal{K}_i .

At the start of our algorithm, we have the following guarantee:

Lemma 3.1. *At the start of our algorithm, we are guaranteed the existence of the following.*

1. *Explicit convex sets $\mathcal{K}_{in} \stackrel{\text{def}}{=} \mathcal{K}_{in,1} \times \mathcal{K}_{in,2} \times \cdots \times \mathcal{K}_{in,n}$ and $\mathcal{K}_{out} \stackrel{\text{def}}{=} \mathcal{K}_{out,1} \times \mathcal{K}_{out,2} \times \cdots \times \mathcal{K}_{out,n}$ such that $\mathcal{K}_{in} \subseteq \mathcal{K} \stackrel{\text{def}}{=} \mathcal{K}_1 \times \cdots \times \mathcal{K}_n \subseteq \mathcal{K}_{out}$,*
2. *An initial $\mathbf{x}_{initial} \in \mathcal{K}_{in}$ such that $\mathbf{A}\mathbf{x}_{initial} = \mathbf{b}$.*

We show how to construct such a set \mathcal{K}_{in} in [Section 5.1](#) and how to find such a \mathcal{K}_{out} and $\mathbf{x}_{initial}$ in [Section 5.2](#).

3.1 Details of our algorithm

In this section, we explain our main algorithm ([Algorithm 1](#)).

The inputs to [Algorithm 1](#) are: initial sets \mathcal{K}_{in} and \mathcal{K}_{out} satisfying $\mathcal{K}_{in} \subseteq \mathcal{K} \subseteq \mathcal{K}_{out}$, an initial point $\mathbf{x} \in \mathcal{K}_{in}$ satisfying $\mathbf{A}\mathbf{x} = \mathbf{b}$, a separation oracle \mathcal{O}_i for each \mathcal{K}_i , the objective vector \mathbf{c} , and scalar parameters m, n, R, r , and ϵ . All the parameters are set in the proof of [Theorem 4.10](#).

Throughout the algorithm, we maintain a central path parameter t for IPM-inspired updates, the current solution \mathbf{x} , and convex sets $\mathcal{K}_{in,i}$ and $\mathcal{K}_{out,i}$ satisfying $\mathcal{K}_{in,i} \subseteq \mathcal{K}_i \subseteq \mathcal{K}_{out,i}$ for each $i \in [n]$. To run IPM-style updates, we choose the entropic barrier on \mathcal{K}_{out} and the universal barrier on \mathcal{K}_{in} .

Given the current set \mathcal{K}_{out} , the current t , and the entropic barrier ψ_{out} defined on $\widehat{\mathcal{K}}_{out} \stackrel{\text{def}}{=} \mathcal{K}_{out} \cap \{\mathbf{u} : \mathbf{A}\mathbf{u} = \mathbf{b}\}$, we define the point

$$\mathbf{x}_{out}^* \stackrel{\text{def}}{=} \arg \min_{\mathbf{x} \in \widehat{\mathcal{K}}_{out}} \{t \langle \mathbf{c}, \mathbf{x} \rangle + \psi_{out}(\mathbf{x})\}. \quad (3.2)$$

Per the IPM paradigm, for the current value of t , this point serves as a target to “chase” when optimizing $\langle \mathbf{c}, \mathbf{x} \rangle$ over the set $\widehat{\mathcal{K}}_{out}$. Although our overall goal in [Problem 3.1](#) is to optimize over $\mathcal{K} \cap \{\mathbf{u} : \mathbf{A}\mathbf{u} = \mathbf{b}\}$, we do not know \mathcal{K} explicitly and therefore must use its known proxies, \mathcal{K}_{in} or \mathcal{K}_{out} ; we choose \mathcal{K}_{out} because $\mathcal{K}_{out} \supseteq \mathcal{K}$ ensures we do not miss a potential optimal point.

Having computed the current target \mathbf{x}_{out}^* , we move the current solution \mathbf{x} towards it by taking a Newton step, provided certain conditions of feasibility and minimum progress are satisfied. If either one of these conditions is violated, we update either \mathcal{K}_{in} , \mathcal{K}_{out} , or the parameter t .

Updating \mathbf{x} . In order to move \mathbf{x} towards \mathbf{x}_{out}^* , we require two conditions to hold: $\mathbf{x}_{out}^* \in \mathcal{K}_{in}$ and $\langle \mathbf{c}, \mathbf{x} \rangle \geq \langle \mathbf{c}, \mathbf{x}_{out}^* \rangle + O(1/t)$.

The first condition implies $\mathbf{x}_{out}^* \in \mathcal{K}$, which would in turn ensure feasibility of the resulting \mathbf{x} after a Newton step. To formally check this condition, we check if the following inequality is satisfied for all $i \in [n]$ and for a fixed constant η :

$$\langle \nabla \psi_{in,i}(\mathbf{x}_i), \mathbf{x}_{out,i}^* - \mathbf{x}_i \rangle + \eta \cdot \|\mathbf{x}_{out,i}^* - \mathbf{x}_i\|_{\mathbf{x}_i} \leq 4d_i. \quad (3.3)$$

The intuition is that since any point within the domain of a self-concordant barrier satisfies the inequalities in [Theorem 2.8](#) and [Theorem 2.9](#), violating [Inequality \(3.3\)](#) implies that $\mathbf{x}_{out,i}^*$ is far from $\mathcal{K}_{in,i}$, and as a result, \mathbf{x}_{out}^* is not a good candidate to move \mathbf{x} towards.

The second condition we impose, one of “sufficient suboptimality”, ensures significant progress in the objective value can be made when updating \mathbf{x} . Formally, we check if

$$\mathbf{c}^\top \mathbf{x}_{out}^* + \frac{4m}{t} \leq \mathbf{c}^\top \mathbf{x}. \quad (3.4)$$

If the inequality holds, then there is still “room for progress” to lower the value of $\langle \mathbf{c}, \mathbf{x} \rangle$ by updating \mathbf{x} ; if the inequality is violated, we update t instead.

Algorithm 1 Minimizing Decomposable Convex Function

```

1: ▷ Solving Problem 3.1
2: Input.  $\epsilon, \mathbf{A}, \mathbf{b}, \mathbf{c}, R, r, m, n, \mathbf{x}, \mathcal{K}_{\text{in}}, \mathcal{K}_{\text{out}},$  and  $\mathcal{O}_i$  for each  $i \in [n]$ .
3: Initialization.  $t = \frac{m \log m}{\sqrt{n} \|\mathbf{c}\|_2 R}$  and  $t_{\text{end}} = \frac{8m}{\epsilon \|\mathbf{c}\|_2 R}$ ,  $\eta = \frac{1}{100}$ , and  $\mathbf{x}_{\text{out}}^*$  (via Equation \(3.2\))
4: while true do
5:   if  $\langle \mathbf{c}, \mathbf{x} \rangle \leq \langle \mathbf{c}, \mathbf{x}_{\text{out}}^* \rangle + \frac{4m}{t}$  then                                ▷ Either update  $t$  or end the algorithm
6:     if  $t \geq t_{\text{end}}$  then
7:       return  $\arg \min_{\mathbf{x}: \mathbf{x} \in \mathcal{K}_{\text{in}}, \mathbf{Ax}=\mathbf{b}} \{t \langle \mathbf{c}, \mathbf{x} \rangle + \sum_{i=1}^n \psi_{\text{in},i}(\mathbf{x}_i)\}$ .                                ▷ End the algorithm
8:     end if
9:      $t \leftarrow t \cdot \left[1 + \frac{\eta}{4m}\right]$                                                                     ▷ Update  $t$ 
10:    Update  $\mathbf{x}_{\text{out}}^*$  and jump to Line 4                                                                    ▷  $\mathbf{x}_{\text{out}}^*$  computed as as per Equation \(3.2\)
11:  end if
12:  for all  $i \in [n]$  do
13:    if  $\langle \nabla \psi_{\text{in},i}(\mathbf{x}_i), \mathbf{x}_{\text{out},i}^* - \mathbf{x}_i \rangle + \eta \|\mathbf{x}_{\text{out},i}^* - \mathbf{x}_i\|_{\mathbf{x}_i} \geq 4d_i$  then
14:      if  $\mathbf{x}_{\text{out},i}^* \in \mathcal{K}_i$  then                                                                    ▷ Query  $\mathcal{O}_i$ 
15:         $\mathcal{K}_{\text{in},i} = \text{conv}(\mathcal{K}_{\text{in},i}, \mathbf{x}_{\text{out},i}^*)$                                                                     ▷ Update  $\mathcal{K}_{\text{in},i}$ 
16:      else
17:         $\mathcal{K}_{\text{out},i} = \mathcal{K}_{\text{out},i} \cap \mathcal{H}_i$ , where  $\mathcal{H}_i = \mathcal{O}_i(\mathbf{x}_{\text{out},i}^*)$                                 ▷ Update  $\mathcal{K}_{\text{out},i}$ 
18:      end if
19:      Update  $\mathbf{x}_{\text{out}}^*$  and jump to Line 4                                                                    ▷  $\mathbf{x}_{\text{out}}^*$  computed as per Equation \(3.2\)
20:    end if
21:  end for
22:  Set  $\delta_{\mathbf{x}} \stackrel{\text{def}}{=} \frac{\eta}{2} \cdot \frac{\mathbf{x}_{\text{out}}^* - \mathbf{x}}{\|\mathbf{x}_{\text{out}}^* - \mathbf{x}\|_{\mathbf{x},1}}$ , where  $\|\mathbf{u}\|_{\mathbf{x},1} \stackrel{\text{def}}{=} \sum_{i=1}^n \|\mathbf{u}\|_{\mathbf{x}_i}$ .
23:   $\mathbf{x} \leftarrow \mathbf{x} + \delta_{\mathbf{x}}$                                                                     ▷ Move  $\mathbf{x}$  towards  $\mathbf{x}_{\text{out}}^*$ 
24: end while

```

Given the two conditions hold, we move \mathbf{x} towards $\mathbf{x}_{\text{out}}^*$ in [Line 23](#). The update step is normalized by the distance between \mathbf{x} and $\mathbf{x}_{\text{out}}^*$ measured in the local norm, which enforces $\mathbf{x} \in \mathcal{K}$ (since by the definition of self-concordance, the unit radius Dikin ball lies inside the domain of the self-concordance barrier), and also helps bound certain first-order error terms ([Inequality \(4.14\)](#) in [Section 4.3](#)).

The rest of this section details the procedure for when either of these conditions is violated.

Updating the inner and outer convex sets. Suppose [Inequality \(3.3\)](#) is violated for some $i \in [n]$. Then $\mathbf{x}_{\text{out},i}^* \notin \mathcal{K}_{\text{in},i}$, which in turn means $\mathbf{x}_{\text{out}}^*$ *might* not be in the feasible set \mathcal{K} . To reestablish [Inequality \(3.3\)](#) for i , we can either update $\mathcal{K}_{\text{in},i}$, or update $\mathcal{K}_{\text{out},i}$ and compute a new $\mathbf{x}_{\text{out},i}^*$ by [Equation \(3.2\)](#).

To decide which option to take, we query \mathcal{O}_i at the point $\mathbf{x}_{\text{out},i}^*$: if the oracle indicates that $\mathbf{x}_{\text{out},i}^* \in \mathcal{K}_i$, then we incorporate $\mathbf{x}_{\text{out},i}^*$ into $\mathcal{K}_{\text{in},i}$ by redefining $\mathcal{K}_{\text{in},i} = \text{conv}(\mathcal{K}_{\text{in},i}, \mathbf{x}_{\text{out},i}^*)$ to be the convex hull of the current $\mathcal{K}_{\text{in},i}$ and $\mathbf{x}_{\text{out},i}^*$ ([Line 15](#)). If, on the other hand, $\mathbf{x}_{\text{out},i}^* \notin \mathcal{K}_i$, the oracle \mathcal{O}_i will return a halfspace \mathcal{H}_i satisfying $\mathcal{H}_i \supseteq \mathcal{K}_i$. Then we redefine $\mathcal{K}_{\text{out},i} = \mathcal{K}_{\text{out},i} \cap \mathcal{H}_i$ ([Line 17](#)). After processing this update of the sets, the algorithm recomputes $\mathbf{x}_{\text{out}}^*$ and returns to the main loop since updating the sets does not necessarily imply that the new $\mathbf{x}_{\text{out}}^*$ satisfies $\mathbf{x}_{\text{out}}^* \in \mathcal{K}_{\text{in}}$.

This update rule for the sets is exactly where our novelty lies: *we do not arbitrarily update sets, rather, we update one only after checking the very specific condition $\mathbf{x}_{\text{out},i}^* \notin \mathcal{K}_{\text{in},i}$* . Since the separation oracle is called only in this part of the algorithm, performing this check enables us to dramatically reduce the number of calls we make to the separation oracle, thereby improving our oracle complexity.

Further, this update rule shows that even when we cannot update the current \mathbf{x} , we make progress by using all the information from the oracles. Over the course of the algorithm, we gradually expand \mathcal{K}_{in} and shrink \mathcal{K}_{out} , until they well-approximate \mathcal{K} . To formally quantify the change in volume due to the above operations, we consider the following alternative view of $\mathbf{x}_{\text{out}}^*$.

Proposition 3.2 (Section 3 in [\[BE15\]](#); Section 3 of [\[Kla06\]](#)). *Let $\theta \in \mathbb{R}^n$, and let p_θ be defined as $p_\theta(\mathbf{x}) \stackrel{\text{def}}{=} \dots$*

$\exp(\langle \theta, \mathbf{x} \rangle - f(\theta))$, where $f(\theta) \stackrel{\text{def}}{=} \log \left[\int_{\mathcal{K}} \exp(\langle \theta, \mathbf{u} \rangle) d\mathbf{u} \right]$. Then,

$$\mathbb{E}_{\mathbf{x} \sim p_\theta}[\mathbf{x}] = \arg \min_{\mathbf{x} \in \text{int}(\mathcal{K})} \{f^*(\mathbf{x}) - \langle \theta, \mathbf{x} \rangle\}.$$

By this proposition, $\mathbf{x}_{\text{out}}^*$ defined in Equation (3.2) satisfies

$$\mathbf{x}_{\text{out}}^* \stackrel{\text{def}}{=} \mathbb{E}_{\mathbf{x} \sim \exp\{-t\langle \mathbf{c}, \mathbf{x} \rangle - \log \left[\int_{\widehat{\mathcal{K}}_{\text{out}}} \exp(-t\langle \mathbf{c}, \mathbf{u} \rangle) d\mathbf{u} \right]\}}[\mathbf{x}], \quad (3.5)$$

that is, $\mathbf{x}_{\text{out}}^*$ is the centroid of some exponential distribution over $\widehat{\mathcal{K}}_{\text{out}}$. As a result, if $\mathbf{x}_{\text{out},i}^* \notin \mathcal{K}_i$, the hyperplane cutting $\widehat{\mathcal{K}}_{\text{out}}$ through $\mathbf{x}_{\text{out}}^*$ will yield a large decrease in volume of $\widehat{\mathcal{K}}_{\text{out}}$, per Remark 2.15. Therefore, the query result in a large change in volume in either \mathcal{K}_{in} or \mathcal{K}_{out} , allowing us to approximate \mathcal{K} with a bounded number of iterations.

Updating t . If Inequality (3.4) is violated, then the current \mathbf{x} is “as optimal as one can get” for the current parameter t . This could mean one of two things:

The first possibility is that we have already reached an approximate optimum, which we verify by checking whether $t \geq O(1/\epsilon)$ in Line 6: If true, this indicates that we have attained our desired suboptimality, and the algorithm terminates by returning

$$\mathbf{x}^{\text{ret}} = \arg \min_{\mathbf{x}: \mathbf{x} \in \mathcal{K}_{\text{in}}, \mathbf{A}\mathbf{x}=\mathbf{b}} \left\{ t \cdot \langle \mathbf{c}, \mathbf{x} \rangle + \sum_{i=1}^n \psi_{\text{in},i}(\mathbf{x}_i) \right\}.$$

The point \mathbf{x}^{ret} is feasible because it is in \mathcal{K}_{in} by definition, and the suboptimality of $O(1/t_{\text{end}}) = O(\epsilon)$ ensures it is an approximate optimum for the original problem.

The second possibility is that we need to increase t to set the next “target suboptimality”. The value of t is increased by a scaling factor of $1 + O(1/m)$ in Line 9. This scaling factor ensures, like in the standard IPM framework, that the next optimum is not too far from the current one. Following the update to t , we recompute $\mathbf{x}_{\text{out}}^*$ by Equation (3.2). Since $\langle \mathbf{c}, \mathbf{x} \rangle > \langle \mathbf{c}, \mathbf{x}_{\text{out}}^* \rangle + O(1/t)$ is not guaranteed with the new t and $\mathbf{x}_{\text{out}}^*$, the algorithm jumps back to the start of the main loop.

4 Our analysis

To analyze Algorithm 1, we define the following potential function that captures the changes in $\mathcal{K}_{\text{in},i}$, $\mathcal{K}_{\text{out},i}$, t , and \mathbf{x} in each iteration:

$$\Phi \stackrel{\text{def}}{=} \underbrace{t\langle \mathbf{c}, \mathbf{x} \rangle + \log \left[\int_{\widehat{\mathcal{K}}_{\text{out}}} \exp(-t\langle \mathbf{c}, \mathbf{u} \rangle) d\mathbf{u} \right]}_{\text{entropic terms}} + \underbrace{\sum_{i \in [n]} \psi_{\text{in},i}(\mathbf{x}_i)}_{\text{universal terms}}, \quad (4.1)$$

where $\log \left[\int_{\widehat{\mathcal{K}}_{\text{out}}} \exp(-t\langle \mathbf{c}, \mathbf{u} \rangle) d\mathbf{u} \right]$ is related to the entropic barrier on $\widehat{\mathcal{K}}_{\text{out}}$ (see Section 4.1) and $\psi_{\text{in},i}$ is the universal barrier on \mathcal{K}_i . In the subsequent sections, we study the changes in each of these potential functions along with obtaining bounds on the initial and final potentials and combine them to bound the algorithm’s separation oracle complexity.

4.1 Potential change for the entropic terms

In this section, we study the changes in the entropic terms of Equation (4.1) upon updating the outer convex set $\widehat{\mathcal{K}}_{\text{out}}$ as well as t . These two changes are lumped together in this section because both updates affect the term $\log \left[\int_{\widehat{\mathcal{K}}_{\text{out}}} \exp(-t \cdot \langle \mathbf{c}, \mathbf{x} \rangle) d\mathbf{x} \right]$, albeit in different ways: the update in $\widehat{\mathcal{K}}_{\text{out}}$ affects it via Grünbaum’s Theorem; the update in t affects it via the fact that, by duality with respect to the entropic barrier (Definition 2.12), $\log \left[\int_{\mathbf{x} \in \widehat{\mathcal{K}}_{\text{out}}} \exp(\langle \mathbf{x}, \theta \rangle) d\mathbf{x} \right]$ is also self-concordant. We detail these two potential changes below.

Lemma 4.1 (Potential analysis for outer set). *Let $\widehat{\mathcal{K}}_{out} \stackrel{def}{=} \{\mathbf{x} : \mathbf{x}_i \in \mathcal{K}_{out,i} \cap \{\mathbf{y} : \mathbf{A}\mathbf{y} = \mathbf{b}\}\}$, and let $\Phi = \log \left[\int_{\widehat{\mathcal{K}}_{out}} \exp(-t \langle \mathbf{c}, \mathbf{u} \rangle) d\mathbf{u} \right] + t \langle \mathbf{c}, \mathbf{x} \rangle + \sum_{i \in [n]} \psi_{in,i}(\mathbf{x}_i)$. Let \mathcal{H}_i be the halfspace generated by the separation oracle \mathcal{O}_i queried at $\mathbf{x}_{out,i}^*$ as shown in [Line 17 of Algorithm 1](#). Then the new potential $\Phi^{(new)} = \log \left[\int_{\widehat{\mathcal{K}}_{out} \cap \mathcal{H}_i} \exp(-t \langle \mathbf{c}, \mathbf{u} \rangle) d\mathbf{u} \right] + t \langle \mathbf{c}, \mathbf{x} \rangle + \sum_{i \in [n]} \psi_{in,i}(\mathbf{x}_i)$ is bounded from above as follows.*

$$\Phi^{(new)} \leq \Phi + \log(1 - 1/e).$$

Proof. The change in potential is given by

$$\Phi^{(new)} - \Phi = \log \left[\frac{\int_{\widehat{\mathcal{K}}_{out} \cap \mathcal{H}_i} \exp(-t \cdot \langle \mathbf{c}, \mathbf{x} \rangle) d\mathbf{x}}{\int_{\widehat{\mathcal{K}}_{out}} \exp(-t \cdot \langle \mathbf{c}, \mathbf{x} \rangle) d\mathbf{x}} \right].$$

We now apply [Theorem 2.14](#) to the right hand side, with the function $f(\mathbf{x}) = \exp(-t \cdot \langle \mathbf{c}, \mathbf{x} \rangle - A(t\mathbf{c}))$, where $A(\theta) = \log \left[\int_{\widehat{\mathcal{K}}_{out}} \exp(-\langle \theta, \mathbf{x} \rangle) d\mathbf{x} \right]$. Noting that each halfspace \mathcal{H}_i passes directly through $\mathbf{x}_{out,i}^*$, where \mathbf{x}_{out}^* is the centroid of $\widehat{\mathcal{K}}_{out}$ with respect to f (by the definition of \mathbf{x}_{out}^* in [Equation \(3.5\)](#)), [Remark 2.15](#) applies and gives the claimed volume change. \square

To capture the change in potential due to the update in t , we recall the alternative perspective to the function $\log \left[\int_{\widehat{\mathcal{K}}_{out}} \exp(-t \langle \mathbf{c}, \mathbf{x} \rangle) d\mathbf{x} \right]$ given by [Definition 2.12](#) and derive properties of self-concordant barriers.

Lemma 4.2. *Consider a ν -self-concordant barrier $\psi : \text{int}(\mathcal{K}) \rightarrow \mathbb{R}$ over the interior of a convex set $\mathcal{K} \subseteq \mathbb{R}^d$. Define*

$$\xi_t^\psi \stackrel{def}{=} \min_{\mathbf{x}} [t \cdot \langle \mathbf{c}, \mathbf{x} \rangle + \psi(\mathbf{x})] \quad \text{and} \quad \mathbf{x}_t \stackrel{def}{=} \arg \min_{\mathbf{x}} [t \cdot \langle \mathbf{c}, \mathbf{x} \rangle + \psi(\mathbf{x})]. \quad (4.2)$$

Then for $0 \leq h \leq \frac{1}{\sqrt{\nu}}$, we have

$$\xi_t^\psi + th \cdot \langle \mathbf{x}_t, \mathbf{c} \rangle \geq \xi_{t(1+h)}^\psi \geq \xi_t^\psi + ht \cdot \langle \mathbf{c}, \mathbf{x}_t \rangle - h^2\nu.$$

Proof. Note that here the first inequality is fairly generic and holds for any function ψ . By definition of $\xi_{t(1+h)}^\psi$ and ξ_t^ψ in [Equation \(4.2\)](#) and using the fact that the value on the right hand side of [Equation \(4.3\)](#) is smaller than the expression evaluated at a fixed $\mathbf{x} = \mathbf{x}_t$, we have

$$\begin{aligned} \xi_{t(1+h)}^\psi &= \min_{\mathbf{x}} [t(1+h) \cdot \langle \mathbf{x}, \mathbf{c} \rangle + \psi(\mathbf{x})] \\ &\leq t(1+h) \cdot \langle \mathbf{x}_t, \mathbf{c} \rangle + \psi(\mathbf{x}_t) \\ &= \xi_t^\psi + th \cdot \langle \mathbf{x}_t, \mathbf{c} \rangle. \end{aligned} \quad (4.3)$$

We now prove the second inequality of the lemma. This one specifically uses the self-concordance of ψ . Observe first, by definition,

$$\xi_t^\psi = -\psi^*(-t\mathbf{c}). \quad (4.4)$$

Since ψ is a self-concordant barrier (and hence, a self-concordant function), [Theorem 2.10](#) implies that ψ^* is a self-concordant function as well. Then, by applying [Theorem 2.11](#) to ψ^* under the assumption $\| -th\mathbf{c} \|_{-t\mathbf{c}} \leq 1$ yields

$$\psi^*(-t\mathbf{c} - th\mathbf{c}) \leq \psi^*(-t\mathbf{c}) + \langle \nabla \psi^*(-t\mathbf{c}), -th\mathbf{c} \rangle + \left[\frac{1}{2} \| -th\mathbf{c} \|_{-t\mathbf{c}}^2 + \frac{\| -th\mathbf{c} \|_{-t\mathbf{c}}^3}{3(1 - \| -th\mathbf{c} \|_{-t\mathbf{c}})} \right]. \quad (4.5)$$

By applying the first-order optimality condition to the definition of \mathbf{x}_t in [Equation \(4.2\)](#), we see that

$$\nabla \psi(\mathbf{x}_t) = -t\mathbf{c}. \quad (4.6)$$

Next, evaluating $a \stackrel{\text{def}}{=} \|-th\mathbf{c}\|_{-t\mathbf{c}}$ to check the assumption $\|-th\mathbf{c}\|_{-t\mathbf{c}} \leq 1$, we get

$$\begin{aligned} a^2 &= h^2 \langle -t\mathbf{c}, \nabla^2 \psi^*(-t\mathbf{c}) \cdot (-t\mathbf{c}) \rangle = h^2 \langle \nabla \psi(\mathbf{x}_t), \nabla^2 \psi^*(\nabla \psi(\mathbf{x}_t)) \cdot \nabla \psi(\mathbf{x}_t) \rangle \\ &= h^2 \langle \nabla \psi(\mathbf{x}_t), (\nabla^2 \psi(\mathbf{x}_t))^{-1} \cdot \nabla \psi(\mathbf{x}_t) \rangle \\ &\leq h^2 \nu \end{aligned}$$

where we used Equation (4.6) and Lemma 2.4, in the first two equations and Definition 2.7 and the complexity value of ψ in the last step. Our range of h proves that $a \leq 1$, which is what we need for Inequality (4.5) to hold. We continue our computation to get

$$\left[\frac{1}{2} \|-th\mathbf{c}\|_{-t\mathbf{c}}^2 + \frac{\|-th\mathbf{c}\|_{-t\mathbf{c}}^3}{3(1 - \|-th\mathbf{c}\|_{-t\mathbf{c}})} \right] \leq \frac{1}{2} h^2 \nu + \frac{1}{3} h^3 \nu^{3/2} \leq \frac{1}{2} h^2 \nu + \frac{1}{3} h^2 \nu \leq h^2 \nu. \quad (4.7)$$

Applying Lemma 2.3 to Equation (4.6) gives

$$\nabla \psi^*(-t\mathbf{c}) = \mathbf{x}_t. \quad (4.8)$$

Plugging Equation (4.8) and Inequality (4.7) into the first and second-order terms, respectively, of Inequality (4.5) gives

$$\psi^*(-t\mathbf{c} - th\mathbf{c}) \leq \psi^*(-t\mathbf{c}) + \langle \mathbf{x}_t, -th\mathbf{c} \rangle + h^2 \nu.$$

Plugging in Equation (4.4) gives the desired inequality and completes the proof. \square

To finally compute the potential change due to t , we need the following result about the self-concordance parameter of the entropic barrier. While [BE15] prove that this barrier on a set in \mathbb{R}^d is $(1 + \epsilon_d)d$ -self-concordant, the recent work of [Che21] remarkably improves this complexity to exactly d .

Theorem 4.3 ([Che21]). *The entropic barrier on any convex body $\mathcal{K} \subseteq \mathbb{R}^d$ is a d -self-concordant barrier.*

We may now compute the potential change due to change in t in Line 9.

Lemma 4.4. *When t is updated to $t \cdot [1 + \frac{\eta}{4m}]$ in Line 9 of Algorithm 1, the potential Φ Equation (4.1) increases to $\Phi^{(\text{new})}$ as follows:*

$$\Phi^{(\text{new})} \leq \Phi + \eta + \eta^2.$$

Proof. Recall that the barrier function we use for the set $\widehat{\mathcal{K}}_{\text{out}}$ is the entropic barrier ψ_{out} . By Equation (4.2) and the definition of conjugate, we have

$$-\xi_t^{\psi_{\text{out}}} = \max_{\mathbf{v}} [\langle -t\mathbf{c}, \mathbf{v} \rangle - \psi_{\text{out}}(\mathbf{v})] = \psi_{\text{out}}^*(-t\mathbf{c}).$$

Applying Definition 2.12, taking the conjugate on both sides of the preceding equation, and using Lemma 2.2 then gives

$$-\xi_t^{\psi_{\text{out}}} = \log \left[\int_{\widehat{\mathcal{K}}_{\text{out}}} \exp(-t \cdot \langle \mathbf{c}, \mathbf{u} \rangle) d\mathbf{u} \right]. \quad (4.9)$$

From Equation (4.1), the change in potential by changing t to $t \cdot (1 + h)$ for some $h > 0$ may be expressed as

$$\Phi^{(\text{new})} - \Phi = \log \left[\int_{\widehat{\mathcal{K}}_{\text{out}}} \exp\langle -t(1+h)\mathbf{c}, \mathbf{v} \rangle d\mathbf{v} \right] - \log \left[\int_{\widehat{\mathcal{K}}_{\text{out}}} \exp\langle -t\mathbf{c}, \mathbf{v} \rangle d\mathbf{v} \right] + \langle th \cdot \mathbf{c}, \mathbf{x} \rangle.$$

By applying $h = \frac{\eta}{4m}$ and $\nu = m$ (via a direct application of Theorem 4.3), we have $h = \frac{\eta}{4m} \leq \frac{1}{\sqrt{m}} = \frac{1}{\sqrt{\nu}}$, and so we may now apply Equation (4.9) and Lemma 4.2 in the preceding equation to obtain the following bound.

$$\Phi^{(\text{new})} - \Phi \leq th \langle \mathbf{c}, \mathbf{x} \rangle - th \langle \mathbf{c}, \mathbf{x}_t \rangle + h^2 \nu.$$

From Equation (3.2) and Equation (4.2), we see that \mathbf{x}_t for the entropic barrier satisfies the equation $\mathbf{x}_t = \mathbf{x}_{\text{out}}^*$, and applying the guarantee $\langle \mathbf{c}, \mathbf{x} \rangle \leq \langle \mathbf{c}, \mathbf{x}_{\text{out}}^* \rangle + \frac{4m}{t}$ to this inequality, we obtain

$$\Phi^{(\text{new})} - \Phi \leq th \cdot \frac{4m}{t} + h^2 \nu = \eta + \left(\frac{\eta}{4m} \right)^2 \nu \leq \eta + \eta^2. \quad \square$$

4.2 Potential change for the universal terms

In this section, we study the change in volume on growing the inner convex set $\mathcal{K}_{\text{in},i}$ in [Line 15](#). As mentioned in [Section 3](#), our barrier of choice on this set is the universal barrier introduced in [\[NN94a\]](#) (see also [\[Gül97\]](#)). This barrier was constructed to demonstrate that *any* convex body in \mathbb{R}^n admits an $O(n)$ -self-concordant barrier, and its complexity parameter was improved to exactly n in [\[LY21\]](#).

Conceptually, we choose the universal barrier for the inner set because the operation we perform on the inner set (i.e., generating its convex hull with an external point $\mathbf{x}_{\text{out}}^*$) is dual to the operation of intersecting the outer set with the separating halfspace containing $\mathbf{x}_{\text{out}}^*$ (see [Lemma 2.6](#)), which suggests the use of a barrier dual to the entropic barrier used on the outer set. As explained in [\[BE15\]](#), for the special case of convex cones, the universal barrier is precisely one such barrier.

We now state a technical property of the universal barrier, which we use in the potential argument for this section.

Lemma 4.5 ([\[LY21, Lemma 1\]](#), [\[NN94a; Gül97\]](#)). *Given a convex set $\mathcal{K} \in \mathbb{R}^d$ and $\mathbf{x} \in \mathcal{K}$, let $\psi_{\mathcal{K}}(\mathbf{x}) \stackrel{\text{def}}{=} \log \text{vol}(\mathcal{K} - \mathbf{x})^\circ$ be the universal barrier defined on \mathcal{K} with respect to \mathbf{x} . Let $\mu \in \mathbb{R}^d$ be the center of gravity and $\Sigma \in \mathbb{R}^{d \times d}$ be the covariance matrix of the body $(\mathcal{K} - \mathbf{x})^\circ$, where $(\mathcal{K} - \mathbf{x})^\circ = \{\mathbf{y} \in \mathbb{R}^n : \mathbf{y}^\top(\mathbf{z} - \mathbf{x}) \leq 1, \forall \mathbf{z} \in \mathcal{K}\}$ is the polar set of \mathcal{K} with respect to \mathbf{x} . Then, we have that*

$$\nabla \psi_{\mathcal{K}}(\mathbf{x}) = (d+1)\mu, \quad \nabla^2 \psi_{\mathcal{K}}(\mathbf{x}) = (d+1)(d+2)\Sigma + (d+1)\mu\mu^\top.$$

Lemma 4.6. *Given a convex set $\mathcal{K} \subseteq \mathbb{R}^d$ and a point $\mathbf{x} \in \mathcal{K}$. Let $\psi_{\mathcal{K}} \stackrel{\text{def}}{=} \log \text{vol}(\mathcal{K} - \mathbf{x})^\circ$ be the universal barrier defined on \mathcal{K} with respect to \mathbf{x} . Let $\eta \leq 1/4$ and $\mathbf{y} \in \mathcal{K}$ be a point satisfying the following condition*

$$\langle \nabla \psi_{\mathcal{K}}(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \eta \|\mathbf{y} - \mathbf{x}\|_{\mathbf{x}} \geq 4d, \quad (4.10)$$

and construct the new set $\text{conv}\{\mathcal{K}, \mathbf{y}\}$. Then, the value of the universal barrier defined on this new set with respect to \mathbf{x} satisfies the following inequality.

$$\psi_{\mathcal{K}, \text{new}}(\mathbf{x}) \stackrel{\text{def}}{=} \psi_{\text{conv}\{\mathcal{K}, \mathbf{y}\}}(\mathbf{x}) = \log \text{vol}(\text{conv}(\mathcal{K}, \mathbf{y}) - \mathbf{x})^\circ \leq \psi_{\mathcal{K}}(\mathbf{x}) + \log(1 - 1/e + \eta).$$

Proof. By [Lemma 2.6](#), we have that

$$(\text{conv}(\mathcal{K}, \mathbf{y}) - \mathbf{x})^\circ \subseteq (\mathcal{K} - \mathbf{x})^\circ \cap \mathcal{H},$$

where $\mathcal{H} = \{\mathbf{z} \in \mathbb{R}^n : \langle \mathbf{z}, \mathbf{y} - \mathbf{x} \rangle \leq 1\}$. Our strategy to computing the deviation of $\psi_{\mathcal{K}, \text{new}}(\mathbf{x}) \stackrel{\text{def}}{=} \psi_{\text{conv}(\mathcal{K}, \mathbf{y})}(\mathbf{x}) = \log \text{vol}(\text{conv}(\mathcal{K}, \mathbf{y}) - \mathbf{x})^\circ$ from $\psi_{\mathcal{K}}(\mathbf{x})$ is to compute the change in $\text{vol}(\text{conv}(\mathcal{K}, \mathbf{y}) - \mathbf{x})^\circ \leq \text{vol}[(\mathcal{K} - \mathbf{x})^\circ \cap \mathcal{H}]$ from $\text{vol}(\mathcal{K} - \mathbf{x})^\circ$, for which it is immediate that one may apply an appropriate form of Grünbaum's Theorem.

Let μ be the center of gravity of the body $(\mathcal{K} - \mathbf{x})^\circ$. If $\mu \notin \mathcal{H}$, then [Corollary 2.16](#) (with $\eta = 0$) gives

$$\text{vol}[(\mathcal{K} - \mathbf{x})^\circ \cap \mathcal{H}] \leq \text{vol}(\mathcal{K} - \mathbf{x})^\circ \cdot (1 - 1/e),$$

and taking the logarithm on both sides gives the claimed bound. We now consider the case in which $\mu \in \mathcal{H}$, and the variance matrix of the body $(\mathcal{K} - \mathbf{x})^\circ$ is Σ . Define $\mathbf{v} = \mathbf{y} - \mathbf{x}$, and consider the point

$$\mathbf{z} = \mu + \frac{1 - \langle \mathbf{v}, \mu \rangle}{\|\mathbf{v}\|_{\Sigma}^2} \cdot \Sigma \mathbf{v}.$$

This point satisfies $\langle \mathbf{v}, \mathbf{z} \rangle = 1$, which implies $\mathbf{z} \in \mathcal{H}$. Specifically, \mathbf{z} lies on the separating hyperplane. We show that \mathbf{z} is sufficiently close to μ , so that even though $\mu \in \mathcal{H}$, the subset of $(\mathcal{K} - \mathbf{x})^\circ$ cut out by the halfspace \mathcal{H} is not too large. By applying [Lemma 4.5](#) to compute $\|\mathbf{v}\|_{\mathbf{x}}^2 = (d+1)(d+2)\|\mathbf{v}\|_{\Sigma}^2 + (d+1)\langle \mathbf{v}, \mu \rangle^2$, we may compute the following quantity.

$$\begin{aligned} \|\mathbf{z} - \mu\|_{\Sigma^{-1}} &= \frac{1 - \langle \mathbf{v}, \mu \rangle}{\sqrt{\frac{1}{(d+1)(d+2)}\|\mathbf{v}\|_{\mathbf{x}}^2 - \frac{1}{d+2} \cdot \langle \mathbf{v}, \mu \rangle^2}} \\ &= \sqrt{(d+1)(d+2)} \cdot \frac{1 - \langle \mathbf{v}, \mu \rangle}{\sqrt{\frac{1}{2}\|\mathbf{v}\|_{\mathbf{x}}^2 + \frac{1}{2}\|\mathbf{v}\|_{\mathbf{x}}^2 - (d+1)\langle \mathbf{v}, \mu \rangle^2}}. \end{aligned} \quad (4.11)$$

Applying the expression for gradient from [Lemma 4.5](#) in [Equation \(4.10\)](#), we have

$$\eta \|\mathbf{v}\|_{\mathbf{x}} \geq 4d - (d+1)\langle \mathbf{v}, \mu \rangle \geq 2d\langle \mathbf{v}, \mu \rangle,$$

where we used the fact that $\mu \in \mathcal{H}$ implies $\langle \mathbf{v}, \mu \rangle \leq 1$. Since $\eta \leq 1/4$, we have $\frac{1}{2}\|\mathbf{v}\|_{\mathbf{x}}^2 \geq (d+1)\langle \mathbf{v}, \mu \rangle^2$. Plugging this in [Equation \(4.11\)](#) gives

$$\begin{aligned} \|\mathbf{z} - \mu\|_{\mathbf{A}^{-1}} &\leq \sqrt{(d+1)(d+2)} \cdot \frac{1 - \langle \mathbf{v}, \mu \rangle}{\sqrt{\frac{1}{2}\|\mathbf{v}\|_{\mathbf{x}}^2}} \leq 4d \frac{1 - \langle \mathbf{v}, \mu \rangle}{\|\mathbf{v}\|_{\mathbf{x}}} \\ &\leq 4d \cdot \frac{1 - \langle \mathbf{v}, \mu \rangle}{4d(1 - \langle \mathbf{v}, \mu \rangle)/\eta} \leq \eta, \end{aligned}$$

which implies [Corollary 2.16](#) applies, giving us the desired volume reduction. \square

4.3 Potential change for the update of \mathbf{x}

In this section, we quantify the amount of progress made in [Line 22](#) of [Algorithm 1](#) by computing the change in the potential Φ as defined in [Equation \(4.1\)](#).

Lemma 4.7. *Consider the potential Φ [Equation \(4.1\)](#) and the update step $\delta_{\mathbf{x}} = \frac{\eta}{2} \cdot \frac{\mathbf{x}_{out}^* - \mathbf{x}}{\|\mathbf{x}_{out}^* - \mathbf{x}\|_{\mathbf{x},1}}$ as in [Line 22](#). Assume the guarantees in [Inequality \(3.3\)](#) and [Inequality \(3.4\)](#). Then the potential Φ incurs the following minimum decrease.*

$$\Phi^{(new)} \leq \Phi - \frac{\eta^2}{4}.$$

Proof. Taking the gradient of Φ with respect to \mathbf{x} and rearranging the terms gives

$$t\mathbf{c} = \nabla_{\mathbf{x}}\Phi - \sum_{i=1}^n \nabla\psi_{in,i}(\mathbf{x}_i). \quad (4.12)$$

By applying the expression for $t\mathbf{c}$ from the preceding equation, we get

$$\begin{aligned} \Phi^{(new)} - \Phi &= t\langle \mathbf{c}, \mathbf{x} + \delta_{\mathbf{x}} \rangle + \sum_{i=1}^n \psi_{in,i}(\mathbf{x}_i + \delta_{\mathbf{x},i}) - t\langle \mathbf{c}, \mathbf{x} \rangle - \sum_{i=1}^n \psi_{in,i}(\mathbf{x}_i) \\ &= \langle \nabla_{\mathbf{x}}\Phi, \delta_{\mathbf{x}} \rangle + \sum_{i=1}^n \underbrace{[\psi_{in,i}(\mathbf{x}_i + \delta_{\mathbf{x},i}) - \psi_{in,i}(\mathbf{x}_i) - \langle \nabla\psi_{in,i}(\mathbf{x}_i), \delta_{\mathbf{x},i} \rangle]}_{q_{\psi_{in,i}}(\mathbf{x}_i)}. \end{aligned} \quad (4.13)$$

The term $q_{\psi_{in,i}}(\mathbf{x}_i)$ measures the error due to first-order approximation of $\psi_{in,i}$ around \mathbf{x}_i . Since $\psi_{in,i}(\mathbf{x}_i)$ is self-concordant functions and $\|\delta_{\mathbf{x},i}\|_{\mathbf{x}_i} \leq \|\delta_{\mathbf{x}}\|_{\mathbf{x},1} \leq \eta \leq 1/4$, [Theorem 2.11](#) shows that

$$\psi_{in,i}(\mathbf{x}_i + \delta_{\mathbf{x},i}) - \psi_{in,i}(\mathbf{x}_i) - \langle \nabla\psi_{in,i}(\mathbf{x}_i), \delta_{\mathbf{x},i} \rangle \leq \|\delta_{\mathbf{x},i}\|_{\mathbf{x}_i}^2. \quad (4.14)$$

Plugging in [Inequality \(4.14\)](#) into [Equation \(4.13\)](#), we get

$$\Phi^{(new)} - \Phi \leq \langle \nabla_{\mathbf{x}}\Phi, \delta_{\mathbf{x}} \rangle + \|\delta_{\mathbf{x}}\|_{\mathbf{x},1}^2. \quad (4.15)$$

We now bound the two terms on the right hand side one at a time. Using the definition of $\delta_{\mathbf{x}}$ (as given in

the statement of the lemma) and of $\nabla_{\mathbf{x}}\Phi$ from Equation (4.12) gives

$$\begin{aligned}
\langle \nabla_{\mathbf{x}}\Phi, \delta_{\mathbf{x}} \rangle &= \frac{\eta}{2} \frac{1}{\|\mathbf{x}_{\text{out}}^* - \mathbf{x}\|_{\mathbf{x},1}} \langle \nabla_{\mathbf{x}}\Phi, \mathbf{x}_{\text{out}}^* - \mathbf{x} \rangle \\
&= \frac{\eta}{2} \frac{1}{\|\mathbf{x}_{\text{out}}^* - \mathbf{x}\|_{\mathbf{x},1}} \left[\langle t\mathbf{c}, \mathbf{x}_{\text{out}}^* - \mathbf{x} \rangle + \sum_{i=1}^n \langle \nabla\psi_{\text{in},i}(\mathbf{x}_i), \mathbf{x}_{\text{out},i}^* - \mathbf{x}_i \rangle \right] \\
&\leq \frac{\eta}{2} \frac{1}{\|\mathbf{x}_{\text{out}}^* - \mathbf{x}\|_{\mathbf{x},1}} \left[\langle t\mathbf{c}, \mathbf{x}_{\text{out}}^* - \mathbf{x} \rangle + \sum_{i=1}^n (4r_i - \eta\|\mathbf{x}_{\text{out},i}^* - \mathbf{x}_i\|_{\mathbf{x}_i}) \right] \\
&= \frac{\eta}{2} \frac{1}{\|\mathbf{x}_{\text{out}}^* - \mathbf{x}\|_{\mathbf{x},1}} [\langle t\mathbf{c}, \mathbf{x}_{\text{out}}^* - \mathbf{x} \rangle + 4m - \eta\|\mathbf{x}_{\text{out}}^* - \mathbf{x}\|_{\mathbf{x},1}] \\
&\leq \frac{\eta}{2} \frac{1}{\|\mathbf{x}_{\text{out}}^* - \mathbf{x}\|_{\mathbf{x},1}} \cdot (-\eta\|\mathbf{x}_{\text{out}}^* - \mathbf{x}\|_{\mathbf{x},1}) \\
&= -\eta^2/2.
\end{aligned} \tag{4.16}$$

where the third step follows from Inequality (3.3), the fourth step follows from $\sum_{i=1}^n d_i = m$, and the fifth step follows from Inequality (3.4). To bound the second term, we note from Line 22 that

$$\|\delta_{\mathbf{x}}\|_{\mathbf{x},1}^2 = \left(\frac{\eta}{2} \cdot \frac{\|\mathbf{x}_{\text{out}}^* - \mathbf{x}\|_{\mathbf{x},1}}{\|\mathbf{x}_{\text{out}}^* - \mathbf{x}\|_{\mathbf{x},1}} \right)^2 = \eta^2/4. \tag{4.17}$$

Hence, we may plug in Inequality (4.16) and Equation (4.17) into Inequality (4.15) to get the desired result. \square

4.4 Total oracle complexity

Before we bound the total oracle complexity of the algorithm, we first bound the total potential change throughout the algorithm.

Lemma 4.8. *Consider the potential function $\Phi = t\langle \mathbf{c}, \mathbf{x} \rangle + \log \left[\int_{\widehat{\mathcal{K}}_{\text{out}}} \exp(-t\langle \mathbf{c}, \mathbf{u} \rangle) d\mathbf{u} \right] + \sum_{i \in [n]} \psi_{\text{in},i}(\mathbf{x}_i)$ as defined in Equation (4.1) associated with Algorithm 1. Let Φ_{init} be the potential at $t = t_{\text{init}}$ of this algorithm, and let Φ_{end} be the potential at $t = t_{\text{end}}$. Suppose at $t = t_{\text{init}}$ in Algorithm 1, we have $\mathcal{B}_m(\mathbf{x}, \bar{r}) \subseteq \mathcal{K}_{\text{in}}$ with $\bar{r} = r/\text{poly}(m)$ and $\mathcal{K}_{\text{out}} \subseteq \mathcal{B}_m(0, R)$ with $R = O(\sqrt{n}R)$. Then we have, under the assumptions of Theorem 4.10, that*

$$\Phi_{\text{init}} - \Phi_{\text{end}} \leq O \left(m \log \left(\frac{mR}{\epsilon r} \right) \right).$$

Proof. For this proof, we introduce the following notation: let $\text{vol}_{\mathbf{A}}(\cdot)$ denote the volume restricted to the subspace $\{\mathbf{x} : \mathbf{A}\mathbf{x} = \mathbf{b}\}$. We also invoke Fact 2.17. We now bound the change in the potential term by term, starting with the entropic terms

$$t\langle \mathbf{c}, \mathbf{x} \rangle + \log \left[\int_{\widehat{\mathcal{K}}_{\text{out}}} \exp(-t\langle \mathbf{c}, \mathbf{u} \rangle) d\mathbf{u} \right] \tag{4.18}$$

at $t = t_{\text{init}}$ and a lower bound on it at $t = t_{\text{end}}$. We start with bounding Equation (4.18) evaluated at $t = t_{\text{end}} = \frac{8m}{\epsilon\|\mathbf{c}\|_2 R}$.

Let $\bar{\mathbf{x}} = \text{argmin}_{\mathbf{x} \in \widehat{\mathcal{K}}_{\text{out}}} \langle \mathbf{c}, \mathbf{x} \rangle$ and $\alpha = \langle \mathbf{c}, \bar{\mathbf{x}} \rangle$. By optimality of $\bar{\mathbf{x}}$, we know that $\bar{\mathbf{x}} \in \partial \widehat{\mathcal{K}}_{\text{out}}$. Denote $\mathcal{B}_{\mathbf{A}}(\mathbf{z}, \bar{r})$ to be $\mathcal{B}(\mathbf{z}, \bar{r})$ restricted to the subspace $\{\mathbf{x} : \mathbf{A}\mathbf{x} = \mathbf{b}\}$. Note that $\widehat{\mathcal{K}}_{\text{out}} \supseteq \mathcal{B}_{\mathbf{A}}(\mathbf{z}, \bar{r})$. Consider the cone \mathcal{C} and halfspace \mathcal{H} defined by

$$\mathcal{C} = \bar{\mathbf{x}} + \{\lambda \mathbf{y} : \lambda > 0, \mathbf{y} \in \mathcal{B}_{\mathbf{A}}(\mathbf{z} - \bar{\mathbf{x}}, \bar{r})\} \text{ and } \mathcal{H} \stackrel{\text{def}}{=} \left\{ \mathbf{x} : \langle \mathbf{c}, \mathbf{x} \rangle \leq \alpha + \frac{1}{t_{\text{end}}} \right\}.$$

Then, by a similarity argument, we note that $\mathcal{C} \cap \mathcal{H}$ contains a cone with height $\frac{1}{t_{\text{end}}\|\mathbf{c}\|_2}$ and base radius $\frac{\bar{r}}{Rt_{\text{end}}\|\mathbf{c}\|_2}$, which means

$$\text{vol}_{\mathbf{A}}(\mathcal{C} \cap \mathcal{H}) \geq \frac{1}{m - \text{rank}(\mathbf{A})} \cdot \frac{1}{t_{\text{end}}\|\mathbf{c}\|_2} \cdot \left(\frac{\bar{r}}{Rt_{\text{end}}\|\mathbf{c}\|_2} \right)^{m - \text{rank}(\mathbf{A}) - 1} \cdot \text{vol}(\mathcal{B}_{m - \text{rank}(\mathbf{A}) - 1}(0, 1)).$$

Then, we have

$$\begin{aligned} \log \left[\int_{\widehat{\mathcal{K}}_{\text{out}}} \exp(-t_{\text{end}}\langle \mathbf{c}, \mathbf{u} \rangle) d\mathbf{u} \right] + t_{\text{end}}\langle \mathbf{c}, \mathbf{x} \rangle &\geq \log \left[\int_{\widehat{\mathcal{K}}_{\text{out}}} \exp(-t_{\text{end}}\langle \mathbf{c}, \mathbf{u} \rangle) d\mathbf{u} \right] + t_{\text{end}} \min_{\mathbf{x} \in \widehat{\mathcal{K}}_{\text{out}}} \langle \mathbf{c}, \mathbf{x} \rangle \\ &\geq \log \left[\int_{\mathcal{C} \cap \mathcal{H}} \exp(-t_{\text{end}}\langle \mathbf{c}, \mathbf{u} \rangle) d\mathbf{u} \right] + t_{\text{end}}\alpha \\ &\geq \log \left[\int_{\mathcal{C} \cap \mathcal{H}} \exp(-t_{\text{end}}\alpha - 1) d\mathbf{u} \right] + t_{\text{end}}\alpha \\ &= \log \left[\frac{1}{e} \cdot \text{vol}_{\mathbf{A}}(\mathcal{C} \cap \mathcal{H}) \exp(-t_{\text{end}}\alpha) \right] + t_{\text{end}}\alpha \\ &= \log \left[\text{vol}_{\mathbf{A}}(\mathcal{C} \cap \mathcal{H}) \cdot \frac{1}{e} \right] \\ &\geq -(m - \text{rank}(\mathbf{A}) - 1) \cdot \log(\bar{R}t_{\text{end}}\|\mathbf{c}\|_2/\bar{r}) \\ &\quad + \log(\text{vol}(\mathcal{B}_{m - \text{rank}(\mathbf{A}) - 1}(0, 1))) \\ &\quad - \log(m - \text{rank}(\mathbf{A})) - \log(t_{\text{end}}\|\mathbf{c}\|_2) - 1. \end{aligned} \quad (4.19)$$

Next, to bound Equation (4.18) at $t = t_{\text{init}}$, we may express these terms as follows.

$$\begin{aligned} \log \left[\int_{\widehat{\mathcal{K}}_{\text{out}}} \exp(-t_{\text{init}} \cdot \langle \mathbf{c}, \mathbf{u} \rangle) d\mathbf{u} \right] + t_{\text{init}} \cdot \langle \mathbf{c}, \mathbf{x} \rangle &\leq \log \left[\text{vol}_{\mathbf{A}}(\widehat{\mathcal{K}}_{\text{out}}) \right] + t_{\text{init}} \cdot \max_{\mathbf{u} \in \widehat{\mathcal{K}}_{\text{out}}} \langle \mathbf{c}, \mathbf{x} - \mathbf{u} \rangle \\ &\leq \log(\text{vol}(\mathcal{B}_{m - \text{rank}(\mathbf{A})}(0, \bar{R}))) + t_{\text{init}} \cdot 2\bar{R}\|\mathbf{c}\|_2 \\ &\leq \log(\text{vol}(\mathcal{B}_{m - \text{rank}(\mathbf{A})}(0, 1))) \\ &\quad + (m - \text{rank}(\mathbf{A})) \log \bar{R} + O(m \log m), \end{aligned} \quad (4.20)$$

where the second step is by $\widehat{\mathcal{K}}_{\text{out}} \subseteq \mathcal{K}_{\text{out}} \subseteq \mathcal{B}_{\sum_{i \in [n]} d_i}(0, \bar{R})$ (here, the second inclusion is by assumption), and the third step is by $\text{vol}(\mathcal{B}_q(0, \bar{R})) = \frac{\pi^{q/2}}{\Gamma(1+q/2)} \bar{R}^q$ and our choice of $t_{\text{init}} \stackrel{\text{def}}{=} \frac{m \log m}{\sqrt{n}\|\mathbf{c}\|_2 \bar{R}}$.

We now compute the change in the entropic barrier $\sum_{i \in [n]} \psi_{\text{in},i}(\mathbf{x}_i)$, where

$$\psi_{\text{in},i}(\mathbf{x}_i) = \log \text{vol}(\mathcal{K}_{\text{in},i}^{\circ}(\mathbf{x}_i)).$$

Define $\mathcal{B}_d(0, r)$ to be the d -dimensional Euclidean ball centred at the origin and with radius r . We note by the radius assumption of Theorem 4.10 that $\mathcal{K}_{\text{in},i} \subseteq \mathcal{K}_i \subseteq \mathcal{B}_{d_i}(0, \bar{R})$ throughout the algorithm. By the assumption made in this lemma's statement, we have that at the start of Algorithm 1, $\mathcal{K}_{\text{in},i} \supseteq \mathcal{B}_{d_i}(\mathbf{x}, \bar{r})$. These give us the following bounds.

$$\psi_{\text{in},i}^{\text{end}}(\mathbf{x}_i) \geq \log(\text{vol}(\mathcal{B}_{d_i}^{\circ}(0, \bar{R}))) \text{ and } \psi_{\text{in},i}^{\text{init}}(\mathbf{x}_i) \leq \log(\text{vol}(\mathcal{B}_{d_i}^{\circ}(\mathbf{x}_i, \bar{r}))).$$

Applying the fact that $\text{vol}(\mathcal{B}_d(0, r)) \propto r^d$ and summing over all $i \in [n]$ gives

$$\begin{aligned} \sum_{i \in [n]} [\psi_{\text{in}, i}^{\text{init}}(\mathbf{x}_i) - \psi_{\text{in}, i}^{\text{end}}(\mathbf{x}_i)] &\leq \sum_{i \in [n]} \log \left(\frac{\text{vol}(\mathcal{B}_{d_i}(\mathbf{x}_i, 1/\bar{r}))}{\text{vol}(\mathcal{B}_{d_i}(0, 1/\bar{R}))} \right) \\ &= \sum_{i \in [n]} d_i \log(\bar{R}/\bar{r}) = m \log(\bar{R}/\bar{r}). \end{aligned} \quad (4.21)$$

Combining [Inequality \(4.20\)](#), [Inequality \(4.19\)](#), and [Inequality \(4.21\)](#), we have

$$\begin{aligned} \Phi_{\text{init}} - \Phi_{\text{end}} &\leq m \log(mR/r) \\ &\quad + [\log(\text{vol}(\mathcal{B}_{m-\text{rank}(\mathbf{A})}(0, 1))) + (m - \text{rank}(\mathbf{A})) \log \bar{R} + O(m \log m)] \\ &\quad + (m - \text{rank}(\mathbf{A}) - 1) \cdot \log(\bar{R} t_{\text{end}} \|\mathbf{c}\|_2 / \bar{r}) - \log(\text{vol}(\mathcal{B}_{m-\text{rank}(\mathbf{A})-1}(0, 1))) \\ &\quad + \log(m - \text{rank}(\mathbf{A})) + \log(t_{\text{end}} \|\mathbf{c}\|_2) + 1 \\ &\leq m \log(mR/\epsilon r) \\ &\quad + O(m \log m) \\ &\quad + O((m - \text{rank}(\mathbf{A})) \log(mR/\epsilon r)) \leq O(m \log(mR/\epsilon r)). \end{aligned}$$

□

Lemma 4.9. *[Total oracle complexity] Suppose the inputs \mathcal{K}_{in} and \mathcal{K}_{out} to [Algorithm 1](#) satisfy $\mathcal{K}_{\text{out}} \subseteq \mathcal{B}_m(0, \bar{R})$ with $\bar{R} = O(\sqrt{n}R)$ and $\mathcal{K}_{\text{in}} \supseteq \mathcal{B}(\mathbf{z}, \bar{r})$ with $\bar{r} = r/\text{poly}(m)$. Then, when [Algorithm 1](#) terminates at $t \geq t_{\text{end}}$, it outputs a solution \mathbf{x} that satisfies*

$$\mathbf{c}^\top \mathbf{x} \leq \min_{\mathbf{x} \in \mathcal{K}, \mathbf{A}\mathbf{x}=\mathbf{b}} \mathbf{c}^\top \mathbf{x} + \epsilon \cdot \|\mathbf{c}\|_2 R$$

using at most $\mathcal{N}_{\text{sep}} = O\left(m \log\left(\frac{mR}{\epsilon r}\right)\right)$ separation oracle calls.

Proof. Let \mathcal{N}_t be the number of times t is updated; \mathcal{N}_{in} the number of times \mathcal{K}_{in} is updated; \mathcal{N}_{out} the number of times \mathcal{K}_{out} is updated; $\mathcal{N}_{\mathbf{x}}$ the number of times \mathbf{x} is updated, and $\mathcal{N}_{\text{total}}$ the total number of iterations of the `while` loop before termination of [Algorithm 1](#). Then, combining [Lemma 4.1](#), [Lemma 4.4](#), [Lemma 4.6](#), and [Lemma 4.7](#) gives

$$\Phi_{\text{end}} \leq \Phi_{\text{init}} + \mathcal{N}_{\text{out}} \cdot \log(1 - 1/e) + \mathcal{N}_t \cdot (\eta + \eta^2) + \mathcal{N}_{\text{in}} \cdot \log(1 - 1/e + \eta) + \mathcal{N}_{\mathbf{x}} \cdot \left(-\frac{\eta^2}{4}\right). \quad (4.22)$$

The initialization step of [Algorithm 1](#) chooses $\eta = 1/100$, $t_{\text{end}} = \frac{8m}{\epsilon \|\mathbf{c}\|_2 R}$, and $t_{\text{init}} = \frac{m \log(m)}{\sqrt{n} \|\mathbf{c}\|_2 R}$, and we always update t by a multiplicative factor of $1 + \frac{\eta}{4m}$ (see [Line 9](#)); therefore, we have

$$\mathcal{N}_t = O(m \log(mR/(\epsilon r))).$$

From [Algorithm 1](#), the only times the separation oracle is invoked is when updating \mathcal{K}_{in} or \mathcal{K}_{out} in [Line 15](#) and [Line 17](#), respectively. Therefore, the total separation oracle complexity is $\mathcal{N}_{\text{sep}} = \mathcal{N}_{\text{in}} + \mathcal{N}_{\text{out}}$. Therefore, we have

$$\mathcal{N}_{\text{sep}} = \mathcal{N}_{\text{in}} + \mathcal{N}_{\text{out}} \leq O(1) \cdot [\Phi_{\text{init}} - \Phi_{\text{end}} + \mathcal{N}_t] = O(m \log(mR/(\epsilon r)))$$

This gives the claimed separation oracle complexity.

We now prove the guarantee on approximation. Let $\mathbf{x}_{\text{output}}$ be the output of [Algorithm 1](#) and \mathbf{x} be the point which entered [Line 5](#) right before termination. Note that the termination of [Algorithm 1](#) implies, by [Line 5](#), that

$$\mathbf{c}^\top \mathbf{x}_{\text{output}} \leq \mathbf{c}^\top \mathbf{x} + \frac{\nu}{t_{\text{end}}} \leq \mathbf{c}^\top \mathbf{x}_{\text{out}}^* + \frac{4(n+m)}{t_{\text{end}}} \leq \min_{\mathbf{x} \in \mathcal{K}, \mathbf{A}\mathbf{x}=\mathbf{b}} \mathbf{c}^\top \mathbf{x} + \epsilon \cdot \|\mathbf{c}\|_2 \cdot R$$

where the first step is by the second inequality in [Lemma 5.8](#) (using the universal barrier) and the last step follows by our choice of t_{end} and the definition of $\mathbf{x}_{\text{out}}^*$ and $\mathcal{K}_{\text{out}} \supseteq \mathcal{K}$. □

Theorem 4.10 (Main theorem of [Problem 3.1](#)). *Given the convex program*

$$\begin{aligned} & \text{minimize} && \langle \mathbf{c}, \mathbf{x} \rangle, \\ & \text{subject to} && \mathbf{x}_i \in \mathcal{K}_i \subseteq \mathbb{R}^{d_i+1} \forall i \in [n], \\ & && \mathbf{A}\mathbf{x} = \mathbf{b}. \end{aligned}$$

Denote $\mathcal{K} = \mathcal{K}_1 \times \mathcal{K}_2 \times \dots \times \mathcal{K}_n$. Assuming we have

- outer radius R : For any $\mathbf{x}_i \in \mathcal{K}_i$, we have $\|\mathbf{x}_i\|_2 \leq R$, and
- inner radius r : There exists a $\mathbf{z} \in \mathbb{R}^d$ such that $\mathbf{A}\mathbf{z} = \mathbf{b}$ and $\mathcal{B}(\mathbf{z}, r) \subset \mathcal{K}$,

then, for any $0 < \epsilon < \frac{1}{2}$, we can find a point $\mathbf{x} \in \mathcal{K}$ satisfying $\mathbf{A}\mathbf{x} = \mathbf{b}$ and

$$\langle \mathbf{c}, \mathbf{x} \rangle \leq \min_{\substack{\mathbf{x}_i \in \mathcal{K}_i \subseteq \mathbb{R}^{d_i+1} \forall i \in [n], \\ \mathbf{A}\mathbf{x} = \mathbf{b}}} \langle \mathbf{c}, \mathbf{x} \rangle + \epsilon \cdot \|\mathbf{c}\|_2 \cdot R,$$

in $O(\text{poly}(m \log(mR/\epsilon r)))$ time and using

$$O(m \log(mR/(\epsilon r)))$$

gradient oracle calls, where $m = \sum_{i=1}^n d_i$.

Proof. We apply [Theorem 5.1](#) for each \mathcal{K}_i separately to find a solution \mathbf{z}_i . Then $\mathbf{z} = (\mathbf{z}_1, \dots, \mathbf{z}_n) \in \mathbb{R}^{m+n}$ satisfies $\mathcal{B}_{m+n}(\mathbf{z}, \bar{r}) \subset \mathcal{K}$ with $\bar{r} = \frac{r}{6m^{3.5}}$. Then, we modified convex problem as in [Definition 5.5](#) with $s = 2^{16} \frac{m^{2.5} R}{r\epsilon}$ and obtaining the following:

$$\begin{aligned} & \text{minimize} && \langle \bar{\mathbf{c}}, \bar{\mathbf{x}} \rangle \\ & \text{subject to} && \mathbf{A}\bar{\mathbf{x}} = \bar{\mathbf{b}}, \\ & && \bar{\mathbf{x}} \in \bar{\mathcal{K}} \stackrel{\text{def}}{=} \mathcal{K} \times \mathbb{R}_{\geq 0}^{m+n} \times \mathbb{R}_{\geq 0}^{m+n} \end{aligned} \tag{4.23}$$

with

$$\bar{\mathbf{A}} = [\mathbf{A} \mid \mathbf{A} \mid -\mathbf{A}], \bar{\mathbf{b}} = \mathbf{b}, \bar{\mathbf{c}} = \left(\mathbf{c}, \frac{\|\mathbf{c}\|_{2s}}{\sqrt{m+n}} \cdot \mathbf{1}, \frac{\|\mathbf{c}\|_{2s}}{\sqrt{m+n}} \cdot \mathbf{1} \right)^\top$$

We solve the linear system $\mathbf{A}\mathbf{y} = \mathbf{b} - \mathbf{A}\mathbf{z}$ for \mathbf{y} . Then, we construct the initial $\bar{\mathbf{x}}$ by set $\bar{\mathbf{x}}^{(1)} = \mathbf{z}$,

$$\bar{\mathbf{x}}_i^{(2)} = \begin{cases} \mathbf{y}_i & \text{if } \mathbf{y}_i \geq 0, \\ 0 & \text{otherwise.} \end{cases} \quad \text{and} \quad \bar{\mathbf{x}}_i^{(3)} = \begin{cases} -\mathbf{y}_i & \text{if } \mathbf{y}_i < 0, \\ 0 & \text{otherwise.} \end{cases}$$

Then, we run [Algorithm 1](#) on the [Problem 4.23](#), with initial $\bar{\mathbf{x}}$ set above, $\bar{m} = 3(m+n)$, $\bar{n} = n+2$, $\bar{\epsilon} = \frac{\epsilon}{6\sqrt{ns}}$, $\bar{\mathcal{K}}_{\text{in}} = \{\mathbf{x}^{(1)} \in \mathcal{B}(\mathbf{z}, \bar{r}), (\mathbf{x}^{(2)}, \mathbf{x}^{(3)}) \in \mathbb{R}_{\geq 0}^{2n}\}$ and $\hat{\mathcal{K}}_{\text{out}} = \mathcal{B}_{\bar{m}}(\mathbf{0}, \sqrt{\bar{n}}R)$.

By our choice of t_{end} , we have

$$\bar{t}_{\text{end}} = \frac{8\bar{m}}{\bar{\epsilon}\|\bar{\mathbf{c}}\|_2\bar{R}} \leq \frac{48m}{\epsilon\|\mathbf{c}\|_2R}.$$

First, we check the condition that $s \geq 48\bar{v}\bar{t}_{\text{end}}\sqrt{m+n}\frac{R^2}{r}\|\mathbf{c}\|_2$, we note that

$$48\bar{v}\bar{t}_{\text{end}}\sqrt{m+n}\frac{R^2}{r}\|\mathbf{c}\|_2 \leq 27648 \frac{m^{2.5}R}{er} \leq 2^{16} \frac{m^{2.5}R}{r\epsilon} = s.$$

Let $\bar{\mathbf{x}}_{\text{output}} = (\mathbf{x}_{\text{output}}^{(1)}, \mathbf{x}_{\text{output}}^{(2)}, \mathbf{x}_{\text{output}}^{(3)})$ be the output of [Algorithm 1](#). Then, let $\mathbf{x}_{\text{output}} = \mathbf{x}_{\text{output}}^{(1)} + \mathbf{x}_{\text{output}}^{(2)} - \mathbf{x}_{\text{output}}^{(3)}$ as defined in [Theorem 5.6](#). By [Lemma 4.9](#), we have

$$\min_{\mathbf{x} \in \mathcal{P}_{\text{in}}} \bar{\mathbf{c}}^\top \bar{\mathbf{x}} \leq \min_{\mathbf{x} \in \mathcal{P}} \mathbf{c}^\top \bar{\mathbf{x}} + \gamma$$

where $\gamma = \bar{\epsilon} \cdot \|\bar{\mathbf{c}}\|_2 \cdot \bar{R}$.

Applying (3) of [Theorem 5.6](#), we have

$$\mathbf{c}^\top \mathbf{x}_{\text{output}} \leq \frac{\bar{\nu} + 1}{\bar{t}_{\text{end}}} + \gamma + \min_{x \in \mathcal{K}, \mathbf{A}\mathbf{x} = \mathbf{b}} \mathbf{c}^\top \mathbf{x} \leq \min_{x \in \mathcal{K}, \mathbf{A}\mathbf{x} = \mathbf{b}} \mathbf{c}^\top \mathbf{x} + \epsilon \cdot \|\mathbf{c}\|_2 \cdot R.$$

The last inequality follows by our choice of $\bar{\epsilon}$ and \bar{t}_{end} , we have $\gamma \leq \frac{\epsilon}{2} \|\mathbf{c}\|_2 R$ and $\frac{\bar{\nu} + 1}{\bar{t}_{\text{end}}} \leq \frac{\epsilon}{2} \|\mathbf{c}\|_2 R$. Plug this $\bar{\epsilon}$ in [Lemma 4.9](#), it gives the claimed oracle complexity. \square

Theorem 4.11 (Main Result). *Given [Problem 1.1](#) and $\theta^{(0)}$ such that $\|\theta^* - \theta^{(0)}\|_2 \leq R$. Assuming all the f_i 's are L -Lipschitz, then there is an algorithm that in time $\text{poly}(m \log(1/\epsilon))$, using $O(m \log(m/\epsilon))$ gradient oracle calls, outputs a vector $\theta \in \mathbb{R}^d$ such that*

$$\sum_{i=1}^n f_i(\theta) \leq \sum_{i=1}^n f_i(\theta^*) + \epsilon \cdot LR.$$

Proof. First, we reformulate (1.1) using a change of variables and the epigraph trick. Suppose each f_i depends on d_i coordinates of θ given by $\{i_1, \dots, i_{d_i}\} \subseteq [d]$. Then, symbolically define $\mathbf{x}_i = [x_{i_1}^{(i)}; x_{i_2}^{(i)}; \dots; x_{i_{d_i}}^{(i)}] \in \mathbb{R}^{d_i}$ for each $i \in [n]$. Since each f_i is convex and supported on d_i variables, its epigraph is convex and $d_i + 1$ dimensional. So we may define the convex set

$$\mathcal{K}_i^{\text{unbounded}} = \{(\mathbf{x}_i, z_i) \in \mathbb{R}^{d_i+1} : f_i(\mathbf{x}_i) \leq Lz_i\}.$$

Finally, we add linear constraints of the form $x_k^{(i)} = x_k^{(j)}$ for all i, j, k where f_i and f_j both depend on θ_k . We denote these by the matrix constraint $\mathbf{A}\mathbf{x} = \mathbf{b}$. Then, [Problem 1.1](#) is equivalent to

$$\begin{aligned} & \text{minimize} && \sum_{i=1}^n Lz_i \\ & \text{subject to} && \mathbf{A}\mathbf{x} = \mathbf{b} \\ & && (\mathbf{x}_i, z_i) \in \mathcal{K}_i^{\text{unbounded}} \text{ for each } i \in [n]. \end{aligned} \tag{4.24}$$

Since we are given $\theta^{(0)}$ satisfying $\|\theta^{(0)} - \theta^*\|_2 \leq R$, we define $\mathbf{x}_i^{(0)} = [\theta_{i_1}^{(0)}; \dots; \theta_{i_{d_i}}^{(0)}]$ and $z_i^{(0)} = f_i(\theta^{(0)})/L$. Then, we can restrict the search space $\mathcal{K}_i^{\text{unbounded}}$ to

$$\mathcal{K}_i = \mathcal{K}_i^{\text{unbounded}} \cap \{(\mathbf{x}_i, z_i) \in \mathbb{R}^{d_i+1} : \|\mathbf{x}_i - \mathbf{x}_i^{(0)}\|_2 \leq R \text{ and } z_i^{(0)} - 2R \leq z_i \leq z_i^{(0)} + 2R\}.$$

It's easy to check that \mathcal{K}_i is contained in a ball of radius $5R$ centered at $(\mathbf{x}_i^{(0)}, z_i^{(0)})$, and contains a ball of radius R centered at $(\mathbf{x}_i^{(0)}, z_i^{(0)})$. The subgradient oracle for f_i translates to a separation oracle for \mathcal{K}_i . Then, we apply [Theorem 4.10](#) to (4.24) with $\mathcal{K}_i^{\text{unbounded}}$ replaced by \mathcal{K}_i to get the error guarantee and oracle complexity directly. \square

Finally, we have the matching lower bound.

Theorem 1.3. *There exist functions $f_1, \dots, f_n : \mathbb{R}^d \mapsto \mathbb{R}$ for which a total of $\Omega(m \log(1/\epsilon))$ gradient queries are required to solve [Problem 1.1](#).*

Proof. [[Nes04](#)] shows that for any d_i , there exists $f_i : \mathbb{R}^{d_i} \mapsto \mathbb{R}$ for which $\Omega(d_i \log(1/\epsilon))$ total gradient queries are required. We define f_1, \dots, f_n to be such functions on disjoint coordinates of θ . It follows that $\Omega(\sum_{i=1}^n d_i \log(1/\epsilon)) = \Omega(m \log(1/\epsilon))$ gradient queries are required in total. \square

5 Initialization

5.1 Constructing an initial $\mathcal{K}_{\text{in},i}$

In this section, we discuss how to construct an initial set $\mathcal{K}_{\text{in},i}$ to serve as an input to [Algorithm 1](#). In particular, we will prove the following theorem.

Theorem 5.1. *Suppose we are given separation oracle access to a convex set \mathcal{K} that satisfies $\mathcal{B}(\mathbf{z}, r) \subseteq \mathcal{K} \subseteq \mathcal{B}(\mathbf{0}, R)$ for some $\mathbf{z} \in \mathbb{R}^d$. Then, [Algorithm 2](#), in $O(d \log(R/r))$ separation oracle calls to \mathcal{K} , outputs a point \mathbf{x} such that $\mathcal{B}(\mathbf{x}, \frac{r}{6d^{3.5}}) \subseteq \mathcal{K}$.*

Algorithm 2 Inner Ball Finding

```

1:  $\mathcal{K}_{\text{out}} \leftarrow B(\mathbf{0}, R)$ 
2: while true do
3:   Let  $\mathbf{v}$  be the center of gravity of  $\mathcal{K}_{\text{out}}$ 
4:   Sample  $\mathbf{u}$  from  $B(\mathbf{v}, r/(6d))$  uniformly
5:   if  $\mathbf{u} \in \mathcal{K}$  then
6:     Let  $S = \{\mathbf{v} \pm \frac{r}{6d^3} \mathbf{e}_i : i \in [d]\}$ 
7:     if  $S \subset \mathcal{K}$  then
8:       return the inscribed ball of  $\text{conv}(S)$ 
9:     end if
10:  end if
11:  Let  $\mathcal{K}_{\text{out}} \leftarrow \mathcal{K}_{\text{out}} \cap \mathcal{H}$  where  $\mathcal{H} = \mathcal{O}(\mathbf{u})$ 
12: end while

```

Before we prove the preceding theorem, we need the following facts about the self-concordant barrier and convex sets.

Theorem 5.2 ([Nes04, Theorem 4.2.6]). *Let $\psi : \text{int}(\mathcal{K}) \rightarrow \mathbb{R}$ be a ν -self-concordant barrier with the minimizer \mathbf{x}_ψ^* . Then, for any $\mathbf{x} \in \text{int}(\mathcal{K})$ we have:*

$$\|\mathbf{x}_\psi^* - \mathbf{x}\|_{\mathbf{x}_\psi^*} \leq \nu + 2\sqrt{\nu}.$$

On the other hand, for any $\mathbf{x} \in \mathbb{R}^d$ such that $\|\mathbf{x} - \mathbf{x}_\psi^*\|_{\mathbf{x}_\psi^*} \leq 1$, we have $x \in \text{int}(\mathcal{K})$.

Theorem 5.3 ([KLS95, Theorem 4.1]). *Let $\mathcal{K} \subseteq \mathbb{R}^d$ be a convex set with center of gravity μ and covariance matrix Σ . Then,*

$$\{\mathbf{x} : \|\mathbf{x} - \mu\|_{\Sigma^{-1}} \leq \sqrt{(d+2)/d}\} \subseteq \mathcal{K} \subseteq \{\mathbf{x} : \|\mathbf{x} - \mu\|_{\Sigma^{-1}} \leq \sqrt{d(d+2)}\}.$$

Theorem 5.4 ([BGVV14, Section 1.4.2]). *Let \mathcal{K} be a convex set with $\mathcal{K} \subset B(\mathbf{u}, R)$ for some R . Let $\mathcal{K}_{-\delta} = \{\mathbf{x} : B(\mathbf{x}, \delta) \subset \mathcal{K}\}$. Then, we have*

$$\text{vol}\mathcal{K}_{-\delta} \geq \text{vol}\mathcal{K} - (1 - (1 - \frac{\delta}{R})^d) \cdot \text{vol}B(\mathbf{u}, R)$$

Proof of [Theorem 5.1](#). We note that by the description of the [Algorithm 2](#), the returned ball is the inscribed ball of $\text{conv}(S)$ and we have $\mathbf{v} \in \mathcal{K}$ for each $\mathbf{v} \in S$. Then, we must have $\text{conv}(S) \subseteq \mathcal{K}$. We note that $\text{conv}(S)$ is a ℓ_1 ball with ℓ_1 radius $\frac{r}{6d^3}$, then the inscribed ball has ℓ_2 radius $\frac{r}{6d^{3.5}}$.

First, we prove the sample complexity of the algorithm above. We use \mathcal{K}_t to denote the \mathcal{K}_{out} at the t -th iteration. We first observe that throughout the algorithm, \mathcal{K}_t is obtained by intersection of halfspaces and $B(\mathbf{0}, R)$. This implies

$$B(\mathbf{z}, r) \subseteq \mathcal{K} \subseteq \mathcal{K}_t \quad \forall t.$$

Since \mathcal{K}_t contains a ball of radius r , let A_t be the covariance matrix of \mathcal{K}_t . By [Theorem 5.3](#), we have

$$A_t \succeq \frac{r^2}{d(d+2)} I.$$

Let \mathcal{H}_t be the halfspace returned by the oracle at iteration t . We note that \mathbf{u} is sampled uniform from $B(\mathbf{v}, r/(6d))$, so we have

$$\|\mathbf{v} - \mathbf{u}\|_{A^{-1}} \leq \frac{\sqrt{d(d+2)}}{r} \cdot \frac{r}{6d} \leq \frac{1}{3}.$$

Apply the inequality above to [Corollary 2.16](#), we have

$$\text{vol}(\mathcal{K}_t) \leq (1 - 1/e + 1/3)^t \text{vol}(\mathcal{K}_0) \leq (1 - 1/30)^t \text{vol}(B(0, R)).$$

Then, since $B(\mathbf{z}, r) \subseteq \mathcal{K}_t$ for all the t , this implies the algorithm at most takes $O(d \log(R/r))$ many iterations.

Now, we consider the number of oracle calls within each iterations. There are three possible cases to consider:

1. $\mathbf{u} \in \mathcal{K}_{-\delta}$ with $\delta = \frac{r}{6d^3}$ (see the definition of $\mathcal{K}_{-\delta}$ in [Theorem 5.4](#)). In this case, we have $S \subset \mathcal{K}$ and this is the last iteration. We can pay this $O(d)$ oracle calls for the last iteration.
2. $\mathbf{u} \in \mathcal{K} \setminus \mathcal{K}_{-\delta}$.

Since \mathbf{u} is uniformly sampled from $B(\mathbf{v}, r/(6d))$, [Theorem 5.4](#) shows that $\mathbf{u} \in \mathcal{K} \setminus \mathcal{K}_{-\delta}$ with probability at most

$$1 - \left(1 - \frac{\delta}{r/(6d)}\right)^d \leq \frac{1}{d}.$$

Hence, this case only happens with probability only at most $1/d$. Since the cost of checking $S \subset \mathcal{K}$ takes $O(d)$ oracle calls. The expected calls for this case is only $O(1)$.

3. $\mathbf{u} \notin K$. The cost is just 1 call.

Combining all the cases, the expected calls is $O(1)$ per iteration. □

5.2 Initial point reduction

In this section, we will show how to obtain an initial feasible point for the algorithm.

Definition 5.5. Given a convex program $\min_{\mathbf{Ax}=\mathbf{b}, \mathbf{x} \in \mathcal{K} \subseteq \mathbb{R}^d} \mathbf{c}^\top \mathbf{x}$ and some $s > 0$, we define $\mathbf{c}_1 = \mathbf{c}$, $\mathbf{c}_2 = \mathbf{c}_3 = \frac{s \|\mathbf{c}\|_2}{\sqrt{d}} \cdot \mathbf{1}$ and $\mathcal{P} = \{\mathbf{x}^{(1)} \in \mathcal{K}, (\mathbf{x}^{(2)}, \mathbf{x}^{(3)}) \in \mathbb{R}_{\geq 0}^{2d} : \mathbf{A}(\mathbf{x}^{(1)} + \mathbf{x}^{(2)} - \mathbf{x}^{(3)}) = \mathbf{b}\}$. We then define the *modified convex program* by

$$\min_{(\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \mathbf{x}^{(3)}) \in \mathcal{P}} \mathbf{c}_1^\top \mathbf{x}^{(1)} + \mathbf{c}_2^\top \mathbf{x}^{(2)} + \mathbf{c}_3^\top \mathbf{x}^{(3)}.$$

We denote $(\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3)$ by $\bar{\mathbf{c}}$.

Theorem 5.6. Given a convex program $\min_{\mathbf{Ax}=\mathbf{b}, \mathbf{x} \in \mathcal{K} \subseteq \mathbb{R}^d} \mathbf{c}^\top \mathbf{x}$ with outer radius R and some convex set \mathcal{K}_{in} with $\mathcal{K}_{in} \subseteq \mathcal{K}$ and inner radius r . For any modified convex program as in [Definition 5.5](#) with $s \geq 48\nu t \sqrt{d} \cdot \frac{R}{r} \cdot \|\mathbf{c}\|_2 R$. For an arbitrary $t \in \mathbb{R}_{\geq 0}$, we define the function

$$f_t(\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \mathbf{x}^{(3)}) = t(\mathbf{c}_1^\top \mathbf{x}^{(1)} + \mathbf{c}_2^\top \mathbf{x}^{(2)} + \mathbf{c}_3^\top \mathbf{x}^{(3)}) + \psi_{\mathcal{P}_{in}}(\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \mathbf{x}^{(3)})$$

where $\psi_{\mathcal{P}_{in}}$ is some ν self-concordant barrier for the set

$$\mathcal{P}_{in} = \{\mathbf{x}^{(1)} \in \mathcal{K}_{in}, (\mathbf{x}^{(2)}, \mathbf{x}^{(3)}) \in \mathbb{R}_{\geq 0}^{2d} : \mathbf{A}(\mathbf{x}^{(1)} + \mathbf{x}^{(2)} - \mathbf{x}^{(3)}) = \mathbf{b}\}.$$

Given $\bar{\mathbf{x}}_t \stackrel{\text{def}}{=} (\mathbf{x}_t^{(1)}, \mathbf{x}_t^{(2)}, \mathbf{x}_t^{(3)}) = \arg \min_{(\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \mathbf{x}^{(3)}) \in \mathcal{P}_{in}} f_t(\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \mathbf{x}^{(3)})$, we denote $\mathbf{x}_{in} = \mathbf{x}_t^{(1)} + \mathbf{x}_t^{(2)} - \mathbf{x}_t^{(3)}$. Suppose $\min_{\bar{\mathbf{x}} \in \mathcal{P}_{in}} \bar{\mathbf{c}}^\top \bar{\mathbf{x}} \leq \min_{\bar{\mathbf{x}} \in \mathcal{P}} \bar{\mathbf{c}}^\top \bar{\mathbf{x}} + \gamma$, we have the following

1. $\mathbf{Ax}_{in} = \mathbf{b}$,
2. $\mathbf{x}_{in} \in \mathcal{K}_{in}$,
3. $\mathbf{c}^\top \mathbf{x}_{in} \leq \min_{\mathbf{x} \in \mathcal{K}, \mathbf{Ax}=\mathbf{b}} \mathbf{c}^\top \mathbf{x} + \frac{\nu+1}{t} + \gamma$.

First, we show that $\mathbf{x}_t^{(1)}$ is not too close to the boundary. Before we proceed, we need the following lemmas.

Lemma 5.7 (Theorem 4.2.5 [Nes04]). *Let ψ be a ν -self-concordant barrier. Then, for any $\mathbf{x} \in \text{dom}(\psi)$ and $\mathbf{y} \in \text{dom}(\psi)$ such that*

$$\langle \psi'(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle \geq 0,$$

we have

$$\|\mathbf{y} - \mathbf{x}\|_{\mathbf{x}} \leq \nu + 2\sqrt{\nu}.$$

Lemma 5.8 (Theorem 2 of [ZLY22]). *Given a convex set⁴ Ω with a ν -self-concordant barrier ψ_{Ω} and inner radius r . Let $\mathbf{x}_t = \arg \min_{\mathbf{x}} t \cdot \mathbf{c}^{\top} \mathbf{x} + \psi_{\Omega}(\mathbf{x})$. Then, for any $t > 0$,*

$$\min \left\{ \frac{1}{2t}, \frac{r\|\mathbf{c}\|_2}{4\nu + 4\sqrt{\nu}} \right\} \leq \mathbf{c}^{\top} \mathbf{x}_t - \mathbf{c}^{\top} \mathbf{x}_{\infty} \leq \frac{\nu}{t}.$$

Consider the optimization problem restricted in the subspace $\{(\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \mathbf{x}^{(3)}) : A(\mathbf{x}^{(1)} + \mathbf{x}^{(2)} - \mathbf{x}^{(3)}) = \mathbf{b}\}$, as a direct corollary of theorem above we have the following:

Corollary 5.9. *Let $\bar{\mathbf{x}}_t$ be as the same as defined in Theorem 5.6. For $t \geq \frac{4\nu}{r\|\mathbf{c}\|_2}$, we have $\text{dist}(\mathbf{x}_t^{(1)}, \mathbf{x}_{\infty}^{(1)}) \geq \frac{1}{2t\|\mathbf{c}\|_2}$.*

Now, we are ready to show $\text{dist}(\mathbf{x}_t^{(1)}, \partial\mathcal{K}_{\text{in}})$ is not too small.

Theorem 5.10. *Let $\bar{\mathbf{x}}_t$ be the same as defined in Theorem 5.6. For $t \geq \frac{4\nu}{r\|\mathbf{c}\|_2}$, we have $\text{dist}(\mathbf{x}_t^{(1)}, \partial\mathcal{K}_{\text{in}}) \geq \frac{r}{12\nu t\|\mathbf{c}\|_2 R}$.*

Proof. We consider the domain restricted in the subspace $\{(\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \mathbf{x}^{(3)}) : \mathbf{A}(\mathbf{x}^{(1)} + \mathbf{x}^{(2)} - \mathbf{x}^{(3)}) = \mathbf{b}\}$. By the optimality of $\bar{\mathbf{x}}_t$ and Lemma 5.7, we have

$$\mathcal{K}_{\mathcal{H}} \subseteq \{\mathbf{x} : \|\mathbf{x} - \mathbf{x}_t^{(1)}\|_{\mathbf{x}_t^{(1)}} \leq \nu + 2\sqrt{\nu}\},$$

where $\mathcal{H} = \{\mathbf{x} : \mathbf{c}^{\top}(\mathbf{x}_t^{(1)} - \mathbf{x}) \geq 0\}$ and $\mathcal{K}_{\mathcal{H}} \stackrel{\text{def}}{=} \mathcal{H} \cap \mathcal{K}_{\text{in}}$.

Recall that \mathcal{K}_{in} contains a ball of radius r , we denote it by B . We note that $\text{conv}(\mathbf{x}_{\infty}^{(1)}, B)$ is a union of a ball and a convex cone \mathcal{C} with diameter at most $2R$. We observe that the set $\text{conv}(\mathbf{x}_{\infty}^{(1)}, B) \cap \mathcal{H}$ contains a ball of radius at least $\frac{r}{4t\|\mathbf{c}\|_2 R}$ since $\text{dist}(\mathbf{x}_{\infty}^{(1)}, \partial\mathcal{H}) \geq \frac{1}{2t\|\mathbf{c}\|_2}$.

We note that

$$\text{conv}(\mathbf{x}_{\infty}^{(1)}, B) \cap \mathcal{H} \subseteq \mathcal{K}_{\text{in}} \subseteq \{\mathbf{x} : \|\mathbf{x} - \mathbf{x}_t^{(1)}\|_{\mathbf{x}_t^{(1)}} \leq \nu + 2\sqrt{\nu}\},$$

this implies $\{\mathbf{x} : \|\mathbf{x} - \mathbf{x}_t^{(1)}\|_{\mathbf{x}_t^{(1)}} \leq \nu + 2\sqrt{\nu}\}$ contains a ball of radius at least $\frac{r}{4t\|\mathbf{c}\|_2 R}$, and then by Theorem 5.2, we have $B(\mathbf{x}_t^{(1)}, \frac{r}{4(\nu + 2\sqrt{\nu})t\|\mathbf{c}\|_2 R}) \subseteq \mathcal{K}_{\text{in}}$. \square

Lemma 5.11. *Let $(\mathbf{x}_t^{(1)}, \mathbf{x}_t^{(2)}, \mathbf{x}_t^{(3)}) \in \mathbb{R}^{3d}$ be the same as defined in Theorem 5.6. If $t > \frac{\nu}{\|\mathbf{c}\|_2 R}$, then we have $\|\mathbf{x}_t^{(2)} - \mathbf{x}_t^{(3)}\|_2 \leq \frac{4\sqrt{d}}{s} R$.*

Proof. Let $\mathbf{x}_{\text{in}}^* = \arg \min_{\mathbf{x} \in \mathcal{K}_{\text{in}}, \mathbf{A}\mathbf{x}=\mathbf{b}} \mathbf{c}^{\top} \mathbf{x}$ and $\bar{\mathbf{x}}_{\text{in}}^* = \arg \min_{\mathbf{x} \in \mathcal{P}_{\text{in}}} \bar{\mathbf{c}}^{\top} \bar{\mathbf{x}}$. Since $\mathbf{x}^* \in \mathcal{B}(0, R)$, we have

$$\mathbf{c}^{\top} \mathbf{x}_{\text{in}}^* \leq \|\mathbf{c}\|_2 R.$$

Note that $(\mathbf{x}_{\text{in}}^*, \mathbf{0}, \mathbf{0}) \in \mathcal{P}_{\text{in}}$, this means we have

$$\bar{\mathbf{c}}^{\top} \bar{\mathbf{x}}_{\text{in}}^* \leq \mathbf{c}^{\top} \mathbf{x}_{\text{in}}^* \leq \|\mathbf{c}\|_2 R.$$

Combining this with the second inequality in Lemma 5.8, we get

$$\bar{\mathbf{c}}^{\top} \bar{\mathbf{x}}_t \leq \bar{\mathbf{c}}^{\top} \bar{\mathbf{x}}_{\text{in}}^* + \frac{\nu}{t} \leq \|\mathbf{c}\|_2 R + \frac{\nu}{t} \leq 2\|\mathbf{c}\|_2 R.$$

⁴The original theorem is stated only for polytopes, but their proof works for general convex sets.

We further note that

$$\mathbf{c}_2^\top \mathbf{x}_t^{(2)} \leq \bar{\mathbf{c}}^\top \bar{\mathbf{x}}_t \leq 2\|\mathbf{c}\|R.$$

This shows

$$\max\{\|\mathbf{x}_t^{(2)}\|_2, \|\mathbf{x}_t^{(3)}\|_2\} \leq \frac{2\sqrt{d}\|\mathbf{c}\|_2 R}{\|\mathbf{c}\|_2^s} \leq \frac{2\sqrt{d}R}{s}.$$

Hence, we have

$$\|\mathbf{x}_t^{(2)} - \mathbf{x}_t^{(3)}\|_2 \leq \frac{4\sqrt{d}}{s}R.$$

□

Now, we are ready to prove [Theorem 5.6](#).

Proof of [Theorem 5.6](#). We note that \mathbf{x}_{in} satisfies (1), directly follows by definition of \mathcal{P} . By assumption, we have $s \geq 48\nu t\sqrt{d} \cdot \frac{R}{r} \cdot \|\mathbf{c}\|_2 R$; using this in [Lemma 5.11](#), we have

$$\|\mathbf{x}_t^{(2)} - \mathbf{x}_t^{(3)}\|_2 \leq \frac{r}{12\nu t\|\mathbf{c}\|_2 R}.$$

This means $\mathbf{x}_{\text{in}} = \mathbf{x}_t^{(1)} + \mathbf{x}_t^{(2)} - \mathbf{x}_t^{(3)} \in \mathcal{K}_{\text{in}}$ since $\text{dist}(\mathbf{x}_t^{(1)}, \partial\mathcal{K}_{\text{in}}) \geq \frac{r}{12\nu t\|\mathbf{c}\|_2 R}$.

Now, we show $\mathbf{c}^\top \mathbf{x}_{\text{in}}$ is close to $\mathbf{c}^\top \mathbf{x}^*$.

Let $\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathcal{K}, \mathbf{Ax}=\mathbf{b}} \mathbf{c}^\top \mathbf{x}$ and $\bar{\mathbf{x}}^* = \arg \min_{\mathbf{x} \in \mathcal{P}} \bar{\mathbf{c}}^\top \bar{\mathbf{x}}$. By [Lemma 5.8](#), we have

$$\bar{\mathbf{c}}^\top \bar{\mathbf{x}}_t - \frac{\nu}{t} \leq \bar{\mathbf{c}}^\top \bar{\mathbf{x}}_{\text{in}} \leq \bar{\mathbf{c}}^\top \bar{\mathbf{x}}^* + \gamma \leq \mathbf{c}^\top \mathbf{x}^* + \gamma.$$

This implies

$$\mathbf{c}^\top \mathbf{x}_t^{(1)} \leq \bar{\mathbf{c}}^\top \bar{\mathbf{x}}_t \leq \mathbf{c}^\top \mathbf{x}^* + \frac{\nu}{t} + \gamma.$$

We have

$$\mathbf{c}^\top \mathbf{x}_{\text{in}} = \mathbf{c}^\top (\mathbf{x}_t^{(1)} + \mathbf{x}_t^{(2)} - \mathbf{x}_t^{(3)}) \leq \mathbf{c}^\top \mathbf{x}^* + \frac{\nu}{t} + \frac{4}{s}\|\mathbf{c}\|_2 R \leq \mathbf{c}^\top \mathbf{x}^* + \frac{\nu+1}{t} + \gamma.$$

□

6 Acknowledgements

We thank Ian Covert for helpful discussions about the problem applications and Ewin Tang for helpful feedback on the paper.

References

- [AB15] Alekh Agarwal and Leon Bottou. “A lower bound for the optimization of finite sums”. In: *International conference on machine learning*. PMLR. 2015, pp. 78–86.
- [AKMSV21] Kyriakos Axiotis, Adam Karczmarz, Anish Mukherjee, Piotr Sankowski, and Adrian Vladu. “Decomposable submodular function minimization via maximum flow”. In: *International Conference on Machine Learning*. PMLR. 2021, pp. 446–456.
- [All17] Zeyuan Allen-Zhu. “Katyusha: The first direct acceleration of stochastic gradient methods”. In: *The Journal of Machine Learning Research* 18.1 (2017), pp. 8194–8244.
- [AY16] Zeyuan Allen-Zhu and Yang Yuan. “Improved SVRG for non-strongly-convex or sum-of-non-convex objectives”. In: *International conference on machine learning*. PMLR. 2016, pp. 1080–1089.
- [BBV04] Stephen Boyd, Stephen P Boyd, and Lieven Vandenbergh. *Convex optimization*. Cambridge university press, 2004.
- [BC03] Léon Bottou and Yann Cun. “Large scale online learning”. In: *Advances in neural information processing systems* 16 (2003).
- [BE15] Sébastien Bubeck and Ronen Eldan. “The entropic barrier: a simple and optimal universal self-concordant barrier”. In: *Conference on Learning Theory*. 2015, pp. 279–279.
- [BGVV14] Silouanos Brazitikos, Apostolos Giannopoulos, Petros Valettas, and Beatrice-Helen Vritsiou. *Geometry of isotropic convex bodies*. Vol. 196. American Mathematical Soc., 2014.
- [BKLLS20] Sébastien Bubeck, Bo’az Klartag, Yin Tat Lee, Yuanzhi Li, and Mark Sellke. “Chasing nested convex bodies nearly optimally”. In: *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms*. 2020, pp. 1496–1508.
- [Bot12] Léon Bottou. “Stochastic gradient descent tricks”. In: *Neural networks: Tricks of the trade*. Springer, 2012, pp. 421–436.
- [BV02] Dimitris Bertsimas and Santosh Vempala. “Solving convex programs by random walks”. In: *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing (STOC)*. ACM. 2002, pp. 109–115.
- [Cau+47] Augustin Cauchy et al. “Méthode générale pour la résolution des systemes d’équations simultanées”. In: *Comp. Rend. Sci. Paris* 25.1847 (1847), pp. 536–538.
- [Che21] Sinho Chewi. “The entropic barrier is n -self-concordant”. In: *arXiv preprint arXiv:2112.10947* (2021).
- [CL11] Chih-Chung Chang and Chih-Jen Lin. “LIBSVM: A library for support vector machines”. In: *ACM Transactions on Intelligent Systems and Technology* 2 (3 2011). Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>, 27:1–27:27.
- [DBL14] Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. “SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives”. In: *Advances in neural information processing systems* 27 (2014).
- [DLY21] Sally Dong, Yin Tat Lee, and Guanghao Ye. “A nearly-linear time algorithm for linear programs with small treewidth: a multiscale representation of robust central path”. In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 2021, pp. 1784–1797.
- [FGKS15] Roy Frostig, Rong Ge, Sham Kakade, and Aaron Sidford. “Un-regularizing: approximate proximal point and faster stochastic algorithms for empirical risk minimization”. In: *International Conference on Machine Learning*. PMLR. 2015, pp. 2540–2548.
- [FJPZ13] Alexander Fix, Thorsten Joachims, Sung Min Park, and Ramin Zabih. “Structured learning of sum-of-submodular higher order energy functions”. In: *Proceedings of the IEEE International Conference on Computer Vision*. 2013, pp. 3104–3111.
- [GLS88] Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric algorithms and combinatorial optimization*. Springer, 1988.

- [Grü60] Branko Grünbaum. “Partitions of mass-distributions and of convex bodies by hyperplanes.” In: *Pacific Journal of Mathematics* 10.4 (1960), pp. 1257–1261.
- [Gül97] Osman Güler. “On the self-concordance of the universal barrier function”. In: *SIAM Journal on Optimization* 7.2 (1997), pp. 295–303.
- [HL16] Elad Hazan and Haipeng Luo. “Variance-reduced and projection-free stochastic optimization”. In: *International Conference on Machine Learning*. PMLR. 2016, pp. 1263–1271.
- [JLLV21] He Jia, Aditi Laddha, Yin Tat Lee, and Santosh Vempala. “Reducing isotropy and volume to KLS: an $\mathcal{O}(n^3 \psi^2)$ volume algorithm”. In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 2021, pp. 961–974.
- [JLSW20] Haotian Jiang, Yin Tat Lee, Zhao Song, and Sam Chiu-wai Wong. “An improved cutting plane method for convex optimization, convex-concave games, and its applications”. In: *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*. 2020, pp. 944–953.
- [JZ13] Rie Johnson and Tong Zhang. “Accelerating stochastic gradient descent using predictive variance reduction”. In: *Advances in neural information processing systems* 26 (2013).
- [Kla06] Boas Klartag. “On convex perturbations with a bounded isotropic constant”. In: *Geometric & Functional Analysis GAFA* 16.6 (2006), pp. 1274–1290.
- [KLS95] Ravi Kannan, László Lovász, and Miklós Simonovits. “Isoperimetric problems for convex bodies and a localization lemma”. In: *Discrete & Computational Geometry* 13.3 (1995), pp. 541–559.
- [KLT09] Pushmeet Kohli, Lubor Ladicky, and Philip H. S. Torr. “Robust higher order potentials for enforcing label consistency”. In: *International Journal of Computer Vision* 82.3 (2009), pp. 302–324.
- [KT10] Alex Kulesza and Ben Taskar. “Structured determinantal point processes”. In: *Advances in neural information processing systems* 23 (2010).
- [KTE88] Leonid G Khachiyan, Sergei Pavlovich Tarasov, and I. I. Erlikh. “The method of inscribed ellipsoids”. In: *Soviet Math. Dokl.* Vol. 37. 1988, pp. 226–230.
- [LMH15] Hongzhou Lin, Julien Mairal, and Zaid Harchaoui. “A universal catalyst for first-order optimization”. In: *Advances in neural information processing systems* 28 (2015).
- [LSW15] Yin Tat Lee, Aaron Sidford, and Sam Chiu-wai Wong. “A faster cutting plane method and its implications for combinatorial and convex optimization”. In: *56th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 2015, pp. 1049–1065.
- [LV21] Yin Tat Lee and Santosh S Vempala. “Tutorial on the Robust Interior Point Method”. In: *arXiv preprint arXiv:2108.04734* (2021).
- [LY21] Yin Tat Lee and Man-Chung Yue. “Universal barrier is n -self-concordant”. In: *Mathematics of Operations Research* 46.3 (2021), pp. 1129–1148.
- [Mai15] Julien Mairal. “Incremental majorization-minimization optimization with application to large-scale machine learning”. In: *SIAM Journal on Optimization* 25.2 (2015), pp. 829–855.
- [MZJ13] Mehrdad Mahdavi, Lijun Zhang, and Rong Jin. “Mixed optimization for smooth functions”. In: *Advances in neural information processing systems* 26 (2013).
- [Nes04] Yurii E. Nesterov. *Introductory Lectures on Convex Optimization - A Basic Course*. Vol. 87. Applied Optimization. Springer, 2004. DOI: 10.1007/978-1-4419-8853-9.
- [Nes83] Yurii E Nesterov. “A method for solving the convex programming problem with convergence rate $\mathcal{O}(1/k^2)$ ”. In: *Dokl. akad. nauk Sssr*. Vol. 269. 1983, pp. 543–547.
- [NN89] YE Nesterov and AS Nemirovskii. “Self-concordant functions and polynomial time methods in convex programming. preprint, Central Economic & Mathematical Institute, USSR Acad”. In: *Sci. Moscow, USSR* (1989).
- [NN94a] Yurii Nesterov and Arkadii Nemirovskii. *Interior-point polynomial algorithms in convex programming*. SIAM, 1994.

- [NN94b] Yurii E. Nesterov and Arkadii Nemirovskii. *Interior-point polynomial algorithms in convex programming*. Vol. 13. Siam studies in applied mathematics. SIAM, 1994. ISBN: 978-0-89871-319-0. DOI: 10.1137/1.9781611970791. URL: <https://doi.org/10.1137/1.9781611970791>.
- [Ren01] James Renegar. *A mathematical view of interior-point methods in convex optimization*. SIAM, 2001.
- [RM51] Herbert Robbins and Sutton Monro. “A stochastic approximation method”. In: *The annals of mathematical statistics* (1951), pp. 400–407.
- [Roc70] R Tyrrell Rockafellar. *Convex Analysis*. Vol. 36. Princeton University Press, 1970.
- [RRWN11] Benjamin Recht, Christopher Re, Stephen Wright, and Feng Niu. “Hogwild!: A lock-free approach to parallelizing stochastic gradient descent”. In: *Advances in neural information processing systems* 24 (2011).
- [RSB12] Nicolas Roux, Mark Schmidt, and Francis Bach. “A stochastic gradient method with an exponential convergence rate for finite training sets”. In: *Advances in neural information processing systems* 25 (2012).
- [RW09] R Tyrrell Rockafellar and Roger J-B Wets. *Variational analysis*. Vol. 317. Springer Science & Business Media, 2009.
- [SK10] Peter Stobbe and Andreas Krause. “Efficient minimization of decomposable submodular functions”. In: *Advances in Neural Information Processing Systems* 23 (2010).
- [SLB17] Mark Schmidt, Nicolas Le Roux, and Francis Bach. “Minimizing finite sums with the stochastic average gradient”. In: *Mathematical Programming* 162.1 (2017), pp. 83–112.
- [SZ13a] Shai Shalev-Shwartz and Tong Zhang. “Accelerated mini-batch stochastic dual coordinate ascent”. In: *Advances in Neural Information Processing Systems* 26 (2013).
- [SZ13b] Shai Shalev-Shwartz and Tong Zhang. “Stochastic dual coordinate ascent methods for regularized loss minimization.” In: *Journal of Machine Learning Research* 14.2 (2013).
- [Vai89] Pravin M Vaidya. “A new algorithm for minimizing convex functions over convex sets”. In: *30th Annual Symposium on Foundations of Computer Science*. 1989, pp. 338–343.
- [VBK20] Nate Veldt, Austin R Benson, and Jon Kleinberg. “Minimizing localized ratio cut objectives in hypergraphs”. In: *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2020, pp. 1708–1718.
- [VKR09] Sara Vicente, Vladimir Kolmogorov, and Carsten Rother. “Joint optimization of segmentation and appearance models”. In: *2009 IEEE 12th international conference on computer vision*. IEEE. 2009, pp. 755–762.
- [WS16] Blake E Woodworth and Nati Srebro. “Tight complexity bounds for optimizing composite objectives”. In: *Advances in neural information processing systems* 29 (2016).
- [Zha04] Tong Zhang. “Solving large scale linear prediction problems using stochastic gradient descent algorithms”. In: *Proceedings of the twenty-first international conference on Machine learning*. 2004, p. 116.
- [ZL15] Yuchen Zhang and Xiao Lin. “Stochastic primal-dual coordinate method for regularized empirical risk minimization”. In: *International Conference on Machine Learning*. PMLR. 2015, pp. 353–361.
- [ZLY22] Manru Zong, Yin Tat Lee, and Man-Chung Yue. “Short-step Methods Are Not Strongly Polynomial-Time”. In: *arXiv preprint arXiv:2201.02768* (2022).

A Decomposable submodular function minimization

A.1 Preliminaries

Throughout, V denotes the ground set of elements. A set function $f : 2^V \rightarrow \mathbb{R}$ is *submodular* if it satisfies the following *diminishing marginal differences* property:

Definition A.1 (Submodularity). A function $f : 2^V \rightarrow \mathbb{R}$ is submodular if $f(T \cup \{i\}) - f(T) \leq f(S \cup \{i\}) - f(S)$, for any subsets $S \subseteq T \subseteq V$ and $i \in V \setminus T$.

We may assume without loss of generality that $f(\emptyset) = 0$ by replacing $f(S)$ by $f(S) - f(\emptyset)$. We assume that f is accessed by an *evaluation oracle* and use EO to denote the time to compute $f(S)$ for a subset S . Our algorithm for decomposable SFM is based on the Lovász extension [GLS88], a standard convex extension of a submodular function.

Definition A.2 (Lovász extension [GLS88]). The Lovász extension $\hat{f} : [0, 1]^V \rightarrow \mathbb{R}$ of a submodular function f is defined as

$$\hat{f}(x) = \mathbb{E}_{t \sim [0, 1]} [f(\{i \in V : x_i \geq t\})],$$

where $t \sim [0, 1]$ is drawn uniformly at random from $[0, 1]$.

The Lovász extension \hat{f} of a submodular function f has many desirable properties. In particular, \hat{f} is a convex relaxation of f and it can be evaluated efficiently.

Theorem A.3 (Properties of Lovász extension [GLS88]). *Let $f : 2^V \rightarrow \mathbb{R}$ be a submodular function and \hat{f} be its Lovász extension. Then,*

- (a) \hat{f} is convex and $\min_{x \in [0, 1]^V} \hat{f}(x) = \min_{S \subseteq V} f(S)$;
- (b) $f(S) = \hat{f}(I_S)$ for any subset $S \subseteq V$, where I_S is the indicator vector for S ;
- (c) Suppose $x \in [0, 1]^V$ satisfies $x_1 \geq \dots \geq x_{|V|}$, then $\hat{f}(x) = \sum_{i=1}^{|V|} (f([i]) - f([i-1]))x_i$.

Property (c) in [Theorem A.3](#) allows us to implement a sub-gradient oracle for \hat{f} by evaluating f .

Theorem A.4 (Sub-gradient oracle implementation for Lovász extension, Theorem 61 of [LSW15]). *Let $f : 2^V \rightarrow \mathbb{R}$ be a submodular function and \hat{f} be its Lovász extension. Then a sub-gradient for \hat{f} can be implemented in time $O(|V| \cdot \text{EO} + |V|^2)$.*

A.2 Decomposable submodular function minimization proofs

In this subsection, we prove the following more general version of [Theorem 1.4](#).

Theorem A.5 (Decomposable SFM). *Let $F : V \rightarrow [-1, 1]$ be given by $F(S) = \sum_{i=1}^n F_i(S \cap V_i)$, where each $F_i : 2^{V_i} \rightarrow \mathbb{R}$ is a submodular function on $V_i \subseteq V$ with $|V_i| = d_i$. Let $m = \sum_{i=1}^n d_i$ and $d_{\max} := \max_{i \in [n]} d_i$. Then we can find an ϵ -approximate minimizer of f using at most $O(d_{\max} m \log(m/\epsilon))$ evaluation oracle calls.*

Proof. Let \hat{f}_i be the Lovász extension of each f_i , then $\hat{f} = \sum_{i=1}^n \hat{f}_i$ is the Lovász extension of f . Note that \hat{f} is 2-Lipschitz since the range of f is $[-1, 1]$. Also, the diameter of the range $[0, 1]^{V_i}$ for each Lovász extension \hat{f}_i is at most $\sqrt{|V_i|} \leq \sqrt{d_{\max}}$. Thus using [Theorem 4.11](#), we can find a vector $x \in [0, 1]^V$ such that $\hat{f}(x) \leq \min_{x^* \in [0, 1]^V} \hat{f}(x^*) + \epsilon$ in $\text{poly}(m \log(1/\epsilon))$ time and $O(m \log(m\sqrt{d_{\max}}/\epsilon)) = O(m \log(m/\epsilon))$ subgradients of the \hat{f}_i 's. By [Theorem A.4](#), each sub-gradient of \hat{f}_i can be computed by making at most $d_i \leq d_{\max}$ queries to the evaluation oracle for f_i . Thus the total number of evaluation oracle calls we make in finding an ϵ -additive approximate minimizer $x \in [0, 1]^V$ of \hat{f} is at most $O(d_{\max} m \log(m/\epsilon))$.

Next we turn the ϵ -additive approximate minimizer x of \hat{f} into an ϵ -additive approximate minimizer $S \subseteq V$ for f . Without loss of generality, assume that $x_1 \geq \dots \geq x_{|V|}$. Then by property (c) in [Theorem A.3](#), we have

$$\hat{f}(x) = \sum_{i=1}^{|V|} (f([i]) - f([i-1]))x_i = f(V) \cdot x_{|V|} + \sum_{i=1}^{|V|-1} f([i]) \cdot (x_i - x_{i+1}).$$

Since $x_i - x_{i+1} \geq 0$, the above implies that $\min_{i \in \{1, \dots, |V|\}} f([i]) \leq \hat{f}(x)$. Thus we can find a subset $S \subseteq V$ among $f([i])$ for all $i \in \{1, \dots, |V|\}$ such that $f(S) \leq \hat{f}(x)$. Then by property (a) in [Theorem A.3](#), the set S is an ϵ -additive approximate minimizer of f . This proves the theorem. \square