
Log-concave Sampling from a Convex Body with a Barrier: a Robust and Unified Dikin Walk

Yuzhou Gu
New York University
yuzhougu@nyu.edu

Nikki Lijing Kuang
University of California, San Diego
11kuang@ucsd.edu

Yi-An Ma
University of California, San Diego
yianma@ucsd.edu

Zhao Song
Simons Institute for the Theory of Computing, UC Berkeley
magic.linuxkde@gmail.com

Lichen Zhang
MIT CSAIL
lichenz@csail.mit.edu

Abstract

We consider the problem of sampling from a d -dimensional log-concave distribution $\pi(\theta) \propto \exp(-f(\theta))$ for L -Lipschitz f , constrained to a convex body with an efficiently computable self-concordant barrier function, contained in a ball of radius R with a w -warm start.

We propose a *robust* sampling framework that computes spectral approximations to the Hessian of the barrier functions in each iteration. We prove that for polytopes that are described by n hyperplanes, sampling with the Lee-Sidford barrier function mixes within $\tilde{O}((d^2 + dL^2R^2) \log(w/\delta))$ steps with a per step cost of $\tilde{O}(nd^{\omega-1})$, where $\omega \approx 2.37$ is the fast matrix multiplication exponent. Compared to the prior work of Mangoubi and Vishnoi, our approach gives faster mixing time as we are able to design a generalized soft-threshold Dikin walk beyond log-barrier.

We further extend our result to show how to sample from a d -dimensional spectrahedron, the constrained set of a semidefinite program, specified by the set $\{x \in \mathbb{R}^d : \sum_{i=1}^d x_i A_i \succeq C\}$ where A_1, \dots, A_d, C are $n \times n$ real symmetric matrices. We design a walk that mixes in $\tilde{O}((nd + dL^2R^2) \log(w/\delta))$ steps with a per iteration cost of $\tilde{O}(n^\omega + n^2d^{3\omega-5})$. We improve the mixing time bound of prior best Dikin walk due to Narayanan and Rakhlin that mixes in $\tilde{O}((n^2d^3 + n^2dL^2R^2) \log(w/\delta))$ steps.

1 Introduction

Given a convex body, generate samples from the body according to structured densities is a fundamental problem in computer science and machine learning. It has extensive applications in constrained convex optimization [Lovász and Vempala, 2006, Narayanan, 2016], differentially private learning [Wang et al., 2015, Lin et al., 2024] and online optimization [Narayanan and Rakhlin, 2017]. A central theme in the theory of sampling is to leverage stochasticity and reduce per iteration costs without having to proportionally increase the number of iterations. This theme has played out in continuous optimization for both first and second order methods. For the first order methods, focus is on reducing the variance of the gradient estimators [Johnson and Zhang, 2013, Shalev-Shwartz and Zhang, 2013, Defazio et al., 2014]. For the second order methods, matrix sketching is often used to

reduce computation and storage related to the Hessian matrix [Lee et al., 2019, Jiang et al., 2020b, 2021, Song and Yu, 2021, Qin et al., 2023].

In Markov chain Monte Carlo (MCMC), much of the progress is on the first order methods. Non-asymptotic analyses are performed for the stochastic gradient Langevin algorithms and their variance reduced extensions [Dubey et al., 2016, Raginsky et al., 2017, Brosse et al., 2018, Chatterji et al., 2018, Zou et al., 2018, Dalalyan and Karagulyan, 2019, Li et al., 2019, Ding and Li, 2021]. For second order methods, there has been a paucity when it comes to applying the sketching techniques.

In this paper, we focus on sampling from a log-concave distribution constrained to a d -dimensional convex body \mathcal{K} that is described by n constraints, where second order information is proven essential for capturing the geometry of the convex body and consequently for achieving fast convergence rates [Narayanan, 2016, Narayanan and Rakhlin, 2017, Chen et al., 2018, Laddha et al., 2020, Mangoubi and Vishnoi, 2023, 2024]. In particular, we associate the convex body with a *self-concordant barrier function* [Nesterov and Nemirovskii, 1994, Renegar, 1988, Vaidya, 1989] and utilize the Hessian matrix $H(x)$ of the barrier function in the sampling algorithm. We then use the Hessian matrix to propose samples and compute the probability to accept or reject the proposed samples. We note that this problem involves subtleties beyond the scope of constrained optimization. In continuous optimization, which focuses on finding the descent directions, unbiased estimators with reasonable variance oftentimes suffice [Vaidya, 1989, Lee et al., 2015, Huang et al., 2022]. In MCMC, on the other hand, the target probability distribution needs to be maintained along the entire trajectory. This poses significant challenges to speed up the sampling algorithms. In the scenario of uniform sampling over a polytope, Laddha et al. [2020] shows that for a simple logarithmic barrier¹, an unbiased estimator in fact suffices. However, more complicated barriers such as the Lee-Sidford barrier [Lee and Sidford, 2014, 2019] can only be approximately computed², an unbiased estimator is extremely hard to be obtained. Moreover, for sampling from more sophisticated log-concave distributions and convex bodies, a more general approach is needed.

We therefore propose to obtain a high precision estimator to the desired acceptance rate with an improved per step running time and without sacrificing the rapid mixing time. In particular, we ask the following question:

Can we significantly reduce per iteration cost, while preserving the convergence rate of the log-concave sampling algorithms over various convex bodies?

We answer the above question in the affirmative. To this end, we present a slew of results regarding log-concave sampling. For polytopes, we provide a walk that mixes in $\tilde{O}(d^2 + dL^2R^2)$ steps³ with per iteration cost $\tilde{O}(nd^{\omega-1})$. Prior state-of-the-art results are due to Mangoubi and Vishnoi [2023, 2024], for which their walks mix in $\tilde{O}(nd + dL^2R^2)$ with a per iteration cost⁴ $\tilde{O}(\text{nnz}(A) + d^2)$. Our walk mixes faster whenever $n \geq d$. Our result partially resolves an open problem in Mangoubi and Vishnoi [2023] as they asked whether it's possible to design a Dikin walk whose mixing time is only depends on d and independent of L, R . We remove the L and R dependence on the dominating term d^2 , hence for the case of sampling from uniform distribution ($L = 0$), we recover the state-of-art result of Laddha et al. [2020] which mixes in $\tilde{O}(d^2)$ steps with a per iteration cost $\tilde{O}(nd^{\omega-1})$.

We also consider sampling from convex bodies beyond polytopes. In semidefinite programming (SDP), one often focuses on the dual program of an SDP, where the constraint set is defined as a spectrahedron $\mathcal{K} = \{x \in \mathbb{R}^d : \sum_{i=1}^d x_i A_i \succeq C\}$ ⁵ for $n \times n$ symmetric matrices A_i and C [Jiang et al., 2020a, Huang et al., 2022]. We propose a walk that samples from a log-concave distribution over a spectrahedron \mathcal{K} using the Hessian information of the log-barrier. The walk mixes in $\tilde{O}(nd + dL^2R^2)$ steps. For log-barrier of an SDP, explicitly forming the Hessian matrix takes a prohibitively large

¹Logarithmic barriers, or log-barriers have been extensively studied in the context of sampling and convex optimization. For a d -dimensional polytope with n constraints, it typically leads to a convergence rate polynomially depends on n .

²Lee-Sidford barriers are the first polynomial-time computable barriers with complexity parameter depends only logarithmic on n and in extension, the convergence rate.

³We use $\tilde{O}(\cdot)$ to hide polylogarithmic dependence on n, d and δ where δ is the TV distance between the target distribution and our Markov chain μ .

⁴We use $\text{nnz}(A)$ to denote the number of nonzero entries in matrix A .

⁵We use $A \succeq 0$ to denote A is a positive semidefinite matrix, and $A \succeq B$ to denote $A - B \succeq 0$.

$O(n^\omega d + n^2 d^{\omega-1})$ time. We utilize our robust sampling framework to approximately compute this Hessian via tensor-based sketching techniques, achieving a runtime of $\tilde{O}(n^\omega + n^2 d^{3\omega-5})$. As long as $n \geq d^{\frac{3\omega-4}{\omega-2}}$ (which is a usual setting for SDP, where $d \ll n$), this provides a speedup.

1.1 Related works

In this work, we focus on Dikin walk [Kannan and Narayanan, 2012], a refined variant of ball walk [Lovász and Simonovits, 1993]. Roughly speaking, ball walk progresses by moving to a random point in the ball centered at the current point with the obvious downside that when the convex body is flat, ball walk progresses slowly. Dikin walk overcomes this problem by instead moving to a random point in a good ellipse centered at the current point, in an effort to better utilize the local geometry of the convex body.

For sampling over polytopes, a number of Dikin walks use the ellipse induced by the log-barrier functions [Kannan and Narayanan, 2012]. The work of Laddha et al. [2020] shows that uniform sampling over a polytope can mix in $O(\nu d)$ steps, where ν is the self-concordant parameter of the barrier function⁶. Going beyond uniform distributions, Narayanan and Rakhlin [2017] proposes a walk that samples from a log-concave distribution on a polytope in $\tilde{O}(\nu^2 d)$ steps. This bound is later improved by Mangoubi and Vishnoi [2023], in which they show that for logarithmic barrier, a mixing time of $\tilde{O}(nd + dL^2 R^2)$ is attainable, where L is the Lipschitz constant of f and R is the radius of the ball that contains \mathcal{K} . A more recent work by Kook and Vempala [2024] has obtained better bound when the function f is α -relative strongly-convex and the density π is β -Lipschitz. They manage to obtain a mixing bound of $O(\nu d \beta \log(w/\delta))$. However, their algorithm requires stronger assumptions on f and hence is incomparable to our result. In most previous works, the focus has been on improving the mixing time, rather than on per iteration costs.

For any convex body, it is well-known that a universal barrier with $\nu = d$ exists [Nesterov and Nemirovskii, 1989, Lee and Yue, 2021], however it is computationally hard to construct the Hessian of the universal barrier. In short, the universal barrier requires to compute the volume of the polar set of a high dimensional body, which is even hard to approximate deterministically [Furedi and Barany, 1986, Nesterov and Nemirovskii, 1994]. The seminal work of Lee and Sidford [Lee and Sidford, 2014] presents a nearly-universal barrier with $\nu = O(d \log^5 n)$ and the Hessian can be approximated in $O(nd^{\omega-1})$ time. The Lee-Sidford barrier was originally designed to solve linear programs in $O(\sqrt{d})$ iterations, and it has been leveraged lately for Dikin walks with rapid mixing time. For uniform sampling, Chen et al. [2018] gives an analysis with a walk that mixes in $\tilde{O}(d^{2.5})$ steps. Subsequently, Laddha et al. [2020] improves the mixing to $\tilde{O}(d^2)$ steps. In a pursuit to better leverage the local geometry, Lee and Vempala [2017] proposes a walk relying on Riemannian metric that mixes in $\tilde{O}(nd^{3/4})$ steps. The mixing rate is later improved to $\tilde{O}(nd^{2/3})$ via Riemannian Hamiltonian Monte Carlo (RHMC) with log-barrier [Lee and Vempala, 2018] and $\tilde{O}(n^{1/3} d^{4/3})$ with a mixed Lee-Sidford and log-barrier [Gatmiry et al., 2024]. In this work, we show that log-concave sampling can also leverage Lee-Sidford barrier to obtain a mixing time of $\tilde{O}(d^2 + dL^2 R^2)$. In comparison, the hit-and-run algorithm [Lovász and Vempala, 2007] mixes in $\tilde{O}(d^2 R^2 / r^2)$ steps where r is the radius of the inscribed ball inside \mathcal{K} . For the regime where $L = O(1)$ and $r = O(1)$, our walk mixes strictly faster than that of hit-and-run.

We note that for sampling over more general convex bodies, a recent work [Chen and Eldan, 2022] proves that for uniform sampling over an isotropic convex body, the mixing time bound is $\tilde{O}(d^2 / \psi_d^2)$ where ψ_d is the KLS constant [Kannan et al., 1995]. However, it is unclear how to generalize their result to non-isotropic convex bodies or log-concave densities.

1.2 Our results

Our results concern log-concave sampling from polytopes and spectrahedra. For polytopes, we state the result in its full generality.

⁶Technically, what they have shown is that if the Hessian is a $\bar{\nu}$ -symmetric self-concordant barrier function, then the walk mixes in $O(\bar{\nu}d)$ steps, and they subsequently prove for various barriers of interest, $\bar{\nu} = \nu$.

Theorem 1.1 (Robust sampling for log-concave distribution over polytopes). *Let $\delta \in (0, 1)$ and $R \geq 1$. Given a constraint matrix $A \in \mathbb{R}^{n \times d}$ with a vector $b \in \mathbb{R}^n$, let $\mathcal{K} := \{x \in \mathbb{R}^d : Ax \leq b\}$ be the corresponding polytope. Suppose \mathcal{K} is enclosed in a ball of radius R with non-empty interior. Let $f : \mathcal{K} \rightarrow \mathbb{R}$ be an L -Lipschitz, convex function with an evaluation oracle. Finally, let π be the distribution such that $\pi \propto e^{-f}$.*

Suppose we are given an initial point from \mathcal{K} that is sampled from a w -warm start distribution⁷ with respect to π for some $w > 0$, then there is an algorithm (Algorithm 1) that outputs a point from a distribution μ where $\text{TV}(\mu, \pi) \leq \delta$.

Let $g : \mathcal{K} \rightarrow \mathbb{R}$ be a ν -self-concordant barrier function such that in time \mathcal{C}_g , a spectral approximation \tilde{H}_g of the Hessian of g denoted by H_g can be computed satisfying

$$(1 - \epsilon) \cdot H_g \preceq \tilde{H}_g \preceq (1 + \epsilon) \cdot H_g$$

for $\epsilon = \Theta(1/d)$. Then, Algorithm 1 takes at most

$$\tilde{O}((\nu d + dL^2 R^2) \log(w/\delta))$$

Markov chain steps. It uses $O(1)$ function evaluations and an extra $\mathcal{C}_g + d^\omega$ time at each step.

Let us pause and make some remarks regarding Theorem 1.1. As long as the Hessian matrix can be approximately generated with an error inversely depends on d , then our algorithm is guaranteed to converge. Moreover, if the approximate Hessian can be generated more efficiently, then this directly implies an improvement of our algorithm. On the mixing time side, Theorem 1.1 is nearly-optimal up to polylogarithmic factors and the dependence on L and R . In particular, the νd mixing time bound is also achieved by Laddha et al. [2020] in the case of uniform sampling. We instantiate the meta result of Theorem 1.1 into the following theorem.

Theorem 1.2 (Robust sampling with nearly-universal barrier). *Under the conditions of Theorem 1.1, let g be the Lee-Sidford barrier with $\nu = O(d \log^5 n)$. Then, we have*

$$\mathcal{C}_g = \tilde{O}(nd^{\omega-1}).$$

The algorithm takes at most

$$\tilde{O}((d^2 + dL^2 R^2) \log(w/\delta))$$

Markov chain steps.

The Lee-Sidford barrier [Lee and Sidford, 2014, 2019] is the first polynomial-time computable barrier with a nearly-optimal self-concordance parameter. Several prior works [Chen et al., 2018, Laddha et al., 2020] utilize this barrier for uniform sampling. For log-concave sampling, the walk of Kook and Vempala [2024] requires strong-convexity-like assumption on f in order to attain an $\tilde{O}(d^2)$ mixing time. Our work is the first to obtain an $\tilde{O}(d^2)$ mixing time for log-concave sampling over polytopes, when L and R are small.

References	Mixing time	Per iteration cost
Lovász and Vempala [2006]	$d^2 R^2 / r^2$	nd
Narayanan and Rakhlin [2017]	$d^5 + d^3 L^2 R^2$	$nd^{\omega-1}$
Mangoubi and Vishnoi [2023]	$nd + dL^2 R^2$	$nd^{\omega-1}$
Mangoubi and Vishnoi [2024]	$nd + dL^2 R^2$	$\text{nnz}(A) + d^2$
This work	$d^2 + dL^2 R^2$	$nd^{\omega-1}$

Table 1: Comparison of algorithms for sampling from a log-concave density e^{-f} over a d -dimensional polytope with n constraints, where f is L -Lipschitz. We omit $\tilde{O}(\cdot)$ to mixing time and per iteration cost. We assume the evaluation of f can be done in unit time. The first row is the hit-and-run walk. Our algorithm has the fastest mixing time among all Dikin walks, and for $L = O(1)$, $r = O(1)$, our walk mixes faster than hit-and-run.

⁷We say a distribution ρ is a w -warm start of with respect to a distribution π if $\sup_{x \in \mathcal{K}} \frac{\rho(x)}{\pi(x)} \leq w$.

Our next result is closely related to semidefinite programming. Given the set $\mathcal{K} = \{x \in \mathbb{R}^d : \sum_{i=1}^d x_i A_i \succeq C\}$ for symmetric $n \times n$ matrices A_1, \dots, A_d, C , we consider sampling from a log-concave distribution over \mathcal{K} . We use the popular log-barrier for \mathcal{K} studied in Nesterov and Nemirovskii [1994], Jiang et al. [2020a], Huang et al. [2022]. To further utilize our robust sampling framework, we develop a novel approach for generating a spectral approximation of the Hessian. Compared to Jiang et al. [2020a], Huang et al. [2022] in which the Hessian is maintained in a very careful way in conjunction with the interior point method, our spectral approximation does not require to control the changes over iterations and is highly efficient when $n \gg d$, which is a popular regime for semidefinite programs.

Theorem 1.3 (Robust sampling for log-concave distribution over spectrahedra). *Let $\delta \in (0, 1)$ and $R \geq 1$. Given a collection of symmetric matrices $A_1, \dots, A_d \in \mathbb{R}^{n \times n}$ and a target symmetric matrix $C \in \mathbb{R}^{n \times n}$, let $\mathcal{K} := \{x \in \mathbb{R}^d : \sum_{i=1}^d x_i A_i \succeq C\}$ be the corresponding spectrahedron. Suppose \mathcal{K} is enclosed in a ball of radius R with non-empty interior. Let $f : \mathcal{K} \rightarrow \mathbb{R}$ be an L -Lipschitz, convex function with an evaluation oracle. Finally, let π be the distribution such that $\pi \propto e^{-f}$.*

Suppose we are given an initial point from \mathcal{K} that is sampled from a w -warm start distribution with respect to π for some $w > 0$, then there is an algorithm (Algorithm 1) that outputs a point from a distribution μ where $\text{TV}(\mu, \pi) \leq \delta$.

Let $g : \mathcal{K} \rightarrow \mathbb{R}$ be the logarithmic barrier function with $\nu = n$. There is an algorithm that uses

$$\tilde{O}((nd + dL^2 R^2) \log(w/\delta))$$

Markov chain steps. In each step it uses $O(1)$ function evaluations and an extra cost of

$$\tilde{O}(n^\omega + n^2 d^{3\omega-5}).$$

Prior work due to Narayanan and Rakhlin [2017] provides a walk that mixes in $\tilde{O}(n^2 d^3 + n^2 dL^2 R^2)$ steps and each step could be implemented in $O(dn^\omega + n^2 d^{\omega-1})$ time, our algorithm improves upon both mixing and per iteration cost. On the front of per iteration cost, computing the Hessian of the log-barrier of \mathcal{K} exactly takes time $O(dn^\omega + n^2 d^{\omega-1})$, for the regime where $n \gg d$ the dominating term is dn^ω . Note that whenever $n \gg d^{\frac{3\omega-4}{\omega-2}}$, our algorithm is more efficient than that of exact computation. This is a common regime in semidefinite program in which the number of constraints d is small compared to the size of the primal solution.

References	Mixing time	Per iteration cost
Lovász and Vempala [2006]	$d^2 R^2 / r^2$	$n^\omega + n^2 d$
Narayanan and Rakhlin [2017]	$n^2 d^3 + n^2 dL^2 R^2$	$n^\omega d + n^2 d^{\omega-1}$
This work	$nd + dL^2 R^2$	$n^\omega + n^2 d^{3\omega-5}$

Table 2: Comparison of algorithms for sampling from a log-concave density e^{-f} over a d -dimensional spectrahedron with $n \times n$ symmetric matrix constraints, where f is L -Lipschitz. We omit $\tilde{O}(\cdot)$ to mixing time and per iteration cost. We assume evaluation of f can be done in unit time. The first row is the hit-and-run walk. Our algorithm mixes faster than Narayanan and Rakhlin [2017] in all parameter regimes, and for $L = O(1)$, $r = O(1)$, $d^2 R^2 \geq n$, our walk mixes faster than hit-and-run.

2 Technical Overview

Before diving into our techniques, we first illustrate an informal version of our algorithm. The algorithm itself is a variant of the Dikin walk [Kannan and Narayanan, 2012] called the soft-threshold Dikin walk [Mangoubi and Vishnoi, 2023]. Starting with a w -warm start point, the algorithm approximately computes the Hessian of the barrier function, then proposes a new point z from a Gaussian with mean at the current point x and variance $\tilde{\Phi}(x)^{-1}$, where $\tilde{\Phi}$ is a multiple of the approximate Hessian plus a multiple of the identity. The proposal is then accepted via a Metropolis filter. Note that compared to the standard Dikin walk framework, our algorithm crucially allows the Hessian to be approximated with a good precision. This means any improvements on generating spectral approximation leads to runtime improvement of our algorithm in a black-box fashion.

Algorithm 1 Sampler for log-concave distribution $\pi \propto \exp(-f)$ over convex body with a barrier function g . \mathcal{A} is the description of the convex set, function f is an L -Lipschitz concave function and g is a ν -self-concordant barrier function, $x_0 \in \mathcal{K}$ is a w -warm start point, and δ is the total variance distance between our Markov chain and the target distribution. Let $\mathcal{K} \subseteq \mathbb{R}^d$ enclosed by a ball of radius R .

```

1: procedure SAMPLINGCONVEXBODY( $\mathcal{A}, f, g, x_0, \delta$ )
2:   Initialize  $x$  to be a point in  $\mathcal{K}$ 
3:    $T \leftarrow \tilde{O}((\nu d + dL^2R^2) \log(w/\delta))$ 
4:   for  $t = 1 \rightarrow T$  do
5:     Sample a point  $\xi \sim \mathcal{N}(0, I_d)$ 
6:     //  $H(x)$  is the exact Hessian matrix at  $x$ , we never explicitly form it
7:     Compute  $\tilde{H}(x)$  where  $(1 - 1/d)H(x) \preceq \tilde{H}(x) \preceq (1 + 1/d)H(x)$ 
8:      $\tilde{\Phi}(x) = d \cdot \tilde{H}(x) + dL^2 \cdot I_d$ 
9:      $z \leftarrow x + \tilde{\Phi}(x)^{-1/2} \xi$ 
10:    if  $z \in \text{Int}(\mathcal{K})$  then
11:      //  $H(z)$  is the exact Hessian matrix at  $z$ 
12:      Compute  $\hat{H}(z)$  such that  $(1 - 1/d)H(z) \preceq \hat{H}(z) \preceq (1 + 1/d)H(z)$ 
13:       $\hat{\Phi}(z) = d \cdot \hat{H}(z) + dL^2 \cdot I_d$ 
14:       $\tau \leftarrow \frac{\exp(-f(z)) \cdot (\det(\hat{\Phi}(z)))^{1/2} \cdot \exp(-0.5\|x-z\|_{\hat{\Phi}(z)}^2)}{\exp(-f(x)) \cdot (\det(\hat{\Phi}(x)))^{1/2} \cdot \exp(-0.5\|x-z\|_{\hat{\Phi}(x)}^2)}$ 
15:      // Our walk is lazy, i.e., it only moves to a new point with half probability
16:      accept  $x \leftarrow z$  with probability  $0.5 \cdot \min\{\tau, 1\}$ 
17:    else
18:      reject  $z$ 
19:    end if
20:  end for
21: end procedure

```

2.1 Soft-threshold Dikin walk for log-concave sampling via a barrier-based argument

For uniform sampling from structured convex bodies such as polytopes, Kannan and Narayanan [2012] introduce a refined ball walk that utilizes a *self-concordant barrier function* associated with the convex body. To extend the Dikin walk framework for log-concave sampling, Narayanan and Rakhlin [2017] adds an additional coefficient $\frac{\exp(-f(x))}{\exp(-f(z))}$ to the acceptance probability of Metropolis filter. This achieves a sub-optimal mixing time bound of $\tilde{O}(\nu^2 d)$. By introducing a soft-threshold regularizer to the Hessian, Mangoubi and Vishnoi [2023] proves that for log-barrier function, it is possible to approach the optimality with an $\tilde{O}(nd + dL^2R^2)$ mixing time. Unfortunately, the method used by Mangoubi and Vishnoi [2023] is specialized to log-barrier, making it hard to generalize to other barriers. They show that under proper scaling, the true Hessian $H(x)$ under a sequence of polytopes converges to the regularized Hessian under the given polytope in the limit. They define the j -th polytope as

$$A_j = \begin{bmatrix} A \\ I_d \\ \vdots \\ I_d \end{bmatrix}, b_j = \begin{bmatrix} b \\ j \cdot \mathbf{1}_d \\ \vdots \\ j \cdot \mathbf{1}_d \end{bmatrix}$$

the number of copies of I_d and $j \cdot \mathbf{1}_d$ is a function of j . While this polytope-based method works well for log-barrier, it functions poorly for more intricate barriers, such as volumetric barrier and Lee-Sidford barrier. Take volumetric barrier as an example, recall that $H_{\text{vol}}(x) = A^\top \Sigma(S(x)^{-1}A)S(x)^{-2}A$ where $\Sigma(S(x)^{-1}A)$ is the statistical leverage score [Drineas et al., 2006, Spielman and Srivastava, 2011]. Leverage score is a numerical quantity that measures *how important a row of a matrix is, compared to other rows*. This means that, if one duplicates a row infinitely many times, the row will be assigned a score of 0. More concretely, suppose we are given an $n \times d$ matrix with n identical rows, then the leverage score of each row is $\frac{d}{n}$. Taking $n \rightarrow \infty$, it's easy to see that all rows have score 0, and the Hessian of volumetric barrier will zero out the rows contributed to the

infinitely many copies of identities. Thus, for volumetric barrier, extra constraints of the polytope sequence vanish. On the other hand, the [Mangoubi and Vishnoi \[2023\]](#) construction heavily relies on the infinitely many occurrences of I_d for it to converge to the regularized Hessian.

To circumvent this issue and provide a simpler argument, we develop a proof strategy that is *barrier-based*, instead of polytope-based. We example a class of popular barrier functions for polytopes that can be classified as *weighted log-barrier functions* [[Vaidya, 1989](#)]. Roughly speaking, these functions take in the form of $g(x) = -\sum_{i=1}^n w_i \log(S(x)_i)$, with $w = \mathbf{1}_n$ being the log-barrier, $w = \sigma(S(x)^{-1}A)$ being the volumetric barrier [[Vaidya, 1989](#)] and if we choose w to be a proper power of the Lewis weights, then the barrier is precisely the Lee-Sidford barrier [[Lee and Sidford, 2019](#), [Laddha et al., 2020](#)]. These weights are stable, meaning that if two points are close in certain local norm induced by proper ellipses given by these weights, then their weights are also close [[Lee and Sidford, 2019](#)]. This property proves to be particularly useful when proving the mixing rate of our Dikin walk. In addition, we carve out three sufficient conditions for a weighed log-barrier barrier function to work for log-concave distribution, and has $\tilde{O}(\nu d)$ mixing rate:

- **$\bar{\nu}$ -symmetry.** The Hessian $H(x)$ is $\bar{\nu}$ -symmetry, that is, for any $x \in \mathcal{K}$, $E_x(1) \subseteq \mathcal{K} \cap (2x - \mathcal{K}) \subseteq E_x(\sqrt{\bar{\nu}})$ where $E_x(r) = \{y \in \mathbb{R}^d : (y - x)^\top H(x)(y - x) \leq r^2\}$. We show that $\bar{\nu} = \nu$ for log-barrier, volumetric barrier and Lee-Sidford barrier for polytopes, and log-barrier for spectrahedra.
- **Bounded local norm.** The variance term $\|H(x)^{-1/2} \cdot \nabla \log \det(H(x))\|_2^2 \leq d \text{poly} \log n$. We show this condition holds for all barriers of interest in this paper.
- **Convexity of regularized barrier.** The function $F(x) = \log \det(H(x) + I_d)$ is convex at x for any $x \in \mathcal{K}$. We note that for log-barrier, volumetric barrier and Lee-Sidford barrier, it has been shown that $\log \det(H(x))$ is indeed convex at x . We further prove that this still holds when a copy of identity is blended in.

We want to highlight that our approach is more robust and generic than that of [Mangoubi and Vishnoi \[2023\]](#), as it solely depends on the structure of the Hessian matrix. Our argument should be treated as a generalization of [Sachdeva and Vishnoi \[2016\]](#), in which they utilize the bounded local norm and convexity of $\log \det(H(x))$ for log-barrier, together with the Gaussianity to conclude that the difference between $\log \det$ are small. We go beyond log-barrier for polytopes and uniform sampling.

2.2 Approximation preserves the mixing time

The core premise of our algorithm is we allow the Hessian to be approximated, so that each step can be implemented efficiently. This poses significant challenges in proving mixing, as even if we have a good approximation of one-step in the Markov chain, the approximate chain does not necessarily mixes as fast as the original chain. For example, if every step is approximated within ϵ -TV distance to the original Markov chain, then in T steps, the TV distance between the resulting distribution under approximate walk and the distribution under original walk could be as large as $T\epsilon$. This means we have to take ϵ very small ($O(T_{\text{mix}}^{-1})$) to guarantee same convergence property as original chain, which is unacceptable. A related issue is that the stationary distribution under approximate walk may not be the same as the target distribution. Therefore we need a way to control the mixing properties of the approximate walk. We resolve this problem by exactly computing the acceptance probability under approximate walk. That is, after proposing the next step z from x , $\Phi(x)$, we sample $\hat{\Phi}(z)$, and accept with probability $\min\{1, \frac{\pi(z)\tilde{G}_z(x)}{\pi(x)\tilde{G}_x(z)}\}$ (where $\tilde{G}_x(z)$ is the probability of proposing z starting from x under the approximate walk). Recall that in the exact walk, we accept with probability $\min\{1, \frac{\pi(z)G_z(x)}{\pi(x)G_x(z)}\}$ (where $G_x(z)$ is the probability of proposing z starting from x under the exact walk). If we use this acceptability for the approximate walk, then we only have one-step approximation guarantee, and will suffer from the problem mentioned previously.

By using this modified acceptance probability, we can prove reversibility of the approximate walk, and that stationary distribution of the approximate walk is indeed our target distribution. For the mixing rate, we bound the conductance of the approximate walk, which requires us to prove the following properties:

- (1) The proposed step is accepted with decent probability.

- (2) Starting from two points moderately close to each other (i.e., $\|x - z\|_{\Phi(x)} \leq \epsilon$), the approximate steps taken P_x and P_z are close in TV distance (i.e., $\text{TV}(P_x, P_z) \leq 1 - \epsilon'$) for some absolute constants $\epsilon, \epsilon' > 0$.

We are able to prove these properties by comparison with the original chain. That is, assuming the original chain satisfies properties (1)(2), then the approximate chain also satisfies the same properties, as long as the approximation is good enough. This enables us to prove these key properties in a hassle-free way. To prove these comparison results, we need to prove, for example, $\tilde{G}_x(z) \approx G_x(z)$. Because both $G_x(z)$ and $\tilde{G}_x(z)$ have nice factorizations

$$G_x(z) \propto \det(\Phi(x))^{1/2} \exp\left(-\frac{1}{2}\|x - z\|_{\Phi(x)}^2\right),$$

$$\tilde{G}_x(z) \propto \det(\tilde{\Phi}(x))^{1/2} \exp\left(-\frac{1}{2}\|x - z\|_{\tilde{\Phi}(x)}^2\right),$$

we only need to prove $\det(\Phi(x)) \approx \det(\tilde{\Phi}(x))$ and $\exp\left(-\frac{1}{2}\|x - z\|_{\Phi(x)}^2\right) \approx \exp\left(-\frac{1}{2}\|x - z\|_{\tilde{\Phi}(x)}^2\right)$ separately. Both of these can be handled via $\Phi(x) \approx \tilde{\Phi}(x)$, using our approximation procedure.

2.3 Approximation of Lewis weights for rapid mixing

Given our robust and generic Dikin walk framework, we realize it with the Lee-Sidford barrier whose complexity $\nu = d \log^5 n$. This ensures the walk mixes in $\tilde{O}(d^2 + dL^2R^2)$ steps, improving upon the log-barrier-based walk of Mangoubi and Vishnoi [2023, 2024] that mixes in $\tilde{O}(nd + dL^2R^2)$ steps, as the log-barrier has complexity $\nu = n$. To implement the Lee-Sidford barrier, it is imperative to compute the ℓ_p (for $p > 0$) Lewis weights, which is defined as the following convex program:

$$\min_{M \succeq 0} -\log \det M, \text{ subject to } \sum_{i=1}^n (a_i^\top M a_i)^{p/2} \leq d, \quad (1)$$

and the ℓ_p Lewis weights is a vector $w_p \in \mathbb{R}^d$ where $(w_p)_i = a_i^\top M_* a_i$ and M_* is the optimum of Program (1). Solving the program exactly is not efficient, and a long line of works [Cohen and Peng, 2015, Lee and Sidford, 2014, 2019, Cohen et al., 2019, Fazel et al., 2022, Jambulapati et al., 2022] provide fast algorithms that approximate all weights up to $(1 \pm \epsilon)$ -factor. The Hessian of Lee-Sidford barrier has the form (up to scaling) $H_{\text{Lewis}}(x) = A^\top S(x)^{-1} W_p (S(x)^{-1} A)^{1-2/p} S(x)^{-1} A$ where $W_p(S(x)^{-1} A)$ is the diagonal matrix of ℓ_p Lewis weights with respect to $S(x)^{-1} A$. When clear from context, we use W_p as a shorthand for $W_p(S(x)^{-1} A)$. As W_p could be approximated in $\tilde{O}(nd^{\omega-1})$ time, we could then perform subsequent operations exactly to form the approximate Hessian. We instead provide a spectral approximation procedure that runs in $\tilde{O}(\text{nnz}(A) + \mathcal{T}_{\text{mat}}(d, d^3, d))$ given the approximate Lewis weights where $\mathcal{T}_{\text{mat}}(a, b, c)$ is the time of multiplying an $a \times b$ matrix with a $b \times c$ matrix and $\mathcal{T}_{\text{mat}}(d, d^3, d) \approx d^{4.2}$. This is important as even though it doesn't lead to direct runtime improvement, if approximate Lewis weights can be computed in $\tilde{O}(\text{nnz}(A) + \text{poly}(d))$ time, then we obtain a per iteration cost upgrade in a black-box fashion. Our approach is based on spectral sparsification via random sampling: given a matrix $B \in \mathbb{R}^{n \times d}$, one can construct a matrix $\tilde{B} \in \mathbb{R}^{s \times d}$ consists of re-scaled rows of B , such that $(1 - \epsilon) \cdot B^\top B \preceq \tilde{B}^\top \tilde{B} \preceq (1 + \epsilon) \cdot B^\top B$. If we sample according to the leverage scores [Drineas et al., 2006, Spielman and Srivastava, 2011], one can set $s = O(\epsilon^{-2} d \log d)$. In our case, we set $B = W_p^{1/2-1/p} S(x)^{-1} A$, and perform the sampling. However, computing leverage score exactly is as expensive as forming the Hessian exactly, which requires $O(nd^{\omega-1})$ time. To speed up the crucial step, we use randomized sketching technique to reduce the row count from n to $\text{poly}(d)$, then approximately estimate all leverage scores with this small matrix. One can either use the sparse embedding matrix [Nelson and Nguyen, 2013] to approximate these scores in time $O(\text{nnz}(A) \log n + \mathcal{T}_{\text{mat}}(d, \epsilon^{-2} d, d))$, or use the subsampled randomized Hadamard transform [Lu et al., 2013] in time $O(nd \log n + \mathcal{T}_{\text{mat}}(d, \epsilon^{-2} d, d))$. We remark that approximating leverage scores for subsampling has many applications in numerical linear algebra, such as low rank approximation [Boutsidis et al., 2016, Song et al., 2017, 2019]. Utilizing this framework, we provide an algorithm that approximates the Hessian of Lee-Sidford barrier in time $\tilde{O}(nd^{\omega-1}) + \tilde{O}(\text{nnz}(A) + \mathcal{T}_{\text{mat}}(d, d^3, d))$, as advertised.

2.4 Approximation of log-barrier Hessian over a spectrahedron

For a spectrahedron induced by the dual semidefinite program, it is also natural to define its log-barrier with its corresponding Hessian being

$$H_{\log}(x) = A(S(x)^{-1} \otimes S(x)^{-1})A^\top,$$

where $A \in \mathbb{R}^{d \times n^2}$ with the i -th row being $\text{vec}(A_i)$ for $n \times n$ symmetric matrix A_i , and $S(x) = \sum_{i=1}^d x_i A_i - C$ where C is also an $n \times n$ symmetric matrix. Nesterov and Nemirovskii [1994] shows that by smartly arranging the organization of $H_{\log}(x)$, it can be computed exactly in time $O(dn^\omega + d^{\omega-1}n^2)$. Beyond exact computation, Jiang et al. [2020a], Huang et al. [2022] present algorithms that approximately *maintain* the Hessian matrix under low-rank updates. These maintenance approaches are crafted towards solving semidefinite programs in which the trajectory can be carefully controlled. For our soft-threshold Dikin walk, though the proposal generates a point that is relatively close to the starting point, due to its Gaussianity nature, much less structure can be exploited and thus maintained. Instead, we propose a maintenance-free approach that uses randomized sketching to generate a spectral approximation of $H_{\log}(x)$.

To illustrate the algorithm, let's first define a matrix

$$B = A(S(x)^{-1/2} \otimes S(x)^{-1/2}).$$

It is not hard to see that $H_{\log}(x) = BB^\top$. Due to the Kronecker product structure of $(S(x)^{-1/2} \otimes S(x)^{-1/2})$, it is natural to consider sketches for Kronecker product of matrices [Diao et al., 2018, 2019, Ahle et al., 2020, Song et al., 2021]. Following the standard procedure for using sketch to generate spectral approximation, we can choose a TensorSRHT matrix [Ahle et al., 2020, Song et al., 2021] T and compute $R := (S(x)^{-1/2} \otimes S(x)^{-1/2})T$ then form AR . This is unfortunately, not efficient enough, as multiplying A with R might be too slow. To further optimize the runtime efficiency, we use a more intrinsic approach based on the matrix T and B . It is well-known that the i -th row of B can be first computed as $S(x)^{-1/2} A_i S(x)^{-1/2}$ then vectorize. If we can manage to apply the sketch in a row-by-row fashion, then fast matrix multiplication can be utilized for even faster algorithms. We recall that $T = P \cdot (HD_1 \otimes HD_2)$ where H is the Hadamard matrix and P is a row sampling matrix. To compute a row, we can first apply HD_1 and HD_2 individually to $S(x)^{-1/2}$ in $\tilde{O}(n^2)$ time since H is a Hadamard matrix. We can then form two matrices X and Y with rows being the corresponding sampled rows from HD_1 and HD_2 . Finally, we compute $X A_i Y^\top$ to form one row. Using fast matrix multiplication, we can form each row in $\mathcal{T}_{\text{mat}}(d^3, n, n) = O(n^2 d^{3(\omega-2)})$ time. Applying this procedure to d rows leads to an overall $O(n^2 d^{3\omega-5})$ time, which beats the $O(dn^\omega)$ time as long as $n \gg d$. Our runtime is somewhat slow when n is not that larger than d , this is due to we have to set $\epsilon = O(1/d)$. For more popular regimes where it suffices to set $\epsilon = O(1)$, our algorithm outperforms exact computation as long as $n \geq d$.

3 Applications

There are many applications of our Dikin walk, such as differentially private learning, simulated annealing and regret minimization. For a more comprehensive overview of the applications, we refer readers to Narayanan and Rakhlin [2017].

Differentially private learning. Let $\mathcal{X}^m = \{x_1, \dots, x_m\}$ denote an m -point dataset, associate a convex loss function ℓ_i to point x_i , the learning problem attempts to find an optimal parameter $\theta_* \in \mathcal{K}$ that minimizes $\ell(\theta) = \sum_{i=1}^m \ell_i(\theta)$. We could further enforce *privacy* constraints to learning: we say a randomized algorithm $\mathcal{M} : \mathcal{X}^m \rightarrow \mathbb{R}$ is ϵ -DP if for any datasets $X, X' \in \mathcal{X}^m$ differ by a single point and for any $S \subseteq \mathbb{R}$, $\Pr[\mathcal{M}(X) \subseteq S] \leq e^\epsilon \Pr[\mathcal{M}(X') \subseteq S]$. The differentially private learning seeks to solve the learning problem under DP guarantees, and it has been shown that if one allows for *approximate DP*, then it can be achieved via the exponential mechanism and sampling from the distribution $\pi(\theta) \propto e^{-\ell(\theta)}$ [Wang et al., 2015]. If \mathcal{K} is a polytope or spectrahedron, and in addition ℓ is L -Lipschitz, one can implement our walk to obtain an $\tilde{O}(d^2 + dL^2 R^2)$ mixing for polytopes or $\tilde{O}(nd + dL^2 R^2)$ mixing for spectrahedra. Via a standard reduction from TV distance bound to infinity distance bound [Mangoubi and Vishnoi, 2022], our walk can also be adapted for a *pure DP* guarantee [Lin et al., 2024].

Simulated annealing for convex optimization. Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be an L -Lipschitz and convex function, we consider the problem of minimizing f over \mathcal{K} when we can only access the function value of f . This is also referred to as zeroth-order convex optimization. One could solve the problem via the *simulated annealing framework*, in which one needs to sample from the distribution $\pi(x) \propto e^{-f(x)/T}$ where T is the temperature parameter [Kalai and Vempala, 2006]. Our walk can be adapted to the framework for polytopes, it mixes in $\tilde{O}(d^{2.5} + d^{1.5}L^2R^2)$ steps and for spectrahedra, it mixes in $\tilde{O}(nd^{1.5} + d^{1.5}L^2R^2)$ steps, improving upon prior best algorithms that mix in $\min\{d^{5.5} + d^{3.5}L^2R^2, nd^{1.5} + d^{1.5}L^2R^2\}$ steps for polytopes and $n^2d^{3.5} + n^2d^{1.5}L^2R^2$ steps for spectrahedra [Narayanan and Rakhlin, 2017, Mangoubi and Vishnoi, 2023].

Online convex optimization. Consider the following online optimization problem: let \mathcal{K} be a convex set that we could choose our actions from and let ℓ_1, \dots, ℓ_T be a sequence of unknown convex cost functions with $\ell_t : \mathcal{K} \rightarrow \mathbb{R}$. The goal is to design a good strategy that chooses from \mathcal{K} so that the total cost is small, compared to the offline optimal cost, in which one could choose the strategy after seeing the entire sequence of ℓ_t 's. The algorithm works as follows: at round t , we choose a distribution μ_{t-1} supported on \mathcal{K} and play the action $Y_t \sim \mu_{t-1}$. Our goal is to minimize the expected regret, defined as $\text{Reg}_T(U) = \mathbb{E}[\sum_{t=1}^T \ell_t(Y_t) - \sum_{t=1}^T \ell_t(U)]$ with respect to all randomized strategies defined by distribution p_U . It turns out if one sets $s_t(x) = \eta \sum_{s=1}^t \ell_s(x)$ where $\eta > 0$ is a step size parameter and sets $\mu_t \propto e^{-s_t}$, then this update rule reflects the multiplicative weights update algorithm [Arora et al., 2012]. Moreover, one could prove a nearly-optimal regret bound of $O(\sqrt{T})$ if the KL divergence between μ_0 and all possible p_U 's are bounded. The algorithmic problem of sampling from μ_t is equivalent to the log-concave sampling problem we study in this paper, and we could efficiently generate strategies with good regret, similar to Narayanan and Rakhlin [2017]. Let ℓ_t 's be 1-Lipschitz linear functions, Narayanan and Rakhlin [2017] uses their time-varying Dikin walk to achieve a regret of $O(d^{2.5}\sqrt{T})$, while we could significantly improve this bound to $O(d\sqrt{T})$.

4 Conclusion

In this paper, we design a class of error-robust Dikin walks for sampling from a log-concave and log-Lipschitz density over a convex body. The key features of our walks are that their mixing time depends *linearly* on the complexity of self-concordant barrier function of the convex body, and they allow computationally intensive quantities to be *approximated* rather than computed exactly. For polytopes, our walk mixes in $\tilde{O}(d^2 + dL^2R^2)$ steps with a per iteration cost of $\tilde{O}(nd^{\omega-1})$, and for spectrahedra, our walk mixes in $\tilde{O}(nd + dL^2R^2)$ steps with a per iteration cost of $\tilde{O}(n^\omega + n^2d^{3\omega-5})$. For polytopes, our walk is the first successful adaptation of the Lee-Sidford barrier for log-concave sampling under the minimal assumption that f is Lipschitz, improving upon the mixing of prior works [Narayanan and Rakhlin, 2017, Mangoubi and Vishnoi, 2023, 2024]. For spectrahedra, we improve the mixing of Narayanan and Rakhlin [2017] from $n^2d^3 + n^2dL^2R^2$ to $nd + dL^2R^2$, for the term that depends on only n and d , we obtain an improvement of nd^2 , and for the term that depends on L and R , we shave off the quadratic dependence on n . Moreover, we adapt our error-robust framework and present a sketching algorithm for approximating the Hessian matrix in $\tilde{O}(n^\omega + n^2d^{3\omega-5})$ time. Our results have deep implications, as it could be leveraged for differentially private learning, convex optimization and regret minimization.

While our framework offers the fastest mixing Dikin walk for log-concave sampling, its mixing time has a rather unsatisfactory dependence on the radius of the bounding box, R . It would be interesting if one could design a walk that *does not* depend on R . We also note that we require the density of be log-Lipschitz rather than Lipschitz. Kook and Vempala [2024] shows that if f in addition satisfies the *relative strongly-convex property*, then there exists a walk whose mixing does not depend on R and one only needs the density of be Lipschitz. Another important direction is to further improve the mixing rate of sampling over a spectrahedron using barrier functions such as volumetric barrier and hybrid barrier [Anstreicher, 2000], while maintaining a small per iteration cost. Finally, extending RHMC to log-concave sampling would be essentially, as it has the potential to give the fastest mixing walk by utilizing Riemannian metrics instead of Euclidean. As our work is theoretical in nature, we don't foresee any potential negative societal impact. It has the potential positive impact to reduce energy consumption and carbon emission when deployed in practice.

Acknowledgement

We would like to thank Jonathan Kelner, Yin Tat Lee, Santosh Vempala and Yunbum Kook for valuable discussions and anonymous reviewers for helpful comments. The research is partially supported by the NSF awards: SCALE MoDL-2134209, CCF-2112665 (TILOS). It is also supported by the U.S. Department of Energy, the Office of Science, the Facebook Research Award, as well as CDC-RFA-FT-23-0069 from the CDC's Center for Forecasting and Outbreak Analytics. Lichen Zhang is supported by NSF awards CCF-1955217 and DMS-2022448. For more information related to the paper and adjacent topics, see <https://www.youtube.com/@zhaosong2031> and <https://space.bilibili.com/3546587376650961>.

References

- Thomas D Ahle, Michael Kapralov, Jakob BT Knudsen, Rasmus Pagh, Ameya Velingker, David P Woodruff, and Amir Zandieh. Oblivious sketching of high-degree polynomial kernels. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 141–160, 2020.
- Josh Alman and Virginia Vassilevska Williams. A refined laser method and faster matrix multiplication. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 522–539. SIAM, 2021.
- Kurt M Anstreicher. The volumetric barrier for semidefinite programming. *Mathematics of Operations Research*, 25(3):365–380, 2000.
- Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.
- Christos Boutsidis, David P Woodruff, and Peilin Zhong. Optimal principal component analysis in distributed and streaming models. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing (STOC)*, pages 236–249, 2016.
- Nicolas Brosse, Alain Durmus, and Eric Moulines. The promises and pitfalls of stochastic gradient langevin dynamics. *Advances in Neural Information Processing Systems*, 31, 2018.
- Niladri Chatterji, Nicolas Flammarion, Yian Ma, Peter Bartlett, and Michael Jordan. On the theory of variance reduction for stochastic gradient monte carlo. In *International Conference on Machine Learning*, pages 764–773. PMLR, 2018.
- Yuansi Chen and Ronen Eldan. Hit-and-run mixing via localization schemes, 2022.
- Yuansi Chen, Raaz Dwivedi, Martin J. Wainwright, and Bin Yu. Fast mcmc sampling algorithms on polytopes. *J. Mach. Learn. Res.*, 2018.
- Michael B. Cohen and Richard Peng. Lp row sampling by lewis weights. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, STOC '15, 2015.
- Michael B Cohen, Ben Cousins, Yin Tat Lee, and Xin Yang. A near-optimal algorithm for approximating the john ellipsoid. In *Conference on Learning Theory*, pages 849–873. PMLR, 2019.
- Arnak S Dalalyan and Avetik Karagulyan. User-friendly guarantees for the langevin monte carlo with inaccurate gradient. *Stochastic Processes and their Applications*, 129(12):5278–5311, 2019.
- Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. *Advances in neural information processing systems*, 27, 2014.
- Huaian Diao, Zhao Song, Wen Sun, and David Woodruff. Sketching for kronecker product regression and p-splines. In *International Conference on Artificial Intelligence and Statistics*, pages 1299–1308. PMLR, 2018.

- Huaian Diao, Rajesh Jayaram, Zhao Song, Wen Sun, and David Woodruff. Optimal sketching for kronecker product regression and low rank approximation. *Advances in neural information processing systems*, 32:4737–4748, 2019.
- Zhiyan Ding and Qin Li. Langevin monte carlo: random coordinate descent and variance reduction. *J. Mach. Learn. Res.*, 22:205–1, 2021.
- Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. Sampling algorithms for l2 regression and applications. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithm*, SODA '06, page 1127–1136, USA, 2006. Society for Industrial and Applied Mathematics.
- Ran Duan, Hongxun Wu, and Renfei Zhou. Faster matrix multiplication via asymmetric hashing. In *FOCS*, 2023.
- Kumar Avinava Dubey, Sashank J Reddi, Sinead A Williamson, Barnabas Poczos, Alexander J Smola, and Eric P Xing. Variance reduction in stochastic gradient langevin dynamics. *Advances in neural information processing systems*, 29, 2016.
- Maryam Fazel, Yin Tat Lee, Swati Padmanabhan, and Aaron Sidford. Computing lewis weights to high precision. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2723–2742. SIAM, 2022.
- Z Furedi and I Barany. Computing the volume is difficult. In *Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing*, STOC '86, page 442–447, New York, NY, USA, 1986. Association for Computing Machinery.
- François Le Gall and Florent Urrutia. Improved rectangular matrix multiplication using powers of the coppersmith-winograd tensor. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '18, page 1029–1046, 2018.
- Francois Le Gall. Faster rectangular matrix multiplication by combination loss analysis. In *Proceedings of the Thirty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA'24, 2024.
- Khashayar Gatzmiry, Jonathan Kelner, and Santosh S. Vempala. Sampling polytopes with riemannian hmc: Faster mixing via the lewis weights barrier. In Shipra Agrawal and Aaron Roth, editors, *Proceedings of Thirty Seventh Conference on Learning Theory*, volume 247 of *Proceedings of Machine Learning Research*, pages 1796–1881. PMLR, 30 Jun–03 Jul 2024.
- Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, 1990.
- Baihe Huang, Shunhua Jiang, Zhao Song, Runzhou Tao, and Ruizhe Zhang. Solving sdp faster: A robust ipm framework and efficient implementation. In *FOCS*, 2022.
- Arun Jambulapati, Yang P. Liu, and Aaron Sidford. Improved iteration complexities for overconstrained p-norm regression. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2022, 2022.
- Haotian Jiang, Tarun Kathuria, Yin Tat Lee, Swati Padmanabhan, and Zhao Song. A faster interior point method for semidefinite programming. In *2020 IEEE 61st annual symposium on foundations of computer science (FOCS)*, pages 910–918. IEEE, 2020a.
- 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*, pages 944–953, 2020b.
- Shunhua Jiang, Zhao Song, Omri Weinstein, and Hengjie Zhang. A faster algorithm for solving general lps. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 823–832, 2021.
- Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. *Advances in neural information processing systems*, 26, 2013.

- William B Johnson and Joram Lindenstrauss. Extensions of lipschitz mappings into a hilbert space. *Contemporary mathematics*, 26(189-206):1, 1984.
- Adam Tauman Kalai and Santosh Vempala. Simulated annealing for convex optimization. *Mathematics of Operations Research*, 31(2):253–266, 2006.
- R. Kannan, L. Lovász, and M. Simonovits. Isoperimetric problems for convex bodies and a localization lemma. *Discrete Comput. Geom.*, 1995.
- Ravindran Kannan and Hariharan Narayanan. Random walks on polytopes and an affine interior point method for linear programming. *Mathematics of Operations Research*, 37(1):1–20, 2012.
- Yunbum Kook and Santosh S. Vempala. Gaussian cooling and Dikin walks: The interior-point method for logconcave sampling. In Shipra Agrawal and Aaron Roth, editors, *Proceedings of Thirty Seventh Conference on Learning Theory*, volume 247 of *Proceedings of Machine Learning Research*, pages 3137–3240. PMLR, 30 Jun–03 Jul 2024.
- Aditi Laddha, Yin Tat Lee, and Santosh Vempala. Strong self-concordance and sampling. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 1212–1222, 2020.
- Beatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, pages 1302–1338, 2000.
- Yin Tat Lee and Aaron Sidford. Path finding methods for linear programming: Solving linear programs in $\tilde{O}(\sqrt{\text{rank}})$ iterations and faster algorithms for maximum flow. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 424–433, 2014.
- Yin Tat Lee and Aaron Sidford. Solving linear programs with $\sqrt{\text{rank}}$ linear system solves. *arXiv preprint arXiv:1910.08033*, 2019.
- Yin Tat Lee and Santosh S. Vempala. Geodesic walks in polytopes. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017*, 2017.
- Yin Tat Lee and Santosh S Vempala. Convergence rate of riemannian hamiltonian monte carlo and faster polytope volume computation. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1115–1121, 2018.
- Yin Tat Lee and Man-Chung Yue. Universal barrier is n -self-concordant. *Mathematics of Operations Research*, 46(3):1129–1148, 2021.
- Yin Tat Lee, Aaron Sidford, and Sam Chiu-wai Wong. A faster cutting plane method and its implications for combinatorial and convex optimization. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 1049–1065. IEEE, 2015.
- Yin Tat Lee, Zhao Song, and Qiuyi Zhang. Solving empirical risk minimization in the current matrix multiplication time. In *Conference on Learning Theory*, pages 2140–2157. PMLR, 2019.
- Zhize Li, Tianyi Zhang, Shuyu Cheng, Jun Zhu, and Jian Li. Stochastic gradient hamiltonian monte carlo with variance reduction for bayesian inference. *Machine Learning*, 108(8):1701–1727, 2019.
- Yingyu Lin, Yian Ma, Yu-Xiang Wang, Rachel Emily Redberg, and Zhiqi Bu. Tractable MCMC for private learning with pure and gaussian differential privacy. In *The Twelfth International Conference on Learning Representations*, 2024.
- László Lovász and Miklós Simonovits. Random walks in a convex body and an improved volume algorithm. *Random structures & algorithms*, 4(4):359–412, 1993.
- László Lovász and Santosh Vempala. Hit-and-run is fast and fun. *preprint, Microsoft Research*, 2003.
- László Lovász and Santosh Vempala. Simulated annealing in convex bodies and an $O^*(n^4)$ volume algorithm. *Journal of Computer and System Sciences*, 72(2):392–417, 2006.

- László Lovász and Santosh Vempala. The geometry of logconcave functions and sampling algorithms. *Random Structures & Algorithms*, 30(3):307–358, 2007.
- Yichao Lu, Paramveer Dhillon, Dean P Foster, and Lyle Ungar. Faster ridge regression via the subsampled randomized hadamard transform. In *Advances in neural information processing systems*, pages 369–377, 2013.
- Oren Mangoubi and Nisheeth Vishnoi. Sampling from log-concave distributions with infinity-distance guarantees. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 12633–12646. Curran Associates, Inc., 2022.
- Oren Mangoubi and Nisheeth K Vishnoi. Sampling from structured log-concave distributions via a soft-threshold dikin walk. In *Advances in Neural Information Processing Systems*, NeurIPS’23, 2023.
- Oren Mangoubi and Nisheeth K. Vishnoi. Faster sampling from log-concave densities over polytopes via efficient linear solvers. In *The Twelfth International Conference on Learning Representations*, 2024.
- Hariharan Narayanan. Randomized interior point methods for sampling and optimization. *Ann. Appl. Probab.*, 2016.
- Hariharan Narayanan and Alexander Rakhlin. Efficient sampling from time-varying log-concave distributions. *The Journal of Machine Learning Research*, 18(1):4017–4045, 2017.
- Jelani Nelson and Huy L Nguyễn. OSNAP: Faster numerical linear algebra algorithms via sparser subspace embeddings. In *54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 117–126. IEEE, 2013.
- Yurii Nesterov and Arkadii Nemirovskii. Self-concordant functions and polynomial-time methods in convex programming. *USSR Academy of Sciences*, 1989.
- Yurii Nesterov and Arkadii Nemirovskii. *Interior-point polynomial algorithms in convex programming*. SIAM, 1994.
- Lianke Qin, Zhao Song, Lichen Zhang, and Danyang Zhuo. An online and unified algorithm for projection matrix vector multiplication with application to empirical risk minimization. In *AISTATS*, 2023.
- Maxim Raginsky, Alexander Rakhlin, and Matus Telgarsky. Non-convex learning via stochastic gradient langevin dynamics: a nonasymptotic analysis. In *Conference on Learning Theory*, pages 1674–1703. PMLR, 2017.
- James Renegar. A polynomial-time algorithm, based on newton’s method, for linear programming. *Math. Program.*, 1988.
- Sushant Sachdeva and Nisheeth K Vishnoi. The mixing time of the dikin walk in a polytope—a simple proof. *Operations Research Letters*, 44(5):630–634, 2016.
- Shai Shalev-Shwartz and Tong Zhang. Stochastic dual coordinate ascent methods for regularized loss minimization. *Journal of Machine Learning Research*, 14(1), 2013.
- Zhao Song and Zheng Yu. Oblivious sketching-based central path method for solving linear programming problems. In *38th International Conference on Machine Learning (ICML)*, 2021.
- Zhao Song, David P Woodruff, and Peilin Zhong. Low rank approximation with entrywise ℓ_1 -norm error. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 688–701, 2017.
- Zhao Song, David P Woodruff, and Peilin Zhong. Relative error tensor low rank approximation. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2772–2789. SIAM, 2019.

- Zhao Song, David Woodruff, Zheng Yu, and Lichen Zhang. Fast sketching of polynomial kernels of polynomial degree. In *International Conference on Machine Learning*, pages 9812–9823. PMLR, 2021.
- Daniel A Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. *SIAM Journal on Computing*, 40(6):1913–1926, 2011.
- Joel A Tropp. Improved analysis of the subsampled randomized hadamard transform. *Advances in Adaptive Data Analysis*, 3(01n02):115–126, 2011.
- Pravin M Vaidya. A new algorithm for minimizing convex functions over convex sets. In *30th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 338–343. IEEE Computer Society, 1989.
- Santosh Vempala. Geometric random walks: a survey. *Combinatorial and computational geometry*, 52(573-612):2, 2005.
- Yu-Xiang Wang, Stephen Fienberg, and Alex Smola. Privacy for free: Posterior sampling and stochastic gradient monte carlo. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2493–2502, Lille, France, 07–09 Jul 2015. PMLR.
- Virginia Vassilevska Williams. Multiplying matrices faster than coppersmith-winograd. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing (STOC)*, pages 887–898. ACM, 2012.
- Virginia Vassilevska Williams, Yinzhan Xu, Zixuan Xu, and Renfei Zhou. New bounds for matrix multiplication: from alpha to omega. In *Proceedings of the Thirty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA'24*, 2024.
- David P. Woodruff. Sketching as a tool for numerical linear algebra. *Foundations and Trends in Theoretical Computer Science*, 10(1–2):1–157, 2014.
- Difan Zou, Pan Xu, and Quanquan Gu. Stochastic variance-reduced hamilton monte carlo methods. In *International Conference on Machine Learning*, pages 6028–6037. PMLR, 2018.

Appendix

Roadmap. Since there are many technical details in the appendix, we provide a roadmap. The appendix can be divided into 5 parts: the first part conveys some preliminary information and states the sufficient conditions we require barrier functions to have, in Section A and B. The second part proves the mixing time of the Dikin walk when barrier functions satisfy the conditions, and the proofs are divided into Section C and D. The next part focuses on the runtime complexity of sampling from polytopes, including how to generate a spectral approximation of the Hessian, approximate leverage scores and Lewis weights to high precision, in Section E and how to incorporate these algorithmic prototypes to implement an efficient sampling algorithm in Section F. We dedicate Section G to study the log-barrier for sampling from a spectrahedron, as it is relatively less explored before. Finally, we prove the convexity of the function $\log \det(H(x) + I_d)$ for log-barrier, volumetric barrier and in extension Lee-Sidford barrier in Section H and I.

A Preliminaries

For any positive integer, we use $[n]$ to denote the set $\{1, 2, \dots, n\}$. For a vector $x \in \mathbb{R}^n$, we use $\|x\|_2$ to denote its ℓ_2 norm, i.e., $\|x\|_2 := (\sum_{i=1}^n x_i^2)^{1/2}$. We use $\|x\|_1$ to denote its ℓ_1 norm, $\|x\|_1 := \sum_{i=1}^n |x_i|$. We use $\|x\|_\infty$ to denote its ℓ_∞ norm, i.e., $\|x\|_\infty := \max_{i \in [n]} |x_i|$. For a random variable X , we use $\mathbb{E}[X]$ to denote its expectation. We use $\Pr[\cdot]$ to denote probability.

We use $\mathbf{0}_d$ to denote a length- d vector where every entry is 0. We use $\mathbf{1}_d$ to denote a length- d vector where every entry is 1. We use I_d to denote an identity matrix which has size $d \times d$ or simply I when dimension is clear from context.

For a matrix A , we use A^\top to denote the transpose of matrix A . For a square and non-singular matrix A , we use A^{-1} to denote the inverse of matrix A . For a real square matrix A , we say it is positive definite (PD, i.e., $A \succ 0$) if for all vectors $x \in \mathbb{R}^n$ (except for $\mathbf{0}_n$), we have $x^\top A x > 0$. For a real square matrix A , we say it is positive semi-definite (PSD, i.e., $A \succeq 0$) if for all vectors $x \in \mathbb{R}^n$, we have $x^\top A x \geq 0$. For a square matrix, we use $\det(A)$ to denote the determinant of matrix A . For a matrix A , we use $\|A\|$ to denote its spectral norm, use $\|A\|_F$ to denote its Frobenius norm, i.e., $\|A\|_F := (\sum_{i=1}^n \sum_{j=1}^d A_{i,j}^2)^{1/2}$ and use $\text{nnz}(A)$ to denote the number of nonzero entries in A .

Given a function $f : \mathcal{K} \rightarrow \mathbb{R}$, we say it's convex if for any $x, y \in \mathcal{K}$, $f(x) \geq f(y) + \nabla f(y)^\top (x - y)$. We say it's L -Lipschitz if $|f(x) - f(y)| \leq L \cdot \|x - y\|_2$ for a fixed parameter $L > 0$.

For two distributions P_1 and P_2 , we use $\text{TV}(P_1, P_2)$ to denote the total variation (TV) distance between P_1 and P_2 .

For a convex body $\mathcal{K} \subseteq \mathbb{R}^d$, we use $\text{Int}(\mathcal{K})$ to denote the interior of \mathcal{K} .

We use \otimes to denote Kronecker product. We use $A \otimes_S B$ to denote $A \otimes I + I \otimes B$. Given a matrix A , we use $\text{vec}(A)$ to denote its vectorization, since we will always apply $\text{vec}(\cdot)$ to a symmetric matrix, whether the vectorization is row-major or column-major doesn't matter. We use \circ to denote Hadamard or element-wise product.

For two vectors $x, y \in \mathbb{R}^d$, we use $\langle x, y \rangle = x^\top y$ to denote the standard Euclidean inner product over \mathbb{R}^d , and for two symmetric matrices $A, B \in \mathbb{R}^{d \times d}$, we use $\langle A, B \rangle = \text{tr}[A^\top B]$ to denote the trace inner product.

In Section A.1, we present several definitions and notations related to fast matrix multiplication. In Section A.2, we provide some backgrounds about convex geometry. In Section A.3, we state several basic probability tools. In Section A.4, we state several basic facts related to trace, and Kronecker product. In Section A.5, we present the standard notion about self-concordance barrier.

A.1 Fast matrix multiplication

Definition A.1 (Fast matrix multiplication). *Given three positive integers a, b, c , we use $\mathcal{T}_{\text{mat}}(a, b, c)$ to denote the time of multiplying an $a \times b$ matrix with another $b \times c$ matrix.*

For convenience, we also define the $\omega(\cdot, \cdot, \cdot)$ function as follows:

Definition A.2 (Fast matrix multiplication, an alternative notation). *Given x, y, z , we use $d^{\omega(x,y,z)}$ to denote the time of multiplying a $d^x \times d^y$ matrix with another $d^y \times d^z$ matrix.*

Lemma A.3 (Gall and Urrutia [2018]). *We have the following bounds of $\omega(\cdot, \cdot, \cdot)$:*

- $\omega(1, 1, 1) = \omega$,
- $\omega(1, 3, 1) = 4.199712$ (see Table 3 in Gall and Urrutia [2018]).

Here, ω denotes the exponent of matrix multiplication, currently $\omega \approx 2.373$ [Williams, 2012, Gall and Urrutia, 2018, Alman and Williams, 2021, Duan et al., 2023, Williams et al., 2024, Gall, 2024].

A.2 Convex geometry

We define $B(x, R) \subset \mathbb{R}^d$ as

$$B(x, R) := \{y \in \mathbb{R}^d : \|y - x\|_2 \leq R\}.$$

Define polytope \mathcal{K} as

$$\mathcal{K} := \{x \in \mathbb{R}^d : Ax \leq b\}.$$

Define spectrahedron \mathcal{K} as

$$\mathcal{K} := \{x \in \mathbb{R}^d : \sum_{i=1}^d x_i A_i \succeq C\},$$

where $A_1, \dots, A_d, C \in \mathbb{R}^{n \times n}$ are symmetric matrices.

We will often use the notion of a *Dikin ellipsoid*:

Definition A.4 (Dikin ellipsoid). *We define the Dikin ellipsoid $D_\theta \subset \mathbb{R}^d$ as follows*

$$D_\theta := \{w \in \mathbb{R}^d : w^\top H^{-1}(\theta)w \leq 1\}$$

where $H(\theta) \in \mathbb{R}^{d \times d}$ is the Hessian of self-concordant barrier at θ

We define the standard cross-ratio distance (see Definition E.1 in Mangoubi and Vishnoi [2023] for example).

Definition A.5 (Cross-ratio distance). *Let $u, v \in \mathcal{K}$. If $u \neq v$, let p, q be two endpoints of the chord in \mathcal{K} which passes through u and v such that the four points lie in the order of p, u, v, q , let*

$$\sigma(u, v) := \frac{\|u - v\|_2 \cdot \|p - q\|_2}{\|p - u\|_2 \cdot \|v - q\|_2}.$$

We additionally set $\sigma(u, v) = 0$ for $u = v$.

For two subsets $S_1, S_2 \subseteq \mathcal{K}$, we define

$$\sigma(S_1, S_2) := \min\{\sigma(u, v) : u \in S_1, v \in S_2\}.$$

We state the standard isoperimetric inequality for cross-ratio distance.

Lemma A.6 (Lovász and Vempala [2003]). *Let $\pi : \mathbb{R}^d \rightarrow \mathbb{R}$ be a log-concave density, with support on a convex body \mathcal{K} . Then for any partition of \mathbb{R}^d into measurable sets S_1, S_2, S_3 , the induced measure π^* satisfies*

$$\pi^*(S_3) \geq \sigma(S_1, S_2) \cdot \pi^*(S_1) \cdot \pi^*(S_2).$$

A.3 Probability tools

We state some useful concentration inequalities.

Lemma A.7 (Lemma 1 on page 1325 of [Laurent and Massart \[2000\]](#)). Let $X \sim \mathcal{X}_k^2$ be a chi-squared distributed random variable with k degrees of freedom. Each random variable has zero mean and σ^2 variance. Then

$$\Pr[X - k\sigma^2 \geq (2\sqrt{kt} + 2t)\sigma^2] \leq \exp(-t),$$

$$\Pr[k\sigma^2 - X \geq 2\sqrt{kt}\sigma^2] \leq \exp(-t).$$

Lemma A.8 (Matrix Chernoff bound [[Tropp, 2011](#)]). Let X_1, \dots, X_s be independent copies of a symmetric random matrix $X \in \mathbb{R}^{d \times d}$ with $\mathbb{E}[X] = 0$, $\|X\| \leq \gamma$ almost surely and $\|\mathbb{E}[X^\top X]\| \leq \sigma^2$. Let $W = \frac{1}{s} \sum_{i \in [s]} X_i$. For any $\epsilon \in (0, 1)$,

$$\Pr[\|W\| \geq \epsilon] \leq 2d \cdot \exp\left(-\frac{s\epsilon^2}{\sigma^2 + \gamma\epsilon/3}\right).$$

A.4 Basic algebra facts

Lemma A.9. Let A and B be $m \times m$ matrices. Then we have

- $\text{tr}[AB] = \text{tr}[BA]$.
- If A is symmetric, then $\text{tr}[AB] = \text{tr}[AB^\top]$.
- If A and B are PSD, then $\langle A, B \rangle \geq 0$, and $\langle A, B \rangle = 0$ if and only if $AB = \mathbf{0}_{m \times m}$.
- If $A \succeq 0$ and $B \succeq C$, then $\langle A, B \rangle \geq \langle A, C \rangle$.

Lemma A.10. Let A, B, C, D be conforming matrices. Then

- $(A \otimes B)(C \otimes D) = AC \otimes BD$.
- $(A \otimes_S B)(C \otimes_S D) = \frac{1}{2}(AC \otimes_S BD + AD \otimes_S BC)$.
- $(A \otimes B)^\top = A^\top \otimes B^\top$.
- If A and B are non-singular, then $A \otimes B$ is non-singular, and $(A \otimes B)^{-1} = A^{-1} \otimes B^{-1}$.
- $\text{vec}(ABC) = (C^\top \otimes A) \cdot \text{vec}(B)$.

A.5 Self-concordant barrier

We provide a standard definition about self-concordance barrier,

Definition A.11 (Self-concordant barrier). A real-valued function $F : \text{Int}(\mathcal{K}) \rightarrow \mathbb{R}$, is a regular self-concordant barrier if it satisfies the conditions stated below. For convenience, if $x \notin \text{Int}(\mathcal{K})$, we define $F(x) = \infty$.

1. (Convex, Smooth) F is a convex thrice continuously differentiable function on $\text{Int}(\mathcal{K})$.
2. (Barrier) For every sequence of points $\{x_i\} \in \text{Int}(\mathcal{K})$ converging to a point $x \notin \text{Int}(\mathcal{K})$, $\lim_{i \rightarrow \infty} f(x_i) = \infty$.
3. (Differential Inequalities) For all $h \in \mathbb{R}^d$ and all $x \in \text{Int}(\mathcal{K})$, the following inequalities hold.

(a) $D^2F(x)[h, h]$ is 2-Lipschitz continuous with respect to the local norm, which is equivalent to

$$D^3F(x)[h, h, h] \leq 2(D^2F(x)[h, h])^{\frac{3}{2}}.$$

(b) $F(x)$ is ν -Lipschitz continuous with respect to the local norm defined by F ,

$$\|\nabla_h F(x)\|_2^2 \leq \nu \cdot \|h\|_{H(x)}^2.$$

We call the smallest positive integer ν for which this holds, the self-concordance parameter of the barrier.

B Sufficient Conditions for Log-concave Sampling

To prove that our algorithm works for log-concave distribution, we state three key assumptions (see Section B.1) which are the sufficient conditions to prove the mixing rate of log-concave distribution, for different barrier functions.

B.1 Conditions: $\bar{\nu}$ -symmetry, convexity and bounded local norm

We state three key conditions in order to lower bound the conductance for log-concave sampling.

Assumption B.1. *Given a convex body \mathcal{K} , let $H : \mathcal{K} \rightarrow \mathbb{R}^{d \times d}$ be a self-concordant matrix function, we assume that*

- (i) **$\bar{\nu}$ -symmetry.** *For any $x \in \mathcal{K}$, we have $E_x(1) \subseteq \mathcal{K} \cap (2x - \mathcal{K}) \subseteq E_x(\sqrt{\bar{\nu}})$ where $E_x(r) = \{y \in \mathbb{R}^d : (y - x)^\top H(x)(y - x) \leq r^2\}$.*
- (ii) **Convexity.** *Let $F : \mathcal{K} \rightarrow \mathbb{R}$ be defined as $F(x) = \log(\det(H(x) + I_d))$, then F is convex in x .*
- (iii) **Bounded local norm.** *Let $\nabla \log(\det(H(x)))$ denote the gradient of $\log(\det(H(x)))$ in x . For any $x \in \mathcal{K}$, we have*

$$\|(H(x))^{-1/2} \cdot \nabla \log(\det(H(x)))\|_2^2 \leq \tilde{O}(d).$$

For the $\bar{\nu}$ -symmetry assumption, we will prove that for barriers of the concern (log-barrier, Lee-Sidford barrier for polytopes and log-barrier for spectrahedra), $\bar{\nu} = \nu$. This characteristic is proved in Laddha et al. [2020] for polytopes, and we prove that for spectrahedra in Section G. For convexity, due to the large amount of calculations, we defer to Section H and I. For bounded local norm, the results for polytopes are similarly shown in Laddha et al. [2020], and we will prove for log-barrier over spectrahedra.

B.2 Local PSD approximation for any self-concordant matrix function

In this section, we prove a generalization of the matrix function self-concordance to regularized matrix function. When H is Hessian of the log barrier, this fact was proved in [Mangoubi and Vishnoi, 2023, Lemma E.3]. We generalize this fact to general barriers.

Lemma B.2. *Let $\alpha \in (0, 1)$. Let $H : \mathcal{K} \rightarrow \mathbb{R}$ be a self-concordant matrix function and $\Phi(u) := \alpha^{-1}H(u) + \eta^{-1}I_d$. For any $u, v \in \mathcal{K}$ such that $\|u - v\|_{\Phi(u)} \leq \frac{1}{\alpha^{1/2}}$ we have*

$$(1 - \alpha^{1/2} \cdot \|u - v\|_{\Phi(u)})^2 \cdot \Phi(v) \preceq \Phi(u) \preceq (1 + \alpha^{1/2} \cdot \|u - v\|_{\Phi(u)})^2 \cdot \Phi(v).$$

Proof. By [Laddha et al., 2020, Lemma 1.1], we have

$$(1 - \|u - v\|_{H(u)})^2 \cdot H(v) \preceq H(u) \preceq (1 + \|u - v\|_{H(u)})^2 \cdot H(v).$$

Because $H(u) \preceq \alpha \cdot \Phi(u)$, we get

$$(1 - \alpha^{1/2} \cdot \|u - v\|_{\Phi(u)})^2 \cdot H(v) \preceq H(u) \preceq (1 + \alpha^{1/2} \cdot \|u - v\|_{\Phi(u)})^2 \cdot H(v).$$

Using $\Phi(u) = \alpha^{-1}H(u) + \eta^{-1}I_d$ again we get

$$(1 - \alpha^{1/2} \cdot \|u - v\|_{\Phi(u)})^2 \cdot \Phi(v) \preceq \Phi(u) \preceq (1 + \alpha^{1/2} \cdot \|u - v\|_{\Phi(u)})^2 \cdot \Phi(v).$$

This completes the proof. □

B.3 Bounded local norm for regularized Hessian function

We note that the bounded local norm condition generally holds for standard barrier functions, but in our core argument, we will instead rely on the bounded local norm of the *regularized Hessian function*. We prove that the bounded local norm condition implies the bounded local norm on the regularized Hessian function.

Lemma B.3. Let $H : \mathcal{K} \rightarrow \mathbb{R}^{d \times d}$ be a matrix function, define $F(x) = \log \det(H(x) + L^2 I_d)$ for $L \in \mathbb{R}$, then we have the following inequality:

$$\|(H(x) + L^2 I_d)^{-1/2} \nabla F(x)\|_2^2 \leq \|H(x)^{-1/2} \nabla \log \det(H(x))\|_2^2.$$

Proof. Let $H(x) = U \Lambda U^\top$ be the eigendecomposition of the matrix $H(x)$, let $y = U^\top \nabla H(x)$, then

$$\begin{aligned} \|(H(x) + L^2 I_d)^{-1/2} \nabla F(x)\|_2^2 &= \|(H(x) + L^2 I_d)^{-3/2} \nabla H(x)\|_2^2 \\ &= \|U(\Lambda + L^2 I_d)^{-3/2} U^\top \nabla H(x)\|_2^2 \\ &= \|(\Lambda + L^2 I_d)^{-3/2} y\|_2^2, \end{aligned}$$

the given condition implies

$$\begin{aligned} \|H(x)^{-1/2} \nabla \log \det(H(x))\|_2^2 &= \|H(x)^{-3/2} \nabla H(x)\|_2^2 \\ &= \|U \Lambda^{-3/2} U^\top \nabla H(x)\|_2^2 \\ &= \|\Lambda^{-3/2} y\|_2^2. \end{aligned}$$

As $L^2 \geq 0$, we can expand the regularized squared ℓ_2 norm as

$$\begin{aligned} \|(\Lambda + L^2 I_d)^{-3/2} y\|_2^2 &= \sum_{i=1}^d \frac{1}{(\lambda_i + L^2)^3} y_i^2 \\ &\leq \sum_{i=1}^d \frac{1}{\lambda_i^3} y_i^2 \\ &= \|\Lambda^{-3/2} y\|_2^2. \end{aligned}$$

This completes the proof. □

C Key Tools for Robust Sampling

In this section, we list several tools which can be shared in the core proofs of all the main theorems. This section is organized as follows. In Section C.1, we provide a tool for Gaussian concentration. In Section C.2, we show that PSD approximation implies the determinant approximation (by losing a factor of d). One of the major idea in this work is instead of using exact accept probability, we will only need to use approximate accept probability. In Section C.3, we prove several useful robust properties for approximate accept probability. In Section C.4, we present a general lemma for bounding the TV distance between exact process and approximate process. In Section C.5, we provide some specific choices for parameters and then bound the TV distance. In Section C.6, we provide a definition which will be used in proof of lower bound on conductance. In Section C.7, we show how to lower bound the conductance by cross ratio.

C.1 Gaussian concentration lemma

Lemma C.1. Let $\xi \sim \mathcal{N}(0, I_d)$, then for any $t > \sqrt{2d}$, we have

$$\Pr[\|\xi\|_2 \geq t] \leq \exp(-(t^2 - d)/8).$$

Proof. We consider the squared ℓ_2 norm of ξ : $\|\xi\|_2^2 = \sum_{i=1}^d \xi_i^2$, which is a χ^2 distribution with degree of freedom d . By Lemma A.7, we know that

$$\Pr[\|\xi\|_2^2 \geq 2\sqrt{dk} + 2k^2 + d] \leq \exp(-k^2),$$

set $k^2 = (t^2 - d)/8$, we obtain the desired bound. □

C.2 Spectral approximation to determinant approximation

Lemma C.2. Let $\epsilon_\Phi \in (0, 1)$. Given four psd matrices satisfy the following conditions

$$\begin{aligned}(1 - \epsilon_\Phi)\Phi(x) &\preceq \tilde{\Phi}(x) \preceq (1 + \epsilon_\Phi)\Phi(x), \\ (1 - \epsilon_\Phi)\Phi(z) &\preceq \hat{\Phi}(z) \preceq (1 + \epsilon_\Phi)\Phi(z).\end{aligned}$$

Then we have

$$(1 - 10 \cdot \epsilon_\Phi \cdot d) \cdot \frac{\det(\Phi(x))}{\det(\Phi(z))} \leq \frac{\det(\tilde{\Phi}(x))}{\det(\hat{\Phi}(z))} \leq (1 + 10 \cdot \epsilon_\Phi \cdot d) \cdot \frac{\det(\Phi(x))}{\det(\Phi(z))}$$

Proof. Since $\tilde{\Phi}(x) \in (1 \pm \epsilon_\Phi)\Phi(x)$, we have the following spectral bound:

$$(1 - \epsilon_\Phi)I_d \preceq \Phi(x)^{-1/2}\tilde{\Phi}(x)\Phi(x)^{-1/2} \preceq (1 + \epsilon_\Phi)I_d,$$

we can take the determinant of the middle term:

$$\begin{aligned}\det(\Phi(x)^{-1/2}\tilde{\Phi}(x)\Phi(x)^{-1/2}) &= \det(\Phi(x)^{-1/2}) \det(\tilde{\Phi}(x)) \det(\Phi(x)^{-1/2}) \\ &= \frac{\det(\tilde{\Phi}(x))}{\det(\Phi(x))} \\ &\in \det((1 \pm \epsilon_\Phi)I_d) \\ &\in (1 \pm \epsilon_\Phi)^d \\ &\in 1 + 10 \cdot \epsilon_\Phi \cdot d.\end{aligned}$$

Thus, we complete the proof. □

C.3 Approximate accept probability

We start with some definitions:

Definition C.3. We define

$$\begin{aligned}\hat{G}_u(x) &:= \det(\hat{\Phi}(u))^{1/2} \cdot \exp(-0.5\|u - x\|_{\hat{\Phi}(u)}^2) \\ G_x(u) &:= \det(\Phi(x))^{1/2} \cdot \exp(-0.5\|u - x\|_{\Phi(x)}^2) \\ \tilde{G}_x(u) &:= \det(\tilde{\Phi}(x))^{1/2} \cdot \exp(-0.5\|u - x\|_{\tilde{\Phi}(x)}^2)\end{aligned}$$

Definition C.4. We define

$$\begin{aligned}P_{x, \hat{\Phi}(x)}^{\text{accept}}(u) &:= \mathbb{E}_{\hat{\Phi}(u)} \left[\min\left\{1, \frac{\hat{G}_u(x) \exp(-f(u))}{G_x(u) \exp(-f(x))}\right\} \right] \\ P_{x, \tilde{\Phi}(x)}^{\text{accept}}(u) &:= \mathbb{E}_{\tilde{\Phi}(u)} \left[\min\left\{1, \frac{\hat{G}_u(x) \exp(-f(u))}{\tilde{G}_x(u) \exp(-f(x))}\right\} \right]\end{aligned}$$

We will use P_x, \tilde{P}_x as a shorthand for the above notations, and we call P_x the exact process and \tilde{P}_x the approximate process.

Lemma C.5. Fix x and $\tilde{\Phi}(x)$. Suppose $\tilde{\Phi}(x)$ is (ϵ_H, δ_H) -good approximation to $\Phi(x)$. Then

$$\epsilon_p := \mathbb{E}_{u \sim \mathcal{N}(x, \Phi^{-1}(x))} [|P_{x, \Phi(x)}^{\text{accept}}(u) - P_{x, \tilde{\Phi}(x)}^{\text{accept}}(u)|] \leq 0.001.$$

Proof. Note that

$$\frac{G_x(u)}{\tilde{G}_x(u)} = \frac{(\det \Phi(x))^{1/2} \exp(-\frac{1}{2}\|x - u\|_{\Phi(x)}^2)}{(\det \tilde{\Phi}(x))^{1/2} \exp(-\frac{1}{2}\|x - u\|_{\tilde{\Phi}(x)}^2)}.$$

Because $\tilde{\Phi}(x)$ is a good approximation to $\Phi(x)$, by Lemma C.2 we have

$$(1 - \epsilon_H)^{d/2} \leq \frac{(\det \Phi(x))^{1/2}}{(\det \tilde{\Phi}(x))^{1/2}} \leq (1 + \epsilon_H)^{d/2}$$

and

$$(1 - \epsilon_H)^2 \leq \frac{\|x - u\|_{\Phi(x)}^2}{\|x - u\|_{\tilde{\Phi}(x)}^2} \leq (1 + \epsilon_H)^2.$$

By choosing step size sufficiently small, we have $\|u - x\|_{\Phi(x)}^2 < d$ with probability 0.9999. Therefore

$$\Pr \left[\left| \|x - u\|_{\Phi(x)}^2 - \|x - u\|_{\tilde{\Phi}(x)}^2 \right| \leq 3\epsilon_H d \right] \geq 0.9999.$$

Then

$$\Pr \left[1 - 3\epsilon_H d \leq \frac{\exp(-0.5\|x - u\|_{\Phi(x)}^2)}{\exp(-0.5\|x - u\|_{\tilde{\Phi}(x)}^2)} \leq 1 + 3\epsilon_H d \right] \geq 0.9999.$$

For $\epsilon_H d$ small enough, we have

$$\Pr \left[0.9999 \leq \frac{P_{x, \Phi(x)}^{\text{accept}}(u)}{P_{x, \tilde{\Phi}(x)}^{\text{accept}}(u)} \leq 1.0001 \right] \geq 0.9998,$$

and therefore

$$\Pr \left[\left| P_{x, \Phi(x)}^{\text{accept}}(u) - P_{x, \tilde{\Phi}(x)}^{\text{accept}}(u) \right| \leq 0.0001 \right] \geq 0.9998.$$

So

$$\begin{aligned} \mathbb{E}_{u \sim \mathcal{N}(x, \Phi^{-1}(x))} \left[\left| P_{x, \Phi(x)}^{\text{accept}}(u) - P_{x, \tilde{\Phi}(x)}^{\text{accept}}(u) \right| \right] &\leq 0.9998 \cdot 0.0001 + 0.0002 \cdot 1 \\ &\leq 0.001. \end{aligned}$$

This completes the proof. □

C.4 TV distance between exact process and approximate process

Lemma C.6 (TV distance between exact process and approximate process). *Let $\tilde{\Phi}$ denote the $(1 \pm \epsilon_H)$ approximation of Φ . Let δ_H denote the failure probability. Let ϵ_p be defined as Lemma C.5. Let P_x denote the exact process. Let \tilde{P}_x denote the approximate process. If $x \in \mathcal{K}$, then*

$$\text{TV}(P_x, \tilde{P}_x) \leq \delta_H + \epsilon_p + \sqrt{d\epsilon_H}.$$

Proof. We say $\tilde{\Phi}(x)$ is good if it is an (ϵ_H, δ_H) -approximation to $\Phi(x)$, and $\tilde{\Phi}(x)$ is bad otherwise. We can upper bound $\text{TV}(P_x, \tilde{P}_x)$ as follows:

$$\begin{aligned} \text{TV}(P_x, \tilde{P}_x) &\leq \mathbb{E}[\text{TV}(P_{x, \Phi(x)}, P_{x, \tilde{\Phi}(x)})] \\ &= \mathbb{E}[\text{TV}(P_{x, \Phi(x)}, P_{x, \tilde{\Phi}(x)}) \mid \tilde{\Phi}(x) \text{ is bad}] \\ &\quad + \mathbb{E}[\text{TV}(P_{x, \Phi(x)}, P_{x, \tilde{\Phi}(x)}) \mid \tilde{\Phi}(x) \text{ is good}] \\ &\leq \Pr[\tilde{\Phi}(x) \text{ is bad}] + \mathbb{E}[\text{TV}(P_{x, \Phi(x)}, P_{x, \tilde{\Phi}(x)}) \mid \tilde{\Phi}(x) \text{ is good}] \\ &\leq \delta_H + \mathbb{E}[\text{TV}(P_{x, \Phi(x)}, P_{x, \tilde{\Phi}(x)}) \mid \tilde{\Phi}(x) \text{ is good}]. \end{aligned}$$

where the third step follows from $\text{TV}(\cdot, \cdot) \leq 1$, the last step follows from Lemma E.3.

It remains to compute the second term in the above equation. By Pinsker's inequality, when $\tilde{\Phi}(x)$ is good, we have

$$\text{TV}(\mathcal{N}(x, \tilde{\Phi}^{-1}(x)), \mathcal{N}(x, \Phi^{-1}(x)))^2$$

$$\begin{aligned}
&\leq 0.5 \cdot D(\mathcal{N}(x, \tilde{\Phi}^{-1}(x)) \parallel \mathcal{N}(x, \Phi^{-1}(x))) \\
&= 0.25 \cdot (\text{tr}[\tilde{\Phi}(x)^{-1/2} \Phi(x) \tilde{\Phi}(x)^{-1/2}] - d + \log \frac{\det \tilde{\Phi}(x)}{\det \Phi(x)}) \\
&\leq 0.25 \cdot ((1 \pm \epsilon_H)d - d + \log(1 \pm \epsilon_H)^d) \\
&\leq 0.5 \cdot d\epsilon_H.
\end{aligned} \tag{2}$$

where the first step follows from Pinsker inequality.

Now, we can upper bound $\text{TV}(P_x, \tilde{P}_x)$ in the following sense,

$$\begin{aligned}
\text{TV}(P_x, \tilde{P}_x) &\leq \delta_H + \mathbb{E}[\text{TV}(P_{x, \Phi(x)}, P_{x, \tilde{\Phi}(x)}) \mid \tilde{\Phi}(x) \text{ is good}] \\
&\leq \delta_H + \mathbb{E}_{\substack{u \sim \mathcal{N}(x, \Phi^{-1}(x)) \\ u' \sim \mathcal{N}(x, \tilde{\Phi}^{-1}(x)) \\ \text{any coupling}}} [|P_{x, \Phi(x)}^{\text{accept}}(u) - P_{x, \tilde{\Phi}(x)}^{\text{accept}}(u)| \cdot \mathbb{1}\{u = u'\} + \mathbb{1}\{u \neq u'\}] \\
&\leq \delta_H + \mathbb{E}_{u \sim \mathcal{N}(x, \Phi^{-1}(x))} [|P_{x, \Phi(x)}^{\text{accept}}(u) - P_{x, \tilde{\Phi}(x)}^{\text{accept}}(u)|] + \text{TV}(\mathcal{N}(x, \tilde{\Phi}^{-1}(x)), \mathcal{N}(x, \Phi^{-1}(x))) \\
&= \delta_H + \epsilon_p + \text{TV}(\mathcal{N}(x, \tilde{\Phi}^{-1}(x)), \mathcal{N}(x, \Phi^{-1}(x))) \\
&\leq \delta_H + \epsilon_p + \sqrt{d\epsilon_H}
\end{aligned}$$

The fourth step follows from definition of ϵ_p , the fifth step follows from Eq. (2). \square

C.5 TV distance between exact and approximate process: instantiation

Lemma C.7. *If $d\epsilon_H < 0.001$ and $\delta_H < 0.001$. Then, we have*

$$\text{TV}(P_x, \tilde{P}_x) \leq 0.01$$

Proof. We can upper bound it as follows

$$\begin{aligned}
\text{TV}(P_x, \tilde{P}_x) &\leq \delta_H + \epsilon_p + \sqrt{d\epsilon_H} \\
&\leq 0.001 + \epsilon_p + 0.001 \\
&\leq 0.01
\end{aligned}$$

where the first step follows from Lemma C.6, the second step follows from choice ϵ_H, δ_H , and the last step follows from Lemma C.5. \square

C.6 Lower bound on conductance, definition

Definition C.8. *Let $\beta \in (0, 0.1)$ denote some fixed constant. Let $S_1 \subset \mathcal{K}$ and $S_2 := \mathcal{K} \setminus S_1$ and $\pi(S_1) \leq 1/2$. We define S'_1 and S'_2 as follows*

$$\begin{aligned}
S'_1 &:= \{x \in S_1 : \tilde{P}_x(S_2) \leq \beta\}, \\
S'_2 &:= \{z \in S_2 : \tilde{P}_z(S_1) \leq \beta\}.
\end{aligned}$$

C.7 Lower bound on conductance, lemma

Lemma C.9 (Lower bound the conductance by cross ratio). *Let S'_1 and S'_2 be defined as Definition C.8. The conductance ϕ satisfies*

$$\phi \geq \frac{1}{16} \beta \cdot \sigma(S'_1, S'_2).$$

Proof. The proof follows the general format for conductance proofs for geometric Markov chains (see e.g. Section 5 of Vempala [2005]).

Let $S_1 \subseteq \mathcal{K}$ and let $S_2 = \mathcal{K} \setminus S_1$. Let S'_1 and S'_2 be defined as Definition C.8. We define S'_3

$$S'_3 := (\mathcal{K} \setminus S'_1) \setminus S'_2.$$

By Lemma D.12, we have that

$$\int_{x \in S_1} \pi(x) \tilde{P}_x(S_2) dx = \int_{x \in S_2} \pi(x) \tilde{P}_x(S_1) dx. \quad (3)$$

Thus, by Lemma A.6, we have

$$\pi^*(S'_3) \geq \sigma(S'_1, S'_2) \pi^*(S'_1) \pi^*(S'_2). \quad (4)$$

Case 1. First, we assume that both $\pi^*(S'_1) \geq \frac{1}{4} \pi^*(S_1)$ and $\pi^*(S'_2) \geq \frac{1}{4} \pi^*(S_2)$. In this case we have

$$\begin{aligned} \int_{S_1} \tilde{P}_x(S_2) \pi(x) dx &= \frac{1}{2} \int_{S_1} \tilde{P}_x(S_2) \pi(x) dx + \frac{1}{2} \int_{S_2} \tilde{P}_x(S_1) \pi(x) dx \\ &\geq \frac{\beta}{2} \cdot \pi^*(S'_3) \\ &\geq \frac{\beta}{2} \cdot \sigma(S'_1, S'_2) \cdot \pi^*(S'_1) \pi^*(S'_2) \\ &\geq \frac{\beta}{4} \cdot \sigma(S'_1, S'_2) \cdot \min(\pi^*(S'_1), \pi^*(S'_2)) \\ &\geq \frac{\beta}{16} \cdot \sigma(S'_1, S'_2) \cdot \min(\pi^*(S_1), \pi^*(S_2)). \end{aligned} \quad (5)$$

where the first step follows from Eq. (3), the second step follows from Definition C.8, the third step follows from Eq. (4), the fifth step follows from one of $\pi^*(S'_1)$ and $\pi^*(S'_2)$ is at least $1/2$.

Case 2. Now suppose that instead either $\pi^*(S'_1) < \frac{1}{4} \pi^*(S_1)$ or $\pi^*(S'_2) < \frac{1}{4} \pi^*(S_2)$.

Case 2a. If $\pi^*(S'_1) < \frac{1}{4} \pi^*(S_1)$ then we have

$$\begin{aligned} \int_{S_1} \tilde{P}_x(S_2) \pi(x) dx &= \frac{1}{2} \int_{S_1} \tilde{P}_x(S_2) \pi(x) dx + \frac{1}{2} \int_{S_2} \tilde{P}_x(S_1) \pi(x) dx \\ &\geq \frac{1}{2} \int_{S_1 \setminus S'_1} \tilde{P}_x(S_2) \pi(x) dx \\ &\geq \frac{1}{2} \cdot \frac{3}{4} \cdot \beta \pi^*(S_1) \\ &\geq \frac{3}{8} \beta \min(\pi^*(S_1), \pi^*(S_2)). \end{aligned} \quad (6)$$

where the first step follows from Eq. (3), and the third step follows from Definition C.8.

Case 2b. Similarly, if $\pi^*(S'_2) < \frac{1}{4} \pi^*(S_2)$ we have

$$\begin{aligned} \int_{S_1} \tilde{P}_x(S_2) \pi(x) dx &= \frac{1}{2} \int_{S_1} \tilde{P}_x(S_2) \pi(x) dx + \frac{1}{2} \int_{S_2} \tilde{P}_x(S_1) \pi(x) dx \\ &\geq \frac{1}{2} \int_{S_2 \setminus S'_2} \tilde{P}_x(S_1) \pi(x) dx \\ &\geq \frac{1}{2} \cdot \frac{3}{4} \cdot \beta \pi^*(S_2) \\ &\geq \frac{3}{8} \beta \min(\pi^*(S_1), \pi^*(S_2)). \end{aligned} \quad (7)$$

where the first step follows from Eq. (3), and the third step follows from Definition C.8.

Therefore, Eq. (5), (6), and (7) together imply that

$$\frac{1}{\min(\pi^*(S_1), \pi^*(S_2))} \int_{S_1} \tilde{P}_x(S_2) \pi(x) dx \geq \frac{\beta}{16} \cdot \sigma(S'_1, S'_2) \quad (8)$$

for every partition $S_1 \cup S_2 = \mathcal{K}$.

Hence, Eq. (8) implies that

$$\phi = \inf_{S \subseteq \mathcal{K}: \pi^*(S) \leq \frac{1}{2}} \frac{1}{\pi^*(S)} \int_S \tilde{P}_x(\mathcal{K} \setminus S) \pi(x) dx \geq \frac{\beta}{16} \sigma(S'_1, S'_2). \quad \square$$

D Correctness for General Barrier Functions with Regularization

In Section D.1, we show how to bound the density ratios. In Section D.2, we show how to bound the determinant. In Section D.3, we show how to bound the local norm. In Section D.4, we explain how to bound the TV distance between two exact distributions. In Section D.5, we show how to bound the TV distance between two approximate processes. In Section D.6, we show how to bound the TV distance between two Gaussians. In Section D.7, we prove reversibility and stationary distribution. In Section D.8, we prove a lower bound on cross ratio distance. In Section D.9, we relate the cross ratio distance to local norm by utilizing $\bar{\nu}$ -symmetry.

Algorithm 2 Our algorithm for sampling from polytope with log-barrier (formal version of Algorithm 1). As the only differences between log-barrier for polytopes and others are the choice of ν and how to generate the approximate the Hessian, we only present this algorithm.

```

1: procedure MAIN( $A \in \mathbb{R}^{n \times d}, b \in \mathbb{R}^n, \delta \in (0, 1), x_0 \in \mathbb{R}^d, \alpha > 0, \eta > 0$ ) ▷ Theorem 1.1
2: ▷  $\mathcal{K} := \{x \in \mathbb{R}^d : Ax \leq b\}$ 
3:   Let  $g$  be the log-barrier for  $\mathcal{K}$  with  $\nu = n$ 
4:    $\alpha \leftarrow \Theta(1/d)$ 
5:    $\eta \leftarrow \Theta(1/(dL^2))$ 
6:    $T \leftarrow (\nu\alpha^{-1} + \eta^{-1}R^2) \cdot \log(w/\delta)$ 
7:    $x \leftarrow x_0$  ▷ We are given  $x_0 \in \text{Int}(\mathcal{K})$ 
8:    $\epsilon \leftarrow \Theta(1)$ 
9:    $\epsilon_H \leftarrow \Theta(\epsilon/d)$ 
10:  for  $t = 1 \rightarrow T$  do
11:    Sample a point  $\xi \sim \mathcal{N}(0, I_d)$ 
12:    ▷ Let  $H$  denote the Hessian of barrier function  $g$ 
13:     $\tilde{H}(x) \leftarrow \text{SUBSAMPLE}(A, b, x, n, d, \epsilon_H)$  ▷ Algorithm 3
14:    ▷  $\tilde{H}(x)$  can be written  $\sum_{i \in S} \tilde{\sigma}_i a_i a_i^\top$  where  $|S| = \epsilon_H^{-2} d \log d$ 
15:    ▷ The above step takes  $\text{nnz}(A) \log n + \mathcal{T}_{\text{mat}}(d, \epsilon_H^{-2} d, d)$  time
16:     $\tilde{\Phi}(x) \leftarrow \alpha^{-1} \tilde{H}(x) + \eta^{-1} I_d$ 
17:     $z \leftarrow x + \tilde{\Phi}(x)^{-1/2} \xi$ 
18:    if  $z \in \text{Int}(\mathcal{K})$  then
19:       $\hat{H}(z) \leftarrow \text{SUBSAMPLE}(A, b, z, n, d, \epsilon_H)$  ▷ Algorithm 3
20:       $\hat{\Phi}(z) \leftarrow \alpha^{-1} \hat{H}(z) + \eta^{-1} I_d$ 
21:      accept  $x \leftarrow z$  with probability
          
$$\frac{1}{2} \cdot \min \left\{ \frac{\exp(-f(z)) \cdot (\det(\hat{\Phi}(z)))^{1/2} \cdot \exp(-0.5\|x - z\|_{\hat{\Phi}(z)}^2)}{\exp(-f(x)) \cdot (\det(\tilde{\Phi}(x)))^{1/2} \cdot \exp(-0.5\|x - z\|_{\tilde{\Phi}(x)}^2)}, 1 \right\}$$

22:    else
23:      reject  $z$ 
24:    end if
25:  end for
26:  return  $x$ 
27: end procedure

```

D.1 Bounding the density ratio

Lemma D.1 (Bounding the density ratio). *Let $x \in \text{Int}(\mathcal{K})$. Suppose f is L -Lipschitz and $\eta \leq 1/(10dL^2)$, then*

$$\Pr_{z \sim \mathcal{N}(x, \Phi^{-1}(x))} \left[\frac{\exp(-f(z))}{\exp(-f(x))} \geq \frac{1}{2} \right] \geq 0.999.$$

If in addition, f is β -smooth and $\eta \leq 1/(10d\beta)$, then we further have

$$\Pr_{z \sim \mathcal{N}(x, \Phi^{-1}(x))} \left[\frac{\exp(-f(x))}{\exp(-f(z))} \geq \frac{1}{2} \right] \geq 0.499.$$

Proof. Proof of Part 1. Since $z \sim \mathcal{N}(x, \Phi(x)^{-1})$, we have that

$$\begin{aligned} z &= x + \Phi(x)^{-1/2} \xi \\ &= x + (\alpha^{-1} \cdot H(x) + \eta^{-1} I_d)^{-1/2} \end{aligned}$$

for some $\xi \sim \mathcal{N}(0, I_d)$.

Since $\alpha^{-1} H(x) + \eta^{-1} I_d \succeq \eta^{-1} I_d$, and $H(x)$ and I_d are both positive definite, we have that

$$\eta \cdot I_d \succeq (\alpha^{-1} H(x) + \eta^{-1} I_d)^{-1} \quad (9)$$

Thus, we can upper bound $\|z - x\|_2$ as follows:

$$\begin{aligned} \|z - x\|_2 &= \|(\alpha^{-1} H(x) + \eta^{-1} I_d)^{-1/2} \xi\|_2 \\ &= \sqrt{\xi^\top (\alpha^{-1} H(x) + \eta^{-1} I_d)^{-1} \xi} \\ &\leq \sqrt{\xi^\top \eta I_d \xi} \\ &= \sqrt{\eta} \cdot \|\xi\|_2 \end{aligned}$$

where the third step follows from Eq. (9),

Recall that ξ is sampled from $\mathcal{N}(0, I_d)$, using Lemma C.1, we can show

$$\Pr[\|\xi\|_2 > t] \leq \exp(-(t^2 - d)/8), \quad \forall t > \sqrt{2d}.$$

Combining the above two equations, we have

$$\Pr[\|z - x\|_2 > \sqrt{\eta} \sqrt{20d}] \leq \exp(-\frac{19}{8}d) < 0.001. \quad (10)$$

Using $\eta \leq 1/(80dL^2)$, we have

$$\Pr[\|z - x\|_2 > 1/(2L)] < 0.001. \quad (11)$$

Since f is L -Lipschitz, then we have

$$\frac{\exp(-f(z))}{\exp(-f(x))} = \exp(-(f(z) - f(x))) \geq \exp(-L\|z - x\|_2).$$

Therefore,

$$\begin{aligned} \Pr_{z \sim \mathcal{N}(x, \Phi^{-1}(x))} \left[\frac{\exp(-f(z))}{\exp(-f(x))} \geq 1/2 \right] &\geq \Pr_{z \sim \mathcal{N}(x, \Phi^{-1}(x))} \left[\exp(-L\|z - x\|_2) \geq 1/2 \right] \\ &= \Pr_{z \sim \mathcal{N}(x, \Phi^{-1}(x))} \left[\|z - x\|_2 \leq \log(2)/L \right] \\ &\geq \Pr_{z \sim \mathcal{N}(x, \Phi^{-1}(x))} \left[\|z - x\|_2 \leq 1/(2L) \right] \\ &\geq 0.999 \end{aligned}$$

where the last inequality holds by Eq. (11).

Proof of Part 2. Moreover, in the setting where f is differentiable and β -smooth, we have that, since $z - x$ is a multivariate Gaussian random variable,

$$\Pr[(z - x)^\top \nabla f(x) \leq 0] = \frac{1}{2}.$$

If $(z - x)^\top \nabla f(x) \leq 0$, we have that

$$\begin{aligned} f(z) - f(x) &\leq (z - x)^\top \nabla f(x) + \beta \|z - x\|_2^2 \\ &\leq \beta \|z - x\|_2^2. \end{aligned}$$

Therefore,

$$\begin{aligned}
& \Pr_{z \sim N(x, \tilde{\Phi}^{-1}(x))} \left[\frac{\pi(z)}{\pi(x)} \geq \frac{1}{2} \right] \\
& \geq \Pr_{z \sim N(x, \tilde{\Phi}^{-1}(x))} \left[\frac{\pi(z)}{\pi(x)} \geq \frac{1}{2} \text{ and } (z-x)^\top \nabla f(x) \leq 0 \right] - \Pr_{z \sim N(x, \tilde{\Phi}^{-1}(x))} \left[(z-x)^\top \nabla f(x) > 0 \right] \\
& \geq \Pr_{z \sim N(x, \tilde{\Phi}^{-1}(x))} \left[e^{-\beta \|z-x\|_2^2} \geq \frac{1}{2} \right] - 0.5 \\
& = \Pr_{z \sim N(x, \tilde{\Phi}^{-1}(x))} \left[\|z-x\|_2 \leq \frac{\sqrt{\log(2)}}{\sqrt{\beta}} \right] - 0.5 \\
& \geq 0.999 - 0.5 \\
& = 0.499,
\end{aligned}$$

where the last Inequality holds by Eq. (10) since $\eta \leq \frac{1}{20d\beta}$. \square

Lemma D.2. Let $\|x-z\|_{\Phi(x)} < 0.001$. Let $\eta \leq 0.01/L^2$. Then we have

$$|f(x) - f(z)| \leq 0.1$$

Proof. We have

$$\begin{aligned}
|f(x) - f(z)| & \leq L \cdot \|x-z\|_2 \\
& \leq L \cdot \sqrt{\eta} \cdot \|x-z\|_{\Phi(x)} \\
& \leq 0.1
\end{aligned}$$

where the second step follows from $\Phi(x) \succeq \eta^{-1}I_d$, the last step follows from $\eta \leq 0.01/L^2$. \square

D.2 Bounding the determinant

Lemma D.3 (Bounding the determinant). Consider any $x \in \text{Int}(\mathcal{K})$, and $\xi \sim \mathcal{N}(0, I_d)$. Let $z = x + (\Phi(x))^{-1/2}\xi$. Then

$$\Pr_{\xi} \left[\log(\det(\Phi(z))) - \log(\det(\Phi(x))) \geq -0.1 \right] \geq 0.999.$$

Proof. First, note that

$$\begin{aligned}
\log(\det(\Phi(z))) - \log(\det(\Phi(x))) & = \log\left(\frac{\det(\Phi(z))}{\det(\Phi(x))}\right) \\
& = \log\left(\frac{\det(d \cdot H(z) + dL^2 \cdot I_d)}{\det(d \cdot H(x) + dL^2 \cdot I_d)}\right) \\
& = \log\left(\frac{\det(H(z) + L^2 \cdot I_d)}{\det(H(x) + L^2 \cdot I_d)}\right),
\end{aligned}$$

thus, it suffices to consider the function $F(x) = \log(\det(H(x) + L^2 \cdot I_d))$ and bound $F(z) - F(x)$. By Assumption (ii), we know that $F(x)$ is convex.

Thus,

$$F(z) - F(x) \geq (z-x)^\top \nabla F(x).$$

We know that

$$z = x + (\Phi(x))^{-1/2}\xi$$

where $\xi \sim \mathcal{N}(0, I_d)$.

Thus,

$$F(z) - F(x) \geq \xi^\top (\Phi(x))^{-1/2} \nabla F(x)$$

By property of Gaussian distribution, we know that

$$\xi^\top (\Phi(x))^{-1/2} \nabla F(x)$$

is a Gaussian with mean 0 and variance $\|(\Phi(x))^{-1/2} \nabla F(x)\|_2^2$.

By definition of Φ and Assumption (iii), we have

$$\|(H(x) + L^2 \cdot I_d)^{-1/2} \nabla F(x)\|_2^2 = O(d)$$

and therefore

$$\|\Phi(x)^{-1/2} \nabla F(x)\|_2^2 = O(1).$$

Thus, rescaling the function and applying the standard concentration inequality, we complete the proof. \square

Lemma D.4 (Bounding the determinant, shifted). *Let $\alpha < 0.001/d$. For fixed $x, z \in \mathcal{K}$ such that $\|z - x\|_{\Phi(x)} < 0.001$. Then, we have*

$$\log(\det(\Phi(z))) - \log(\det(\Phi(x))) \geq -0.01.$$

Proof. Let $F(x) = \log \det(H(x) + L^2 \cdot I_d)$. Then

$$\log(\det(\Phi(z))) - \log(\det(\Phi(x))) = F(z) - F(x) \geq (z - x)^\top \nabla F(x)$$

where the second step is because $F(x)$ is convex (by Assumption (ii)).

We have

$$\begin{aligned} ((z - x)^\top \nabla F(x))^2 &= ((z - x)^\top \Phi(x)^{1/2} \cdot \Phi(x)^{-1/2} \nabla F(x))^2 \\ &\leq \|(z - x)^\top \Phi(x)^{1/2}\|_2 \cdot \|\Phi(x)^{-1/2} \nabla F(x)\|_2 \\ &\leq 0.001 \cdot \|\Phi(x)^{-1/2} \nabla F(x)\|_2 \\ &\leq 0.001 \cdot 0.1. \end{aligned}$$

where the second step is by Cauchy-Schwarz, the third step is by assumption $\|z - x\|_{\Phi(x)} \leq 0.001$, the fourth step is by $\alpha < 0.001/d$ and Assumption (iii). \square

D.3 Bounding the local norm

Lemma D.5 (Bounding the local norm). *Consider any $x \in \text{Int}(\mathcal{K})$, and $\xi \sim \mathcal{N}(0, I_d)$. Let $\eta \leq 0.001/d$. Let $z = x + (\Phi(x))^{-1/2} \xi$. Then*

$$\Pr_{\xi} \left[\left| \|z - x\|_{\Phi(z)}^2 - \|z - x\|_{\Phi(x)}^2 \right| \leq 0.01 \right] \geq 0.999$$

Proof Sketch. Our proof is a generalization of [Sachdeva and Vishnoi, 2016, Proposition 7], where they prove the result for non-regularized log-barrier. Note that their proof also works for regularized barriers because two points are close in the $\Phi(x)$ norm indicates that they are close in the $H(x)$ norm. The only difference is that when computing the Gaussian polynomial as in Sachdeva and Vishnoi [2016], we need to handle non-uniform weights instead of the uniform weights of log-barrier. This could also be handled by observing that if two points are close in $H(x)$ norm, then their weights are close in the sense that $\|w_p(x)^{-1}(w_p(x) - w_p(z))\|_{\infty} \leq c_p$ where $c_p < 1$ is some small constant that depends on p (see [Lee and Sidford, 2019, Lemma 34]). One could then check that the argument of [Sachdeva and Vishnoi, 2016, Proposition 7] is robust under small perturbation to the weights, and conclude similar conclusions. \square

Lemma D.6 (Bounding the local norm, shifted). *Consider any $x \in \text{Int}(\mathcal{K})$, and $\xi \sim \mathcal{N}(0, I_d)$. Let $\eta \leq 0.001/d$. Let $\|z - x\|_{\Phi(x)} < 0.001$. Let $u = x + (\Phi(x))^{-1/2} \xi$. Then*

$$\Pr_{\xi} \left[\left| \|u - z\|_{\Phi(u)}^2 - \|u - x\|_{\Phi(u)}^2 \right| \leq 0.1 \right] \geq 0.999$$

Proof. Consider the following inequalities:

$$\begin{aligned} & \left| \|u - z\|_{\Phi(u)}^2 - \|u - x\|_{\Phi(u)}^2 \right| \\ & \leq \|u - x\|_{\Phi(u)}^2 + 2|\langle u - x, x - z \rangle_{\Phi(u)}| \\ & \leq 2\|u - x\|_{\Phi(x)}^2 + 4|\langle u - x, x - z \rangle_{\Phi(x)}| \\ & = 2\|u - x\|_{\Phi(x)}^2 + 4|w| \end{aligned}$$

where the second step follows from Lemma B.2, in the last step $w \sim \mathcal{N}(0, \|\Phi(x)^{1/2}(x - z)\|_2^2)$.

First, using η is small, we can show $\|u - x\|_{\Phi(x)}^2 < 0.01$ with probability 0.9999.

Second, using $\|z - x\|_{\Phi(x)} < 0.001$ and Gaussian concentration, we can show that

$$\Pr[|w| < 0.01] \geq 0.9999.$$

Put them together and union bound, we obtain the desired result. \square

Lemma D.7 (Bounding the local norm, shifted). *Consider any $x \in \text{Int}(\mathcal{K})$, and $\xi \sim \mathcal{N}(0, I_d)$. Let $\eta \leq 0.001/d$. Let $\|z - x\|_{\Phi(x)} < 0.001$. Let $u = x + (\Phi(x))^{-1/2}\xi$. We have*

$$\Pr[\|u - z\|_{\Phi(z)}^2 - \|u - x\|_{\Phi(u)}^2 < 0.1] \geq 0.998$$

Proof. Let $u' \sim z + \Phi(z)^{-1/2}\xi'$ where $\xi' \sim \mathcal{N}(0, I_d)$. We have

$$\begin{aligned} & \left| \|u - z\|_{\Phi(z)}^2 - \|u - x\|_{\Phi(u)}^2 \right| \\ & \leq \left| \|u - z\|_{\Phi(z)}^2 - \|u' - z\|_{\Phi(z)}^2 \right| + \left| \|u' - z\|_{\Phi(z)}^2 - \|u' - z\|_{\Phi(u')}^2 \right| \\ & \quad + \left| \|u' - z\|_{\Phi(u')}^2 - \|u - z\|_{\Phi(u)}^2 \right| + \left| \|u - z\|_{\Phi(u)}^2 - \|u - x\|_{\Phi(u)}^2 \right| \end{aligned}$$

So

$$\begin{aligned} & \Pr[\|u - z\|_{\Phi(z)}^2 - \|u - x\|_{\Phi(u)}^2 > 0.1] \\ & \leq \underbrace{\Pr[\|u' - z\|_{\Phi(z)}^2 - \|u' - z\|_{\Phi(u')}^2 > 0.05]}_{\text{Lemma D.5}} \\ & \quad + \underbrace{\Pr[\|u - z\|_{\Phi(u)}^2 - \|u - x\|_{\Phi(u)}^2 > 0.05]}_{\text{Lemma D.6}} + \underbrace{\Pr[u \neq u']}_{\text{Lemma D.10}} \\ & \leq 0.001 + 0.001 + 0.001 = 0.003. \end{aligned}$$

where first step is by union bound, second step is by Lemma D.5, Lemma D.6, and by $\text{TV}(u, u') \leq 0.001$ when $\|x - z\|_{\Phi(x)} \leq 0.001$ (Lemma D.10). \square

D.4 Bounding TV distance between exact processes

Recall that we define $P_x(S) = \Pr[x \in S]$. We prove a bound between exact distributions of two points x and z .

Lemma D.8 (TV distance between the exact distribution). *For any $x, z \in \mathcal{K}$ such that $\|x - z\|_{\Phi(x)} \leq 0.001$, we have that*

$$\text{TV}(P_x, P_z) \leq 0.99.$$

Proof. We define

$$G_{x, \Phi(x)}(u) := (\det(\Phi(x)))^{1/2} \cdot \exp(-0.5\|u - x\|_{\Phi(x)}^2).$$

Then

$$\text{TV}(P_{x, \Phi(x)}, P_{z, \Phi(z)}) = 1 - p_{\min},$$

where

$$p_{\min} := \mathbb{E}_{u \sim \mathcal{N}(x, \Phi^{-1}(x))} [\min\{1, \frac{G_z(u)}{G_x(u)}, \frac{\exp(-f(u))G_u(x)}{\exp(-f(x))G_x(u)}, \frac{\exp(-f(u))G_u(z)}{\exp(-f(z))G_z(u)} \cdot \frac{G_z(u)}{G_x(u)}\}].$$

For convenience, we define

$$\begin{aligned} a_1 &:= \frac{G_z(u)}{G_x(u)} \\ a_2 &:= \frac{\exp(-f(u))G_u(x)}{\exp(-f(x))G_x(u)} \\ a_3 &:= \frac{\exp(-f(u))G_u(z)}{\exp(-f(z))G_z(u)} \cdot \frac{G_z(u)}{G_x(u)} \end{aligned}$$

Let $\theta \in (0, 1)$. We have the following:

$$\begin{aligned} \text{TV}(P_x, P_z) &= 1 - p_{\min} \\ &= 1 - \mathbb{E}[\min\{1, a_1, a_2, a_3\}] \\ &= 1 - \min\{1, \theta\} \cdot \Pr[\min\{a_1, a_2, a_3\} \geq \theta] - \min\{a_1, a_2, a_3\} \cdot \Pr[\min\{a_1, a_2, a_3\} < \theta] \\ &\leq 1 - \min\{1, \theta\} \cdot \Pr[\min\{a_1, a_2, a_3\} \geq \theta] \\ &\leq 1 - \theta \cdot \Pr[\min\{a_1, a_2, a_3\} \geq \theta] \\ &\leq 1 - \theta \cdot \theta_0 \\ &\leq 1 - 0.01 \cdot 0.95 \\ &\leq 0.99 \end{aligned}$$

where the sixth step follows from Eq. (12), and the last step follows from choice of θ and θ_0 .

Let $a_i = e^{-b_i}$. It suffices to prove that for some $\theta \in (0, 1)$, we have

$$\Pr[\min\{a_1, a_2, a_3\} \geq \theta] \geq \theta_0. \quad (12)$$

Suppose for some $\tau \in [1, 5]$ we can show that

$$\begin{aligned} \Pr[b_1 < \tau] &\geq \theta_1 \\ \Pr[b_2 < \tau] &\geq \theta_2 \\ \Pr[b_3 < \tau] &\geq \theta_3 \end{aligned}$$

Then by union bound, we can know that

$$\begin{aligned} &\Pr[\max\{b_1, b_2, b_3\} < \tau] \\ &= \Pr[\min\{a_1, a_2, a_3\} > e^{-\tau}] \\ &= \Pr[\min\{a_1, a_2, a_3\} > \theta] \\ &\geq 1 - (1 - \theta_1) - (1 - \theta_2) - (1 - \theta_3) \\ &\geq 0.95 \end{aligned}$$

where the second step follows from $e^{-\tau} = \theta \geq 0.01$, and the last step follows from $\theta_1 = \theta_2 = \theta_3 = 0.99$. In the following, we will establish the three bounds of interest.

Part 1. By definition, we have

$$\begin{aligned} a_1 &= \frac{G_z(u)}{G_x(u)} \\ &= \frac{(\det(\Phi(z)))^{1/2}}{(\det(\Phi(x)))^{1/2}} \cdot \exp(-0.5\|u - z\|_{\Phi(z)}^2 + 0.5\|u - x\|_{\Phi(x)}^2) \\ &= \exp(-0.5\|u - z\|_{\Phi(z)}^2 + 0.5\|u - x\|_{\Phi(x)}^2 + 0.5 \log \det(\Phi(z)) - 0.5 \log(\det(\Phi(x)))) \end{aligned}$$

Using Lemma D.7 (on the term $0.5\|u - z\|_{\Phi(z)}^2 - 0.5\|u - x\|_{\Phi(x)}^2$) and Lemma D.5 (on the term $0.5\|u - x\|_{\Phi(x)}^2 - 0.5\|u - x\|_{\Phi(x)}^2$), we have

$$\Pr_u \left[\underbrace{0.5\|u - z\|_{\Phi(z)}^2 - 0.5\|u - x\|_{\Phi(x)}^2}_{\text{Lemma D.7}} + \underbrace{0.5\|u - x\|_{\Phi(x)}^2 - 0.5\|u - x\|_{\Phi(x)}^2}_{\text{Lemma D.5}} \leq 0.1\tau \right] \geq 0.999,$$

From Lemma D.4, we know

$$0.5 \log(\det(\Phi(x))) - 0.5 \log(\det(\Phi(z))) \leq 0.1\tau.$$

Thus,

$$\Pr[b_1 \leq \tau] \geq 0.99$$

Part 2. By definition, for a_2 we have

$$\begin{aligned} a_2 &= \frac{\exp(-f(u))G_u(x)}{\exp(-f(x))G_x(u)} \\ &= \frac{\exp(-f(u))}{\exp(-f(x))} \cdot \frac{(\det(\Phi(u)))^{1/2}}{(\det(\Phi(x)))^{1/2}} \cdot \exp(-0.5\|x-u\|_{\Phi(u)}^2 + 0.5\|u-x\|_{\Phi(x)}^2) \\ &= \exp(-f(u) + f(x) + 0.5 \log(\det(\Phi(u))) - 0.5 \log(\det(\Phi(x))) - 0.5\|x-u\|_{\Phi(u)}^2 + 0.5\|u-x\|_{\Phi(x)}^2) \end{aligned}$$

By Lemma D.1, we have

$$\Pr_u[f(u) - f(x) \leq 0.1\tau] \geq 0.999$$

By Lemma D.3, we have

$$\Pr_u[0.5 \log(\det(\Phi(x))) - 0.5 \log(\det(\Phi(u))) \leq 0.1\tau] \geq 0.999$$

By Lemma D.5, we have

$$\Pr_u[0.5\|x-u\|_{\Phi(u)}^2 - 0.5\|u-x\|_{\Phi(x)}^2 \leq 0.1\tau] \geq 0.999$$

Then, combining the above three equations, we can get the following result:

$$\Pr[b_2 \leq \tau] \geq 0.99$$

Part 3.

$$\begin{aligned} a_3 &= \frac{\exp(-f(u))G_u(z)}{\exp(-f(z))G_z(u)} \cdot \frac{G_z(u)}{G_x(u)} \\ &= \frac{\exp(-f(u))}{\exp(-f(z))} \cdot \frac{(\det(\Phi(u)))^{1/2}}{(\det(\Phi(x)))^{1/2}} \cdot \exp(-0.5\|z-u\|_{\Phi(u)}^2 + 0.5\|u-x\|_{\Phi(x)}^2) \\ &= \exp(-f(u) + f(z) + 0.5 \log(\det(\Phi(u))) - 0.5 \log(\det(\Phi(x))) - 0.5\|z-u\|_{\Phi(u)}^2 + 0.5\|u-x\|_{\Phi(x)}^2) \end{aligned}$$

By Lemma D.1 (on the term $f(u) - f(x)$) and Lemma D.2 (on the term $f(x) - f(z)$), we have

$$\Pr_u[\underbrace{f(u) - f(x)}_{\text{Lemma D.1}} + \underbrace{f(x) - f(z)}_{\text{Lemma D.2}} \leq 0.1\tau] \geq 0.999.$$

By Lemma D.3, we have

$$\Pr_u[0.5 \log(\det(\Phi(x))) - 0.5 \log(\det(\Phi(u))) \leq 0.1\tau] \geq 0.999.$$

By Lemma D.5 (on the term $0.5\|u-x\|_{\Phi(u)}^2 - 0.5\|u-x\|_{\Phi(x)}^2$) and Lemma D.6 (on the term $0.5\|z-u\|_{\Phi(u)}^2 - 0.5\|u-x\|_{\Phi(x)}^2$), we have

$$\Pr_u[\underbrace{0.5\|z-u\|_{\Phi(u)}^2 - 0.5\|u-x\|_{\Phi(x)}^2}_{\text{Lemma D.6}} + \underbrace{0.5\|u-x\|_{\Phi(u)}^2 - 0.5\|u-x\|_{\Phi(x)}^2}_{\text{Lemma D.5}} \leq 0.2\tau] \geq 0.999$$

Thus, combining the above three equations, we obtain the following:

$$\Pr[b_3 \leq \tau] \geq 0.99. \quad \square$$

D.5 Bounding TV distance between approximate processes

Lemma D.9 (Robust version of Lemma E.10 in Mangoubi and Vishnoi [2023]). *If $\delta_H < 0.001$, $d\alpha < 0.001$, $d\epsilon_H < 0.001$ and for any $x, z \in \mathcal{K}$, $\|x - z\|_{\Phi(x)} < 0.001$, then we have*

$$\text{TV}(\tilde{P}_x, \tilde{P}_z) \leq 0.99.$$

Proof. We can upper bound $\text{TV}(\tilde{P}_x, \tilde{P}_y)$ as follows:

$$\begin{aligned} \text{TV}(\tilde{P}_x, \tilde{P}_y) &\leq \text{TV}(\tilde{P}_x, P_x) + \text{TV}(P_x, P_y) + \text{TV}(P_y, \tilde{P}_y) \\ &\leq \text{TV}(\tilde{P}_x, P_x) + 0.9 + \text{TV}(P_y, \tilde{P}_y) \\ &\leq 0.01 + 0.9 + 0.01 \\ &\leq 0.99 \end{aligned}$$

First step is by triangle inequality. Second step is by Lemma D.8. The third step is by Lemma C.7. \square

D.6 Bounding TV distance between Gaussians

Lemma D.10 (Lemma E.6 in Mangoubi and Vishnoi [2023]). *For any $x, z \in \mathcal{K}$ such that $\|x - z\|_{\Phi(x)} \leq 0.001$, we have*

$$\text{TV}(\mathcal{N}(x, \Phi^{-1}(x)), \mathcal{N}(z, \Phi^{-1}(z))) \leq \sqrt{3d\alpha + 1/2} \cdot \|x - z\|_{\Phi(x)}.$$

Further, if $\alpha d < 0.001$, then we have

$$\text{TV}(\mathcal{N}(x, \Phi^{-1}(x)), \mathcal{N}(z, \Phi^{-1}(z))) \leq 0.001.$$

Lemma D.11 (Robust version of Lemma D.10). *For any $x, z \in \mathcal{K}$ such that $\|x - z\|_{\Phi(x)} \leq \frac{1}{4\alpha^{1/2}}$, with probability at least $1 - 1/1000$, we have*

$$\text{TV}(\mathcal{N}(x, \tilde{\Phi}^{-1}(x)), \mathcal{N}(z, \hat{\Phi}^{-1}(z))) \leq \sqrt{2d\epsilon_H} + \sqrt{3d\alpha + 1/2} \cdot \|x - z\|_{\Phi(x)}$$

Proof. By triangle inequality,

$$\begin{aligned} &\text{TV}(\mathcal{N}(x, \tilde{\Phi}^{-1}(x)), \mathcal{N}(z, \hat{\Phi}^{-1}(z))) \\ &\leq \text{TV}(\mathcal{N}(x, \tilde{\Phi}^{-1}(x)), \mathcal{N}(x, \Phi^{-1}(x))) \\ &\quad + \text{TV}(\mathcal{N}(x, \Phi^{-1}(x)), \mathcal{N}(z, \Phi^{-1}(z))) \\ &\quad + \text{TV}(\mathcal{N}(z, \Phi^{-1}(z)), \mathcal{N}(z, \hat{\Phi}^{-1}(z))). \end{aligned}$$

Second term is bounded using Lemma D.10. First and third term are bounded using similar ways. We have

$$\begin{aligned} &\text{TV}(\mathcal{N}(x, \tilde{\Phi}^{-1}(x)), \mathcal{N}(x, \Phi^{-1}(x)))^2 \\ &\leq \frac{1}{2} D(\mathcal{N}(x, \tilde{\Phi}^{-1}(x)) \| \mathcal{N}(x, \Phi^{-1}(x))) \\ &= \frac{1}{4} (\text{tr}[\tilde{\Phi}(x)^{-1/2} \Phi(x) \tilde{\Phi}(x)^{-1/2}] - d + \log \frac{\det \tilde{\Phi}(x)}{\det \Phi(x)}) \\ &\leq \frac{1}{4} ((1 \pm \epsilon_H)d - d + \log(1 \pm \epsilon_H)^d) \\ &\leq \frac{1}{2} d\epsilon_H. \end{aligned}$$

Thus, we complete the proof. \square

D.7 Reversibility and stationary distribution

For any $x \in \mathcal{K}$, we define the random variable Z_x to be the step taken by the Markov chain in Algorithm 2 from the point x , that is, set $z = x + \tilde{\Phi}(x)^{-1/2}\xi$ where $\xi \sim \mathcal{N}(0, I_d)$. If $z \in \mathcal{K}$, set $Z_x = z$ with probability

$$\frac{1}{2} \cdot \min \left\{ \frac{\exp(-f(z)) \cdot (\det(\tilde{\Phi}(z)))^{1/2} \cdot \exp(-0.5\|x - z\|_{\tilde{\Phi}(z)}^2)}{\exp(-f(x)) \cdot (\det(\tilde{\Phi}(x)))^{1/2} \cdot \exp(-0.5\|x - z\|_{\tilde{\Phi}(x)}^2)}, 1 \right\}.$$

Else, we set $z = x$.

We provide a modified version of Proposition E.12 in Mangoubi and Vishnoi [2023] where our Markov chain differs from theirs.

Lemma D.12 (Reversibility and stationary distribution). *For any $S_1, S_2 \subseteq \mathcal{K}$ we have that*

$$\int_{x \in S_1} \pi(x) \tilde{P}_x(S_2) dx = \int_{y \in S_2} \pi(y) \tilde{P}_y(S_1) dy.$$

Proof. Let a, b be two i.i.d. random variables (one can view a and b as random coins to generate the corresponding sparsifiers) such that $\tilde{\Phi}_{a,x}$ is a function of a, x and $\tilde{\Phi}_{b,y}$ is a function of b, y .

Let $p_{a,x}(\cdot)$ be pdf of $\mathcal{N}(0, \tilde{\Phi}_{a,x}^{-1})$ and $p_{b,y}(\cdot)$ be the pdf of $\mathcal{N}(0, \tilde{\Phi}_{b,y}^{-1})$. Because a and b are i.i.d, they are interchangeable, i.e., (a, b) has the same distribution as (b, a) .

$$\begin{aligned} & \int_{x \in S_1} \pi(x) \tilde{P}_x(S_2) dx \\ &= \int_{x \in S_1} \pi(x) \int_a q(a) \int_{y \in S_2} p_{a,x}(y) \int_b q(b) \cdot \min \left\{ \frac{\pi(y) p_{b,y}(x)}{\pi(x) p_{a,x}(y)}, 1 \right\} db dy da dx \\ &= \int_a q(a) \int_b q(b) \int_{x \in S_1} \int_{y \in S_2} \min \{ \pi(y) p_{b,y}(x), \pi(x) p_{a,x}(y) \} dy dx db da \\ &= \int_a q(a) \int_b q(b) \int_{x \in S_1} \int_{y \in S_2} \min \{ \pi(y) p_{a,y}(x), \pi(x) p_{b,x}(y) \} dy dx db da \\ &= \int_{y \in S_2} \int_a q(a) \int_{x \in S_1} \int_b q(b) \min \{ \pi(y) p_{a,y}(x), \pi(x) p_{b,x}(y) \} db dx da dy \\ &= \int_{y \in S_2} \pi(y) \tilde{P}_y(S_1) dy. \end{aligned}$$

The first and fifth steps are by definition of \tilde{P}_x . The second and fourth steps are by changing order of integration. The third step is by interchangeability of a and b .

This proves reversibility and stationary distribution. □

D.8 Lower bound on cross ratio distance

Definition D.13. *We define*

$$F(n, \alpha, \eta, R) := 0.1 / \sqrt{n\alpha^{-1} + \eta^{-1}R^2}.$$

For simplicity, we use F to denote $F(n, \alpha, \beta, R)$.

Lemma D.14. *Let S'_1 and S'_2 be defined as Definition C.8. Let F be defined as Definition D.13. We have*

$$\sigma(S'_1, S'_2) > 1000F.$$

Proof. Using Lemma D.17, we have that for any $u, v \in \mathcal{K}$,

$$\sigma(u, v) \geq F \cdot \|u - v\|_{\Phi(u)}.$$

We will try to prove $\sigma(u, v) > 1000F$ in the following. To do that, we will prove it by making a contradiction.

Suppose $\sigma(u, v) \leq 1000F$, then we have to require $\|u - v\|_{\Phi(u)} \leq 0.001$.

Once we had that $\|u - v\|_{\Phi(u)} \leq 0.001$, using Lemma D.9, for any $u, v \in \mathcal{K}$ we have that

$$\text{TV}(\tilde{P}_u, \tilde{P}_v) \leq 0.99 \quad (13)$$

On the other hand, Definition C.8 implies that, for any $u \in S'_1, v \in S'_2$ we have that

$$\text{TV}(\tilde{P}_u, \tilde{P}_v) \geq 1 - \tilde{P}_u(S_2) - \tilde{P}_v(S_1) \geq 1 - 2\beta > 0.99. \quad (14)$$

where the last step follows from $1 - 2\beta > 0.99$.

Thus, we obtain a contradiction, which means that

$$\sigma(S'_1, S'_2) \geq \sigma(u, v) > 1000F. \quad \square$$

D.9 Bounding cross ratio distance through \bar{v} -symmetry

In Lemma E.2 of Mangoubi and Vishnoi [2023], they only prove the cross ratio distance bound for log-barrier function. Here, we generalize it to arbitrary barrier function. The key concept our proof relies on is the notion of \bar{v} -symmetry (Assumption (i)).

Lemma D.15. *Let $H \in \mathcal{K} \rightarrow \mathbb{R}^{d \times d}$ be a \bar{v} -symmetric function, and $u, v \in \mathcal{K}$. Consider the chord that passes through u, v and intersects with \mathcal{K} . Let p, q be the two endpoints of the chord, with the order p, u, v, q . Then, we have*

$$\|p - u\|_{H(u)} \leq \sqrt{\bar{v}}.$$

Proof. Suppose we can show that $p \in \mathcal{K} \cap (2u - \mathcal{K})$, then it naturally follows that $p \in E_u(\sqrt{\bar{v}})$ and $\|p - u\|_{H(u)} \leq \sqrt{\bar{v}}$.

$p \in \mathcal{K}$ trivially follows by the construction of the chord. To see $p \in 2u - \mathcal{K}$, it is enough to show that for some $y \in \mathcal{K}$, we have $p = 2u - y$. We claim $y = 2u - p$ is such a choice. To see $2u - p \in \mathcal{K}$, we need to show that $2Au - Ap \leq b$. We partition the constraints into two sets.

Case 1. Suppose for a_i , we have $a_i^\top u \leq a_i^\top p$, then $2a_i^\top u \leq 2a_i^\top p \leq a_i^\top p + b_i$, therefore, $2a_i^\top u - a_i^\top p \leq b_i$.

Case 2. Suppose otherwise, $a_i^\top u > a_i^\top p$, then consider

$$\begin{aligned} 2a_i^\top u - a_i^\top p - b_i &= (a_i^\top u - a_i^\top p) + (a_i^\top u - b_i) \\ &= \underbrace{(a_i^\top u - b_i) - (a_i^\top p - b_i)}_{d_1} - \underbrace{(b_i - a_i^\top u)}_{d_2}. \end{aligned}$$

We shall show that $d_1 \leq d_2$ via a geometric argument. Consider the projection of both p and u onto the hyperplane $a_i x = b_i$, and the two chords that pass through p and u respectively, orthogonal to the hyperplane. Let us denote them with l_p, l_u respectively. Observe that l_p and l_u are parallel, and $\|l_p - l_u\|_2 = (a_i^\top u - b_i) - (a_i^\top p - b_i) = d_1$. It should then be obvious that $d_1 \leq d_2$: let $\text{proj}(u)$ denote the orthogonal projection of u onto the hyperplane, then

$$\underbrace{\|l_u - \text{proj}(u)\|_2}_{d_2} = \underbrace{\|l_u - l_p\|_2}_{d_1} + \|l_p - \text{proj}(p)\|_2$$

since $\|l_p - \text{proj}(p)\|_2 \geq 0$, we have that $d_1 \leq d_2$, and consequently,

$$2a_i^\top u - a_i^\top p \leq b_i.$$

This proves that $2u - p \in \mathcal{K}$, and we can conclude that $p \in 2u - \mathcal{K}$, which yields our desired result. \square

Remark D.16. *The above argument can be extended to spectrahedra and more general convex sets.*

Let $u, v \in \mathcal{K}$ be two arbitrary points, and consider the chord that passes through u, v with two endpoints being p, q . Note that $p, q \in \partial\mathcal{K}$.

Lemma D.17. For any $u, v \in \mathcal{K}$, let H be any matrix function satisfying Assumption (i), (ii) and (iii). We assume that \mathcal{K} is contained in ball of radius R and has nonempty interior.

For any parameter $\alpha \in (0, 1), \eta \in (0, 1)$, we define matrix function $\Phi(u) := \alpha^{-1} \cdot H(u) + \eta^{-1}I_d$.

Then, we have

$$\sigma(u, v) \geq \frac{1}{\sqrt{2\nu\alpha^{-1} + \eta^{-1}R^2}} \cdot \|u - v\|_{\Phi(u)}$$

Proof. Without loss of generality, we can assume that

$$\|p - u\|_2 \leq \|u - q\|_2. \tag{15}$$

We can lower bound $\sigma^2(u, v)$ as follows:

$$\begin{aligned} \sigma^2(u, v) &= \left(\frac{\|u - v\|_2^2 \cdot \|p - q\|_2^2}{\|p - u\|_2^2 \cdot \|v - q\|_2^2} \right) \\ &\geq \max\left\{ \frac{\|u - v\|_2^2}{\|p - u\|_2^2}, \frac{\|u - v\|_2^2}{\|v - q\|_2^2} \right\} \\ &\geq \max\left\{ \frac{\|u - v\|_2^2}{\|p - u\|_2^2}, \frac{\|u - v\|_2^2}{\|u - q\|_2^2} \right\} \\ &\geq \max\left\{ \frac{\|u - v\|_2^2}{\|p - u\|_2^2}, \frac{\|u - v\|_2^2}{\|u - q\|_2^2}, \frac{\|u - v\|_2^2}{\|p - q\|_2^2} \right\} \\ &= \max\left\{ \frac{\|u - v\|_2^2}{\|p - u\|_2^2}, \frac{\|u - v\|_2^2}{\|p - q\|_2^2} \right\} \\ &\geq \frac{1}{2} \frac{\|u - v\|_2^2}{\|p - u\|_2^2} + \frac{1}{2} \frac{\|u - v\|_2^2}{\|p - q\|_2^2} \\ &\geq \frac{1}{2} \frac{\|u - v\|_H^2}{\|p - u\|_H^2} + \frac{1}{2} \frac{\|u - v\|_2^2}{R^2} \\ &\geq \frac{1}{2} \frac{\|u - v\|_H^2}{\nu} + \frac{1}{2} \frac{\|u - v\|_2^2}{R^2} \\ &= (u - v)^\top \left(\frac{1}{2\nu\alpha^{-1}} \times \alpha^{-1}H + \frac{1}{2R^2\eta^{-1}} \times \eta^{-1}I_d \right) (u - v) \\ &\geq \frac{1}{2\nu\alpha^{-1} + 2\eta^{-1}R^2} \|u - v\|_{\Phi(u)}^2 \end{aligned}$$

where the second step is by $\|p - q\|_2 \geq \max\{\|p - u\|_2, \|v - q\|_2\}$, the third step is by $\|v - q\|_2 \leq \|u - q\|_2$, the fourth step is by $\|p - q\|_2 \geq \|p - u\|_2$, the fifth step follows from $\|p - u\|_2 \leq \|u - q\|_2$ (see Eq. (15)), sixth step is by $\max\{A, B\} \geq 0.5 \cdot (A + B)$, where the seventh step follows from $\|p - q\|_2 \leq R$ and Eq. (16). The eighth step is due to Lemma D.15.

It remains to show Eq. (16). Due to the fact that p, u, v are on the same chord, vector $u - v$ and $p - u$ are on the same direction. This means we can write $u - v = c \cdot (p - u)$ for some $c \in \mathbb{R}$, and

$$\begin{aligned} \frac{\|u - v\|_H}{\|p - u\|_H} &= \frac{\|c \cdot (p - u)\|_H}{\|p - u\|_H} \\ &= c \\ &= \frac{\|u - v\|_2}{\|p - u\|_2}. \end{aligned} \tag{16}$$

This completes the proof. □

E Computing Approximate Lewis Weights and Leverage Scores

In Section E.1, we present an algorithm to approximate the Hessian. In Section E.2, we show how to compute leverage score to low precision. In Section E.3, we show how to compute Lewis weights to high precision.

E.1 Subsampling to approximate the Hessian

Algorithm 3 Subsampling Algorithm

```

1: procedure SUBSAMPLE( $A \in \mathbb{R}^{n \times d}$ ,  $b \in \mathbb{R}^n$ ,  $x, n, d, \epsilon_H \in (0, 1)$ )
2:    $\epsilon_\sigma \leftarrow 0.001$ 
3:    $m \leftarrow \epsilon_H^{-2} d \log d$ 
4:   Form  $B \in \mathbb{R}^{n \times d}$  where  $i$ -th row of  $B$  is  $a_i \cdot (\langle a_i, x \rangle - b_i)^{-1}$ 
5:   Compute the  $O(1 \pm \epsilon_\sigma)$ -approximation to the leverage score of  $B$ 
6:   Sample a matrix  $\tilde{B} \in \mathbb{R}^{m \times d}$  such that  $(1 - \epsilon_H)B^\top B \preceq \tilde{B}^\top \tilde{B} \preceq (1 + \epsilon_H)B^\top B$ 
7:   return  $\tilde{B}^\top \tilde{B}$ 
8: end procedure

```

In this section, we provide a sparsification result for $H(w)$ using leverage score sampling.

To better monitor the whole process, it is useful to write $H(w)$ as $A^\top S(w)^{-2}A$, where $A \in \mathbb{R}^{n \times d}$ is the constraint matrix and $S(w)$ is a diagonal matrix with $S(w)_i = a_i^\top w - b_i$. The sparsification process is then sample the rows from the matrix $S(w)^{-1}A$.

Definition E.1. Let $B \in \mathbb{R}^{n \times d}$ be a full rank matrix. We define the leverage score of the i -th row of B as

$$\sigma_i(B) := b_i^\top (B^\top B)^{-1} b_i,$$

where b_i is the i -th row of B .

Definition E.2 (Sampling process). For any $w \in K$, let $H(w) = A^\top S(w)^{-2}A$. Let $p_i \geq \frac{\beta \cdot \sigma_i(S^{-1}(w)A)}{d}$, suppose we sample with replacement independently for s rows of matrix $S(w)^{-1}A$, with probability p_i of sampling row i for some $\beta \geq 1$. Let $i(j)$ denote the index of the row sampled in the j -th trial. Define the generated sampling matrix as

$$\tilde{H}(w) := \frac{1}{s} \sum_{j=1}^s \frac{1}{p_{i(j)}} \frac{a_{i(j)} a_{i(j)}^\top}{(a_{i(j)}^\top w - b_{i(j)})^2}.$$

Lemma E.3 (Sample using Matrix Chernoff). Let $\epsilon, \delta \in (0, 1)$ be precision and failure probability parameters, respectively. Suppose $\tilde{H}(w)$ is generated as in Definition E.2, then with probability at least $1 - \delta$, we have

$$(1 - \epsilon) \cdot H(w) \preceq \tilde{H}(w) \preceq (1 + \epsilon) \cdot H(w).$$

Moreover, the number of rows $s = \Theta(\beta \cdot \epsilon^{-2} d \log(d/\delta))$.

Proof. The proof will be designing a family of random matrices X . Let $y_i = (A^\top S(w)^{-2}A)^{-1/2} S(w)_{i,i}^{-1} \cdot a_i$ be the i -th sampled row and set $Y_i = \frac{1}{p_i} y_i y_i^\top$. Let $X_i = Y_i - I_d$. Note that

$$\begin{aligned} \sum_{i=1}^n y_i y_i^\top &= \sum_{i=1}^n (A^\top S(w)^{-2}A)^{-1/2} S(w)_{i,i}^{-2} \cdot a_i a_i^\top (A^\top S(w)^{-2}A)^{-1/2} \\ &= (A^\top S(w)^{-2}A)^{-1/2} \left(\sum_{i=1}^n S(w)_{i,i}^{-2} a_i a_i^\top \right) (A^\top S(w)^{-2}A)^{-1/2} \\ &= (A^\top S(w)^{-2}A)^{-1/2} (A^\top S(w)^{-2}A) (A^\top S(w)^{-2}A)^{-1/2} \\ &= I_d. \end{aligned} \tag{17}$$

Also, the norm of y_i connects directly to the leverage score:

$$\begin{aligned}\|y_i\|_2^2 &= S(w)_{i,i}^{-1} a_i^\top (A^\top S(w)^{-2} A)^{-1} S(w)_{i,i}^{-1} a_i \\ &= \sigma_i(S^{-1}A).\end{aligned}\tag{18}$$

We use $i(j)$ to denote the index of row that has been sampled during j -th trial.

Unbiased Estimator. Note that

$$\begin{aligned}\mathbb{E}[X] &= \mathbb{E}[Y] - I_d \\ &= \left(\sum_{i=1}^n p_i \cdot \frac{1}{p_i} y_i y_i^\top\right) - I_d \\ &= 0.\end{aligned}$$

Bound on $\|X\|$. To bound $\|X\|$, we provide a bound for any $\|X_i\|$ as follows:

$$\begin{aligned}\|X_i\| &= \|Y_i - I_d\| \\ &\leq 1 + \|Y_i\| \\ &= 1 + \frac{\|y_i y_i^\top\|}{p_i} \\ &\leq 1 + \frac{\|y_i\|_2^2}{\beta \cdot \sigma_i(S^{-1}(w)A)} \\ &= \frac{d}{\beta \cdot \sigma_i(S^{-1}(w)A)} \|(S(w)A)_i\|_2^2 + 1. \\ &= 1 + \frac{d}{\beta}.\end{aligned}$$

Bound on $\|\mathbb{E}[X^\top X]\|$. We compute the spectral norm of the covariance matrix:

$$\begin{aligned}\mathbb{E}[X_{i(j)}^\top X_{i(j)}] &= I_d + \mathbb{E}\left[\frac{y_{i(j)} y_{i(j)}^\top y_{i(j)} y_{i(j)}^\top}{p_i^2}\right] - 2 \mathbb{E}\left[\frac{y_{i(j)} y_{i(j)}^\top}{p_i}\right] \\ &= I_d + \left(\sum_{i=1}^n \frac{\sigma_i(S(w)^{-1}A)}{p_i} y_i y_i^\top\right) - 2I_d \\ &\leq \sum_{i=1}^n \frac{d}{\beta} y_i y_i^\top - I_d \\ &= \left(\frac{d}{\beta} - 1\right) I_d,\end{aligned}$$

the spectral norm is then

$$\|\mathbb{E}[X_{i(j)}^\top X_{i(j)}]\| \leq \frac{d}{\beta} - 1.$$

Put things together. Set $\gamma = 1 + \frac{d}{\beta}$ and $\sigma^2 = \frac{d}{\beta} - 1$, we apply Matrix Chernoff Bound as in Lemma A.8:

$$\begin{aligned}\Pr[\|W\| \geq \epsilon] &\leq 2d \cdot \exp\left(-\frac{s\epsilon^2}{d/\beta - 1 + (1 + d/\beta)\epsilon/3}\right) \\ &= 2d \cdot \exp(-s\epsilon^2 \cdot \Theta(\beta/d)) \\ &\leq \delta\end{aligned}$$

where we choose $s = \Theta(\beta \cdot \epsilon^{-2} d \log(d/\delta))$. Finally, we notice that

$$W = \frac{1}{s} \left(\sum_{j=1}^s \frac{1}{p_{i(j)}} y_{i(j)} y_{i(j)}^\top - I_d\right)$$

$$\begin{aligned}
&= (A^\top S(w)^{-2}A)^{-1/2} \left(\frac{1}{s} \sum_{j=1}^s \frac{1}{p_{i(j)}} \frac{a_{i(j)} a_{i(j)}^\top}{(a_{i(j)}^\top w - b_{i(j)})^2} \right) (A^\top S(w)^{-2}A)^{-1/2} - I_d \\
&= H(w)^{-1/2} \tilde{H}(w) H(w)^{-1/2} - I_d.
\end{aligned}$$

Therefore, we can conclude the desired result via $\|W\| \geq \epsilon$. \square

E.2 Computing leverage score

In this section, we provide an algorithm that approximates leverage scores in near-input sparsity time. Our algorithm makes use of the sparse embedding matrix of Nelson and Nguyen [2013] that has $O(\log(n/\epsilon))$ nonzero entries per column.

Lemma E.4 (Approximate leverage scores). *Let $A \in \mathbb{R}^{n \times d}$, there exists an algorithm that runs in time $\tilde{O}(\epsilon^{-1}(\text{nnz}(A) + \epsilon^{-1}d^\omega))$ that outputs a vector $\tilde{w}_2(A) \in \mathbb{R}^n$, such that $\tilde{w}_2(A) \approx_\epsilon w_2(A)$.*

Proof. We follow an approach of Woodruff [2014], but instead we use a high precision sparse embedding matrix Nelson and Nguyen [2013] and prove a two-sided bound on leverage score.

Let S be a sparse embedding matrix with $r = O(\epsilon^{-2}d \text{poly} \log(d/(\epsilon\delta)))$ rows and each column has $O(\epsilon^{-1} \log(nd/\epsilon))$ nonzero entries. We first compute SA in time $\tilde{O}(\epsilon^{-1} \text{nnz}(A))$, then compute the QR decomposition $SA = QR$ in time $\tilde{O}(\epsilon^{-2}d^\omega)$. Note that $R \in \mathbb{R}^{d \times d}$ hence R^{-1} can be computed in $O(d^\omega)$ time.

Now, let $G \in \mathbb{R}^{d \times t}$ matrix with $t = O(\epsilon^{-2} \log(n/\delta))$, each entry of G is i.i.d. $\mathcal{N}(0, 1/t)$ random variables. Set $q_i = \|e_i^\top AR^{-1}G\|_2^2$ for all $i \in [n]$. We argue q_i is a good approximation to $(w_2(A))_i$.

First, with failure probability at most δ/n , we have that $q_i \approx_\epsilon \|e_i^\top AR^{-1}\|_2^2$ via Johnson-Lindenstrauss lemma [Johnson and Lindenstrauss, 1984]. Now, it suffices to argue that $\|e_i^\top AR^{-1}\|_2^2$ approximates $\|e_i^\top U\|_2^2$ well, where $U \in \mathbb{R}^{n \times d}$ is the left singular vectors of A . To see this, first observe that for any $x \in \mathbb{R}^d$,

$$\begin{aligned}
\|AR^{-1}x\|_2^2 &= (1 \pm \epsilon) \cdot \|SAR^{-1}x\|_2^2 \\
&= (1 \pm \epsilon) \cdot \|Qx\|_2^2 \\
&= (1 \pm \epsilon) \cdot \|x\|_2^2,
\end{aligned}$$

where the last step is due to Q has orthonormal columns. This means that all singular values of AR^{-1} are in the range $[1 - \epsilon, 1 + \epsilon]$. Now, since U is an orthonormal basis for the column space of A , AR^{-1} and U has the same column space (since R is full rank). This means that there exists a change of basis matrix $T \in \mathbb{R}^{d \times d}$ with $AR^{-1}T = U$. Our goal is to provide a bound on all singular values of T . For the upper bound, we claim the largest singular value is at most $1 + 2\epsilon$, to see this, suppose for the contradiction that the largest singular is larger than $1 + 2\epsilon$ and let v be its corresponding (unit) singular vector. Since the smallest singular value of AR^{-1} is at least $1 - \epsilon$, we have

$$\begin{aligned}
\|AR^{-1}Tv\|_2^2 &\geq (1 - \epsilon)\|Tv\|_2^2 \\
&> (1 - \epsilon)(1 + 2\epsilon) \\
&> 1,
\end{aligned}$$

however, recall $AR^{-1}T = U$, therefore $\|AR^{-1}Tv\|_2^2 = \|Uv\|_2^2 = \|v\|_2^2 = 1$, a contradiction. One can similarly establish a lower bound of $1 - 2\epsilon$. Hence, the singular values of T are in the range of $[1 - 2\epsilon, 1 + 2\epsilon]$. This means that

$$\begin{aligned}
\|e_i^\top AR^{-1}\|_2^2 &= \|e_i^\top UT^{-1}\|_2^2 \\
&= (1 \pm 2\epsilon)\|e_i^\top U\|_2^2 \\
&= (1 \pm 2\epsilon)(w_2(A))_i,
\end{aligned}$$

as desired. Scaling ϵ to $\epsilon/2$ yields the approximation result.

Now, regarding the running time of computing q_i , note that we can first multiply R^{-1} with G in time $\tilde{O}(\epsilon^{-2}d^2)$, this gives a matrix of size $d \times t$. Multiplying this matrix with A takes $\tilde{O}(\epsilon^{-1} \text{nnz}(A))$ time. Hence, the overall time for computing $\tilde{w}_2(A) = q \in \mathbb{R}^n$ is $\tilde{O}(\epsilon^{-1} \text{nnz}(A) + \epsilon^{-2}d^\omega)$. \square

Remark E.5. It suffices to choose $\epsilon = \frac{1}{10}$ for generating approximate leverage scores and subsampling the target matrix $H(x)$.

E.3 Computing Lewis weights

Before stating the main technical tool for this section, we want to make several remarks regarding the Lewis weights.

Since its introduction by Cohen and Peng [Cohen and Peng, 2015] as an algorithmic tool for ℓ_p row sampling, Lewis weights, as its natural formulation follows from a convex program, is *not known to be solved exactly in polynomial time*. All known algorithms [Cohen and Peng, 2015, Lee and Sidford, 2019, Fazel et al., 2022, Cohen et al., 2019, Jambulapati et al., 2022] can only compute ϵ -approximate Lewis weights.

Definition E.6 (ℓ_p Lewis weights). Let $A \in \mathbb{R}^{n \times d}$. The ℓ_p Lewis weights of A is a vector $w_p(A) \in \mathbb{R}^n$ satisfying

$$w_p(A) = \sigma(W_p(A)^{\frac{1}{2} - \frac{1}{p}} A),$$

where $W_p(A) \in \mathbb{R}^{n \times n}$ is the diagonal matrix that puts $w_p(A)$ on the diagonal, $\sigma : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^n$ is the operation that outputs all leverage scores of input matrix.

We are now posed to state a main tool from Lee and Sidford [2019].

Theorem E.7 (Theorem 46 in Lee and Sidford [2019]). Let $P = \{x : Ax > b\}$ denote the interior of non-empty polytope for non-degenerate $A \in \mathbb{R}^{n \times d}$. There is an $O(d \log^5 n)$ -self concordant barrier ψ defined by ℓ_p Lewis weight with $p = \Theta(\log n)$ satisfying

$$A_x^\top W_x A_x \preceq \nabla^2 \psi(x) \preceq (q+1) A_x^\top W_x A_x$$

where $A_x = \text{diag}(Ax - b)$ and w_x is the ℓ_p Lewis weight of the matrix A_x . Furthermore, we can compute or update the w_x , $\nabla \psi(x)$ and $\nabla^2 \psi(x)$ as follows:

- **Initial Weight:** For any $x \in \mathbb{R}^d$, one can compute a vector \tilde{w}_x such that $(1 - \epsilon)w_x \leq \tilde{w}_x \leq (1 + \epsilon)w_x$ in $O(nd^{\omega-1/2} \cdot \log^3 n \log(n/\epsilon))$ time.
- **Update Weight and Compute Gradient/Hessian:** Given a vector \tilde{w}_x such that $\tilde{w}_x = (1 \pm \frac{1}{100})w_x$ for any y with $\|x - y\|_{A_x^\top W_x A_x} \leq \frac{c}{\log^2 n}$ for some small constant $c > 0$, we can compute \tilde{w}_y , v and H such that $\tilde{w}_y = (1 \pm \epsilon)w_y$,

$$\|v - \nabla \psi(x)\|_{\nabla^2 \psi(x)^{-1}} \leq \epsilon, \quad \text{and} \quad (1 - \epsilon)\nabla^2 \psi(x) \preceq H \preceq (1 + \epsilon)\nabla^2 \psi(x)$$

in $O(nd^{\omega-1} \cdot \log n \cdot \log(n/\epsilon))$ time.

F Fast Approximate Sampling from Polytopes: Complexity

In this section, we provide the runtime analysis for sampling from polytopes via log-barrier and Lee-Sidford barrier. The result for log-barrier is in Section F.1 and for Lee-Sidford barrier is in Section F.2.

F.1 Log-barrier in nearly-linear time

Throughout this section, we let $\alpha \in (0, 1/(10^5 d))$ and $\eta \in (0, 1/(20dL^2))$.

Lemma F.1. If g is log-barrier function with $\nu = n$, then each iteration of Algorithm 2 can be implemented in $O(\text{nnz}(A) + T_{\text{mat}}(d, d^3, d))$ time plus $O(1)$ calls to the oracle of f .

Proof. We go through each step of Algorithm 2 and add up the time and oracle calls for each step:

We first need to sample a d -dimensional Gaussian random vector $\xi \sim N(0, I_d)$, which can be performed in $O(d)$ time.

At each iteration, by the definition of log-barrier function, we can write the Hessian as follows:

$$H(x) := \sum_{j=1}^n \frac{a_j a_j^\top}{(a_j^\top x - b_j)^2}.$$

We define matrix $B \in \mathbb{R}^{n \times d}$ as follows

$$B := S^{-1}A$$

where $s_i = (\langle a_i, x \rangle - b_i)$, $\forall i \in [n]$. Then we can write $H(x) = B^\top B \in \mathbb{R}^{d \times d}$.

For matrix B , we can compute a constant approximation of its leverage score, then we can sample according to leverage, and generate a diagonal matrix D such that

$$(1 - \epsilon)B^\top B \preceq B^\top D D B \preceq (1 + \epsilon)B^\top B.$$

This step takes $\tilde{O}(\text{nnz}(A) + d^{\omega(1,3,1)})$ time.

Computing $\tilde{\Phi}(x)$ can be done in d^2 time and computing the vector z

$$z = x + \tilde{\Phi}(x)^{-\frac{1}{2}} \xi$$

requires to invert the matrix $\tilde{\Phi}(x)$, taking the square root and multiplying with a vector ξ . Inversion and square root can be performed in time $O(d^\omega)$ by computing its spectral decomposition, and the matrix-vector product can be done in time $O(d^2)$. Hence, the overall time of this step is $O(d^\omega)$.

To determine whether $z \in K$, one can simply verify $Az \leq b$, in time $O(\text{nnz}(A))$.

We also need to compute $\hat{H}(z)$ and $\hat{\Phi}(z)$. This is similar to computing $\tilde{H}(x)$ and $\tilde{\Phi}(x)$.

We then need to compute the determinant $\det(\tilde{\Phi}(x))$ and $\det(\hat{\Phi}(z))$ which can be done in $O(d^\omega)$ time via a spectral decomposition.

We then need $O(1)$ oracle queries to the function values of f .

Therefore, adding up the number of arithmetic operations and oracle calls from all the different steps of Algorithm 2, we get that each iteration of Algorithm 2 can be computed in $\tilde{O}(\text{nnz}(A) + \mathcal{T}_{\text{mat}}(d, d^3, d))$ arithmetic operations plus $O(1)$ calls to the oracle of f . \square

Remark F.2. We note a recent work *Mangoubi and Vishnoi [2024]* has provided an algorithm that runs in $\text{nnz}(A) + d^2$ time per iteration, by adapting techniques from *Laddha et al. [2020]*. Their per iteration cost improvement only works for log-barrier, and they could only manage to obtain an $\tilde{O}(nd + dL^2 R^2)$ mixing time. Our framework on the other hand generalizes to other barriers such as Lee-Sidford barrier, which provides a nearly-optimal mixing rate. We also adapt the framework to sampling from spectrahedra.

F.2 Lee-Sidford barrier via approximation scheme

Lemma F.3. If g is the Lee-Sidford barrierfunction with $\nu = d \log^5 n$, then, each iteration of Algorithm 2 can be implemented in $\tilde{O}(nd^{\omega-1})$ time plus $O(1)$ calls to the oracle of f .

Proof. For Lee-Sidford barrier, it suffices to use $H(x) = A_x^\top W_x A_x$ as where W_x is the ℓ_p Lewis weights for $p = \Theta(\log m)$. However, as Lewis weights cannot be computed exactly, we shall use the algorithm of *Lee and Sidford [2019]* to compute an ϵ -approximation.

To invoke Theorem E.7, we need to make sure that z and x satisfy $\|z - x\|_{A_x^\top W_x A_x} \leq \frac{\epsilon}{\log^2 n}$. We can achieve so by scaling down α by a factor of $\log^2 n$. This only blows up the convergence by $\tilde{O}(1)$ factor, so it is acceptable.

The runtime analysis is then similar to Lemma F.1. Computing the initial Lewis weights takes $\tilde{O}(nd^{\omega-1/2} \log(1/\epsilon))$ time owing to Theorem E.7. As the algorithm requires $\tilde{O}(d^2)$ iterations, this time can be amortized. For each iteration, it needs to query the Lewis weights data structure that outputs an ϵ -Lewis weights in $\tilde{O}(nd^{\omega-1} \log(1/\epsilon))$ time. As we will choose ϵ as $O(1/d)$, this time is $\tilde{O}(nd^{\omega-1})$. We can then use this approximate Hessian to progress the algorithm. All subsequent operations take $O(d^\omega)$ time. Thus, the cost per iteration is $\tilde{O}(nd^{\omega-1})$. \square

G Approximate Sampling from a Spectrahedron

We present an algorithm that efficiently and approximately samples from a spectrahedron, which is a popular convex body utilized by semidefinite programs.

G.1 Definitions

Definition G.1. Let $A_1, \dots, A_d \in \mathbb{R}^{n \times n}$ be a collection of symmetric matrices and $C \in \mathbb{R}^{n \times n}$ be symmetric. We define the corresponding spectrahedron as

$$\mathcal{K} = \{x \in \mathbb{R}^d : \sum_{i=1}^d x_i A_i \succeq C\}.$$

We define the log-barrier for spectrahedron and its corresponding Hessian.

Definition G.2 (Nesterov and Nemirovskii [1994]). Let \mathcal{K} be a spectrahedron described by symmetric matrices $\{A_1, \dots, A_d\} \subseteq \mathbb{R}^{n \times n}$ and $C \in \mathbb{R}^{n \times n}$. The log barrier $\phi_{\log}(x)$ is defined as

$$\phi_{\log}(x) = -\log \det(S(x))$$

and its corresponding Hessian is defined as

$$H_{\log}(x) = A(S(x)^{-1} \otimes S(x)^{-1})A^\top,$$

where $S(x) := \sum_{i=1}^d x_i A_i - C \in \mathbb{R}^{n \times n}$ and $A \in \mathbb{R}^{d \times n^2}$ and the i -th row of A is the vectorization of $A_i \in \mathbb{R}^{n \times n}$.

To compute $H_{\log}(x)$, it is handy to define a matrix $B \in \mathbb{R}^{d \times n^2}$.

Definition G.3. We define a matrix $B \in \mathbb{R}^{d \times n^2}$ as $B = A(S(x)^{-1/2} \otimes S(x)^{-1/2})$. Consequently, $H_{\log}(x) = BB^\top$.

We state a useful lemma for computing the matrix B .

Lemma G.4. Let matrix $B \in \mathbb{R}^{d \times n^2}$ be defined as in Def. G.3. Then, the i -th row of B can be computed as $\text{vec}(S(x)^{-1/2} A_i S(x)^{-1/2})$.

Proof. The proof relies on a simple fact of Kronecker product and vectorization:

$$\text{vec}(S(x)^{-1/2} A_i S(x)^{-1/2}) = (S(x)^{-1/2} \otimes S(x)^{-1/2}) \text{vec}(A_i),$$

which is the definition of the i -th row of B . □

G.2 n -symmetry of H_{\log} for spectrahedra

In this section, we prove that H_{\log} is n -symmetry for spectrahedra.

Lemma G.5. $H_{\log}(x)$ is n -symmetry, that is, for any $x \in \mathcal{K}$,

$$E_x(1) \subseteq \mathcal{K} \cap (2x - \mathcal{K}) \subseteq E_x(\sqrt{n}).$$

Proof. Pick a point $y \in E_x(1)$, then we know that $(y - x)^\top A(S(x)^{-1} \otimes S(x)^{-1})A^\top (y - x) \leq 1$. We note that the set $\mathcal{K} \cap (2x - \mathcal{K})$ can be characterized as the set of all points $u \in \mathbb{R}^d$ such that

$$\begin{aligned} \sum_{i=1}^d u_i A_i &\succeq C, \\ \sum_{i=1}^d (2x - u)_i A_i &\succeq C. \end{aligned}$$

These two conditions imply

$$-S(x) \preceq \sum_{i=1}^d (u_i - x_i) A_i \preceq S(x),$$

$$-I_n \preceq \sum_{i=1}^d (u_i - x_i) S(x)^{-1/2} A_i S(x)^{-1/2} \preceq I_n.$$

We will prove y satisfies these conditions. Note that $y \in E_x(1)$ implies that

$$\begin{aligned} \left\| \sum_{i=1}^d (y_i - x_i) \text{vec}(S(x)^{-1/2} A_i S(x)^{-1/2}) \right\|_2^2 &\leq 1, \\ \left\| \sum_{i=1}^d (y_i - x_i) S(x)^{-1/2} A_i S(x)^{-1/2} \right\|_F^2 &\leq 1 \end{aligned}$$

if we let M denote $\sum_{i=1}^d (y_i - x_i) S(x)^{-1/2} A_i S(x)^{-1/2}$, then the above condition implies that $M^\top M \preceq I_n$, meaning all the eigenvalues of M lie between $[-1, 1]$, thus we have shown that $E_x(1) \subseteq \mathcal{K} \cap (2x - \mathcal{K})$.

For the other direction, we know that

$$-I_n \preceq \sum_{i=1}^d (y_i - x_i) S(x)^{-1/2} A_i S(x)^{-1/2} \preceq I_n,$$

and henceforth

$$\begin{aligned} &(y - x)^\top \mathbf{A} (S(x)^{-1} \otimes S(x)^{-1}) \mathbf{A}^\top (y - x) \\ &= \left\| \sum_{i=1}^d (y_i - x_i) S(x)^{-1/2} A_i S(x)^{-1/2} \right\|_F^2 \\ &\leq \|I_n\|_F^2 \\ &= n, \end{aligned}$$

where we use the fact that $-I_n \preceq M \preceq I_n$ hence all eigenvalues of M lie between $[-1, 1]$, and the squared eigenvalues of M have its magnitude at most 1. We then use the squared Frobenius norm is equivalent to squared ℓ_2 norm of singular values of M , as desired. \square

G.3 Bounded local norm of $H_{\log}(x)$

We prove that $H_{\log}(x)$ has the bounded local norm property.

Lemma G.6. *Let $H_{\log}(x) \in \mathbb{R}^{d \times d}$ be defined as in Def. G.2. Let $F(x) = \log \det(H_{\log}(x))$. Then, $F(x)$ is convex in x .*

Proof. The proof is by observe that the function $F(x)$ is actually the volumetric barrier function for the SDP spectrahedron. Thus, the function F is convex, for more details, see [Nesterov and Nemirovskii \[1994\]](#). \square

Lemma G.7. *Let $F(x)$ be defined as in Lemma G.6. Then, we have*

$$\nabla F(x) = \begin{bmatrix} -\text{tr}[\mathbf{P}((S(x)^{-1/2} A_1 S(x)^{-1/2}) \otimes_S (S(x)^{-1/2} A_1 S(x)^{-1/2}))] \\ -\text{tr}[\mathbf{P}((S(x)^{-1/2} A_2 S(x)^{-1/2}) \otimes_S (S(x)^{-1/2} A_2 S(x)^{-1/2}))] \\ \vdots \\ -\text{tr}[\mathbf{P}((S(x)^{-1/2} A_m S(x)^{-1/2}) \otimes_S (S(x)^{-1/2} A_m S(x)^{-1/2}))] \end{bmatrix}$$

where $\mathbf{P} = \mathbf{B}^\top H(x)^{-1} \mathbf{B}$.

Proof. For simplicity, we let H to denote H_{\log} . For each coordinate $j \in [d]$, we have

$$\begin{aligned} \frac{\partial F}{\partial x_j} &= \text{tr}[H(x)^{-1} \frac{\partial H(x)}{\partial x_j}] \\ &= \text{tr}[H(x)^{-1} \mathbf{A} \frac{\partial (S(x)^{-1} \otimes S(x)^{-1})}{\partial x_j} \mathbf{A}^\top] \end{aligned}$$

$$\begin{aligned}
&= \text{tr}[H(x)^{-1}A(\frac{\partial S(x)^{-1}}{\partial x_j} \otimes S(x)^{-1} + S(x)^{-1} \otimes \frac{\partial S(x)^{-1}}{\partial x_j})A^\top] \\
&= -\text{tr}[H(x)^{-1}A((S(x)^{-1}\frac{\partial S(x)}{\partial x_j}S(x)^{-1}) \otimes S(x)^{-1} + S(x)^{-1} \otimes (S(x)^{-1}\frac{\partial S(x)}{\partial x_j}S(x)^{-1}))A^\top] \\
&= -\text{tr}[H(x)^{-1}A((S(x)^{-1}A_jS(x)^{-1}) \otimes S(x)^{-1} + S(x)^{-1} \otimes (S(x)^{-1}A_jS(x)^{-1}))A^\top],
\end{aligned}$$

examine the following term:

$$\begin{aligned}
&\text{tr}[H(x)^{-1}A((S(x)^{-1}A_jS(x)^{-1}) \otimes S(x)^{-1})A^\top] \\
&= \text{tr}[A^\top H(x)^{-1}A((S(x)^{-1}A_j) \otimes I)(S(x)^{-1} \otimes S(x)^{-1})] \\
&= \text{tr}[(S(x)^{-1/2} \otimes S(x)^{-1/2})A^\top H(x)^{-1}A(S(x)^{-1/2} \otimes S(x)^{-1/2}) \\
&\quad \cdot ((S(x)^{-1/2}A_j) \otimes S(x)^{1/2})(S(x)^{-1/2} \otimes S(x)^{-1/2})] \\
&= \text{tr}[P((S(x)^{-1/2}A_j) \otimes S(x)^{1/2})(S(x)^{-1/2} \otimes S(x)^{-1/2})] \\
&= \text{tr}[P((S(x)^{-1/2}A_jS(x)^{-1/2}) \otimes I_n)],
\end{aligned}$$

where $P \in \mathbb{R}^{n^2 \times n^2}$ is the projection of B^\top , i.e., $P = B^\top H(x)^{-1}B$. Thus,

$$\nabla F(x) = \begin{bmatrix} -\text{tr}[P((S(x)^{-1/2}A_1S(x)^{-1/2}) \otimes_S (S(x)^{-1/2}A_1S(x)^{-1/2}))] \\ -\text{tr}[P((S(x)^{-1/2}A_2S(x)^{-1/2}) \otimes_S (S(x)^{-1/2}A_2S(x)^{-1/2}))] \\ \vdots \\ -\text{tr}[P((S(x)^{-1/2}A_mS(x)^{-1/2}) \otimes_S (S(x)^{-1/2}A_mS(x)^{-1/2}))] \end{bmatrix}.$$

This completes the proof. \square

Lemma G.8. *The matrix function $H(x) = A(S^{-1} \otimes S^{-1})A^\top \in \mathbb{R}^{d \times d}$ has the following bounded norm property: for any direction $h \in \mathbb{R}^d$,*

$$\|H(x)^{-1/2}DH(x)[h]H(x)^{-1/2}\|_F^2 \leq 4d\|h\|_{H(x)}^2.$$

Proof. Let $x_t = x + t \cdot h$ for some fixed vector $h \in \mathbb{R}^d$, $S_t = \sum_{i=1}^m x_{t,i}A_i - C \in \mathbb{R}^{n \times n}$, $A_t = A(S_t^{-1/2} \otimes S_t^{-1/2}) \in \mathbb{R}^{d \times n^2}$, $P_t = A_t^\top (A_t A_t^\top)^{-1} A_t \in \mathbb{R}^{n^2 \times n^2}$ and $\Sigma_t = \sum_{i=1}^n (e_i^\top \otimes I_n)P_t(e_i \otimes I_n) = \sum_{i=1}^n (I_n \otimes e_i^\top)P_t(I_n \otimes e_i)$. Note that for any matrix $M \in \mathbb{R}^{n \times n}$, $\text{tr}[P_t \cdot (M \otimes I_n)] = \text{tr}[P_t \cdot (I_n \otimes M)] = \text{tr}[\Sigma_t \cdot M]$. Also note that $\text{tr}[\Sigma_t] = \text{tr}[P_t] = d$.

Let $H_t = A^\top (S_t^{-1} \otimes S_t^{-1})A \in \mathbb{R}^{d \times d}$. We have

$$\begin{aligned}
&\|H_t^{-1/2}(\frac{\partial}{\partial t}H_t)H_t^{-1/2}\|_F^2 \\
&= \text{tr}[H_t^{-1} \cdot (\frac{\partial}{\partial t}H_t) \cdot H_t^{-1} \cdot (\frac{\partial}{\partial t}H_t)] \\
&= \text{tr}[H_t^{-1} \cdot A^\top \frac{\partial S_t^{-1} \otimes S_t^{-1}}{\partial t} A \cdot H_t^{-1} A^\top \frac{\partial (S_t^{-1} \otimes S_t^{-1})}{\partial t} A] \\
&= \text{tr}[A H_t^{-1} A^\top \cdot \frac{\partial S_t^{-1} \otimes S_t^{-1}}{\partial t} \cdot A H_t^{-1} A^\top \cdot \frac{\partial (S_t^{-1} \otimes S_t^{-1})}{\partial t}] \\
&= \text{tr}[P_t (S_t^{-1/2} \otimes S_t^{-1/2})^{-1} \frac{\partial S_t^{-1} \otimes S_t^{-1}}{\partial t} (S_t^{-1/2} \otimes S_t^{-1/2})^{-1} \\
&\quad \cdot P_t (S_t^{-1/2} \otimes S_t^{-1/2})^{-1} \frac{\partial S_t^{-1} \otimes S_t^{-1}}{\partial t} (S_t^{-1/2} \otimes S_t^{-1/2})^{-1}] \\
&\leq \text{tr}[P_t \cdot (S_t^{-1/2} \otimes S_t^{-1/2})^{-1} \frac{\partial S_t^{-1} \otimes S_t^{-1}}{\partial t} (S_t^{-1} \otimes S_t^{-1})^{-1} \frac{\partial S_t^{-1} \otimes S_t^{-1}}{\partial t} (S_t^{-1/2} \otimes S_t^{-1/2})^{-1}] \\
&= \text{tr}[P_t \cdot (S_t^{-1/2} (\sum_{i=1}^d h_i A_i) S_t^{-1/2} \otimes_S I_n)^2]
\end{aligned}$$

$$\begin{aligned}
&= \text{tr}[P_t \cdot ((S_t^{-1/2}(\sum_{i=1}^d h_i A_i)S_t^{-1/2})^2 \otimes_S I_n + (S_t^{-1/2}(\sum_{i=1}^d h_i A_i)S_t^{-1/2} \otimes S_t^{-1/2}(\sum_{i=1}^d h_i A_i)S_t^{-1/2}))] \\
&= 2 \text{tr}[S_t^{-1/2}(\sum_{i=1}^d h_i A_i)S_t^{-1/2} \cdot \Sigma_t \cdot S_t^{-1/2}(\sum_{i=1}^d h_i A_i)S_t^{-1/2}] \\
&+ \text{tr}[P_t \cdot (S_t^{-1/2}(\sum_{i=1}^d h_i A_i)S_t^{-1/2} \otimes S_t^{-1/2}(\sum_{i=1}^d h_i A_i)S_t^{-1/2})] \\
&\leq 4 \text{tr}[S_t^{-1/2}(\sum_{i=1}^d h_i A_i)S_t^{-1/2} \cdot \Sigma_t \cdot S_t^{-1/2}(\sum_{i=1}^d h_i A_i)S_t^{-1/2}] \\
&\leq 4\|\Sigma_t\| \text{tr}[(S_t^{-1/2}(\sum_{i=1}^d h_i A_i)S_t^{-1/2})^2] \\
&\leq 4 \text{tr}[\Sigma_t] \text{tr}[(S_t^{-1/2}(\sum_{i=1}^d h_i A_i)S_t^{-1/2})^2] \\
&= 4d\|h\|_{H_t}^2,
\end{aligned}$$

the fifth step is by for PSD matrix M , we have $MPM \preceq M^2$ where P is a projection matrix, the sixth step is by the below derivation, the seventh step is by $\text{tr}[P_t \cdot (M^2 \otimes_S I_n)] = 2 \text{tr}[M\Sigma_t M]$, the eighth step is by for PSD matrix M , $2M \otimes_S I_n \succeq M \otimes M$, the ninth step is by Hölder's inequality where $\langle A, B \rangle \leq \|A\| \cdot \|B\|_{s_1}$ for $\|B\|_{s_1}$ be its Schatten-1 norm. If B is PSD, $\|B\|_1 = \text{tr}[B]$, the tenth step is by for PSD matrix A , $\|A\| \leq \text{tr}[A]$.

Note that

$$\begin{aligned}
\frac{\partial S_t^{-1}}{\partial t} &= -S_t^{-1} \frac{\partial S_t}{\partial t} S_t^{-1} \\
&= -S_t^{-1} (\sum_{i=1}^d h_i A_i) S_t^{-1}
\end{aligned}$$

Thus,

$$\begin{aligned}
\frac{\partial S_t^{-1} \otimes S_t^{-1}}{\partial t} &= -S_t^{-1} (\sum_{i=1}^d h_i A_i) S_t^{-1} \otimes S_t^{-1} \\
&\quad - S_t^{-1} \otimes S_t^{-1} (\sum_{i=1}^d h_i A_i) S_t^{-1}
\end{aligned}$$

and we have

$$\begin{aligned}
&(S_t^{-1/2} \otimes S_t^{-1/2})^{-1} \frac{\partial S_t^{-1} \otimes S_t^{-1}}{\partial t} (S_t^{-1} \otimes S_t^{-1})^{-1} \frac{\partial S_t^{-1} \otimes S_t^{-1}}{\partial t} (S_t^{-1/2} \otimes S_t^{-1/2})^{-1} \\
&= (S_t^{-1/2} \otimes S_t^{-1/2})^{-1} (S_t^{-1} (\sum_{i=1}^d h_i A_i) S_t^{-1} \otimes S_t^{-1} + S_t^{-1} \otimes S_t^{-1} (\sum_{i=1}^d h_i A_i) S_t^{-1}) (S_t^{-1/2} \otimes S_t^{-1/2})^{-1} \\
&\quad \cdot (S_t^{-1/2} \otimes S_t^{-1/2})^{-1} (S_t^{-1} (\sum_{i=1}^d h_i A_i) S_t^{-1} \otimes S_t^{-1} + S_t^{-1} \otimes S_t^{-1} (\sum_{i=1}^d h_i A_i) S_t^{-1}) (S_t^{-1/2} \otimes S_t^{-1/2})^{-1} \\
&= (S_t^{-1/2} (\sum_{i=1}^d h_i A_i) S_t^{-1/2} \otimes_S I_n)^2,
\end{aligned}$$

plug in the above derivation we obtain the desired result. Finally, we send $t \rightarrow 0$ and conclude the proof. \square

Lemma G.9. Let $H(x)$ be the matrix function $A(S^{-1} \otimes S^{-1})A^\top$, then we have

$$\|H(x)^{-1/2} \nabla F(x)\|_2 \leq \tilde{O}(d).$$

Proof. We have

$$\begin{aligned}
 \|H(x)^{-1/2}\nabla F(x)\|_2 &= \max_{v:\|v\|_2=1} (H(x)^{-1/2}\nabla F(x))^\top v \\
 &= \max_{v:\|v\|_2=1} \text{tr}[H(x)^{-1}DH(x)[H(x)^{-1/2}v]] \\
 &= \max_{u:\|u\|_{H(x)}=1} \text{tr}[H(x)^{-1/2}DH(x)[u]H(x)^{-1/2}] \\
 &\leq \max_{u:\|u\|_{H(x)}=1} \sqrt{d}\|H(x)^{-1/2}DH(x)[u]H(x)^{-1/2}\|_F \\
 &\leq 2d \cdot \|u\|_{H(x)} \\
 &= 2d,
 \end{aligned}$$

where the fourth step is by $\text{tr}[M] \leq \sqrt{d} \cdot \|M\|_F$ for $d \times d$ PSD matrix M , the fifth step follows from Lemma G.8, the last step follows from $\|u\|_{H(x)} = 1$. \square

G.4 Fast Hessian approximation

We need a particular type of sketch for Kronecker product of matrices.

Definition G.10 (TensorSRHT [Song et al., 2021]). *The TensorSRHT $\Pi : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^s$ is defined as $S := \frac{1}{\sqrt{s}}P \cdot (HD_1 \otimes HD_2)$, where each row of $P \in \{0, 1\}^{s \times n^2}$ contains only one 1 at a random coordinate and one can view P as a sampling matrix. H is a $n \times n$ Hadamard matrix, and D_1, D_2 are two $n \times n$ independent diagonal matrices with diagonals that are each independently set to be a Rademacher random variable (uniform in $\{-1, 1\}$).*

Lemma G.11 (Lemma 2.12 in Song et al. [2021]). *Let Π be a TensorSRHT matrix defined in Definition G.10. If $s = O(\epsilon^{-2}d \log^3(nd/(\epsilon\delta)))$, then for any orthonormal basis $U \in \mathbb{R}^{n^2 \times d}$, we have that with probability at least $1 - \delta$, the singular values of ΠU lie in the range of $[1 - \epsilon, 1 + \epsilon]$.*

The following result provides an efficient embedding for B .

Lemma G.12. *Let $B \in \mathbb{R}^{d \times n^2}$ be defined as in Def. G.3, $\epsilon \in (0, 1/10)$ denote an accuracy parameter and $\delta \in (0, 1/10)$ denote a failure probability.*

Let $\Pi \in \mathbb{R}^{s \times n^2}$ be a TensorSRHT matrix with $s = \Theta(\epsilon^{-2}d \log^3(nd/(\epsilon\delta)))$, then we have

$$\Pr[\|\Pi B^\top x\|_2 = (1 \pm \epsilon)\|B^\top x\|_2, \forall x \in \mathbb{R}^d] \geq 1 - \delta.$$

Moreover, ΠB^\top can be computed in time

$$O(d \cdot \mathcal{T}_{\text{mat}}(n, n, s)).$$

Proof. The correctness part follows directly from Lemma G.11. It remains to argue for the running time. We need to unravel the construction of both S and B (see Definition G.2). Recall that $\Pi = \frac{1}{\sqrt{s}}P \cdot (HD_1 \otimes HD_2)$ and

$$\begin{aligned}
 &\Pi(S(x)^{-1/2} \otimes S(x)^{-1/2}) \text{vec}(A_i) \\
 &= \frac{1}{\sqrt{s}}P \cdot (HD_1 \otimes HD_2) \cdot (S(x)^{-1/2} \otimes S(x)^{-1/2}) \text{vec}(A_i) \\
 &= \frac{1}{\sqrt{s}}P \cdot (HD_1 S(x)^{-1/2} \otimes HD_2 S(x)^{-1/2}) \text{vec}(A_i) \\
 &= \frac{1}{\sqrt{s}}P \cdot \text{vec}(HD_2 S(x)^{-1/2} A_i S(x)^{-1/2} D_1 H^\top)
 \end{aligned}$$

since P is a row sampling matrix, the product can be computed as follows:

- First compute $HD_1 S(x)^{-1/2}$ and $HD_2 S(x)^{-1/2}$. Since H is a Hadamard matrix, this step can be carried out in $O(n^2 \log n)$ time.

- Applying P to the vector can be interpreted as sampling s coordinates from the matrix $HD_2S(x)^{-1/2}A_iS(x)^{-1/2}D_1H^\top$. Let $(i_1, j_1), \dots, (i_s, j_s)$ denote the coordinates sampled by P . We form two matrices $X, Y \in \mathbb{R}^{s \times n}$ where the k -th row of X is $(HD_2S(x)^{-1/2})_{i_k, *}$ and the k -th row of Y is $(HD_1S(x)^{-1/2})_{j_k, *}$. It is easy to verify that the (i_k, j_k) -th entry of $XA_iY^\top \in \mathbb{R}^{s \times s}$ is the corresponding entry of $HD_2S(x)^{-1/2}A_iS(x)^{-1/2}D_1H^\top$. This step therefore takes $\mathcal{T}_{\text{mat}}(n, n, s)$ time.

As we need to apply the second step to d rows, the total runtime is

$$O(d \cdot \mathcal{T}_{\text{mat}}(n, n, s)). \quad \square$$

Remark G.13. For comparison, compute B straightforwardly using the Kronecker product identity takes $O(dn^\omega)$ time. As long as the row count s is independent of n , Lemma G.12 approximately computes B in time $\tilde{O}(n^2 \cdot \text{poly}(d))$. For the regime where $d \ll n$, our algorithm is much more efficient.

H Convexity of Regularized Log-Barrier Function

In this section, we prove the convexity of $\log \det(H_{\log}(x) + I_d)$ for both polytopes and spectrahedra.

H.1 Convexity of $\log \det(H_{\log}(x) + I_d)$: polytopes

We prove that for $H(x) = \sum_{i=1}^n \frac{a_i a_i^\top}{(a_i^\top x - b_i)^2}$, the Hessian of the log barrier, we have the convexity.

Definition H.1 (Ridge leverage score). Let $B \in \mathbb{R}^{n \times d}$, we define the λ -ridge leverage score of B as

$$\tilde{\sigma}_{\lambda, i}(B) := b_i^\top (B^\top B + \lambda I_d)^{-1} b_i.$$

If $\lambda = 1$, we abbreviate as $\tilde{\sigma}_i(B)$.

For convenient of the proof, we also define

$$\tilde{\sigma}_{\lambda, i, j}(B) := b_i^\top (B^\top B + \lambda I_d)^{-1} b_j.$$

We also define $\hat{\sigma}_{\lambda, i}(B)$ as follows:

$$\hat{\sigma}_{\lambda, i}(B) := b_i^\top (B^\top B + \lambda I_d)^{-2} b_i.$$

Fact H.2. Let X be $\mathbb{R}^{d \times d}$ be invertible and symmetric. Then we have

$$\frac{\partial X^{-1}}{\partial t} = -X^{-1} \frac{\partial X}{\partial t} X^{-1}.$$

Lemma H.3. Let $X \in \mathbb{R}^{d \times d}$ be invertible and symmetric. Then,

$$\partial \log \det X = X^{-1} \partial X.$$

Proof. The proof is by chain rule:

$$\partial \log \det X = \frac{1}{\det X} \partial \det X,$$

to compute $\partial \det X$, recall the adjugate of X , denoted by $\text{adj}(X)$ where $X^{-1} = \frac{1}{\det X} \text{adj}(X)$, and it remains to show that $\partial \det X = \text{adj}(X)$, this can be derived using cofactor expansion of a matrix, which states that fix a row i , we have

$$\det X = \sum_{k=1}^d X_{i, k} \cdot \text{adj}(X)_{i, k},$$

by product rule,

$$\frac{\partial \det X}{\partial X_{i, j}} = \sum_{k=1}^d \frac{\partial X_{i, k}}{\partial X_{i, j}} \cdot \text{adj}(X)_{i, k} + X_{i, k} \cdot \frac{\partial \text{adj}(X)_{i, k}}{\partial X_{i, j}}$$

$$= \text{adj}(X)_{i,j},$$

where $\frac{\partial \text{adj}_{i,k}}{\partial X_{i,j}} = 0$ for all k since the cofactor of i, k is the principal minor by removing row i and column k . Therefore,

$$\partial \det X = \text{adj}(X)$$

and consequently,

$$\begin{aligned} \partial \log \det X &= \frac{1}{\det X} \partial \det X \\ &= \frac{1}{\det X} \text{adj}(X) \\ &= X^{-1}. \end{aligned} \quad \square$$

The goal of this section is to prove the $d \times d$ matrix $\nabla^2 F(x) \succ 0$. We start with computing $\nabla F(x) \in \mathbb{R}^d$.

Lemma H.4. We define $s_{x,i} = a_i^\top x - b_i$ for each $i \in [n]$. Let S_x denote a diagonal matrix such that $S_x = \text{diag}(s_x)$. Let $H(x) = A^\top S_x^{-2} A \in \mathbb{R}^{d \times d}$. Let $F(x) = \log \det(H(x) + I_d)$, then

$$\nabla F(x) = -2 \sum_{i=1}^n \tilde{\sigma}_i(S_x^{-1} A) \frac{a_i}{s_{x,i}}.$$

Proof. First, we know that

$$\frac{\partial s_{x,i}^{-2}}{\partial x_j} = \frac{-2a_{i,j}}{s_{x,i}^3}. \quad (19)$$

For each $j \in [d]$, we can write $\frac{\partial F}{\partial x_j}$ as follows:

$$\begin{aligned} \frac{\partial F}{\partial x_j} &= \text{tr}[(H(x) + I_d)^{-1} \frac{\partial(H(x) + I_d)}{\partial x_j}] \\ &= \text{tr}[(H(x) + I_d)^{-1} \frac{\partial H(x)}{\partial x_j}] \\ &= \text{tr}[(H(x) + I_d)^{-1} \sum_{i=1}^n a_i a_i^\top \frac{\partial s_{x,i}^{-2}}{\partial x_j}] \\ &= \text{tr}[(H(x) + I_d)^{-1} \sum_{i=1}^n -2a_i a_i^\top \frac{a_{i,j}}{s_{x,i}^3}] \\ &= -2 \sum_{i=1}^n \text{tr}[\frac{a_{i,j}}{s_{x,i}} \frac{a_i^\top (H(x) + I_d)^{-1} a_i}{s_{x,i}^2}] \\ &= -2 \sum_{i=1}^n \tilde{\sigma}_i(S_x^{-1} A) \frac{a_{i,j}}{s_{x,i}}. \end{aligned}$$

where the forth step follows from Eq. (19), the fifth step follows from $\text{tr}[AB] = \text{tr}[BA]$, and the last step follows from $\tilde{\sigma}$.

We use chain rule:

$$\begin{aligned} \frac{\partial F}{\partial x} &= [\frac{\partial F}{\partial x_1} \quad \frac{\partial F}{\partial x_2} \quad \cdots \quad \frac{\partial F}{\partial x_d}]^\top \\ &= -2 \sum_{i=1}^n \tilde{\sigma}_i(S_x^{-1} A) \frac{a_i}{s_{x,i}}. \end{aligned} \quad \square$$

Fact H.5. We have the following partial derivative:

$$\frac{\partial(\tilde{\sigma}_i(S_x^{-1} A) s_{x,i}^2)}{\partial x_l} = 2 \sum_{j=1}^n \tilde{\sigma}_{i,j}^2(S_x^{-1} A) \frac{a_{j,l}}{s_{x,j}}$$

Proof. We analyze the partial derivative term

$$\begin{aligned}
 \frac{\partial}{\partial x_l} a_i^\top (H(x) + I_d)^{-1} a_i &= a_i^\top \left(\frac{\partial(H(x) + I_d)^{-1}}{\partial x_l} \right) a_i \\
 &= a_i^\top \left(-(H(x) + I_d)^{-1} \frac{\partial H(x)}{\partial x_l} (H(x) + I_d)^{-1} \right) a_i \\
 &= a_i^\top \left(2(H(x) + I_d)^{-1} \left(\sum_{j=1}^n a_j a_j^\top \frac{a_{j,l}}{s_{x,j}^3} \right) (H(x) + I_d)^{-1} \right) a_i \\
 &= 2 \sum_{j=1}^n \frac{(a_i^\top (H(x) + I_d)^{-1} a_j)^2}{s_{x,j}^2} \frac{a_{j,l}}{s_{x,j}} \\
 &= 2 \sum_{j=1}^n \tilde{\sigma}_{i,j}^2 (S_x^{-1} A) \frac{a_{j,l}}{s_{x,j}}
 \end{aligned} \tag{20}$$

where the second step follows from Fact H.2. □

Fact H.6. We have

$$\tilde{\sigma}_i(S_x^{-1} A) = \hat{\sigma}_i(S_x^{-1} A) + \sum_{j=1}^n \tilde{\sigma}_{i,j}^2(S_x^{-1} A).$$

Proof. Note that

$$\begin{aligned}
 (H(x) + I_d)^{-1} &= (H(x) + I_d)^{-1} (H(x) + I_d) (H(x) + I_d)^{-1} \\
 &= (H(x) + I_d)^{-2} + \sum_{j=1}^n \frac{(H(x) + I_d)^{-1} a_j a_j^\top (H(x) + I_d)^{-1}}{s_{x,j}^2},
 \end{aligned}$$

therefore,

$$\begin{aligned}
 \tilde{\sigma}_i(S_x^{-1} A) &= a_i^\top (H(x) + I_d)^{-1} a_i \frac{1}{s_{x,i}^2} \\
 &= \frac{a_i^\top (H(x) + I_d)^{-2} a_i}{s_{x,i}^2} + \sum_{j=1}^n \frac{(a_i^\top (H(x) + I_d)^{-1} a_j)^2}{s_{x,i}^2 s_{x,j}^2} \\
 &= \hat{\sigma}_i(S_x^{-1} A) + \sum_{j=1}^n \tilde{\sigma}_{i,j}^2(S_x^{-1} A)
 \end{aligned}$$

Thus we complete the proof. □

Lemma H.7. Let $F(x) = \log \det(H(x) + I_d)$, then

$$\nabla^2 F(x) \succeq 0.$$

Thus, $F(x)$ is convex.

Proof. Note that

$$\begin{aligned}
 \frac{\partial}{\partial x_l} \frac{\partial}{\partial x_k} F(x) &= - \left(\sum_{i=1}^n a_i^\top (H(x) + I_d)^{-1} a_i a_i \frac{\partial}{\partial x_l} \left(\frac{a_{i,k}}{s_{x,i}^3} \right) + \frac{a_{i,k}}{s_{x,i}^3} \frac{\partial}{\partial x_l} a_i^\top (H(x) + I_d)^{-1} a_i \right) \\
 &= 3 \sum_{i=1}^n \frac{a_i^\top (H(x) + I_d)^{-1} a_i}{s_{x,i}^2} \frac{a_{i,k} a_{i,l}}{s_{x,i}^2} - \frac{a_{i,k}}{s_{x,i}^3} \frac{\partial}{\partial x_l} a_i^\top (H(x) + I_d)^{-1} a_i.
 \end{aligned} \tag{21}$$

Plug in Eq. (20) into Eq. (21), we have

$$3 \sum_{i=1}^n \tilde{\sigma}_i(S_x^{-1} A) \frac{a_{i,k} a_{i,l}}{s_{x,i}^2} - 2 \sum_{i=1}^n \sum_{j=1}^n \tilde{\sigma}_{i,j}(S_x^{-1} A)^2 \frac{a_{i,k} a_{j,l}}{s_{x,j} s_{x,i}},$$

and

$$\begin{aligned}
& 3 \sum_{i=1}^n \tilde{\sigma}_i(S_x^{-1}A) \frac{a_i a_i^\top}{s_{x,i}^2} - 2 \sum_{i=1}^n \sum_{j=1}^n \tilde{\sigma}_{i,j}^2(S_x^{-1}A) \frac{a_i a_j^\top}{s_{x,j} s_{x,i}} \\
&= 3 \sum_{i=1}^n (\hat{\sigma}_i(S_x^{-1}A) + \sum_{j=1}^n \tilde{\sigma}_{i,j}^2(S_x^{-1}A)) \frac{a_i a_i^\top}{s_{x,i}^2} - 2 \sum_{i=1}^n \sum_{j=1}^n \tilde{\sigma}_{i,j}^2(S_x^{-1}A) \frac{a_i a_j^\top}{s_{x,j} s_{x,i}} \\
&= 3 \sum_{i=1}^n \hat{\sigma}_i(S_x^{-1}A) \frac{a_i a_i^\top}{s_{x,i}^2} + 3 \sum_{i=1}^n \sum_{j=1}^n \tilde{\sigma}_{i,j}^2(S_x^{-1}A) \frac{a_i a_i^\top}{s_{x,i}^2} - 2 \sum_{i=1}^n \sum_{j=1}^n \tilde{\sigma}_{i,j}^2(S_x^{-1}A) \cdot \frac{a_i a_j^\top}{s_{x,j} s_{x,i}} \\
&= 3 \sum_{i=1}^n \hat{\sigma}_i(S_x^{-1}A) \frac{a_i a_i^\top}{s_{x,i}^2} + \sum_{i=1}^n \tilde{\sigma}_i(S_x^{-1}A) \frac{a_i a_i^\top}{s_{x,i}^2} + 2 \sum_{i=1}^n \sum_{j=1}^n \tilde{\sigma}_{i,j}^2(S_x^{-1}A) \cdot \left(\frac{a_i a_i^\top}{s_{x,i}^2} - \frac{a_i a_j^\top}{s_{x,j} s_{x,i}} \right) \\
&:= B_1 + B_2 + B_3.
\end{aligned}$$

where the first step follows from Fact H.6, the last step follows from

$$\begin{aligned}
B_1 &:= 3 \sum_{i=1}^n \hat{\sigma}_i(S_x^{-1}A) \frac{a_i a_i^\top}{s_{x,i}^2}, \\
B_2 &:= \sum_{i=1}^n \tilde{\sigma}_i(S_x^{-1}A) \frac{a_i a_i^\top}{s_{x,i}^2}, \\
B_3 &:= 2 \sum_{i=1}^n \sum_{j=1}^n \tilde{\sigma}_{i,j}^2(S_x^{-1}A) \cdot \left(\frac{a_i a_i^\top}{s_{x,i}^2} - \frac{a_i a_j^\top}{s_{x,j} s_{x,i}} \right).
\end{aligned}$$

It is obvious that $B_1 \succ 0$ and $B_2 \succ 0$.

It is not hard to see that

$$\frac{a_i a_i^\top}{s_{x,i}^2} + \frac{a_j a_j^\top}{s_{x,j}^2} \succeq 2 \frac{a_i a_j^\top}{s_{x,j} s_{x,i}}.$$

Thus we have $B_3 \succeq 0$ and we complete the proof. \square

H.2 Convexity of $\log \det(H_{\log}(x) + I_d)$: spectrahedra

In this section, we prove that $\log \det(H_{\log}(x) + I_d)$ is a convex function on x . Combined with the bounded variance property, we provide an algorithm that samples from the SDP spectrahedron.

Definition H.8. We define the ridge-projection matrix $P \in \mathbb{R}^{n^2 \times n^2}$ as

$$P := (S^{-1/2} \otimes S^{-1/2}) A^\top (H + I_d)^{-1} A (S^{-1/2} \otimes S^{-1/2}).$$

We define projection matrix $\bar{P} \in \mathbb{R}^{n^2 \times n^2}$ as follows

$$\bar{P} := (S^{-1/2} \otimes S^{-1/2}) A^\top H^{-1} A (S^{-1/2} \otimes S^{-1/2}).$$

Remark H.9. It is not hard to see that P is a PSD matrix. Note that the matrix $H + I_d$ is PSD, since H is PSD, therefore $(H + I_d)^{-1}$ is also PSD. Thus, the ridge-projection is PSD, as its two ‘‘arms’’ are symmetric.

Lemma H.10. Let $\bar{P} \in \mathbb{R}^{n^2 \times n^2}$ be the projection matrix and let $P \in \mathbb{R}^{n^2 \times n^2}$ be the ridge-projection matrix defined in Def. H.8. Then, we have $P \preceq \bar{P}$.

Proof. Let $H = U \Sigma U^\top$ be its eigendecomposition, we need to show that

$$(H + I_d)^{-1} \preceq H^{-1}.$$

Let us expand the LHS:

$$\begin{aligned}
(U \Sigma U^\top + I_d)^{-1} &= (U(\Sigma + I_d)U^\top)^{-1} \\
&= U(\Sigma + I_d)^{-1}U^\top \\
&\preceq U \Sigma^{-1} U^\top,
\end{aligned}$$

where the last step follows from for any non-negative number λ , $(1 + \lambda)^{-1} \leq \lambda^{-1}$. \square

The following lemma is a generalization of Theorem 4.1 of Anstreicher [2000].

Lemma H.11. Let $H_{\log}(x) \in \mathbb{R}^{d \times d}$ be defined as Definition G.6. Let $F(x) = \log \det(H_{\log}(x) + I_d)$. Then, $F(x)$ is convex in x .

Proof. Using standard algebra and chain rule computations, we have

$$\nabla^2 F(x) = 2Q(x) + R(x) - 2T(x).$$

For simplicity, we drop x . The $Q, R, T \in \mathbb{R}^{d \times d}$ can be defined as follows

$$\begin{aligned} Q_{i,j} &:= A(H + I_d)^{-1} A^\top \cdot (S^{-1} A_i S^{-1} A_j S^{-1} \oplus S^{-1}), \\ R_{i,j} &:= A(H + I_d)^{-1} A^\top \cdot (S^{-1} A_i S^{-1} \oplus S^{-1} A_j S^{-1}), \\ T_{i,j} &:= A(H + I_d)^{-1} A^\top \cdot (S^{-1} A_i S^{-1} \oplus S^{-1}) \cdot A(H + I_d)^{-1} A^\top \cdot (S^{-1} A_j S^{-1} \oplus S^{-1}). \end{aligned}$$

Using Lemma H.12, we know that

$$0 \preceq Q(x) \preceq \nabla^2 F(x) \preceq 3Q(x)$$

Thus, $F(x)$ is convex. □

Lemma H.12. For any x , if $S(x) \succ 0$, then we have

- Part 1. $Q \succeq 0$;
- Part 2. $T \succeq 0$;
- Part 3. $T \preceq \frac{1}{2}(Q + R)$;
- Part 4. $\nabla^2 F(x) \succeq 0$;
- Part 5. $R \preceq Q$;
- Part 6. $\nabla^2 F(x) \preceq 3Q(x)$.

Proof. Proof of Part 1 and 2. Let $\xi \in \mathbb{R}^d$ and $\xi \neq 0$.

Then we have

$$\begin{aligned} \xi^\top Q \xi &= \sum_{i,j} Q_{i,j} \xi_i \xi_j \\ &= A(H + I_d)^{-1} A^\top \cdot (S^{-1} B S^{-1} B S^{-1} \oplus S^{-1}) \\ &= P \cdot (\bar{B}^2 \oplus I_d) \end{aligned}$$

where $B = B(\xi) = \sum_{i=1}^d \xi_i A_i$, and $\bar{B} = S^{-1/2} B S^{-1/2}$.

Similarly, we have

$$\begin{aligned} \xi^\top R \xi &= A(H + I_d)^{-1} A^\top \cdot (S^{-1} B S^{-1} \otimes S^{-1} B S^{-1}) \\ &= P \cdot (\bar{B} \otimes \bar{B}) \end{aligned}$$

We can rewrite $\xi^\top T \xi$ as follows:

$$\begin{aligned} \xi^\top T \xi &= A^\top (H + I)^{-1} A \cdot (S^{-1} B S^{-1} \oplus S^{-1}) A (H + I)^{-1} A^\top \cdot (S^{-1} B S^{-1} \oplus S^{-1}) \\ &= P \cdot (\bar{B} \oplus I_d) P (\bar{B} \oplus I_d) \end{aligned}$$

It is obvious that I_d, P and \bar{B}^2 are all PSD matrices.

Using Part 3 of Lemma A.9 and Part 6 of Lemma A.10, we have

$$\begin{aligned} \xi^\top Q \xi &\geq 0 \\ \xi^\top T \xi &\geq 0 \end{aligned}$$

Since ξ is arbitrary, then we know that

$$\begin{aligned} Q &\succeq 0 \\ T &\succeq 0 \end{aligned}$$

Proof of Part 3. Note that \bar{P} is a projection matrix, and $P \preceq \bar{P}$ by Lemma H.10.

$$\begin{aligned} (\bar{B} \oplus I_d)P(\bar{B} \oplus I_d) &\preceq (\bar{B} \oplus I_d)\bar{P}(\bar{B} \oplus I_d) \\ &\preceq (\bar{B} \oplus I_d)I_d(\bar{B} \oplus I_d) \\ &= \frac{1}{2}((\bar{B}^2 \oplus I_d) + (\bar{B} \otimes \bar{B})) \end{aligned}$$

where the last step follows from Part 2 of Lemma A.10.

Applying Part 1 of Lemma A.9, we have

$$\langle P, (\bar{B} \oplus I_d)P(\bar{B} \oplus I_d) \rangle \leq \frac{1}{2} \langle P, (\bar{B}^2 \oplus I_d) + (\bar{B} \otimes \bar{B}) \rangle$$

The above equation implies the following

$$\xi^\top T \xi \leq \frac{1}{2} \xi^\top (Q + R) \xi$$

Note that, here ξ is arbitrary, thus we know that

$$T \preceq \frac{1}{2}(Q + R)$$

Proof of Part 4. We have

$$\begin{aligned} \nabla^2 F(x) &= 2Q + R - 2T \\ &\succeq 2Q + R - 2 \cdot \frac{1}{2}(Q + R) \\ &= Q \end{aligned}$$

Proof of Part 5. Let v_i be orthonormal eigenvectors of \bar{B} with corresponding eigenvalues λ_i .

Then using [Horn and Johnson, 1990, Theorem 4.4.5], $\bar{B}^2 \oplus I$ has orthonormal eigenvectors $v_i \otimes v_j$ with corresponding eigenvalues $\frac{1}{2}(\lambda_i^2 + \lambda_j^2)$, while (see Horn and Johnson [1990]) $\bar{B} \otimes \bar{B}$ has the same eigenvectors $v_i \otimes v_j$, with corresponding eigenvalues $\lambda_i \lambda_j$.

It then follows from $(\lambda_i - \lambda_j)^2$ for each i, j that

$$\bar{B}^2 \oplus I_d \succeq \bar{B} \otimes \bar{B}.$$

Using Part 4 of Lemma A.9, we know

$$\langle P, \bar{B}^2 \oplus I_d \rangle \geq \langle P, \bar{B} \otimes \bar{B} \rangle.$$

Which is

$$\xi^\top Q \xi \geq \xi^\top R \xi$$

Since it's arbitrary ξ , then we have

$$Q \succeq R$$

Proof of Part 6. We have

$$\begin{aligned} \nabla^2 F(x) &= 2Q + R - 2T \\ &\preceq 2Q + R \\ &\preceq 3Q \end{aligned}$$

where the second step follows from $T \succeq 0$, the last step follows from Part 5. □

H.3 Discussion

It is well-known that for popular barrier functions, such as volumetric barrier and Lee-Sidford barrier [Lee and Sidford, 2014, 2019], the function $\log \det H(x)$ is convex in x . However, given a PSD matrix function $H(x)$ for which $\log \det H(x)$ is convex in x , it is generally not true that $\log \det(H(x) + I_d)$ is convex.

Fact H.13 (Folklore). *Suppose that $H(x)$ is a positive semi-definite matrix for x in the domain. Suppose $\log \det H(x)$ is convex. It is possible $\log \det(H(x) + I)$ is not convex.*

Proof. If $\log \det(H(x) + I)$ is convex in x , then for any fixed $x \in K$, $v \in \mathbb{R}^d$, $\log \det(H(x + tv) + I_d)$ is convex in t (for t sufficiently close to 0).

For t sufficiently close to 0, define $f_i(t) := \lambda_i(H(x + tv))$. Then

$$\log \det H(x + tv) = \sum_{i \in [d]} \log f_i(t)$$

and the condition that $\log \det H(x)$ is convex is equivalent to

$$\frac{\partial^2}{\partial t^2} \log \det H(x + tv) = \sum_{i \in [d]} \frac{f_i''(t)f_i(t) - f_i'(t)^2}{f_i(t)^2} \geq 0.$$

Now

$$\log \det(H(x + tv) + I_d) = \sum_{i \in [d]} \log(f_i(t) + 1).$$

The condition that $\log \det(H(x) + I_d)$ is convex is equivalent to

$$\frac{\partial^2}{\partial t^2} \log \det(H(x + tv) + I_d) = \sum_{i \in [d]} \frac{f_i''(t)(f_i(t) + 1) - f_i'(t)^2}{(f_i(t) + 1)^2} \geq 0.$$

Suppose $d = 2$, $f_1(0) = 1$, $f_1'(0) = 0$, $f_1''(0) = 1$, $f_2(0) = 4$, $f_2'(0) = 0$, $f_2''(0) = -3$. Then

$$\begin{aligned} \frac{\partial^2}{\partial t^2} \Big|_{t=0} \log \det H(x + tv) &= \frac{f_1''(0)}{f_1(0)} + \frac{f_2''(0)}{f_2(0)} > 0. \\ \frac{\partial^2}{\partial t^2} \Big|_{t=0} \log \det(H(x + tv) + I) &= \frac{f_1''(0)}{f_1(0) + 1} + \frac{f_2''(0)}{f_2(0) + 1} < 0. \end{aligned}$$

This gives an example for which $\log \det H(x)$ is convex in a region but $\log \det(H(x) + I_d)$ is not. \square

I Convexity of Regularized Volumetric Barrier Function

In this section, we prove the convexity of $\log \det$ of the regularized volumetric barrier function. This is crucial, as the convexity proof for regularized Lee-Sidford barrier is identical up to replacing leverage score by Lewis weights.

I.1 Definitions

Definition I.1. *Let $A \in \mathbb{R}^{n \times d}$. Let $b \in \mathbb{R}^n$. For each $i \in [n]$, we define*

$$\begin{aligned} s_{x,i} &:= (a_i^\top x - b_i) \\ s_x &:= Ax - b \end{aligned}$$

For each $i \in [n]$, we define

$$\sigma_{i,i}(A_x) := a_{x,i}^\top (A_x^\top A_x)^{-1} a_{x,i}$$

For each $i \in [n]$, for each $l \in [n]$

$$\sigma_{i,l}(A_x) := a_{x,i}^\top (A_x^\top A_x)^{-1} a_{x,l}$$

Let $\Sigma = \sigma_{*,*}(A_x) \circ I_n$.

Definition I.2. We define $H(x) \in \mathbb{R}^{d \times d}$ as follows

$$H(x) := \underbrace{A_x^\top}_{d \times n} \underbrace{\Sigma(A_x)}_{n \times n} \underbrace{A_x}_{n \times d}$$

Definition I.3. For each $j \in [d]$, we define $P_j(x) \in \mathbb{R}^{n \times n}$ as follows

$$P_j(x) := \underbrace{\frac{\partial \Sigma(A_x)}{\partial x_j}}_{n \times n}$$

For each $j \in [d]$, for each $k \in [d]$, we define $P_{j,k}(x) \in \mathbb{R}^{n \times n}$ as follows:

$$P_{j,k}(x) := \underbrace{\frac{\partial^2 \Sigma(A_x)}{\partial x_k \partial x_j}}_{n \times n}$$

In the previous sections, we mainly use notation $\sigma_{*,*}(A_x)$. For simplicity, from this section, we will use notation $Q(x)$ instead.

Definition I.4. We define $Q(x) \in \mathbb{R}^{n \times n}$ as follows

$$Q(x) := \underbrace{\sigma_{*,*}(A_x)}_{n \times n}.$$

Fact I.5. We have

$$\sigma_{i,i}(A_x) = \langle \sigma_{*,i}^2(A_x), \mathbf{1}_n \rangle.$$

Proof. The proof is straightforward from definition. □

We list a number of standard calculus results.

Fact I.6. We define the following quantities:

- Let $s_x = Ax - b \in \mathbb{R}^n$;
- Let $S_x = \text{diag}(s_x) \in \mathbb{R}^{n \times n}$ denote the diagonal matrix by putting s_x on the diagonal;
- Let $A_{*,j}$ denote the j -th column of matrix $A \in \mathbb{R}^{n \times d}$;
- Let $A_x = S_x^{-1}A$;
- Let $a_{x,i}^\top$ denote the i -th row of A_x for each $i \in [n]$.

Then, we have for each $j \in [d]$

- **Part 1.**

$$\underbrace{\frac{\partial s_x}{\partial x_j}}_{n \times 1} = \underbrace{A_{*,j}}_{n \times 1}$$

- **Part 2.**

$$\underbrace{\frac{\partial s_x^{-1}}{\partial x_j}}_{n \times 1} = - \underbrace{s_x^{-2}}_{n \times 1} \circ \underbrace{A_{*,j}}_{n \times 1}$$

- **Part 3.**

$$\underbrace{\frac{\partial s_x^{-2}}{\partial x_j}}_{n \times 1} = -2 \underbrace{s_x^{-3}}_{n \times 1} \circ \underbrace{A_{*,j}}_{n \times 1}$$

• **Part 4.**

$$\underbrace{\frac{\partial S_x^{-2}}{\partial x_j}}_{n \times n} = 2 \underbrace{\text{diag}(-s_x^{-3} \circ A_{*,j})}_{n \times n} = 2 \text{diag}(-S_x^{-3} A_{*,j})$$

• **Part 5.**

$$\underbrace{\frac{\partial A^\top S_x^{-2} A}{\partial x_j}}_{d \times d} = 2 \underbrace{A^\top}_{d \times n} \underbrace{\text{diag}(-S_x^{-3} A_{*,j})}_{n \times n} \underbrace{A}_{n \times d}$$

• **Part 6.**

$$\frac{\partial A_x^\top A_x}{\partial x_j} = 2 A_x^\top \text{diag}(-A_{x,*,j}) A_x$$

• **Part 7.**

$$\frac{\partial (A_x^\top A_x)^{-1}}{\partial x_j} = 2 (A_x^\top A_x)^{-1} \cdot A_x^\top \text{diag}(A_{x,*,j}) A_x \cdot (A_x^\top A_x)^{-1}$$

• **Part 8.** For each $i \in [n]$

$$\frac{\partial a_{x,i}}{\partial x_j} = - \underbrace{A_{x,i,j}}_{\text{scalar}} \cdot \underbrace{a_{x,i}}_{d \times 1}$$

• **Part 9.** For each $i \in [n]$

$$\frac{\partial a_{x,i} a_{x,i}^\top}{\partial x_j} = -2 \cdot \underbrace{A_{x,i,j}}_{\text{scalar}} \cdot \underbrace{a_{x,i} a_{x,i}^\top}_{d \times d}$$

• **Part 10.**

$$\underbrace{\frac{\partial A_{x,i,j}}{\partial x_j}}_{\text{scalar}} = - \underbrace{A_{x,i,j}^2}_{\text{scalar}}$$

• **Part 11.**

$$\underbrace{\frac{\partial A_{x,i,j}}{\partial x_k}}_{\text{scalar}} = - \underbrace{A_{x,i,j}}_{\text{scalar}} \underbrace{A_{x,i,k}}_{\text{scalar}}$$

• **Part 12.**

$$\frac{\partial A_{x,*,j}}{\partial x_j} = -A_{x,*,j}^{\circ 2}$$

• **Part 13.**

$$\frac{\partial A_{x,*,j}}{\partial x_k} = -A_{x,*,j} \circ A_{x,*,k}$$

• **Part 14.**

$$\underbrace{\frac{\partial A_x}{\partial x_j}}_{n \times d} = - \underbrace{\text{diag}(A_{x,*,j})}_{n \times n} \underbrace{A_x}_{n \times d}$$

Definition I.7. We define $\tilde{\Sigma}(A_x) \in \mathbb{R}^{n \times n}$ and $\tilde{\sigma}_{*,*}(A_x) \in \mathbb{R}^{n \times n}$ as

$$\begin{aligned}\tilde{\Sigma}(A_x) &:= (A_x(H(x) + I_d)^{-1}A_x^\top) \circ I_n \\ \tilde{\sigma}_{*,*}(A_x) &:= A_x(H(x) + I_d)^{-1}A_x^\top\end{aligned}$$

Definition I.8. We define $F(x)$ as follows:

$$F(x) := \log(\det(H(x) + I_d))$$

Definition I.9. We define $\frac{\partial}{\partial x} \frac{\partial F}{\partial x^\top} \in \mathbb{R}^{d \times d}$ to be

$$\frac{\partial}{\partial x} \frac{\partial F}{\partial x^\top} = \left[\frac{\partial}{\partial x} \frac{\partial F}{\partial x^\top} \right]_1 + \left[\frac{\partial}{\partial x} \frac{\partial F}{\partial x^\top} \right]_2$$

where

$$\begin{aligned}\left[\frac{\partial}{\partial x} \frac{\partial F}{\partial x^\top} \right]_{1,j,k} &= + \operatorname{tr}[(H(x) + I_d)^{-1} \frac{\partial^2 H(x)}{\partial x_j \partial x_k}] \\ \left[\frac{\partial}{\partial x} \frac{\partial F}{\partial x^\top} \right]_{2,j,k} &= - \operatorname{tr}[(H(x) + I_d)^{-1} \frac{\partial H(x)}{\partial x_k} (H(x) + I_d)^{-1} \cdot \frac{\partial H(x)}{\partial x_j}]\end{aligned}$$

I.2 Gradient of $\sigma_{i,j}$

Fact I.10 (First derivative of leverage score). We define the following quantities:

- For each $j \in [d]$, let $A_{*,j} \in \mathbb{R}^n$ denote j -th column of A ;
- Let $A_x := S_x^{-1}A \in \mathbb{R}^{n \times d}$;
- Let $A_{x,*,j} \in \mathbb{R}^n$ denote the j -th column of $A_x \in \mathbb{R}^{n \times d}$;
- Let $\Sigma(A_x) \in \mathbb{R}^{n \times n}$ denote a diagonal matrix where (i, i) -th entry is $\sigma_{i,i}(A_x)$;
- Let $\sigma_{*,*}^{\circ 2}(A_x) = \sigma_{*,*}(A_x) \circ \sigma_{*,*}(A_x) \in \mathbb{R}^{n \times n}$;
- Let $\sigma_{*,i}(A_x) \in \mathbb{R}^n$ denote a column vector of $\sigma_{*,*}(A_x) \in \mathbb{R}^{n \times n}$.

Then we have for each $j \in [d]$

- Part 1. For each $i \in [n]$,

$$\frac{\partial \sigma_{i,i}(A_x)}{\partial x_j} = 2 \langle \sigma_{*,i}(A_x) \circ \sigma_{*,i}(A_x), A_{x,*,j} \rangle - 2 \sigma_{i,i}(A_x) \cdot A_{x,i,j}$$

- Part 2. For each $i \in [n]$, $l \in [n]$

$$\frac{\partial \sigma_{i,l}(A_x)}{\partial x_j} = 2 \langle \sigma_{i,*}(A_x) \circ \sigma_{l,*}(A_x), A_{x,*,j} \rangle - \sigma_{i,l}(A_x) \cdot (A_{x,i,j} + A_{x,l,j})$$

Proof. We know

$$\begin{aligned}\frac{\partial \sigma_{i,i}(A_x)}{\partial x_j} &= \frac{\partial a_{x,i}^\top (A_x^\top A_x)^{-1} a_{x,i}}{\partial x_j} \\ &= \frac{\partial \langle a_{x,i} a_{x,i}^\top, (A_x^\top A_x)^{-1} \rangle}{\partial x_j} \\ &= \left\langle \frac{\partial a_{x,i} a_{x,i}^\top}{\partial x_j}, (A_x^\top A_x)^{-1} \right\rangle + \langle a_{x,i} a_{x,i}^\top, \frac{\partial (A_x^\top A_x)^{-1}}{\partial x_j} \rangle\end{aligned}$$

For the first term in the above, we have

$$-2A_{x,i,j} \cdot \langle a_{x,i} a_{x,i}^\top, (A_x^\top A_x)^{-1} \rangle = -2A_{x,i,j} \cdot \sigma_{i,i}(A_x)$$

For the second term, we have

$$\begin{aligned}
 \langle a_{x,i} a_{x,i}^\top, \frac{\partial (A_x^\top A_x)^{-1}}{\partial x_j} \rangle &= 2 \langle a_{x,i} a_{x,i}^\top, (A_x^\top A_x)^{-1} \cdot A_x^\top \text{diag}(A_{x,*,j}) A_x \cdot (A_x^\top A_x)^{-1} \rangle \\
 &= 2 a_{x,i}^\top (A_x^\top A_x)^{-1} \cdot A_x^\top \text{diag}(A_{x,*,j}) A_x \cdot (A_x^\top A_x)^{-1} a_{x,i} \\
 &= 2 a_{x,i}^\top (A_x^\top A_x)^{-1} \cdot \left(\sum_{l=1}^n a_{x,l} a_{x,l}^\top A_{x,l,j} \right) \cdot (A_x^\top A_x)^{-1} a_{x,i} \\
 &= 2 \sum_{l=1}^n \sigma_{l,i}(A_x)^2 A_{x,l,j} \\
 &= 2 \langle \sigma_{*,i}^2(A_x), A_{x,*,j} \rangle.
 \end{aligned}$$

Similarly, we can prove for $\frac{\partial \sigma_{i,i}(A_x)}{\partial x_j}$. □

I.3 Gradient of σ

Lemma I.11. *Let σ and Σ be defined as Definition I.1. Then, we have*

• **Part 1**

$$\frac{\partial \sigma_{*,i}(A_x)}{\partial x_j} = 2 \underbrace{\sigma_{*,*}(A_x)}_{n \times n} \underbrace{(\sigma_{*,i} \circ A_{x,*,j})}_{n \times 1} - 2 \underbrace{\sigma_{*,i}(A_x)}_{n \times 1} \underbrace{A_{x,*,j}}_{n \times 1}$$

• **Part 2.**

$$\begin{aligned}
 &\frac{\partial \sigma_{*,*}(A_x)}{\partial x_j} \\
 &= 2 \underbrace{\sigma_{*,*}(A_x)}_{n \times n} \text{diag}(A_{x,*,j}) \sigma_{*,*}(A_x) - \text{diag}(A_{x,*,j}) \sigma_{*,*}(A_x) - \sigma_{*,*}(A_x) \text{diag}(A_{x,*,j})
 \end{aligned}$$

• **Part 3.**

$$\begin{aligned}
 \frac{\partial \Sigma(A_x)}{\partial x_j} &= 2 \text{diag}(\underbrace{\sigma_{*,*}^{\circ 2}(A_x)}_{n \times n} \underbrace{A_{x,*,j}}_{n \times 1}) - 2 \text{diag}(\underbrace{\Sigma(A_x)}_{n \times n} \underbrace{A_{x,*,j}}_{n \times 1}) \\
 &= 2 \text{diag}((\sigma_{*,*}^{\circ 2}(A_x) - \Sigma(A_x)) A_{x,*,j})
 \end{aligned}$$

Proof. It follows from Fact I.10. □

I.4 Gradient for $H(x)$

Lemma I.12. *Recall the definitions of the following quantities:*

- Let $H(x) \in \mathbb{R}^{d \times d}$ be defined as Definition I.2.
- For each $j \in [d]$, let $P_j(x) \in \mathbb{R}^{n \times n}$ be defined as Definition I.3.
- For each $j \in [d]$, for each $k \in [d]$, let $P_{j,k}(x) \in \mathbb{R}^{n \times n}$ be defined as Definition I.3.

Then, we have

- **Part 1.** For each $j \in [d]$

$$\begin{aligned}
 \frac{\partial H(x)}{\partial x_j} &= -2 A_x^\top \text{diag}(\Sigma(A_x) A_{x,*,j}) A_x \\
 &\quad + A_x^\top P_j(x) A_x
 \end{aligned}$$

- **Part 2.** For each $j \in [d]$, for each $k \in [d]$

$$\begin{aligned} \frac{\partial}{\partial x_k} \left(\frac{\partial H(x)}{\partial x_j} \right) &= + 6A_x^\top \text{diag}(A_{x,*,k}) \Sigma(A_x) \text{diag}(A_{x,*,j}) A_x \\ &\quad - 2A_x^\top P_k(x) \text{diag}(A_{x,*,j}) A_x \\ &\quad - 2A_x^\top P_j(x) \text{diag}(A_{x,*,k}) A_x \\ &\quad + A_x^\top P_{j,k}(x) A_x \end{aligned}$$

Proof. **Proof of Part 1.** We have

$$\begin{aligned} \frac{\partial H(x)}{\partial x_j} &= \frac{\partial (A_x^\top \Sigma(A_x) A_x)}{\partial x_j} \\ &= \left(\frac{\partial A_x^\top}{\partial x_j} \right) \cdot \Sigma(A_x) A_x + A_x^\top \cdot \left(\frac{\partial \Sigma(A_x)}{\partial x_j} \right) \cdot A_x + A_x^\top \Sigma(A_x) \cdot \left(\frac{\partial A_x}{\partial x_j} \right) \\ &= -A_x \text{diag}(A_{x,*,j}) \Sigma(A_x) A_x \\ &\quad + A_x^\top P(x)_j A_x \\ &\quad - A_x \Sigma(A_x) \text{diag}(A_{x,*,j}) A_x \\ &= -2A_x^\top \text{diag}(\Sigma(A_x) A_{x,*,j}) A_x \\ &\quad + A_x^\top P_j(x) A_x \end{aligned}$$

Proof of Part 2. Then, we have

$$\begin{aligned} \frac{\partial}{\partial x_k} \left(\frac{\partial H(x)}{\partial x_j} \right) &= \frac{\partial}{\partial x_k} (-2A_x^\top \Sigma(A_x) \text{diag}(A_{x,*,j}) A_x + A_x^\top P_j(x) A_x) \\ &= \frac{\partial}{\partial x_k} (-2A_x^\top \Sigma(A_x) \text{diag}(A_{x,*,j}) A_x) + \frac{\partial}{\partial x_k} (A_x^\top P_j(x) A_x), \end{aligned} \quad (22)$$

where the first step follows from **Part 1**, the second step follows from the sum rule.

For the first term of Eq. (22), we have

$$\begin{aligned} &\frac{\partial}{\partial x_k} (-2A_x^\top \Sigma(A_x) \text{diag}(A_{x,*,j}) A_x) \\ &= -2 \frac{\partial}{\partial x_k} (A_x^\top) \Sigma(A_x) \text{diag}(A_{x,*,j}) A_x - 2A_x^\top \frac{\partial}{\partial x_k} (\Sigma(A_x)) \text{diag}(A_{x,*,j}) A_x \\ &\quad - 2A_x^\top \Sigma(A_x) \frac{\partial}{\partial x_k} (\text{diag}(A_{x,*,j})) A_x - 2A_x^\top \Sigma(A_x) \text{diag}(A_{x,*,j}) \frac{\partial}{\partial x_k} (A_x) \\ &= -2 \frac{\partial}{\partial x_k} (A_x^\top) \Sigma(A_x) \text{diag}(A_{x,*,j}) A_x - 2A_x^\top \frac{\partial}{\partial x_k} (\Sigma(A_x)) \text{diag}(A_{x,*,j}) A_x \\ &\quad - 2A_x^\top \Sigma(A_x) \text{diag} \left(\frac{\partial}{\partial x_k} (A_{x,*,j}) \right) A_x - 2A_x^\top \Sigma(A_x) \text{diag}(A_{x,*,j}) \frac{\partial}{\partial x_k} (A_x) \\ &= + 2A_x^\top \text{diag}(A_{x,*,k}) \Sigma(A_x) \text{diag}(A_{x,*,j}) A_x \\ &\quad - 2A_x^\top P_k(x) \text{diag}(A_{x,*,j}) A_x \\ &\quad + 2A_x^\top \Sigma(A_x) \text{diag}(A_{x,*,j}) \text{diag}(A_{x,*,k}) A_x \\ &\quad + 2A_x^\top \Sigma(A_x) \text{diag}(A_{x,*,j}) \text{diag}(A_{x,*,k}) A_x, \end{aligned} \quad (23)$$

For the second term of Eq. (22), we have

$$\begin{aligned} &\frac{\partial}{\partial x_k} (A_x^\top P_j(x) A_x) \\ &= \frac{\partial}{\partial x_k} (A_x^\top) P_j(x) A_x + A_x^\top \frac{\partial}{\partial x_k} (P_j(x)) A_x + A_x^\top P_j(x) \frac{\partial}{\partial x_k} (A_x) \\ &= -(A_x^\top) \text{diag}(A_{x,*,k}) P_j(x) A_x + A_x^\top P_{j,k}(x) A_x - A_x^\top P_j(x) \text{diag}(A_{x,*,k}) A_x \end{aligned}$$

$$\begin{aligned}
&= - (A_x^\top) \text{diag}(A_{x,*,k})P_j(x)A_x + A_x^\top P_{j,k}(x)A_x - A_x^\top \text{diag}(A_{x,*,k})P_j(x)A_x \\
&= - 2(A_x^\top) \text{diag}(A_{x,*,k})P_j(x)A_x + A_x^\top P_{j,k}(x)A_x.
\end{aligned} \tag{24}$$

By combining Eq. (22), Eq. (23), and Eq. (24), we have

$$\begin{aligned}
\frac{\partial}{\partial x_k} \left(\frac{\partial H(x)}{\partial x_j} \right) &= + 2A_x^\top \text{diag}(A_{x,*,k})\Sigma(A_x) \text{diag}(A_{x,*,j})A_x \\
&\quad - 2A_x^\top P_k(x) \text{diag}(A_{x,*,j})A_x \\
&\quad + 2A_x^\top \Sigma(A_x) \text{diag}(A_{x,*,j}) \text{diag}(A_{x,*,k})A_x \\
&\quad + 2A_x^\top \Sigma(A_x) \text{diag}(A_{x,*,j}) \text{diag}(A_{x,*,k})A_x \\
&\quad - 2A_x^\top \text{diag}(A_{x,*,k})P_j(x)A_x \\
&\quad + A_x^\top P_{j,k}(x)A_x \\
&= + 6A_x^\top \text{diag}(A_{x,*,k})\Sigma(A_x) \text{diag}(A_{x,*,j})A_x \\
&\quad - 2A_x^\top P_k(x) \text{diag}(A_{x,*,j})A_x \\
&\quad - 2A_x^\top P_j(x) \text{diag}(A_{x,*,k})A_x \\
&\quad + A_x^\top P_{j,k}(x)A_x.
\end{aligned}$$

Thus, we complete the proof. □

I.5 Hessian of $H(x)$

Lemma I.13. Recall the definitions of following quantities:

- Let $H(x) := A_x^\top \Sigma(A_x)A_x \in \mathbb{R}^{d \times d}$ be as definition I.2.
- Let $Q(x) := \sigma_{*,*}(A_x) \in \mathbb{R}^{n \times n}$ be defined as Definition I.4.

Then, we have

- **Part 1.** For each $j \in [d]$ and for each $k \in [d]$

$$\frac{\partial^2 H(x)}{\partial x_j \partial x_k} = C_1 + C_2 + C_3 + C_4 + C_5$$

where we define some local $d \times d$ size matrix variables

- $C_1 = +20A_x^\top \text{diag}(A_{x,*,j})\Sigma(A_x) \text{diag}(A_{x,*,k})A_x$
- $C_2 = -6A_x^\top \text{diag}(Q^{\circ 2}(x)(A_{x,*,k} \circ A_{x,*,j}))A_x$
- $C_3 = -8A_x^\top \text{diag}(A_{x,*,k}) \text{diag}(Q^{\circ 2}(x)A_{x,*,j})A_x$
- $C_4 = -8A_x^\top \text{diag}(A_{x,*,j}) \text{diag}(Q^{\circ 2}(x)A_{x,*,k})A_x$
- $C_5 = +8A_x^\top ((Q(x) \text{diag}(A_{x,*,k})Q(x) \text{diag}(A_{x,*,j})Q(x)) \circ I_n)A_x$

Proof. By **Part 2** of Lemma I.12, we can show that

$$\frac{\partial^2 H(x)}{\partial x_j \partial x_k} = B_1 + B_2 + B_3 + B_4$$

where

$$\begin{aligned}
B_1 &= + 6A_x^\top \text{diag}(A_{x,*,k})\Sigma(A_x) \text{diag}(A_{x,*,j})A_x \\
B_2 &= - 2A_x^\top P_k(x) \text{diag}(A_{x,*,j})A_x \\
B_3 &= - 2A_x^\top P_j(x) \text{diag}(A_{x,*,k})A_x \\
B_4 &= + A_x^\top P_{j,k}(x)A_x
\end{aligned}$$

where B_2 can be decomposed further as

$$B_2 = - 2A_x^\top 2 \text{diag}((Q^{\circ 2}(x) - \Sigma(A_x))A_{x,*,k}) \text{diag}(A_{x,*,j})A_x$$

$$\begin{aligned}
&= -4A_x^\top \text{diag}(Q^{\circ 2}(x)A_{x,*,k} - \Sigma(A_x)A_{x,*,k}) \text{diag}(A_{x,*,j})A_x \\
&= -4A_x^\top (\text{diag}(Q^{\circ 2}(x)A_{x,*,k}) - \text{diag}(\Sigma(A_x)A_{x,*,k})) \text{diag}(A_{x,*,j})A_x \\
&= -4A_x^\top \text{diag}(Q^{\circ 2}(x)A_{x,*,k}) \text{diag}(A_{x,*,j})A_x + 4A_x^\top \text{diag}(\Sigma(A_x)A_{x,*,k}) \text{diag}(A_{x,*,j})A_x \\
&= -4A_x^\top \text{diag}(Q^{\circ 2}(x)A_{x,*,k}) \text{diag}(A_{x,*,j})A_x + 4A_x^\top \text{diag}(A_{x,*,k})\Sigma(A_x) \text{diag}(A_{x,*,j})A_x,
\end{aligned}$$

where the second step follows from simple algebra, the third step follows from the definition of $\text{diag}(\cdot)$, the fourth step follows from simple algebra, and the last step follows from the definition of $\text{diag}(\cdot)$. For B_3 , consider the following:

$$\begin{aligned}
B_3 &= -2A_x^\top \cdot 2 \text{diag}((Q^{\circ 2}(x) - \Sigma(A_x))A_{x,*,j}) \text{diag}(A_{x,*,k})A_x \\
&= -4A_x^\top \text{diag}(Q^{\circ 2}(x)A_{x,*,j} - \Sigma(A_x)A_{x,*,j}) \text{diag}(A_{x,*,k})A_x \\
&= -4A_x^\top (\text{diag}(Q^{\circ 2}(x)A_{x,*,j}) - \text{diag}(\Sigma(A_x)A_{x,*,j})) \text{diag}(A_{x,*,k})A_x \\
&= -4A_x^\top \text{diag}(Q^{\circ 2}(x)A_{x,*,j}) \text{diag}(A_{x,*,k})A_x + 4A_x^\top \text{diag}(\Sigma(A_x)A_{x,*,j}) \text{diag}(A_{x,*,k})A_x \\
&= -4A_x^\top \text{diag}(Q^{\circ 2}(x)A_{x,*,j}) \text{diag}(A_{x,*,k})A_x + 4A_x^\top \text{diag}(A_{x,*,j})\Sigma(A_x) \text{diag}(A_{x,*,k})A_x,
\end{aligned}$$

where the second step follows from simple algebra, the third step follows from the definition of $\text{diag}(\cdot)$, the fourth step follows from simple algebra, and the last step follows from the definition of $\text{diag}(\cdot)$.

Finally for B_4 ,

$$\begin{aligned}
B_4 &= A_x^\top (\\
&\quad + 8(Q(x) \text{diag}(A_{x,*,k})Q(x) \text{diag}(A_{x,*,j})Q(x)) \circ I_n \\
&\quad - 6 \text{diag}(Q^{\circ 2}(x) \cdot (A_{x,*,k} \circ A_{x,*,j})) \\
&\quad - 4 \text{diag}(Q^{\circ 2}(x) \cdot A_{x,*,j}) \text{diag}(A_{x,*,k}) \\
&\quad - 4 \text{diag}(Q^{\circ 2}(x) \cdot A_{x,*,k}) \text{diag}(A_{x,*,j}) \\
&\quad + 6\Sigma(A_x) \text{diag}(A_{x,*,j}) \text{diag}(A_{x,*,k}) \\
&\quad)A_x
\end{aligned}$$

Eventually, we can show Hessian is $C_1 + \dots + C_5$. The reason is following:

For the term C_1 , we have

$$\begin{aligned}
C_1 &= B_1 + B_{2,2} + B_{3,2} + B_{4,5} \\
&= (6 + 4 + 4 + 6) \cdot A_x^\top \text{diag}(A_{x,*,j})\Sigma(A_x) \text{diag}(A_{x,*,i}) \\
&= 20 \cdot A_x^\top \text{diag}(A_{x,*,j})\Sigma(A_x) \text{diag}(A_{x,*,i})
\end{aligned}$$

where $B_{2,2}$ is the second term of B_2 , $B_{3,2}$ is the second term of B_3 and $B_{4,5}$ is the last term of B_4 .

For the term C_2 , we have

$$C_2 = B_{4,2}$$

For the term C_3 , we have

$$\begin{aligned}
C_3 &= B_{3,1} + B_{4,3} \\
&= -(4 + 4) \cdot A_x^\top \text{diag}(A_{x,*,k}) \text{diag}(Q^{\circ 2}(x)A_{x,*,j})A_x \\
&= -8 \cdot A_x^\top \text{diag}(A_{x,*,k}) \text{diag}(Q^{\circ 2}(x)A_{x,*,j})A_x
\end{aligned}$$

For the term C_4 , we have

$$\begin{aligned}
C_4 &= B_{2,1} + B_{4,4} \\
&= -4A_x^\top \text{diag}(Q^{\circ 2}(x)A_{x,*,k}) \text{diag}(A_{x,*,j})A_x - 4A_x^\top \text{diag}(Q^{\circ 2}(x) \cdot A_{x,*,k}) \text{diag}(A_{x,*,j})A_x \\
&= -8A_x^\top \text{diag}(Q^{\circ 2}(x)A_{x,*,k}) \text{diag}(A_{x,*,j})A_x.
\end{aligned}$$

For the term C_5 , we have

$$C_5 = B_{4,1}.$$

□

I.6 Gradient and Hessian of $F(x)$

Lemma I.14. Let $F(x) \in \mathbb{R}$ be defined as Definition I.8. Then, we have

• **Part 1.**

$$\frac{\partial F(x)}{\partial x_j} = \text{tr}[(H(x) + I_d)^{-1} \cdot \frac{\partial H(x)}{\partial x_j}]$$

• **Part 2**

$$\begin{aligned} \frac{\partial}{\partial x_k} \frac{\partial F(x)}{\partial x_j} &= + \text{tr}[(H(x) + I_d)^{-1} \frac{\partial^2 H}{\partial x_j \partial x_k}] \\ &\quad - \text{tr}[(H(x) + I_d)^{-1} \frac{\partial H(x)}{\partial x_k} (H(x) + I_d)^{-1} \cdot \frac{\partial H(x)}{\partial x_j}] \end{aligned}$$

Proof. Proof of Part 1. We have

$$\begin{aligned} \frac{\partial F(x)}{\partial x_j} &= \frac{\partial \log(\det(H(x) + I_d))}{\partial x_j} \\ &= \text{tr}[(H(x) + I_d)^{-1} \frac{\partial(H(x) + I_d)}{\partial x_j}] \\ &= \text{tr}[(H(x) + I + d)^{-1} (\frac{\partial H(x)}{\partial x_j} + \frac{\partial I_d}{\partial x_j})] \\ &= \text{tr}[(H(x) + I_d)^{-1} \cdot \frac{\partial H(x)}{\partial x_j}], \end{aligned}$$

where the first step follows from the definition of $F(x)$ (see Definition I.8), the third step follows from the basic calculus rule, and the last step follows from $\frac{\partial I_d}{\partial x_j} = 0$.

Proof of Part 2. We have

$$\begin{aligned} &\frac{\partial}{\partial x_k} \frac{\partial F(x)}{\partial x_j} \\ &= \frac{\partial}{\partial x_k} (\text{tr}[(H(x) + I_d)^{-1} \cdot \frac{\partial H(x)}{\partial x_j}]) \\ &= \text{tr}[\frac{\partial}{\partial x_k} ((H(x) + I_d)^{-1}) \cdot \frac{\partial H(x)}{\partial x_j}] + \text{tr}[(H(x) + I_d)^{-1} \cdot \frac{\partial}{\partial x_k} (\frac{\partial H(x)}{\partial x_j})] \\ &= \text{tr}[-(H(x) + I_d)^{-1} \frac{\partial}{\partial x_k} (H(x)) (H(x) + I_d)^{-1} \cdot \frac{\partial H(x)}{\partial x_j}] + \text{tr}[(H(x) + I_d)^{-1} \cdot \frac{\partial}{\partial x_k} (\frac{\partial H(x)}{\partial x_j})] \\ &= + \text{tr}[(H(x) + I_d)^{-1} \frac{\partial^2 H}{\partial x_j \partial x_k}] - \text{tr}[(H(x) + I_d)^{-1} \frac{\partial H(x)}{\partial x_k} (H(x) + I_d)^{-1} \cdot \frac{\partial H(x)}{\partial x_j}], \end{aligned}$$

where the first step follows from **Part 1**, the second step follows from the product rule, and the last step follows from simple algebra. \square

Using standard algebraic tools in literature [Lee and Sidford, 2019], we can show that,

Lemma I.15. We can rewrite it as follows

$$\begin{aligned} \frac{\partial}{\partial x} \frac{\partial F}{\partial x} &= 20H_1 - 6H_2 - 8H_3 - 8H_4 + 8H_5 \\ &\quad - 4G_1 + 8G_2 + 8G_3 - 16G_4 \end{aligned}$$

and

$$\frac{\partial}{\partial x} \frac{\partial F}{\partial x} \succ 0.$$

Thus, the function F is convex in x .

I.7 Quantity I for Hessian of F

Lemma I.16. *We have*

$$\left[\frac{\partial}{\partial x} \frac{\partial F}{\partial x^\top}\right]_1 = 20H_1 - 6H_2 - 8H_3 - 8H_4 + 8H_5$$

where

- $H_1 = A_x^\top \tilde{\Sigma}(A_x) \Sigma(A_x) A_x$
- $H_2 = A_x^\top \text{diag}(Q^{\circ 2}(x) \tilde{\Sigma}(A_x) \mathbf{1}_n) A_x$
- $H_3 = A_x^\top Q^{\circ 2}(x) \tilde{\Sigma}(A_x) A_x$
- $H_4 = A_x^\top \tilde{\Sigma}(A_x) Q^{\circ 2}(x) A_x$
- $H_5 = A_x^\top (Q(x) \tilde{\Sigma}(A_x) Q(x)) \circ Q(x) A_x$

Proof. It comes from $P_{i,j}$ and the following lemma. □

Lemma I.17. *Let $H_1 \in \mathbb{R}^{d \times d}$ be Hessian where each entry is*

$$(H_1)_{j,k} = \text{tr}[(H(x) + I_d)^{-1} A_x^\top \text{diag}(A_{x,*,j}) \Sigma(A_x) \text{diag}(A_{x,*,k}) A_x]$$

Then, we have

$$H_1 = A_x^\top \tilde{\Sigma}(A_x) \Sigma(A_x) A_x$$

Proof. We can rewrite $(H_1)_{j,k}$ as follows

$$\begin{aligned} (H_1)_{j,k} &= \text{tr}[(H(x) + I_d)^{-1} A_x^\top \text{diag}(A_{x,*,j}) \Sigma(A_x) \text{diag}(A_{x,*,k}) A_x] \\ &= \text{tr}[A_x (H(x) + I_d)^{-1} A_x^\top \text{diag}(A_{x,*,j}) \Sigma(A_x) \text{diag}(A_{x,*,k})] \\ &= \text{tr}[\tilde{\sigma}_{*,*}(A_x) \text{diag}(A_{x,*,j}) \Sigma(A_x) \text{diag}(A_{x,*,k})] \\ &= A_{x,*,j}^\top \tilde{\Sigma}(A_x) \Sigma(A_x) A_{x,*,k}, \end{aligned}$$

where the first step follows from the Lemma statement, the second step follows from the fact that all of $(H(x) + I_d)^{-1}$, $\text{diag}(A_{x,*,j})$, $\Sigma(A_x)$, and $\text{diag}(A_{x,*,k})$ are diagonal matrices, the third step follows from the definition of $\tilde{\sigma}_{*,*}$ (see Definition I.7), and the last step follows from the definition of $\tilde{\Sigma}(A_x)$ (see Definition I.7).

Thus, we have

$$H_1 = A_x^\top \tilde{\Sigma}(A_x) \Sigma(A_x) A_x. \quad \square$$

Lemma I.18. *Let H_2 be defined as*

$$(H_2)_{j,k} = \text{tr}[(H(x) + I_d)^{-1} A_x^\top \text{diag}(Q^{\circ 2}(x) (A_{x,*,k} \circ A_{x,*,j})) A_x]$$

Then, we have

$$H_2 = A_x^\top \text{diag}(Q^{\circ 2}(x) \tilde{\Sigma}(A_x) \mathbf{1}_n) A_x$$

Proof. We have

$$\begin{aligned} (H_2)_{j,k} &= \text{tr}[(H(x) + I_d)^{-1} A_x^\top \text{diag}(Q^{\circ 2}(x) (A_{x,*,k} \circ A_{x,*,j})) A_x] \\ &= \text{tr}[\tilde{\Sigma}(A_x) \text{diag}(Q^{\circ 2}(x) (A_{x,*,k} \circ A_{x,*,j}))] \\ &= \mathbf{1}_n^\top \tilde{\Sigma}(A_x) Q^{\circ 2}(x) (A_{x,*,k} \circ A_{x,*,j}) \\ &= A_{x,*,k}^\top \text{diag}(Q^{\circ 2}(x) \tilde{\Sigma}(A_x) \mathbf{1}_n) A_{x,*,j} \end{aligned}$$

where the first step follows from the Lemma statement, the second step follows from the definition of $\tilde{\Sigma}(A_x)$ (see Definition I.7), the third step follows from the definition of $\mathbf{1}_n^\top$, and the last step follows from the definition of $\Sigma(A_x)$.

Thus, we have

$$H_2 = A_x^\top \text{diag}(Q^{\circ 2}(x) \tilde{\Sigma}(A_x) \mathbf{1}_n) A_x. \quad \square$$

Lemma I.19. Let H_3 be defined as

$$(H_3)_{j,k} = \text{tr}[(H(x) + I_d)^{-1} A_x^\top \text{diag}(A_{x,*,k}) \text{diag}(Q^{\circ 2}(x)(A_{x,*,j})) A_x]$$

Then, we

$$H_3 = A_x^\top Q^{\circ 2}(x) \tilde{\Sigma}(A_x) A_x$$

Proof. Then, we have

$$\begin{aligned} (H_3)_{j,k} &= \text{tr}[(H(x) + I_d)^{-1} A_x^\top \text{diag}(A_{x,*,k}) \text{diag}(Q^{\circ 2}(x)(A_{x,*,j})) A_x] \\ &= \text{tr}[\tilde{\Sigma}(A_x) \text{diag}(A_{x,*,k}) \text{diag}(Q^{\circ 2}(x)(A_{x,*,j}))] \\ &= A_{x,*,k}^\top \tilde{\Sigma}(A_x) Q^{\circ 2}(x) A_{x,*,j}, \end{aligned}$$

where the first step follows from the Lemma statement, the second step follows from the definition of $\tilde{\Sigma}(A_x)$ (see Definition I.7), and the last step follows from the property of $\text{diag}(\cdot)$.

Thus, we have

$$H_3 = A_x^\top Q^{\circ 2}(x) \tilde{\Sigma}(A_x) A_x. \quad \square$$

Lemma I.20. Let H_4 be defined as

$$(H_4)_{j,k} = \text{tr}[(H(x) + I_d)^{-1} A_x^\top \text{diag}(A_{x,*,j}) \text{diag}(Q^{\circ 2}(x)(A_{x,*,k})) A_x]$$

Then, we

$$H_4 = A_x^\top \tilde{\Sigma}(A_x) Q^{\circ 2}(x) A_x$$

Proof. This proof is very similar to Lemma I.19, so we omit the details here. □

Lemma I.21. Let H_5 be defined as

$$(H_5)_{j,k} = \text{tr}[(H(x) + I_d)^{-1} A_x^\top (Q(x) \text{diag}(A_{x,*,k}) Q(x) \text{diag}(A_{x,*,j}) Q(x)) \circ I_n] A_x]$$

Then, we have

$$H_5 = A_x^\top (Q(x) \tilde{\Sigma}(A_x) Q(x)) \circ Q(x) A_x$$

Proof. We have

$$\begin{aligned} (H_5)_{j,k} &= \text{tr}[(H(x) + I_d)^{-1} A_x^\top (Q(x) \text{diag}(A_{x,*,k}) Q(x) \text{diag}(A_{x,*,j}) Q(x)) \circ I_n] A_x] \\ &= \text{tr}[\tilde{\Sigma}(A_x) (Q(x) \text{diag}(A_{x,*,k}) Q(x) \text{diag}(A_{x,*,j}) Q(x)) \circ I_n] \\ &= \sum_{l=1}^n \tilde{\Sigma}_{l,l}(A_x) Q_{*,l}(x)^\top \text{diag}(A_{x,*,k}) Q(x) \text{diag}(A_{x,*,j}) Q_{*,l}(x) \\ &= A_{x,*,k}^\top \sum_{l=1}^n \tilde{\Sigma}_{l,l}(A_x) \text{diag}(Q_{*,l}(x)) Q(x) \text{diag}(Q_{*,l}(x)) A_{x,*,j} \\ &= A_{x,*,k}^\top \left(\sum_{l=1}^n \tilde{\Sigma}_{l,l}(A_x) (Q_{*,l}(x) Q_{*,l}(x)^\top \circ Q(x)) \right) A_{x,*,j} \\ &= A_{x,*,k}^\top ((Q(x) \tilde{\Sigma}(A_x) Q(x)) \circ Q(x)) A_{x,*,j}, \end{aligned}$$

where the first step follows from the Lemma statement, the second step follows from the definition of $\tilde{\Sigma}(A_x)$ (see Definition I.7), the third step follows from the definition of $\text{tr}[\cdot]$, and the last step follows from the fact that $\tilde{\Sigma}(A_x)$ is a diagonal matrix (see Definition I.7).

Thus, we have

$$H_5 = A_x^\top (Q(x) \tilde{\Sigma}(A_x) Q(x)) \circ Q(x) A_x. \quad \square$$

I.8 Quantity II for Hessian of F

Recall that

$$\begin{aligned} \frac{\partial H}{\partial x_j} &= -2A_x^\top \text{diag}(\Sigma(A_x)A_{x,*,j})A_x \\ &\quad + A_x^\top \frac{\partial \Sigma(A_x)}{\partial x_j} A_x \\ &= -2A_x^\top \text{diag}(\Sigma(A_x)A_{x,*,j})A_x \\ &\quad + 2A_x^\top \text{diag}((Q^{\circ 2}(x) - \Sigma(A_x))A_{x,*,j})A_x \\ &= A_x^\top \text{diag}(2Q^{\circ 2}(x) - 4\Sigma(A_x))A_{x,*,j}A_x \end{aligned}$$

Lemma I.22. For the second item, we have

$$\left[\frac{\partial}{\partial x} \frac{\partial F}{\partial x} \right]_2 = -4G_1 + 8G_2 + 8G_3 - 16G_4$$

- $G_1 = A_x^\top Q^{\circ 2}(x) \tilde{\Sigma}(A_x)^2 Q^{\circ 2}(x) A_x$
- $G_2 = A_x^\top Q^{\circ 2}(x) \tilde{\Sigma}(A_x)^2 \Sigma(A_x) A_x$
- $G_3 = A_x^\top \Sigma(A_x) \tilde{\Sigma}(A_x)^2 Q^{\circ 2}(x) A_x$
- $G_4 = A_x^\top \Sigma(A_x) \tilde{\Sigma}(A_x)^2 \Sigma(A_x) A_x$

Proof. The proof follows by combining the following lemmas. □

Lemma I.23. Let $(G_1)_{j,k}$

$$(G_1)_{j,k} = \text{tr}[(H + I_d)^{-1} A_x^\top \text{diag}(Q^{\circ 2}(x)A_{x,*,j})A_x (H + I_d)^{-1} A_x^\top \text{diag}(Q^{\circ 2}(x)A_{x,*,k})A_x]$$

Then, we have

$$G_1 = A_x^\top Q^{\circ 2}(x) \tilde{\Sigma}(A_x)^2 Q^{\circ 2}(x) A_x$$

Proof. We can show

$$\begin{aligned} (G_1)_{j,k} &= \text{tr}[(H + I_d)^{-1} A_x^\top \text{diag}(Q^{\circ 2}(x)A_{x,*,j})A_x (H + I_d)^{-1} A_x^\top \text{diag}(Q^{\circ 2}(x)A_{x,*,k})A_x] \\ &= \text{tr}[\tilde{\Sigma}(A_x) \text{diag}(Q^{\circ 2}A_{x,*,j})A_x (H + I_d)^{-1} A_x^\top \text{diag}(Q^{\circ 2}(x)A_{x,*,k})] \\ &= \text{tr}[\tilde{\Sigma}(A_x) \text{diag}(Q^{\circ 2}A_{x,*,j}) \tilde{\Sigma}(A_x) \text{diag}(Q^{\circ 2}(x)A_{x,*,k})] \\ &= A_x^\top Q^{\circ 2}(x) \tilde{\Sigma}(A_x)^2 Q^{\circ 2}(x) A_{x,*,k} \end{aligned}$$

Thus, we have

$$G_1 = A_x^\top Q^{\circ 2}(x) \tilde{\Sigma}(A_x)^2 Q^{\circ 2}(x) A_x. \quad \square$$

The proofs of G_2, G_3, G_4 are similar to G_1 . So we omit the details here.

Lemma I.24. Let $(G_2)_{j,k}$

$$(G_2)_{j,k} = \text{tr}[(H + I_d)^{-1} A_x^\top \text{diag}(Q^{\circ 2}(x)A_{x,*,j})A_x (H + I_d)^{-1} A_x^\top \text{diag}(\Sigma(A_x)A_{x,*,k})A_x]$$

Then, we have

$$G_2 = A_x^\top Q^{\circ 2}(x) \tilde{\Sigma}(A_x)^2 \Sigma(A_x) A_x.$$

Lemma I.25. Let $(G_3)_{j,k}$

$$(G_3)_{j,k} = \text{tr}[(H + I_d)^{-1} A_x^\top \text{diag}(\Sigma(A_x)A_{x,*,j})A_x (H + I_d)^{-1} A_x^\top \text{diag}(Q^{\circ 2}(x)A_{x,*,k})A_x]$$

Then, we have

$$G_3 = A_x^\top \Sigma(A_x) \tilde{\Sigma}(A_x)^2 Q^{\circ 2}(x) A_x.$$

Lemma I.26. Let $(G_4)_{j,k}$

$$(G_4)_{j,k} = \text{tr}[(H + I_d)^{-1} A_x^\top \text{diag}(\Sigma(A_x)A_{x,*,j})A_x (H + I_d)^{-1} A_x^\top \text{diag}(\Sigma(A_x)A_{x,*,k})A_x]$$

Then, we have

$$G_4 = A_x^\top \Sigma(A_x) \tilde{\Sigma}(A_x)^2 \Sigma(A_x) A_x.$$

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The main result of this paper is a class of Dikin walks for log-concave sampling from convex bodies with self-concordant barrier functions. The abstract and introduction reflect the complexity of these walks, and the remainder of the paper is dedicated to prove these assertions.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The limitations are discussed in the last paragraph of Section 4.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: The assumptions are clearly stated in the statement of the theorems: convex body contained in a ball of radius R , density is log-Lipschitz with parameter L , the walk starts from a w -warm starting point. See Theorem 1.1, 1.2 and 1.3. The proofs of Theorem 1.1 and 1.2 results can be found in Section C and D for mixing time, and Section F for per iteration cost. Proofs of Theorem 1.3 can be found in Section G.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [NA]

Justification: The results in this paper are theoretical with no experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification: The results in this paper are theoretical with no experiments.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [NA]

Justification: The results in this paper are theoretical with no experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: The results in this paper are theoretical with no experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [NA]

Justification: The results in this paper are theoretical with no experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: We have carefully read Code of Ethics and the paper is adhered to those principles.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: The potential societal impact is discussed in the last paragraph of Section 4.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The results in this paper are theoretical with no experiments.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: The results in this paper are theoretical with no experiments.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The results in this paper are theoretical with no experiments.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The results in this paper are theoretical with no experiments.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The results in this paper are theoretical with no experiments.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.