

Approaching Cubic Oracle Calls for Isotropic Rounding

October 4, 2026

Abstract

The main results are:

1. For arbitrary convex bodies, an improvement of the previous best query bound from $\tilde{O}(n^{3.5})$ to $\tilde{O}(n^{3.25})$, while also giving a simpler cooling path. This is obtained by directly prompting GPT 5.6 Sol and in Kimi Code with K3max (but not through prompt), and is in [Section 4](#). Its sharpness in [Section A](#) is obtained by directly prompting GPT 5.6 Sol.
2. For symmetric convex bodies, a matrix-regularized Gaussian cooling algorithm improves the rounding cost to $\tilde{O}(n^3)$, and hence gives volume estimation in $\tilde{O}(n^3\epsilon^{-2})$ membership queries ([Section 5](#)). This is obtained by prompting GPT 5.6 Sol to improve an algorithm from a previous submission with the above improvement in asymmetric settings. We have been unable to reproduce it by directly prompting / running Codex, including running a GPT 6 Astra agent swarm up to a budget of \$100.
3. The counter examples in [Section B](#) and [Section C](#) were obtained by directly prompting GPT 5.6 Sol with the statements of those conjectures.

Note. This document contains intermediate results of a project that started around February 2025 involving first Yang P. Liu and Colin Tang, then September 2025 Alicia Stepin, and July 2026 Richard Peng. These results were obtained using a combination of GPT 5.6 Sol, FizzyProver (a heavy harness written by Junzhao Yang) running GPT 5.6 Sol, and Kimi K3 / Kimi Code.

Contents

1	Introduction	3
1.1	Volume Estimation of Convex Bodies	4
1.2	High-Dimensional Geometry	5
2	Preliminaries	5
2.1	Probability	6
2.2	Convex geometry	7
2.3	Log-Concave Distributions	7
2.4	Conductance of Log-Concave Distributions	9
2.5	Additional conventions	13
2.6	Matrices, moments, and affine changes	14
2.7	Oracle models and adaptive randomized calls	14
2.8	Finite-order change of measure	16
2.9	Standard log-concave facts	16
3	Overview	17
3.1	General-Body Scalar Cooling Algorithm	18
3.2	Cubic Symmetric Algorithm	18
3.3	Counterexamples to Previously Stated Conjectures	19
4	General-Body Algorithm: Scalar Cooling	21
4.1	The Full-Density Rounding Interface	22
4.2	Cooling Estimates	28
4.3	Endpoint Comparison and Raw Reconstruction	31
4.4	Algorithm, Induction, and Complexity	33
4.5	Remarks on the Step Size	35
5	Cubic Algorithm for the Symmetric Case	35
5.1	Proof Setup for Even Log-Concave Laws	37
5.2	Quadratic Weights and Analytic Inputs	38
5.3	Regularized Update and Invariant	50
5.4	Contraction of D_{sym} and Comparison with the Original Distribution	51
5.5	Algorithm and Complexity	55
A	Proof of the Sharp Trace Growth Theorem	61
A.1	The lower-bound construction	61
A.1.1	A truncated-Gaussian response function	61
A.1.2	The polytope and its phase switch	62
A.1.3	Centering, endpoint concentration, and whitening	63
B	Counterexamples to The Coordinate Moment Conjecture	65
C	Asymmetric Counterexample to the Ellipsoid-Growth Conjecture	70
C.1	The raw body and its thin direction	70
C.2	The enlarged ellipsoid creates a plateau	72

1 Introduction

Volume estimation of a convex set is a fundamental problem in computer science with a long history [DFK91, LS93, KLS97, LV06b, LV06a, LV07, CV18]. There is a standard model for accessing convex bodies which most of these previous works use. For a convex body K , we assume that there is a routine $\text{QUERY}(x)$ which for a point $x \in \mathbb{R}^n$ answers whether $x \in K$. We call QUERY a *membership oracle*, and we measure the complexity of our algorithm in the number of times we call this membership oracle. All other running time overheads can be made to be polynomial in n , though we do not carefully track it in this article. Also, it is standard to assume that $B_n(0, r) \subseteq K \subseteq B_n(0, R)$ where $B_n(x, r)$ is the n -dimensional ball centered at x of radius r . Generally, the complexities of our algorithms (and previous algorithms) depend logarithmically on R/r . Some assumption like this is necessary to design nontrivial algorithms: otherwise, K could be a very small convex set inside a ball of radius R , and it would take exponential time to even locate a single point inside K .

The work of [DFK91] proved that the volume can be estimated to within a constant factor using polynomially many membership oracle queries. More recent work has focused on obtaining faster fine-grained runtimes for volume estimation. Generally, volume estimation algorithms can be broken down into two main phases, with most works focusing on one phase or another. The first phase is *isotropic rounding*. In this phase, the goal of the algorithm is to compute a linear transformation T on $\mathbb{R}^n$ such that the covariance matrix of the body TK is approximately the identity matrix, i.e., to make the convex body nearly isotropic. The second phase is to design algorithms to estimate the volume of a nearly-isotropic convex body.

The complexity of the second phase, i.e., estimating the volume of a nearly-isotropic convex body, is reasonably well-understood by now. In this case, the *Gaussian cooling* algorithm of Cousins-Vempala [CV18] gives a $(1 + \varepsilon)$ -multiplicative approximation to the volume in $\tilde{O}(n^3 \varepsilon^{-2})$ membership queries. The complexity of the first isotropic rounding step has historically been larger. Before this work, the best general bound was due to [JLLV26], using $\tilde{O}(n^{3.5} \log(nR/r))$ membership queries. In the remainder of the introduction, we assume that $R/r \leq \text{poly}(n)$ for simplicity, and use $\tilde{O}(\cdot)$ to suppress $\text{poly}(\log n)$ factors throughout the paper.

The main algorithmic results of this paper are two near-cubic rounding algorithms. The first works for arbitrary convex bodies and estimates the uncentered second moment in $\tilde{O}(n^{3.25})$ membership queries. The second uses symmetry more strongly and obtains a cubic rounding bound for symmetric convex bodies.

The next theorem concerns the general-body rounding.

Theorem 1. *The algorithm `ASYMMETRICSECOND MOMENT` (Figure 2), given membership query access to a convex set $K \subseteq \mathbb{R}^n$ satisfying $B_n(x_0, r) \subseteq K \subseteq B_n(x_0, R)$ and a failure parameter $\beta \in (0, 1)$, computes a positive semidefinite matrix $\widehat{M}$ satisfying*

$$0.99\widehat{M} \preceq \mathbb{E}_{x \sim K} [xx^\top] \preceq 1.01\widehat{M}.$$

It succeeds with probability at least $1 - \beta$ and makes

$$O\left(n^{13/4} \log\left(2 + \frac{nR}{r\beta}\right)^{O(1)}\right)$$

membership query calls. On the same success event, the algorithm also returns an invertible affine map $\mathcal{T}$ such that the uniform law on $\mathcal{T}(K)$ has covariance in $[1/1.01, 1.01]I$ and mean of norm at most $1/250$; in particular $B(0, 0.99) \subseteq \mathcal{T}(K)$ and $\mathbb{E}_{x \sim \mathcal{T}(K)} [\|x\|^2] \leq 1.02n$.

The next theorem concerns the cubic symmetric rounding.

Theorem 2. *The algorithm SYMMETRIC-COVARIANCE (Figure 3), given membership query access to a symmetric convex set $K \subseteq \mathbb{R}^n$ satisfying $B_n(0, r) \subseteq K \subseteq B_n(0, R)$ and a failure parameter $\beta \in (0, 1)$, computes a matrix Σ satisfying*

$$0.99\Sigma \preceq \mathbb{E}_{x \sim K} [xx^\top] \preceq 1.01\Sigma.$$

It succeeds with probability at least $1 - \beta$ and makes

$$O\left(n^3 \log\left(2 + \frac{nR}{r\beta}\right)^{O(1)}\right)$$

membership query calls.

The next corollary concerns the volume estimation.

Corollary 1.1. *There is an algorithm that, given membership query access to a convex set $K \subseteq \mathbb{R}^n$ satisfying $B_n(x_0, r) \subseteq K \subseteq B_n(x_0, R)$ and an error parameter $\varepsilon \in (0, 1)$, outputs with probability at least $3/4$ a number within a factor $1 \pm \varepsilon$ of the volume of K after making*

$$O\left(n^{13/4} \log(2 + nR/r)^{O(1)} + n^3 \varepsilon^{-2} \log(n/\varepsilon)^{O(1)}\right)$$

membership queries. If K is symmetric and the two balls are centered at the origin, the query bound improves to

$$O\left(n^3 \log(2 + nR/r)^{O(1)} + n^3 \varepsilon^{-2} \log(n/\varepsilon)^{O(1)}\right).$$

The corollary is proved at the end of Section 4 from the two theorems, the inner-ball bound of Fact 2.24, and the Gaussian cooling volume algorithm of [CV18] recalled in Fact 2.25.

The Gaussian cooling perspective gives a simpler route to rounding than the previous ellipsoid-growth framework and improves the general membership-query bound from $\tilde{O}(n^{3.5})$ to $\tilde{O}(n^{13/4})$. The algorithms are discussed in Section 3 and presented in Sections 4 and 5.

1.1 Volume Estimation of Convex Bodies

In this section we discuss volume estimation algorithms for polytopes and how various techniques (including ours) fit within the context of prior work. The main primitive for volume estimation is *sampling* from a convex body K using membership oracle queries. Sampling from a convex body is done with Markov chain methods, for example the *Ball-Walk* or *hit-and-run* algorithm. The mixing time of these methods depends on the conductance of the polytope K , where conductance refers to minimum surface area to small volume ratio of a partition of K into two equal-size pieces. So, in discussion of these kinds of volume algorithms, it was important to understand the conductance of convex bodies in high dimensions. The famous KLS conjecture asks whether this is always at least $\Omega(1)$, if K is *isotropic*, i.e., the uniform distribution has covariance equal to I . Exciting recent progress has led to a bound of $\Omega(\log^{-1/4} n)$ [Let26] (see Section 1.2 for more discussion on these lines of work), which in turn implies that the Ball-Walk mixes in $\tilde{\tilde{O}}(n^2)$ steps for isotropic bodies. Volume estimation requires sampling several points, and it is known that the volume of an isotropic (even well-rounded) convex body can be estimated in $\tilde{O}(n^3)$ membership queries, using the Gaussian cooling algorithm of Cousins-Vempala [CV18], as mentioned above. However, the

question remains on how to transform any convex body to isotropic position efficiently, and that is the focus of this paper.

There are several related lines of work on sampling general log-concave distributions in different settings, such as sampling and volume estimation in polytopes (that is, convex bodies explicitly described by their m faces) [KN12, Nar16, NR17, SV16, CDWY17, CDWY18, LLV20, LV22, LV18, KLSV22, KLSV23, GKV24, MV23, GKM⁺24, JC25, GC11]. Several different algorithms are known in this setting, often motivated by barrier functions from interior point methods.

1.2 High-Dimensional Geometry

The KLS conjecture of Kannan-Lovasz-Simonovits [KLS95] asks about the minimal conductance of a cut of a convex body K . The conjecture claims that if K is isotropic, then the conductance is $\Omega(1)$. This is an important conjecture, because the KLS constant not only governs the mixing time of Markov chains for sampling from K , but captures how well functions distributed according to K concentrate. For example, the thin-shell bound (concentration of the norm)

$$\mathbb{E}_{x \sim K} \left[(\|x\|_2 - n^{1/2})^2 \right] \leq O(1)$$

would follow from the KLS conjecture. Our sampling bounds use proved estimates toward the KLS conjecture. The cooling analysis uses Brascamp–Lieb and Paouris inequalities in the general case, and an even-sector Poincaré inequality in the symmetric case.

Recent progress towards the KLS conjecture started when Eldan introduced a new method for proving lower bounds on the conductance called *stochastic localization*. This initiated a chain of work [Eld13, LV24, Che21, KL22, JLV22, Kla23, KL25b, Let26] culminating in the current $\Omega(\log^{-1/4} n)$ KLS lower bound. Related breakthroughs resolved the thin-shell and slicing conjectures [Kla06, Gua24, KL25a, Fle10, GM11, EL14]. These analytic inputs are developed in Sections 3 to 5.

2 Preliminaries

General notation. The notation of $\tilde{O}(\cdot)$ throughout indicates hidden dependence on the error parameter and logarithmic factors. We use $\mathbb{R}_{\geq 0}$ to denote the nonnegative real numbers. For a twice-differentiable function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ we let ∇f and $\nabla^2 f$ denote its gradient and Hessian respectively. We call a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ *locally Lipschitz* if, for every $x \in \mathbb{R}^n$, there is a neighborhood N_x of x and a constant $L_x < \infty$ such that

$$|f(y) - f(z)| \leq L_x \|y - z\|_2 \quad \text{for all } y, z \in N_x.$$

We first restate some notions and facts fundamental to high dimensional geometry.

For positive semidefinite (PSD) matrices $A, B \in \mathbb{R}^{n \times n}$ we say that A and B are α -spectral-approximations (or α -approximations where it does not create confusion) for $\alpha \geq 1$ if $\alpha^{-1}A \preceq B \preceq \alpha A$ where $\preceq$ denotes the Loewner ordering of PSD matrices. For a PSD matrix A we let $\|A\|_{\text{op}} = \max_{\|v\|_2=1} v^\top A v$ denote its operator norm.

Furthermore, we abbreviate two sided multiplicative approximations using α . For non-negative scalars $a, b \in \mathbb{R}_{\geq 0}$, we use $a \approx_\alpha b$ to denote $\frac{1}{\alpha}a \leq b \leq \alpha a$, and similarly for matrices $A, B \in \mathbb{R}^{n \times n}$, we use $A \approx_\alpha B$ to denote $\frac{1}{\alpha}A \preceq B \preceq \alpha A$.

The Frobenius norm of a matrix A is denoted as $\|A\|_F := \text{Tr}(A^\top A)^{1/2}$. The operator norm of matrix A is denoted by $\|A\|_{\text{op}}$ and defined to be

$$\|A\|_{\text{op}} := \sup_{u \in \mathbb{R}^n: \|u\|_2=1} \|Au\|_2.$$

When the domain of integration is omitted, integrals are taken over $\mathbb{R}^n$.

2.1 Probability

A probability density function, or distribution, on $\mathbb{R}^n$ is a measure $p : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ such that $\int_{\mathbb{R}^n} p(x) dx = 1$. We use *with high probability* (abbreviated w.h.p.) to mean that an event holds with probability $1 - n^{-C}$ for any constant $C > 0$ (possibly by adjusting constants in the definition of the event).

We measure distances between two probability density functions using total variation distance. The next definition concerns the total variation distance.

Definition 2.1. For two probability density functions p, q on $\mathbb{R}^n$ we define their total-variation distance as

$$\text{TV}(p, q) = \sup_A \left| \int_A (p(x) - q(x)) dx \right|,$$

where $A \subseteq \mathbb{R}^n$ ranges over all measurable subsets of $\mathbb{R}^n$. Alternatively, we can write

$$\text{TV}(p, q) = \frac{1}{2} \int_{\mathbb{R}^n} |p(x) - q(x)| dx = \frac{1}{2} \|p - q\|_1.$$

Given a probability density function p on $\mathbb{R}^n$, and a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, we use $\mathbb{E}_{x \sim p}[f(x)]$, $\text{Var}_p[f(x)]$ to denote expectation and variance of $f(x)$. Precisely, $\mathbb{E}_{x \sim p}[f(x)] = \int_{\mathbb{R}^n} f(x)p(x)dx$ and

$$\text{Var}_{x \sim p}[f(x)] = \int_{\mathbb{R}^n} (f(x) - \mathbb{E}_{x \sim p}[f(x)])^2 p(x) dx = \mathbb{E}_{x \sim p}[f(x)^2] - \mathbb{E}_{x \sim p}[f(x)]^2.$$

Sometimes to avoid notational clutter we will just write $\mathbb{E}_p[f(x)]$ instead of $\mathbb{E}_{x \sim p}[f(x)]$ when the variable being drawn from p is clear from context.

Now we define the mean and covariance matrix of a probability density function on $\mathbb{R}^n$

Definition 2.2. Given a probability density function p on $\mathbb{R}^n$, we denote its mean m_p as

$$\mathbb{R}^n \ni m_p := \mathbb{E}_{x \sim p}[x]$$

and covariance matrix Cov_p as

$$\mathbb{R}^{n \times n} \ni \text{Cov}_p := \mathbb{E}_{x \sim p}[(x - m_p)(x - m_p)^\top] = \mathbb{E}_{x \sim p}[xx^\top] - m_p m_p^\top.$$

Matrix concentration bounds state that a sum of random matrices is close to the sum of their expectations, assuming the norm of each matrix in their sum is not too large. Below we use the following version of Matrix Bernstein.

The next fact concerns the matrix Bernstein inequality, symmetric case of Theorem 1.6.2 of [Tro15]¹.

Fact 2.3. Let $Y_1, \dots, Y_m \in \mathbb{R}^{n \times n}$ be independent mean 0 random Hermitian matrices such that for each $i \in [m]$, $\|Y_i\|_{\text{op}} \leq R$ with probability 1. Then

$$\mathbb{P}_{Y_1 \dots Y_m} \left[\left\| \sum_{i=1}^m Y_i \right\|_{\text{op}} \geq t \right] \leq 2n \cdot \exp \left(- \frac{t^2/2}{\|\sum_i \mathbb{E}_{Y_i}[Y_i^2]\|_{\text{op}} + Rt/3} \right)$$

¹[arXiv:1501.01571v1](https://arxiv.org/abs/1501.01571v1)

2.2 Convex geometry

Let $B_n(x, \delta)$ denote the n -dimensional Euclidean ball with radius δ centered at x . Alternatively $B_n(x, \delta) = \{z \in \mathbb{R}^n : \|x - z\| \leq \delta\}$.

We say that a probability distribution (or density function) p on $\mathbb{R}^n$ is *isotropic* if $m_p = \vec{0}$ and $\text{Cov}_p = I_n$ where $I_n \in \mathbb{R}^{n \times n}$ is the $n \times n$ identity matrix. We say that p is *near-isotropic* up to a factor of α if $m_p = \vec{0}$ and $\text{Cov}_p \approx_\alpha I_n$.

Let K be a convex body in $\mathbb{R}^n$. Its associated indicator function $1_K : \mathbb{R}^n \rightarrow \mathbb{R}$ is given by

$$1_K(x) := \begin{cases} 1 & \text{if } x \in K \\ 0 & \text{otherwise} \end{cases},$$

and its volume is given by $\text{Vol}(K) := \int_{\mathbb{R}^n} 1_K(x) dx$. We say that K is isotropic if its associated probability density function $p_K := \frac{1}{\text{Vol}(K)} 1_K$ is isotropic.

2.3 Log-Concave Distributions

In this section we collect some standard facts on log-concave distributions. We first define a log-concave function. The next definition concerns the log-concave.

Definition 2.4. A function $f : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ is *log-concave* if for all $x, y \in \mathbb{R}^n$ and $\alpha \in [0, 1]$ it satisfies

$$f(\alpha x + (1 - \alpha)y) \geq f(x)^\alpha f(y)^{1-\alpha}$$

Throughout we refer to a log-concave probability density function as a *log-concave distribution* for brevity. We now define a strongly log-concave density.

The next definition concerns M -strongly log-concave.

Definition 2.5. For a PSD matrix M , we call a log-concave function $f(x)$ *M -strongly log-concave* if f is twice-differentiable and satisfies $\nabla^2(-\log(f(x))) \succeq M$ for all $x \in \mathbb{R}^n$. For a positive real number t , we say that f is *t -strongly log-concave* if $\nabla^2(-\log(f(x))) \succeq tI$ for all $x \in \mathbb{R}^n$.

We note that log-concave distributions have sub-exponential tails. The next lemma states that log-concave distributions have sub-exponential tails.

Lemma 2.6. *There exists a universal constant $c > 0$ such that for every isotropic log-concave distribution μ , for every unit vector $v \in \mathbb{R}^n$, and for every $t \geq 0$, the following holds:*

$$\mathbb{P}_{z \sim \mu} [|\langle v, z \rangle| \geq t] \leq \exp(-ct).$$

We require the following bound on the values achieved by log-concave distributions due to Fradelizi (we cite the version in [Pao12, Theorem 3.3]).

Lemma 2.7. *Let $\mu(x) \propto e^{-\phi(x)}$ be log-concave with mean 0. Let $x^* := \arg\min_{x \in \mathbb{R}^n} \phi(x)$. Then $\mu(x^*) \leq e^n \mu(0)$. Equivalently, $\phi(0) \leq n + \phi(x^*)$.*

We require the Brascamp-Lieb inequality. This will be applied to bound the covariance of M -strongly log-concave distributions. The next lemma concerns the Brascamp-Lieb inequality.

Lemma 2.8. *For any log-concave distribution $p : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ of the form $p(x) = \exp(-\phi(x))$, and any differentiable function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ it holds that*

$$\text{Var}_{x \sim p} [f(x)] \leq \mathbb{E}_{x \sim p} [\langle \nabla f(x), \nabla^2 \phi(x)^{-1} \nabla f(x) \rangle].$$

The next corollary concerns the Brascamp-Lieb inequality for covariances.

Corollary 2.9. *Let $p(x) \propto e^{-\phi(x)}$ be a probability density on $\mathbb{R}^n$ and suppose p is M -strongly log-concave for some positive definite matrix M . Then*

$$\text{Cov}_p \preceq M^{-1}.$$

Proof. We fix $v \in \mathbb{R}^n$ and set $f(x) = \langle x, v \rangle$. Then $\nabla f(x) = v$, so by the Brascamp-Lieb inequality (Lemma 2.8)

$$\text{Var}_{x \sim p} [f(x)] \leq \mathbb{E}_{x \sim p} [\langle \nabla f(x), (\nabla^2 \phi(x))^{-1} \nabla f(x) \rangle] = \mathbb{E}_{x \sim p} [v^\top (\nabla^2 \phi(x))^{-1} v] \leq v^\top M^{-1} v$$

where the final inequality comes from the definition of M -strongly log-concave. Since $v^\top \text{Cov}_p v = \text{Var}_{x \sim p} [\langle x, v \rangle]$, this yields the claimed spectral bound $\text{Cov}_p \preceq M^{-1}$. $\square$

We require the standard fact that log-concavity is preserved under marginals [Pr673]. The next lemma concerns the Prekopa-Leindler inequality.

Lemma 2.10. *Let p be a log-concave probability density on $\mathbb{R}^n$ and let $v \in \mathbb{R}^n$. Then the distribution $\langle x, v \rangle$ on $\mathbb{R}$ (where $x \sim p$) is a log-concave probability measure.*

For $v \neq 0$, this measure has density

$$\hat{p}(t) = \frac{1}{\|v\|_2} \int_{\{x \in \mathbb{R}^n : \langle x, v \rangle = t\}} p(x) d\mathcal{H}^{n-1}(x),$$

where $\mathcal{H}^{n-1}$ is the induced Euclidean surface measure. For $v = 0$, it is the point mass at zero.

We require the Paouris Inequality which allows us to bound the tail behavior of the Euclidean norm of samples drawn from an isotropic log-concave distribution [Pao06].

The next lemma concerns the Paouris inequality.

Lemma 2.11. *There is an absolute constant $C > 0$ such that for any isotropic log-concave distribution p on $\mathbb{R}^n$ and every $t \geq 1$ it holds that*

$$\mathbb{P}_{x \sim p} [\|x\|_2 \geq Ctn^{1/2}] \leq \exp(-tn^{1/2}).$$

We need the following inequality to bound moments of log-concave variables [LV07, Theorem 5.22].

The next lemma concerns the reverse Hölder.

Lemma 2.12. *Let p be a log-concave probability density on $\mathbb{R}^n$. For every positive integer k ,*

$$(\mathbb{E}_p [\|x\|_2^k])^{1/k} \leq 2k \mathbb{E}_p [\|x\|_2] \leq 2k (\mathbb{E}_p [\|x\|_2^2])^{1/2}.$$

We use the following Gaussian-source form of Caffarelli's contraction theorem [Caf00, Theorem 11].

Theorem 3. *Let $V : \mathbb{R}^n \rightarrow \mathbb{R}$ be twice continuously differentiable with $\nabla^2 V(x) \succeq 2I$ for every $x \in \mathbb{R}^n$, and define*

$$d\gamma(x) = \pi^{-n/2} e^{-\|x\|_2^2} dx, \quad d\mu(x) = Z^{-1} e^{-V(x)} dx, \quad Z = \int_{\mathbb{R}^n} e^{-V(x)} dx.$$

The quadratic-cost Brenier map $T = \nabla \varphi$ satisfying $T_\# \gamma = \mu$ is unique up to a γ -null set and has a 1-Lipschitz representative. In particular,

$$\|DT(x)\|_{\text{op}} \leq 1 \quad \text{for } \gamma\text{-almost every } x.$$

2.4 Conductance of Log-Concave Distributions

We start by defining the conductance (KLS constant) of a log-concave distribution.

Definition 2.13. Let p be a log-concave distribution on $\mathbb{R}^n$. We define its KLS constant ψ_p as

$$\frac{1}{\psi_p} := \inf_{S \subseteq \mathbb{R}^n} \frac{\int_{\partial S} p(x) dx}{\min \left\{ \int_S p(x) dx, \int_{\mathbb{R}^n \setminus S} p(x) dx \right\}}.$$

Here, ∂S denotes the *boundary* of the set S . The KLS conjecture [KLS95] asks whether $\psi_p = O(\|\text{Cov}_p\|_{\text{op}}^{1/2})$ for all log-concave distributions p . While this has not been proven, there has been significant recent work on tightening the bound. The current best bound is due to Letwin [Let26].

Theorem 4. *If p is a log-concave distribution, then $\psi_p \leq O(\|\text{Cov}_p\|_{\text{op}}^{1/2} \log^{1/4} n)$.*

The next inequality converts conductance bounds into variance bounds for smooth test functions. The next lemma concerns the Poincare inequality.

Lemma 2.14. *Let p be a log-concave distribution on $\mathbb{R}^n$ and let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a differentiable function. Then*

$$\text{Var}_{x \sim p} [f(x)] \leq O(\psi_p^2) \cdot \mathbb{E}_{x \sim p} [\|\nabla f(x)\|_2^2].$$

Two algorithmic black boxes from [KV25] drive the full-density rounding primitive of [Section 4.1](#): the near-cubic isotropic rounding of a well-rounded log-concave law, and the estimation of the mean and covariance of a well-rounded law with bounded covariance from one run of a Markov chain. We state both with the exact hypotheses under which they are proved in [KV25]; the reductions of our inputs to these hypotheses are carried out in [Section 4.1](#). The constants in the following definition are those of [KV25, Assumption 4.4].

The next definition concerns well-rounded log-concave distributions.

Definition 2.15. Let $\pi \propto e^{-W}$ be a log-concave probability law on $\mathbb{R}^n$, let $z \in \mathbb{R}^n$, and let $\rho, H, B > 0$. We say that π is (ρ, H, B) -*well-rounded around* z if

$$B(z, \rho) \subseteq \{x \in \mathbb{R}^n : W(x) \leq \inf W + Hn\} \quad \text{and} \quad \mathbb{E}_{X \sim \pi} [\|X - z\|^2] \leq Bn.$$

We say that π is *grounded around z with respect to W* if $\rho = 1$ and $H = 10$ can be taken, and in addition $\min W \geq 0$ and $W \leq 10n$ on the support of π , so that the support of π equals the level set $\{W \leq \inf W + 10n\}$.

Note that well-roundedness is a relaxation of near-isotropy.

The following adapts the well-rounded isotropizer of [KV25, Algorithm 4 and Lemma 4.11] under their Assumption 4.4. We record its output through an affine map and give the implementation and proof below.

Theorem 5. *There is a universal constant $c_{\text{iso}} > 0$ and, for every $C \geq 1$, there are constants $A_C, a_C > 0$ such that the following holds. Let $\pi \propto e^{-W}$ be a log-concave law on $\mathbb{R}^n$ that is grounded around a known point z with respect to W ([Definition 2.15](#)), with $\mathbb{E}_{X \sim \pi} [\|X - z\|^2] \leq C^2 n$ for the known constant C , and let W be given by an evaluation oracle. There is a randomized algorithm which, on an event of probability at least $1 - c_{\text{iso}} n^{-1/2}$, returns an invertible affine map $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$ with*

$$C_{F_{\#}\pi} \approx_{1.01} I_n \quad \text{and} \quad \|m_{F_{\#}\pi}\| \leq 0.01,$$

and the expected number of evaluation queries it makes on this event is at most $A_C n^3 \log^{a_C}(2n)$.

Writing $F(x) = A(x - \bar{x})$ with A invertible and $\Sigma = (A^\top A)^{-1}$, the conclusion reads $\Sigma \approx_{1.01} C_\pi$ and $\|\Sigma^{-1/2}(\bar{x} - m_\pi)\| \leq 0.01$, because $C_{F\#\pi} = AC_\pi A^\top$ and $\|\Sigma^{-1/2}v\| = \|Av\|$ for every v .

Implementation and proof of Theorem 5. We specify the rounding primitive using the sampling and stationary covariance-estimation results of [KV25, KZ24]. Translate the original input by its known point z , so its potential is $W_0(x) = W(x + z)$ and $\mathbb{E}_{\pi_0}[\|X\|^2] \leq C^2 n$. All affine maps below are recorded and composed with this translation. Set

$$L = \max\{4, \log(2n)\}, \quad q = L^{-1}, \quad d_0 = \frac{nq^2}{100}, \quad r_* = \frac{q^2\sqrt{n}}{100}.$$

Start with π_0 and $r_0 = 1$. At stage i , write m_i, C_i for the true mean and covariance of the current law π_i . While $r_i < r_*$, estimate its covariance by a positive semidefinite empirical matrix S_i satisfying

$$-\frac{1}{10}C_i - d_0I \preceq S_i - C_i \preceq \frac{1}{10}C_i + d_0I. \quad (2.1)$$

Let P_i be the spectral projection of S_i onto eigenvalues at most n , put $M_i = I + P_i$, and update

$$\pi_{i+1} = (M_i)_{\#}\pi_i, \quad r_{i+1} = 2(1 - q)r_i.$$

The estimates, their probability of success, and their query cost are justified below. Their required accuracy differs from a constant additive multiple of n only by a logarithmic factor.

We first give the deterministic geometric argument on the event that (2.1) holds at every completed stage. The covariance obeys $C_{i+1} = M_i C_i M_i$, whence $\text{tr } C_{i+1} \leq 4 \text{tr } C_i$. The lower estimate in (2.1) gives $P_i C_i P_i \preceq 2nP_i$. Consequently,

$$\|C_{i+1}\|_{\text{op}} = \|C_i^{1/2}(I + 3P_i)C_i^{1/2}\|_{\text{op}} \leq \|C_i\|_{\text{op}} + 6n.$$

There are $O(L)$ stages, because $2(1 - q) \geq 3/2$ and $r_* \leq \sqrt{n}$. Since $(1 - q)^{-O(L)} = O(1)$, these observations imply

$$\text{tr } C_i = O_C(r_i^2 n), \quad \|C_i\|_{\text{op}} = O_C(nL). \quad (2.2)$$

We next prove that the support of π_i contains a ball $B(c_i, r_i)$ for some center c_i ; the algorithm need not know this center. The assertion holds initially. The support also contains $m_i + e^{-1}C_i^{1/2}B(0, 1)$ by the log-concave support-ball bound [KV25, Lemma A.10]. Suppress the index i and put $M = I + P$. After applying M , the support contains the two ellipsoids

$$A = MB(c_i, r_i), \quad B = M(m_i + e^{-1}C_i^{1/2}B(0, 1)),$$

and hence their weighted Minkowski sum $(1 - q/2)A + (q/2)B$. Center this sum at $c' = (1 - q/2)Mc_i + (q/2)Mm_i$. For a unit vector u , write $a = \|(I - P)u\|$. The centered support function of A is $r_i\|Mu\| = r_i\sqrt{4 - 3a^2} \geq r_i(2 - \frac{3}{2}a^2)$. Since $S_i \succeq n(I - P)$, the upper estimate in (2.1) gives

$$u^\top MC_i Mu \geq \frac{na^2 - 4d_0}{1.1}.$$

If $a^2 \leq 8d_0/n$, the contribution of A alone is at least $2r_i - qr_i - 0.12q^2r_i \geq 2(1 - q)r_i$. Otherwise the centered support function of B is at least $\sqrt{n}a/(e\sqrt{2.2})$. Its weight $q/2$ compensates the term $\frac{3}{2}r_ia^2$, because

$$r_i < r_* = \frac{q^2\sqrt{n}}{100} \leq \frac{q\sqrt{n}}{3e\sqrt{2.2}}.$$

The support function of the sum is therefore at least $2(1 - q)r_i$ in every unit direction. Thus the sum contains $B(c', 2(1 - q)r_i)$, proving the induction. This argument uses support functions of the full ellipsoids and does not require their principal sections to be large.

Here is the sampling implementation. Before each covariance-estimation stage, including the final one, independently apply [KV25, Theorem 3.16] to the original input W_0 , with accuracy $\log 2$. Its hypotheses hold with known center 0 and $R^2 = C^2 n$. The construction uses evaluations of W_0 , including $W_0(0)$, and does not require the value of $\inf W_0$. It generates a fresh ideal sample from a law bounded by twice the original exponential lift. Push this sample forward by the current recorded linear map, leaving the lift coordinate t unchanged. Conditioned on the preceding stages, the resulting ideal starting law is bounded by twice the lift of π_i .

For sampling at stage i , rescale only the x -coordinates by r_i^{-1} . The ground set then has unit inradius. By [KV25, Lemma 2.5], the lifted covariance has operator norm at most

$$2(\|C_i\|_{\text{op}}/r_i^2 + 160).$$

The exact proximal kernel is reversible. Its χ^2 contraction, the Poincare bound, and [KV25, Lemma 2.14] provide a thinned chain with spectral gap at least 0.99, using

$$\tilde{O}(n^2 (\|C_i\|_{\text{op}}/r_i^2 + 1))$$

queries per retained sample, with logarithmic dependence on its abort budget. The mixing lengths and sample counts use deterministic upper bounds from (2.2), computed from the known C, n, r_i and stage index; they do not require the unknown covariance C_i . Undoing the known coordinate rescaling preserves this spectral gap. The resulting stationary law is the unscaled exponential lift, whose covariance has trace at most $\text{tr } C_i + 160$ and operator norm at most $2(\|C_i\|_{\text{op}} + 160)$. Apply [KZ24, Theorem 8] to this lifted chain with multiplicative accuracy $1/10$ and additive accuracy d_0 . Using (2.2), its sample requirement is $k_i = \tilde{O}_C(r_i^2)$. Compressing the empirical covariance to the x -coordinates gives (2.1) with stationary-start failure probability $O(1/n)$. Since $r_i^2 \leq n$ at every updating stage, the total sampling cost of that stage is $\tilde{O}_C(n^3)$.

At termination, $r_i \geq \max\{1, r_*\}$. The per-sample bound above is now $\tilde{O}_C(n^2)$. Use $\tilde{O}(n)$ retained samples for the final empirical mean $\bar{m}$ and covariance S . If S is not positive definite, declare failure. For their analysis, whiten the x -coordinates by $C_i^{-1/2}$ and leave t unchanged; this is an invertible change of variables of the stationary chain and preserves its spectral gap. The resulting lifted covariance has operator norm at most 322. The stationary covariance and mean estimates [KZ24, Corollaries 10 and 33], with sufficiently small fixed accuracy, therefore give

$$(1 - \varepsilon)C_i \preceq S \preceq (1 + \varepsilon)C_i, \quad \|C_i^{-1/2}(\bar{m} - m_i)\| \leq \varepsilon, \quad \varepsilon = 10^{-3},$$

except with probability $O(1/n)$. Whitening here is used only in the analysis: the algorithm computes the ordinary empirical statistics in its current coordinates. The final map $x \mapsto S^{-1/2}(x - \bar{m})$ has covariance between $(1 + \varepsilon)^{-1}I$ and $(1 - \varepsilon)^{-1}I$, and its mean has norm at most $\varepsilon/\sqrt{1 - \varepsilon} < 0.01$. Compose it with all recorded maps to obtain the required F .

Finally, we account for failure and queries without conditioning a sampler on successful rejection steps. Fix a deterministic bound $J = \lceil 4L \rceil$ on the number of covariance-estimation stages. Assign every warm-start run and every sampling run an abort budget $n^{-2}/(2J)$. Ideal processes use rejection sampling without a cutoff. Couple each implemented run to its ideal run using the same proposals until a cutoff is exceeded. The per-step estimates in [KV25, Lemmas 2.14 and 3.10], summed along the ideal processes on accurate preceding stages, bound the probability of an abort

before any earlier statistical failure by n^{-2} . For any event E of a stage's ideal chain, its conditional probability, given the preceding stages, is at most twice the corresponding stationary-start probability: if $\nu_i \leq 2\hat{\pi}_i$, then

$$\mathbb{P}_{\nu_i}(E) = \int \mathbb{P}_x(E) \nu_i(dx) \leq 2 \int \mathbb{P}_x(E) \hat{\pi}_i(dx).$$

Thus each first statistical failure, conditional on accurate preceding estimates, has probability $O(1/n)$. A union bound gives overall failure probability $O(J/n) + n^{-2} = O(n^{-1/2})$. Each fresh start costs $\tilde{O}_C(n^3)$, as does each updating stage and the final stage. The same per-step estimates bound the expected capped query cost along accurate prefixes; summing over $J = O(L)$ stages gives the stated expected $\tilde{O}_C(n^3)$ query bound on the success event. As usual, fixed small dimensions are absorbed into the universal constants; the later application explicitly pads to a sufficiently large dimension. $\square$

The second black box is the final step of [KV25, Algorithm 4], run on a law whose covariance is already known to be bounded: a warm start from [KV25, Theorem 3.16], one run of the proximal sampler of [KV25, Section 2], and the dependent-sample mean and covariance estimator of [KZ24, Corollaries 10 and 33], as recorded in [KV25, Lemma 4.6]. We state the resulting guarantee and then record how it is assembled from the cited statements.

Theorem 6. *There is a universal constant $c_{\text{est}} > 0$ and, for every $c_+ \geq 1$ and every $\xi \in (0, 1)$, there are constants $A, a > 0$ such that the following holds. Let $\pi \propto e^{-W}$ be a log-concave law on $\mathbb{R}^n$ with $\|C_\pi\|_{\text{op}} \leq c_+$, let W be given by an evaluation oracle, and let a point z and a radius $R > 0$ be known with*

$$B(z, 1) \subseteq \{W \leq \inf W + 10n\} \quad \text{and} \quad \mathbb{E}_{X \sim \pi} [\|X - z\|^2] \leq R^2.$$

For every $\eta \in (0, 1)$ there is a randomized algorithm which, on an event of probability at least $1 - 2\eta - c_{\text{est}}/n$, returns $(\hat{m}, \hat{C})$ with

$$(1 - \xi)C_\pi \preceq \hat{C} \preceq (1 + \xi)C_\pi \quad \text{and} \quad \|\hat{m} - m_\pi\| \leq \xi,$$

and the expected number of evaluation queries it makes on this event is at most $A(n^2 R^2 + n^3) \log^a(2n/\eta)$.

The theorem is extracted from [KV25] as follows. Let $\hat{\pi}(x, t) \propto e^{-nt}$ on $\mathcal{K} = \{(x, t) \in \mathbb{R}^n \times \mathbb{R} : W(x) \leq nt\}$ be the exponential reduction (exp-red) of [KV25, Section 2.1]: it is log-concave, its x -marginal is π , and its covariance $\hat{\Sigma}$ satisfies $\|\hat{\Sigma}\|_{\text{op}} \leq 2(\|C_\pi\|_{\text{op}} + 160) \leq 2c_+ + 320$ by [KV25, Lemma 2.5]. The two displayed hypotheses are those of [KV25, Problems 2.1 and 3.1] with $x_0 = z$. First, use [KV25, Theorem 3.16] with accuracy $\log 2$ and abort budget η to generate a warm start. Its ideal process uses rejection sampling without a cutoff and has output law $\nu \leq 2\hat{\pi}$; the implemented algorithm agrees with that output except on an abort event of probability at most η . The expected query bound is $\tilde{O}(n^2(R^2 \vee n) \log^8(1/\eta))$.

Let P be the exact proximal-sampler kernel for $\hat{\pi}$. It is reversible, and [KV25, Lemma 2.9], applied in dimension $n + 1$, gives

$$\chi^2(\lambda P \|\hat{\pi}) \leq (1 + h/C_{\text{PI}}(\hat{\pi}))^{-2} \chi^2(\lambda \|\hat{\pi}).$$

The covariance bound above and Theorem 4 and Lemma 2.14 imply

$$C_{\text{PI}}(\hat{\pi}) = O((c_+ + 1) \log^{1/2}(n + 2)).$$

Choose N so that P^N has spectral gap at least 0.99, using this known upper bound in place of the unknown Poincaré constant. Set

$$\xi' = \frac{\xi}{\sqrt{2c_+ + 320}}, \quad k \asymp (\xi')^{-2}(n+1) \log^2 \frac{n+2}{\xi'} \log^2(n+2).$$

Choose the step size and rejection cutoff by [KV25, Lemma 2.14] for kN steps, warmness $M = 2$, and total abort budget η . These choices can be made with $N = O((c_+ + 1)n^2 \log^{O(1)}((n+2)/\eta))$, where constants may depend on ξ and c_+ .

For the ideal chain started from ν , every marginal is 2-warm because $\nu P^j \leq 2\hat{\pi} P^j = 2\hat{\pi}$. The per-step estimates in that lemma, summed along the ideal chain, bound the probability that any rejection cutoff is exceeded by η and the expected implemented query count by $\tilde{O}(kN \log^2(kN/\eta))$. Couple the implemented and ideal chains using the same proposals until the first abort. This argument does not condition the chain on successful rejection steps.

Let $Z_1, \dots, Z_k$ be every N -th ideal iterate. Under a stationary start, [KV25, Lemma 4.6] gives empirical mean $\bar{m}$ and covariance $\bar{\Sigma}$ satisfying

$$\begin{aligned} -\xi' \hat{\Sigma} &\preceq \bar{\Sigma} - \hat{\Sigma} \preceq \xi' \hat{\Sigma}, \\ \|\bar{m} - m_{\hat{\pi}}\|^2 &\leq (\xi')^2 \|\hat{\Sigma}\|_{\text{op}} \leq \xi^2, \end{aligned}$$

except with probability $O(1/n)$. The same conclusion holds from ν with at most twice this failure probability: for every path event A ,

$$\mathbb{P}_{\nu}(A) = \int \mathbb{P}_z(A) \nu(dz) \leq 2 \int \mathbb{P}_z(A) \hat{\pi}(dz) = 2\mathbb{P}_{\hat{\pi}}(A).$$

Taking the x -blocks of $\bar{m}$ and $\bar{\Sigma}$ gives $\hat{m}$ and $\hat{C}$ with the claimed bounds, since $\xi' \leq \xi$. Outside the two abort events, these are exactly the implemented outputs. A union bound gives total failure probability at most $2\eta + c_{\text{est}}/n$. Adding the warm-start and chain costs gives the stated query bound.

2.5 Additional conventions

The algorithmic parts of the paper distinguish centered covariance from the uncentered second moment, and the technical upper bounds are stated through an evaluation oracle for a convex potential before being specialized back to membership oracles. We record the relevant notation here so that the main body can remain focused on the main membership-oracle convex-body second moment estimation.

We will also use the following layer-cake consequence of Tonelli's theorem. The next lemma concerns the layer-cake identity.

Lemma 2.16. *Let Y, Y' be random variables on a common probability space satisfying $Y, Y' \geq a$ almost surely, let $s, t > -a$, and let $k, \ell > 0$. Then, with the expectation and the probability taken over the joint law of (Y, Y') ,*

$$\mathbb{E} \left[(Y + s)^{-k} (Y' + t)^{-\ell} \right] = \int_a^\infty \int_a^\infty \mathbb{P}(Y \leq x, Y' \leq y) k(x + s)^{-k-1} \ell(y + t)^{-\ell-1} dx dy.$$

The corresponding one-variable identity is

$$(Y + s)^{-k} = \int_a^\infty \mathbb{1}_{\{Y \leq x\}} k(x + s)^{-k-1} dx.$$

Proof. The one-variable identity follows by integrating the derivative of $u \mapsto (u + s)^{-k}$ from Y to ∞ . Multiplying the two identities and applying Tonelli's theorem to the resulting nonnegative integrand gives the two-variable formula. $\square$

2.6 Matrices, moments, and affine changes

Vectors are columns, $\|\cdot\|$ is the Euclidean norm, and we define a ball of radius r centered at x as $B(x, r) = \{y : \|y - x\| \leq r\}$; we abbreviate $B_2^n = B(0, 1)$. For symmetric matrices, $A \preceq B$ means that $B - A$ is positive semidefinite. We write $A^{-\top} = (A^{-1})^\top$ for an invertible matrix A .

The next definition concerns the moment matrices and relative Loewner approximation.

Definition 2.17. Let π be a probability law on $\mathbb{R}^n$ with finite second moment, and let $X \sim \pi$. Its mean, uncentered second moment, and centered covariance are

$$m_\pi := \mathbb{E}_{X \sim \pi} [X], \quad M_\pi := \mathbb{E}_{X \sim \pi} [XX^\top], \quad C_\pi := \mathbb{E}_{X \sim \pi} [(X - m_\pi)(X - m_\pi)^\top].$$

Thus

$$M_\pi = C_\pi + m_\pi m_\pi^\top. \quad (2.3)$$

For positive definite A, B and $\kappa \geq 1$, write

$$A \approx_\kappa B \iff \kappa^{-1}B \preceq A \preceq \kappa B.$$

When an accuracy is given additively, we state the oriented inequalities explicitly; for example, $(1 - \varepsilon)B \preceq A \preceq (1 + \varepsilon)B$. For a convex body K , let $\pi_K = \text{Unif}(K)$ and abbreviate $m_K := m_{\pi_K}$, $C_K := C_{\pi_K}$, and $M_K := M_{\pi_K}$.

If a law is even, then its mean is zero and its uncentered second moment is its covariance. In particular, this holds for π_K when $K = -K$. This identity explains why the symmetric result may speak of covariance, whereas the asymmetric results must retain the phrase *uncentered second moment*.

The next definition concerns the affine pushforwards.

Definition 2.18. Let $\mathcal{T}(x) = B_{\text{aff}}(x - b)$, where B_{aff} is known and invertible, and write $\mathcal{T}_\# \pi$ for the law of $\mathcal{T}(X)$ when $X \sim \pi$. Then

$$m_{\mathcal{T}_\# \pi} = B_{\text{aff}}(m_\pi - b), \quad (2.4)$$

$$C_{\mathcal{T}_\# \pi} = B_{\text{aff}} C_\pi B_{\text{aff}}^\top, \quad (2.5)$$

$$M_{\mathcal{T}_\# \pi} = B_{\text{aff}} \left(C_\pi + (m_\pi - b)(m_\pi - b)^\top \right) B_{\text{aff}}^\top. \quad (2.6)$$

In particular, Loewner inequalities between covariance matrices are preserved by affine congruence. Raw second moments obey the same congruence under a linear map; under a translation one must use Equation (2.6), rather than treating a raw moment as a covariance.

2.7 Oracle models and adaptive randomized calls

This subsection separates the binary membership model used for convex bodies from the exact-real evaluation model used for general log-concave laws. It also records the conditional-success and deterministic query-limit conventions needed when the later algorithms choose their oracle calls adaptively.

The next definition concerns the membership and evaluation oracles.

Definition 2.19. A membership oracle for a convex body K returns the bit $\mathbb{1}_{\{x \in K\}}$ on input $x \in \mathbb{R}^n$. An exact-real evaluation oracle for an extended-valued convex potential $V : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ returns $V(x)$ on input x . The associated log-concave law has density

$$\pi(dx) = Z_V^{-1} e^{-V(x)} dx, \quad Z_V = \int_{\mathbb{R}^n} e^{-V(x)} dx.$$

Neither Z_V nor the irrelevant additive normalization of V is assumed as known. For a convex body K , its indicator potential

$$I_K(x) = \begin{cases} 0, & x \in K, \\ +\infty, & x \notin K \end{cases}$$

satisfies $e^{-I_K} = \mathbb{1}_K$; hence evaluation access to I_K is exactly membership access to K . Computation between oracle calls is not charged unless explicitly stated.

General log-concave densities require their potential values, while the uniform-body specialization $V = I_K$ uses only membership. This is a model distinction, not an impossibility result.

The next definition concerns the conditional guarantees and query limits.

Definition 2.20. Suppose a randomized routine is called on inputs chosen from earlier outputs. A failure guarantee p *under adaptive conditioning* means that, conditional on every realized pre-call history for which the routine's displayed input promises hold, fresh internal randomness makes the call succeed with probability at least $1 - p$. A deterministic query limit is a bound that holds for every valid input, every history, and every realization of the internal randomness, including executions that return a failure symbol.

This convention allows failure budgets to be assigned chronologically and summed by conditional union bounds. It also keeps correctness events separate from the all-outcomes query limits used in the algorithmic theorems.

The next definition concerns the grounded evaluation-oracle input.

Definition 2.21. Fix a universal constant $h_0 > 0$. A grounded input in dimension $n \geq 1$ consists of a proper lower-semicontinuous convex function $V : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$, exact-real evaluation access to V , numbers $r, R > 0$, and a known point $x_0 \in \mathbb{R}^n$ such that

$$0 < \int_{\mathbb{R}^n} e^{-V(x)} dx < \infty, \quad B(x_0, r) \subseteq \{x : V(x) \leq \inf V + h_0 n\},$$

and, for $X \sim \pi \propto e^{-V}$,

$$\mathbb{E}_{X \sim \pi} [\|X - x_0\|^2] \leq R^2.$$

Constants hidden by asymptotic notation may depend on h_0 , but not on n, r, R , or a caller-supplied failure probability.

Here and below, a sublevel set of V is a set of the form $\{x : V(x) \leq a\}$. The promise concerns one such set, not necessarily the entire support. It gives a known ball in a low sublevel set and a global second-moment bound; the algorithm need not know $\inf V$.

2.8 Finite-order change of measure

This subsection fixes the change-of-measure quantity used to pass between successive laws on the cooling path. Distinguishing the logarithmic divergence from its exponentiated form will keep the warmness assumptions and one-step estimates in the overview and upper-bound proof unambiguous.

The next definition concerns the Rényi divergence and finite-order warmness.

Definition 2.22. Let ν and π be probability laws and let $q > 1$. If $\nu \ll \pi$, define

$$D_q(\nu \parallel \pi) := \frac{1}{q-1} \log \int \left(\frac{d\nu}{d\pi} \right)^q d\pi, \quad \mathcal{R}_q(\nu \parallel \pi) := \exp(D_q(\nu \parallel \pi));$$

otherwise both quantities are $+\infty$. In particular,

$$\mathcal{R}_2(\nu \parallel \pi) = \int \left(\frac{d\nu}{d\pi} \right)^2 d\pi.$$

We use D_q for the logarithmic divergence and $\mathcal{R}_q$ for its exponentiated finite-order warmness.

The order-two orientation matters: in the cooling argument, the new target is compared to the law from which the rounding call inherits its state. Writing both arguments in [Definition 2.22](#) prevents this direction from being suppressed.

2.9 Standard log-concave facts

The following proposition records exactly the background consequences used later. Its constants are universal and may change from one occurrence to the next.

The next proposition concerns the log-concave toolkit.

Proposition 2.23. Let ν be a full-dimensional log-concave probability law on $\mathbb{R}^n$ with potential W , mean m , and covariance C , and let $Y \sim \nu$.

(i) The barycenter inequality gives

$$W(m) \leq \inf W + n.$$

(ii) If $m = 0$ and $\frac{1}{2}I \preceq C \preceq 2I$, then there is a universal $c_P > 0$ such that, for every $u \geq 1$,

$$\Pr\left(\|Y\| \geq c_P \cdot un^{1/2}\right) \leq \exp(-un^{1/2}/c_P).$$

(iii) There is a universal $C_4 > 0$ such that every affine functional ℓ satisfies

$$\mathbb{E}_{Y \sim \nu} [|\ell(Y)|^4] \leq C_4 (\mathbb{E}_{Y \sim \nu} [|\ell(Y)|^2])^2.$$

(iv) There are universal $\rho_L, H_L > 0$ such that, if $m = 0$ and $\frac{1}{2}I \preceq C \preceq 2I$, then

$$B(0, \rho_L) \subseteq \{W \leq \inf W + H_L n\}.$$

Under the same covariance bounds, the conclusion with radius $\rho_L/2$ still holds when the centering hypothesis is relaxed to $\|m\| \leq \rho_L/2$.

(v) Every affine functional ℓ obeys

$$(\mathbb{E}_{Y \sim \nu} [|\ell(Y)|^2])^{1/2} \leq 4\mathbb{E}_{Y \sim \nu} [|\ell(Y)|].$$

Moreover, if $m = 0$ and $s^{-1}I \preceq C \preceq sI$ for $s \geq 1$, then

$$B\left(0, \frac{1}{es^{1/2}}\right) \subseteq \text{supp } \nu.$$

(vi) If W is twice continuously differentiable and strictly convex on its support, then every locally Lipschitz g in the associated Sobolev space satisfies the Brascamp–Lieb variance inequality

$$\text{Var}_{Y \sim \nu} [g] \leq \mathbb{E}_{Y \sim \nu} \left[\nabla g(Y)^\top (\nabla^2 W(Y))^{-1} \nabla g(Y) \right].$$

Consequently, $\nabla^2 W \succeq \alpha I$ implies $C \preceq \alpha^{-1}I$ and $\text{Var}_{Y \sim \nu} [g] \leq \alpha^{-1} \mathbb{E}_{Y \sim \nu} [\|\nabla g\|^2]$.

The barycenter estimate in part (i) is due to Fradelizi [Fra97]. Part (ii) is Paouris’ norm concentration [Pao06]. The one-dimensional moment and level-set estimates of Lovász and Vempala [LV07] imply the precise reverse-Hölder, support, and level-set consequences in parts (iii)–(v); their short reductions are retained with the full-density-rounder proof in Section 4.1. Part (vi) is the classical Brascamp–Lieb inequality [BL76]. The later cooling proof establishes the nonsmooth extended-valued form it needs, rather than silently applying part (vi) beyond its hypotheses.

For orientation, parts (iv)–(v) are used to pass from a bounded restriction to the full law in Section 4.1; parts (i), (ii), and (vi) enter the cooling estimates in Section 4.2; and the fourth-moment and reverse-Hölder bounds in parts (iii) and (v) are used in the moment-transfer and endpoint arguments in Sections 4.2 and 4.3.

The next fact is Theorem 4.1 of [KLS95], the inner ball of an isotropic convex body, in the scaled form used for the volume corollary.

Fact 2.24. *Let $K \subseteq \mathbb{R}^n$ be a convex body with centroid m and covariance C . If $C \succeq \lambda I$ for some $\lambda > 0$, then $B(m, \lambda^{1/2}) \subseteq K$.*

Proof. The body $C^{-1/2}(K - m)$ has centroid 0 and covariance I , so Theorem 4.1 of [KLS95] gives $B(0, ((n+2)/n)^{1/2}) \subseteq C^{-1/2}(K - m)$. Hence $K - m \supseteq C^{1/2}B(0, 1) \supseteq B(0, \lambda^{1/2})$, because $\|C^{-1/2}u\| \leq \lambda^{-1/2}\|u\|$ for every u . $\square$

The next fact is Theorem 1.1 of [CV18], in the form stated there for a general second moment.

Fact 2.25. *There is an algorithm that, for any $\varepsilon, p \in (0, 1)$ and any convex body $K \subseteq \mathbb{R}^n$ containing the unit ball and satisfying $\mathbb{E}_{x \sim K} [\|x\|^2] \leq R^2$, with probability at least $1-p$ outputs a number within a factor $1 \pm \varepsilon$ of $\text{Vol}(K)$ after*

$$O\left(\frac{\max\{R^2 n^2, n^3\}}{\varepsilon^2} \log^2 n \log^2 \frac{1}{\varepsilon} \log^2 \frac{n}{\varepsilon} \log \frac{1}{p}\right)$$

membership queries.

3 Overview

We begin with the general-body scalar Gaussian-cooling algorithm. Previous general rounding algorithms followed an ellipsoid-growth framework, and the best known membership-query bound had remained at $\tilde{O}(n^{3.5})$. Our scalar cooling path gives a more direct analysis and improves the query bound to $\tilde{O}(n^{13/4})$. After describing this general-body algorithm, we explain the symmetric matrix-cooling algorithm, which uses symmetry to reach cubic complexity. We then explain how the counterexamples to natural conjectural shortcuts fit with this algorithmic picture.

3.1 General-Body Scalar Cooling Algorithm

The scalar cooling path has an inherent limitation: the single temperature t relaxes the Gaussian factor by the same amount in every direction. If only a few directions are close to becoming poorly conditioned, the scalar update must slow down all directions anyway. The general-body algorithm still gets a speedup by combining this scalar cooling path with a stronger near-cubic rounding primitive for already well-rounded log-concave distributions. This lets the algorithm decrease the temperature by a factor of $1 - \Theta(n^{-1/4})$ while preserving enough finite-order warmth between consecutive targets.

For an asymmetric body, the target matrix is naturally the uncentered second moment

$$M_K = \mathbb{E}_{x \sim K} [xx^\top] = C_K + m_K m_K^\top.$$

At each temperature the algorithm estimates both the centered covariance and the mean, then uses the covariance estimate to define the next affine coordinates. The key comparison between consecutive temperatures is a finite-order change-of-measure estimate. If $s = (1 - \delta)t$ and $\mu_u(x) \propto e^{-u\|x\|_2^2} 1_K(x)$, then

$$D_2(\mu_s \| \mu_t) = \Lambda(t) + \Lambda(2s - t) - 2\Lambda(s), \quad \Lambda(u) = \log \int_K e^{-u\|x\|_2^2} dx.$$

Using the variance bound $\Lambda''(u) = \text{Var}_{X \sim \mu_u} [\|X\|_2^2] \leq O(n/u^2)$ gives $D_2(\mu_s \| \mu_t) = O(n\delta^2)$. Choosing $\delta = \Theta(n^{-1/4})$ makes this divergence $O(n^{1/2})$. The full-density primitive can tolerate this finite-order warmth because, after truncating to the Paouris event for the old rounded law, the complement contributes negligibly and the second moment of the new target remains $O(n)$ in the old coordinates. This is the source of the $n^{1/4}$ factor in the $O(n^{1/4} \log(nR/r))$ cooling-stage bound.

Thus the algorithm makes $O(n^{1/4} \log(nR/r))$ calls to the near-cubic full-density rounding primitive. Each call returns an approximate mean and covariance in the current coordinates; the next coordinates are obtained by subtracting the empirical mean and applying the inverse square root of the empirical covariance. At the endpoint the Gaussian tilt is weak enough that the cooled law is comparable to the uniform law on K , and one final call estimates the desired covariance and mean. Combining them as $\widehat{M} = \widehat{C} + \widehat{m}\widehat{m}^\top$ gives [Theorem 1](#).

Compared with the previous $\widetilde{O}(n^{3.5})$ rounding framework, this is both faster and structurally simpler: the proof follows an explicit Gaussian-cooling path rather than tracking intersections with a growing ellipsoid. The symmetric cubic algorithm in the next section goes further by replacing tI with a matrix precision, so the algorithm can continue cooling the well-conditioned directions while separately regularizing the troublesome ones.

For symmetric bodies this scalar algorithm is still valid, with the simplification that all means are zero. The next section describes a genuinely symmetric improvement: because multiplying the density by any even quadratic weight preserves its center, the algorithm can choose an anisotropic matrix precision and remove the extra $n^{1/4}$ factor.

3.2 Cubic Symmetric Algorithm

For a symmetric body, the uncentered second moment and covariance coincide throughout the cooling path. More importantly, symmetry lets us replace the scalar temperature by a matrix-valued quadratic penalty without having to track a moving mean. For a positive semidefinite matrix H , define

$$\mu_H(x) \propto e^{-x^\top H x} 1_K(x), \quad C_H = \mathbb{E}_{x \sim \mu_H} [xx^\top].$$

The matrix H specifies the strength of the quadratic penalty in each direction; it is also called the precision matrix of the Gaussian factor. We start from $H_0 = r^{-2}I$ and drive H toward 0. At every successful iteration the algorithm maintains an estimate $W \approx C_H$. In the normalized coordinates $z = W^{-1/2}x$, the current covariance is

$$W^{-1/2}C_HW^{-1/2} \approx I.$$

The update is the main new feature. Set

$$H^+ = (1 - \delta)H + \eta_{\text{ridge}} \frac{\text{Tr}(WH)}{n} W^{-1}.$$

If $B = W^{1/2}HW^{1/2}$, then the quadratic-penalty matrix in the normalized coordinates changes to

$$(1 - \delta)B + \eta_{\text{ridge}} \frac{\text{Tr}(B)}{n} I.$$

Thus every eigenvalue is decreased by the scalar cooling step and receives the same additive ridge. The ridge has the largest relative effect on directions in which the current quadratic penalty is smallest. The interpolation estimates in the proof show that the next transformed law still has $O(n)$ second moment and a fixed-radius ball in an $O(n)$ -height sublevel set. These are the hypotheses of the near-cubic covariance estimator used at each step.

The analysis tracks

$$D_{\text{sym}}(H) = \log \frac{\text{Vol}(K)}{\int_K e^{-x^\top H x} dx}, \quad q_{\text{sym}}(H) = \text{Tr}(C_H H).$$

If X is uniform on K , then

$$e^{-D_{\text{sym}}(H)} = \mathbb{E}_{X \sim K} \left[e^{-X^\top H X} \right].$$

Thus D_{sym} measures the reduction in the normalizing integral caused by the quadratic weight, while $q_{\text{sym}}(H) = \mathbb{E}_{X \sim \mu_H} [X^\top H X]$ is the expected quadratic penalty under μ_H . With $L = \Theta(1 + \log(2 + nR/r))$, $\delta = L^{-2}$, and $\eta_{\text{ridge}} = L^{-3}$, the proof gives

$$D_{\text{sym}}(H^+) \leq D_{\text{sym}}(H) - \frac{\delta}{2} q_{\text{sym}}(H), \quad q_{\text{sym}}(H) \leq D_{\text{sym}}(H) \leq O(L) q_{\text{sym}}(H).$$

Hence D_{sym} contracts by a factor $1 - \Theta(L^{-3})$ at each successful update, and $O(L^4)$ updates suffice. Once D_{sym} is a small constant, a terminal comparison shows that the uniform measure on K is well-rounded in the current coordinates, so one final near-cubic covariance call estimates C_K .

This gives the cubic symmetric algorithm of [Theorem 2](#). Sampling-based membership-oracle volume algorithms already spend $\tilde{O}(n^2)$ membership queries per well-rounded sample and use $O(n)$ samples for stable cooling ratios. Thus we do not expect a sampling-based approach to beat cubic query complexity except by changing the model or avoiding the bottleneck invoked by sampling procedures.

3.3 Counterexamples to Previously Stated Conjectures

We also include several refutations of previously conjectured directions for further improving these algorithms. Specifically, we refute the possibility of constant-size scalar Gaussian-cooling steps from the algorithm in this paper, and the possibility of large ellipsoid-growth steps from an already rounded intersection referenced in [\[JLLV26\]](#).

For the scalar-cooling obstruction, an astute reader may observe that if

$$\left\| A_t^{-1/2} \frac{dA_t}{dt} A_t^{-1/2} \right\|_{\text{op}} \leq \tilde{O}(1/t),$$

then one could simply decrease $t \rightarrow (1 - 1/\tilde{O}(1))t$ per iteration, and obtain a cubic scalar-cooling algorithm even for asymmetric bodies. For a compact convex body $K \subset \mathbb{R}^n$ and positive definite Q , write

$$\mu_{K,Q}(\mathrm{d}x) = \frac{1}{Z_{K,Q}} \mathbf{1}_K(x) \exp\left(-\frac{1}{2}x^\top Q^{-1}x\right) \mathrm{d}x,$$

and say that $\mu_{K,Q}$ is isotropic if it has mean zero and covariance I_n . For $0 \leq \delta \leq 1$, define

$$R_n(\delta) = \sup \left\{ \text{Tr Cov}(\mu_{K,(1+\delta)Q}) : K \subset \mathbb{R}^n \text{ a convex body, } Q \succ 0, \mu_{K,Q} \text{ isotropic} \right\}.$$

The next theorem concerns the sharp trace growth under anisotropic Gaussian warming.

Theorem 7. *There is a universal constant $c > 0$ such that, for every $n \geq 1$ and $0 \leq \delta \leq 1$,*

$$R_n(\delta) \geq c(n + n^2\delta^4).$$

For this lower bound, K may be taken to be a convex polytope.

The proof is given in [Section A](#). Together with [Theorem 9](#), this rules out the shortcut above. In particular, [Theorem 7](#) constructs restricted Gaussians for which the trace after a small anisotropic warming is $\Omega(n + n^2\delta^4)$. Thus a constant warming step can make the next target have trace $\Omega(n^2)$ in the old coordinates, far outside the well-rounded regime needed by the near-cubic primitive. The scale $\delta = \Theta(n^{-1/4})$ is exactly where this obstruction stops being superlinear in n . Similarly, [Theorem 9](#) gives a symmetric log-concave counterexample to the coordinate-moment simplification behind a constant-size derivative bound, so this shortcut fails even before one has to handle moving means.

For the ellipsoid-growth obstruction, the framework of [\[JLLV26\]](#) suggests another possible simplification: if $K \cap E$ is already approximately isotropic, then modestly enlarging E might keep $K \cap (1 + \delta)E$ well-rounded. Such a statement would allow the older ball/ellipsoid-growth method to take much larger growth steps. The following theorem shows that this route fails as well.

The next theorem concerns the counterexample to Conjecture 3.5 of [\[JLLV26\]](#).

Theorem 8. *There are absolute constants $c, C > 0$ such that, for every sufficiently large n , one can construct a convex body $K \subset \mathbb{R}^n$ and an ellipsoid $E \subset \mathbb{R}^n$ such that $K \cap E$ has mean zero and covariance I_n , but for every $0 < \delta \leq 1$ satisfying $(n - 1)\delta \geq C$,*

$$\text{Tr}(\text{Cov}(K \cap (1 + \delta)E)) \geq c(n - 1)^2\delta^2.$$

[Theorem 8](#) is proved in [Section C](#); the quantitative form used there is [Theorem 12](#). This limitation still leaves open the possibility of a future asymmetric volume algorithm that avoids full second-moment estimation.

We now state the two conjectures refuted in [Sections B](#) and [C](#). Recall the notation from [Section 2](#): a convex body K has uniform law π_K , and its mean, covariance, and uncentered second moment are m_K , C_K , and M_K respectively. When $K = -K$, we have $m_K = 0$ and $C_K = M_K$.

We first give a counterexample to a coordinate-moment conjecture that would make the scalar Gaussian-cooling derivative uniformly bounded. If true, the resulting bound would permit constant-size temperature decreases in the scalar cooling method of [Section 4](#), removing its $n^{1/4}$ stage overhead and yielding a cubic asymmetric rounding algorithm up to logarithmic factors.

The next conjecture concerns the coordinate-moment conjecture.

Conjecture 3.1. Let $f : \mathbb{R}^n \rightarrow [0, +\infty)$ be an isotropic log-concave function. Then

$$\left| \mathbb{E}_{x \sim f} \left[x_1^2 (\|x\|^2 - n) \right] \right| = O(1),$$

where $x_1 = \langle x, e_1 \rangle$ and the implicit constant is absolute.

The second obstruction is the ellipsoid-growth statement that would let the previous method of [JLLV26] take substantially larger growth steps. If true, it would allow that method to replace its $1 + \Theta(n^{-1/2})$ updates by constant-factor updates, reducing the number of stages from $\tilde{O}(n^{1/2} \log(R/r))$ to $\tilde{O}(\log(R/r))$ and giving a cubic asymmetric rounding algorithm.

The next conjecture concerns the ellipsoid growth, Conjecture 3.5 from [JLLV26].

Conjecture 3.2. Suppose that for a convex body K and an ellipsoid E , the set $K \cap E$ is 2-isotropic. Then $K \cap 2E$ is well-rounded.

For this body-level conjecture, we say that a convex body L is 2-isotropic if

$$\frac{1}{2}I_n \preceq M_L \preceq 2I_n,$$

and well-rounded if $\text{Tr}(C_L) = \tilde{O}(n)$ (the $\tilde{O}$ hides polylogarithmic factors) and $B(0, 1) \subseteq L$. This is the body-level specialization of the well-roundedness terminology used in the main text.

The next theorem concerns the failure of the coordinate-moment conjecture.

Theorem 9. *There is a universal constant $c > 0$ such that, for every sufficiently large n , there is an even isotropic log-concave density f on $\mathbb{R}^n$ satisfying*

$$\left| \mathbb{E}_{x \sim f} \left[x_1^2 (\|x\|^2 - n) \right] \right| \geq cn^{1/2}.$$

The conjecture counterexamples indicate that the scalar Gaussian-cooling route is unlikely to be pushed to constant-size steps using the moment simplifications discussed in Section 4.5, and the previous ellipsoid-growth route cannot rely on Conjecture 3.5 of [JLLV26]. We therefore view cubic complexity as the natural limit for these second-moment methods in the membership-oracle model, while a different idea might still achieve cubic asymmetric volume estimation.

4 General-Body Algorithm: Scalar Cooling

In this section we prove the $\tilde{O}(n^{13/4})$ rounding algorithm for a general convex body, specifically Theorem 1.

The proof replaces the ellipsoid-growth chain of previous rounding algorithms by an explicit Gaussian-cooling path. Rather than intersecting the body with a growing ball or ellipsoid and tracking how the covariance changes under that geometric operation, we keep the body fixed and change the density: with

$$\Phi_t(x) = V(x) + t\|x - x_0\|_2^2, \quad \pi_t \propto e^{-\Phi_t},$$

we cool the temperature t from r^{-2} down to $\Theta(R^{-2})$ by multiplicative steps $t \leftarrow (1 - \delta)t$ with $\delta = \Theta(n^{-1/4})$. At each temperature a near-cubic full-density rounding primitive, FULLDENSITY-ROUND, estimates the mean and covariance of the current cooled law; the estimate defines the affine coordinates in which the next, slightly warmer, law is again well-rounded. The step size is dictated by a finite-order change-of-measure estimate: the order-two Rényi divergence between consecutive

cooled laws is $O(n\delta^2) = O(n^{1/2})$, exactly the scale at which the Paouris tail bound transfers the $O(n)$ second-moment guarantee from one temperature to the next. After $O(n^{1/4} \log(2 + R/r))$ stages the Gaussian tilt is negligible, one final call estimates the mean and covariance of the target law, and the uncentered second moment is reconstructed as $\widehat{M} = \widehat{C} + \widehat{m}\widehat{m}^\top$.

The section is organized as follows. [Section 4.1](#) constructs the full-density rounding interface ([Lemma 4.5](#)) from the near-cubic well-rounded rounding black box recorded in [Theorem 5](#) and the moment-estimation black box of [Theorem 6](#); the short reductions of [Proposition 2.23](#) parts (iii)–(v) promised in [Section 2](#) are also given there. [Section 4.2](#) proves the cooling estimates: the nonsmooth extended Brascamp–Lieb inequality ([Lemma 4.6](#)), the quadratic energy and variance bound ([Lemma 4.7](#)), the one-step Rényi cost, and the well-roundedness induction at the next temperature. [Section 4.3](#) compares the cooled moments to the target moments and reconstructs the raw second moment. [Section 4.4](#) gives the algorithm ([Figure 2](#)), the well-roundedness induction, and the query accounting, and proves [Theorem 1](#). [Section 4.5](#) records why the scalar route cannot be pushed to constant-size steps.

4.1 The Full-Density Rounding Interface

The two black boxes. This subsection rests on the two imported statements of [Section 2](#): the isotropic rounding of a grounded law ([Theorem 5](#), [[KV25](#), Lemma 4.11]) and the moment estimation of a well-rounded law with bounded covariance ([Theorem 6](#), assembled from [[KV25](#), Theorem 3.16, Lemma 2.14, Lemma 2.9, and Lemma 4.6]). Both are stated with the exact constants of [[KV25](#), Assumption 4.4]: unit inradius, level-set height $10n$, and (for the rounding) support equal to that level set. An input of FULLDENSITYROUND is only (ρ, H, B) -well-rounded ([Definition 2.15](#)) for universal ρ, H, B , so the proof of [Lemma 4.5](#) first brings it to these exact constants by an explicit reduction ([Lemma 4.2](#)): a rescaling, a lift by independent Gaussian coordinates to a dimension $N = O(n)$ in which the level-set height Hn is below N , and a restriction to a level set of mass at least $1 - 10^{-5}$. Constants hidden by asymptotic notation below may depend on (ρ, H, B) , but not on n or on a caller-supplied failure probability. Throughout, $n_0 \geq 4$ is a universal integer fixed in the proof of [Lemma 4.5](#).

Reductions of the toolkit parts (iii)–(v). We discharge the promise made in [Section 2](#) that the one-dimensional moment and level-set estimates of [[LV07](#)] imply parts (iii)–(v) of [Proposition 2.23](#).

Proof of [Proposition 2.23](#) parts (iii)–(v). Part (iii). Let ℓ be an affine functional and write $L = |\ell(Y)|$. If ℓ is constant, the moment claims are immediate; otherwise its law is a one-dimensional log-concave density. Applying [Lemma 2.12](#) with $k = 4$ and with $k = 2$ to the one-dimensional log-concave law of $\ell(Y)$ ([Lemma 2.10](#)) gives

$$\mathbb{E}[L^4] \leq 8^4 (\mathbb{E}[L])^4 \leq 8^4 (\mathbb{E}[L^2])^2,$$

which is part (iii) with $C_4 = 8^4$.

Part (v), first claim. The same bound with $k = 2$ reads $(\mathbb{E}[L^2])^{1/2} \leq 4\mathbb{E}[L]$.

Part (v), second claim, and the support ball. For an isotropic log-concave law, each one-dimensional marginal has density at most 1, and each halfspace containing the mean has probability at least $1/e$ [[LV07](#), Lemmas 5.5(a) and 5.12]. Consequently, every supporting hyperplane of its support is at distance at least $1/e$ from its mean. Applying the inverse whitening map gives the radius $1/(es^{1/2})$ when $s^{-1}I \preceq C \preceq sI$; see also [[KV25](#), Lemma A.10].

Part (iv). The isotropic density lower bound of [[LV07](#), Theorem 5.14(a)] gives universal $\rho_0, H_0 > 0$ such that every isotropic log-concave density e^{-W} obeys $e^{-W(x)} \geq e^{-H_0 n} e^{-W(0)}$ on $B(0, 2\rho_0)$.

Suppose first that $m = 0$ and $\frac{1}{2}I \preceq C \preceq 2I$. Pushing to isotropic coordinates preserves the conclusion up to a factor 2, so $W(x) \leq W(0) + H_0 n$ on $B(0, 2^{1/2} \rho_0)$; the barycenter inequality (part (i)) gives $W(0) \leq \inf W + n$, and hence $B(0, \rho_L) \subseteq \{W \leq \inf W + H_L n\}$ with $\rho_L = 2^{1/2} \rho_0$ and $H_L = H_0 + 1$. If instead $\|m\| \leq \rho_L/2$, apply the centered statement to $\tilde{W}(x) = W(x + m)$, whose law is centered and still has covariance in $[\frac{1}{2}I, 2I]$: then $B(m, \rho_L) \subseteq \{W \leq \inf W + H_L n\}$, and $B(0, \rho_L/2) \subseteq B(m, \rho_L)$. $\square$

The rest of this subsection constructs the rounding interface. The first step is a level-set mass estimate, which will let us truncate a log-concave law to a sublevel set of its potential while barely changing it.

The next lemma concerns the level-set mass.

Lemma 4.1. *Let $W : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ be proper, lower-semicontinuous, and convex with $0 < \int e^{-W} < \infty$, let $a = \inf W$, and let $\mu \propto e^{-W}$ be the associated probability law. Then for every $\bar{H} \geq 2$ and every $s \geq 2$,*

$$\mathbb{P}_{X \sim \mu} [W(X) \geq a + (\bar{H} + s)n] \leq \frac{4}{\bar{H}} e^{-sn/2}.$$

Proof. Replacing W by $W(\cdot - x^*) - a$ for a minimizer x^* (which exists: integrability forces W to tend to $+\infty$ along every ray, so its sublevel sets are bounded), assume $a = 0$ and $W(0) = 0$. For $t \geq 0$ let $v(t) = \text{Vol}(\{W \leq t\})$. For $0 < s < t$ and $x \in \{W \leq t\}$, convexity gives $W(\frac{s}{t}x) \leq \frac{s}{t}W(x) + (1 - \frac{s}{t})W(0) \leq s$, that is, $\frac{s}{t}\{W \leq t\} \subseteq \{W \leq s\}$, and therefore

$$v(s) \geq \left(\frac{s}{t}\right)^n v(t) \quad (0 < s < t).$$

Layer cake gives $Z = \int e^{-W} dx = \int_0^\infty e^{-t} v(t) dt$, hence with $T_0 = \bar{H}n$,

$$Z \geq v(T_0) T_0^{-n} \int_0^{T_0} e^{-t} t^n dt \geq v(T_0) T_0^{-n} e^{-T_0} \frac{T_0^{n+1}}{n+1},$$

while for $T > T_0$, using $(1 + u/T)^n \leq e^{nu/T}$,

$$\int_T^\infty e^{-t} v(t) dt \leq v(T_0) T_0^{-n} e^{-T} T^n \int_0^\infty e^{-u+nu/T} du = v(T_0) T_0^{-n} e^{-T} T^n \left(1 - \frac{n}{T}\right)^{-1}.$$

Taking $T = (\bar{H} + s)n$ with $s \geq 2$ and dividing, and using $(1 + s/\bar{H})^n \leq e^{sn/\bar{H}}$, $\bar{H} \geq 2$, and $(1 - n/T)^{-1} \leq 2$,

$$\mu\{W \geq T\} \leq \frac{(n+1) e^{\bar{H}n}}{\bar{H}n(1 - n/T)} \left(1 + \frac{s}{\bar{H}}\right)^n e^{-(\bar{H}+s)n} \leq \frac{2(n+1)}{\bar{H}n} e^{-sn(1-1/\bar{H})} \leq \frac{4}{\bar{H}} e^{-sn/2},$$

since $s \geq 2$. $\square$

The next lemma concerns the lifted grounded restriction.

Lemma 4.2. *Let $\mu \propto e^{-U}$ on $\mathbb{R}^n$ be (ρ, H, B) -well-rounded around z , let $k = \lceil Hn \rceil + n_0$ and $N = n + k$, and define on $\mathbb{R}^N = \mathbb{R}^n \times \mathbb{R}^k$*

$$\bar{U}(x, y) = U(z + \rho x) + \frac{1}{2}\|y\|^2, \quad \bar{\mu} \propto e^{-\bar{U}},$$

and

$$K = \{(x, y) : \bar{U}(x, y) \leq \bar{U}(0) + 9N\}, \quad \nu = \bar{\mu}(\cdot \mid K).$$

Then $\bar{\mu} = \mu_\rho \otimes \gamma_k$, where μ_ρ is the law of $\rho^{-1}(Y - z)$ for $Y \sim \mu$ and γ_k is the standard Gaussian law on $\mathbb{R}^k$, and the following hold with $\bar{B} = B/\rho^2 + 1$.

1. $B(0, 1) \subseteq \{\bar{U} \leq \inf \bar{U} + N\} \subseteq K$ and $\mathbb{E}_{X \sim \bar{\mu}} [\|X\|^2] \leq \bar{B}N$.
2. $\bar{\mu}(K) \geq q_0 := 1 - 10^{-5}$ and $d\nu/d\bar{\mu} \leq q_0^{-1}$ on K .
3. $\inf_K \bar{U} = \inf \bar{U}$ and $\sup_K \bar{U} - \inf \bar{U} \leq 10N$.
4. With $U' = \bar{U} + I_K - \bar{U}(0) + N$, the law $\nu \propto e^{-U'}$ is grounded around 0 with respect to U' ([Definition 2.15](#)), and $\mathbb{E}_{X \sim \nu} [\|X\|^2] \leq 2\bar{B}N$.
5. One evaluation of U' costs one evaluation of U after $U(z)$ has been computed once.

Proof. The product form is the change of variables $x = \rho^{-1}(y_1 - z)$ on the first factor, the Gaussian factor being $e^{-\|y\|^2/2}$. Write $\bar{a} = \inf \bar{U} = \inf U$, the infimum of the Gaussian part being 0. For $\|(x, y)\| \leq 1$ we have $z + \rho x \in B(z, \rho)$ and $\frac{1}{2}\|y\|^2 \leq \frac{1}{2}$, so $\bar{U}(x, y) \leq \inf U + Hn + \frac{1}{2} \leq \bar{a} + N$, because $N \geq Hn + n_0 \geq Hn + 1$; this is the first inclusion of Item [4.2\(1\)](#). Since $\bar{a} \leq \bar{U}(0)$, the set $\{\bar{U} \leq \bar{a} + N\}$ lies in $\{\bar{U} \leq \bar{U}(0) + 9N\} = K$. The second moment is $\mathbb{E}_{\bar{\mu}} [\|X\|^2] = \mathbb{E}_{\mu} [\|\rho^{-1}(Y - z)\|^2] + k \leq Bn/\rho^2 + k \leq \bar{B}N$.

For Item [4.2\(2\)](#), $K \supseteq \{\bar{U} \leq \bar{a} + 9N\}$, so [Lemma 4.1](#) in dimension N with $\bar{H} = 2$ and $s = 7$ gives

$$\bar{\mu}(K^c) \leq \mathbb{P}_{X \sim \bar{\mu}} [\bar{U}(X) \geq \bar{a} + 9N] \leq 2e^{-7N/2} \leq 2e^{-14} \leq 10^{-5},$$

using $N \geq n_0 \geq 4$; and $d\nu/d\bar{\mu} = \mathbb{1}_K/\bar{\mu}(K) \leq q_0^{-1}$.

For Item [4.2\(3\)](#), K contains $\{\bar{U} \leq \bar{a} + 9N\}$, which contains points of $\bar{U}$ -value arbitrarily close to $\bar{a}$, so $\inf_K \bar{U} = \bar{a}$; and $\sup_K \bar{U} \leq \bar{U}(0) + 9N \leq \bar{a} + 10N$ because $0 \in B(0, 1) \subseteq \{\bar{U} \leq \bar{a} + N\}$.

For Item [4.2\(4\)](#), $\nu \propto e^{-U'}$ since U' differs from $\bar{U} + I_K$ by a constant. We have $\min U' = \inf_K \bar{U} - \bar{U}(0) + N = \bar{a} - \bar{U}(0) + N \geq 0$ by Item [4.2\(3\)](#), and on K , $U' \leq \bar{U}(0) + 9N - \bar{U}(0) + N = 10N$. Hence $\text{supp } \nu = K \subseteq \{U' \leq \min U' + 10N\}$, while $\{U' \leq \min U' + 10N\} \subseteq \{U' < \infty\} = K$; so the support of ν equals the level set $\{U' \leq \inf U' + 10N\}$, which contains $B(0, 1)$ by Item [4.2\(1\)](#). Finally $\mathbb{E}_{\nu} [\|X\|^2] = \bar{\mu}(K)^{-1} \mathbb{E}_{\bar{\mu}} [\|X\|^2 \mathbb{1}_K] \leq \bar{B}N/q_0 \leq 2\bar{B}N$.

For Item [4.2\(5\)](#), membership in K is decided by comparing $\bar{U}(x, y)$ with $\bar{U}(0) + 9N = U(z) + 9N$, and $\bar{U}(x, y)$ is one evaluation of U at $z + \rho x$ plus explicit arithmetic. $\square$

The next lemma transfers rounding information from a large-mass restriction back to the full law; it is the content of Steps 2 and 3 of [\[KV25, Section 4.3\]](#), recorded here in the form we use.

The next lemma concerns the restricted-to-full transfer.

Lemma 4.3. *Let λ be a log-concave probability law on $\mathbb{R}^n$, let $S \subseteq \mathbb{R}^n$ be measurable with $\lambda(S) \geq 1 - \varepsilon$ for some $\varepsilon \in (0, 10^{-4}]$, and let $\nu = \lambda(\cdot \mid S)$. Then for every unit vector v and every center $c \in \mathbb{R}^n$,*

$$\mathbb{E}_{\lambda} [|v^\top (X - c)|] \leq 1.05 \mathbb{E}_{\nu} [|v^\top (X - c)|],$$

and, writing the variances as $\sigma_{\lambda}^2 = \text{Var}_{\lambda} [v^\top X]$ and $\sigma_{\nu}^2 = \text{Var}_{\nu} [v^\top X]$,

$$(1 - \varepsilon) \sigma_{\nu}^2 \leq \sigma_{\lambda}^2 \leq 18.1 \sigma_{\nu}^2, \quad (v^\top (m_{\lambda} - m_{\nu}))^2 \leq \frac{\varepsilon}{1 - 2\varepsilon} \sigma_{\lambda}^2.$$

In particular, if an affine map F makes $F_{\#}\nu$ satisfy $C_{F_{\#}\nu} \approx_{1.01} I$ and $\|m_{F_{\#}\nu}\| \leq 0.01$, then $0.99 I \preceq C_{F_{\#}\lambda} \preceq 18.5 I$ and $\|m_{F_{\#}\lambda}\| \leq 0.06$.

Proof. Set $f(x) = |v^\top (x - c)|$ and note the χ^2 -warmness of a conditional law: since $d\nu = \mathbb{1}_S d\lambda/\lambda(S)$,

$$\chi^2(\nu \parallel \lambda) = \mathbb{E}_{\lambda} \left[\left(\frac{d\nu}{d\lambda} - 1 \right)^2 \right] = \frac{1}{\lambda(S)} - 1 \leq \frac{\varepsilon}{1 - \varepsilon}.$$

Cauchy–Schwarz and [Proposition 2.23\(v\)](#) applied to the affine functional $v^\top(x - c)$ under the log-concave law λ give

$$|\mathbb{E}_\lambda[f] - \mathbb{E}_\nu[f]| \leq \mathbb{E}_\lambda[f^2]^{1/2} (\chi^2(\nu\|\lambda))^{1/2} \leq 4 \mathbb{E}_\lambda[f] (\varepsilon/(1 - \varepsilon))^{1/2},$$

so, since $4(\varepsilon/(1 - \varepsilon))^{1/2} \leq 0.05$ for $\varepsilon \leq 10^{-4}$,

$$\mathbb{E}_\lambda[f] \leq (1 - 4(\varepsilon/(1 - \varepsilon))^{1/2})^{-1} \mathbb{E}_\nu[f] \leq 1.05 \mathbb{E}_\nu[f].$$

Squaring after another application of [Proposition 2.23\(v\)](#),

$$\mathbb{E}_\lambda[f^2] \leq 16(\mathbb{E}_\lambda[f])^2 \leq 16(1.05)^2 (\mathbb{E}_\nu[f])^2 \leq 18 \mathbb{E}_\nu[f^2],$$

since $16(1.05)^2 = 17.64 \leq 18$. For the centered bounds, let $g(x) = v^\top(x - m_\nu)$ and $\Delta = (v^\top(m_\lambda - m_\nu))^2$. The same Cauchy–Schwarz step with g in place of f gives $\Delta = |\mathbb{E}_\lambda[g] - \mathbb{E}_\nu[g]|^2 \leq \mathbb{E}_\lambda[g^2] \chi^2(\nu\|\lambda) \leq (\sigma_\lambda^2 + \Delta) \varepsilon/(1 - \varepsilon)$, hence $\Delta \leq \frac{\varepsilon}{1-2\varepsilon} \sigma_\lambda^2$. The lower bound follows from $\mathbb{E}_\lambda[(X - m_\lambda)(X - m_\lambda)^\top] \succeq \lambda(S) \mathbb{E}_\nu[(X - m_\lambda)(X - m_\lambda)^\top] \succeq (1 - \varepsilon)C_\nu$. For the upper bound, take $c = m_\lambda$ above: $\sigma_\lambda^2 = \mathbb{E}_\lambda[f^2] \leq 18 \mathbb{E}_\nu[f^2]$ with $\mathbb{E}_\nu[f^2] = \sigma_\nu^2 + \Delta \leq \sigma_\nu^2 + \frac{\varepsilon}{1-2\varepsilon} \sigma_\lambda^2$; rearranging, $\sigma_\lambda^2 \leq \frac{18}{1-18\varepsilon/(1-2\varepsilon)} \sigma_\nu^2 \leq 18.1 \sigma_\nu^2$, since $18\varepsilon/(1 - 2\varepsilon) \leq 1.9 \cdot 10^{-3}$.

For the final claim, in F -coordinates $\sigma_\nu^2 \in [1/1.01, 1.01]$ for every unit v . Hence

$$\sigma_\lambda^2 \in [0.9999/1.01, 18.1 \cdot 1.01] \subseteq [0.99, 18.5].$$

Moreover,

$$\|m_{F\#\lambda} - m_{F\#\nu}\|^2 = \sup_v (v^\top(m_{F\#\lambda} - m_{F\#\nu}))^2 \leq \frac{\varepsilon}{1 - 2\varepsilon} \sup_v \sigma_\lambda^2 \leq \frac{10^{-4}}{1 - 2 \cdot 10^{-4}} \cdot 18.5 \leq 0.002,$$

so $\|m_{F\#\lambda}\| \leq \|m_{F\#\nu}\| + 0.002^{1/2} \leq 0.01 + 0.05 \leq 0.06$. $\square$

The next lemma concerns the robust selection.

Lemma 4.4. *Fix $\xi \in (0, 1/50)$ and $p \in (0, 1/10)$, and let $L = \lceil 512 \log(2/p) \rceil$. Fix $m \in \mathbb{R}^n$ and $C \succ 0$. Let L independent random outputs each be either a failure symbol or a pair $(\hat{m}_i, \hat{C}_i)$ with $\hat{C}_i \succ 0$. Suppose each output is ξ -accurate for (m, C) with probability at least $17/32$, meaning $(1 - \xi)C \preceq \hat{C}_i \preceq (1 + \xi)C$ and $\|C^{-1/2}(\hat{m}_i - m)\| \leq \xi$. Failure symbols have no partners. For nonfailure outputs, declare $i \sim j$ when*

$$(1 - 3\xi)\hat{C}_i \preceq \hat{C}_j \preceq (1 + 3\xi)\hat{C}_i, \quad \|\hat{C}_i^{-1/2}(\hat{m}_i - \hat{m}_j)\| \leq 3\xi.$$

Select a nonfailure index i^* with the most partners, returning a failure symbol if none exists. With probability at least $1 - p/2$, the selected pair is 5ξ -accurate.

Proof. Accurate pairs are mutually compatible: for accurate i, j , $\hat{C}_i \preceq (1 + \xi)C \preceq \frac{1+\xi}{1-\xi}\hat{C}_j \preceq (1 + 3\xi)\hat{C}_j$ since $\xi \leq 1/50$, and symmetrically; and

$$\|\hat{C}_i^{-1/2}(\hat{m}_i - \hat{m}_j)\| \leq \|\hat{C}_i^{-1/2}C^{1/2}\| (\|C^{-1/2}(\hat{m}_i - m)\| + \|C^{-1/2}(m - \hat{m}_j)\|) \leq 2(1 - \xi)^{-1/2} \xi \leq 3\xi,$$

using $\|\hat{C}_i^{-1/2}C^{1/2}\|^2 = \|\hat{C}_i^{-1/2}C\hat{C}_i^{-1/2}\|_{\text{op}} \leq (1 - \xi)^{-1}$. By Hoeffding's inequality, the number of accurate indices exceeds $L/2$ with probability at least $1 - e^{-L/512} \geq 1 - p/2$. On that event, every

accurate index has more than $L/2$ partners, so i^* also has more than $L/2$ partners. Since fewer than $L/2$ indices are inaccurate, i^* has an accurate partner j . Compatibility then gives

$$\widehat{C}_{i^*} \preceq \frac{1}{1-3\xi} \widehat{C}_j \preceq \frac{1+\xi}{1-3\xi} C \preceq (1+5\xi)C, \quad \widehat{C}_{i^*} \succeq \frac{1}{1+3\xi} \widehat{C}_j \succeq \frac{1-\xi}{1+3\xi} C \succeq (1-5\xi)C,$$

and

$$\|C^{-1/2}(\widehat{m}_{i^*} - m)\| \leq \|C^{-1/2}\widehat{C}_{i^*}^{1/2}\| \|\widehat{C}_{i^*}^{-1/2}(\widehat{m}_{i^*} - \widehat{m}_j)\| + \|C^{-1/2}(\widehat{m}_j - m)\| \leq (1+5\xi)^{1/2} \cdot 3\xi + \xi \leq 5\xi. \quad \square$$

FULLDENSITYROUND($U, z, \rho, H, B, \varepsilon, p$)

Input: evaluation oracle for a proper lower-semicontinuous convex potential U ; a known point $z \in \mathbb{R}^n$; universal constants $\rho, H, B > 0$; a fixed accuracy $\varepsilon \in (0, 1/10)$; a failure tolerance $p \in (0, 1/10)$.

Output: a failure symbol $\perp$, or a pair $(\widehat{m}, \widehat{C})$ with $\widehat{C} \succ 0$.

Assumptions: $\mu \propto e^{-U}$ is a probability law with $B(z, \rho) \subseteq \{U \leq \inf U + Hn\}$ and $\mathbb{E}_{Y \sim \mu} [\|Y - z\|^2] \leq Bn$.

1. Set $L = \lceil 512 \log(2/p) \rceil$, $\xi = \varepsilon/5$, $k = \lceil Hn \rceil + n_0$, $N = n + k$, and compute $U(z)$. Implement the grounded potential U' on $\mathbb{R}^N$ of [Lemma 4.2](#); each evaluation costs one evaluation of U . Let $Q_{\text{iso}}(N) = A_C N^3 \log^{ac}(2N)$ be the expected-query bound of [Theorem 5](#) for $C^2 = 2\bar{B} = 2B/\rho^2 + 2$, and let $Q_{\text{est}}(N) = 168AN^3 \log^a(128N)$ be that of [Theorem 6](#) for $c_+ = 167$, accuracy ξ , $R^2 = 167N$, and $\eta = 1/64$.
2. For each $i \in [L]$, with fresh randomness:
 - (a) Run the rounding algorithm of [Theorem 5](#) on U' with the known point 0 and the known constant C , capped at $16Q_{\text{iso}}(N)$ evaluations. If it fails or times out, mark trial i failed and continue with the next trial. Otherwise it returns an invertible affine map $F_i(x) = A_i(x - \bar{x}_i)$ on $\mathbb{R}^N$.
 - (b) Run the estimator of [Theorem 6](#) on the potential $W_i(u) = \bar{U}(F_i^{-1}(u/3))$ on $\mathbb{R}^N$ with the known point 0, $R^2 = 167N$, accuracy ξ , and $\eta = 1/64$, capped at $16Q_{\text{est}}(N)$ evaluations. If it fails or times out, mark trial i failed and continue with the next trial. Otherwise it returns $(\widehat{m}'_i, \widehat{C}'_i)$; set $\widehat{m}_i = \widehat{m}'_i/3$ and $\widehat{C}_i = \widehat{C}'_i/9$, pull back to the lifted coordinates by $\widetilde{m}_i = F_i^{-1}\widehat{m}_i$ and $\widetilde{C}_i = A_i^{-1}\widehat{C}_i A_i^{-\top}$, and let $\widetilde{m}_i^x \in \mathbb{R}^n$ and $\widetilde{C}_i^x \in \mathbb{R}^{n \times n}$ be their blocks on the first n coordinates. If $\widetilde{C}_i^x$ is not positive definite, mark trial i failed.
3. Among the successful trials, select an index i^* whose pair $(\widetilde{m}_{i^*}^x, \widetilde{C}_{i^*}^x)$ has the most $\sim$ -partners in the sense of [Lemma 4.4](#); if no trial succeeded, return $\perp$.
4. Return $\widehat{m} = z + \rho \widetilde{m}_{i^*}^x$ and $\widehat{C} = \rho^2 \widetilde{C}_{i^*}^x$.

Figure 1: The full-density rounding primitive used by the scalar cooling algorithm.

The next lemma concerns the full-density well-rounded rounding.

Lemma 4.5. Fix universal constants $\rho, H, B > 0$ and a fixed accuracy $\varepsilon \in (0, 1/10)$. The algorithm FULLDENSITYROUND (Figure 1), given exact-real evaluation access to a proper lower-semicontinuous convex U whose law $\mu \propto e^{-U}$ is a probability, a known point z with

$$B(z, \rho) \subseteq \{U \leq \inf U + Hn\}, \quad \mathbb{E}_{Y \sim \mu} [\|Y - z\|^2] \leq Bn,$$

and a failure tolerance $p \in (0, 1/10)$, returns either a failure symbol or a pair $(\hat{m}, \hat{C})$ with $\hat{C} \succ 0$. Conditional on every realized pre-call history for which these input promises hold (Definition 2.20), with probability at least $1 - p$ it returns a pair satisfying

$$(1 - \varepsilon)C_\mu \preceq \hat{C} \preceq (1 + \varepsilon)C_\mu, \quad \|C_\mu^{-1/2}(\hat{m} - m_\mu)\| \leq \varepsilon.$$

The number of evaluation queries is at most $O(n^3 \log(n/p)^{O(1)})$ on every execution.

Proof. We analyze Figure 1. Fix the universal integer $n_0 \geq 4$ so large that $c_{\text{iso}} n_0^{-1/2} \leq 1/16$ and $c_{\text{est}}/n_0 \leq 1/16$, with c_{iso} and c_{est} from Theorems 5 and 6. Then $N = n + \lceil Hn \rceil + n_0$ satisfies $n_0 \leq N \leq (H + 1)n + n_0 + 1 = O(n)$. Step 1 is Lemma 4.2: the lifted law $\bar{\mu} = \mu_\rho \otimes \gamma_k$ on $\mathbb{R}^N$ has the restriction $\nu = \bar{\mu}(\cdot \mid K)$ with $\bar{\mu}(K) \geq 1 - 10^{-5}$, and $\nu \propto e^{-U'}$ is grounded around 0 with $\mathbb{E}_\nu [\|X\|^2] \leq 2\bar{B}N = C^2N$ for the universal constant $C^2 = 2B/\rho^2 + 2$. Fix a trial i and condition on the pre-call history and on the other trials; the hypotheses on μ are deterministic properties of the input, so the two black boxes apply with their stated guarantees over the fresh randomness of trial i .

Step 2(a). Theorem 5 applies to ν with the known point 0 and the known constant C , and returns, on an event E_1 of probability at least $1 - c_{\text{iso}}N^{-1/2} \geq 1 - 1/16$, an invertible affine map $F_i(x) = A_i(x - \bar{x}_i)$ with $C_{F_i\#\nu} \approx_{1.01} I_N$ and $\|m_{F_i\#\nu}\| \leq 0.01$. On E_1 the expected number of queries is at most $Q_{\text{iso}}(N)$, so by Markov's inequality the run capped at $16Q_{\text{iso}}(N)$ queries times out on E_1 with probability at most $1/16$. Hence Step 2(a) succeeds with probability at least $1 - 1/8$.

Step 2(b). Apply Lemma 4.3 with $\lambda = \bar{\mu}$, $S = K$, and $\varepsilon = 10^{-5}$: the law $\pi_i = F_i\#\bar{\mu}$ satisfies $0.99I \preceq C_{\pi_i} \preceq 18.5I$ and $\|m_{\pi_i}\| \leq 0.06$. Its potential is $\bar{U} \circ F_i^{-1}$ up to an additive constant, and $\text{supp } F_i\#\nu = F_i(K)$. Since $F_i\#\nu$ is 1.01-isotropic with mean of norm at most 0.01, Proposition 2.23(v) applied to its centering gives

$$B(0, 0.35) \subseteq B\left(m_{F_i\#\nu}, \frac{1}{e^{1.01^{1/2}}}\right) \subseteq F_i(K) \subseteq \{\bar{U} \circ F_i^{-1} \leq \inf(\bar{U} \circ F_i^{-1}) + 10N\},$$

the last inclusion by Item 4.2(3). The potential $W_i(u) = \bar{U}(F_i^{-1}(u/3))$ is, up to an additive constant, that of the law $\pi'_i = (3 \cdot)\#\pi_i$, for which

$$B(0, 1) \subseteq B(0, 1.05) \subseteq \{W_i \leq \inf W_i + 10N\}, \quad 8.9I \preceq C_{\pi'_i} = 9C_{\pi_i} \preceq 166.5I,$$

and

$$\mathbb{E}_{X \sim \pi'_i} [\|X\|^2] = \text{Tr}(C)_{\pi'_i} + \|m_{\pi'_i}\|^2 \leq 166.5N + 0.18^2 \leq 167N.$$

Thus Theorem 6 applies to π'_i with the known point 0, $R^2 = 167N$, $c_+ = 167$, accuracy ξ , and $\eta = 1/64$: on an event E_2 of probability at least $1 - 1/32 - c_{\text{est}}/N \geq 1 - 3/32$ it returns $(\hat{m}'_i, \hat{C}'_i)$ with $(1 - \xi)C_{\pi'_i} \preceq \hat{C}'_i \preceq (1 + \xi)C_{\pi'_i}$ and $\|\hat{m}'_i - m_{\pi'_i}\| \leq \xi$, and on E_2 the run capped at $16Q_{\text{est}}(N)$ queries times out with probability at most $1/16$ by Markov's inequality. Each evaluation of W_i is one evaluation of U by Item 4.2(5), the Gaussian part being explicit. On E_2 without time-out, since $C_{\pi'_i} \succeq 8.9I$,

$$\|C_{\pi'_i}^{-1/2}(\hat{m}'_i - m_{\pi'_i})\| \leq 8.9^{-1/2}\|\hat{m}'_i - m_{\pi'_i}\| \leq \xi,$$

so $(\widehat{m}'_i, \widehat{C}'_i)$ is ξ -accurate for $(m_{\pi'_i}, C_{\pi'_i})$ in the sense of [Lemma 4.4](#), and dividing by 3 and by 9 gives a ξ -accurate pair $(\widehat{m}_i, \widehat{C}_i)$ for (m_{π_i}, C_{π_i}) , both accuracy notions being invariant under scaling. Altogether trial i produces a ξ -accurate pair for (m_{π_i}, C_{π_i}) with probability at least

$$1 - \frac{1}{16} - \frac{1}{16} - \frac{3}{32} - \frac{1}{16} = \frac{23}{32} \geq \frac{17}{32}$$

over its fresh randomness, conditional on the pre-call history and on the other trials.

Step 3. Since $C_{\pi_i} = A_i C_{\bar{\mu}} A_i^\top$ and $m_{\pi_i} = F_i(m_{\bar{\mu}})$, congruence invariance (compare [Definition 2.18](#)) makes the pulled-back pair $(\widetilde{m}_i, \widetilde{C}_i)$ of a successful trial ξ -accurate for $(m_{\bar{\mu}}, C_{\bar{\mu}})$: indeed $\widetilde{C}_i = A_i^{-1} \widehat{C}_i A_i^{-\top}$ shares the Loewner window of $\widehat{C}_i$ relative to C_{π_i} , and $\|C_{\bar{\mu}}^{-1/2}(\widetilde{m}_i - m_{\bar{\mu}})\| = \|C_{\pi_i}^{-1/2}(\widehat{m}_i - m_{\pi_i})\|$. As $\bar{\mu} = \mu_\rho \otimes \gamma_k$, we have $m_{\bar{\mu}} = (m_{\mu_\rho}, 0)$ and $C_{\bar{\mu}} = \text{diag}(C_{\mu_\rho}, I_k)$. Compressing $(1 - \xi)C_{\bar{\mu}} \preceq \widetilde{C}_i \preceq (1 + \xi)C_{\bar{\mu}}$ to the first n coordinates gives $(1 - \xi)C_{\mu_\rho} \preceq \widetilde{C}_i^x \preceq (1 + \xi)C_{\mu_\rho}$, and since $C_{\bar{\mu}}^{-1/2}$ is block diagonal,

$$\|C_{\mu_\rho}^{-1/2}(\widetilde{m}_i^x - m_{\mu_\rho})\| \leq \|C_{\bar{\mu}}^{-1/2}(\widetilde{m}_i - m_{\bar{\mu}})\| \leq \xi,$$

so $(\widetilde{m}_i^x, \widetilde{C}_i^x)$ is ξ -accurate for $(m_{\mu_\rho}, C_{\mu_\rho})$. Every accurate covariance block is positive definite, so the additional check rejects no accurate trial. Treat failed trials as failure symbols in [Lemma 4.4](#). The trials use fresh randomness, so conditional on the pre-call history their outcomes are independent, each ξ -accurate with probability at least $17/32$. [Lemma 4.4](#) with $L = \lceil 512 \log(2/p) \rceil$ then gives, except with probability $p/2 \leq p$, a selected pair that is $5\xi = \varepsilon$ -accurate for $(m_{\mu_\rho}, C_{\mu_\rho})$.

Step 4. Since $C_{\mu_\rho} = \rho^{-2}C_\mu$ and $m_\mu = z + \rho m_{\mu_\rho}$, the returned pair $(z + \rho \widetilde{m}_{i*}^x, \rho^2 \widetilde{C}_{i*}^x)$ is ε -accurate for (m_μ, C_μ) on the selection event, again by congruence invariance.

Query cap. Each trial costs at most $16Q_{\text{iso}}(N) + 16Q_{\text{est}}(N) = O(n^3 \log^{O(1)} n)$ evaluations deterministically, since $N = O(n)$ and the constants A_C, a_C, A, a depend only on $(\rho, H, B, \varepsilon)$; with $L = \lceil 512 \log(2/p) \rceil$ trials the total is $O(n^3 \log(n/p)^{O(1)})$ on every execution. Each query is one evaluation of U plus arithmetic. $\square$

4.2 Cooling Estimates

For the rest of this section, fix a grounded input (V, x_0, r, R) in the sense of [Definition 2.21](#): V is proper, lower-semicontinuous, and convex with $0 < \int e^{-V} < \infty$, $B(x_0, r) \subseteq \{V \leq \inf V + h_0 n\}$, and $\mathbb{E}_{X \sim \pi} [\|X - x_0\|^2] \leq R^2$ for $\pi \propto e^{-V}$. For $t > 0$ set

$$\Phi_t(x) = V(x) + t\|x - x_0\|^2, \quad Z_t = \int e^{-\Phi_t(x)} dx, \quad \pi_t = Z_t^{-1} e^{-\Phi_t},$$

and write $m_t = m_{\pi_t}$, $C_t = C_{\pi_t}$, and $a_t = \inf \Phi_t$. Note $Z_t \leq e^{-\inf V} \int e^{-t\|x - x_0\|^2} dx < \infty$ and $Z_t \geq \text{Vol}(B(x_0, r)) e^{-\inf V - h_0 n - tr^2} > 0$, so π_t is a well-defined log-concave probability law. Evaluating $\Phi_t \circ F^{-1}$ for a known invertible affine F costs one evaluation of V plus arithmetic.

The first estimate is the nonsmooth form of Brascamp–Lieb promised in [Section 2](#); it is used with $G = I_K$ in [Section 5](#).

The next lemma concerns the extended Brascamp–Lieb consequence.

Lemma 4.6. *Let $G : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ be proper, lower-semicontinuous, and convex, let $t > 0$, and suppose $\nu_t \propto e^{-G - t\|x\|^2}$ is a probability law. Then*

$$C_{\nu_t} \preceq \frac{1}{2t} I, \quad \text{Var}_{\nu_t}[g] \leq \frac{1}{2t} \mathbb{E}_{\nu_t} [\|\nabla g\|^2]$$

for every locally Lipschitz g with $\|\nabla g(x)\| \leq c_g(1 + \|x\|)$ almost everywhere.

Proof. Suppose first that G is twice continuously differentiable. Then $W(x) = G(x) + t\|x\|^2$ satisfies $\nabla^2 W \succeq 2tI$, so [Proposition 2.23](#)(vi) applies and gives both claims: $C_{\nu_t} \preceq (2t)^{-1}I$ and $\text{Var}_{\nu_t}[g] \leq (2t)^{-1}\mathbb{E}_{\nu_t}[\|\nabla g\|^2]$.

For the general case, fix an affine minorant ℓ of G (which exists by properness) and approximate in two steps. First, the Moreau envelope $G_\eta(x) = \inf_y \{G(y) + \frac{1}{2\eta}\|x - y\|^2\}$ is convex, continuously differentiable with Lipschitz gradient, increases to G pointwise as $\eta \downarrow 0$, and satisfies $G_\eta \geq \ell - C_\eta$ for some $C_\eta < \infty$. Second, the Gaussian mollification $G_{\eta,\sigma} = G_\eta * \rho_\sigma$ is smooth and convex, and by Jensen $G_{\eta,\sigma}(x) = \mathbb{E}[G_\eta(x + \sigma\xi)] \geq G_\eta(x)$. The smooth case applies to $G_{\eta,\sigma}$, giving both inequalities for $\nu_{\eta,\sigma} \propto e^{-G_{\eta,\sigma} - t\|x\|^2}$. The densities are dominated uniformly in (η, σ) : with $\tilde{\ell}(x) = \ell(x) - C_1$,

$$e^{-G_{\eta,\sigma}(x) - t\|x\|^2} \leq e^{-G_\eta(x) - t\|x\|^2} \leq e^{-\tilde{\ell}(x) - t\|x\|^2},$$

whose right side is integrable together with $\|x\|^2$ times itself. As $\sigma \downarrow 0$, $G_{\eta,\sigma} \rightarrow G_\eta$ locally uniformly, so dominated convergence identifies $\nu_{\eta,\sigma} \rightarrow \nu_\eta$ and convergence of the covariance and of both sides of the variance inequality (using the growth bound on $\|\nabla g\|$); then $\eta \downarrow 0$ gives $G_\eta \uparrow G$ and the same passage yields both claims for ν_t . $\square$

The next lemma concerns the quadratic energy and variance.

Lemma 4.7. *There is a constant $K_E = K_E(h_0)$ such that, for every $t \in (0, r^{-2}]$,*

$$t \mathbb{E}_{\pi_t} [\|X - x_0\|^2] \leq K_E n, \quad \text{Var}_{\pi_t} [\|X - x_0\|^2] \leq \frac{2K_E n}{t^2}.$$

Proof. Covariance is unaffected by translation, so applying [Lemma 4.6](#) with $G(x) = V(x + x_0)$ gives $\text{Tr}(C_t) \leq \frac{n}{2t}$. For the mean, the barycenter inequality ([Proposition 2.23](#)(i)) gives $\Phi_t(m_t) \leq a_t + n$, and any $y \in B(x_0, r)$ with $V(y) \leq \inf V + h_0 n$ gives $a_t \leq \Phi_t(y) \leq \inf V + h_0 n + tr^2$; hence

$$t\|m_t - x_0\|^2 \leq \Phi_t(m_t) - \inf V \leq a_t - \inf V + n \leq (h_0 + 1)n + tr^2 \leq (h_0 + 2)n,$$

using $tr^2 \leq 1$. Summing, $t \mathbb{E}_{\pi_t} [\|X - x_0\|^2] = t \text{Tr}(C_t) + t\|m_t - x_0\|^2 \leq \frac{n}{2} + (h_0 + 2)n =: K_E n$. The variance claim is the same application of [Lemma 4.6](#) to $g(x) = \|x - x_0\|^2$ with $\|\nabla g(x)\| = 2\|x - x_0\|$: $\text{Var}_{\pi_t}[g] \leq \frac{1}{2t} \mathbb{E}_{\pi_t} [4\|X - x_0\|^2] \leq \frac{2K_E n}{t^2}$. $\square$

The next lemma concerns the one-step Rényi cost.

Lemma 4.8. *For $0 < \delta \leq 1/4$, $s = (1 - \delta)t$, and $t \leq r^{-2}$,*

$$D_2(\pi_s \| \pi_t) \leq K_D n \delta^2, \quad K_D = 8K_E.$$

Proof. Since $2\Phi_s - \Phi_t = \Phi_{2s-t}$ with $2s - t = (1 - 2\delta)t > 0$,

$$\mathcal{R}_2(\pi_s \| \pi_t) = \int \frac{\pi_s^2}{\pi_t} = \frac{Z_t Z_{2s-t}}{Z_s^2}, \quad \text{so with } \Lambda(u) = \log Z_u, \quad D_2(\pi_s \| \pi_t) = \Lambda(t) + \Lambda(2s - t) - 2\Lambda(s),$$

a centered second difference. Differentiating twice, $\Lambda''(u) = \text{Var}_{\pi_u} [\|X - x_0\|^2] \leq 2K_E n/u^2$ by [Lemma 4.7](#). Taylor's theorem with $h = \delta t$, noting $s - h = (1 - 2\delta)t \geq t/2$, gives

$$D_2(\pi_s \| \pi_t) = \int_0^h (h - w) [\Lambda''(s + w) + \Lambda''(s - w)] dw \leq h^2 \cdot \frac{2K_E n}{(t/2)^2} = 8K_E n \delta^2. \quad \square$$

The next lemma concerns the Rényi-to-trace transfer.

Lemma 4.9. *There are universal constants $\gamma \in (0, 1/4]$ and $K_T = K_T(h_0)$ such that, with $\delta = \gamma n^{-1/4}$ and $s = (1 - \delta)t$, the following holds for every $t \leq r^{-2}$. If F is an invertible affine map for which $Z = F(X)$, $X \sim \pi_t$, satisfies*

$$\frac{1}{2}I \preceq C_{F\#\pi_t} \preceq 2I, \quad \|\mathbb{E}_{F\#\pi_t}[Z]\| \leq \rho_L/2,$$

then $\mathbb{E}_{\pi_s}[\|F(X)\|^2] \leq K_T n$.

Proof. Write $L = \frac{d\pi_s}{d\pi_t}(X)$ and split $\mathbb{E}_{\pi_s}[\|Z\|^2] = \mathbb{E}_{\pi_t}[\|Z\|^2 L]$ on the Paouris event $A = \{\|Z\| \leq u c_P n^{1/2}\}$, with c_P from [Proposition 2.23\(ii\)](#) and $u \geq 2$ universal, chosen below. On A the integrand is at most $u^2 c_P^2 n L$, contributing at most $u^2 c_P^2 n$. Off A , Hölder with exponents $(4, 2, 4)$ gives

$$\mathbb{E}_{\pi_t}[\|Z\|^2 L \mathbb{1}_{A^c}] \leq (\mathbb{E}_{\pi_t}[\|Z\|^8])^{1/4} (\mathbb{E}_{\pi_t}[L^2])^{1/2} \mathbb{P}_{\pi_t}[A^c]^{1/4}.$$

The three factors are: $(\mathbb{E}_{\pi_t}[\|Z\|^8])^{1/4} \leq 16^2 \mathbb{E}_{\pi_t}[\|Z\|^2] \leq 768n$ by [Lemma 2.12](#) with $k = 8$ and $\mathbb{E}_{\pi_t}[\|Z\|^2] = \text{Tr}(C_{F\#\pi_t}) + \|m_{F\#\pi_t}\|^2 \leq 2n + \rho_L^2/4 \leq 3n$; $(\mathbb{E}_{\pi_t}[L^2])^{1/2} = \mathcal{R}_2(\pi_s\|\pi_t)^{1/2} \leq e^{K_D \gamma^2 \sqrt{n}/2}$ by [Lemma 4.8](#); and $\mathbb{P}_{\pi_t}[A^c]^{1/4} \leq e^{-u\sqrt{n}/(8c_P)}$ by [Proposition 2.23\(ii\)](#) applied to $Z - m_{F\#\pi_t}$, with the mean shift $\rho_L/2$ absorbed by the choice $u \geq 2$. Choosing $u = \max(2, 8c_P K_D \gamma^2)$ makes the combined exponent at most $-K_D \gamma^2 n^{1/2}/4$ in either case, so

$$\mathbb{E}_{\pi_t}[\|Z\|^2 L \mathbb{1}_{A^c}] \leq 768n \cdot e^{-K_D \gamma^2 \sqrt{n}/4} \leq 768n.$$

In total,

$$\mathbb{E}_{\pi_s}[\|Z\|^2] \leq (u^2 c_P^2 + 768)n = K_T n, \quad K_T := u^2 c_P^2 + 768.$$

□

The next lemma concerns the moving minimum.

Lemma 4.10. *For $s = (1 - \delta)t$ and $t \leq r^{-2}$,*

$$0 \leq a_t - a_s \leq \frac{\delta}{1 - \delta} (h_0 + 1)n.$$

Proof. Since $\Phi_t - \Phi_s = (t - s)\|x - x_0\|^2 \geq 0$, we have $a_t \geq a_s$. Let y minimize Φ_s . As in [Lemma 4.7](#), $a_s \leq \inf V + h_0 n + s r^2$, so $s\|y - x_0\|^2 \leq \Phi_s(y) - \inf V \leq (h_0 + 1)n$. Then

$$a_t \leq \Phi_t(y) = a_s + (t - s)\|y - x_0\|^2 \leq a_s + \frac{t - s}{s} (h_0 + 1)n = a_s + \frac{\delta}{1 - \delta} (h_0 + 1)n. \quad \square$$

We now fix the universal constants used by the cooling induction:

$$\rho_* = \rho_L/2, \quad H_* = H_L + h_0 + 1, \quad B_* = \max(K_T, K_E/\rho_*^2), \quad \varepsilon_{\text{cool}} = \min(1/500, \rho_L/3),$$

where ρ_L, H_L are from [Proposition 2.23\(iv\)](#) and K_E, K_T from [Lemmas 4.7](#) and [4.9](#).

The next lemma concerns the well-roundedness at the next temperature.

Lemma 4.11. *Let $s = (1 - \delta)t$ with $\delta = \gamma n^{-1/4}$ and $t \leq r^{-2}$. Suppose a call to FULLDENSITY-ROUND on π_t (in affine coordinates in which π_t is (ρ_*, H_*, B_*) -well-rounded around 0) succeeded and returned $(\hat{m}, \hat{C})$, and set $F(x) = \hat{C}^{-1/2}(x - \hat{m})$. Then in these coordinates π_s is (ρ_*, H_*, B_*) -well-rounded around 0:*

$$B(0, \rho_*) \subseteq \{\Phi_s \circ F^{-1} \leq a_s + H_* n\}, \quad \mathbb{E}_{\pi_s}[\|F(X)\|^2] \leq B_* n.$$

Proof. On the success event, $(1 - \varepsilon_{\text{cool}})C_t \preceq \widehat{C} \preceq (1 + \varepsilon_{\text{cool}})C_t$ and $\|C_t^{-1/2}(\widehat{m} - m_t)\| \leq \varepsilon_{\text{cool}}$. Hence

$$C_{F_{\#}\pi_t} = \widehat{C}^{-1/2}C_t\widehat{C}^{-1/2} \in [(1 + \varepsilon_{\text{cool}})^{-1}, (1 - \varepsilon_{\text{cool}})^{-1}]I \subseteq [\tfrac{1}{2}I, 2I],$$

and $\|m_{F_{\#}\pi_t}\| = \|\widehat{C}^{-1/2}(m_t - \widehat{m})\| \leq (1 - \varepsilon_{\text{cool}})^{-1/2} \varepsilon_{\text{cool}} \leq \rho_L/2$. [Proposition 2.23](#)(iv) (the relaxed-centering form) now gives $B(0, \rho_L/2) \subseteq \{\Phi_t \circ F^{-1} \leq a_t + H_L n\}$. For z in that ball, since $\Phi_s \leq \Phi_t$ pointwise and by [Lemma 4.10](#),

$$\Phi_s(F^{-1}z) \leq a_t + H_L n \leq a_s + H_L n + \frac{\delta}{1 - \delta}(h_0 + 1)n \leq a_s + (H_L + h_0 + 1)n = a_s + H_* n,$$

using $\frac{\delta}{1 - \delta} \leq 1$. The second-moment bound is [Lemma 4.9](#): $\mathbb{E}_{\pi_s} [\|F(X)\|^2] \leq K_T n \leq B_* n$. $\square$

The next lemma concerns the initialization.

Lemma 4.12. *For $t_0 = r^{-2}$ and $F_{\text{base}}(x) = (x - x_0)/(r\rho_*)$, the law π_{t_0} in F_{base} -coordinates is (ρ_*, H_*, B_*) -well-rounded around 0.*

Proof. By [Lemma 4.7](#) with $t_0 r^2 = 1$,

$$\mathbb{E}_{\pi_{t_0}} [\|F_{\text{base}}(X)\|^2] = \frac{\mathbb{E}_{\pi_{t_0}} [\|X - x_0\|^2]}{r^2 \rho_*^2} \leq \frac{K_E n}{t_0 r^2 \rho_*^2} = \frac{K_E n}{\rho_*^2} \leq B_* n.$$

For $\|z\| \leq \rho_*$ (≤ 1 without loss of generality, since the constants may be shrunk), the point $F_{\text{base}}^{-1}z = x_0 + r\rho_* z$ lies in $B(x_0, r)$, so $V(F_{\text{base}}^{-1}z) \leq \inf V + h_0 n$ and

$$\Phi_{t_0}(F_{\text{base}}^{-1}z) \leq \inf V + h_0 n + t_0 r^2 \rho_*^2 \|z\|^2 \leq \inf V + (h_0 + 1)n \leq a_{t_0} + H_* n,$$

using $a_{t_0} \geq \inf V$ and $H_* \geq h_0 + 1$. $\square$

4.3 Endpoint Comparison and Raw Reconstruction

At the end of the cooling path the Gaussian tilt is weak enough that the cooled moments and the target moments are comparable, uniformly in the centering. Write

$$M(b) = \mathbb{E}_{\pi} [\langle (X - b)(X - b)^{\top} \rangle], \quad M_t(b) = \mathbb{E}_{\pi_t} [\langle (X - b)(X - b)^{\top} \rangle].$$

The next lemma concerns the Reverse-Hölder endpoint comparison.

Lemma 4.13. *There is a universal K_H such that, whenever $\eta = tR^2 \leq 1/2$,*

$$(1 - K_H \eta) M(b) \preceq M_t(b) \preceq (1 + 2\eta) M(b) \quad \text{for every } b \in \mathbb{R}^n.$$

Proof. The likelihood ratio is $\frac{d\pi_t}{d\pi}(x) = \frac{1}{c_t} e^{-t\|x - x_0\|^2}$ with $c_t = Z_t/Z_V = \mathbb{E}_{\pi} [e^{-t\|X - x_0\|^2}]$. Jensen and the grounded moment bound give $c_t \geq e^{-t\mathbb{E}_{\pi} [\|X - x_0\|^2]} \geq e^{-\eta}$, while $c_t \leq 1$. The upper bound is immediate: $M_t(b) = \frac{1}{c_t} \mathbb{E}_{\pi} [\langle (X - b)(X - b)^{\top} e^{-t\|X - x_0\|^2} \rangle] \preceq c_t^{-1} M(b) \preceq e^{\eta} M(b) \preceq (1 + 2\eta) M(b)$ for $\eta \leq 1/2$. For the lower bound, since $c_t \leq 1$ and $1 - e^{-u} \leq u$, for every unit vector v ,

$$v^{\top} (M(b) - M_t(b)) v \leq \mathbb{E}_{\pi} [\langle X - b, v \rangle^2 (1 - e^{-t\|X - x_0\|^2})] \leq t \mathbb{E}_{\pi} [\langle X - b, v \rangle^2 \|X - x_0\|^2].$$

Cauchy-Schwarz, [Proposition 2.23](#)(iii) applied to the affine functional $\langle X - b, v \rangle$, and [Lemma 2.12](#) with $k = 4$ give

$$t (\mathbb{E}_{\pi} [\langle X - b, v \rangle^4])^{1/2} (\mathbb{E}_{\pi} [\|X - x_0\|^4])^{1/2} \leq t C_4^{1/2} (v^{\top} M(b) v) \cdot 64 R^2 = 64 C_4^{1/2} \eta v^{\top} M(b) v,$$

so $M_t(b) \succeq (1 - K_H \eta) M(b)$ with $K_H = 64 C_4^{1/2}$. $\square$

The next lemma states that exact target is well-rounded.

Lemma 4.14. *Set $\eta_0 = \min(1/2, 1/(4K_H))$. Suppose $tR^2 \leq \eta_0$ and that, in some invertible affine coordinates F , the cooled law π_t is (ρ_*, H_*, B_*) -well-rounded around 0. Then the target law π in the same coordinates is $(\rho_*, H'_*, 2B_*)$ -well-rounded around 0 with $H'_* = H_* + h_0 + 1$:*

$$B(0, \rho_*) \subseteq \{V \circ F^{-1} \leq \inf(V \circ F^{-1}) + H'_*n\}, \quad \mathbb{E}_\pi [\|F(X)\|^2] \leq 2B_*n.$$

Proof. Write $F(x) = A(x - b)$. By Lemma 4.13 with $K_H\eta_0 \leq 1/4$, $M(b) \preceq (1 - K_H\eta_0)^{-1}M_t(b) \preceq \frac{4}{3}M_t(b)$, so

$$\mathbb{E}_\pi [\|F(X)\|^2] = \text{Tr} \left(A^\top A M(b) \right) \leq \frac{4}{3} \text{Tr} \left(A^\top A M_t(b) \right) = \frac{4}{3} \mathbb{E}_{\pi_t} [\|F(X)\|^2] \leq 2B_*n.$$

For the ball, since $\Phi_t = V + t\|x - x_0\|^2 \geq V$, on $B(0, \rho_*)$ we have $V \circ F^{-1} \leq \Phi_t \circ F^{-1} \leq a_t + H_*n$, and $a_t \leq \inf V + h_0n + tr^2 \leq \inf V + (h_0 + 1)n$ because $tr^2 \leq t_0r^2 = 1$; hence $V \circ F^{-1} \leq \inf V + (H_* + h_0 + 1)n$ on $B(0, \rho_*)$. $\square$

The next lemma concerns the raw second-moment reconstruction.

Lemma 4.15. *Let $C \succ 0$, $m \in \mathbb{R}^n$, and fix any center $c \in \mathbb{R}^n$; write $M(c) = C + (m - c)(m - c)^\top$. If $(1 - \varepsilon)C \preceq \widehat{C} \preceq (1 + \varepsilon)C$ and $\|C^{-1/2}(\widehat{m} - m)\| \leq \zeta$, then $\widehat{M}(c) = \widehat{C} + (\widehat{m} - c)(\widehat{m} - c)^\top$ is positive semidefinite and*

$$-(\varepsilon + \zeta + \zeta^2) M(c) \preceq \widehat{M}(c) - M(c) \preceq (\varepsilon + \zeta + \zeta^2) M(c).$$

For $\varepsilon = \zeta = 1/500$ this gives $0.99\widehat{M}(c) \preceq M(c) \preceq 1.01\widehat{M}(c)$.

Proof. Let $d = \widehat{m} - m$. Then $\widehat{M}(c) - M(c) = (\widehat{C} - C) + dd^\top + (m - c)d^\top + d(m - c)^\top$. The first term obeys $-\varepsilon C \preceq \widehat{C} - C \preceq \varepsilon C \preceq \varepsilon M(c)$; the second $0 \preceq dd^\top \preceq \zeta^2 C \preceq \zeta^2 M(c)$ from $\|C^{-1/2}d\| \leq \zeta$; and for the cross term, every unit vector v obeys

$$\begin{aligned} |v^\top ((m - c)d^\top + d(m - c)^\top)v| &= 2|\langle v, m - c \rangle| |\langle v, d \rangle| \\ &\leq 2(v^\top (m - c)(m - c)^\top v)^{1/2} (\zeta^2 v^\top C v)^{1/2} \\ &\leq \zeta v^\top M(c) v, \end{aligned}$$

by $2ab \leq a^2 + b^2$. Summing the three gives the Loewner claim. With $\alpha = \varepsilon + \zeta + \zeta^2 = 0.004004 < \frac{1}{249}$, the bounds $(1 - \alpha)M(c) \preceq \widehat{M}(c) \preceq (1 + \alpha)M(c)$ imply

$$(1 + \alpha)^{-1}\widehat{M}(c) \preceq M(c) \preceq (1 - \alpha)^{-1}\widehat{M}(c).$$

Thus $0.99\widehat{M}(c) \preceq M(c) \preceq 1.01\widehat{M}(c)$. $\square$

4.4 Algorithm, Induction, and Complexity




Figure 2: The general-body scalar-cooling rounding algorithm.

The next lemma concerns the cooling induction.

Lemma 4.16. *Conditional on all earlier calls succeeding, every **FULLDENSITYROUND** call made by **ASYMMETRICSECONDMOMENT** ([Figure 2](#)) satisfies the input promises of [Lemma 4.5](#), and each evaluation of the transformed potential costs one evaluation of V .*

Proof. At $j = 0$, the promises hold by [Lemma 4.12](#). If the call at stage $j \geq 0$ succeeds, then [Lemma 4.11](#) applied in the current coordinates shows that $\pi_{t_{j+1}}$ is (ρ_*, H_*, B_*) -well-rounded around 0 in the updated coordinates, which are exactly the input promises for the call at stage $j + 1$. For the final call, [Lemma 4.14](#) applies since $t_N R^2 \leq \eta_0$, so the exact target $V \circ F^{-1}$'s law is $(\rho_*, H'_*, 2B_*)$ -well-rounded around 0. Evaluation of $\Phi_t \circ F^{-1}(z) = V(F^{-1}z) + t\|F^{-1}z - x_0\|^2$ (or of $V \circ F^{-1}$) is one evaluation of V plus arithmetic, and the conditional guarantee of [Lemma 4.5](#) ([Definition 2.20](#)) is exactly what is invoked at each stage. $\square$

The next lemma concerns the radius comparison.

Lemma 4.17. *Every grounded input obeys $R \geq c_{h_0} r$ with $c_{h_0} = \frac{1}{4}e^{-h_0}$.*

Proof. Suppose $r > 4e^{h_0} R$. Since $\mathbb{E}_\pi [\|X - x_0\|^2] \leq R^2$, Markov gives $\pi(B(x_0, 2R)) \geq 3/4$, and

$B(x_0, 2R) \subseteq B(x_0, r)$. Writing $Z_V = \int e^{-V}$, on $B(x_0, r)$ we have $V \leq \inf V + h_0 n$, so

$$\frac{3}{4} \leq \pi(B(x_0, 2R)) \leq \frac{\text{Vol}(B(x_0, 2R)) e^{-\inf V}}{Z_V}, \quad \pi(B(x_0, r)) \geq \frac{\text{Vol}(B(x_0, r)) e^{-\inf V - h_0 n}}{Z_V}.$$

Dividing, $\pi(B(x_0, r)) \geq (r/2R)^n e^{-h_0 n} \cdot \frac{3}{4} > (2e^{h_0})^n e^{-h_0 n} \cdot \frac{3}{4} = \frac{3}{4} \cdot 2^n > 1$, a contradiction. $\square$

The next lemma concerns the stage count and total cost.

Lemma 4.18. *The number of cooling stages satisfies*

$$N = O(n^{1/4} \log(2 + R/r)).$$

With $p = \min(1/20, \beta/(N+1))$, all $N+1$ calls of `ASYMMETRICSECONDMOMENT` succeed simultaneously with probability at least $1 - \beta$, and the total number of evaluation queries is

$$O\left(n^{13/4} \log\left(2 + \frac{nR}{r\beta}\right)^{O(1)}\right)$$

on every execution.

Proof. By the definition of N and $(1 - \delta)^j \leq e^{-\delta j}$,

$$N \leq \max\left\{0, \left\lceil \delta^{-1} \log \frac{t_0 R^2}{\eta_0} \right\rceil\right\} = O(n^{1/4} \log(2 + R/r)),$$

since $\delta = \gamma n^{-1/4}$ and η_0 is universal. By the conditional guarantees ([Lemma 4.16](#) and [Definition 2.20](#)), the failure probability of the whole run is at most $(N+1)p \leq \beta$. Each call costs at most $O(n^3 \log(n/p)^{O(1)})$ deterministically, and

$$\log(n/p) = O(\log(n(N+1)/\beta)) = O\left(\log\left(2 + \frac{nR}{r\beta}\right)\right),$$

so the total is $O(n^{1/4} \log(2 + R/r) \cdot n^3 \log(2 + nR/(r\beta))^{O(1)}) = O(n^{13/4} \log(2 + nR/(r\beta))^{O(1)})$. $\square$

The next theorem concerns the grounded-potential rounding.

Theorem 10. *Given a grounded input (V, x_0, r, R) as in [Definition 2.21](#) and $\beta \in (0, 1)$, the algorithm `ASYMMETRICSECONDMOMENT` ([Figure 2](#)) makes $O(n^{13/4} \log(2 + nR/(r\beta))^{O(1)})$ evaluation queries on every execution, and with probability at least $1 - \beta$ returns a positive semidefinite $\widehat{M}$ with*

$$0.99\widehat{M} \preceq M_\pi \preceq 1.01\widehat{M}, \quad \pi \propto e^{-V},$$

together with an affine map $\mathcal{T}$ such that $C_{\mathcal{T}\# \pi} \approx_{1.01} I$ and $\|m_{\mathcal{T}\# \pi}\| \leq 1/250$.

Proof. On the event that all $N+1$ calls succeed, which has probability at least $1 - \beta$ by [Lemma 4.18](#), every call returned accurate estimates in the sense of [Lemma 4.5](#). In particular the final call returned $(\widehat{m}, \widehat{C})$ with $(1 - \varepsilon_{\text{cool}})C_{F\# \pi} \preceq \widehat{C} \preceq (1 + \varepsilon_{\text{cool}})C_{F\# \pi}$ and $\|C_{F\# \pi}^{-1/2}(\widehat{m} - m_{F\# \pi})\| \leq \varepsilon_{\text{cool}}$, where $\varepsilon_{\text{cool}} \leq 1/500$. With $F(x) = A(x - b)$ and center $c = -Ab$, [Definition 2.18](#) gives $M_\pi = A^{-1}(C_{F\# \pi} + (m_{F\# \pi} + Ab)(m_{F\# \pi} + Ab)^\top)A^{-\top}$, which is exactly $A^{-1}M_{F\# \pi}(c)A^{-\top}$, while the returned matrix is $\widehat{M} = A^{-1}\widehat{M}_{F\# \pi}(c)A^{-\top}$. [Lemma 4.15](#) with $\varepsilon = \zeta = \varepsilon_{\text{cool}} \leq 1/500$ and Loewner congruence by A^{-1} give $0.99\widehat{M} \preceq M_\pi \preceq 1.01\widehat{M}$; $\widehat{M}$ is positive semidefinite since $\widehat{C} \succ 0$. The map $\mathcal{T}$ pushes π to a law with covariance $\widehat{C}^{-1/2}C_{F\# \pi}\widehat{C}^{-1/2} \in [(1 + \varepsilon_{\text{cool}})^{-1}, (1 - \varepsilon_{\text{cool}})^{-1}]I \subseteq [1/1.01, 1.01]I$ and mean norm $\|\widehat{C}^{-1/2}(m_{F\# \pi} - \widehat{m})\| \leq (1 - \varepsilon_{\text{cool}})^{-1/2} \varepsilon_{\text{cool}} \leq 1/250$. The query bound is [Lemma 4.18](#). $\square$

Proof of Theorem 1. Set $V = I_K$, the indicator potential of Definition 2.19: evaluation of V is exactly membership querying of K , and $\pi = \pi_K$. The input is grounded for any $h_0 > 0$, since $B(x_0, r) \subseteq K = \{V \leq 0\} = \{V \leq \inf V + h_0 n\}$ and $\mathbb{E}_{\pi_K} [\|X - x_0\|^2] \leq R^2$ from $K \subseteq B(x_0, R)$. Theorem 10 then gives the query bound, the success probability, and $0.99\widehat{M} \preceq \mathbb{E}_{x \sim K} [xx^\top] = M_K \preceq 1.01\widehat{M}$, as required. On the same success event, Theorem 10 also returns $\mathcal{T}$ with $C_{\mathcal{T}_{\#}\pi_K} \in [1/1.01, 1.01]I$ and $\|m_{\mathcal{T}_{\#}\pi_K}\| \leq 1/250$; its linear part is invertible because this covariance is positive definite. Fact 2.24 with $\lambda = 1/1.01$ gives $B(m_{\mathcal{T}_{\#}\pi_K}, 1.01^{-1/2}) \subseteq \mathcal{T}(K)$, and $1.01^{-1/2} - 1/250 \geq 0.99$, so $B(0, 0.99) \subseteq \mathcal{T}(K)$. Finally

$$\mathbb{E}_{x \sim \mathcal{T}(K)} [\|x\|^2] = \text{Tr}(C_{\mathcal{T}_{\#}\pi_K}) + \|m_{\mathcal{T}_{\#}\pi_K}\|^2 \leq 1.01n + 250^{-2} \leq 1.02n.$$

□

Proof of Corollary 1.1. Run ASYMMETRICSECONDMOMENT with $\beta = 1/8$. On its success event, Theorem 1 gives an invertible affine map $\mathcal{T}(x) = Ax + a$ with $B(0, 0.99) \subseteq \mathcal{T}(K)$ and $\mathbb{E}_{x \sim \mathcal{T}(K)} [\|x\|^2] \leq 1.02n$. The body $K' = 0.99^{-1}\mathcal{T}(K)$ contains the unit ball, satisfies $\mathbb{E}_{x \sim K'} [\|x\|^2] \leq 1.02n/0.99^2 \leq 1.05n$, and $\text{Vol}(K) = 0.99^n \text{Vol}(K')/|\det A|$; a membership query to K' is one membership query to K . Fact 2.25 with $R^2 = 1.05n$ and $p = 1/8$ returns, with probability at least $7/8$, a number within a factor $1 \pm \varepsilon$ of $\text{Vol}(K')$ after $O(n^3 \varepsilon^{-2} \log(n/\varepsilon)^{O(1)})$ queries, and the output is $0.99^n |\det A|^{-1}$ times that number. A union bound gives success probability at least $3/4$, and the query count is the sum of the two bounds, with $\log(2 + 8nR/r) = O(\log(2 + nR/r))$.

If K is symmetric and $x_0 = 0$, run SYMMETRIC-COVARIANCE with $\beta = 1/8$ instead. On its success event, π_K has mean 0 and covariance $C_K = \mathbb{E}_{x \sim K} [xx^\top]$ with $0.99\Sigma \preceq C_K \preceq 1.01\Sigma$, so $\mathcal{T} = \Sigma^{-1/2}$ makes the covariance of the uniform law on $\mathcal{T}(K)$ lie in $[0.99, 1.01]I$. Fact 2.24 gives $B(0, 0.99^{1/2}) \subseteq \mathcal{T}(K)$, and $\mathbb{E}_{x \sim \mathcal{T}(K)} [\|x\|^2] \leq 1.01n$, so the argument above applies to $K' = 0.99^{-1/2}\mathcal{T}(K)$, with the rounding cost of Theorem 2. □

4.5 Remarks on the Step Size

The exponent $13/4$ is the cost of one full-density rounding call per cooling stage, times the number $O(n^{1/4} \log(2 + R/r))$ of stages; the stage count is forced by the step size $\delta = \Theta(n^{-1/4})$, which in turn is the largest step the moment transfer of Lemma 4.9 tolerates: the one-step Rényi cost of Lemma 4.8 is $O(n\delta^2)$, hence $O(n^{1/2})$ at this δ , and the Paouris tail can absorb exactly such $e^{O(\sqrt{n})}$ warmth. Two conjectural simplifications would remove the $n^{1/4}$ overhead, and both fail. First, the coordinate-moment bound of Conjecture 3.1 would make the per-coordinate derivative of the cooled trace $O(1)$ -bounded, permitting constant-size steps; it is refuted by Theorem 9. Second, even the trace transfer itself is sharp: Theorem 7 exhibits restricted Gaussians whose trace after a warming step is $\Omega(n + n^2 \delta^4)$, which stays $O(n)$ exactly when $\delta = O(n^{-1/4})$. In this sense the scalar moment bounds used above cannot support larger steps; the matrix-valued cooling of Section 5 is what removes the overhead for symmetric bodies.

5 Cubic Algorithm for the Symmetric Case

We now describe the symmetric cubic algorithm. The algorithm still uses Gaussian cooling, but it replaces the scalar temperature tI by a matrix precision H . This allows the cooling schedule to treat different directions differently: directions that remain well-conditioned can keep cooling, while directions that threaten the invariant receive extra regularization. The rest of this section gives the full proof.

Assume $K = -K$ and $rB_2^n \subseteq K \subseteq RB_2^n$. For $H \succeq 0$, let

$$U_H(x) = I_K(x) + x^\top Hx, \quad \mu_H(x) \propto e^{-x^\top Hx} \mathbf{1}_K(x), \quad C_H = \mathbb{E}_{x \sim \mu_H} [xx^\top].$$

We treat the function U_H as the potential, or negative log-density up to an additive constant, of μ_H . It records both the constraint $x \in K$ and the quadratic weight $x^\top Hx$. The algorithm changes H to remove this weight gradually. At every successful iteration i , it maintains the potential U_{H_i} together with a covariance estimate satisfying

$$W_i \approx_{\kappa_{\text{app}}} C_{H_i}, \quad \kappa_{\text{app}} = 1.01.$$

Initially $H_0 = r^{-2}I$. Under the change of variables $x = rz$, the promised inclusion $rB_2^n \subseteq K$ becomes $B_2^n \subseteq r^{-1}K$, and the initial transformed potential is

$$P_0(z) = U_{H_0}(rz) = I_K(rz) + \|z\|_2^2.$$

If $\|z\|_2 \leq 1$, then $rz \in K$, so

$$P_0(z) = \|z\|_2^2 \leq 1.$$

Moreover, $P_0 \geq 0$ and $P_0(0) = 0$. Therefore

$$B_2^n \subseteq \{z : P_0(z) \leq \inf P_0 + 1\}.$$

We refer to this ball as the sublevel-set ball for P_0 . Since $K = -K$, both $I_K(rz)$ and $\|z\|_2^2$ are even functions of z . Hence P_0 is even, the density proportional to e^{-P_0} is even, and the corresponding random vector Z has mean zero. Applying [Lemma 4.6](#) with $G = I_{r^{-1}K}$ and $t = 1$ gives

$$\text{Cov}(Z) \preceq \frac{1}{2}I, \quad \mathbb{E}_{P_0} [\|Z\|_2^2] \leq \frac{n}{2}.$$

Thus the initial law has the required $O(n)$ second-moment bound. The near-cubic covariance estimator in [Theorem 11](#) can therefore construct W_0 .

Given a pair (H, W) , the next precision is given by

$$H^+ = (1 - \delta)H + \eta_{\text{ridge}} \frac{\text{Tr}(WH)}{n} W^{-1}, \quad \delta = L^{-2}, \quad \eta_{\text{ridge}} = L^{-3},$$

where $L = \Theta(1 + \log(2 + nR/r))$. The vector $z = W^{-1/2}x$ is the whitened version of x obtained from the current covariance estimate W . Since $W \approx_{\kappa_{\text{app}}} C_H$, its covariance satisfies

$$W^{-1/2}C_HW^{-1/2} \approx_{\kappa_{\text{app}}} I,$$

so all directions have comparable variance in the z -coordinates. In these coordinates, the update becomes

$$J = (1 - \delta)B + \eta_{\text{ridge}} \frac{\text{Tr}(B)}{n} I, \quad B = W^{1/2}HW^{1/2}.$$

Thus B is the quadratic penalty in these well-conditioned coordinates, and $J = W^{1/2}H^+W^{1/2}$. If λ is an eigenvalue of B , the update maps it to

$$(1 - \delta)\lambda + \eta_{\text{ridge}} \frac{\text{Tr}(B)}{n}.$$

The additive term therefore has its largest relative effect when λ is small. Moreover,

$$\mathrm{Tr}(B) = \mathrm{Tr}(W H) \approx_{\kappa_{\mathrm{app}}} \mathrm{Tr}(C_H H).$$

The ridge adds total trace $\eta_{\mathrm{ridge}} \mathrm{Tr}(B)$, whereas the cooling term removes $\delta \mathrm{Tr}(B)$ from B .

The invariant rests on two estimates. First, even-sector Poincaré bounds along the interpolation from B to J show that the next law has $O(n)$ second moment in the z -coordinates. The same fixed-radius ball remains in an $O(n)$ -height sublevel set of the next potential. These are the two well-roundedness hypotheses of the near-cubic covariance estimator, so the next covariance estimate costs $\tilde{O}(n^3)$ membership queries. Second, define

$$D_{\mathrm{sym}}(H) = \log \frac{\mathrm{Vol}(K)}{\int_K e^{-x^\top H x} dx}, \quad q_{\mathrm{sym}}(H) = \mathrm{Tr}(C_H H) = \mathbb{E}_{X \sim \mu_H} [X^\top H X],$$

so, for X uniform on K ,

$$e^{-D_{\mathrm{sym}}(H)} = \mathbb{E}_{X \sim K} [e^{-X^\top H X}].$$

We can understand $D_{\mathrm{sym}}(H)$ as measuring how much the quadratic weight decreases the normalizing integral, while $q_{\mathrm{sym}}(H)$ is the expected quadratic penalty under the current distribution μ_H . In particular, $D_{\mathrm{sym}}(H) \geq 0$ and $D_{\mathrm{sym}}(0) = 0$. The regularized update gives

$$D_{\mathrm{sym}}(H^+) \leq D_{\mathrm{sym}}(H) - \frac{\delta}{2} q_{\mathrm{sym}}(H).$$

For every matrix reached by the algorithm, a separate comparison gives

$$q_{\mathrm{sym}}(H) \leq D_{\mathrm{sym}}(H) \leq C_D L q_{\mathrm{sym}}(H),$$

where C_D is a constant. Consequently,

$$D_{\mathrm{sym}}(H^+) \leq \left(1 - \frac{\delta}{2C_D L}\right) D_{\mathrm{sym}}(H) = (1 - \Theta(L^{-3})) D_{\mathrm{sym}}(H).$$

The initial value satisfies $\log(D_{\mathrm{sym}}(H_0)/c_*) = O(L)$ for every fixed terminal threshold $c_* > 0$. Thus $O(L^4)$ successful updates make $D_{\mathrm{sym}}(H)$ at most c_* .

At this point the uniform measure on K is well-rounded in the current coordinates. The terminal comparison shows that C_H and C_K are comparable in relative Loewner order, and removing the nonnegative quadratic term preserves the fixed-radius sublevel-set ball. One final covariance call on K gives the matrix Σ in [Theorem 2](#). The algorithm makes $O(L^4)$ calls, each costing $\tilde{O}(n^3)$ membership queries, so its overall query complexity is cubic up to logarithmic factors.

5.1 Proof Setup for Even Log-Concave Laws

We now prove the symmetric upper bound stated in [Theorem 2](#). Unlike the preceding scalar schedule, the algorithm cools a matrix-valued Gaussian precision and adds an adaptive regularizer at each step. Symmetry keeps every distribution with density proportional to $e^{-V(x) - x^\top H x}$ centered. The remaining subsections give the analytic inputs, invariant, progress argument, and final algorithm.

5.2 Quadratic Weights and Analytic Inputs

The following covariance estimator is used at each cooling step. Given an even log-concave distribution satisfying the two displayed well-roundedness conditions, it returns a 1% relative covariance estimate using $O(n^3 \log(n/\xi)^{O(1)})$ evaluations. Its success guarantee remains valid after conditioning on the preceding calls, and its query count has a deterministic upper bound. We apply the estimator to distributions with density proportional to $e^{-V(x)-x^\top Hx}$ for $H \succeq 0$. The next theorem concerns the near-cubic covariance estimation for even distributions.

Theorem 11. *Fix universal $b_0, c_1, c_2 > 0$ and $\kappa_{\text{app}} = 1.01$. Let $U : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ be a proper, lower-semicontinuous, convex, and even potential whose normalized law ν_U satisfies*

$$\mathbb{E}_{\nu_U} [\|Z\|^2] \leq c_1 n, \quad B(0, b_0) \subseteq \{U \leq \inf U + c_2 n\}.$$

Given an evaluation oracle for U and a failure tolerance $\xi \in (0, 1/2]$, there is an algorithm that returns either a failure symbol or a positive-definite matrix $\hat{C}$, and with probability at least $1 - \xi$ the latter satisfies

$$\hat{C} \approx_{\kappa_{\text{app}}} \text{Cov}_{\nu_U}(Z).$$

It has a deterministic query limit of

$$O(n^3 \log(n/\xi)^{O(1)})$$

evaluations of U , and its guarantee holds under adaptive conditioning in the sense of [Definition 2.20](#).

Proof. Set

$$p = \min\{\xi, 1/20\}, \quad \varepsilon_{\text{prim}} = 1/500.$$

FULLDENSITYROUND ([Lemma 4.5](#)) takes a ball contained in a height- $O(n)$ sublevel set of the potential together with an $O(n)$ second-moment bound and returns relative estimates of the mean and covariance, with a conditional failure guarantee and a deterministic query limit. By [Lemma 4.5](#), FULLDENSITYROUND implements this interface. Apply it with center 0 and $(\rho, H, B) = (b_0, c_2, c_1)$, accuracy $\varepsilon_{\text{prim}}$, and failure tolerance p . If the routine returns a failure symbol, pass that symbol to the caller. Otherwise discard its mean estimate and return its covariance estimate. Evenness makes the true mean zero. Since $p \leq \xi$ and

$$1 - \varepsilon_{\text{prim}} \geq \kappa_{\text{app}}^{-1}, \quad 1 + \varepsilon_{\text{prim}} \leq \kappa_{\text{app}},$$

the routine supplies the stated covariance guarantee, conditional failure bound, and deterministic cap. Replacing ξ by p inside the logarithm changes the cap only by a universal factor. $\square$

The next definition concerns the well-rounded even potential.

Definition 5.1. Let $U : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ be proper, lower-semicontinuous, even, and convex, and suppose $0 < \int_{\mathbb{R}^n} e^{-U(z)} dz < \infty$. For positive constants b_0, c_1, c_2 , call U (b_0, c_1, c_2) -well-rounded if its normalized law ν_U satisfies

$$\mathbb{E}_{\nu_U} [\|Z\|^2] \leq c_1 n \quad \text{and} \quad B(0, b_0) \subseteq \{z : U(z) \leq \inf U + c_2 n\}.$$

For a base potential P and $J \succeq 0$ for which the denominator below is positive and finite, we use the shorthand

$$\nu_{P,J}(dz) = \frac{e^{-P(z)-z^\top Jz} dz}{\int_{\mathbb{R}^n} e^{-P(y)-y^\top Jy} dy}.$$

Here “well-rounded” is a property of the probability distribution defined by a potential, not a matrix property. We call $\nu_{P,J}$ well-rounded when the potential $U(z) = P(z) + z^\top Jz$ satisfies the two preceding conditions: its normalized distribution has second moment at most $c_1 n$, and $B(0, b_0) \subseteq \{U \leq \inf U + c_2 n\}$.

The next lemma supplies the whole-space Poincaré inequality from which the normalized and singular-precision inequalities, and hence the energy bounds used in the cooling analysis, will follow.

The next lemma concerns the whole-space strongly convex even-sector inequality.

Lemma 5.2. *Let $V \in C^\infty(\mathbb{R}^n)$ be even, suppose $\nabla^2 V(y) \succeq 2I$ for every y , and suppose $Z_V = \int_{\mathbb{R}^n} e^{-V(y)} dy < \infty$. For $d\mu_V = Z_V^{-1} e^{-V(y)} dy$ and every $h \in C_c^\infty(\mathbb{R}^n)$,*

$$\text{Var}_{\mu_V} [h] \leq \frac{1}{2} \int_{\mathbb{R}^n} \|\nabla h\|_2^2 d\mu_V. \quad (5.1)$$

If h is even, the factor $1/2$ may be replaced by $1/4$.

Proof. The Brascamp–Lieb inequality, [Lemma 2.8](#), and $\nabla^2 V \succeq 2I$ give

$$\text{Var}_{\mu_V} [h] \leq \int \langle (\nabla^2 V)^{-1} \nabla h, \nabla h \rangle d\mu_V \leq \frac{1}{2} \int \|\nabla h\|_2^2 d\mu_V.$$

This proves (5.1), without using that h is even. It remains to improve the constant when h is even.

Suppose now that h is even, and let $\gamma(dy) = \pi^{-n/2} e^{-\|y\|_2^2} dy$. Caffarelli’s contraction theorem, [Theorem 3](#), gives a transport map T from γ to μ_V that does not increase distances:

$$T_\# \gamma = \mu_V, \quad \|DT(y)\|_{\text{op}} \leq 1 \quad \text{for almost every } y.$$

Here the hypothesis of Caffarelli’s theorem holds because the source potential $\|y\|_2^2$ has Hessian $2I$, while $\nabla^2 V \succeq 2I$. Moreover, T is odd. Indeed, both γ and μ_V are unchanged by $y \mapsto -y$. Writing $T = \nabla \varphi$ as in Caffarelli’s theorem, the reflected map $y \mapsto -T(-y)$ equals $\nabla(\varphi(-y))$, and $\varphi(-y)$ is convex. This reflected map also sends γ to μ_V , so the uniqueness assertion in Caffarelli’s theorem gives $T(-y) = -T(y)$ almost everywhere.

We recall the elementary Gaussian fact that will be used. Hermite polynomials play the same role for Gaussian measure that Fourier characters play on the discrete cube: they give an orthogonal decomposition by degree. If g_k denotes the part of g of total Hermite degree k , then, for the normalization of γ above,

$$\text{Var}_\gamma [g] = \sum_{k \geq 1} \|g_k\|_{L_2(\gamma)}^2, \quad \int \|\nabla g\|_2^2 d\gamma = \sum_{k \geq 1} 2k \|g_k\|_{L_2(\gamma)}^2.$$

Under $y \mapsto -y$, the degree- k part changes by the sign $(-1)^k$. Thus, if g is even, all of its odd-degree parts vanish. Every nonconstant term in the sums above then has $k \geq 2$, and therefore

$$\text{Var}_\gamma [g] \leq \frac{1}{4} \int \|\nabla g\|_2^2 d\gamma. \quad (5.2)$$

Apply (5.2) to $g = h \circ T$. Because h is even and T is odd, g is even. The chain rule and the contraction bound give

$$\|\nabla(h \circ T)(y)\|_2 \leq \|\nabla h(T(y))\|_2 \quad \text{for almost every } y.$$

Finally, $T_\# \gamma = \mu_V$, so

$$\text{Var}_{\mu_V} [h] = \text{Var}_\gamma [h \circ T] \leq \frac{1}{4} \int \|\nabla(h \circ T)\|_2^2 d\gamma \leq \frac{1}{4} \int \|\nabla h\|_2^2 d\mu_V.$$

□

The next lemma concerns the normalized even-sector Poincaré inequality.

Lemma 5.3. *Let $P : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ be proper, lower-semicontinuous, even, and convex, and suppose*

$$0 < Z_P := \int_{\mathbb{R}^n} e^{-P(z) - \|z\|_2^2} dz < \infty.$$

Define $d\gamma_P(y) = Z_P^{-1} e^{-P(y) - \|y\|_2^2} dy$. For every locally Lipschitz h ,

$$\text{Var}_{\gamma_P}[h] \leq \frac{1}{2} \mathbb{E}_{\gamma_P} [\|\nabla h\|_2^2]. \quad (5.3)$$

If h is even almost everywhere, then the factor $1/2$ in (5.3) may be replaced by $1/4$:

$$\text{Var}_{\gamma_P}[h] \leq \frac{1}{4} \mathbb{E}_{\gamma_P} [\|\nabla h\|_2^2]. \quad (5.4)$$

The next lemma concerns the Poincaré inequality for possibly singular precision.

Lemma 5.4. *Let $P : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ be proper, lower-semicontinuous, even, and convex, let $H \succeq 0$, and let*

$$0 < \int_{\mathbb{R}^n} e^{-P(x) - x^\top H x} dx < \infty.$$

Let

$$d\nu(x) = \frac{e^{-P(x) - x^\top H x} dx}{\int_{\mathbb{R}^n} e^{-P(z) - z^\top H z} dz}.$$

If f is locally Lipschitz and $\nabla f(x) \in \text{im}(H)$ almost everywhere, then

$$\text{Var}_\nu[f] \leq \frac{1}{2} \mathbb{E}_\nu [\nabla f(X)^\top H^\dagger \nabla f(X)]. \quad (5.5)$$

If f is even, the factor $1/2$ may be replaced by $1/4$.

Proof of Lemma 5.3. After adding a constant, assume $P(0) = \inf P = 0$. Approximate P pointwise by finite, smooth, even, convex functions $P_j \geq 0$; thus

$$P_j(x) \rightarrow P(x) \quad \text{for every } x \in \mathbb{R}^n.$$

Set $Z_j = \int e^{-P_j(y) - \|y\|^2} dy$ and let γ_j be the corresponding normalized measure. Since

$$0 \leq e^{-P_j(y) - \|y\|^2} \leq e^{-\|y\|^2},$$

dominated convergence gives $Z_j \rightarrow Z$ and, for every $h \in C_c^\infty$,

$$\int h^k d\gamma_j \rightarrow \int h^k d\gamma_P \quad (k = 1, 2), \quad \int \|\nabla h\|^2 d\gamma_j \rightarrow \int \|\nabla h\|^2 d\gamma_P.$$

Apply Lemma 5.2 to $V_j = P_j + \|\cdot\|^2$. Passing to the limit proves both constants for C_c^∞ functions. Now let h be compactly supported and Lipschitz. Choose a smooth, even, nonnegative function ρ with compact support and $\int \rho = 1$, define

$$\rho_\varepsilon(x) = \varepsilon^{-n} \rho(x/\varepsilon),$$

and let

$$h_\varepsilon = h * \rho_\varepsilon.$$

Then $h_\varepsilon \in C_c^\infty$, and as $\varepsilon \downarrow 0$,

$$h_\varepsilon \rightarrow h, \quad \nabla h_\varepsilon \rightarrow \nabla h \quad \text{almost everywhere.}$$

For $\varepsilon \leq 1$, these functions are supported in a fixed compact set, and $\|\nabla h_\varepsilon\|$ is bounded by the Lipschitz constant of h . Dominated convergence therefore gives convergence of the variances and of the gradient integrals. If h is even, then h_ε is even because the kernel ρ_ε is even. Applying the smooth inequality to h_ε and sending $\varepsilon \downarrow 0$ proves the claim for compactly supported Lipschitz functions.

For a locally Lipschitz h , let $T_M(t) = \max\{-M, \min\{t, M\}\}$, and choose a radial $\chi_R \in C_c^\infty$ such that $\chi_R = 1$ on B_R , $\chi_R = 0$ outside B_{2R} , and $\|\nabla \chi_R\| \leq C/R$. Define

$$h_{R,M} = \chi_R T_M(h), \quad \nabla h_{R,M} = \chi_R \mathbf{1}_{\{|h| < M\}} \nabla h + T_M(h) \nabla \chi_R \quad \text{a.e.}$$

The radial choice makes $h_{R,M}$ even whenever h is even. Moreover,

$$\|\nabla h_{R,M}\| \leq \mathbf{1}_{\{|h| < M\}} \|\nabla h\| + \frac{CM}{R} \mathbf{1}_{B_{2R} \setminus B_R}.$$

Suppose first that

$$\mathbb{E}_{\gamma_P} [\|\nabla h\|^2] < \infty.$$

For fixed M , apply the compactly supported inequality to $h_{R,M}$ and let $R \rightarrow \infty$. The preceding gradient bound and dominated convergence give

$$\text{Var}_{\gamma_P} [T_M(h)] \leq c \mathbb{E}_{\gamma_P} [\mathbf{1}_{\{|h| < M\}} \|\nabla h\|^2],$$

where $c = 1/2$ in general and $c = 1/4$ for even h . If X, X' are independent with law γ_P , then Fatou's lemma gives

$$\text{Var}_{\gamma_P} [h] = \frac{1}{2} \mathbb{E}_{X, X' \sim \gamma_P \otimes \gamma_P} [(h(X) - h(X'))^2] \leq \liminf_{M \rightarrow \infty} \text{Var}_{\gamma_P} [T_M(h)].$$

Monotone convergence on the right-hand side proves the claim. If the gradient integral is infinite, the claimed inequality is immediate. $\square$

Proof of Lemma 5.4. First suppose $H \succ 0$. Set

$$y = H^{1/2}x, \quad g(y) = f(H^{-1/2}y), \quad \tilde{P}(y) = P(H^{-1/2}y).$$

The law of $Y = H^{1/2}X$ has density proportional to $e^{-\tilde{P}(y) - \|y\|^2}$, so Lemma 5.3 applies. Moreover,

$$\nabla_y g(y) = H^{-1/2} \nabla_x f(x), \quad \|\nabla_y g(y)\|^2 = \nabla_x f(x)^\top H^{-1} \nabla_x f(x).$$

Therefore, with $c = 1/2$ in general and $c = 1/4$ when f is even,

$$\text{Var}_\nu [f] = \text{Var}_{\gamma_{\tilde{P}}} [g(Y)] \leq c \mathbb{E}_{\gamma_{\tilde{P}}} [\|\nabla g(Y)\|^2] = c \mathbb{E}_\nu [\nabla f(X)^\top H^{-1} \nabla f(X)].$$

Since $H \succ 0$ implies $H^\dagger = H^{-1}$, this proves both constants for positive-definite H .

Now suppose H is singular and let P_0 be the orthogonal projector onto $\ker H$. Define $H_\varepsilon = H + \varepsilon P_0$. Then $H_\varepsilon \succ 0$ and

$$H_\varepsilon^{-1} = H^\dagger + \varepsilon^{-1} P_0.$$

The densities with precision H_ε converge to the density of ν in total variation, because

$$0 \leq e^{-P(x)-x^\top H_\varepsilon x} \leq e^{-P(x)-x^\top H x}$$

and the right-hand side is integrable. Since $\nabla f \in \text{im}(H) = \ker(P_0)$ almost everywhere,

$$\nabla f^\top H_\varepsilon^{-1} \nabla f = \nabla f^\top H^\dagger \nabla f.$$

For bounded, compactly supported Lipschitz f , the positive-definite result and total-variation convergence give

$$\text{Var}_\nu[f] \leq c \mathbb{E}_\nu \left[\nabla f^\top H^\dagger \nabla f \right].$$

For the stated locally Lipschitz class, write $\mathbb{R}^n = \text{im}(H) \oplus \ker(H)$, and let P_H be the orthogonal projector onto $\text{im}(H)$. The gradient condition implies that f is constant along $\ker(H)$ -directions. With the truncation T_M and radial cutoff χ_R used above, define

$$f_{R,M}(x) = \chi_R(P_H x) T_M(f(x)).$$

Then $f_{R,M}$ and its gradient are bounded, $\nabla f_{R,M} \in \text{im}(H)$ almost everywhere, and the positive-definite inequality passes to the limit $\varepsilon \downarrow 0$ by total-variation convergence. Letting first $R \rightarrow \infty$ and then $M \rightarrow \infty$, exactly as in the preceding proof, gives the stated inequality. If $H = 0$, the gradient condition makes f constant, so the conclusion is immediate. $\square$

The next lemma concerns the energy domination and cooling monotonicity.

Lemma 5.5. *Let $V : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ be proper, lower-semicontinuous, even, and convex, suppose $0 < \int_{\mathbb{R}^n} e^{-V(x)} dx < \infty$, let $B, D \succeq 0$, and let $a > 0$. Define*

$$A_{aB+D} = \frac{\int_{\mathbb{R}^n} x x^\top e^{-V(x)-x^\top (aB+D)x} dx}{\int_{\mathbb{R}^n} e^{-V(x)-x^\top (aB+D)x} dx}.$$

Let P_B be the orthogonal projector onto $\text{im}(B)$. Then

$$aB^{1/2} A_{aB+D} B^{1/2} \preceq \frac{1}{2} P_B, \quad a \text{Tr}(B A_{aB+D}) \leq \frac{1}{2} \text{rank}(B), \quad (5.6)$$

and the function $a \mapsto a \text{Tr}(B A_{aB+D})$ is nondecreasing on $(0, \infty)$.

Proof. The distribution in the lemma is even, so its covariance is A_{aB+D} . Set $H = aB + D$. For $v \in \mathbb{R}^n$, let $u = B^{1/2}v$. The linear function $x \mapsto u^\top x$ is generally not even, so we apply the ordinary factor 1/2 in Lemma 5.4; also $u \in \text{im}(B) \subseteq \text{im}(H)$. This inclusion verifies the gradient condition in that lemma. We obtain

$$v^\top B^{1/2} A_{aB+D} B^{1/2} v \leq \frac{1}{2} v^\top B^{1/2} H^\dagger B^{1/2} v \leq \frac{1}{2a} v^\top P_B v.$$

Indeed, $(H^\dagger)^{1/2} B (H^\dagger)^{1/2} \preceq a^{-1} I$; hence $B^{1/2} H^\dagger B^{1/2}$ has operator norm at most a^{-1} and is supported on $\text{im}(B)$, which proves the last inequality. Thus the first matrix inequality in (5.6) follows, and taking traces gives the second. If $e(a) = \text{Tr}(B A_{aB+D})$, differentiation gives

$$\frac{d}{da}(ae(a)) = e(a) - a \text{Var}_{\nu_a} [X^\top B X].$$

The function $q(x) = x^\top Bx$ is even and $\nabla q(x) = 2Bx \in \text{im}(H)$. The even factor $1/4$ in [Lemma 5.4](#) therefore gives

$$\text{Var}_{\nu_a}[q(X)] \leq \frac{1}{4} \mathbb{E}_{\nu_a} [4X^\top BH^\top BX] = \mathbb{E}_{\nu_a} [X^\top BH^\top BX] \leq \frac{1}{a} \mathbb{E}_{\nu_a} [X^\top BX] = \frac{e(a)}{a}.$$

Substituting this bound into the derivative identity gives

$$\frac{d}{da}(ae(a)) \geq e(a) - a \frac{e(a)}{a} = 0.$$

It remains to justify differentiation when V is nonsmooth or may take the value $+\infty$. Let $q(x) = x^\top Bx$, let ν_a denote the normalized distribution with density proportional to $e^{-V(x) - x^\top Dx - aq(x)}$, and define

$$Z(a) = \int e^{-V(x) - x^\top Dx - aq(x)} dx, \quad N(a) = \int q(x) e^{-V(x) - x^\top Dx - aq(x)} dx.$$

The integrable log-concave density proportional to $e^{-V(x) - x^\top Dx}$ has finite polynomial moments. Hence q and q^2 are integrable against this density, and differentiation under the integral gives

$$Z'(a) = -N(a), \quad N'(a) = - \int q(x)^2 e^{-V(x) - x^\top Dx - aq(x)} dx.$$

Since $e(a) = N(a)/Z(a)$, the quotient rule yields

$$e'(a) = -\text{Var}_{\nu_a}[q(X)].$$

The Poincaré inequality used above already allows proper lower-semicontinuous, extended-valued V , so the same calculation proves the result without a smoothness assumption. $\square$

The next lemma concerns the even log-concave marginals.

Lemma 5.6. *Let $f : \mathbb{R}^n \rightarrow [0, \infty)$ be integrable, even, and log-concave in the extended-valued sense. For a unit vector u , define*

$$g(t) = \int_{u^\perp} f(tu + y) dy,$$

where in dimension $n = 1$ the integral over the zero-dimensional space is evaluation at its unique point. Then g is even and log-concave. In particular, $g(t) \leq g(0)$ for every $t \in \mathbb{R}$.

Proof. Evenness follows from the change of variables $y \mapsto -y$ in $u^\perp$, and log-concavity follows from [Lemma 2.10](#). An even log-concave function on $\mathbb{R}$ attains its maximum at zero. $\square$

The next lemma concerns the grounded central-section comparison.

Lemma 5.7. *Let V be proper, lower-semicontinuous, even, and convex, and suppose*

$$0 < \int_{\mathbb{R}^n} e^{-V(x)} dx < \infty, \quad B(0, r) \subseteq \{V \leq V(0) + h_0 n\}.$$

Let $H \succeq 0$, let $m = \|H\|_{\text{op}}$, fix a unit vector u , and define

$$g(t) = \int_{u^\perp} e^{-V(tu+y) - (tu+y)^\top H(tu+y)} dy.$$

For every $|t| < r$,

$$g(t) \geq \left(1 - \frac{|t|}{r}\right)^{n-1} \exp\left(-h_0 n \frac{|t|}{r} - mt^2\right) g(0). \quad (5.7)$$

The factor with exponent zero is interpreted as 1 when $n = 1$.

Proof. Evenness and convexity imply $V(0) = \inf V$. Write $\alpha = |t|/r$ and $\sigma = \text{sgn}(t)$, with an arbitrary choice of σ when $t = 0$. For $y \in u^\perp$, set $x_\pm = \pm tu + (1 - \alpha)y$. Convexity, the points $\pm \sigma ru$ in the promised sublevel-set ball, and $V(0) \leq V(y)$ give

$$V(\pm \sigma ru) \leq V(0) + h_0 n \leq V(y) + h_0 n.$$

Since

$$x_+ = (1 - \alpha)y + \alpha(\sigma ru), \quad x_- = (1 - \alpha)y + \alpha(-\sigma ru),$$

convexity yields

$$\frac{V(x_+) + V(x_-)}{2} \leq (1 - \alpha)V(y) + \frac{\alpha}{2}(V(\sigma ru) + V(-\sigma ru)) \leq V(y) + \alpha h_0 n.$$

For the quadratic terms,

$$\begin{aligned} x_+^\top H x_+ &= (1 - \alpha)^2 y^\top H y + 2t(1 - \alpha)u^\top H y + t^2 u^\top H u, \\ x_-^\top H x_- &= (1 - \alpha)^2 y^\top H y - 2t(1 - \alpha)u^\top H y + t^2 u^\top H u. \end{aligned}$$

The cross terms cancel upon averaging. Since $(1 - \alpha)^2 \leq 1$ and $u^\top H u \leq \|H\|_{\text{op}} = m$,

$$\frac{x_+^\top H x_+ + x_-^\top H x_-}{2} = (1 - \alpha)^2 y^\top H y + t^2 u^\top H u \leq y^\top H y + m t^2.$$

Define

$$F_\pm(y) = \exp(-V(x_\pm) - x_\pm^\top H x_\pm).$$

Changing transverse variables by $z = (1 - \alpha)y$ gives

$$g(\pm t) = (1 - \alpha)^{n-1} \int_{u^\perp} F_\pm(y) dy.$$

Evenness gives $g(t) = g(-t)$, and the arithmetic–geometric mean inequality therefore gives

$$\begin{aligned} g(t) &= \frac{g(t) + g(-t)}{2} \\ &\geq (1 - \alpha)^{n-1} \int_{u^\perp} (F_+(y) F_-(y))^{1/2} dy \\ &\geq (1 - \alpha)^{n-1} e^{-\alpha h_0 n - m t^2} \int_{u^\perp} e^{-V(y) - y^\top H y} dy. \end{aligned}$$

The last integral is $g(0)$, proving the claim. $\square$

The next lemma concerns the grounded directional second moment.

Lemma 5.8. *Under the assumptions of Lemma 5.7, let C_H be the second moment of the law proportional to $e^{-V(x) - x^\top H x}$. There is a constant $c_{h_0} > 0$, depending only on h_0 , such that for every unit vector u ,*

$$u^\top C_H u \geq \frac{c_{h_0}}{\|H\|_{\text{op}} + n^2/r^2}. \quad (5.8)$$

Proof. The marginal density in direction u is proportional to the function g in [Lemma 5.7](#). Its maximum is at zero by [Lemma 5.6](#). Set

$$a = c'_{h_0} (\|H\|_{\text{op}} + n^2/r^2)^{-1/2},$$

where $c'_{h_0} > 0$ is sufficiently small. Then $a \leq c'_{h_0} r/n$ and, when $\|H\|_{\text{op}} > 0$, $a \leq c'_{h_0} \|H\|_{\text{op}}^{-1/2}$. For $|t| \leq a$, write $\alpha = |t|/r$. Choosing $c'_{h_0} \leq 1/2$ gives

$$(1 - \alpha)^{n-1} \geq e^{-2n\alpha} \geq e^{-2c'_{h_0}}, \quad h_0 n \alpha \leq h_0 c'_{h_0}, \quad \|H\|_{\text{op}} t^2 \leq (c'_{h_0})^2.$$

Thus Equation (5.7) gives $g(t) \geq c''_{h_0} g(0)$ on $[-a, a]$. If p is the normalized marginal density, then

$$1 \geq \int_{-a}^a p(t) dt \geq 2ac''_{h_0} p(0), \quad p(0) \leq \frac{1}{2c''_{h_0} a}.$$

For completeness, any probability density bounded by M satisfies $\int t^2 p(t) dt \geq 1/(12M^2)$. Indeed, for $s \geq 0$,

$$\mathbb{P}(|T| > s) \geq 1 - \int_{-s}^s p(t) dt \geq (1 - 2Ms)_+.$$

Consequently,

$$\int t^2 p(t) dt = \int_0^\infty 2s \mathbb{P}(|T| > s) ds \geq \int_0^{1/(2M)} 2s(1 - 2Ms) ds = \frac{1}{12M^2}.$$

Apply this with $M = p(0)$ to obtain

$$u^\top C_H u = \int t^2 p(t) dt \geq c_{h_0} a^2 \geq \frac{c_{h_0}}{\|H\|_{\text{op}} + n^2/r^2},$$

as claimed. □

The next lemma concerns the inverse covariance versus precision.

Lemma 5.9. *Under the grounded assumptions above, for every $H \succeq 0$,*

$$\|C_H^{-1}\|_{\text{op}} \leq C_{h_0} \left(\|H\|_{\text{op}} + \frac{n^2}{r^2} \right),$$

where $C_{h_0} = c_{h_0}^{-1}$ depends only on h_0 .

Proof. [Lemma 5.8](#) gives $C_H \succeq c_{h_0} (\|H\|_{\text{op}} + n^2/r^2)^{-1} I$. Inverting this positive-definite matrix inequality proves the claim. □

The next lemma concerns the covariance-scale sublevel-set ball.

Lemma 5.10. *Let U be an even convex potential with normalized law of covariance C . If $W \approx_{\kappa_{\text{app}}} C$, then for universal $b_g, H_g > 0$,*

$$B(0, b_g) \subseteq \{z : U(W^{1/2}z) \leq \inf U + H_g n\}.$$

Proof. The transformed law is centered and has covariance between $\kappa_{\text{app}}^{-1} I$ and $\kappa_{\text{app}} I$. Apply item 4 of [Proposition 2.23](#), the covariance-to-level-set implication, with a universal rescaling to absorb $\kappa_{\text{app}} = 1.01$. □

Standing parameters. Fix $\kappa_{\text{app}} = 1.01$. Choose C_L sufficiently large, depending only on h_0 and the universal constants below, and throughout the symmetric argument set

$$L_{\text{sym}} = C_L \max\{1, \log(2 + nR/r)\}, \quad \delta_{\text{sym}} = L_{\text{sym}}^{-2}, \quad \eta_{\text{ridge}} = L_{\text{sym}}^{-3}.$$

The next definition concerns the quadratically weighted distributions.

Definition 5.11. Fix the even potential V from [Theorem 2](#). For $H \succeq 0$, define

$$Z_V(H) = \int_{\mathbb{R}^n} e^{-V(x) - x^\top H x} dx, \quad \nu_H(dx) = \frac{e^{-V(x) - x^\top H x}}{Z_V(H)} dx, \quad C_H = \mathbb{E}_{\nu_H} [XX^\top].$$

Also set

$$D_{\text{sym}}(H) = \log \frac{Z_V(0)}{Z_V(H)}, \quad q_{\text{sym}}(H) = \text{Tr}(C_H H).$$

Equivalently,

$$e^{-D_{\text{sym}}(H)} = \frac{Z_V(H)}{Z_V(0)} = \mathbb{E}_{\nu_0} [e^{-X^\top H X}], \quad q_{\text{sym}}(H) = \mathbb{E}_{\nu_H} [X^\top H X].$$

Thus $D_{\text{sym}}(H)$ records the reduction in the normalizing integral caused by the quadratic weight. We refer to ν_H as the quadratic tilt of ν_0 by H . Every ν_H is even, so C_H is both its covariance and its uncentered second moment.

Standing analytic estimate. The singular-precision Poincaré inequality in [Lemma 5.4](#) controls the variance of even scalar functions by their gradient energy in the image of H . Thus, for every $H \succeq 0$ and every even locally Lipschitz function g satisfying $\nabla g(z) \in \text{im}(H)$ almost everywhere, we use

$$\text{Var}_{\nu_H} [g] \leq \frac{1}{4} \mathbb{E}_{\nu_H} [\nabla g(Z)^\top H^\dagger \nabla g(Z)]. \quad (5.9)$$

The next definition concerns the trace and energy along the interpolation.

Definition 5.12. Suppose $B \succeq 0$ and $b = \text{Tr}(B)/n > 0$. Let $\gamma = \eta_{\text{ridge}} b$, and for $s \in [0, 1]$ define

$$a_s = 1 - \delta_{\text{sym}} s, \quad H_s = a_s B + s \gamma I, \quad F(s) = \text{Tr}(C_{H_s}), \quad q(s) = \text{Tr}(B C_{H_s}).$$

Here the covariance is computed with any fixed even convex base potential in the current coordinates. Also set $f(s) = F(s)/n$ and $u(s) = q(s)/(bn)$.

The next lemma concerns the coupled differential control.

Lemma 5.13. *In the setting of [Definition 5.12](#), suppose $0 < \delta_{\text{sym}} < 1$ and use the even-sector Poincaré estimate (5.9). For every $s \in (0, 1]$,*

$$f'(s) \leq \frac{\delta_{\text{sym}}^2}{4a_s \eta_{\text{ridge}}} u(s), \quad (5.10)$$

$$u'(s) \leq \frac{\delta_{\text{sym}}}{a_s} u(s) + \left(\frac{\eta_{\text{ridge}} u(s) f(s)}{a_s s} \right)^{1/2}. \quad (5.11)$$

Moreover, these inequalities may be integrated down to $s = 0$.

Proof. For brevity write

$$\delta = \delta_{\text{sym}}, \quad \eta = \eta_{\text{ridge}}, \quad C_s = C_{H_s},$$

and let $G(x) = \|x\|^2$ and $K(x) = x^\top Bx$. Differentiating the expectations with respect to s uses

$$H'_s = -\delta B + \gamma I, \quad \frac{d}{ds} \mathbb{E}_s[\varphi(X)] = -\text{Cov}_s(\varphi(X), X^\top H'_s X).$$

Taking $\varphi = G$ and $\varphi = K$ gives

$$F'(s) = \delta \text{Cov}(G, K) - \gamma \text{Var}_s[G], \quad q'(s) = \delta \text{Var}_s[K] - \gamma \text{Cov}(K, G).$$

For $s > 0$, applying the even-sector variance bound (5.9) to K and G , and using $H_s \succeq a_s B$ and $H_s \succeq s\gamma I$, gives

$$\begin{aligned} \text{Var}_s[K] &\leq \frac{1}{4} \mathbb{E}_s \left[(2BX)^\top H_s^{-1} (2BX) \right] \\ &= \text{Tr}(C_s B H_s^{-1} B) \leq \frac{q(s)}{a_s}, \\ \text{Var}_s[G] &\leq \frac{1}{4} \mathbb{E}_s \left[(2X)^\top H_s^{-1} (2X) \right] \\ &= \text{Tr}(C_s H_s^{-1}) \leq \frac{F(s)}{s\gamma}. \end{aligned}$$

Indeed, restricted to $\text{im}(B)$, the inequality $H_s \succeq a_s B$ gives $H_s^{-1} \preceq a_s^{-1} B^\dagger$. Therefore

$$B H_s^{-1} B \preceq a_s^{-1} B B^\dagger B = a_s^{-1} B,$$

which is the bound used above. The scalar Cauchy–Schwarz inequality

$$|\text{Cov}(U, V)| \leq (\text{Var}_s[U] \text{Var}_s[V])^{1/2},$$

and the inequality

$$\delta xy - \gamma x^2 \leq \frac{\delta^2 y^2}{4\gamma}$$

imply

$$\begin{aligned} F'(s) &\leq \delta (\text{Var}_s[G] \text{Var}_s[K])^{1/2} - \gamma \text{Var}_s[G] \\ &\leq \frac{\delta^2}{4\gamma} \text{Var}_s[K] \leq \frac{\delta^2}{4a_s\gamma} q(s). \end{aligned}$$

Since $\gamma = \eta b$, $F = nf$, and $q = bnu$, division by n proves Equation (5.10). Similarly,

$$\begin{aligned} q'(s) &\leq \delta \text{Var}_s[K] + \gamma (\text{Var}_s[K] \text{Var}_s[G])^{1/2} \\ &\leq \frac{\delta}{a_s} q(s) + \left(\frac{\gamma q(s) F(s)}{a_s s} \right)^{1/2}. \end{aligned}$$

Dividing by bn proves Equation (5.11).

Finally, let P denote the fixed base potential in Definition 5.12. Since $H_s \succeq 0$,

$$0 \leq (1 + \|x\|^2) e^{-P(x) - x^\top H_s x} \leq (1 + \|x\|^2) e^{-P(x)}.$$

The function on the right is integrable: after an affine normalization, this follows from the log-concave tail bound in Lemma 2.6. Dominated convergence, applied to the normalizing integral and

the second-moment integrals, shows that f and u extend continuously to $s = 0$. In particular, for some $C < \infty$ and all sufficiently small s ,

$$u(s) \leq C, \quad f(s) \leq C, \quad a_s \geq 1 - \delta_{\text{sym}} > 0.$$

The right-hand side of Equation (5.10) is then bounded, and that of Equation (5.11) is $O(1 + s^{-1/2})$, which is integrable at zero. Integrating over $[\varepsilon, t]$ and letting $\varepsilon \downarrow 0$ proves the final assertion. $\square$

The next lemma concerns the trace bound after one update.

Lemma 5.14. *Let P be a fixed proper, lower-semicontinuous, even convex base potential with $0 < \int_{\mathbb{R}^n} e^{-P(z)} dz < \infty$, and let C_H denote covariance under the law proportional to $e^{-P - z^\top H z}$. Fix $\kappa_{\text{app}} = 1.01$. Suppose $B \succeq 0$ and*

$$\kappa_{\text{app}}^{-1} I \preceq C_B \preceq \kappa_{\text{app}} I, \quad \|B\|_{\text{op}} \leq \kappa_{\text{app}}/2, \quad \text{Tr}(B) \leq \kappa_{\text{app}} n/2.$$

Let $\delta_{\text{sym}} = L_{\text{sym}}^{-2}$ and $\eta_{\text{ridge}} = L_{\text{sym}}^{-3}$, and define

$$J_{\text{sym}} = (1 - \delta_{\text{sym}})B + \eta_{\text{ridge}} \frac{\text{Tr}(B)}{n} I.$$

There is a universal $L_{\text{sym},0}$ such that, for every $L_{\text{sym}} \geq L_{\text{sym},0}$,

$$\text{Tr}(C_{J_{\text{sym}}}) \leq 2\kappa_{\text{app}} n.$$

Proof. Write

$$\kappa = \kappa_{\text{app}}, \quad L = L_{\text{sym}}, \quad \delta = L^{-2}, \quad \eta = L^{-3}.$$

If $\text{Tr}(B) = 0$, then $B = J_{\text{sym}} = 0$, and the conclusion follows from $C_B \preceq \kappa I$. Otherwise set $b = \text{Tr}(B)/n > 0$ and use the interpolation from Definition 5.12, which parametrizes the path from B to J_{sym} . Then $H_0 = B$, $H_1 = J_{\text{sym}}$, and

$$f(0) = \frac{\text{Tr}(C_B)}{n} \leq \kappa, \quad u(0) = \frac{\text{Tr}(BC_B)}{\text{Tr}(B)} \leq \kappa.$$

Define $f_* = \sup_{0 \leq s \leq 1} f(s)$ and $v(s) = u(s)^{1/2}$. Since $B \neq 0$ and $C_{H_s} \succ 0$,

$$q(s) = \text{Tr}(BC_{H_s}) > 0, \quad v(s) > 0 \quad (0 \leq s \leq 1).$$

The second inequality in Lemma 5.13 therefore implies

$$v'(s) \leq \frac{\delta}{2a_s} v(s) + \frac{1}{2} \left(\frac{\eta f_*}{a_s s} \right)^{1/2}.$$

Since $a_s = 1 - \delta s$,

$$(a_s^{1/2} v(s))' \leq \frac{1}{2} \left(\frac{\eta f_*}{s} \right)^{1/2}.$$

Consequently

$$a_s^{1/2} v(s) \leq v(0) + (\eta f_* s)^{1/2} \leq \kappa^{1/2} + (\eta f_*)^{1/2}.$$

For all sufficiently large L , $a_s \geq 1/2$, and hence

$$u(s) \leq 4\kappa + 4\eta f_* \quad (0 \leq s \leq 1).$$

The first differential inequality, which controls the trace growth along this path, now gives

$$f'(s) \leq \frac{\delta^2}{4a_s\eta} u(s) \leq \frac{1}{2L} u(s).$$

Hence, for every $s \in [0, 1]$,

$$\begin{aligned} f(s) &\leq f(0) + \frac{1}{2L} \int_0^s u(t) dt \\ &\leq \kappa + \frac{2\kappa}{L} + \frac{2}{L^4} f_*. \end{aligned}$$

Taking the supremum over s yields

$$f_* \leq \kappa + \frac{2\kappa}{L} + \frac{2}{L^4} f_*.$$

For every sufficiently large universal L , this implies $f_* \leq 2\kappa$. Since $H_1 = J_{\text{sym}}$, $\text{Tr}(C_{J_{\text{sym}}}) = nf(1) \leq 2\kappa_{\text{app}}n$. $\square$

The next lemma concerns the normalized sublevel set and quadratic penalty.

Lemma 5.15. *Let $H \succeq 0$, let $A = C_H$, and let $W \approx_{\kappa_{\text{app}}} A$. Define*

$$B = W^{1/2} H W^{1/2}, \quad P_{H,W}(z) = V(W^{1/2}z) + z^\top B z.$$

Then $P_{H,W}$ is a well-rounded even potential with universal constants, and

$$\|B\|_{\text{op}} \leq \kappa_{\text{app}}/2, \quad \text{Tr}(B) \leq \kappa_{\text{app}}n/2.$$

Proof. The covariance of the normalized distribution is $W^{-1/2} C_H W^{-1/2}$, which lies between $\kappa_{\text{app}}^{-1} I$ and $\kappa_{\text{app}} I$. Its second moment is therefore at most $\kappa_{\text{app}}n$, and [Lemma 5.10](#) supplies a universal ball in an $O(n)$ -height sublevel set of $P_{H,W}$. Apply [Lemma 5.5](#) with $B = H$, $D = 0$, and $a = 1$ to obtain

$$H^{1/2} C_H H^{1/2} \preceq \frac{1}{2} P_{\text{im}(H)}.$$

Since $W \preceq \kappa_{\text{app}} C_H$, a congruence transformation by $H^{1/2}$ gives

$$H^{1/2} W H^{1/2} \preceq \kappa_{\text{app}} H^{1/2} C_H H^{1/2} \preceq \frac{\kappa_{\text{app}}}{2} I.$$

The nonzero eigenvalues of $W^{1/2} H W^{1/2}$ and $H^{1/2} W H^{1/2}$ agree, so $\|B\|_{\text{op}} \leq \kappa_{\text{app}}/2$. Taking traces in the same inequality gives

$$\text{Tr}(B) = \text{Tr}(W H) \leq \kappa_{\text{app}} \text{Tr}(C_H H) \leq \frac{\kappa_{\text{app}}n}{2}.$$

$\square$

The next lemma concerns the initial quadratic tilt.

Lemma 5.16. *Write*

$$H_0 = r^{-2} I, \quad P_0(z) = V(rz) + \|z\|^2.$$

Then P_0 is a well-rounded even potential with constants depending only on h_0 . Consequently, a successful invocation of the covariance estimator in [Theorem 11](#) on P_0 returns $\hat{C}_0 \approx_{\kappa_{\text{app}}} \text{Cov}_{P_0}(Z)$, and

$$W_0 = r^2 \hat{C}_0 \approx_{\kappa_{\text{app}}} C_{H_0}.$$

Proof. Evenness gives $V(0) = \inf V$. For $\|z\| \leq 1$, the grounded promise gives

$$P_0(z) \leq V(0) + h_0 n + 1 \leq \inf P_0 + (h_0 + 1)n.$$

Thus the unit ball lies in an $O(n)$ -height sublevel set. Applying the energy-domination estimate [Lemma 5.5](#) with base potential $z \mapsto V(rz)$, $B = I$ (the identity matrix), $D = 0$, and $a = 1$ gives $\mathbb{E}_{P_0} [\|Z\|^2] \leq n/2$. Under $x = rz$, the covariance transforms as

$$\text{Cov}_{P_0}(Z) = r^{-2} C_{H_0}.$$

Thus the covariance estimate returned for P_0 satisfies the displayed relation $W_0 = r^2 \hat{C}_0 \approx_{\kappa_{\text{app}}} C_{H_0}$. $\square$

5.3 Regularized Update and Invariant

Using the distributions and analytic estimates from [Section 5.2](#), we define the matrix-regularized update and prove that every accepted state remains within the well-rounded regime required by the covariance estimator. The next lemma concerns the adaptive-ridge well-roundedness.

Lemma 5.17. *Let $H \succeq 0$, let $W \approx_{\kappa_{\text{app}}} C_H$, and define*

$$\hat{q} = \text{Tr}(W H), \quad H^+ = (1 - \delta_{\text{sym}})H + \eta_{\text{ridge}} \frac{\hat{q}}{n} W^{-1}.$$

In the coordinates $x = W^{1/2}z$, define

$$B = W^{1/2} H W^{1/2}, \quad J_{\text{sym}} = W^{1/2} H^+ W^{1/2},$$

and define the next transformed potential

$$P^+(z) = V(W^{1/2}z) + z^\top J_{\text{sym}} z.$$

Then P^+ is (b_0, c_1, c_2) -well-rounded for universal positive constants b_0, c_1, c_2 .

Proof. The covariance of the current distribution in z -coordinates is

$$C = W^{-1/2} C_H W^{-1/2}, \quad \kappa_{\text{app}}^{-1} I \preceq C \preceq \kappa_{\text{app}} I.$$

By [Lemma 5.15](#), which supplies its sublevel-set ball and quadratic-penalty bounds, the current transformed potential

$$P(z) = V(W^{1/2}z) + z^\top B z$$

has a universal ball in an $O(n)$ -height sublevel set, while

$$\|B\|_{\text{op}} \leq \kappa_{\text{app}}/2, \quad \text{Tr}(B) \leq \kappa_{\text{app}} n/2.$$

Moreover $\hat{q} = \text{Tr}(B)$ and

$$J_{\text{sym}} = (1 - \delta_{\text{sym}})B + \eta_{\text{ridge}} \frac{\text{Tr}(B)}{n} I.$$

If $\text{Tr}(B) = 0$, then $B = 0$, the next potential equals the current one, and the conclusion already follows from [Lemma 5.15](#). Now assume that $\text{Tr}(B) > 0$. The hypotheses of [Lemma 5.14](#), the one-step trace-growth bound for the adaptive ridge, hold with base potential $z \mapsto V(W^{1/2}z)$, so the law of P^+ has second moment at most $2\kappa_{\text{app}}n$.

Let $B(0, b_g)$ be the ball supplied for P by [Lemma 5.15](#), so that

$$P(z) - P(0) \leq H_g n \quad (\|z\| \leq b_g).$$

On this ball,

$$P^+(z) - P(z) = -\delta_{\text{sym}} z^\top B z + \eta_{\text{ridge}} \frac{\text{Tr}(B)}{n} \|z\|^2 \leq \frac{\kappa_{\text{app}} \eta_{\text{ridge}} b_g^2}{2}.$$

Evenness and convexity give $V(0) = \inf V$, and both quadratic matrices are positive semidefinite. Hence $P(0) = \inf P$ and $P^+(0) = \inf P^+$. Since $P^+(0) - P(0) = 0$, for $\|z\| \leq b_g$,

$$P^+(z) - \inf P^+ \leq H_g n + \frac{\kappa_{\text{app}} \eta_{\text{ridge}} b_g^2}{2} \leq H'_g n$$

for a universal H'_g . Together with the second-moment bound, this proves both well-roundedness clauses. $\square$

The pullback of a successful transformed covariance estimate is exact under congruence.

The next lemma concerns the invariant propagation.

Lemma 5.18. *In the setting of [Lemma 5.17](#), let $C^+ = W^{-1/2} C_H W^{-1/2}$. If a covariance call returns $\hat{C}^+ \approx_{\kappa_{\text{app}}} C^+$ and $W^+ = W^{1/2} \hat{C}^+ W^{1/2}$, then $W^+ \approx_{\kappa_{\text{app}}} C_{H^+}$.*

Proof. The covariance guarantee is

$$\kappa_{\text{app}}^{-1} C^+ \preceq \hat{C}^+ \preceq \kappa_{\text{app}} C^+.$$

Multiplying all three matrices on the left and right by $W^{1/2}$ preserves Loewner order. Since

$$W^{1/2} C^+ W^{1/2} = C_{H^+}, \quad W^{1/2} \hat{C}^+ W^{1/2} = W^+,$$

we obtain

$$\kappa_{\text{app}}^{-1} C_{H^+} \preceq W^+ \preceq \kappa_{\text{app}} C_{H^+}.$$

$\square$

5.4 Contraction of D_{sym} and Comparison with the Original Distribution

We prove one-step progress and reachable-state control, then show that when D_{sym} is sufficiently small, the original distribution ν_0 is well-rounded and comparable in relative Loewner order to the final distribution ν_H . These results prove the contraction and show that the final covariance-estimation call in the next subsection satisfies its hypotheses. The next lemma concerns the one-step decrease of D_{sym} .

Lemma 5.19. *Let $H \succeq 0$, let $W \approx_{\kappa_{\text{app}}} C_H$, and let H^+ be the update in [Lemma 5.17](#). If C_L is sufficiently large, then*

$$D_{\text{sym}}(H^+) \leq D_{\text{sym}}(H) - \frac{\delta_{\text{sym}}}{2} q_{\text{sym}}(H).$$

Proof. Set $\Delta = H^+ - H$. Since $H, H^+ \succeq 0$, both normalizing constants are positive and finite. Jensen's inequality gives

$$\begin{aligned} D_{\text{sym}}(H^+) - D_{\text{sym}}(H) &= -\log \mathbb{E}_{X \sim \nu_H} \left[e^{-X^\top \Delta X} \right] \\ &\leq \mathbb{E}_{X \sim \nu_H} \left[X^\top \Delta X \right] = \text{Tr}(C_H \Delta). \end{aligned}$$

The quadratic expectation is finite because ν_H has finite second moments. Substituting the update formula gives

$$\begin{aligned} D_{\text{sym}}(H^+) - D_{\text{sym}}(H) &\leq \text{Tr}(C_H(H^+ - H)) \\ &= -\delta_{\text{sym}} q_{\text{sym}}(H) + \eta_{\text{ridge}} \frac{\hat{q}}{n} \text{Tr}(C_H W^{-1}). \end{aligned}$$

The covariance comparison gives

$$\hat{q} = \text{Tr}(W H) \leq \kappa_{\text{app}} q_{\text{sym}}(H), \quad \text{Tr}(C_H W^{-1}) \leq \kappa_{\text{app}} n.$$

Substitution gives

$$D_{\text{sym}}(H^+) - D_{\text{sym}}(H) \leq -(\delta_{\text{sym}} - \eta_{\text{ridge}} \kappa_{\text{app}}^2) q_{\text{sym}}(H).$$

Increasing C_L ensures $\eta_{\text{ridge}} \kappa_{\text{app}}^2 \leq \delta_{\text{sym}}/2$. □

The next lemma combines energy domination, the grounded inverse-covariance bound, and the global moment promise to control the reachable precision and compare $D_{\text{sym}}(H)$ with $q_{\text{sym}}(H)$.

The next lemma concerns the reachable cap, comparison, and contraction.

Lemma 5.20. *Start from $H_0 = r^{-2}I$ and from a successful initial estimate $W_0 \approx_{\kappa_{\text{app}}} C_{H_0}$ supplied by [Lemma 5.16](#). Suppose that along the first j updates every covariance call succeeds, so $W_i \approx_{\kappa_{\text{app}}} C_{H_i}$ for $0 \leq i \leq j$. Then a constant C_{prec} depending only on h_0 satisfies*

$$\|H_i\|_{\text{op}} \leq C_{\text{prec}} \frac{n^2}{r^2} \quad (0 \leq i \leq j).$$

Every such reachable H satisfies, for a constant C_D depending only on h_0 ,

$$q_{\text{sym}}(H) \leq D_{\text{sym}}(H) \leq C_D L_{\text{sym}} q_{\text{sym}}(H).$$

Consequently

$$D_{\text{sym}}(H_{i+1}) \leq \left(1 - \frac{\delta_{\text{sym}}}{2C_D L_{\text{sym}}}\right) D_{\text{sym}}(H_i).$$

For every sufficiently small universal $c_* > 0$, there is a sufficiently large constant C_T such that, if $T = \lceil C_T L_{\text{sym}}^4 \rceil$ and the first T update calls succeed, then

$$D_{\text{sym}}(H_T) \leq c_*.$$

Proof. Let $C_p = C_{h_0}$ be the constant from the grounded lower covariance estimate in [Lemma 5.9](#), and define $s = n^2/r^2$ and $u_i = \|H_i\|_{\text{op}}$. By the energy bound in [Lemma 5.5](#), $q_{\text{sym}}(H_i) \leq n/2$, and therefore $\hat{q}_i \leq \kappa_{\text{app}} n/2$. From $W_i \succeq \kappa_{\text{app}}^{-1} C_{H_i}$ and the resulting inverse-covariance bound in [Lemma 5.9](#),

$$\|W_i^{-1}\|_{\text{op}} \leq \kappa_{\text{app}} C_p (u_i + s).$$

For $c_0 = \kappa_{\text{app}}^2 C_p/2$, the update satisfies

$$\begin{aligned} u_{i+1} &\leq (1 - \delta_{\text{sym}}) u_i + \eta_{\text{ridge}} \frac{\hat{q}_i}{n} \|W_i^{-1}\|_{\text{op}} \\ &\leq (1 - \delta_{\text{sym}} + c_0 \eta_{\text{ridge}}) u_i + c_0 \eta_{\text{ridge}} s. \end{aligned}$$

Because $\eta_{\text{ridge}}/\delta_{\text{sym}} = L_{\text{sym}}^{-1}$, increasing C_L makes $c_0\eta_{\text{ridge}} \leq \delta_{\text{sym}}/2$. Since $u_0 = r^{-2} \leq s$, we have

$$u_{i+1} \leq \left(1 - \frac{\delta_{\text{sym}}}{2}\right) u_i + c_0\eta_{\text{ridge}}s.$$

Let $C_{\text{prec}} = \max\{1, 2c_0\}$. If $u_i \leq C_{\text{prec}}s$, then

$$u_{i+1} \leq C_{\text{prec}}s - \left(\frac{C_{\text{prec}}\delta_{\text{sym}}}{2} - c_0\eta_{\text{ridge}}\right)s \leq C_{\text{prec}}s,$$

where the last inequality uses $\eta_{\text{ridge}} \leq \delta_{\text{sym}}$. Induction proves $u_i \leq C_{\text{prec}}s$.

Fix a reachable H and define

$$Q_H(x) = x^\top Hx, \quad Z(t) = \int e^{-V(x)-tQ_H(x)} dx, \quad N(t) = \int Q_H(x) e^{-V(x)-tQ_H(x)} dx.$$

The moment promise and the fourth-moment bound in [Proposition 2.23](#)(iii) give

$$\mathbb{E}_\pi [\|X\|^4] \leq n \sum_{j=1}^n \mathbb{E}_\pi [X_j^4] \leq C_4 n \left(\sum_{j=1}^n \mathbb{E}_\pi [X_j^2] \right)^2 \leq C_4 n R^4.$$

Since

$$0 \leq Q_H(x) \leq \|H\|_{\text{op}} \|x\|^2,$$

both Q_H and Q_H^2 are integrable under π . For $t \in [0, 1]$, $e^{-tQ_H} \leq 1$, so these same integrable functions dominate the derivatives of $Z(t)$ and $N(t)$. Differentiation under the integral gives

$$Z'(t) = -N(t), \quad N'(t) = - \int Q_H(x)^2 e^{-V(x)-tQ_H(x)} dx.$$

Moreover,

$$e(t) := \text{Tr}(HC_{tH}) = \frac{N(t)}{Z(t)}.$$

It follows that

$$D_{\text{sym}}(H) = \int_0^1 e(t) dt, \quad e'(t) = -\text{Var}_{\nu_{tH}} [X^\top HX] \leq 0.$$

It follows that $D_{\text{sym}}(H) \geq e(1) = q_{\text{sym}}(H)$. For the reverse inequality set

$$\tau = \min \left\{ 1, \frac{r^2}{n^2 R^2} \right\} = \min \{1, (sR^2)^{-1}\}.$$

The moment promise gives $C_0 \preceq R^2 I$ and hence

$$\int_0^\tau e(t) dt \leq \tau e(0) \leq \tau R^2 \text{Tr}(H) \leq \frac{\text{Tr}(H)}{s}.$$

The inverse-covariance bound and the precision cap imply

$$C_H \succeq \frac{1}{C_p(C_{\text{prec}} + 1)s} I, \quad q_{\text{sym}}(H) \geq \frac{\text{Tr}(H)}{C_p(C_{\text{prec}} + 1)s}.$$

Therefore

$$\int_0^\tau e(t) dt \leq C_p(C_{\text{prec}} + 1)q_{\text{sym}}(H).$$

If $\tau < 1$, energy monotonicity in [Lemma 5.5](#) says that $te(t)$ is nondecreasing, so for $t \in [\tau, 1]$,

$$e(t) \leq q_{\text{sym}}(H)/t.$$

Consequently

$$\int_{\tau}^1 e(t) dt \leq q_{\text{sym}}(H) \log(1/\tau) \leq 2L_{\text{sym}} q_{\text{sym}}(H),$$

where $\log(1/\tau) = 2 \log(nR/r) \leq 2L_{\text{sym}}$ when $\tau < 1$. Adding the two integrals proves the comparison between D_{sym} and q_{sym} with $C_D = C_p(C_{\text{prec}} + 1) + 2$, including $H = 0$.

Combining this comparison with the one-step decrease in [Lemma 5.19](#) gives the explicit contraction

$$D_{\text{sym}}(H_{i+1}) \leq \left(1 - \frac{\delta_{\text{sym}}}{2C_D L_{\text{sym}}}\right) D_{\text{sym}}(H_i).$$

Finally, Jensen's inequality and $H_0 = r^{-2}I$ give

$$D_{\text{sym}}(H_0) = -\log \mathbb{E}_{\pi} \left[e^{-\|X\|^2/r^2} \right] \leq \frac{\mathbb{E}_{\pi} [\|X\|^2]}{r^2} \leq \frac{R^2}{r^2}.$$

The logarithm of the ratio between this bound and c_* is $O(L_{\text{sym}})$. Since $\delta_{\text{sym}} = L_{\text{sym}}^{-2}$, iterating the contraction gives

$$D_{\text{sym}}(H_T) \leq \exp \left(-\frac{T\delta_{\text{sym}}}{2C_D L_{\text{sym}}} \right) D_{\text{sym}}(H_0).$$

Since $T\delta_{\text{sym}}/L_{\text{sym}} = \Theta(C_T L_{\text{sym}})$, choosing C_T sufficiently large makes the last expression at most c_* . $\square$

Comparison with the original distribution. The next lemma concerns the comparison with the original distribution.

Lemma 5.21. *Let H be reachable in the sense of [Lemma 5.20](#), let $W \approx_{\kappa_{\text{app}}} C_H$, and suppose*

$$D_{\text{sym}}(H) \leq c_*$$

for a sufficiently small universal $c_* > 0$. Then the exact transformed potential

$$P_*(z) = V(W^{1/2}z)$$

is (b_0, c_1, c_2) -well-rounded for universal constants. Moreover, if $C_0 = \mathbb{E}_{\pi} [XX^{\top}]$, then

$$e^{-D_{\text{sym}}(H)} C_H \preceq C_0 \preceq \frac{e^{-D_{\text{sym}}(H)}}{1 - (C_4 D_{\text{sym}}(H))^{1/2}} C_H,$$

where C_4 is the fourth-moment constant from part (iii) of [Proposition 2.23](#).

Proof. Let $Q(X) = X^{\top} H X$. Splitting the original-law second moment according to $1 = e^{-Q} + (1 - e^{-Q})$ gives

$$C_0 = e^{-D_{\text{sym}}(H)} C_H + E_H, \quad \text{where} \quad E_H = \mathbb{E}_{\pi} \left[(1 - e^{-Q(X)}) X X^{\top} \right].$$

For every vector u , Cauchy–Schwarz and item 3 of [Proposition 2.23](#), the fourth-moment estimate for affine functionals, give

$$\begin{aligned} u^\top E_H u &\leq (\mathbb{E}_\pi \left[(1 - e^{-Q(X)})^2 \right])^{1/2} (\mathbb{E}_\pi \left[(u^\top X)^4 \right])^{1/2} \\ &\leq C_4^{1/2} (1 - e^{-D_{\text{sym}}(H)})^{1/2} u^\top C_0 u \\ &\leq (C_4 D_{\text{sym}}(H))^{1/2} u^\top C_0 u. \end{aligned}$$

Here

$$0 \leq 1 - e^{-Q(X)} \leq 1, \quad \mathbb{E}_\pi \left[1 - e^{-Q(X)} \right] = 1 - e^{-D_{\text{sym}}(H)},$$

so $\mathbb{E}_\pi \left[(1 - e^{-Q(X)})^2 \right] \leq 1 - e^{-D_{\text{sym}}(H)}$. Thus $0 \preceq E_H \preceq (C_4 D_{\text{sym}}(H))^{1/2} C_0$. Let $\alpha = (C_4 D_{\text{sym}}(H))^{1/2}$. Substitution into the exact identity gives

$$e^{-D_{\text{sym}}(H)} C_H \preceq C_0, \quad (1 - \alpha) C_0 \preceq e^{-D_{\text{sym}}(H)} C_H.$$

Since $\alpha < 1$, division by $1 - \alpha$ proves the two stated relative-Loewner bounds.

Choose c_* so that $(C_4 c_*)^{1/2} \leq 1/4$. Since $W \approx_{\kappa_{\text{app}}} C_H$, the covariance of the exact law in W -coordinates satisfies

$$cI \preceq W^{-1/2} C_0 W^{-1/2} \preceq CI$$

for universal $c, C > 0$. In particular, its second moment is at most Cn . Item 4 of [Proposition 2.23](#), the covariance-to-level-set implication, supplies a universal-radius ball in an $O(n)$ -height sublevel set of P_* . Hence P_* satisfies both clauses of [Definition 5.1](#). $\square$

5.5 Algorithm and Complexity

We assemble the regularized update into the full symmetric algorithm and use the preceding invariant and contraction results for its correctness, failure accounting, deterministic evaluation-query limit, and well-rounded linear output. Indicator potentials recover the membership-oracle convex-body specialization. The next definition concerns the adaptive state update.

Definition 5.22. For a positive-definite W and $H \succeq 0$, define

$$\widehat{q}_{\text{sym}}(H, W) = \text{Tr}(W H), \quad \mathcal{U}(H, W) = (1 - \delta_{\text{sym}})H + \eta_{\text{ridge}} \frac{\widehat{q}_{\text{sym}}(H, W)}{n} W^{-1}.$$

For $H^+ = \mathcal{U}(H, W)$ define

$$J_{\text{sym}}(H^+, W) = W^{1/2} H^+ W^{1/2}, \quad P(H^+, W)(z) = V(W^{1/2} z) + z^\top J_{\text{sym}}(H^+, W) z.$$

Fix once and for all a single triple of constants for the estimator in [Theorem 11](#): take its radius constant no larger than the smallest radius certified by [Lemmas 5.16, 5.17](#) and [5.21](#), and take its moment and height constants no smaller than the largest of the corresponding three certificates. Since h_0 is fixed, this common triple is universal. Every estimator call below uses this same choice.

SYMMETRIC-COVARIANCE.

Input: An even grounded input $(V, 0, r, R)$ in the sense of [Definition 2.21](#) (the ground point of an even potential is $x_0 = 0$), and a failure probability $\beta \in (0, 1)$.

Parameters:

$$L_{\text{sym}} = C_L \max\{1, \log(2 + nR/r)\}, \quad T = \lceil C_T L_{\text{sym}}^4 \rceil, \quad \xi = \frac{\beta}{2(T+2)}.$$

In particular, $\xi \in (0, 1/2]$ for every $n \geq 1$. The update map $\mathcal{U}$ of [Definition 5.22](#) is used with $\delta_{\text{sym}} = L_{\text{sym}}^{-2}$ and $\eta_{\text{ridge}} = L_{\text{sym}}^{-3}$, and every call to the covariance estimator of [Theorem 11](#) runs at its fixed accuracy $\kappa_{\text{app}} = 1.01$.

Output: A positive-definite matrix Σ and a well-rounding linear map T_* , or a failure symbol.

1. Set $H_0 = r^{-2}I$ and define the rescaled initial potential

$$P_0(z) = V(rz) + \|z\|^2.$$

Invoke the covariance estimator of [Theorem 11](#) on (P_0, ξ) . If the call fails, return the failure symbol. Otherwise, denote its output by $\widehat{C}_0$. If a required inverse or positive square root does not exist, return the failure symbol; otherwise, set $W_0 = r^2 \widehat{C}_0$.

2. For $j = 0, \dots, T-1$, form

$$H_{j+1} = \mathcal{U}(H_j, W_j), \quad P_{j+1} = P(H_{j+1}, W_j).$$

Invoke the covariance estimator of [Theorem 11](#) on (P_{j+1}, ξ) . If the call fails, return the failure symbol. Otherwise, denote its output by $\widehat{C}_{j+1}$. If a required inverse or positive square root does not exist, return the failure symbol; otherwise, set

$$W_{j+1} = W_j^{1/2} \widehat{C}_{j+1} W_j^{1/2}.$$

3. Define the final exact transformed potential

$$P_*(z) = V(W_T^{1/2} z).$$

Invoke the covariance estimator of [Theorem 11](#) on (P_*, ξ) . If the call fails or a required inverse or positive square root does not exist, return the failure symbol. Otherwise, denote its output by $\widehat{C}_*$ and return

$$\Sigma = W_T^{1/2} \widehat{C}_* W_T^{1/2}, \quad T_* = W_T^{-1/2}.$$

Figure 3: Near-cubic covariance estimation for even potentials.

For any transformed distribution, evaluating the displayed potential at z uses one query at the single point $W^{1/2}z$ (or rz in the initial call), plus evaluation of a known quadratic. Thus each estimator evaluation uses exactly one evaluation of V . When $V = I_K$, this is exactly one membership query to K .

Proof of Theorem 2. The algorithm is SYMMETRIC-COVARIANCE (Figure 3). By Lemma 5.16, the initial quadratic tilt has the fixed-radius sublevel-set ball and bounded second moment required by the covariance estimator, so P_0 is well-rounded. On success of the initial call, $W_0 \approx_{\kappa_{\text{app}}} C_{H_0}$. Condition on the event that all calls before update j succeeded. Repeated use of Lemma 5.18, which transports a successful covariance estimate back by congruence, gives $W_j \approx_{\kappa_{\text{app}}} C_{H_j}$, and Lemma 5.17, which controls the next trace and sublevel-set ball, verifies both clauses of Definition 5.1 for the explicit potential P_{j+1} . The covariance estimator is therefore applicable with failure tolerance ξ , and on success the same congruence argument establishes the next invariant.

Conditioned on success through the first T update calls, Lemma 5.20, whose contraction estimate gives $D_{\text{sym}}(H_T) \leq c_*$, shows that the terminal threshold is reached. Thus Lemma 5.21, the comparison between the moments under ν_{H_T} and ν_0 , verifies both well-roundedness clauses for the original transformed potential P_* . A successful final call gives

$$\hat{C}_* \approx_{\kappa_{\text{app}}} W_T^{-1/2} C_0 W_T^{-1/2},$$

and congruence by $W_T^{1/2}$ yields $\Sigma \approx_{\kappa_{\text{app}}} C_0$. Since $\kappa_{\text{app}}^{-1} > 0.99$, this is the required matrix inequality. The law π is full dimensional because its density has positive integral with respect to n -dimensional Lebesgue measure, so C_0 and hence Σ are positive definite. The well-rounded output clauses are exactly those proved for P_* in Lemma 5.21.

The covariance-estimator guarantee applies conditionally on every preceding successful history. There are $T + 2$ calls, so a union bound gives

$$\mathbb{P}(\text{some call fails}) \leq (T + 2)\xi = \frac{\beta}{2} \leq \beta.$$

This argument does not condition the cost bound on success. Each call is stopped at the estimator's deterministic query limit. At $H = 0$, the grounded lower covariance estimate in Lemma 5.9 gives

$$C_0 \succeq c_{h_0} \frac{r^2}{n^2} I.$$

Consequently the moment promise implies

$$R^2 \geq \text{Tr}(C_0) \geq c_{h_0} \frac{r^2}{n}, \quad \log n = O_{h_0}(L_{\text{sym}}).$$

Since $T + 2 = O(L_{\text{sym}}^4)$, it follows that

$$\log(n/\xi) = O_{h_0}(\log(2 + nR/(r\beta))).$$

Multiplying the $T + 2$ calls by the estimator's per-call $n^3 \log(n/\xi)^{O(1)}$ query limit gives the deterministic total query limit

$$O\left(n^3 \log(2 + nR/(r\beta))^{O(1)}\right)$$

evaluations. The one-query transformed-oracle implementation completes the resource accounting and its membership-oracle specialization. $\square$

References

- [BL76] Herm Jan Brascamp and Elliott H. Lieb. On Extensions of the Brunn-Minkowski and Prékopa-Leindler Theorems, Including Inequalities for Log-Concave Functions, and with an Application to the Diffusion Equation. *Journal of Functional Analysis*, 22(4):366–389, 1976.

- [Caf00] Luis A. Caffarelli. Monotonicity properties of optimal transportation and the fkg and related inequalities. *Communications in Mathematical Physics*, 214(3):547–563, 2000.
- [CDWY17] Yuansi Chen, Raaz Dwivedi, Martin J. Wainwright, and Bin Yu. Vaidya walk: A sampling algorithm based on the volumetric barrier. In *2017 55th Annual Allerton Conference on Communication, Control, and Computing*, pages 1220–1227, 2017.
- [CDWY18] Yuansi Chen, Raaz Dwivedi, Martin J. Wainwright, and Bin Yu. Fast mcmc sampling algorithms on polytopes. *Journal of Machine Learning Research*, 19(55):1–86, 2018.
- [Che21] Yuansi Chen. An almost constant lower bound of the isoperimetric coefficient in the KLS conjecture. *Geometric and Functional Analysis*, 31:34–61, 2021.
- [CV18] Ben Cousins and Santosh Vempala. Gaussian cooling and $o^*(n^3)$ algorithms for volume and gaussian volume. *SIAM Journal on Computing*, 47(3):1237–1273, 2018. Available at: <https://arxiv.org/abs/1409.6011>.
- [DFK91] Martin E. Dyer, Alan M. Frieze, and Ravi Kannan. A random polynomial time algorithm for approximating the volume of convex bodies. *Journal of the ACM*, 38(1):1–17, 1991.
- [EL14] Ronen Eldan and Joseph Lehec. Bounding the norm of a log-concave vector via thin-shell estimates. In Bo’az Klartag and Emanuel Milman, editors, *Geometric Aspects of Functional Analysis*, volume 2116 of *Lecture Notes in Mathematics*, pages 107–122. Springer, Cham, 2014.
- [Eld13] Ronen Eldan. Thin shell implies spectral gap up to polylog via a stochastic localization scheme. *Geometric and Functional Analysis*, 23(2):532–569, 2013.
- [Fle10] B. Fleury. Concentration in a thin Euclidean shell for log-concave measures. *Journal of Functional Analysis*, 259(4):832–841, 2010.
- [Fra97] Matthieu Fradelizi. Sections of convex bodies through their centroid. *Archiv der Mathematik*, 69:515–522, 1997.
- [GC11] Mark Girolami and Ben Calderhead. Riemann manifold langevin and hamiltonian monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011.
- [GKM⁺24] Yuzhou Gu, Nikki Lijing Kuang, Yi-An Ma, Zhao Song, and Lichen Zhang. Log-concave sampling from a convex body with a barrier: a robust and unified dikin walk. In *Advances in Neural Information Processing Systems*, volume 37, 2024.
- [GKV24] Khashayar Gatmiry, 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, 2024.
- [GM11] Olivier Guédon and Emanuel Milman. Interpolating thin-shell and sharp large-deviation estimates for isotropic log-concave measures. *Geometric and Functional Analysis*, 21(5):1043–1068, 2011.

- [Gua24] Qingyang Guan. A note on bourgain’s slicing problem, 2024.
- [JC25] Minhui Jiang and Yuansi Chen. Regularized Dikin walks for sampling truncated log-concave measures, mixed isoperimetry and beyond worst-case analysis. In *Proceedings of Thirty Eighth Conference on Learning Theory*, volume 291 of *Proceedings of Machine Learning Research*, pages 3017–3078. PMLR, 2025.
- [JLLV26] He Jia, Aditi Laddha, Yin Tat Lee, and Santosh Vempala. Reducing isotropy and volume to kls: Faster rounding and volume algorithms. *J. ACM*, 73(2), April 2026.
- [JLV22] Arun Jambulapati, Yin Tat Lee, and Santosh S. Vempala. A slightly improved bound for the KLS constant, 2022.
- [KL22] Bo’az Klartag and Joseph Lehec. Bourgain’s slicing problem and KLS isoperimetry up to polylog. *Geometric and Functional Analysis*, 32(5):1134–1159, 2022.
- [KL25a] Bo’az Klartag and Joseph Lehec. Affirmative resolution of bourgain’s slicing problem using guan’s bound. *Geometric and Functional Analysis*, 35:1147–1168, 2025.
- [KL25b] Boaz Klartag and Joseph Lehec. Thin-shell bounds via parallel coupling. *arXiv preprint arXiv:2507.15495*, 2025.
- [Kla06] B. Klartag. On convex perturbations with a bounded isotropic constant. *Geometric and Functional Analysis*, 16:1274–1290, 2006.
- [Kla23] Bo’az Klartag. Logarithmic bounds for isoperimetry and slices of convex sets. *Ars Inveniendi Analytica*, pages Paper No. 4, 17 pp., 2023.
- [KLS95] R. Kannan, L. Lovász, and M. Simonovits. Isoperimetric problems for convex bodies and a localization lemma. *Discrete Comput. Geom.*, 13(3-4):541–559, 1995.
- [KLS97] Ravi Kannan, László Lovász, and Miklós Simonovits. Random walks and an $O^*(n^5)$ volume algorithm for convex bodies. *Random Structures & Algorithms*, 11(1):1–50, 1997.
- [KLSV22] Yunbum Kook, Yin-Tat Lee, Ruoyi Shen, and Santosh Vempala. Sampling with riemannian hamiltonian monte carlo in a constrained space. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 31684–31696. Curran Associates, Inc., 2022.
- [KLSV23] Yunbum Kook, Yin Tat Lee, Ruoyi Shen, and Santosh Vempala. Condition-number-independent convergence rate of riemannian hamiltonian monte carlo with numerical integrators. In Gergely Neu and Lorenzo Rosasco, editors, *Proceedings of Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pages 4504–4569. PMLR, 2023.
- [KN12] 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.
- [KV25] Yunbum Kook and Santosh S. Vempala. Sampling and integration of log-concave functions by algorithmic diffusion. In Michal Koucký and Nikhil Bansal, editors, *Proceedings of the 57th Annual ACM Symposium on Theory of Computing, STOC 2025, Prague*,

- Czechia, June 23-27, 2025*, pages 924–932. ACM, 2025. Full version (used for theorem numbering): <https://arxiv.org/abs/2411.13462v1>.
- [KZ24] Yunbum Kook and Matthew S. Zhang. Covariance estimation using Markov chain Monte Carlo, 2024. <https://arxiv.org/abs/2410.17147v1>.
 - [Let26] Brayden Letwin. The KLS constant is $O(\log^{1/4} n)$, 2026.
 - [LLV20] Aditi Laddha, Yin Tat Lee, and Santosh S. Vempala. Strong self-concordance and sampling. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1212–1222, 2020.
 - [LS93] László Lovász and Miklós Simonovits. Random walks in a convex body and an improved volume algorithm. *Random Structures & Algorithms*, 4(4):359–412, 1993.
 - [LV06a] László Lovász and Santosh S. Vempala. Hit-and-run from a corner. *SIAM Journal on Computing*, 35(4):985–1005, 2006.
 - [LV06b] László Lovász and Santosh S. Vempala. Simulated annealing in convex bodies and an $O^*(n^4)$ volume algorithm. *Journal of Computer and System Sciences*, 72(2):392–417, 2006.
 - [LV07] László Lovász and Santosh S. Vempala. The geometry of log-concave functions and sampling algorithms. *Random Structures & Algorithms*, 30(3):307–358, 2007.
 - [LV18] 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.
 - [LV22] Yin Tat Lee and Santosh S. Vempala. Geodesic walks in polytopes. *SIAM Journal on Computing*, 51(2):STOC17–400–STOC17–488, 2022. Conference version in STOC 2017.
 - [LV24] Yin Tat Lee and Santosh S. Vempala. Eldan’s stochastic localization and the KLS conjecture: Isoperimetry, concentration and mixing. *Annals of Mathematics*, 199(3):1043–1092, 2024.
 - [MV23] Oren Mangoubi and Nisheeth K. Vishnoi. Sampling from structured log-concave distributions via a soft-threshold dikin walk. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 31908–31942. Curran Associates, Inc., 2023.
 - [Nar16] Hariharan Narayanan. Randomized interior point methods for sampling and optimization. *The Annals of Applied Probability*, 26(1):597–641, 2016.
 - [NR17] Hariharan Narayanan and Alexander Rakhlin. Efficient sampling from time-varying log-concave distributions. *Journal of Machine Learning Research*, 18(112):1–29, 2017.
 - [Pao06] G Paouris. Concentration of mass on convex bodies. *GAFa Geom. Funct. Anal.*, 16(5):1021–1049, December 2006.
 - [Pao12] Grigoris Paouris. Small ball probability estimates for log-concave measures. *Transactions of the American Mathematical Society*, 364(1):287–308, 2012.

- [Pré73] András Prékopa. On log-concave measures and functions. *Acta Sci. Math. (Szeged)*, 34:335–343, 1973.
- [SV16] 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.
- [Tro15] Joel A. Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends in Machine Learning*, 8(1–2):1–230, 2015.

A Proof of the Sharp Trace Growth Theorem

This section proves [Theorem 7](#), which shows that the trace growth bounds used in the $\tilde{O}(n^{13/4})$ analysis of the asymmetric algorithm in [Section 4](#) are sharp up to constants. We construct restricted Gaussian examples where anisotropic warming creates trace growth of the same order as those bounds.

A.1 The lower-bound construction

We now construct the matching family. The mechanism is a temperature-driven switch in a one-dimensional control variable. Here a fiber means the slice of the construction obtained by fixing this control variable and looking at the remaining coordinates. One collection of such slices is temperature-sensitive and expands with the control variable; another is almost temperature-insensitive and contracts with it. The parameters are chosen so that these two contributions balance at the original temperature but move in opposite directions after warming.

A.1.1 A truncated-Gaussian response function

For $t > 0$ and $r > 0$, set

$$J_t(r) = \int_{-r}^r e^{-v^2/(2t)} dv, \quad D_t(r) = r \partial_r \log J_t(r). \quad (\text{A.1})$$

Writing $D = D_1$, scaling gives

$$D_t(r) = D(r/t^{1/2}). \quad (\text{A.2})$$

The next lemma concerns the temperature response.

Lemma A.1. *The function D is strictly decreasing on $(0, \infty)$ and*

$$D(r) = 1 - \frac{r^2}{3} + O(r^4) \quad (r \rightarrow 0). \quad (\text{A.3})$$

Consequently, one may fix a sufficiently small universal $b > 0$ and constants $c_0, C_0 > 0$ such that, for every $T = 1 + \delta \in [1, 2]$,

$$c_0 \delta \leq D_T(1)D_1(b) - D_1(1)D_T(b) \leq C_0 \delta. \quad (\text{A.4})$$

Proof. Direct differentiation gives

$$D'(r) = \frac{D(r)}{r} (1 - r^2 - D(r)). \quad (\text{A.5})$$

We claim that $D(r) > 1 - r^2$. This is immediate for $r \geq 1$. For $0 < r < 1$, let

$$H(r) = 2re^{-r^2/2} - (1 - r^2)J_1(r).$$

Then $H(0) = 0$ and $H'(r) = 2rJ_1(r) > 0$. Since $H(r) = J_1(r)[D(r) - (1 - r^2)]$, the claim follows. It follows from Equation (A.5) that $D'(r) < 0$. Expanding the numerator and denominator in the definition of D gives Equation (A.3).

By Equation (A.2),

$$\partial_t D_t(1) = -\frac{1}{2t^{3/2}} D'(t^{-1/2}),$$

which has a strictly positive minimum for $1 \leq t \leq 2$. Hence $D_T(1) - D_1(1) \asymp \delta$. On the other hand, Equation (A.3) gives

$$D_t(b) = 1 - \frac{b^2}{3t} + O(b^4), \quad D_T(b) - D_1(b) = O(b^2\delta),$$

uniformly for $1 \leq t, T \leq 2$; here the remainder is smooth in t and its t -derivative is $O(b^4)$. Taking b sufficiently small proves Equation (A.4). $\square$

A.1.2 The polytope and its phase switch

Assume first that $n \geq 8$. Let

$$m = \lfloor n/4 \rfloor, \quad r = n - 2m - 2. \tag{A.6}$$

Then $m \geq n/8$ and $r \geq n/4$. Use coordinates

$$(x, y_1, \dots, y_m, z_1, \dots, z_m, w, v_1, \dots, v_r).$$

Fix the constant b from Lemma A.1. This lemma gives a uniform linear separation between the temperature responses at 1 and $T = 1 + \delta$, which is used to choose η below. Let $\kappa > 0$ be a sufficiently small universal constant, define

$$\varepsilon = \kappa\delta, \quad \eta = \varepsilon \frac{D_1(1) + D_T(1)}{D_1(b) + D_T(b)}, \quad T = 1 + \delta. \tag{A.7}$$

Uniformly in $0 \leq \delta \leq 1$, we have $\eta \asymp \varepsilon$; after decreasing κ , we may assume $0 \leq \eta \leq 1/2$.

For $s \in [0, 1]$, write $u = x + s$ and define the polytope K_s by

$$\begin{aligned} 0 &\leq u \leq 1, \\ |y_i| &\leq 1 + \varepsilon u \quad (1 \leq i \leq m), \\ |z_i| &\leq b(1 - \eta u) \quad (1 \leq i \leq m), \\ |w| &\leq u, \\ |v_j| &\leq 1 \quad (1 \leq j \leq r). \end{aligned} \tag{A.8}$$

All constraints are affine, so K_s is a full-dimensional compact convex polytope.

Before whitening, take the ambient Gaussian covariance to be tI_n . The u -marginal at scale t is proportional on $[0, 1]$ to

$$H_{t,s}(u)e^{F_t(u)}, \tag{A.9}$$

where

$$H_{t,s}(u) = e^{-(u-s)^2/(2t)} \int_{-u}^u e^{-w^2/(2t)} dw, \quad (\text{A.10})$$

and

$$F_t(u) = m \log J_t(1 + \varepsilon u) + m \log J_t(b(1 - \eta u)). \quad (\text{A.11})$$

For $s, u \in [0, 1]$ and $t \in [1, 2]$,

$$c_1 u \leq H_{t,s}(u) \leq C_1 u. \quad (\text{A.12})$$

The next lemma concerns the uniform slope reversal.

Lemma A.2. *There are universal constants $c, C > 0$ such that, with κ chosen sufficiently small and*

$$\lambda = m\kappa\delta^2, \quad (\text{A.13})$$

one has, for every $0 \leq u \leq 1$,

$$-C\lambda \leq F'_1(u) \leq -c\lambda, \quad c\lambda \leq F'_T(u) \leq C\lambda. \quad (\text{A.14})$$

Proof. Let $L_t(q) = \partial_q \log J_t(q) = D_t(q)/q$. Differentiating Equation (A.11) gives

$$F'_t(u) = m \left[\varepsilon L_t(1 + \varepsilon u) - b\eta L_t(b(1 - \eta u)) \right]. \quad (\text{A.15})$$

At $u = 0$,

$$F'_t(0) = m(\varepsilon D_t(1) - \eta D_t(b)).$$

The choice of η gives

$$F'_T(0) = -F'_1(0) = m\varepsilon \frac{D_T(1)D_1(b) - D_1(1)D_T(b)}{D_1(b) + D_T(b)}. \quad (\text{A.16})$$

By Lemma A.1, the right-hand side is between $cm\kappa\delta^2$ and $Cm\kappa\delta^2$.

The radii in Equation (A.15) stay in fixed compact subintervals of $(0, \infty)$. Differentiating once more and using $\eta \asymp \varepsilon$ gives

$$|F'_t(u) - F'_t(0)| \leq Cm(\varepsilon^2 + \eta^2) \leq Cm\kappa^2\delta^2. \quad (\text{A.17})$$

Choosing κ sufficiently small relative to the lower constant in Equation (A.16) proves Equation (A.14). $\square$

A.1.3 Centering, endpoint concentration, and whitening

We first choose s so that the scale-1 law is centered. Let ν_s be the standard Gaussian restricted to K_s , and define

$$\Phi(s) = \mathbb{E}_{u \sim \nu_s} [u] - s.$$

The function Φ is continuous. Since $u > 0$ almost surely when $s = 0$ and $u < 1$ almost surely when $s = 1$, we have $\Phi(0) > 0$ and $\Phi(1) < 0$. Therefore some $s \in (0, 1)$ satisfies

$$s = \mathbb{E}_{u \sim \nu_s} [u]. \quad (\text{A.18})$$

For this choice, $\mathbb{E}_{x \sim \nu_s} [x] = \mathbb{E}_{u \sim \nu_s} [u - s] = 0$. Every other coordinate has mean zero by sign symmetry. The same conditional sign symmetries imply that all cross-covariances vanish. Thus the scale-1 covariance matrix C is positive definite and diagonal.

For the chosen s , write $\mathbb{P}_t$ for probability under the law μ_{K_s, tI_n} obtained by restricting $N(0, tI_n)$ to K_s .

Suppose first that $\lambda \geq 1$. From Equation (A.12) and the first pair of inequalities in Equation (A.14),

$$\begin{aligned} \mathbb{E}_{u \sim \mu_{K_s, I_n}} [u^2] &= \frac{\int_0^1 u^2 H_{1,s}(u) e^{F_1(u)} du}{\int_0^1 H_{1,s}(u) e^{F_1(u)} du} \\ &\leq C \frac{\int_0^\infty u^3 e^{-c\lambda u} du}{\int_0^{1/\lambda} u e^{-C\lambda u} du} \leq \frac{C}{\lambda^2}. \end{aligned} \quad (\text{A.19})$$

At scale T , Equation (A.14) and Equation (A.12) show, by comparing $[0, 1/2]$ with $[3/4, 1]$, that

$$\mathbb{P}_T(u \geq 1/2) \geq c, \quad \mathbb{E}_{u \sim \mu_{K_s, TI_n}} [u^2] \geq c. \quad (\text{A.20})$$

Conditionally on u , the w -coordinate is the centered $N(0, t)$ law restricted to $[-u, u]$. Uniformly for $0 \leq u \leq 1$ and $1 \leq t \leq 2$,

$$cu^2 \leq \mathbb{E}_{w \sim \mu_{K_s, tI_n}} [w^2 \mid u] \leq Cu^2. \quad (\text{A.21})$$

Since $\mathbb{E}_{w \sim \mu_{K_s, tI_n}} [w \mid u] = 0$, Equations (A.19)–(A.21) imply

$$\frac{\text{Var}_{w \sim \mu_{K_s, TI_n}} [w]}{\text{Var}_{w \sim \mu_{K_s, I_n}} [w]} \geq c\lambda^2 \geq cn^2\delta^4. \quad (\text{A.22})$$

Here and below the universal constants absorb the fixed κ and the fact that $m \geq n/8$.

The filler coordinates v_j are independent of the gadget. At scales $1 \leq t \leq 2$, their one-dimensional densities on $[-1, 1]$ are uniformly comparable. Consequently,

$$\frac{\text{Var}_{v_j \sim \mu_{K_s, TI_n}} [v_j]}{\text{Var}_{v_j \sim \mu_{K_s, I_n}} [v_j]} \geq c. \quad (\text{A.23})$$

Now apply the linear whitening map

$$L = C^{-1/2}, \quad \tilde{K} = LK_s, \quad \tilde{Q} = LI_nL^\top = C^{-1}. \quad (\text{A.24})$$

If $\tilde{X} = LX$, its scale-1 law is exactly $\mu_{\tilde{K}, \tilde{Q}}$, with mean zero and covariance I_n . At scale T , its law is exactly $\mu_{\tilde{K}, T\tilde{Q}}$. No translation has been applied, so the ambient Gaussian remains centered at the origin. Since C and the warmed covariance are diagonal, Equations (A.22) and (A.23) give

$$\text{Tr Cov}(\mu_{\tilde{K}, T\tilde{Q}}) \geq cr + c\lambda^2 \geq c(n + n^2\delta^4) \quad (\text{A.25})$$

whenever $\lambda \geq 1$.

If $\lambda < 1$, the filler coordinates still contribute at least $cr \geq cn$. Moreover, $m \geq n/8$ and Equation (A.13) imply

$$n^2\delta^4 = \frac{n^2\lambda^2}{m^2\kappa^2} \leq \frac{64}{\kappa^2},$$

which is a universal constant and can be absorbed into a constant multiple of n . Thus Equation (A.25) remains valid, after decreasing its universal constant.

For $1 \leq n < 8$, take a product of one-dimensional symmetric truncated Gaussians on $[-1, 1]$ and apply scalar whitening. Warming by a factor in $[1, 2]$ leaves each variance bounded below by a universal constant. Since $n + n^2\delta^4 \leq 8n$ in these dimensions, the lower bound follows after one last adjustment of the constant. This completes the proof of Theorem 7.

B Counterexamples to The Coordinate Moment Conjecture

We prove [Theorem 9](#). Throughout, $x' := (x_2, x_3, \dots, x_n)$ denotes all but the first coordinate of $x \in \mathbb{R}^n$. The counterexample density is a linear image of the restriction of a multivariate Gaussian to a symmetric convex body, defined next.

Assume $n \geq 7$. Define

$$g(x) := \begin{cases} (2\pi)^{-(n-1)/2} \exp(-\|x'\|^2/2) & \text{if } |x_1| \leq 1 \text{ and } \|x'\| \leq (n-2)^{1/2} - |x_1|, \\ 0 & \text{otherwise.} \end{cases}$$

Also, for $|x_1| \leq 1$, define

$$g_1(x_1) := \int_{\mathbb{R}^{n-1}} g(x) dx' = \int_0^{(n-2)^{1/2}-|x_1|} \frac{1}{2^{(n-1)/2-1} \Gamma((n-1)/2)} t^{n-2} e^{-t^2/2} dt$$

$$g_{\text{sq}}(x_1) := \int_{\mathbb{R}^{n-1}} \|x'\|^2 g(x) dx' = \int_0^{(n-2)^{1/2}-|x_1|} t^2 \cdot \frac{1}{2^{(n-1)/2-1} \Gamma((n-1)/2)} t^{n-2} e^{-t^2/2} dt$$

where the last equality uses polar coordinates in x' : if $Z \sim N(0, I_{n-1})$ and $T = \|Z\|$, then T has density

$$\frac{1}{2^{(n-1)/2-1} \Gamma((n-1)/2)} t^{n-2} e^{-t^2/2}.$$

Observe that $g_{\text{sq}}(x_1)/g_1(x_1)$ is a decreasing function of $|x_1|$.

Defining

$$G := \int_{\mathbb{R}^n} g(x) dx = \int_{-1}^1 g_1(x_1) dx_1$$

we have that $g(x)/G$ is a probability density function. It is even and log-concave, being the standard Gaussian density in x' restricted to the symmetric convex body $\{|x_1| \leq 1, \|x'\| \leq (n-2)^{1/2} - |x_1|\}$.

The next lemma concerns the moments of g/G and its isotropic linear image.

Lemma B.1. *Let $X \sim g/G$. Then $\mathbb{E}_{X \sim g/G} [X_1^2] = \Theta(1)$ and $\mathbb{E}_{X \sim g/G} [\|X'\|^2] = \Theta(n)$. With*

$$c_1 := \mathbb{E}_{X \sim g/G} [X_1^2]^{1/2}, \quad c' := \left(\frac{\mathbb{E}_{X \sim g/G} [\|X'\|^2]}{n-1} \right)^{1/2},$$

both of order $\Theta(1)$, the density

$$f^*(x) := c_1(c')^{n-1} (g(c_1 x_1, c' x')/G)$$

of $(X_1/c_1, X'/c')$ is even, isotropic and log-concave, and

$$\mathbb{E}_{x \sim f^*} [x_1^2 (\|x'\|^2 - (n-1))] = \frac{1}{c_1^2(c')^2} \mathbb{E}_{X \sim g/G} [(X_1^2 - \mathbb{E}_{X \sim g/G} [X_1^2]) \|X'\|^2].$$

The next lemma concerns the covariance of X_1^2 and $\|X'\|^2$ under g/G .

Lemma B.2. *Let $X \sim g/G$. Then*

$$\mathbb{E}_{X \sim g/G} [(X_1^2 - \mathbb{E}_{X \sim g/G} [X_1^2]) \|X'\|^2] = -\Theta(n^{1/2}).$$

Proof of Theorem 9. Let $f = f^\star$ from Lemma B.1; it is even, isotropic and log-concave. Since $\|x\|^2 - n = (\|x'\|^2 - (n-1)) + (x_1^2 - 1)$ and $\mathbb{E}_{x \sim f} [x_1^2] = 1$,

$$\mathbb{E}_{x \sim f} [x_1^2 (\|x\|^2 - n)] = \mathbb{E}_{x \sim f} [x_1^2 (\|x'\|^2 - (n-1))] + \mathbb{E}_{x \sim f} [x_1^4] - 1,$$

and $\mathbb{E}_{x \sim f} [x_1^4] \leq C_4$ by Proposition 2.23(iii). The first term on the right is $-\Theta(n^{1/2})/(c_1^2(c')^2) = -\Theta(n^{1/2})$ by Lemmas B.1 and B.2, so the left side has absolute value at least $cn^{1/2}$ for a universal $c > 0$ once n is sufficiently large. $\square$

Proof of Lemma B.1. We first show that $g_1(1)/g_1(0) > \Omega(1)$ and $g_{\text{sq}}(0)/g_1(0), g_{\text{sq}}(1)/g_1(1) = \Theta(n)$.

(i) To show $g_1(1)/g_1(0) > \Omega(1)$, it suffices to show

$$\frac{\int_{(n-2)^{1/2}-2}^{(n-2)^{1/2}-1} t^{n-2} e^{-t^2/2} dt}{\int_{(n-2)^{1/2}-1}^{(n-2)^{1/2}} t^{n-2} e^{-t^2/2} dt} > \Omega(1).$$

But we compute

$$\begin{aligned} \frac{(t-1)^{n-2} e^{-(t-1)^2/2}}{t^{n-2} e^{-t^2/2}} &= \left(1 - \frac{1}{t}\right)^{n-2} e^{t-(1/2)} \\ &= \exp\left(-\frac{n-2}{t} - O\left(\frac{n-2}{t^2}\right) + t - \frac{1}{2}\right) \\ &= \exp\left(-\frac{((n-2)^{1/2}-t)((n-2)^{1/2}+t)}{t} - O(1)\right) \\ &= \exp(-O(1)) \\ &= \Omega(1) \end{aligned}$$

for $t \in [(n-2)^{1/2}-1, (n-2)^{1/2}]$. Integrating yields

$$\int_{(n-2)^{1/2}-2}^{(n-2)^{1/2}-1} t^{n-2} e^{-t^2/2} dt > \Omega(1) \cdot \int_{(n-2)^{1/2}-1}^{(n-2)^{1/2}} t^{n-2} e^{-t^2/2} dt,$$

as desired.

(ii) To show $g_{\text{sq}}(0)/g_1(0) = O(n)$, we note the trivial bound

$$\frac{\int_0^{(n-2)^{1/2}} t^2 \cdot t^{n-2} e^{-t^2/2} dt}{\int_0^{(n-2)^{1/2}} t^{n-2} e^{-t^2/2} dt} \leq ((n-2)^{1/2})^2 < n.$$

(iii) To show $g_{\text{sq}}(1)/g_1(1) = \Omega(n)$, we note that

$$\frac{\int_0^{(n-2)^{1/2}-1} t^2 \cdot t^{n-2} e^{-t^2/2} dt}{\int_0^{(n-2)^{1/2}-1} t^{n-2} e^{-t^2/2} dt} \geq \left(\frac{\int_0^{(n-2)^{1/2}-1} t \cdot t^{n-2} e^{-t^2/2} dt}{\int_0^{(n-2)^{1/2}-1} t^{n-2} e^{-t^2/2} dt} \right)^2$$

so it suffices to show

$$\frac{\int_0^{(n-2)^{1/2}-1} t \cdot t^{n-2} e^{-t^2/2} dt}{\int_0^{(n-2)^{1/2}-1} t^{n-2} e^{-t^2/2} dt} = \Omega(n^{1/2}).$$

Now, we calculate

$$\log(t^{n-2}e^{-t^2/2}) = (n-2)\log t - \frac{t^2}{2}$$

is a concave function of t , and

$$\frac{d}{dt} \log(t^{n-2}e^{-t^2/2}) = \frac{n-2}{t} - t > \frac{n-2}{(n-2)^{1/2}} - t \geq 1$$

for $t \in [0, (n-2)^{1/2} - 1]$. It follows (by a comparison argument) that

$$\frac{\int_0^{(n-2)^{1/2}-1} t \cdot t^{n-2} e^{-t^2/2} dt}{\int_0^{(n-2)^{1/2}-1} t^{n-2} e^{-t^2/2} dt} \geq \frac{\int_{-\infty}^{(n-2)^{1/2}-1} t \cdot e^t dt}{\int_{-\infty}^{(n-2)^{1/2}-1} e^t dt} = (n-2)^{1/2} - 2 = \Omega(n^{1/2})$$

as desired.

The density of X_1 is g_1/G on $[-1, 1]$, and g_1 is nonincreasing in $|x_1|$ with $g_1(1)/g_1(0) > \Omega(1)$, so this density is within constant factors of the uniform density on $[-1, 1]$ and $\mathbb{E}_{X \sim g/G} [X_1^2] = \Theta(1)$. Conditionally on $X_1 = x_1$, the mean of $\|X'\|^2$ is $g_{\text{sq}}(x_1)/g_1(x_1)$, which lies between $g_{\text{sq}}(1)/g_1(1) = \Omega(n)$ and $g_{\text{sq}}(0)/g_1(0) = O(n)$; hence $\mathbb{E}_{X \sim g/G} [\|X'\|^2] = \Theta(n)$, and $c_1, c' = \Theta(1)$. The law g/G is invariant under $x_1 \mapsto -x_1$ and under every rotation of x' , so its mean is 0 and its covariance is $\text{diag}(c_1^2, (c')^2 I_{n-1})$; therefore f^* , the density of $(X_1/c_1, X'/c')$, has mean 0 and covariance I_n , and it is even and log-concave as a linear image of g/G . Finally,

$$\mathbb{E}_{x \sim f^*} [x_1^2 (\|x'\|^2 - (n-1))] = \frac{1}{c_1^2 (c')^2} \mathbb{E}_{X \sim g/G} [X_1^2 (\|X'\|^2 - (c')^2 (n-1))],$$

and $(c')^2 (n-1) = \mathbb{E}_{X \sim g/G} [\|X'\|^2]$, so subtracting the constant $\mathbb{E}_{X \sim g/G} [X_1^2] \mathbb{E}_{X \sim g/G} [\|X'\|^2]$ in both forms gives the stated identity. $\square$

The proof of [Lemma B.2](#) uses the following derivative bound for truncated second moments of a monotone log-concave density.

Lemma B.3. *Let $L_0 > 0$, and let $f : [0, +\infty) \rightarrow \mathbb{R}$ be a log-concave probability density function which is continuous, such that f is nondecreasing on $[0, L_0]$ and $\int_0^{L_0} f(t) dt > 0$. Then*

$$\frac{1}{2} L_0 \leq \frac{d}{dL} \frac{\int_0^L t^2 \cdot f(t) dt}{\int_0^L f(t) dt} \Big|_{L=L_0} \leq 2L_0.$$

Proof. We compute

$$\begin{aligned} \frac{d}{dL} \frac{\int_0^L t^2 \cdot f(t) dt}{\int_0^L f(t) dt} &= \frac{L^2 \cdot f(L)}{\int_0^L f(t) dt} - \frac{f(L) \int_0^L t^2 \cdot f(t) dt}{\left(\int_0^L f(t) dt\right)^2} \\ &= \frac{f(L) \int_0^L (L^2 - t^2) \cdot f(t) dt}{\left(\int_0^L f(t) dt\right)^2} \end{aligned}$$

and thus

$$\frac{d}{dL} \frac{\int_0^L t^2 \cdot f(t) dt}{\int_0^L f(t) dt} \Big|_{L=L_0} = \frac{f(L_0) \int_0^{L_0} (L_0^2 - t^2) \cdot f(t) dt}{\left(\int_0^{L_0} f(t) dt\right)^2}.$$

For the lower bound, we have

$$\begin{aligned}
\left(\int_0^{L_0} f(t) dt \right)^2 &= 2 \int_0^{L_0} \int_t^{L_0} f(t) f(s) ds dt \\
&\leq 2 \int_0^{L_0} \int_t^{L_0} f(t) f(L_0) ds dt \quad (f \text{ is nondecreasing on } [0, L_0]) \\
&= 2f(L_0) \int_0^{L_0} (L_0 - t) \cdot f(t) dt \\
&\leq 2f(L_0) \int_0^{L_0} (L_0 - t) \frac{L_0 + t}{L_0} \cdot f(t) dt \\
&= \frac{2}{L_0} f(L_0) \int_0^{L_0} (L_0^2 - t^2) \cdot f(t) dt
\end{aligned}$$

which implies

$$\left. \frac{d}{dL} \frac{\int_0^L t^2 \cdot f(t) dt}{\int_0^L f(t) dt} \right|_{L=L_0} = \frac{f(L_0) \int_0^{L_0} (L_0^2 - t^2) \cdot f(t) dt}{\left(\int_0^{L_0} f(t) dt \right)^2} \geq \frac{1}{2} L_0,$$

as desired.

For the upper bound, we have

$$\begin{aligned}
\left(\int_0^{L_0} f(t) dt \right)^2 &= \int_0^{L_0} \int_0^{L_0} f(t) f(s) ds dt \\
&\geq \iint_{t,s \in [0, L_0], t+s \geq L_0} f(t) f(s) ds dt \\
&\geq \iint_{t,s \in [0, L_0], t+s \geq L_0} f(L_0) f(t+s-L_0) ds dt \quad (\text{log-concavity of } f) \\
&= f(L_0) \int_0^{L_0} (L_0 - t) \cdot f(t) dt \\
&\geq f(L_0) \int_0^{L_0} (L_0 - t) \frac{L_0 + t}{2L_0} \cdot f(t) dt \\
&= \frac{1}{2L_0} f(L_0) \int_0^{L_0} (L_0^2 - t^2) \cdot f(t) dt
\end{aligned}$$

which implies

$$\left. \frac{d}{dL} \frac{\int_0^L t^2 \cdot f(t) dt}{\int_0^L f(t) dt} \right|_{L=L_0} = \frac{f(L_0) \int_0^{L_0} (L_0^2 - t^2) \cdot f(t) dt}{\left(\int_0^{L_0} f(t) dt \right)^2} \leq 2L_0,$$

as desired. □

Proof of Lemma B.2. Let us rewrite the left-hand side as

$$\begin{aligned}
&\int_{-1}^1 (x_1^2 - \mathbb{E}_{x \sim g(x)/G} [x_1^2]) \cdot \frac{g_{\text{sq}}(x_1)}{g_1(x_1)} \cdot \frac{g_1(x_1)}{G} dx_1 \\
&= 2 \int_0^1 (x_1^2 - \mathbb{E}_{x \sim g(x)/G} [x_1^2]) \cdot \frac{g_{\text{sq}}(x_1)}{g_1(x_1)} \cdot \frac{g_1(x_1)}{G} dx_1 \\
&= -2 \int_0^1 \left(\frac{d}{dx_1} \frac{g_{\text{sq}}(x_1)}{g_1(x_1)} \right) \cdot \left(\int_0^{x_1} (s^2 - \mathbb{E}_{x \sim g(x)/G} [x_1^2]) \cdot \frac{g_1(s)}{G} ds \right) dx_1.
\end{aligned}$$

Since $t^{n-2}e^{-t^2/2}$ is log-concave and nondecreasing on $t \in [0, (n-2)^{1/2}]$, [Lemma B.3](#) gives

$$\frac{d}{dx_1} \frac{g_{\text{sq}}(x_1)}{g_1(x_1)} = -\Theta((n-2)^{1/2} - x_1) = -\Theta(n^{1/2})$$

for $x_1 \in [0, 1]$. Since

$$\int_0^{x_1} (s^2 - \mathbb{E}_{x \sim g(x)/G} [x_1^2]) \cdot \frac{g_1(s)}{G} ds \leq 0$$

for all $x_1 \in [0, 1]$, it suffices to show

$$\int_0^1 \int_0^{x_1} (s^2 - \mathbb{E}_{x \sim g(x)/G} [x_1^2]) \cdot \frac{g_1(s)}{G} ds dx_1 = -\Theta(1).$$

We prove the upper and lower bounds separately.

(i) The bound

$$\left| \int_0^1 \int_0^{x_1} (s^2 - \mathbb{E}_{x \sim g(x)/G} [x_1^2]) \cdot \frac{g_1(s)}{G} ds dx_1 \right| = O(1)$$

follows from the fact that $\int_0^1 g_1(s)/G ds = 1/2 = O(1)$.

(ii) To show

$$\int_0^1 \int_0^{x_1} (s^2 - \mathbb{E}_{x \sim g(x)/G} [x_1^2]) \cdot \frac{g_1(s)}{G} ds dx_1 = -\Omega(1),$$

it suffices to show that there exist $d_1, d_2 = \Omega(1)$ such that

$$\int_0^{x_1} (s^2 - \mathbb{E}_{x \sim g(x)/G} [x_1^2]) \cdot \frac{g_1(s)}{G} ds = -\Omega(1) \quad (\text{B.1})$$

whenever $x_1 \in [d_1, d_2]$.

We know that $g_1(s)/G = \Omega(1)$ for each $s \in [0, 1]$, due to the bound $g_1(1)/g_1(0) > \Omega(1)$ we proved earlier. Hence

$$\mathbb{E}_{x \sim g(x)/G} [x_1^2] = 2 \cdot \int_0^1 s^2 \cdot \frac{g_1(s)}{G} ds = \Omega(1),$$

so if we choose

$$d_1 := \frac{1}{2} \mathbb{E}_{x \sim g(x)/G} [x_1^2]^{1/2}, \quad d_2 := \mathbb{E}_{x \sim g(x)/G} [x_1^2]^{1/2}$$

we will have $d_1, d_2 = \Omega(1)$.

It remains to show Equation (B.1). We only need to check the case when $x_1 = d_1$, i.e. we only need to verify

$$\int_0^{d_1} (s^2 - \mathbb{E}_{x \sim g(x)/G} [x_1^2]) \cdot \frac{g_1(s)}{G} ds = -\Omega(1),$$

which follows from

$$\begin{aligned} \int_0^{d_1} (s^2 - \mathbb{E}_{x \sim g(x)/G} [x_1^2]) \cdot \frac{g_1(s)}{G} ds &\leq \int_0^{d_1} (d_1^2 - \mathbb{E}_{x \sim g(x)/G} [x_1^2]) \cdot \frac{g_1(s)}{G} ds \\ &= -\frac{3}{4} \mathbb{E}_{x \sim g(x)/G} [x_1^2] \int_0^{d_1} \frac{g_1(s)}{G} ds \\ &= -\Omega(1) \end{aligned}$$

using $g_1(s)/G = \Omega(1)$ and $\mathbb{E}_{x \sim g(x)/G} [x_1^2] = \Omega(1)$, proved above.

The two bounds together prove the lemma. $\square$

C Asymmetric Counterexample to the Ellipsoid-Growth Conjecture

Here we provide an asymmetric counterexample to the Ellipsoid-Growth Conjecture, which is Conjecture 3.5 of Jia–Laddha–Lee–Vempala [JLLV26], and stated in Conjecture 3.2. Specifically, to prove Theorem 8, we give an explicit family of convex bodies $K \subset \mathbb{R}^n$ and ellipsoids $E \subset \mathbb{R}^n$ for which $K \cap E$ has mean zero and covariance I_n , but the covariance trace of $K \cap (1 + \delta)E$ is at least $cn^2\delta^2$. Here $(1 + \delta)E = \{(1 + \delta)z : z \in E\}$ is dilation about the origin. Our family works for every $n^{-1/2} \leq \delta \leq 1$, including $n^{-1/4}$.

The proof keeps track of the affine centering step, which does not commute with dilation, and gives explicit estimates. For a convex body A , all means and covariance matrices below refer to the uniform probability measure on A . We use the standard origin-dilation notation

$$sE := \{sz : z \in E\}, \quad s \geq 1.$$

The distinction between origin dilation and dilation about the center of E matters because our final isotropizing map is affine.

Conjecture 3.5 of [JLLV26] asserts that if $K \cap E$ is 2-isotropic, then $K \cap 2E$ is well-rounded. The covariance part of well-roundedness is $\text{tr Cov} = \tilde{O}(n)$. We prove the following stronger negation, using mean-zero, covariance- I_n isotropy rather than merely 2-isotropy.

The next theorem concerns the quantitative counterexample.

Theorem 12. *There are absolute constants $c, C > 0$ such that, for every sufficiently large n , one can construct a convex body $K \subset \mathbb{R}^n$ and an ellipsoid $E \subset \mathbb{R}^n$ with the following properties.*

1. *The center of E lies in $\text{int } K$.*
2. *$K \cap E$ has mean zero and covariance I_n .*
3. *Simultaneously for all $0 < \delta \leq 1$ satisfying $(n - 1)\delta \geq C$,*

$$\text{tr Cov}(K \cap (1 + \delta)E) \geq c(n - 1)^2\delta^2.$$

One may take, for example, $C = 32$ and $c = 1/98304$.

C.1 The raw body and its thin direction

Write $n = d + 1$ and points of $\mathbb{R}^n$ as $(x, y) \in \mathbb{R} \times \mathbb{R}^d$. Define

$$f(x) := \begin{cases} 1 + x/2, & -2 \leq x \leq 0, \\ 1, & 0 \leq x \leq 4, \end{cases}$$

and

$$\begin{aligned} K_0 &:= \{(x, y) : -2 \leq x \leq 4, \|y\|_2 \leq f(x)\}, \\ E_0 &:= \{(x, y) : (x + \tfrac{1}{2})^2 + \|y\|_2^2 \leq \tfrac{5}{4}\}. \end{aligned} \tag{C.1}$$

The function f is nonnegative and concave: its slope drops from $1/2$ to 0 at the origin. Consequently K_0 is convex. Indeed, the triangle inequality and concavity of f show that the segment between any two points of K_0 remains in K_0 . The center $(-1/2, 0)$ of E_0 is strictly inside K_0 , since $f(-1/2) = 3/4$.

Let $L_0 := K_0 \cap E_0$. Let $g(x)$ be the radius of its x -section. The axial support of E_0 is

$$\alpha := -\frac{1 + 5^{1/2}}{2} \leq x \leq \beta := \frac{5^{1/2} - 1}{2},$$

and the section radius of E_0 is

$$e_0(x) = (1 - x - x^2)^{1/2}.$$

For $x \leq 0$, direct subtraction gives

$$e_0(x)^2 - (1 + x/2)^2 = -x \left(2 + \frac{5}{4}x \right). \quad (\text{C.2})$$

It follows that the section function is exactly

$$g(x) = \begin{cases} (1 - x - x^2)^{1/2}, & \alpha \leq x \leq -8/5, \\ 1 + x/2, & -8/5 \leq x \leq 0, \\ (1 - x - x^2)^{1/2}, & 0 \leq x \leq \beta, \\ 0, & \text{otherwise.} \end{cases} \quad (\text{C.3})$$

Thus there is a small secondary cap near $x = \alpha$, but its section radius is at most $1/5$. More importantly, the unique section of radius 1 is at $x = 0$, and the section radius decreases linearly to first order on both sides.

Let X denote the first coordinate of a uniform point of L_0 . After the common factor $\text{vol}(B_2^d)$ is removed, its density is

$$\frac{g(x)^d}{Z_d}, \quad Z_d := \int_{\mathbb{R}} g(x)^d dx. \quad (\text{C.4})$$

The next lemma concerns the inner moments.

Lemma C.1. *For $d \geq 3$, there are absolute constants $c_i, C_i > 0$ such that*

$$\frac{c_0}{d} \leq Z_d \leq \frac{C_0}{d}, \quad |\mathbb{E}_{X \sim g^d/Z_d}[X]| \leq \frac{C_1}{d}, \quad \frac{c_2}{d^2} \leq \text{Var}_{X \sim g^d/Z_d}[X] \leq \frac{C_2}{d^2}.$$

Moreover, by rotational symmetry,

$$\text{Cov}(L_0) = \text{diag}(\sigma_x^2, \tau^2 I_d), \quad \sigma_x^2 = \Theta(d^{-2}), \quad \tau^2 = \Theta(d^{-1}).$$

The explicit bounds $\sigma_x^2 \leq 64/d^2$ and $\sigma_x^2 \geq 9/(2048d^2)$ are valid.

Proof. On the negative half-line, $g(x) \leq 1 + x/2$, while for $x \geq 0$,

$$g(x)^d = (1 - x - x^2)^{d/2} \leq e^{-dx/2}.$$

Hence

$$Z_d \leq \int_{-2}^0 (1 + x/2)^d dx + \int_0^\infty e^{-dx/2} dx \leq \frac{4}{d}.$$

On $[-1/d, 0]$, one has $g(x) = 1 + x/2$ and $g(x)^d \geq 1/2$, so $Z_d \geq 1/(2d)$.

For $k = 1, 2$, the negative-side moments satisfy the exact beta-integral bound

$$\begin{aligned} \int_{-2}^0 |x|^k g(x)^d dx &\leq 2^{k+1} \int_0^1 (1 - u)^k u^d du \\ &= \frac{2^{k+1} k!}{(d+1)(d+2) \cdots (d+k+1)}. \end{aligned} \quad (\text{C.5})$$

The positive-side moments obey

$$\int_0^\infty x^k e^{-dx/2} dx = k! \left(\frac{2}{d}\right)^{k+1}.$$

Dividing by the lower bound for Z_d gives $\mathbb{E}_{X \sim g^d/Z_d} [|X|] = O(d^{-1})$ and $\mathbb{E}_{X \sim g^d/Z_d} [X^2] \leq 64/d^2$, hence the desired upper variance bound.

For the lower bound, take

$$I_- = [-2/d, -1/d], \quad I_+ = [1/(2d), 1/d].$$

For $d \geq 3$, the unnormalized density g^d is bounded below by $1/4$ on both intervals. Since $Z_d \leq 4/d$,

$$\mathbb{P}(X \in I_-) \geq \frac{1}{16}, \quad \mathbb{P}(X \in I_+) \geq \frac{1}{32}.$$

If X' is an independent copy of X , then $\text{Var}_{X \sim g^d/Z_d} [X] = \frac{1}{2} \mathbb{E}_{X, X' \sim g^d/Z_d} [(X - X')^2]$. The two intervals are separated by at least $3/(2d)$, and therefore

$$\text{Var}_{X \sim g^d/Z_d} [X] \geq \frac{1}{16} \frac{1}{32} \left(\frac{3}{2d}\right)^2 = \frac{9}{2048d^2}.$$

Finally, conditional on $X = x$, the variable y is uniform in a d -ball of radius $g(x)$, so

$$\mathbb{E}_{y \sim \text{Unif}(g(x)B_2^d)} [y_i^2 \mid X = x] = \frac{g(x)^2}{d+2}.$$

The upper bound $\tau^2 \leq 1/(d+2)$ is immediate. A central interval of probability bounded below has $g(x)^2$ bounded below, giving $\tau^2 \geq c/(d+2)$. Rotational symmetry eliminates all mixed terms. $\square$

Let

$$m := (\mu, 0), \quad \mu := \mathbb{E}_{X \sim g^d/Z_d} [X], \quad \Sigma := \text{Cov}(L_0),$$

and define the affine map

$$F(z) := \Sigma^{-1/2}(z - m), \quad K := F(K_0), \quad E := F(E_0). \quad (\text{C.6})$$

Since F is bijective,

$$K \cap E = F(L_0).$$

The uniform law is carried to the uniform law under an invertible affine map, so $K \cap E$ has mean zero and covariance I_n exactly. Also, the center of E belongs to $\text{int } K$, because $(-1/2, 0) \in \text{int } K_0$.

C.2 The enlarged ellipsoid creates a plateau

Fix $s = 1 + \delta \in (1, 2]$. This is the step at which the translation in F must be handled carefully. For every $z \in \mathbb{R}^n$,

$$F^{-1}(sF(z)) = m + s(z - m) = sz + (1 - s)m.$$

Consequently

$$F^{-1}(sE) = sE_0 + (1 - s)m =: E_s, \quad (\text{C.7})$$

and hence

$$K \cap sE = F(D_s), \quad D_s := K_0 \cap E_s. \quad (\text{C.8})$$

The set E_s is the Euclidean ball with center $(-s/2 - \delta\mu, 0)$ and radius $s5^{1/2}/2$. Its section radius is

$$e_s(x)^2 = \frac{5s^2}{4} - (x + s/2 + \delta\mu)^2. \quad (\text{C.9})$$

On the cylindrical part $x \geq 0$ of K_0 , the complete unit d -ball is present precisely while $e_s(x) \geq 1$. The positive endpoint of this interval is

$$x_* = -\frac{s}{2} - \delta\mu + A_s, \quad A_s := \left(\frac{5s^2}{4} - 1 \right)^{1/2}. \quad (\text{C.10})$$

The other root is negative: since $A_s \geq s/2$ and $\mu \geq \alpha$, it is at most $-s - \delta\mu \leq -1 + (-1 - \alpha)\delta < 0$. To estimate x_* , define

$$q(s) := A_s - s/2.$$

A direct differentiation gives

$$\frac{3}{4} \leq q'(s) \leq 2, \quad 1 \leq s \leq 2,$$

and $q(1) = 0$. Since the support of X is $[\alpha, \beta]$, its mean also lies in that interval. Therefore

$$\frac{\delta}{8} \leq x_* = q(1 + \delta) - \delta\mu \leq 4\delta, \quad 0 < \delta \leq 1. \quad (\text{C.11})$$

Our choice of the upper endpoint 4 in the definition of K_0 guarantees that this whole interval lies in the cylindrical part. We have thus found an exact radius-one cylindrical plateau

$$\{(x, y) : 0 \leq x \leq x_*, \|y\|_2 \leq 1\} \subset D_s. \quad (\text{C.12})$$

The next lemma concerns the outer axial variance.

Lemma C.2. *If $d\delta \geq 32$, then the first-coordinate variance of the uniform law on D_s satisfies*

$$\text{Var}_{D_s}[x] \geq \frac{\delta^2}{1536}.$$

Proof. After the common factor $\text{vol}(B_2^d)$ is removed, the plateau has mass exactly x_* . All mass with $x < 0$ is at most

$$\int_{-2}^0 (1 + x/2)^d dx = \frac{2}{d+1}.$$

For the positive tail, write $x = x_* + v$, $v \geq 0$. From Equation (C.9)–(C.10),

$$e_s(x)^2 = 1 - 2A_s v - v^2.$$

On the support of this tail, the D_s -section radius is at most $e_s(x)$. Since $A_s \geq 1/2$, its section volume, divided by $\text{vol}(B_2^d)$, is at most $e^{-dA_s v}$. Hence the positive-tail mass is at most

$$\int_0^\infty e^{-dA_s v} dv \leq \frac{2}{d}.$$

Thus the total off-plateau mass is at most $4/d$. By the lower endpoint bound $x_* \geq \delta/8$ in Equation (C.11), $d\delta \geq 32$ implies $x_* \geq 4/d$, so the plateau has probability at least $1/2$ under the uniform law on D_s .

Conditioned on the plateau, the first coordinate is exactly uniform on $[0, x_*]$. Applying the law of total variance to the indicator of the plateau gives

$$\text{Var}_{D_s}[x] \geq \frac{1}{2} \text{Var}_{x \sim \text{Unif}[0, x_*]}[x] = \frac{x_*^2}{24} \geq \frac{\delta^2}{1536}.$$

□

Proof of Theorem 12. The first two assertions, namely mean zero and covariance I_n for $K \cap E$, follow from the affine normalization in Equation (C.6). Rotational symmetry gives

$$\Sigma = \text{diag}(\sigma_x^2, \tau^2 I_d).$$

By Equation (C.8), the first coordinate after applying F is $(x - \mu)/\sigma_x$. Hence its variance under the uniform law on $K \cap sE$ is exactly

$$\frac{\text{Var}_{D_s}[x]}{\sigma_x^2}.$$

Using Lemma C.1 and Lemma C.2,

$$\text{tr Cov}(K \cap sE) \geq \frac{\delta^2/1536}{64/d^2} = \frac{d^2\delta^2}{98304}.$$

The construction of K and E depends on n but not on δ , so the estimate holds simultaneously throughout the stated range. $\square$

The next corollary concerns the scales of failure.

Corollary C.3. *The choice $\delta = 1$ gives $\text{tr Cov}(K \cap 2E) = \Omega(n^2)$ and disproves Conjecture 3.5. Moreover,*

$$\delta = n^{-1/4} \implies \text{tr Cov}(K \cap (1 + \delta)E) = \Omega(n^{3/2}).$$

For every fixed $0 < \varepsilon < 1/2$, the smaller choice $\delta = n^{-1/2+\varepsilon}$ gives trace $\Omega(n^{1+2\varepsilon})$. Finally,

$$\delta(n) = n^{-1/2} \exp((\log n)^{1/2}) = n^{-1/2+o(1)}$$

gives

$$\text{tr Cov}(K \cap (1 + \delta)E) \geq n \exp(2(\log n)^{1/2})/C,$$

which is not $\tilde{O}(n)$.