CS395T: Continuous Algorithms, Part XII
Ball walk

Kevin Tian

1 Sampling from a convex body

In this lecture, we begin our study of sampling from continuous distributions (i.e. those which have absolutely continuous densities with respect to the Lebesgue measure on $\mathbb{R}^d$). To introduce key techniques for analyzing continuous random walks, we use the fundamental problem of producing an approximately-uniform sample from a convex body $\mathcal{K} \subseteq \mathbb{R}^d$ as a running example throughout these notes. Specifically we study uniform sampling from convex bodies under *membership oracle* access, i.e. we assume access to an oracle which takes as input $x \in \mathbb{R}^d$, and tells us whether $x \in \mathcal{K}$.

Sampling from a convex body is a prominent example of the more general problem of sampling from a logconcave density on $\mathbb{R}^d$ (see Section 4, Part I for the relevant definitions). When $\mu \propto \exp(-f)$ is a logconcave density on $\mathbb{R}^d$, so $f : \mathbb{R}^d \to \mathbb{R} \cup \{\infty\}$ is convex, the assumption of value oracle access to f grants us membership oracle access in the special case when $f = \chi_{\mathcal{K}}$ is the 0-∞ characteristic function of $\mathcal{K}$. We will discuss continuous sampling settings beyond uniform densities on convex bodies in the following lectures, which build upon the tools introduced here.

One reason that we base our exploration of continuous sampling in this uniform sampling problem is historical: uniform sampling from a convex body has been a longstanding testbed for the development of new techniques for analyzing continuous Markov chains. Indeed, the fact that polynomial-time sampling was possible even in this special case was a very surprising breakthrough when it was first established by [DFK91]. For example, previous results by [Ele86, BF87] had established that it is impossible to approximate the volume up to an exponential factor in the dimension in deterministic polynomial time, and [DF88] showed that it is #P-hard to exactly compute the volume. On the other hand, by crucially leveraging randomness, [DFK91] showed how to use their convex body sampler to estimate volume up to a $1 \pm \epsilon$ multiplicative factor with high probability in time $\mathrm{poly}(d, \epsilon^{-1})$. To do so, one can follow a standard reduction from counting to sampling. In particular, assuming that $\mathbb{B}(1) \subseteq \mathcal{K} \subseteq \mathbb{B}(d)$, we can estimate $\mathrm{Vol}(\mathcal{K})$ as the product of volume ratios for sets $\mathcal{K}_i := \mathcal{K} \cap \mathbb{B}(2^{\frac{i}{d}})$ over a sequence of indices $0 \le i \le d \log_2 d$, i.e.

$$\mathrm{Vol}(\mathcal{K}) = \mathrm{Vol}(\mathcal{K}_0) \cdot \prod_{0 \le i \le \lfloor d \log_2 d \rfloor} \frac{\mathrm{Vol}(\mathcal{K}_{i+1})}{\mathrm{Vol}(\mathcal{K}_i)} = \mathrm{Vol}(\mathbb{B}(1)) \cdot \prod_{0 \le i \le \lfloor d \log_2 d \rfloor} \frac{\mathrm{Vol}(\mathcal{K}_{i+1})}{\mathrm{Vol}(\mathcal{K}_i)}.$$

Here we used that our containment assumptions imply $\mathcal{K}_0 = \mathbb{B}(1)$ and $\mathcal{K}_{\lfloor d \log_2 d \rfloor + 1} = \mathcal{K}$. Moreover, each of the volume ratios we need to estimate, of the form

$$\frac{\mathrm{Vol}(\mathcal{K}_{i+1})}{\mathrm{Vol}(\mathcal{K}_i)},$$

are $\Theta(1)$, so they can be efficiently estimated to multiplicative error via random sampling, simply by checking what proportion of random samples from $\mathcal{K}_{i+1}$ also fall in $\mathcal{K}_i$. Finally, as we will see in a future lecture, the assumption that $\mathbb{B}(1) \subseteq \mathcal{K} \subseteq \mathbb{B}(d)$ (up to a negligible fraction of the body) can also be achieved via random sampling, by estimating an affine transformation which makes the body isotropic,[1] and applying concentration inequalities for logconcave densities.

The original algorithm of [DFK91], while polynomial-time, used $\approx d^{19}$ membership oracle queries to compute a single high-accuracy sample. Since this seminal work, a long line of simplifications of the [DFK91] framework, as well as new algorithmic tools, have considerably improved the query

---

[1] We say that a density π on $\mathbb{R}^d$ is *isotropic* if $\mathbb{E}_{x \sim \pi}[x] = \mathbb{0}_d$ and $\mathbb{E}_{x \sim \pi}[xx^\top] = \mathbf{I}_d$.

complexity of sampling from a convex body [AK91, LS92, LS93, KLS97, LV06a, LV06b, LV06c, LV07, CV18, JLLV21]. The goal of this lecture and the few following it are to give the tools needed to obtain an improved high-accuracy convex body sampler, which leverages the aforementioned works and a few additional tools to use $\approx d^3$ membership oracle queries.

In Sections 2 and 3, we begin by giving general-purpose tools for analyzing the convergence of random walks in continuous space, which are not specialized to any particular random walk. We specialize Sections 4 and 5 to the analysis of a particular Markov chain, the *ball walk*, which is a basic building block for all the sampling algorithms we develop towards sampling convex bodies.

2 Mixing in the continuous setting

2.1 Preliminaries

Throughout these notes, for any density π over an event space $\mathcal{E}$, we overload $\pi(S)$ to mean $\Pr_{s \sim \pi}[s \in S]$ for any $S \subseteq \mathcal{E}$, e.g. if π is a density over $\mathbb{R}^d$, then $\pi(S)$ denotes the probability a sample from π falls in $S \subseteq \mathbb{R}^d$. We let $\pi^\star$ denote a target stationary density on $\mathbb{R}^d$, which we assume is absolutely continuous with respect to the Lebesgue measure.[2] We consider Markov chains parameterized by *transition distributions* $\{\mathcal{T}_x\}_{x \in \mathbb{R}^d}$, which are also densities on $\mathbb{R}^d$, so

$$\int \mathcal{T}_x(y)\mathrm{d}y = 1 \text{ for all } x \in \mathbb{R}^d. \tag{1}$$

We assume each $\mathcal{T}_x$ is *lazy*, which means that it puts a point mass of at least $\frac{1}{2}$ on x, and that its density on $\mathbb{R}^d \setminus \{x\}$ is absolutely continuous with respect to the Lebesgue measure. This is the analog of Eq. (7), Part XI, and as was the case there, laziness facilitates simpler proofs (but does not meaningfully affect implementation). The $\{\mathcal{T}_x\}_{x \in \mathbb{R}^d}$ induce a random walk via the update

$$x_{k+1} \sim \mathcal{T}_{x_k}, \text{ for all } k \geq 0.$$

We denote the distribution that the random walk is initialized with by π_0, and the distribution of the k^{th} iterate of the walk by π_k. By definition of the random walk, for all $k \geq 0$,

$$\pi_{k+1}(x) = \int \mathcal{T}_y(x)\pi_k(y)\mathrm{d}y. \tag{2}$$

We additionally assume that the $\{\mathcal{T}_x\}_{x \in \mathbb{R}^d}$ are *reversible* with respect to $\pi^\star$, analogously to Definition 3, Part XI. In the continuous setting, this means that for all $(x, y) \in \mathbb{R}^d \times \mathbb{R}^d$,

$$\pi^\star(x)\mathcal{T}_x(y) = \pi^\star(y)\mathcal{T}_y(x). \tag{3}$$

Under the assumption (3), it is simple to verify that $\pi^\star$ is indeed a stationary distribution for the random walk with updates (2). Indeed, supposing that $\pi_k = \pi^\star$ for some k, we have

$$\pi_{k+1}(x) = \int \mathcal{T}_y(x)\pi^\star(y)\mathrm{d}y = \int \mathcal{T}_x(y)\pi^\star(x)\mathrm{d}y = \pi^\star(x),$$

where the second equality used (3) and the third used (1). Moreover, given a set of *proposal distributions* $\{\mathcal{P}_x\}_{x \in \mathbb{R}^d}$ and a target stationary distribution $\pi^\star$, such that (3) is not necessarily satisfied with respect to the $\{\mathcal{P}_x\}_{x \in \mathbb{R}^d}$ one can apply the Metropolis-Hastings correction (Lemma 3, Part XI) to obtain transition distributions which do satisfy (3), by letting

$$\mathcal{T}_x(y) := \mathcal{P}_x(y) \min\left(1, \frac{\pi^\star(y)\mathcal{P}_y(x)}{\pi^\star(x)\mathcal{P}_x(y)}\right) \text{ for all } y \in \mathbb{R}^d, \ y \neq x. \tag{4}$$

As was the case in Part XI, the correction (4) has a simple interpretation as applying a filter which either accepts the proposal distribution or does not move, with an acceptance probability proportional to how "reversible" the proposals are for the relevant pair of points. Finally, when the transition distributions $\{\mathcal{T}_x\}_{x \in \mathbb{R}^d}$ of a Markov chain on $\mathbb{R}^d$ are clear from context, we use $\mathcal{T}\pi$ to denote the density resulting from applying one step of the Markov chain to a point drawn from π.

---

[2]In Sections 4 and 5, we focus on the specific case where $\pi^\star(x) \propto \mathbf{1}_{x \in \mathcal{K}}$, for convex $\mathcal{K} \subseteq \mathbb{R}^d$.

2.2 Mixing from conductance

In this section, we introduce a continuous definition of the *conductance* of a random walk, which is used to analyze mixing, and is patterned off of the discrete analog in Definition 6, Part XI.

Definition 1 (Conductance). *For a random walk on $\mathbb{R}^d$ with stationary distribution $\pi^\star$ and lazy, reversible transition distributions $\{\mathcal{T}_x\}_{x \in \mathbb{R}^d}$, we denote the flow out of a set $S \subseteq \mathbb{R}^d$ by*

$$Q(S) := \int_{x \in S} \mathcal{T}_x(S^c) \pi^\star(x) \mathrm{d}x = \int_{x \in S^c} \mathcal{T}_x(S) \pi^\star(x) \mathrm{d}x, \tag{5}$$

where $S^c := \mathbb{R}^d \setminus S$. We define the conductance *of the random walk by*

$$\Phi := \inf_{0 < \pi^\star(S) \le \frac{1}{2}} \frac{Q(S)}{\pi^\star(S)}. \tag{6}$$

We note that the equality (5) holds because of the calculation

$$\int_{x \in S} \mathcal{T}_x(S^c) \pi^\star(x) \mathrm{d}x = \int_{x \in S} \int_{y \in S^c} \mathcal{T}_x(y) \pi^\star(x) \mathrm{d}y \mathrm{d}x$$
$$= \int_{x \in S} \int_{y \in S^c} \mathcal{T}_y(x) \pi^\star(y) \mathrm{d}y \mathrm{d}x = \int_{y \in S^c} \mathcal{T}_y(S) \pi^\star(y) \mathrm{d}y,$$

where we applied reversibility and Fubini's theorem. As in Part XI, the quantity $\frac{Q(S)}{\pi^\star(S)}$ has a natural interpretation as the probability that one step of a random walk initialized randomly in S leaves the set. The intuition is again that if the conductance Φ is large, then no set S is a bottleneck for the random walk, and hence the walk will rapidly mix. In order to quantify this intuition, we require a few new definitions specialized to the continuous setting.

Definition 2 (Warmness). *For densities $\pi_0, \pi^\star$ on $\mathbb{R}^d$, we say π_0 is β-warm with respect to $\pi^\star$ if*

$$\frac{\pi_0(x)}{\pi^\star(x)} \le \beta \text{ for all } x \in \mathbb{R}^d.$$

Definition 3 (Lovasz-Simonovits curve). *For a random walk on $\mathbb{R}^d$ with stationary distribution $\pi^\star$ and lazy, reversible transition distributions $\{\mathcal{T}_x\}_{x \in \mathbb{R}^d}$, initialized at π_0 and with density π_k after k steps for all $k \in \mathbb{N}$, we define the* Lovasz-Simonovits curve *$h_k : [0, 1] \to [0, 1]$ for all $k \in \mathbb{N}$ by*

$$h_k(\tau) := \sup_{g \in G_\tau} \int g(x) \pi_k(x) \mathrm{d}x - \tau, \text{ where } G_\tau := \left\{ g : \mathbb{R}^d \to [0, 1] \mid \int g(x) \pi^\star(x) \mathrm{d}x = \tau \right\}. \tag{7}$$

Definition 2 is fairly straightforward, and it implies in particular that for any $S \subseteq \mathbb{R}^d$ that $\pi_0(S) \le \beta \pi^\star(x)$. Definition 3 is somewhat more complicated, but we briefly interpret it here. First, it is straightforward to check that if π_0 is absolutely continuous with respect to $\pi^\star$, then all π_k for $k \in \mathbb{N}$ are too, so the supremum over G_τ in (7) is achieved by greedily placing mass on points sorted by the value of $\frac{\pi_k}{\pi}$. In other words, the supremum is achieved by a 0-1 set indicator function, $g(x) = \mathbb{1}_{x \in S}$, for some $S \subseteq \mathbb{R}^d$. Rewriting (7) where the supremum is taken over $S \subseteq \mathbb{R}^d$,

$$h_k(\tau) = \sup_{\substack{S \subseteq \mathbb{R}^d \\ \pi^\star(S) = \tau}} \pi_k(S) - \tau. \tag{8}$$

Next, the maximum of the above expression over $\tau \in [0, 1]$ is exactly the first characterization of the total variation distance from Fact 1, Part XI. Therefore, to show that $\pi_k \to \pi^\star$ in total variation, it suffices to uniformly bound h_k over $[0, 1]$, which we now provide the tools for.

Lemma 1. *Consider a random walk on $\mathbb{R}^d$ with stationary distribution $\pi^\star$ and lazy, reversible transition distributions $\{\mathcal{T}_x\}_{x \in \mathbb{R}^d}$, initialized at π_0 and with density π_k after k steps for all $k \in \mathbb{N}$, and suppose it has conductance Φ. Then the Lovasz-Simonovits curve (7) satisfies:*

$$h_k(\tau) \le \frac{1}{2} h_{k-1}\left((1 - 2\Phi)\tau\right) + \frac{1}{2} h_{k-1}\left((1 + 2\Phi)\tau\right) \text{ for all } \tau \in \left[0, \frac{1}{2}\right],$$

$$h_k(\tau) \le \frac{1}{2} h_{k-1}\left(\tau - 2\Phi(1 - \tau)\right) + \frac{1}{2} h_{k-1}\left(\tau + 2\Phi(1 - \tau)\right) \text{ for all } \tau \in \left[\frac{1}{2}, 1\right].$$

Proof. We only prove the former statement, as the latter follows analogously. We observe that h_k is a concave function of τ for all $k \in \mathbb{N} \cup \{0\}$, since recalling the characterization (8), h_k selects points for the optimal set greedily according to $\frac{\pi_k}{\pi}$, which is a decreasing ratio as the $\pi^\star$-measure of the set τ increases. Next, let S achieve the bound in $h_k(\tau)$, following (8). We define

$$g_1(x) := \begin{cases} 2\mathcal{T}_x(S) - 1 & x \in S \\ 0 & x \notin S \end{cases}, \; g_2(x) := \begin{cases} 1 & x \in S \\ 2\mathcal{T}_x(S) & x \notin S \end{cases}.$$

Because of the laziness assumption, $g_1(x) \in [0, 1]$ pointwise (since $2\mathcal{T}_x(S) \geq 1$ if $x \in S$; we similarly have $g_2(x) \in [0, 1]$ pointwise. Moreover, $\frac{1}{2}(g_1(x) + g_2(x)) = \mathcal{T}_x(S)$ for all $x \in \mathbb{R}^d$. Thus, letting

$$\tau_1 := \int g_1(x)\pi^\star(x)\mathrm{d}x, \; \tau_2 := \int g_2(x)\pi^\star(x)\mathrm{d}x,$$

we have $\frac{1}{2}(\tau_1 + \tau_2) = \int \mathcal{T}_x(S)\pi^\star(x)\mathrm{d}x = \pi^\star(S)$ by the definition of stationary measure. Further,

$$\begin{aligned} h_k(\tau) = \pi_k(S) - \pi^\star(S) &= \int \mathcal{T}_x(S)\pi_{k-1}(x)\mathrm{d}x \\ &= \frac{1}{2}\int g_1(x)\pi_{k-1}(x)\mathrm{d}x + \frac{1}{2}\int g_2(x)\pi_{k-1}(x)\mathrm{d}x \\ &\leq \frac{1}{2}h_{k-1}(\tau_1) + \frac{1}{2}h_{k-1}(\tau_2), \end{aligned}$$

where the first line used (2), and the last used that $g_1 \in G_{\tau_1}$, $g_2 \in G_{\tau_2}$ by definition. Finally, we use the definition of conductance to show that τ_1 and τ_2 are in fact separated from τ:

$$\begin{aligned} \tau_1 = \int_{x \in S}(2\mathcal{T}_x(S) - 1)\pi^\star(x)\mathrm{d}x &= 2\int_{x \in S}\mathcal{T}_x(S)\pi^\star(x)\mathrm{d}x - \tau \\ &= 2\int_{x \in S}\left(1 - \mathcal{T}_x(S^c)\right)\pi^\star(x)\mathrm{d}x - \tau \\ &= \tau - 2\int_{x \in S}\mathcal{T}_x(S^c)\pi^\star(x)\mathrm{d}x = \tau - 2Q(S) \leq (1 - 2\Phi)\tau. \end{aligned}$$

In the last line, we recalled the definitions (5) and (6). This also implies $\tau_2 \geq (1 + 2\Phi)\tau$, since earlier we showed $\frac{1}{2}(\tau_1 + \tau_2) = \tau$. Finally, by concavity of h_{k-1}, we have the conclusion

$$h_k(\tau) \leq \frac{1}{2}h_{k-1}(\tau_1) + \frac{1}{2}h_{k-1}(\tau_2) \leq \frac{1}{2}h_{k-1}\left((1 - 2\Phi)\tau\right) + \frac{1}{2}h_{k-1}\left((1 + 2\Phi)\tau\right).$$

$\square$

As a consequence of Lemma 1, we have the following bound on the decay of the Lovasz-Simonovits curve when the random walk is initialized at a warm start.

Corollary 1. *Consider a random walk on $\mathbb{R}^d$ with stationary distribution $\pi^\star$ and lazy, reversible transition distributions $\{\mathcal{T}_x\}_{x \in \mathbb{R}^d}$, initialized at π_0 and with density π_k after k steps for all $k \in \mathbb{N}$, and suppose it has conductance Φ. Let π_0 be β-warm with respect to $\pi^\star$. Then,*

$$D_{\mathrm{TV}}\left(\pi_k, \pi^\star\right) \leq \epsilon, \; \text{if } k \geq \frac{2}{\Phi^2}\log\left(\frac{\beta}{\epsilon}\right).$$

Proof. We claim that for all $\tau \in [0, 1]$ and all $k \in \mathbb{N} \cup \{0\}$,

$$h_k(\tau) \leq \beta \min\left(\sqrt{\tau}, \sqrt{1 - \tau}\right)\left(1 - \frac{\Phi^2}{2}\right)^k. \tag{9}$$

The conclusion then follows from (9) using (8) and the first definition of total variation distance from Fact 1, Part XI, as well as $\exp(-c) \geq 1 - c$ for $c \in [0, 1]$. For $k = 0$, (9) follows from

warmness, since (following notation (8)) any S with $\pi^\star(S) = \tau$ has $\pi_0(S) \le \min(\beta\tau, 1)$. We proceed inductively: for $\tau \in [0, \frac{1}{2}]$, supposing (9) holds for some $k - 1$, Lemma 1 yields

$$
\begin{aligned}
h_k(\tau) &\le \frac{1}{2} h_{k-1}((1 - 2\Phi)\tau) + \frac{1}{2} h_{k-1}((1 + 2\Phi)\tau) \\
&\le \beta \cdot \frac{1}{2} \left(\sqrt{(1 - 2\Phi)\tau} + \sqrt{(1 + 2\Phi)\tau} \right) \left(1 - \frac{\Phi^2}{2} \right)^{k-1} \\
&\le \beta \min \left(\sqrt{\tau}, \sqrt{1 - \tau} \right) \left(1 - \frac{\Phi^2}{2} \right)^k,
\end{aligned}
$$

where we used that $\sqrt{1 - x} + \sqrt{1 + x} \le 2(1 - \frac{x^2}{2})$.[3] Finally, the case $\tau \in [\frac{1}{2}, 1]$ is handled analogously using the other conclusion in Lemma 1, establishing (9) as desired. $\qquad \square$

The most appropriate comparison to Corollary 1 is Corollary 2, Part XI, the analogous result we showed in the discrete setting. Compared to its discrete counterpart, Corollary 1 incurs a logarithmic dependence on the warmness β rather than $(\min_{i \in [d]} \pi_i^\star)^{-1}$. This latter quantity is typically meaningless in the continuous setting when the target density has no atoms.

Recall that Corollary 2, Part XI was actually established in two steps: we first assumed a spectral gap on the Markov chain transition operator, and then used Cheeger's inequality to lower bound the spectral gap as a function of the conductance. There is an analogous definition in the continuous setting, where one defines the "spectral gap" to be the decay of the variance of the relative density $\frac{\pi_k}{\pi^\star}$, which will be formalized in a following lecture. To introduce these ideas, we mention a few useful definitions here and compare them. First, the total variation convergence result in Corollary 1 can be slightly modified to directly show that

$$
\chi^2(\pi_k \| \pi^\star) := \mathbb{E}_{\pi^\star} \left[\left(\frac{\pi_k}{\pi^\star} - 1 \right)^2 \right] = \mathrm{Var}_{\pi^\star} \left[\frac{\pi_k}{\pi^\star} \right]
$$

also decays by a factor of $\Omega(\Phi^2)$ in each iteration (we state such a result shortly). Here, $\chi^2(\mu \| \pi) := \mathrm{Var}_\pi[\frac{\mu}{\pi}]$ is the *chi-squared divergence*, which satisfies the inequality $\chi^2(\mu \| \pi) \ge D_{\mathrm{KL}}(\mu \| \pi) := \mathbb{E}_\pi[\frac{\mu}{\pi} \log \frac{\mu}{\pi}]$,[4] where $D_{\mathrm{KL}}(\mu \| \pi)$ is the KL divergence. Importantly, the continuous analog of Pinsker's inequality (see Lemma 7 and Remark 5, Part III), which states that

$$
D_{\mathrm{KL}}(\mu \| \pi) \ge \frac{1}{2} \left(\int |\mu(x) - \pi(x)| \mathrm{d}x \right)^2 = 2 D_{\mathrm{TV}}(\mu, \pi)^2,
$$

lets us transfer χ^2 bounds to D_{TV} bounds. Further, under a β-warm start it is simple to see that $\chi^2(\pi_0 \| \pi^\star) \le \beta^2$, so the aforementioned variance decay result (i.e. $\chi^2(\pi_k \| \pi^\star)$ falls by $\Omega(\Phi^2)$) recovers Corollary 1 up to constant factors, and yields a stronger bound in general.

One shortcoming of Corollary 1 is that in continuous settings, the warmness β is typically exponential in the dimension. As a simple example, using the uniform distribution on $\mathbb{B}(\mathbb{0}_d, 1)$ as a warm start for the uniform distribution on $\mathbb{B}(\mathbb{0}_d, 2)$ yields warmness parameter

$$
\beta = \frac{\mathrm{Vol}\,(\mathbb{0}_d, 2)}{\mathrm{Vol}\,(\mathbb{0}_d, 1)} = 2^d.
$$

This phenomenon is quite general when scale parameters on the target distribution cannot be estimated to high precision, and so extraneous $\log \beta$ factors in mixing times result in dimension-dependent overheads. A natural alternative approach is to directly show that the KL divergence decays at a linear rate, because the initial KL divergence for a β-warm start π_0 is bounded by

$$
D_{\mathrm{KL}}\left(\pi_0 \| \pi^\star\right) = \mathbb{E}_{\pi^\star} \left[\frac{\pi_0}{\pi^\star} \log \frac{\pi_0}{\pi^\star} \right] = \mathbb{E}_{\pi_0} \left[\log \frac{\pi_0}{\pi^\star} \right] \le \log \beta.
$$

---

[3] The choice of the potential function $\sqrt{\cdot}$ in the proof is not particularly important. The most important property is that we use a concave potential function, so that the separation of τ values in Lemma 1 induces negative drift. In particular, a first-order Taylor approximation of $\sqrt{\cdot}$ provides intuition for the calculation used here; the first-order terms cancel, and the second-order terms induce a negative quadratic in Φ by concavity.

[4] We prove this in a later lecture, but it follows from the facts that $\chi^2(\mu \| \pi)$ dominates the Rényi divergence of order 2, D_{KL} is the Rényi divergence of order 1, and Rényi divergences are monotone in the order.

Therefore, establishing this type of *entropy decay* (rather than variance decay) results in mixing time overheads of $\log \log \beta$, which is $O(\log d)$ rather than $O(d)$ when $\beta = \exp(\Theta(d))$. These types of entropy decay bounds, which are called modified log-Sobolev inequalities in the literature (see [BT03]), are typically harder to come by than the simpler spectral gap bounds which imply variance decay, though we will see an example of entropy decay later in the course.

Finally, we mention an extension of Corollary 1 which we find useful in applications. In particular, the following extension due to [CDWY20] (building upon *average conductance* frameworks from [LK99, GMT06]) handles situations where uniform conductance bounds do not hold due to small sets with poor behavior, and also sometimes removes extraneous $\log \beta$ factors.

Proposition 1 (Lemma 3, [CDWY20]). *Let $\epsilon \in (0,1)$, and consider a random walk on $\mathbb{R}^d$ with stationary distribution $\pi^\star$ and lazy, reversible transition distributions $\{\mathcal{T}_x\}_{x \in \mathbb{R}^d}$, initialized at π_0 and with density π_k after k steps for all $k \in \mathbb{N}$. Let π_0 be β-warm with respect to $\pi^\star$, and suppose $\Omega \subseteq \mathbb{R}^d$ has $m := \pi^\star(\Omega) \geq 1 - \frac{\epsilon^2}{3\beta^2}$. For all $\tau \in (0, \frac{m}{2}]$, define (following notation (5))*

$$\Phi_\Omega(\tau) := \inf_{\substack{S \subseteq \Omega \\ \pi^\star(S) \leq \tau}} \frac{Q(S)}{\pi^\star(S)}. \tag{10}$$

Then,

$$\chi^2(\pi_k \| \pi^\star) \leq \epsilon, \ \text{if } k \geq \int_{\frac{4}{\beta}}^{\frac{m}{2}} \frac{16}{\tau \Phi_\Omega(\tau)^2} d\tau + \frac{32}{\Phi_\Omega(\frac{m}{2})^2} \log\left(\frac{4}{\epsilon}\right).$$

We pause to interpret Proposition 1. First, because $\Phi_\Omega(\tau)$ is decreasing as τ increases because the definition (10) includes more sets, in the special case $\Omega = \mathbb{R}^d$ and $m = 1$, so $\Phi = \Phi_\Omega(\frac{m}{2})$ is the standard conductance (Definition 1), the bound in Proposition 1 is never worse than

$$32 \int_{\frac{4}{\beta}}^{\frac{4}{\epsilon}} \frac{1}{\tau \Phi^2} d\tau = \Theta\left(\frac{1}{\Phi^2} \log\left(\frac{\beta}{\epsilon}\right)\right).$$

This matches Corollary 1 up to a constant factor, and bounds the χ^2 divergence rather than the smaller D_{TV}, giving a stronger bound. Further, Proposition 1 allows for excluding a small set with poor behavior, which can be useful on random walks in $\mathbb{R}^d$ where "far away" points are difficult to control, but also are unlikely under $\pi^\star$. Finally, if one can prove a bound on $\Phi_\Omega(\tau)$ which improves with τ, Proposition 1 can lead to significantly sharper mixing time estimates. For example, suppose

$$\Phi_\Omega(\tau) = \Omega\left(\Phi \sqrt{\log \frac{1}{\tau}}\right),$$

which is the case for Gaussian densities (and more generally, distributions with strongly convex negative log-densities, see e.g. Theorem 2.7 in [Led99]). Then, the bound in Proposition 1 reads[5]

$$k \gtrsim \int_{\frac{4}{\beta}}^{\frac{m}{2}} \frac{1}{\tau \log(\frac{1}{\tau}) \cdot \Phi^2} d\tau + \frac{1}{\Phi^2} \log \frac{1}{\epsilon} \gtrsim \frac{1}{\Phi^2} \log \log \beta + \log \beta + \frac{1}{\Phi^2} \log \frac{1}{\epsilon},$$

where we used that the antiderivative of $(\tau \log(\frac{1}{\tau}))^{-1}$ is $\log \log \frac{1}{\tau}$. This can improve mixing times by dimension-dependent factors, when initialized at a "cold start" of quality $\beta = \exp(\Theta(d))$.

2.3 Boosting

In this section, we give a generic reduction inspired by the boosting result shown in Corollary 1, Part XI for discrete Markov chains. Specifically, we show that if a Markov chain requires T_{mix} iterations to achieve constant total variation distance, where T_{mix} has a logarithmic overhead in the initial warmness, then it mixes to ϵ total variation distance in roughly $T_{\text{mix}} \log \frac{1}{\epsilon}$ iterations. As in Corollary 1, Part XI, this section will not make use of reversibility, so it applies generically.

kjtian: This section is Homework IV, Problems 1 and 5. I will update when it is due.

---

[5] The additive $\log \beta$ term arises because $\Phi_\Omega(\tau) \leq 1$ for all τ.

3 Conductance via isoperimetry and transition overlaps

We established in Corollary 1 and Proposition 1 that to bound the mixing time (in D_{TV} or χ^2, respectively), it suffices to lower bound the conductance. In this section, we give a geometric argument by [LS93] which reduces conductance bounds to bounding two quantities which are typically simpler to reason about in applications. We define below the first such quantity.

Definition 4 (Isoperimetry). *Let $\pi^\star$ be a density on $\mathbb{R}^d$, and let $\mathrm{m} : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}_{\geq 0}$ be a metric. We say that $\pi^\star$ has* isoperimetric constant ψ *if, for any partition[6] S_1, S_2, S_3 of $\mathbb{R}^d$,*

$$\frac{\pi^\star(S_3)}{\mathrm{m}(S_1, S_2)\pi^\star(S_1)\pi^\star(S_2)} \geq \psi,$$

where we let $\mathrm{m}(S, T) := \min_{x \in S, y \in T} \mathrm{m}(x, y)$ for $S, T \subseteq \mathbb{R}^d$.

Importantly, the isoperimetric constant in Definition 4 is a function of only the target distribution $\pi^\star$, and has no dependence on the transition distributions $\{\mathcal{T}_x\}_{x \in \mathbb{R}^d}$. This definition isolates a geometric property of the distribution, which can be reasoned about separately. To gain some intuition for Definition 4, note that when S_3 is taken to be an infinitesimally small strip on the boundary between S_1 and S_2, then $\frac{\pi^\star(S_3)}{\mathrm{m}(S_1, S_2)}$ is essentially the surface area of the boundary. Therefore, in this case Definition 4 reduces to a lower bound on surface area-to-volume ratios.

The second quantity used in our conductance lower bound framework can informally be thought of as an overlap bound between "nearby" transition distributions, where the notion of nearby is compatible with Definition 4. Specifically, we assume that

$$D_{\mathrm{TV}}(\mathcal{T}_x, \mathcal{T}_y) \leq \frac{1}{2} \text{ for all } x, y \in \mathbb{R}^d \text{ with } \mathrm{m}(x, y) \leq \Delta. \tag{11}$$

The bound (11) is a local property of transition distributions, because it only needs to hold for pairs of nearby points (x, y). We next give a representative result which shows how to lift local bounds of the form (11) and establish conductance lower bounds, a global property, via isoperimetry.

Proposition 2. *Consider a random walk on $\mathbb{R}^d$ with stationary distribution $\pi^\star$ and lazy, reversible transition distributions $\{\mathcal{T}_x\}_{x \in \mathbb{R}^d}$. Suppose that $\pi^\star$ has isoperimetric constant ψ, and that (11) holds. Then the conductance of the random walk satisfies*

$$\Phi \geq \min\left(\frac{1}{8}, \frac{\Delta\psi}{64}\right).$$

Proof. Choose an arbitrary set $S \subseteq \mathbb{R}^d$ with $0 < s := \pi^\star(S) \leq \frac{1}{2}$, so our goal is to lower bound $\frac{Q(S)}{s}$ by the stated quantity. Define the following sets:

$$S_1 := \left\{x \in S \mid \mathcal{T}_x(S^c) < \frac{1}{4}\right\}, \; S_2 := \left\{x \in S^c \mid \mathcal{T}_x(S) < \frac{1}{4}\right\}, \; S_3 := \mathbb{R}^d \setminus (S_1 \cup S_2).$$

Intuitively, S_1 and S_2 are the points deep in S and S^c respectively, since they are unlikely to cross after one random walk step. Note that for $x \in S_1$, $y \in S_2$, we have $D_{\mathrm{TV}}(x, y) > \frac{1}{2}$, and therefore $\mathrm{m}(x, y) > \Delta$ due to the assumption (11). If $\pi^\star(S_1) < \frac{1}{2}\pi^\star(S)$, then

$$Q(S) \geq \int_{x \in S_3 \cap S} \mathcal{T}_x(S^c)\pi^\star(x)\mathrm{d}x \geq \frac{1}{4}\pi^\star(S_3 \cap S) \geq \frac{s}{8}.$$

Similarly, if $\pi^\star(S_2) < \frac{1}{2}\pi^\star(S^c)$, then

$$Q(S) \geq \int_{x \in S_3 \cap S^c} \mathcal{T}_x(S)\pi^\star(x)\mathrm{d}x \geq \frac{1}{4}\pi^\star(S_3 \cap S^c) \geq \frac{1-s}{8} \geq \frac{s}{8}.$$

Finally, if both $\pi^\star(S_1) \geq \frac{1}{2}\pi^\star(S)$ and $\pi^\star(S_2) \geq \frac{1}{2}\pi^\star(S^c)$, we have

$$Q(S) = \frac{1}{2}\left(\int_{x \in S} \mathcal{T}_x(S^c)\pi^\star(x)\mathrm{d}x + \int_{x \in S^c} \mathcal{T}_x(S)\pi^\star(x)\mathrm{d}x\right)$$

$$\geq \frac{1}{8}\int_{x \in S_3} \pi^\star(x)\mathrm{d}x = \frac{\pi^\star(S_3)}{8} \geq \frac{\Delta\psi\pi^\star(S_1)\pi^\star(S_2)}{8} \geq \frac{\Delta\psi s}{64}.$$

---

[6]That is, $\bigcup_{i \in [3]} S_i = \mathbb{R}^d$ and $S_i \cap S_j = \emptyset$ for all $i \neq j$.

In the second-to-last inequality, we used our earlier conclusion $\mathrm{m}(S_1, S_2) \geq \Delta$, and the definition of ψ; in the last inequality, we used our assumptions $\pi^\star(S_1) \geq \frac{1}{2}\pi^\star(S)$ and $\pi^\star(S_2) \geq \frac{1}{2}\pi^\star(S^c) \geq \frac{1}{4}$. $\quad\square$

There are various extensions to the proof strategy in Proposition 2, for example handling sets of smaller measure or excluding ill-behaved regions, as suggested by Proposition 1; we will require one such extension in Section 5. Nonetheless, the overall strategy of lower bounding an appropriate notion of isoperimetry (potentially restricted to specific subsets) and establishing the overlap of nearby transition distributions remains as a common theme. In the rest of this lecture, we give techniques for reasoning about these quantities, specialized to uniform sampling of convex bodies.

4 Localization

Up to this point, our development has not made any structural assumption about the target distribution $\pi^\star$ (or the transition distributions $\{\mathcal{T}_x\}_{x \in \mathbb{R}^d}$). In this section, we provide an isoperimetric inequality which is specific to logconcave densities $\pi^\star$, in order to introduce a powerful tool: the *localization lemma*, first introduced by [LS93, KLS95]. The localization lemma states, at a high level, that the boundary points of the set of logconcave densities satisfying one linear inequality are one-dimensional and logaffine. This characterization is extremely useful when maximizing linear functions of logconcave densities subject to one linear inequality constraint, since it means we only need to establish a one-dimensional inequality. To make these statements precise, we give a convenient formulation of the localization lemma from [FG04], building upon [LS93, KLS95].[7]

Proposition 3 (Theorem 1, [FG04])**.** *Let $\mathcal{K} \subset \mathbb{R}^d$ be compact and convex, and let $f : \mathcal{K} \to \mathbb{R}$ be upper semicontinuous. Let $\mathcal{P}(f)$ be the set of logconcave densities $\pi : \mathcal{K} \to \mathbb{R}_{\geq 0}$ satisfying $\int_{x \in \mathcal{K}} f(x)\pi(x)\mathrm{d}x \geq 0$. All boundary points[8] of $\mathrm{Conv}(\mathcal{P}(f))$ satisfy one of the following.*

1. *π is a Dirac measure at a point $x \in \mathcal{K}$ where $f(x) \geq 0$.*

2. *π is a logaffine distribution satisfying $\int f(x)\pi(x)\mathrm{d}x = 0$, and is supported on a one-dimensional subspace, i.e. for some $\gamma \in \mathbb{R}$ and $a, b \in \mathcal{K}$, we have*

$$\pi(x) \propto \begin{cases} \exp(\gamma t) & x = (1-t)a + tb \text{ for some } t \in [0,1] \\ 0 & \text{otherwise} \end{cases}. \tag{12}$$

Densities of the form (12) are often called exponential needles, because they are truncated exponential distributions over one-dimensional "needles" within $\mathcal{K}$. We first briefly sketch the proof of Proposition 3, deferring more details to [FG04]. Suppose for the sake of contradiction that there is a boundary point π of $\mathrm{Conv}(\mathcal{P}(f))$ with $\dim S \geq 2$, where S is the least affine subspace containing the support of π. Let $x_0 \in S$ and let $E \subseteq S$ be an arbitrary two-dimensional subspace, such that $x_0 + E \subseteq S$. Let $C(E)$ be the unit circle in E, and for any $u \in C(E)$ denote

$$H_u^+ := \{x \in S \mid \langle x - x_0, u \rangle > 0\}, \; H_u^- := \{x \in S \mid \langle x - x_0, u \rangle < 0\}.$$

By choosing x_0 in the relative interior of the support of π, we can ensure that $\pi(H_u^+), \pi(H_u^-) > 0$ for all $u \in C(E)$. Next, define $\phi : C(E) \to \mathbb{R}$ by $\phi(u) := \int_{x \in H_u^+} f(x)\pi(x)\mathrm{d}x - \frac{1}{2}\int f(x)\pi(x)\mathrm{d}x$. Note that ϕ is a continuous function of u, and further $\phi(u) = -\phi(-u)$ for all $u \in C(E)$.[9] The intermediate value theorem then gives $u \in C(E)$ with $\phi(u) = 0$, so that π^+ and π^-, the restrictions of π to H_u^+ and H_u^-, have $\int f(x)\pi^+(x)\mathrm{d}x \geq 0$ and $\int f(x)\pi^-(x)\mathrm{d}x \geq 0$. Finally, we can write π as a nontrivial convex combination of π^+ and π^-, so π is not a boundary point of $\mathrm{Conv}(\mathcal{P}(f))$.

Similar arguments are then used to decompose any one-dimensional elements of $\mathcal{P}(f)$ which do not satisfy $\int f\pi\mathrm{d}x = 0$, or are not logaffine, so they cannot be boundary points. For the former property, one can split the measure along the line at any point giving half the value of $\int f\pi\mathrm{d}x$.

---

[7]The statement in [FG04] is more general than Proposition 3, as it gives a complete characterization of extreme points, allows for multiple constraints, and extends to s-concave distributions, which generalizes logconcavity.

[8]Recall from Lemma 3, Part I that the boundary points of a convex set S are points $x \in S$ such that there do not exist $x', x'' \in S$ and $\lambda \in (0, 1)$ such that $x = (1 - \lambda)x' + \lambda x''$.

[9]To see this, we used that $\pi(\mathbb{R}^d \setminus (H_u^+ \cup H_u^-)) = 0$. This is a nontrivial fact which follows from a characterization of logconcave functions in [Bor75], showing that logconcave functions are absolutely continuous with respect to the Lebesgue measure on the least affine subspace containing their support. Therefore, they place zero measure on any lower-dimensional subspace, so we can exclude the boundary $\mathbb{R}^d \setminus (H_u^+ \cup H_u^-)$ from consideration.

For the latter property, the "largest" possible extension of a logconcave function is logaffine, just as extremal examples of convex functions are affine; one can use this strategy to obtain a logaffine function which dominates any logconcave function, and split it accordingly. For more details, we refer the reader to the (fairly short) proof of Proposition 3 in [FG04]. We also mention that these proofs use relatively little about the structure of logconcave functions; for example, the one-dimensional argument only used absolute continuity and that restricting to convex sets preserves logconcavity. Due to this, analogous arguments can be extended to broader families of distributions; see e.g. Lemma 1 of [GLL$^+$23] for an extension to strongly logconcave densities.

We now give an example use of Proposition 3, in proving a variant of the localization lemma which is often easier to apply. The key idea, as mentioned previously, is that if we want to prove a linear inequality for all logconcave functions in some $\mathcal{P}(f)$, it suffices to prove the inequality for point masses and one-dimensional logaffine functions, by using Lemma 3, Part I.

Lemma 2. *Let $f_1 : \mathbb{R}^d \to \mathbb{R}_{\geq 0}$, $f_2 : \mathbb{R}^d \to \mathbb{R}_{\geq 0}$ be upper semicontinuous and let $f_3 : \mathbb{R}^d \to \mathbb{R}_{\geq 0}$, $f_4 : \mathbb{R}^d \to \mathbb{R}_{\geq 0}$ be lower semicontinuous. Then the following two statements are equivalent.*

1. For every logconcave density $\pi : \mathbb{R}^d \to \mathbb{R}_{\geq 0}$,

$$\left(\int f_1(x)\pi(x)\mathrm{d}x \right) \left(\int f_2(x)\pi(x)\mathrm{d}x \right) \leq \left(\int f_3(x)\pi(x)\mathrm{d}x \right) \left(\int f_4(x)\pi(x)\mathrm{d}x \right).$$

2. For all $x \in \mathbb{R}^d$, $f_1(x)f_2(x) \leq f_3(x)f_4(x)$, and for every $a, b \in \mathbb{R}^d$ and $\gamma \in \mathbb{R}$,

$$\left(\int_0^1 f_1((1-t)a + tb)\exp(\gamma t)\mathrm{d}t \right) \left(\int_0^1 f_1((1-t)a + tb)\exp(\gamma t)\mathrm{d}t \right)$$
$$\leq \left(\int_0^1 f_3((1-t)a + tb)\exp(\gamma t)\mathrm{d}t \right) \left(\int_0^1 f_4((1-t)a + tb)\exp(\gamma t)\mathrm{d}t \right).$$

Proof. It is obvious that Item 1 implies Item 2, since truncated logaffine functions are also logconcave, so we focus on showing Item 2 implies Item 1. By a limiting argument, it suffices to consider the case where π is compactly supported on $\mathcal{K} \subset \mathbb{R}^d$, and $\int f_3(x)\pi(x)\mathrm{d}x > 0$. Then, define

$$f = f_1 - \left(\frac{\int f_1(x)\pi(x)\mathrm{d}x}{\int f_3(x)\pi(x)\mathrm{d}x} \right) f_3, \; g = \left(\frac{\int f_1(x)\pi(x)\mathrm{d}x}{\int f_3(x)\pi(x)\mathrm{d}x} \right) f_2 - f_4.$$

The maximum value of $\int g(x)\mu(x)\mathrm{d}x$ for $\mu \in \mathrm{Conv}(\mathcal{P}(f))$ is achieved by either a Dirac measure or a one-dimensional logaffine function, by Proposition 3. The fact that $\mu \in \mathrm{Conv}(\mathcal{P}(f))$ implies that

$$\int f_1(x)\mu(x)\mathrm{d}x - \left(\frac{\int f_1(x)\pi(x)\mathrm{d}x}{\int f_3(x)\pi(x)\mathrm{d}x} \right) \int f_3(x)\mu(x)\mathrm{d}x \geq 0$$

$$\implies \frac{\int f_1(x)\pi(x)\mathrm{d}x}{\int f_3(x)\pi(x)\mathrm{d}x} \leq \frac{\int f_1(x)\mu(x)\mathrm{d}x}{\int f_3(x)\mu(x)\mathrm{d}x}. \tag{13}$$

Thus, the fact that $\pi \in \mathcal{P}(f)$ by definition of f shows the desired

$$\left(\frac{\int f_1(x)\pi(x)\mathrm{d}x}{\int f_3(x)\pi(x)\mathrm{d}x} \right) \int f_2(x)\pi(x)\mathrm{d}x - \int f_4(x)\pi(x)\mathrm{d}x$$

$$\leq \left(\frac{\int f_1(x)\pi(x)\mathrm{d}x}{\int f_3(x)\pi(x)\mathrm{d}x} \right) \int f_2(x)\mu(x)\mathrm{d}x - \int f_4(x)\mu(x)\mathrm{d}x$$

$$\leq \left(\frac{\int f_1(x)\mu(x)\mathrm{d}x}{\int f_3(x)\mu(x)\mathrm{d}x} \right) \int f_2(x)\mu(x)\mathrm{d}x - \int f_4(x)\mu(x)\mathrm{d}x \leq 0,$$

where the second inequality used (13), and the last used the assumption in Item 2. $\square$

To demonstrate the utility of Lemma 2, we use it to give a bound on the isoperimetric constant (Definition 4) for any logconcave density π, parameterized by the average distance of points to its mean, which we denote (patterning off notation in Theorem 2, Part I) by

$$\bar{x}_\pi := \mathbb{E}_{x \sim \pi}[x] = \int x\pi(x)\mathrm{d}x. \tag{14}$$

Lemma 3. *Let $\pi : \mathbb{R}^d \to \mathbb{R}_{\geq 0}$ be a logconcave density, and let S_1, S_2, S_3 be a partition of $\mathbb{R}^d$ with $\min_{x \in S_1, y \in S_2} \|x - y\|_2 \geq \Delta$. Then,*

$$\pi(S_1)\pi(S_2) \leq \frac{\mathbb{E}_{x \sim \pi}[\|x - \bar{x}_\pi\|_2]}{\Delta \log 2} \cdot \pi(S_3).$$

Proof. Let $f_i(x) := \mathbf{1}_{x \in S_i}$ for $i \in [3]$ be the indicator functions associated to the partition, and let $f_4(x) := \frac{1}{\Delta \log 2} \|x - \bar{x}_\pi\|_2$. Clearly for all $x \in \mathbb{R}^d$, we have $f_1(x)f_2(x) = 0 \leq f_3(x)f_4(x)$. Thus, to obtain the conclusion, Lemma 2 states that it suffices to show that for all $a, b \in \mathbb{R}^d$ and $\gamma \in \mathbb{R}$,

$$\left(\int_0^1 \mathbf{1}_{(1-t)a+tb \in S_1} \exp(\gamma t)\mathrm{d}t\right) \left(\int_0^1 \mathbf{1}_{(1-t)a+tb \in S_2} \exp(\gamma t)\mathrm{d}t\right)$$
$$\leq \frac{1}{\Delta \log 2} \left(\int_0^1 \mathbf{1}_{(1-t)a+tb \in S_3} \exp(\gamma t)\mathrm{d}t\right) \left(\int_0^1 \|(1-t)a + tb - \bar{x}_\pi\|_2 \exp(\gamma t)\mathrm{d}t\right).$$

We next observe that changing $\bar{x}_\pi$ to its projection on the line between a, b, i.e. some point $(1-u)a + ub$ for $u \in [0, 1]$, can only decrease the right-hand side above. Moreover, we can redefine

$$I_i := \{t \in [0, 1] \mid (1-t)a + tb \in S_i\} \text{ for all } i \in [3],$$

so that $\min_{t \in I_1, t' \in I_2} \|((1-t)a + tb) - ((1-t')a + t'b)\|_2 = \min_{t \in I_1, t' \in I_2} |t - t'| \|a - b\|_2 \geq \Delta$. Therefore, rescaling Δ by a factor of $\|a - b\|_2$, our problem reduces to showing that for any partition I_1, I_2, I_3 of $[0, 1]$ with $\min_{t \in I_1, t' \in I_2} |t - t'| \geq \Delta$, and for any $u \in [0, 1]$,

$$\left(\int_{t \in I_1} \exp(\gamma t)\mathrm{d}t\right) \left(\int_{t \in I_2} \exp(\gamma t)\mathrm{d}t\right) \leq \frac{1}{\Delta \log 2} \left(\int_{t \in I_3} \exp(\gamma t)\mathrm{d}t\right) \left(\int_0^1 |u - t| \exp(\gamma t)\mathrm{d}t\right). \quad (15)$$

In the case when I_1, I_2, and I_3 are single intervals (where the length of I_3 is at least Δ), the proof of (15) is carried out using elementary arguments in Theorem 5.2, [KLS95], which we omit here. We instead discuss how to reduce to this particular case for general partitions, which is a core part of localization arguments. First, we may assume that I_3 is open by moving a measure zero set of endpoints which does not affect the inequality. Next, every interval in I_3 has length at least Δ, since any shorter interval must have both endpoints in I_1 or I_2, and hence can be moved to the corresponding partition piece which only makes (15) tighter. Under these simplifications, we can write $I_3 = \bigcup_{j \in [k]}(t_j, u_j)$ for a finite k. Applying (15) to each interval (t_j, u_j) and summing,

$$\sum_{j \in [k]} \left(\int_0^{t_j} \exp(\gamma t)\mathrm{d}t\right) \left(\int_{u_j}^1 \exp(\gamma t)\mathrm{d}t\right) \leq \frac{1}{\Delta \log 2} \left(\int_{t \in I_3} \exp(\gamma t)\mathrm{d}t\right) \left(\int_0^1 |u - t| \exp(\gamma t)\mathrm{d}t\right).$$

The claim follows since every pair of intervals in I_1 and I_2 are separated by some (t_j, u_j), so

$$\sum_{j \in [k]} \left(\int_0^{t_j} \exp(\gamma t)\mathrm{d}t\right) \left(\int_{u_j}^1 \exp(\gamma t)\mathrm{d}t\right) \geq \left(\int_{t \in I_1} \exp(\gamma t)\mathrm{d}t\right) \left(\int_{t \in I_2} \exp(\gamma t)\mathrm{d}t\right).$$

$\square$

We mention that the proof of Lemma 3 uses nothing specific about $\bar{x}_\pi$, so it can be replaced with an arbitrary point. One reason why the bound with $\bar{x}_\pi$ in particular is useful is because

$$\mathbb{E}_{x \sim \pi}\left[\|x - \bar{x}_\pi\|_2\right] \leq \sqrt{\mathbb{E}_{x \sim \pi}\left[\|x - \bar{x}_\pi\|_2^2\right]} = \sqrt{\mathrm{Tr}\left(\underbrace{\mathbb{E}_{x \sim \pi}\left[(x - \bar{x}_\pi)(x - \bar{x}_\pi)^\top\right]}_{:=\mathrm{Cov}_\pi}\right)}. \quad (16)$$

In other words, Lemma 3 proves that the isoperimetric constant (Definition 4) of any logconcave density $\pi : \mathbb{R}^d \to \mathbb{R}_{\geq 0}$ with respect to the ℓ_2 norm is lower bounded by

$$\frac{\log 2}{\sqrt{\mathrm{Tr}(\mathrm{Cov}_\pi)}}.$$

This result was first established by [KLS95], and shows that e.g. if π is isotropic (meaning $\mathrm{Cov}_\pi = \mathbf{I}_d$), the isoperimetric constant is $\Omega(d^{-\frac{1}{2}})$. The authors of [KLS95] conjectured that this bound could be improved to $\Omega(\|\mathrm{Cov}_\pi\|_{\mathrm{op}}^{-1})$, which has since come to be called the KLS conjecture, a deep result in probability theory which has intimate connections with various other results in mathematics (see [LV17b] for an extended discussion of the conjecture and these connections).

More recently, there has been tremendous progress on the KLS conjecture. First, [Eld13] developed a technique known as *stochastic localization*, which simulates the localization process in Proposition 3 which decomposes a measure into low-dimensional pieces, not by a bisection method, but rather by gradually adding a Gaussian component.[10] The reason this is useful is because the KLS conjecture is much simpler to establish for Gaussians, and indeed [LV17a] built upon [Eld13] to show that the isoperimetric constant of logconcave π is bounded by $\Omega(\|\mathrm{Cov}_\pi\|_{\mathrm{F}}^{-1})$, which is $\Omega(d^{-\frac{1}{4}})$ in the isotropic case. By bootstrapping the arguments of [Eld13, LV17a], a breakthrough was made by [Che21], who finally proved that the KLS conjecture is true up to a $d^{o(1)}$ factor. We state here the current state-of-the-art bound on the isoperimetric constant of logconcave distributions.

Proposition 4 (Theorem 1.2, [Kla23]). *There is a universal constant $C > 0$ such that for any logconcave density $\pi : \mathbb{R}^d \to \mathbb{R}_{\geq 0}$, the isoperimetric constant ψ of π with respect to $\|\cdot\|_2$ satisfies*

$$\psi \geq \frac{1}{C\sqrt{\log d}} \cdot \frac{1}{\sqrt{\|\mathrm{Cov}_\pi\|_{\mathrm{op}}}}.$$

5 Ball walk analysis

In this section, we finally apply the results of Sections 2, 3, and 4 to a specific Markov chain for sampling from the uniform density over a compact, convex body $\mathcal{K} \subseteq \mathbb{R}^d$. We fix the notation

$$\pi^\star(x) := \frac{1}{\mathrm{Vol}(\mathcal{K})} \cdot \mathbf{1}_{x \in \mathcal{K}}$$

throughout the section; by observation, $\pi^\star$ is logconcave. Following the notation (14), (16), we also make the simplifying assumptions that $\pi^\star$ is isotropic, meaning

$$\mathrm{Cov}_{\pi^\star} := \mathbb{E}_{x \sim \pi^\star}\left[(x - \bar{x}_\pi)(x - \bar{x}_\pi)^\top \right] = \mathbf{I}_d, \text{ where } \bar{x}_{\pi^\star} := \mathbb{E}_{x \sim \pi^\star}[x]. \tag{17}$$

We will discuss how to relax the assumption (17), and in a later lecture we design a rounding procedure which reduces to the case where (17) is approximately satisfied.

As a basic building block of our algorithm, we define the ball walk transition operators, denoted $\{\mathcal{T}_x\}_{x \in \mathcal{K}}$ throughout the section. The ball walk operator $\mathcal{T}_x$ first returns x with probability $\frac{1}{2}$ (so it is lazy). Otherwise, it uniformly samples $y \sim \mathbb{B}(x, \eta)$ for a step size $\eta > 0$, and then updates x to y if $y \in \mathcal{K}$, else staying at x. This filter in the latter case preserves reversibility, as it is the result of a Metropolis filter (4) applied to the lazy proposal distributions $\mathcal{P}_x$ which uniformly sample a point in $\mathbb{B}(x, \eta)$. To make this definition simpler to state, we define the *local conductance* at x,

$$\ell(x) := \frac{1}{\mathrm{Vol}(\mathbb{B}(x, \eta))} \int_{y \in \mathbb{B}(x,\eta)} \mathbf{1}_{y \in \mathcal{K}} \mathrm{d}y = \frac{\mathrm{Vol}(\mathbb{B}(x, \eta) \cap \mathcal{K})}{\mathrm{Vol}(\mathbb{B}(x, \eta))} \tag{18}$$

to be the probability of accepting a step. Drawing $x' \sim \mathcal{T}_x$ can then be concisely stated as

$$x' \leftarrow \begin{cases} x & \text{with probability } 1 - \frac{1}{2}\ell(x) \\ \text{a uniform draw from } \mathbb{B}(x, \eta) \cap \mathcal{K} & \text{with probability } \frac{1}{2}\ell(x) \end{cases}. \tag{19}$$

Due to the way we set up the Metropolis-Hastings filter, the stationary distribution of the $\{\mathcal{T}_x\}_{x \in \mathbb{R}^d}$ is indeed $\pi^\star$. We finally mention one useful property of $\ell(x)$ which is used later.

Lemma 4. *The local conductance ℓ defined in (19) is logconcave.*

Proof. Logconcavity is preserved by multiplication by a constant, so it suffices to prove that $\int_{y \in \mathbb{B}(x,\eta)} \mathbf{1}_{y \in \mathcal{K}} \mathrm{d}y$ is logconcave. This follows because it is the convolution of the logconcave indicator functions of $\mathbb{B}(\mathbb{0}_d, \eta)$ and $\mathcal{K}$, so the claim follows from Corollary 2, Part I. $\square$

---

[10] This stochastic localization process is actually identical to a popular theoretical framework for diffusion models in generative AI, see the survey article [Mon23] for further discussion.

5.1 Overlap bounds

In order to bound the conductance of the ball walk, Proposition 2 requires an isoperimetric constant bound, as well as a distance Δ such that (11) holds. From Proposition 4 and the assumption (17), we know that we can take $\psi^{-1} = O(\sqrt{\log d})$, so we wish to get a handle on the Δ parameter in (11) where m is the ℓ_2 distance. To this end, we note that for two points x, x',

$$
\begin{aligned}
D_{\mathrm{TV}}\left(\mathcal{T}_x, \mathcal{T}_{x'}\right) &\leq D_{\mathrm{TV}}\left(\mathcal{P}_x, \mathcal{P}_{x'}\right) + D_{\mathrm{TV}}\left(\mathcal{T}_x, \mathcal{P}_x\right) + D_{\mathrm{TV}}\left(\mathcal{T}_{x'}, \mathcal{P}_{x'}\right) \\
&\leq D_{\mathrm{TV}}\left(\mathcal{P}_x, \mathcal{P}_{x'}\right) + \frac{1}{2}\left(1 - \ell(x)\right) + \frac{1}{2}\left(1 - \ell(x')\right),
\end{aligned}
\tag{20}
$$

where $\mathcal{P}_x$ is the aforementioned lazy proposal distribution, which with probability $\frac{1}{2}$ draws a uniform point from $\mathbb{B}(x, \eta)$. Because $\mathcal{P}_x$ and $\mathcal{P}_{x'}$ are uniform distributions over nearby balls when they are not lazy, their total variation distance enjoys a simple closed-form formula:

$$
D_{\mathrm{TV}}\left(\mathcal{P}_x, \mathcal{P}_{x'}\right) = \frac{\mathrm{Vol}\left(\mathbb{B}(x, \eta) \setminus \mathbb{B}(x', \eta)\right)}{2\mathrm{Vol}(\mathbb{B}(x, \eta))} = \frac{\mathrm{Vol}\left(\mathbb{B}(x', \eta) \setminus \mathbb{B}(x, \eta)\right)}{2\mathrm{Vol}(\mathbb{B}(x', \eta))}.
\tag{21}
$$

We therefore provide a simple geometric lemma bounding the overlap of nearby balls.

Lemma 5. *Let $\eta > 0$, and suppose $x, x' \in \mathbb{R}^d$ satisfy $\|x - x'\|_2 = \frac{\alpha\eta}{\sqrt{d}}$. Then,*

$$
\frac{\mathrm{Vol}\left(\mathbb{B}(x, \eta) \setminus \mathbb{B}(x', \eta)\right)}{\mathrm{Vol}(\mathbb{B}(x, \eta))} \leq \alpha.
$$

Proof. The statement is invariant to shifts, rotations, and rescalings, so we may assume $\eta = 1$, $x = \mathbb{0}_d$, and $x' = \beta e_1$ for simplicity, where $\beta := \frac{\alpha}{\sqrt{d}}$. We recall the formula

$$
V_d := \mathrm{Vol}\left(\mathbb{B}(\mathbb{0}_d, 1)\right) = \frac{\pi^{\frac{d}{2}}}{\Gamma\left(\frac{d}{2} + 1\right)},
\tag{22}
$$

where Γ is Euler's gamma function, satisfying $\Gamma(n+1) = n!$ for any $n \in \mathbb{N}$. We can hence compute

$$
\begin{aligned}
\frac{\mathrm{Vol}(\mathbb{B}(\mathbb{0}_d, 1) \setminus \mathrm{Vol}(\mathbb{B}(\beta e_1, 1)))}{\mathrm{Vol}(\mathbb{B}(\mathbb{0}_d, 1))} &= \frac{1}{V_d} \int_{-1}^{\frac{\beta}{2}} V_{d-1} \cdot \left(\sqrt{1 - r^2}\right)^{d-1} \mathrm{d}r \\
&\quad - \frac{1}{V_d} \int_{-1}^{-\frac{\beta}{2}} V_{d-1} \cdot \left(\sqrt{1 - r^2}\right)^{d-1} \mathrm{d}r \\
&= \frac{V_{d-1}}{V_d} \int_{-\frac{\beta}{2}}^{\frac{\beta}{2}} \left(\sqrt{1 - r^2}\right)^{d-1} \mathrm{d}r \leq \frac{\beta V_{d-1}}{V_d}.
\end{aligned}
$$

The first equality above followed because the symmetric difference of the two balls falls left of $\frac{\beta}{2}e_1$ along the first coordinate axis, so we directly computed how much volume of each ball falls left of this point and take the difference. The conclusion follows from the inequality

$$
\frac{V_{d-1}}{V_d} = \frac{\Gamma(\frac{d}{2} + 1)}{\sqrt{\pi}\Gamma(\frac{d-1}{2} + 1)} \leq \sqrt{d}.
$$

$\square$

In conclusion, as long as two points x, x' have local conductances $\ell(x), \ell(x') \geq \frac{3}{4}$, combining (20), (21), and Lemma 5 shows that $\|x - x'\|_2 \leq \frac{\eta}{2\sqrt{d}}$ implies $D_{\mathrm{TV}}(\mathcal{T}_x, \mathcal{T}_{x'}) \leq \frac{1}{2}$, as required by (11). We will discuss the issue of requiring large local conductances in the following sections. For now, we remark that [KLS97] established a strengthening of the argument in this section by reasoning more carefully about the relative volume of balls when intersecting with $\mathcal{K}$. In particular, they showed the following multiplicative variant of Lemma 5 which holds after intersection.

Lemma 6 (Lemma 3.5, [KLS97])**.** *If $x, x' \in \mathcal{K}$ for compact, convex $\mathcal{K} \subseteq \mathbb{R}^d$ have $\|x - x'\|_2 \leq \frac{\eta}{\sqrt{d}}$,*

$$
\frac{\mathrm{Vol}\left(\mathcal{K} \cap \left(\mathbb{B}(x, \eta) \cap \mathbb{B}(x', \eta)\right)\right)}{\mathrm{Vol}(\mathbb{B}(x, \eta))} \geq \frac{1}{4} \min\left\{\ell(x), \ell(x')\right\}.
$$

As a useful corollary of Lemma 6, we show how to bound the "crossing over" probability of nearby points on different sides of a partition, in line with what was used by the proof of Proposition 2.

Corollary 2. *For compact, convex $\mathcal{K} \subseteq \mathbb{R}^d$, let $S \subseteq \mathcal{K}$, and denote $S^c := \mathcal{K} \setminus S$. For any $x \in S$ and $x' \in S^c$ with $\|x - x'\|_2 \leq \frac{\eta}{\sqrt{d}}$ and $\ell(x) \in [\frac{2}{3}\ell(x'), \frac{3}{2}\ell(x')]$,*

$$\frac{\mathcal{T}_x(S^c)}{\ell(x)} + \frac{\mathcal{T}_{x'}(S)}{\ell(x')} \geq \frac{1}{18}.$$

Proof. By the definition of our lazy transition operators $\{\mathcal{T}_x\}_{x \in \mathbb{R}^d}$, we have

$$\mathcal{T}_x(S^c) = \frac{\text{Vol}(S^c \cap \mathbb{B}(x, \eta))}{2\text{Vol}(\mathbb{B}(x, \eta))} \geq \frac{\text{Vol}(S^c \cap (\mathbb{B}(x, \eta) \cap \mathbb{B}(x', \eta)))}{2\text{Vol}(\mathbb{B}(x, \eta))},$$

and similarly,

$$\mathcal{T}_{x'}(S) \geq \frac{\text{Vol}(S \cap (\mathbb{B}(x, \eta) \cap \mathbb{B}(x', \eta)))}{2\text{Vol}(\mathbb{B}(x', \eta))}.$$

By summing these inequalities and applying Lemma 6, we hence have

$$\mathcal{T}_x(S^c) + \mathcal{T}_{x'}(S) \geq \frac{1}{8} \min\left\{\ell(x), \ell(x')\right\}.$$

Dividing both sides by $\ell(x)$ and using the assumptions then gives the claim. $\qquad\square$

5.2 Average local conductance

kjtian: Finished discussion in class, completion of notes TBD. Please check out [KLS97, JLLV21] for the remainder of the ideas discussed in class, and email me any questions. In the meantime, here is a nice isoperimetric result in that we showed in class.

Proposition 5. *Among all compact $\mathcal{K} \subseteq \mathbb{R}^d$ with a fixed volume $\text{Vol}(\mathcal{K})$, the minimum surface area is achieved when $\mathcal{K}$ is a ball of radius $(\frac{\text{Vol}(\mathcal{K})}{V_d})^{\frac{1}{d}}$.*

Proof. It is standard that another way to define the surface area of a set $\mathcal{K} \subseteq \mathbb{R}^d$ is

$$\text{Vol}(\partial \mathcal{K}) = \lim_{\epsilon \to 0^+} \frac{\text{Vol}(\mathcal{K} \oplus \mathbb{B}(\mathbb{0}_d, \epsilon)) - \text{Vol}(\mathcal{K})}{\epsilon},$$

where $\oplus$ is the Minkowski sum. The Brunn-Minkowski inequality (Theorem 5, Part I) then shows

$$\text{Vol}(\mathcal{K} \oplus \mathbb{B}(\mathbb{0}_d, \epsilon))^{\frac{1}{d}} \geq \text{Vol}(\mathcal{K})^{\frac{1}{d}} + \text{Vol}(\mathbb{B}(\mathbb{0}_d, \epsilon))^{\frac{1}{d}}$$

$$\implies \lim_{\epsilon \to 0^+} \frac{\text{Vol}(\mathcal{K} \oplus \mathbb{B}(\mathbb{0}_d, \epsilon)) - \text{Vol}(\mathcal{K})}{\epsilon} \geq \lim_{\epsilon \to 0^+} \frac{(\text{Vol}(\mathcal{K})^{\frac{1}{d}} + \text{Vol}(\mathbb{B}(\mathbb{0}_d, \epsilon))^{\frac{1}{d}})^d - \text{Vol}(\mathcal{K})}{\epsilon}$$

$$= d \cdot \text{Vol}(\mathcal{K})^{\frac{d-1}{d}} \text{Vol}(\mathbb{B}(\mathbb{0}_d, 1))^{\frac{1}{d}}.$$

It is well-known that $\text{Vol}(\partial \mathbb{B}(\mathbb{0}_d, 1)) = d \cdot \text{Vol}(\mathbb{B}(\mathbb{0}_d, 1))$. Hence, for $r := (\frac{\text{Vol}(\mathcal{K})}{V_d})^{\frac{1}{d}}$,

$$\frac{\text{Vol}(\partial \mathcal{K})}{\text{Vol}(\mathcal{K})} \geq d \cdot \left(\frac{V_d}{\text{Vol}(\mathcal{K})}\right)^{\frac{1}{d}} = \frac{\text{Vol}(\partial \mathbb{B}(\mathbb{0}_d, r))}{\text{Vol}(\mathbb{B}(\mathbb{0}_d, r))}.$$

$\qquad\square$

Source material

This lecture is primarily based on material in [LS93, KLS95, KLS97].

References

[AK91]   David L. Applegate and Ravi Kannan. Sampling and integration of near log-concave functions. In *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing*, pages 156–163. ACM, 1991.

[BF87]   Imre Bárány and Zoltán Füredi. Computing the volume is difficult. *Discrete and Computational Geometry*, 2:319–326, 1987.

[Bor75]   Christer Borell. Convex set functions in d-space. *Periodica Mathematica Hungarica*, 6(2):111–136, 1975.

[BT03]   Sergey G. Bobkov and Prasad Tetali. Modified log-sobolev inequalities, mixing and hypercontractivity. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, pages 287–296. ACM, 2003.

[CDWY20] Yuansi Chen, Raaz Dwivedi, Martin J. Wainwright, and Bin Yu. Fast mixing of metropolized hamiltonian monte carlo: Benefits of multi-step gradients. *J. Mach. Learn. Res.*, 21:92:1–92:72, 2020.

[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 S. Vempala. Gaussian cooling and $o(n^3)$ algorithms for volume and gaussian volume. *SIAM J. Comput.*, 47(3):1237–1273, 2018.

[DF88]   Martin E. Dyer and Alan M. Frieze. On the complexity of computing the volume of a polyhedron. *SIAM J. Comput.*, 17(5):967–974, 1988.

[DFK91]   Martin E. Dyer, Alan M. Frieze, and Ravi Kannan. A random polynomial time algorithm for approximating the volume of convex bodies. *J. ACM*, 38(1):1–17, 1991.

[Eld13]   Ronen Eldan. Thin shell implies spectral gap up to polylog via a stochastic localization scheme. *Geometric and Functional Analysis*, 23:532–569, 2013.

[Ele86]   György Elekes. A geometric inequality and the complexity of computing volume. *Discrete and Computational Geometry*, 1:289–292, 1986.

[FG04]   Matthieu Fradelizi and Olivier Guédon. The extreme points of subsets of s-concave probabilities and a geometric localization theorem. *Discrete & Computational Geometry*, 31(2):327–335, 2004.

[GLL+23] Sivakanth Gopi, Yin Tat Lee, Daogao Liu, Ruoqi Shen, and Kevin Tian. Private convex optimization in general norms. In *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023*, pages 5068–5089. SIAM, 2023.

[GMT06]   Sharad Goel, Ravi Montenegro, and Prasad Tetali. Mixing time bounds via the spectral profile. *Electronic Journal of Probability*, 11:1–26, 2006.

[JLLV21] He Jia, Aditi Laddha, Yin Tat Lee, and Santosh S. Vempala. Reducing isotropy and volume to KLS: an $o(n^3\psi^2)$ volume algorithm. In *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, 2021*, pages 961–974. ACM, 2021.

[Kla23]   Boaz Klartag. Logarithmic bounds for isoperimetry and slices of convex sets. *Ars Inveniendi Analytica*, 2023.

[KLS95]   Ravi Kannan, László Lovász, and Miklós Simonovits. Isoperimetric problems for convex bodies and a localization lemama. *Discret. Comput. Geom.*, 13:541–559, 1995.

[KLS97]   Ravi Kannan, László Lovász, and Miklós Simonovits. Random walks and an $o^*(n^5)$ volume algorithm for convex bodies. *Random Struct. Algorithms*, 11(1):1–50, 1997.

[Led99]   Michel Ledoux. *Concentration of measure and logarithmic Sobolev inequalities*. Seminaire de probabilities XXXIII, 1999.

[LK99]   László Lovász and Ravi Kannan. Faster mixing via average conductance. In *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing*, pages 282–287. ACM, 1999.

[LS92]   László Lovász and Miklós Simonovits. On the randomized complexity of volume and diameter. In *33rd Annual Symposium on Foundations of Computer Science*, pages 482–491. IEEE Computer Society, 1992.

[LS93]   László Lovász and Miklós Simonovits. Random walks in a convex body and an improved volume algorithm. *Random Struct. Algorithms*, 4(4):359–412, 1993.

[LST20]   Yin Tat Lee, Ruoqi Shen, and Kevin Tian. Logsmooth gradient concentration and tighter runtimes for metropolized hamiltonian monte carlo. In *Conference on Learning Theory, COLT 2020*, volume 125 of *Proceedings of Machine Learning Research*, pages 2565–2597. PMLR, 2020.

[LV06a]   László Lovász and Santosh S. Vempala. Fast algorithms for logconcave functions: Sampling, rounding, integration and optimization. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS) 2006*, pages 57–68. IEEE Computer Society, 2006.

[LV06b]   László Lovász and Santosh S. Vempala. Hit-and-run from a corner. *SIAM J. Comput.*, 35(4):985–1005, 2006.

[LV06c]   László Lovász and Santosh S. Vempala. Simulated annealing in convex bodies and an $o(n^4)$ volume algorithm. *J. Comput. Syst. Sci.*, 72(2):392–417, 2006.

[LV07]   László Lovász and Santosh S. Vempala. The geometry of logconcave functions and sampling algorithms. *Random Struct. Algorithms*, 30(3):307–358, 2007.

[LV17a]   Yin Tat Lee and Santosh Srinivas Vempala. Eldan's stochastic localization and the KLS hyperplane conjecture: An improved lower bound for expansion. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017*, pages 998–1007. IEEE Computer Society, 2017.

[LV17b]   Yin Tat Lee and Santosh Srinivas Vempala. The kannan-lovász-simonovits conjecture. *Current Developments in Mathematics*, 2017:1–36, 2017.

[Mon23]   Andrea Montanari. Sampling, diffusions, and stochastic localization. *CoRR*, abs/2305.10690, 2023.