# Geometric Random Walks: A Survey* 

Santosh Vempala ${ }^{\dagger}$<br>Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139.


#### Abstract

The developing theory of geometric random walks is outlined here. Three aspects - general methods for estimating convergence (the "mixing" rate), isoperimetric inequalities in $\mathbb{R}^{n}$ and their intimate connection to random walks, and algorithms for fundamental problems (volume computation and convex optimization) that are based on sampling by random walks - are discussed.


## 1 Introduction

A geometric random walk starts at some point in $\mathbb{R}^{n}$ and at each step, moves to a "neighboring" point chosen according to some distribution that depends only on the current point, e.g., a uniform random point within a fixed distance $\delta$. The sequence of points visited is a random walk. The distribution of the current point, in particular, its convergence to a steady state (or stationary) distribution, turns out to be a very interesting phenomenon. By choosing the 1step distribution appropriately, one can ensure that the steady state distribution is, for example, the uniform distribution over a convex body, or indeed any reasonable distribution in $\mathbb{R}^{n}$.

Geometric random walks are Markov chains, and the study of the existence and uniqueness of and the convergence to a steady state distribution is a classical field of mathematics. In the geometric setting, the dependence on the dimension (called $n$ in this survey) is of particular interest. Pólya proved that with probability 1 , a random walk on an $n$-dimensional grid returns to its starting point infinitely often for $n \leq 2$, but only a finite number of times for $n \geq 3$.

Random walks also provide a general approach to sampling a geometric distribution. To sample a given distribution, we set up a random walk whose steady state is the desired distribution. A random (or nearly random) sample is obtained by taking sufficiently many steps of the walk. Basic problems such as optimization and volume computation can be reduced to sampling. This

[^0]connection, pioneered by the randomized polynomial-time algorithm of Dyer, Frieze and Kannan [12] for computing the volume of a convex body, has lead to many new developments in recent decades.

In order for sampling by a random walk to be efficient, the distribution of the current point has to converge rapidly to its steady state. The first part of this survey (Section 3) deals with methods to analyze this convergence, and describes the most widely used method, namely, bounding the conductance, in detail. The next part of the survey is about applying this to geometric random walks and the issues that arise therein. Notably, there is an intimate connection with geometric isoperimetric inequalities. The classical isoperimetric inequality says that among all measurable sets of fixed volume, a ball of this volume is the one that minimizes the surface area. Here one is considering all measurable sets. In contrast, we will encounter the following type of question: Given a convex set $K$, and $0<t<1$, what subset $S$ of volume $t \operatorname{vol}(K)$ has the smallest surface area inside $K$ (i.e., not counting the boundary of $S$ that is part of the boundary of $K)$ ? The inequalities that result are interesting in their own right.

The last two sections describe polynomial-time algorithms for minimizing a quasi-convex function over a convex body and for computing the volume of a convex body. The application to volume computation is rather remarkable in the light of results that no deterministic polynomial-time algorithm can approximate the volume to within an exponential (in $n$ ) factor. In Section 9 , we briefly discuss the history of the problem and describe the latest algorithm.

Several topics related to this survey have been addressed in detail in the literature. For a general introduction to discrete random walks, the reader is referred to the survey by Lovász [31] or the book by Aldous and Fill [1]. For an in-depth account of the volume problem that includes all but the most recent improvements, there is a survey by Simonovits [44] and an earlier article by Bollobás [7].

### 1.1 Three walks

Before we introduce various concepts and tools, let us state precisely three different ways to walk randomly inside a convex body $K$ in $\mathbb{R}^{n}$. It might be useful to keep these examples in mind. Later, we will see generalizations of these walks.

The Grid Walk is restricted to a discrete subset of $K$, namely, all points in $K$ whose coordinates are integer multiples of a fixed parameter $\delta$. These points form a grid, and the neighbors of a grid point $x$ are the points reachable by changing one coordinate by $\pm \delta$. Let $e_{1}, \ldots, e_{n}$ denote the coordinate vectors in $\mathbb{R}^{n}$; then the neighbors of a grid point $x$ are $\left\{x \pm \delta e_{i}\right\}$. The grid walk tries to move to a random neighboring point.

## Grid Walk ( $\delta$ )

- Pick a grid point $y$ uniformly at random from the neighbors of the current point $x$.
- If $y$ is in $K$, go to $y$; else stay at $x$.

The Ball Walk tries to step to a random point within distance $\delta$ of the current point. Its state space is the entire set $K$.

```
Ball Walk ( }\delta\mathrm{ )
    - Pick a uniform random point y from the ball of radius }
        centered at the current point }x\mathrm{ .
    - If }y\mathrm{ is in K, go to }y\mathrm{ ; else stay at }x\mathrm{ .
```

Hit-and-run picks a random point along a random line through the current point. It does not need a "step-size" parameter. The state space is again all of $K$.

## Hit-and-run

- Pick a uniform random line $\ell$ through the current point $x$.
- Go to a uniform random point on the chord $\ell \cap K$.

To implement the first step of hit-and-run, we can generate $n$ independent random numbers, $u_{1}, \ldots, u_{n}$ each from the standard Normal distribution, and then the direction of the vector $u=\left(u_{1}, \ldots, u_{n}\right)$ is uniformly distributed. For the second step, using the membership oracle for $K$, we find an interval $[a, b]$ that contains the chord through $x$ parallel to $u$ so that $|a-b|$ is at most twice (say) the length of the chord (this can be done by a binary search with a logarithmic overhead). Then pick random points in $[a, b]$ till we find one in $K$.

For the first step of the ball walk, in addition to a random direction $u$, we generate a number $r$ in the range $[0, \delta]$ with density $f(x)$ proportional to $x^{n-1}$ and then $z=r u /|u|$ is uniformly distributed in a ball of radius $\delta$.

Do these random walks converge to a steady state distribution? If so, what is it? How quickly do they converge? How does the rate of convergence depend on the convex body $K$ ?

These are some of the questions that we will address in analyzing the walks.

## 2 Basic definitions

### 2.1 Markov chains

A Markov chain is defined using a $\sigma$-algebra $(K, \mathcal{A})$, where $K$ is the state space and $\mathcal{A}$ is a set of subsets of $K$ that is closed under complements and countable unions. For each element $u$ of $K$, we have a probability measure $P_{u}$ on $(K, \mathcal{A})$, i.e., each set $A \in \mathcal{A}$ has a probability $P_{u}(A)$. Informally, $P_{u}$ is the distribution obtained on taking one step from $u$. The triple ( $K, \mathcal{A},\left\{P_{u}: u \in K\right\}$ ) along with a starting distribution $Q_{0}$ defines a Markov chain, i.e., a sequence of elements of $K, w_{0}, w_{1}, \ldots$, where $w_{0}$ is chosen from $Q_{0}$ and each subsequent $w_{i}$ is chosen from $P_{w_{i-1}}$. Thus, the choice of $w_{i+1}$ depends only on $w_{i}$ and is independent of $w_{0}, \ldots, w_{i-1}$.

A distribution $Q$ on $(K, \mathcal{A})$ is called stationary if one step from it gives the same distribution, i.e., for any $A \in \mathcal{A}$,

$$
\int_{A} P_{u}(A) d Q(u)=Q(A)
$$

A distribution $Q$ is atom-free if there is no $x \in K$ with $Q(x)>0$.
The ergodic flow of subset $A$ w.r.t. the distribution $Q$ is defined as

$$
\Phi(A)=\int_{A} P_{u}(K \backslash A) d Q(u) .
$$

It is easy to verify that a distribution $Q$ is stationary iff $\Phi(A)=\Phi(K \backslash A)$. The existence and uniqueness of the stationary distribution $Q$ for general Markov chains is a rather technical issue that is not covered in this survey (the reader is referred to the book by Revuz [42]) ${ }^{1}$. In all the chains we study in this survey, the stationary distribution will be given explicitly and can be easily verified. To avoid the issue of uniqueness of the stationary distribution, we only consider lazy Markov chains. In a lazy version of a given Markov chain, at each step, with probability $1 / 2$, we do nothing; with the rest we take a step according to the Markov chain. The next theorem is folklore and will also be implied by convergence theorems that we present later.

Theorem 2.1 If $Q$ is stationary w.r.t. a lazy Markov chain then it is the unique stationary distribution for that Markov chain.

For some additional properties of lazy Markov chains, see Section 1 of [34]. We will hereforth assume that the distribution in the definition of $\Phi$ is the unique stationary distribution.

The conductance of a subset $A$ is defined as

$$
\phi(A)=\frac{\Phi(A)}{\min \{Q(A), Q(K \backslash A)\}}
$$

[^1]and the conductance of the Markov chain is
$$
\phi=\min _{A} \phi(A)=\min _{0<Q(A) \leq \frac{1}{2}} \frac{\int_{A} P_{u}(K \backslash A) d Q(u)}{Q(A)} .
$$

The local conductance of an element $u$ is $\ell(u)=1-P_{u}(\{u\})$.
The following weaker notion of conductance will also be useful. For any $0 \leq s<\frac{1}{2}$, the $s$-conductance of a Markov chain is defined as

$$
\phi_{s}=\min _{A: s<Q(A) \leq \frac{1}{2}} \frac{\Phi(A)}{Q(A)-s}
$$

### 2.2 Comparing distributions

We will often have to compare two distributions $P$ and $Q$ (typically, the distribution of the current point and the stationary distribution). There are many reasonable ways to do this. Here are three that will come up.

1. Total variation distance of $P$ and $Q$ is

$$
\|P-Q\|_{t v}=\sup _{A \in \mathcal{A}}|P(A)-Q(A)|
$$

2. $L_{2}$ distance of $P$ with respect to $Q$ is

$$
\|P / Q\|=\int_{K} \frac{d P(u)}{d Q(u)} d P(u)=\int_{K}\left(\frac{d P(u)}{d Q(u)}\right)^{2} d Q(u)
$$

3. $P$ is said to be $M$-warm w.r.t. $Q$ if

$$
M=\sup _{A \in \mathcal{A}} \frac{P(A)}{Q(A)}
$$

If $Q_{0}$ is $O(1)$-warm w.r.t. the stationary distribution $Q$ for some Markov chain, then we say that $Q_{0}$ is a warm start for $Q$.

### 2.3 Convexity

Convexity plays a key role in the convergence of geometric random walks. The following notation/concepts will be used.

The unit ball in $\mathbb{R}^{n}$ is $B_{n}$ and its volume is $\operatorname{vol}\left(B_{n}\right)$. For two subsets $A, B$ of $\mathbb{R}^{n}$, their Minkowski sum is

$$
A+B=\{x+y: x \in A, y \in B\}
$$

The Brunn-Minkowski theorem says that if $A, B$ and $A+B$ are measurable, then

$$
\begin{equation*}
\operatorname{vol}(A+B)^{\frac{1}{n}} \geq \operatorname{vol}(A)^{\frac{1}{n}}+\operatorname{vol}(B)^{\frac{1}{n}} \tag{1}
\end{equation*}
$$

Recall that a subset $S$ of $\mathbb{R}^{n}$ is convex if for any two points $x, y \in S$, the interval $[x, y] \subseteq S$. A function $f: \mathbb{R}^{n} \rightarrow \mathbb{R}_{+}$is said to be logconcave if for any two points $x, y \in \mathbb{R}^{n}$ and any $\lambda \in[0,1]$,

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

The product and the minimum of two logconcave functions are both logconcave; the sum is not in general. The following fundamental properties, proved by Dinghas [9], Leindler [27] and Prékopa [41, 40], are often useful.

Theorem 2.2 All marginals as well as the distribution function of a logconcave function are logconcave. The convolution of two logconcave functions is logconcave.

Logconcave functions have many properties that are reminiscent of convexity. For a logconcave density $f$, we denote the induced measure by $\pi_{f}$ and its centroid by $z_{f}=\mathrm{E}_{f}(X)$. The second moment of $f$ refers to $\mathrm{E}_{f}\left(\left|X-z_{f}\right|^{2}\right)$. The next three lemmas are chosen for illustration from [35]. The first one was proved earlier by Grünbaum [16] for the special case of the uniform density over a convex body. We will later see a further refinement of this lemma that is useful for optimization.

Lemma 2.3 Let $f: \mathbb{R}^{n} \rightarrow \mathbb{R}_{+}$be a logconcave density function, and let $H$ be any halfspace containing its centroid. Then

$$
\int_{H} f(x) d x \geq \frac{1}{e}
$$

Lemma 2.4 If $X$ is drawn from a logconcave distribution in $\mathbb{R}^{n}$, then for any integer $k>0$,

$$
\mathrm{E}\left(|X|^{k}\right)^{1 / k} \leq 2 k \mathrm{E}(|X|)
$$

Note that this can be viewed as a converse to Hölder's inequality which says that

$$
\mathrm{E}\left(|X|^{k}\right)^{1 / k} \geq \mathrm{E}(|X|)
$$

Lemma 2.5 Let $X \in \mathbb{R}^{n}$ be a random point from a logconcave distribution with second moment $R^{2}$. Then for any $t>1, \mathrm{P}(|X|>t R)<e^{-t}$.

A density function $f: \mathbb{R}^{n} \rightarrow \mathbb{R}_{+}$is said to be isotropic, if its centroid is the origin, and its covariance matrix is the identity matrix. This latter condition can be expressed in terms of the coordinate functions as

$$
\int_{\mathbb{R}^{n}} x_{i} x_{j} f(x) d x=\delta_{i j}
$$

for all $1 \leq i, j \leq n$. This condition is equivalent to saying that for every vector $v \in \mathbb{R}^{n}$,

$$
\int_{\mathbb{R}^{n}}\left(v^{T} x\right)^{2} f(x) d x=|v|^{2}
$$

In terms of the associated random variable $X$, this means that

$$
\mathrm{E}(X)=0 \quad \text { and } \quad \mathrm{E}\left(X X^{T}\right)=I
$$

We say that $f$ is near-isotropic up to a factor of $C$ or $C$-isotropic, if

$$
\frac{1}{C} \leq \int_{\mathbb{R}^{n}}\left(v^{T} x\right)^{2} d \pi_{f}(x) \leq C
$$

for every unit vector $v$. The notions of "isotropic" and "near-isotropic" extend to non-negative integrable functions $f$, in which case we mean that the density function $f / \int_{\mathbb{R}^{n}} f$ is isotropic. For any full-dimensional integrable function $f$ with bounded second moment, there is an affine transformation of the space bringing it to isotropic position, and this transformation is unique up to an orthogonal transformation of the space. Indeed if $f$ is not isotropic, we can make the centroid be the origin by a translation. Next, compute $A=\mathrm{E}\left(X X^{T}\right)$ for the associated random variable $X$. Now $A$ must be positive semi-definite (since each $X X^{T}$ is) and so we can write $A=B B^{T}$ for some matrix $B$. Then the transformation $B^{-1}$ makes $f$ isotropic.

It follows easily that for an isotropic distribution in $\mathbb{R}^{n}$, the second moment is

$$
\mathrm{E}\left(|X|^{2}\right)=\sum_{i} \mathrm{E}\left(X_{i}^{2}\right)=n
$$

Further, Lemma 2.5 implies that for an isotropic logconcave distribution $f$,

$$
\mathrm{P}(X>t \sqrt{n})<e^{-t}
$$

which means that most of $f$ is contained in a ball of radius $O(\sqrt{n})$, and this is sometimes called its effective diameter.

### 2.4 Computational model

If the input to an algorithm is a convex body $K$ in $\mathbb{R}^{n}$, we assume that it is given by a membership oracle which on input $x \in \mathbb{R}^{n}$ returns "YES" if $x \in K$ and "NO" otherwise. In addition we will have some bounds on $K$ - typically, $B_{n} \subseteq K \subseteq R B_{n}$, i.e., $K$ contains a unit ball around the origin and is contained in a ball of given radius. It is enough to have any point $x$ in $K$ and the guarantee that a ball of radius $r$ around $x$ is contained in $K$ and one of radius $R$ contains $K$ (by translation and scaling this is equivalent to the previous condition). Sometimes, we will need a separation oracle for $K$, i.e., a procedure which either verifies that a given point $x$ is in $K$ or returns a hyperplane that separates $x$ from $K$. The complexity of the algorithm will be measured mainly by the number of oracle queries, but we will also keep track of the number of arithmetic operations.

In the case of a logconcave density $f$, we have an oracle for $f$, i.e., for any point $x$ it returns $C f(x)$ where $C$ is an unknown parameter independent of $x$. This is useful when we know a function proportional to the desired density, but
not its integral, e.g., in the case of the uniform density over a bounded set, all we need is the indicator function of the support. In addition, we have a guarantee that the centroid of $f$ satisfies $\left|z_{f}\right|^{2}<Z$ and the eigenvalues of the covariance matrix of $f$ are bounded from above and below by two given numbers. Again, the complexity is measured by the number of oracle calls. We will say that an algorithm is efficient if its complexity is polynomial in the relevant parameters.

To emphasize the essential dependence on the dimension we will sometimes use the $O^{*}($.$) notation which suppresses logarithmic factors and also the depen-$ dence on error parameters. E.g., $n \log n / \varepsilon=O^{*}(n)$.

### 2.5 Examples

For the ball walk in a convex body, the state space $K$ is the convex body, and $\mathcal{A}$ is the set of all measurable subsets of $K$. Further,

$$
P_{u}(\{u\})=1-\frac{\operatorname{vol}\left(K \cap\left(u+\delta B_{n}\right)\right)}{\operatorname{vol}\left(\delta B_{n}\right)}
$$

and for any measurable subset $A$, such that $u \notin A$,

$$
P_{u}(A)=\frac{\operatorname{vol}\left(A \cap\left(u+\delta B_{n}\right)\right)}{\operatorname{vol}\left(\delta B_{n}\right)}
$$

If $u \in A$, then

$$
P_{u}(A)=P_{u}(A \backslash\{u\})+P_{u}(\{u\})
$$

It is straightforward to verify that the uniform distribution is stationary, i.e.,

$$
Q(A)=\frac{\operatorname{vol}(A)}{\operatorname{vol}(K)}
$$

For hit-and-run, the one-step distribution for a step from $u \in K$ is given as follows. For any measurable subset $A$ of $K$,

$$
\begin{equation*}
P_{u}(A)=\frac{2}{\operatorname{vol}_{n-1}\left(\partial B_{n}\right)} \int_{A} \frac{d x}{\ell(u, x)|x-u|^{n-1}} \tag{2}
\end{equation*}
$$

where $\ell(u, x)$ is the length of the chord in $K$ through $u$ and $x$. The uniform distribution is once again stationary. One way to see this is to note that the one-step distribution has a density function and the density of stepping from $u$ to $v$ is the same as that for stepping from $v$ to $u$.

These walks can be modified to sample much more general distributions. Let $f: \mathbb{R}^{n} \rightarrow \mathbb{R}_{+}$be a nonnegative integrable function. It defines a measure $\pi_{f}$ (on measurable subsets of $\mathbb{R}^{n}$ ):

$$
\pi_{f}(A)=\frac{\int_{A} f(x) d x}{\int_{\mathbb{R}^{n}} f(x) d x}
$$

The following extension of the ball walk, usually called the ball walk with a Metropolis filter has $\pi_{f}$ as its stationary distribution (it is a simple exercise to prove, but quite nice that this works).

## Ball walk with Metropolis filter ( $\delta, f$ )

- Pick a uniformly distributed random point $y$ in the ball of radius $\delta$ centered at the current point $x$.
- Move to $y$ with probability $\min \left\{1, \frac{f(y)}{f(x)}\right\}$; stay at $x$ with the remaining probability.

Hit-and-run can also be extended to sampling from such a general distribution $\pi_{f}$. For any line $\ell$ in $\mathbb{R}^{n}$, let $\pi_{\ell, f}$ be the restriction of $\pi$ to $\ell$, i.e.,

$$
\pi_{\ell, f}(S)=\frac{\int_{p+t u \in S} f(p+t u) d t}{\int_{\ell} f(x) d x}
$$

where $p$ is any point on $\ell$ and $u$ is a unit vector parallel to $\ell$.

```
Hit-and-run (f)
    - Pick a uniform random line \ell through the current point x.
    - Go to a random point }y\mathrm{ along }\ell\mathrm{ chosen from the
        distribution }\mp@subsup{\pi}{\ell,f}{}\mathrm{ .
```

Once again, it is easy to verify that $\pi_{f}$ is the stationary distribution for this walk. One way to carry out the second step is to use a binary search to find the point $p$ on $\ell$ where the function is maximal, and the points $a$ and $b$ on both sides of $p$ on $\ell$ where the value of the function is $\varepsilon f(p)$. We allow a relative error of $\varepsilon$, so the number of oracle calls is only $O(\log (1 / \varepsilon))$. Then select a uniformly distributed random point $y$ on the segment $[a, b]$, and independently a uniformly distributed random real number in the interval $[0,1]$. Accept $y$ if $f(y)>r f(p)$; else, reject $y$ and repeat.

## 3 Convergence and conductance

So far we have seen that random walks can be designed to approach any reasonable distribution in $\mathbb{R}^{n}$. For this to lead to an efficient sampling method, the convergence to the stationary distribution must be fast. This section is devoted to general methods for bounding the rate of convergence.

One way to define the mixing rate of a random walk is the number of steps required to reduce some measure of the distance of the current distribution to the stationary distribution by a factor of 2 (e.g., one of the distance measures from Section 2.2). We will typically use the total variation distance. For a discrete random walk (i.e., the state space is a finite set), the mixing rate is
characterized by the eigenvalues gap of the transition matrix $P$ whose $i j$ th entry is the probability of stepping from $i$ to $j$, conditioned on currently being at $i$. Let $\lambda_{1} \geq \lambda_{2} \ldots \geq \lambda_{n}$ be the eigenvalues of $P$. The top eigenvalue is 1 (by the definition of stationarity) and let $\lambda=\max \left\{\lambda_{2},\left|\lambda_{n}\right|\right\}$ (in the lazy version of any walk, all the $\lambda_{i}$ are nonnegative and $\lambda=\lambda_{2}$ ). Then, for a random walk starting at the point $x$, with $Q_{t}$ being the distribution after $t$ steps, the following bound on the convergence can be derived (see e.g., [31]). For any point $y \in K$,

$$
\begin{equation*}
\left|Q_{t}(y)-Q(y)\right| \leq \sqrt{\frac{Q(y)}{Q(x)}} \lambda^{t} \tag{3}
\end{equation*}
$$

Estimating $\lambda$ can be difficult or impractical even in the discrete setting (if e.g., the state space is too large to write down $P$ explicitly).

Intuitively, a random walk will "mix" slowly if it has a bottleneck, i.e., a partition $S, K \backslash S$ of its state space, such that the probability of stepping from $S$ to $K \backslash S$ (the ergodic flow out of $S$ ) is small compared to the measures of $S$ and $K \backslash S$. Note that this ratio is precisely the conductance of $S, \phi(S)$. It takes about $1 / \phi(S)$ steps in expectation to even go from one side to the other. As we will see in this section, the mixing rate is bounded from above by $2 / \phi^{2}$. Thus, conductance captures the mixing rate upto a quadratic factor. This was first proved for discrete Markov chains by Jerrum and Sinclair [17] who showed that conductance can be related to the eigenvalue gap of the transition matrix. A similar relationship for a related quantity called expansion was found by Alon [2] and by Dodziuk and Kendall [10]. The inequality below is a discrete analogue of Cheeger's inequality in differential geometry.
Theorem 3.1

$$
\frac{\phi^{2}}{2} \leq 1-\lambda \leq 2 \phi
$$

As a consequence of this and (3), we get that for a discrete random walk starting at $x$, and any point $y \in K$,

$$
\begin{equation*}
\left|Q_{t}(y)-Q(y)\right| \leq \sqrt{\frac{Q(y)}{Q(x)}}\left(1-\frac{\phi^{2}}{2}\right)^{t} \tag{4}
\end{equation*}
$$

For the more general continuous setting, Lovász and Simonovits [34] proved the connection between conductance and convergence. Their proof does not use eigenvalues. We will sketch it here since it is quite insightful, but does not seem to be well-known. It also applies to situations where the conductance can be bounded only for subsets of bounded size (i.e., the $s$-conductance, $\phi_{s}$, can be bounded from below for some $s>0$ ). We remind the reader that we have assumed that our Markov chains are lazy.

To show convergence, we need to prove that $\left|Q_{t}(A)-Q(A)\right|$ falls with $t$ for every measurable subset $A$ of $K$. However, this quantity might converge at different rates for different subsets. So we consider

$$
\sup _{A: Q(A)=x} Q_{t}(A)-Q(A)
$$

for each $x \in[0,1]$. A bound for every $x$ would imply what we want. To prove inductively that this quantity decreases with $t$, Lovász and Simonovits define the following formal upper bound. Let $\mathcal{G}_{x}$ be the set of functions defined as

$$
\mathcal{G}_{x}=\left\{g: K \rightarrow[0,1]: \int_{u \in K} g(u) d Q(u)=x\right\}
$$

Using this, define

$$
h_{t}(x)=\sup _{g \in \mathcal{G}_{x}} \int_{u \in K} g(u)\left(d Q_{t}(u)-d Q(u)\right)=\sup _{g \in \mathcal{G}_{x}} \int_{u \in K} g(u) d Q_{t}(u)-x
$$

It is clear that for $A$ with $Q(A)=x, h_{t}(x) \geq Q_{t}(A)-Q(A)$ since the indicator function of $A$ is in $\mathcal{G}_{x}$. The function $h_{t}(x)$ has the following properties.

Lemma 3.2 For any positive integer $t$,
a. The function $h_{t}$ is concave.
b. If $Q$ is atom-free, then $h_{t}(x)=\sup _{A: Q(A)=x} Q_{t}(A)-Q(A)$ and the supremum is achieved by some subset.
c. Let $Q$ be atom-free and $t \geq 1$. For any $0 \leq x \leq 1$, let $y=\min \{x, 1-x\}$. Then,

$$
h_{t}(x) \leq \frac{1}{2} h_{t-1}(x-2 \phi y)+\frac{1}{2} h_{t-1}(x+2 \phi y) .
$$

The first part of the lemma is easily verified. We sketch the second part: to maximize $h_{t}$, we should use a function $g$ that puts high weight on points $u$ with $d Q_{t}(u) / d Q(u)$ as high as possible. Let $A$ be a subset with $Q(A)=x$, so that for any point $y$ not in $A$, the value of $d Q_{t}(y) / d Q(y)$ is no more than the value for any point in $A$ (i.e., $A$ consists of the top $x$ fraction of points according to $\left.d Q_{t}(u) / d Q(u)\right)$. Let $g$ be the corresponding indicator function. These points give the maximum payoff per unit of weight, so it is optimal to put as much weight on them as possible. We mention in passing that the case when $Q$ has atoms is a bit more complicated, namely we might need to include one atom fractionally (so that $Q(A)=x$ ). In the general case, $h_{t}(x)$ can be achieved by a function $g$ that is $0-1$ valued everywhere except for at most one point.

The third part of the lemma, which is the key to convergence, is proved below.
Proof. (of Lemma 3.2(c)) Assume that $0 \leq x \leq \frac{1}{2}$. The other range is proved in a similar way. We will construct two functions, $g_{1}$ and $g_{2}$, and use these to bound $h_{t}(x)$. Let $A$ be a subset to be chosen later with $Q(A)=x$. Let

$$
g_{1}(u)=\left\{\begin{array}{ll}
2 P_{u}(A)-1 & \text { if } u \in A, \\
0 & \text { if } u \notin A,
\end{array} \quad \text { and } \quad g_{2}(u)= \begin{cases}1 & \text { if } u \in A \\
2 P_{u}(A) & \text { if } u \notin A\end{cases}\right.
$$

First, note that $\frac{1}{2}\left(g_{1}+g_{2}\right)(u)=P_{u}(A)$ for all $u \in K$, which means that

$$
\begin{align*}
\frac{1}{2} \int_{u \in K} g_{1}(u) d Q_{t-1}(u)+\frac{1}{2} \int_{u \in K} g_{2}(u) d Q_{t-1}(u) & =\int_{u \in K} P_{u}(A) d Q_{t-1}(u) \\
& =Q_{t}(A) \tag{5}
\end{align*}
$$

By the laziness of the walk $\left(P_{u}(A) \geq \frac{1}{2}\right.$ iff $\left.u \in A\right)$, the range of the functions $g_{1}$ and $g_{2}$ is between 0 and 1 and if we let

$$
x_{1}=\int_{u \in K} g_{1}(u) d Q(u) \quad \text { and } \quad x_{2}=\int_{u \in K} g_{2}(u) d Q(u)
$$

then $g_{1} \in \mathcal{G}_{x_{1}}$ and $g_{2} \in \mathcal{G}_{x_{2}}$. Further,

$$
\begin{aligned}
\frac{1}{2}\left(x_{1}+x_{2}\right) & =\frac{1}{2} \int_{u \in K} g_{1}(u) d Q(u)+\frac{1}{2} \int_{u \in K} g_{2}(u) d Q(u) \\
& =\int_{u \in K} P_{u}(A) d Q(u)=Q(A)=x
\end{aligned}
$$

since $Q$ is stationary. Next, since $Q$ is atom-free, there is a subset $A \subseteq K$ that achieves $h_{t}(x)$. Using this and (5),

$$
\begin{aligned}
h_{t}(x) & =Q_{t}(A)-Q(A) \\
& =\frac{1}{2} \int_{u \in K} g_{1}(u) d Q_{t-1}(u)+\frac{1}{2} \int_{u \in K} g_{2}(u) d Q_{t-1}(u)-Q(A) \\
& =\frac{1}{2} \int_{u \in K} g_{1}(u)\left(d Q_{t-1}(u)-d Q(u)\right)+\frac{1}{2} \int_{u \in K} g_{2}(u)\left(d Q_{t-1}(u)-d Q(u)\right) \\
& \leq \frac{1}{2} h_{t-1}\left(x_{1}\right)+\frac{1}{2} h_{t-1}\left(x_{2}\right)
\end{aligned}
$$

We already know that $x_{1}+x_{2}=2 x$. In fact, $x_{1}$ and $x_{2}$ are separated from $x$.

$$
\begin{aligned}
x_{1} & =\int_{u \in K} g_{1}(u) d Q(u) \\
& =2 \int_{u \in A} P_{u}(A) d Q(u)-\int_{u \in A} d Q(u) \\
& =2 \int_{u \in A}\left(1-P_{u}(K \backslash A)\right) d Q(u)-x \\
& =x-2 \int_{u \in A} P_{u}(K \backslash A) d Q(u) \\
& =x-2 \Phi(A) \\
& \leq x-2 \phi x \\
& =x(1-2 \phi)
\end{aligned}
$$

(In the penultimate step, we used the fact that $x \leq \frac{1}{2}$.) Thus we have,

$$
x_{1} \leq x(1-2 \phi) \leq x \leq x(1+2 \phi) \leq x_{2}
$$



Figure 1: Bounding $h_{t}$.

Since $h_{t-1}$ is concave, the chord from $x_{1}$ to $x_{2}$ on $h_{t-1}$ lies below the chord from $x(1-2 \phi)$ to $x(1+2 \phi)$ (see Figure 1). Therefore,

$$
h_{t}(x) \leq \frac{1}{2} h_{t-1}(x(1-2 \phi))+\frac{1}{2} h_{t-1}(x(1+2 \phi))
$$

which is the conclusion of the lemma.
In fact, a proof along the same lines implies the following generalization of part (c).

Lemma 3.3 Let $Q$ be atom-free and $0 \leq s \leq 1$. For any $s \leq x \leq 1-s$, let $y=\min \{x-s, 1-x-s\}$. Then for any integer $t>0$,

$$
h_{t}(x) \leq \frac{1}{2} h_{t-1}\left(x-2 \phi_{s} y\right)+\frac{1}{2} h_{t-1}\left(x+2 \phi_{s} y\right) .
$$

Given some information about $Q_{0}$, we can now bound the rate of convergence to the stationary distribution. We assume that $Q$ is atom-free in the next theorem and its corollary. These results can be extended to the case when $Q$ has atoms with slightly weaker bounds [34].

Theorem 3.4 Let $0 \leq s \leq 1$ and $C_{0}$ and $C_{1}$ be such that

$$
h_{0}(x) \leq C_{0}+C_{1} \min \{\sqrt{x-s}, \sqrt{1-x-s}\} .
$$

Then

$$
h_{t}(x) \leq C_{0}+C_{1} \min \{\sqrt{x-s}, \sqrt{1-x-s}\}\left(1-\frac{\phi_{s}^{2}}{2}\right)^{t}
$$

Proof. By induction on $t$. The inequality is true for $t=0$ by the hypothesis. Now, suppose it holds for all values less than $t$. Assume $s=0$ (for convenience)
and w.l.o.g. that $x \leq 1 / 2$. From Lemma 3.3, we know that

$$
\begin{aligned}
h_{t}(x) & \leq \frac{1}{2} h_{t-1}(x(1-2 \phi))+\frac{1}{2} h_{t-1}(x(1+2 \phi)) \\
& \leq C_{0}+\frac{1}{2} C_{1}(\sqrt{x(1-2 \phi)}+\sqrt{x(1+2 \phi)})\left(1-\frac{\phi^{2}}{2}\right)^{t-1} \\
& =C_{0}+\frac{1}{2} C_{1} \sqrt{x}(\sqrt{1-2 \phi}+\sqrt{1+2 \phi})\left(1-\frac{\phi^{2}}{2}\right)^{t-1} \\
& \leq C_{0}+C_{1} \sqrt{x}\left(1-\frac{\phi^{2}}{2}\right)^{t}
\end{aligned}
$$

Here we have used $\sqrt{1-2 \phi}+\sqrt{1+2 \phi} \leq 2\left(1-\frac{\phi^{2}}{2}\right)$.
The following corollary, about convergence from various types of "good" starting distributions, gives concrete implications of the theorem.

Corollary 3.5 a. Let $M=\sup _{A} Q_{0}(A) / Q(A)$. Then,

$$
\left\|Q_{t}-Q\right\|_{t v} \leq \sqrt{M}\left(1-\frac{\phi^{2}}{2}\right)^{t}
$$

b. Let $0<s \leq \frac{1}{2}$ and $H_{s}=\sup \left\{\left|Q_{0}(A)-Q(A)\right|: Q(A) \leq s\right\}$. Then,

$$
\left\|Q_{t}-Q\right\|_{t v} \leq H_{s}+\frac{H_{s}}{s}\left(1-\frac{\phi_{s}^{2}}{2}\right)^{t}
$$

c. Let $M=\left\|Q_{0} / Q\right\|$. Then for any $\varepsilon>0$,

$$
\left\|Q_{t}-Q\right\|_{t v} \leq \varepsilon+\sqrt{\frac{M}{\varepsilon}}\left(1-\frac{\phi^{2}}{2}\right)^{t}
$$

Proof. The first two parts are straightforward. For the third, observe that the $L_{2}$ norm,

$$
\left\|Q_{0} / Q\right\|=\mathrm{E}_{Q_{0}}\left(\frac{d Q_{0}(x)}{d Q(x)}\right)
$$

So, for $1-\varepsilon$ of $Q_{0}, d Q_{0}(x) / d Q(x) \leq M / \varepsilon$. We can view the starting distribution as being generated as follows: with probability $1-\varepsilon$ it is a distribution with $\sup Q_{0}(A) / Q(A) \leq M / \varepsilon$; with probability $\varepsilon$ it is some other distribution. Now using part (a) implies part (c).

Conductance and $s$-conductance are not the only known ways to bound the rate of convergence. Kannan et al.[20, 21] have extended conductance to the notion of blocking conductance which is a certain type of average of the conductance over various subset sizes. In some cases, it provides a sharper bound than conductance. Let $\phi(x)$ denote the minimum conductance over all subsets of measure $x$. Then one version of their main theorem is the following.

Theorem 3.6 Let $\pi_{0}$ be the measure of the starting point. Then, after

$$
t>C \ln \left(\frac{1}{\varepsilon}\right) \int_{\pi_{0}}^{\frac{1}{2}} \frac{d x}{x \phi(x)^{2}}
$$

steps, where $C$ is an absolute constant, we have $\left\|Q_{t}-Q\right\| \leq \varepsilon$.
The theorem can be extended to continuous Markov chains. Another way to bound convergence which we do not describe here is via the log-Sobolev inequalities [8].

## 4 Isoperimetry

How to bound the conductance of a geometric random walk? To show that the conductance is large, we have to prove that for any subset $A \subset K$, the probability that a step goes out of $A$ is large compared to $Q(A)$ and $Q(K \backslash A)$. To be concrete, consider the ball walk. For any particular subset $S$, the points that are likely to "cross over" to $K \backslash S$ are those that are "near" the boundary of $S$ inside $K$. So, showing that $\phi(S)$ is large seems to be closely related to showing that there is a large volume of points near the boundary of $S$ inside $K$. This section is devoted to inequalities which will have such implications and will play a crucial role in bounding the conductance.

To formulate an isoperimetric inequality for convex bodies, we consider a partition of a convex body $K$ into three sets $S_{1}, S_{2}, S_{3}$ such that $S_{1}$ and $S_{2}$ are "far" from each other, and the inequality bounds the minimum possible volume of $S_{3}$ relative to the volumes of $S_{1}$ and $S_{2}$. We will consider different notions of distance between subsets. Perhaps the most basic is the Euclidean distance:

$$
d\left(S_{1}, S_{2}\right)=\min \left\{|u-v|: u \in S_{1}, v \in S_{2}\right\}
$$

Suppose $d\left(S_{1}, S_{2}\right)$ is large. Does this imply that the volume of $S_{3}=K \backslash\left(S_{1} \cup S_{2}\right)$ is large? The classic counterexample to such a theorem is a dumbbell - two large subsets separated by very little. Of course, this is not a convex set!

The next theorem, proved in [11] (improving on a theorem in [33] by a factor of 2 ; see also [34]) asserts that the answer is yes.

Theorem 4.1 Let $S_{1}, S_{2}, S_{3}$ be a partition into measurable sets of a convex body $K$ of diameter $D$. Then,

$$
\operatorname{vol}\left(S_{3}\right) \geq \frac{2 d\left(S_{1}, S_{2}\right)}{D} \min \left\{\operatorname{vol}\left(S_{1}\right), \operatorname{vol}\left(S_{2}\right)\right\}
$$

A limiting version of this inequality is the following: For any subset $S$ of a convex body of diameter $D$,

$$
\operatorname{vol}_{n-1}(\partial S \cap K) \geq \frac{2}{D} \min \{\operatorname{vol}(S), \operatorname{vol}(K \backslash S)\}
$$

which says that the surface area of $S$ inside $K$ is large compared to the volumes of $S$ and $K \backslash S$. This is in direct analogy with the classical isoperimetric inequality, which says that the surface area to volume ratio of any measurable set is at least the ratio for a ball.

How does one prove such an inequality? In what generality does it hold? (i.e., for what measures besides the uniform measure on a convex set?) We will address these questions in this section. We first give an overview of known inequalities and then outline the proof technique.

Theorem 4.1 can be generalized to arbitrary logconcave measures. Its proof is very similar to that of 4.1 and we will presently give an outline.

Theorem 4.2 Let $f$ be a logconcave function whose support has diameter $D$ and let $\pi_{f}$ be the induced measure. Then for any partition of $\mathbb{R}^{n}$ into measurable sets $S_{1}, S_{2}, S_{3}$,

$$
\pi_{f}\left(S_{3}\right) \geq \frac{2 d\left(S_{1}, S_{2}\right)}{D} \min \left\{\pi_{f}\left(S_{1}\right), \pi_{f}\left(S_{2}\right)\right\}
$$

In terms of diameter, this inequality is the best possible, as shown by a cylinder. A more refined inequality is obtained in [22, 35] using the average distance of a point to the center of gravity (in place of diameter). It is possible for a convex body to have much larger diameter than average distance to its centroid (e.g., a cone). In such cases, the next theorem provides a better bound.

Theorem 4.3 Let $f$ be a logconcave density in $\mathbb{R}^{n}$ and $\pi_{f}$ be the corresponding measure. Let $z_{f}$ be the centroid of $f$ and define $M(f)=\mathrm{E}_{f}\left(\left|x-z_{f}\right|\right)$. Then, for any partition of $\mathbb{R}^{n}$ into measurable sets $S_{1}, S_{2}, S_{3}$,

$$
\pi_{f}\left(S_{3}\right) \geq \frac{\ln 2}{M(f)} d\left(S_{1}, S_{2}\right) \pi_{f}\left(S_{1}\right) \pi_{f}\left(S_{2}\right)
$$

For an isotropic density, $M(f)^{2} \leq \mathrm{E}_{f}\left(\left|x-z_{f}\right|^{2}\right)=n$ and so $M(f) \leq \sqrt{n}$. The diameter could be unbounded (e.g., an isotropic Gaussian).

A further refinement, based on isotropic position, has been conjectured in [22]. Let $\lambda$ be the largest eigenvalue of the inertia matrix of $f$, i.e.,

$$
\begin{equation*}
\lambda=\max _{v:|v|=1} \int_{\mathbb{R}^{n}} f(x)\left(v^{T} x\right)^{2} d x \tag{6}
\end{equation*}
$$

Then the conjecture says that there is an absolute constant $c$ such that

$$
\pi_{f}\left(S_{3}\right) \geq \frac{c}{\sqrt{\lambda}} d\left(S_{1}, S_{2}\right) \pi_{f}\left(S_{1}\right) \pi_{f}\left(S_{2}\right)
$$

Euclidean distance and isoperimetric inequalities based on it are relevant for the analysis of "local" walks such as the ball walk. Hit-and-run, with its nonlocal moves, is connected with a different notion of distance.

The cross-ratio distance between two points $u, v$ in a convex body $K$ is computed as follows: Let $p, q$ be the endpoints of the chord in $K$ through $u$ and $v$ such that the points occur in the order $p, u, v, q$. Then

$$
d_{K}(u, v)=\frac{|u-v||p-q|}{|p-u||v-q|}=(p: v: u: q) .
$$

where ( $p: v: u: q$ ) denotes the classical cross-ratio. We can now define the cross-ratio distance between two sets $S_{1}, S_{2}$ as

$$
d_{K}\left(S_{1}, S_{2}\right)=\min \left\{d_{K}(u, v): u \in S_{1}, v \in S_{2}\right\}
$$

The next theorem was proved in [30] for convex bodies and extended to logconcave densities in [36].

Theorem 4.4 Let $f$ be a logconcave density in $\mathbb{R}^{n}$ whose support is a convex body $K$ and let $\pi_{f}$ be the induced measure. Then for any partition of $\mathbb{R}^{n}$ into measurable sets $S_{1}, S_{2}, S_{3}$,

$$
\pi_{f}\left(S_{3}\right) \geq d_{K}\left(S_{1}, S_{2}\right) \pi_{f}\left(S_{1}\right) \pi_{f}\left(S_{2}\right)
$$

All the inequalities so far are based on defining the distance between $S_{1}$ and $S_{2}$ by the minimum over pairs of some notion of pairwise distance. It is reasonable to think that perhaps a much sharper inequality can be obtained by using some average distance between $S_{1}$ and $S_{2}$. Such an inequality was proved in [38]. As we will see in Section 6, it leads to a radical improvement in the analysis of random walks.

Theorem 4.5 Let $K$ be a convex body in $\mathbb{R}^{n}$. Let $f: K \rightarrow \mathbb{R}_{+}$be a logconcave density with corresponding measure $\pi_{f}$ and $h: K \rightarrow \mathbb{R}_{+}$, an arbitrary function. Let $S_{1}, S_{2}, S_{3}$ be any partition of $K$ into measurable sets. Suppose that for any pair of points $u \in S_{1}$ and $v \in S_{2}$ and any point $x$ on the chord of $K$ through $u$ and $v$,

$$
h(x) \leq \frac{1}{3} \min \left(1, d_{K}(u, v)\right)
$$

Then

$$
\pi_{f}\left(S_{3}\right) \geq \mathrm{E}_{f}(h(x)) \min \left\{\pi_{f}\left(S_{1}\right), \pi_{f}\left(S_{2}\right)\right\}
$$

The coefficient on the RHS has changed from a "minimum" to an "average". The weight $h(x)$ at a point $x$ is restricted only by the minimum cross-ratio distance between pairs $u, v$ from $S_{1}, S_{2}$ respectively, such that $x$ lies on the line between them (previously it was the overall minimum). In general, it can be much higher than the minimum cross-ratio distance between $S_{1}$ and $S_{2}$.

### 4.1 The localization lemma

The proofs of these inequalities are based on an elegant idea: integral inequalities in $\mathbb{R}^{n}$ can be reduced to one-dimensional inequalities! Checking the latter can
be tedious but is relatively easy. We illustrate the main idea by sketching the proof of Theorem 4.2.

For a proof of Theorem 4.2 by contradiction, let us assume the converse of its conclusion, i.e., for some partition $S_{1}, S_{2}, S_{3}$ of $\mathbb{R}^{n}$ and logconcave density $f$,

$$
\int_{S_{3}} f(x) d x<C \int_{S_{1}} f(x) d x \quad \text { and } \quad \int_{S_{3}} f(x) d x<C \int_{S_{2}} f(x) d x
$$

where $C=2 d\left(S_{1}, S_{2}\right) / D$. This can be reformulated as

$$
\begin{equation*}
\int_{\mathbb{R}^{n}} g(x) d x>0 \quad \text { and } \quad \int_{\mathbb{R}^{n}} h(x) d x>0 \tag{7}
\end{equation*}
$$

where

$$
g(x)=\left\{\begin{array}{ll}
C f(x) & \text { if } x \in S_{1}, \\
0 & \text { if } x \in S_{2}, \\
-f(x) & \text { if } x \in S_{3} .
\end{array} \quad \text { and } \quad h(x)= \begin{cases}0 & \text { if } x \in S_{1} \\
C f(x) & \text { if } x \in S_{2} \\
-f(x) & \text { if } x \in S_{3}\end{cases}\right.
$$

These inequalities are for functions in $\mathbb{R}^{n}$. The main tool to deal with them is the localization lemma [34] (see also [22] for extensions and applications).

Lemma 4.6 Let $g, h: \mathbb{R}^{n} \rightarrow \mathbb{R}$ be lower semi-continuous integrable functions such that

$$
\int_{\mathbb{R}^{n}} g(x) d x>0 \quad \text { and } \quad \int_{\mathbb{R}^{n}} h(x) d x>0
$$

Then there exist two points $a, b \in \mathbb{R}^{n}$ and a linear function $\ell:[0,1] \rightarrow \mathbb{R}_{+}$such that

$$
\int_{0}^{1} \ell(t)^{n-1} g((1-t) a+t b) d t>0 \quad \text { and } \quad \int_{0}^{1} \ell(t)^{n-1} h((1-t) a+t b) d t>0
$$

The points $a, b$ represent an interval $A$ and one may think of $l(t)^{n-1} d A$ as the cross-sectional area of an infinitesimal cone with base area $d A$. The lemma says that over this cone truncated at $a$ and $b$, the integrals of $g$ and $h$ are positive. Also, without loss of generality, we can assume that $a, b$ are in the union of the supports of $g$ and $h$.

The main idea behind the lemma is the following. Let $H$ be any halfspace such that

$$
\int_{H} g(x) d x=\frac{1}{2} \int_{\mathbb{R}^{n}} g(x) d x
$$

Let us call this a bisecting halfspace. Now either

$$
\int_{H} h(x) d x>0 \quad \text { or } \quad \int_{R^{n} \backslash H} h(x) d x>0 .
$$

Thus, either $H$ or its complementary halfspace will have positive integrals for both $g$ and $h$. Thus we have reduced the domains of the integrals from $\mathbb{R}^{n}$ to a
halfspace. If we could repeat this, we might hope to reduce the dimensionality of the domain. But do there even exist bisecting halfspaces? In fact, they are aplenty: for any ( $n-2$ )-dimensional affine subspace, there is a bisecting halfspace with $A$ contained in its bounding hyperplane. To see this, let $H$ be halfspace containing $A$ in its boundary. Rotating $H$ about $A$ we get a family of halfspaces with the same property. This family includes $H^{\prime}$, the complementary halfspace of $H$. Now the function $\int_{H} g-\int_{R^{n} \backslash H} g$ switches sign from $H$ to $H^{\prime}$. Since this is a continuous family, there must be a halfspace for which the function is zero, which is exactly what we want (this is sometimes called the "ham sandwich" theorem).

If we now take all ( $n-2$ )-dimensional affine subspaces given by some $x_{i}=r_{1}$, $x_{j}=r_{2}$ where $r_{1}, r_{2}$ are rational, then the intersection of all the corresponding bisecting halfspaces is a line (by choosing only rational values for $x_{i}$, we are considering a countable intersection). As long as we are left with a two or higher dimensional set, there is a point in its interior with at least two coordinates that are rational, say $x_{1}=r_{1}$ and $x_{2}=r_{2}$. But then there is a bisecting halfspace $H$ that contains the affine subspace given by $x_{1}=r_{1}, x_{2}=r_{2}$ in its boundary, and so it properly partitions the current set. With some additional work, this leads to the existence of a concave function on an interval (in place of the linear function $\ell$ in the theorem) with positive integrals. Simplifying further from concave to linear takes quite a bit of work.

Going back to the proof sketch of Theorem 4.2, we can apply the lemma to get an interval $[a, b]$ and a linear function $\ell$ such that

$$
\begin{equation*}
\int_{0}^{1} \ell(t)^{n-1} g((1-t) a+t b) d t>0 \quad \text { and } \quad \int_{0}^{1} \ell(t)^{n-1} h((1-t) a+t b) d t>0 \tag{8}
\end{equation*}
$$

(The functions $g, h$ as we have defined them are not lower semi-continuous. However, this can be easily achieved by expanding $S_{1}$ and $S_{2}$ slightly so as to make them open sets, and making the support of $f$ an open set. Since we are proving strict inequalities, we do not lose anything by these modifications).

Let us partition $[0,1]$ into $Z_{1}, Z_{2}, Z_{3}$.

$$
Z_{i}=\left\{t \in[0,1]:(1-t) a+t b \in S_{i}\right\}
$$

Note that for any pair of points $u \in Z_{1}, v \in Z_{2},|u-v| \geq d\left(S_{1}, S_{2}\right) / D$. We can rewrite (8) as

$$
\int_{Z_{3}} \ell(t)^{n-1} f((1-t) a+t b) d t<C \int_{Z_{1}} \ell(t)^{n-1} f((1-t) a+t b) d t
$$

and

$$
\int_{Z_{3}} \ell(t)^{n-1} f((1-t) a+t b) d t<C \int_{Z_{2}} \ell(t)^{n-1} f((1-t) a+t b) d t
$$

The functions $f$ and $\ell(.)^{n-1}$ are both logconcave, so $F(t)=\ell(t)^{n-1} f((1-t) a+t b)$ is also logconcave. We get,

$$
\begin{equation*}
\int_{Z_{3}} F(t) d t<C \min \left\{\int_{Z_{1}} F(t) d t, \int_{Z_{2}} F(t) d t\right\} \tag{9}
\end{equation*}
$$

Now consider what Theorem 4.2 asserts for the function $F(t)$ over the interval $[0,1]$ and the partition $Z_{1}, Z_{2}, Z_{3}$ :

$$
\begin{equation*}
\int_{Z_{3}} F(t) d t \geq 2 d\left(Z_{1}, Z_{2}\right) \min \left\{\int_{Z_{1}} F(t) d t, \int_{Z_{2}} F(t) d t\right\} \tag{10}
\end{equation*}
$$

We have substituted 1 for the diameter of the interval [0,1]. Also, $d\left(Z_{1}, Z_{2}\right) \geq$ $d\left(S_{1}, S_{2}\right) / D=C / 2$. Thus, Theorem 4.2 applied to the function $F(t)$ contradicts (9) and to prove the theorem in general, it suffices to prove it in the onedimensional case.

In fact, it will be enough to prove (10) for the case when each $Z_{i}$ is a single interval. Suppose we can do this. Then, for each maximal interval $(c, d)$ contained in $Z_{3}$, the integral of $F$ over $Z_{3}$ is at least $C$ times the smaller of the integrals to its left $[0, c]$ and to its right $[d, 1]$ and so one of these intervals is "accounted" for. If all of $Z_{1}$ or all of $Z_{2}$ is accounted for, then we are done. If not, there is an unaccounted subset $U$ that intersects both $Z_{1}$ and $Z_{2}$. But then, since $Z_{1}$ and $Z_{2}$ are separated by at least $d\left(S_{1}, S_{2}\right) / D$, there is an interval of $Z_{3}$ of length at least $d\left(S_{1}, S_{2}\right) / D$ between $U \cap Z_{1}$ and $U \cap Z_{2}$ which can account for more.

We are left with proving (10) when each $Z_{i}$ is an interval. Without the factor of two, this is trivial by the logconcavity of $F$. To get $C$ as claimed, one can reduce this to the case when $F(t)=e^{c t}$ and verify it for this function [34]. The main step is to show that there is a choice of $c$ so that when the current $F(t)$ is replaced by $e^{c t}$, the LHS of (10) does not increase and the RHS does not decrease.

## 5 Mixing of the ball walk

With the isoperimetric inequalities at hand, we are now ready to prove bounds on the conductance and hence on the mixing time. In this section, we focus on the ball walk in a convex body $K$. Assume that $K$ contains the unit ball.

A geometric random walk is said to be rapidly mixing if its conductance is bounded from below by an inverse polynomial in the dimension. By Corollary 3.5 , this implies that the number of steps to halve the variation distance to stationary is a polynomial in the dimension. The conductance of the ball walk in a convex body $K$ can be exponentially small. Consider, for example, starting at point $x$ near the apex of a rotational cone in $\mathbb{R}^{n}$. Most points in a ball of radius $\delta$ around $x$ will lie outside the cone (if $x$ is sufficiently close to the apex) and so the local conductance is arbitrarily small. So, strictly speaking, the ball walk is not rapidly mixing.

There are two ways to get around this. For the purpose of sampling uniformly from $K$, one can expand $K$ a little bit by considering $K^{\prime}=K+\alpha B_{n}$, i.e., adding a ball of radius $\alpha$ around every point in $K$. Then for $\alpha>2 \delta \sqrt{n}$, it is not hard to see that $\ell(u)$ is at least $1 / 8$ for every point $u \in K^{\prime}$. We can now consider the ball walk in $K^{\prime}$. This fix comes at a price. First, we need a membership oracle for $K^{\prime}$. This can be constructed as follows: given a point $x \in \mathbb{R}^{n}$, we find a point
$y \in K$ such that $|x-y|$ is minimum. This is a convex program and can be solved using the Ellipsoid algorithm [15] and the membership oracle for $K$, Second, we need to ensure that $\operatorname{vol}\left(K^{\prime}\right)$ is comparable to $\operatorname{vol}(K)$. Since $K$ contains a unit ball, $K^{\prime} \subseteq(1+\alpha) K$ and so with $\alpha<1 / n$, we get that $\operatorname{vol}\left(K^{\prime}\right)<e \operatorname{vol}(K)$. Thus, we would need $\delta<1 / 2 n \sqrt{n}$.

Does large local conductance imply that the conductance is also large? We will prove that the answer is yes. The next lemma about one-step distributions of nearby points will be useful.

Lemma 5.1 Let $u, v$ be such that $|u-v| \leq \frac{t \delta}{\sqrt{n}}$ and $\ell(u), \ell(v) \geq \ell$. Then,

$$
\left\|P_{u}-P_{v}\right\|_{t v} \leq 1+t-\ell
$$

Roughly speaking, the lemma says that if two points with high local conductance are close in Euclidean distance, then their one-step distributions have a large overlap. Its proof follows from a computation of the overlap volume of the balls of radius $\delta$ around $u$ and $v$.

We can now state and prove a bound on the conductance of the ball walk.
Theorem 5.2 Let $K$ be a convex body of diameter $D$ so that for every point $u$ in $K$, the local conductance of the ball walk with $\delta$ steps is at least $\ell$. Then,

$$
\phi \geq \frac{\ell^{2} \delta}{16 \sqrt{n} D}
$$

The structure of most proofs of conductance is similar and we will illustrate it by proving this theorem.
Proof. Let $K=S_{1} \cup S_{2}$ be a partition into measurable sets. We will prove that

$$
\begin{equation*}
\int_{S_{1}} P_{x}\left(S_{2}\right) d x \geq \frac{\ell^{2} \delta}{16 \sqrt{n} D} \min \left\{\operatorname{vol}\left(S_{1}\right), \operatorname{vol}\left(S_{2}\right)\right\} \tag{11}
\end{equation*}
$$

Note that since the uniform distribution is stationary,

$$
\int_{S_{1}} P_{x}\left(S_{2}\right) d x=\int_{S_{2}} P_{x}\left(S_{1}\right) d x
$$

Consider the points that are "deep" inside these sets, i.e. unlikely to jump out of the set (see Figure 2):

$$
S_{1}^{\prime}=\left\{x \in S_{1}: P_{x}\left(S_{2}\right)<\frac{\ell}{4}\right\}
$$

and

$$
S_{2}^{\prime}=\left\{x \in S_{2}: P_{x}\left(S_{1}\right)<\frac{\ell}{4}\right\}
$$

Let $S_{3}^{\prime}$ be the rest i.e., $S_{3}^{\prime}=K \backslash S_{1}^{\prime} \backslash S_{2}^{\prime}$.


Figure 2: The conductance proof. The dark line is the boundary between $S_{1}$ and $S_{2}$.

$$
\begin{aligned}
& \text { Suppose } \operatorname{vol}\left(S_{1}^{\prime}\right)<\operatorname{vol}\left(S_{1}\right) / 2 \text {. Then } \\
& \qquad \int_{S_{1}} P_{x}\left(S_{2}\right) d x \geq \frac{\ell}{4} \operatorname{vol}\left(S_{1} \backslash S_{1}^{\prime}\right) \geq \frac{\ell}{8} \operatorname{vol}\left(S_{1}\right)
\end{aligned}
$$

which proves (11).
So we can assume that $\operatorname{vol}\left(S_{1}^{\prime}\right) \geq \operatorname{vol}\left(S_{1}\right) / 2$ and similarly $\operatorname{vol}\left(S_{2}^{\prime}\right) \geq \operatorname{vol}\left(S_{2}\right) / 2$. Now, for any $u \in S_{1}^{\prime}$ and $v \in S_{2}^{\prime}$,

$$
\left\|P_{u}-P_{v}\right\|_{t v} \geq 1-P_{u}\left(S_{2}\right)-P_{v}\left(S_{1}\right)>1-\frac{\ell}{2}
$$

Applying Lemma 5.1 with $t=\ell / 2$, we get that

$$
|u-v| \geq \frac{\ell \delta}{2 \sqrt{n}}
$$

Thus $d\left(S_{1}, S_{2}\right) \geq \ell \delta / 2 \sqrt{n}$. Applying Theorem 4.1 to the partition $S_{1}^{\prime}, S_{2}^{\prime}, S_{3}^{\prime}$, we have

$$
\begin{aligned}
\operatorname{vol}\left(S_{3}^{\prime}\right) & \geq \frac{\ell \delta}{\sqrt{n} D} \min \left\{\operatorname{vol}\left(S_{1}^{\prime}\right), \operatorname{vol}\left(S_{2}^{\prime}\right)\right\} \\
& \geq \frac{\ell \delta}{2 \sqrt{n} D} \min \left\{\operatorname{vol}\left(S_{1}\right), \operatorname{vol}\left(S_{2}\right)\right\}
\end{aligned}
$$

We can now prove (11) as follows:

$$
\begin{aligned}
\int_{S_{1}} P_{x}\left(S_{2}\right) d x & =\frac{1}{2} \int_{S_{1}} P_{x}\left(S_{2}\right) d x+\frac{1}{2} \int_{S_{2}} P_{x}\left(S_{1}\right) d x \\
& \geq \frac{1}{2} \operatorname{vol}\left(S_{3}^{\prime}\right) \frac{\ell}{4} \\
& \geq \frac{\ell^{2} \delta}{16 \sqrt{n} D} \min \left\{\operatorname{vol}\left(S_{1}\right), \operatorname{vol}\left(S_{2}\right)\right\}
\end{aligned}
$$

As observed earlier, by going to $K^{\prime}=K+(1 / n) B_{n}$ and using $\delta=1 / 2 n \sqrt{n}$, we have $\ell \geq 1 / 8$. Thus, for the ball walk in $K^{\prime}, \phi=\Omega\left(1 / n^{2} D\right)$. Using Corollary 3.5 , the mixing rate is $O\left(n^{4} D^{2}\right)$.

We mentioned earlier that there are two ways to get around the fact that the ball walk can have very small local conductance. The second, which we describe next, is perhaps a bit cleaner and also achieves a better bound on the mixing rate. It is based on the observation that only a small fraction of points can have small local conductance. Define the points of high local conductance as

$$
K_{\delta}=\left\{u \in K: \ell(u) \geq \frac{3}{4}\right\}
$$

Lemma 5.3 Suppose that $K$ is a convex body containing a unit ball in $\mathbb{R}^{n}$. Then,
a. $K_{\delta}$ is a convex set.
b. $\operatorname{vol}\left(K_{\delta}\right) \geq(1-2 \delta \sqrt{n}) \operatorname{vol}(K)$.

The first part follows from the Brunn-Minkowski inequality (1). The second is proved by estimating the average local conductance [23] and has the following implication: if we set $\delta \leq \varepsilon / 2 \sqrt{n}$, we get that at least (1- $\varepsilon$ ) fraction of points in $K$ have large local conductance. Using this, we can prove the following theorem.

Theorem 5.4 For any $0 \leq s \leq 1$, we can choose the step-size $\delta$ for the ball walk in a convex body $K$ of diameter $D$ so that

$$
\phi_{s} \geq \frac{s}{200 n D}
$$

Proof. The proof is quite similar to that of Theorem 5.2. Let $S_{1}, S_{2}$ be a partition of $K$. First, since we are proving a bound on the $s$-conductance, we can assume that $\operatorname{vol}\left(S_{1}\right), \operatorname{vol}\left(S_{2}\right) \geq s \operatorname{vol}(K)$. Next, we choose $\delta=s / 4 \sqrt{n}$ so that by Lemma 5.3,

$$
\operatorname{vol}\left(K_{\delta}\right) \geq\left(1-\frac{s}{2}\right) \operatorname{vol}(K)
$$

So only an $s / 2$ fraction of $K$ has small local conductance and we will be able to ignore it. Define

$$
S_{1}^{\prime}=\left\{x \in S_{1} \cap K_{\delta}: P_{x}\left(S_{2}\right)<\frac{3}{16}\right\}
$$

and

$$
S_{2}^{\prime}=\left\{x \in S_{2} \cap K_{\delta}: P_{x}\left(S_{1}\right)<\frac{3}{16}\right\}
$$

As in the proof of Theorem 5.2, these points are "deep" in $S_{1}$ and $S_{2}$ respectively and they are also restricted to be in $K_{\delta}$. Recall that the local conductance of
every point in $K_{\delta}$ is at least $3 / 4$. We can assume that $\operatorname{vol}\left(S_{1}^{\prime}\right) \geq \operatorname{vol}\left(S_{1}\right) / 3$. Otherwise,

$$
\begin{aligned}
\int_{S_{1}} P_{x}\left(S_{2}\right) d x & \geq\left(\frac{2}{3} \operatorname{vol}\left(S_{1}\right)-\frac{s}{2} \operatorname{vol}(K)\right) \frac{3}{16} \\
& \geq \frac{\operatorname{vol}\left(S_{1}\right)}{32}
\end{aligned}
$$

which implies the theorem. Similarly, we can assume that $\operatorname{vol}\left(S_{2}^{\prime}\right) \geq \operatorname{vol}\left(S_{2}\right) / 3$.
For any $u \in S_{1}^{\prime}$ and $v \in S_{2}^{\prime}$,

$$
\left\|P_{u}-P_{v}\right\|_{t v} \geq 1-P_{u}\left(S_{2}\right)-P_{v}\left(S_{1}\right)>1-\frac{3}{8}
$$

Applying Lemma 5.1 with $t=3 / 8$, we get that

$$
|u-v| \geq \frac{3 \delta}{8 \sqrt{n}}
$$

Thus $d\left(S_{1}, S_{2}\right) \geq 3 \delta / 8 \sqrt{n}$. Let $S_{3}^{\prime}=K_{\delta} \backslash\left(S_{1}^{\prime} \cup S_{2}^{\prime}\right)$ Applying Theorem 4.1 to the partition $S_{1}^{\prime}, S_{2}^{\prime}, S_{3}^{\prime}$ of $K_{\delta}$, we have

$$
\begin{aligned}
\operatorname{vol}\left(S_{3}^{\prime}\right) & \geq \frac{3 \delta}{4 \sqrt{n} D} \min \left\{\operatorname{vol}\left(S_{1}^{\prime}\right), \operatorname{vol}\left(S_{2}^{\prime}\right)\right\} \\
& \geq \frac{s}{16 n D} \min \left\{\operatorname{vol}\left(S_{1}\right), \operatorname{vol}\left(S_{2}\right)\right\}
\end{aligned}
$$

The theorem follows:

$$
\begin{aligned}
\int_{S_{1}} P_{x}\left(S_{2}\right) d x & \geq \frac{1}{2} \operatorname{vol}\left(S_{3}^{\prime}\right) \frac{3}{16} \\
& >\frac{s}{200 n D} \min \left\{\operatorname{vol}\left(S_{1}\right), \operatorname{vol}\left(S_{2}\right)\right\}
\end{aligned}
$$

Using Corollary 3.5(b), this implies that from an M-warm start, the variation distance of $Q_{t}$ and $Q$ is smaller than $\varepsilon$ after

$$
\begin{equation*}
t \geq C \frac{M^{2}}{\varepsilon^{2}} n^{2} D^{2} \ln \left(\frac{2 M}{\varepsilon}\right) \tag{12}
\end{equation*}
$$

steps, for some absolute constant $C$.
There is another way to use Lemma 5.3. In [23], the following modification of the ball walk, called the speedy walk, is described. At a point $x$, the speedy walk picks a point uniformly from $K \cap x+\delta B_{n}$. Thus, the local conductance of every point is 1 . However, there are two complications with this. First, the stationary distribution is not uniform, but proportional to $\ell(u)$. Second, each step seems unreasonable - we could make $\delta>D$ and then we would only need one step to get a random point in $K$. We can take care of the first problem
with a rejection step at the end (and using Lemma 5.3). The root of the second problem is the question: how do we implement one step of the speedy walk? The only general way is to get random points from the ball around the current point till one of them is also in $K$. This process is the ball walk and it requires $1 / \ell(u)$ attempts in expectation at a point $u$. However, if we count only the proper steps, i.e., ones that move away from the current point, then it is possible to show that the mixing rate of the walk is in fact $O\left(n^{2} D^{2}\right)$ from any starting point [20]. Again, the proof is based on an isoperimetric inequality which is slightly sharper than Theorem 4.2. For this bound to be useful, we also need to bound the total number of improper or "wasted" steps. If we start at a random point, then this is the number of proper steps times $\mathrm{E}(1 / \ell(u))$, which can be unbounded. But, if we allow a small overall probability of failure, then with the remaining probability, the expected number of wasted steps is bounded by $O\left(n^{2} D^{2}\right)$ as well.

The bound of $O\left(n^{2} D^{2}\right)$ on the mixing rate is the best possible in terms of the diameter, as shown by a cylinder. However, if the convex body is isotropic, then the isoperimetry conjecture (6) implies a mixing rate of $O\left(n^{2}\right)$.

For the rest of this section, we will discuss how these methods can be extended to sampling more general distributions. We saw already that the ball walk can be used along with a Metropolis filter to sample arbitrary density functions. When is this method efficient? In [3] and [14] respectively, it is shown that the ball walk and the lattice walk are rapidly mixing from a warm start, provided that the density is logconcave and it does not vary much locally, i.e., its Lipschitz constant is bounded. In [35], the assumptions on smoothness are eliminated, and it is shown that the ball walk is rapidly mixing from a warm start for any logconcave function in $\mathbb{R}^{n}$. Moreover, the mixing rate is $O\left(n^{2} D^{2}\right)$ (ignoring the dependence on the start), which matches the case of the uniform density on a convex body. Various properties of logconcave functions are developed in [35] with an eye to the proof. In particular, a smoother version of any given logconcave density is defined and used to prove an analogue of Lemma 5.3. For a logconcave density $f$ in $\mathbb{R}^{n}$, the smoother version is defined as

$$
\hat{f}(x)=\min _{C} \frac{1}{\operatorname{vol}(C)} \int_{C} f(x+u) d u
$$

where $C$ ranges over all convex subsets of the ball $x+r B_{n}$ with $\operatorname{vol}(C)=$ $\operatorname{vol}\left(B_{n}\right) / 16$. This function is logconcave and bounded from above by $f$ everywhere (using Theorem 2.2). Moreover, for $\delta$ small enough, its integral is close to the integral of $f$. We get a lemma very similar to Lemma 5.3. The function $\hat{f}$ can be thought of as a generalization of $K_{\delta}$.

Lemma 5.5 Let $f$ be any logconcave density in $\mathbb{R}^{n}$. Then

1. The function $\hat{f}$ is logconcave.
2. If $f$ is isotropic, then $\int_{\mathbb{R}^{n}} \hat{f}(x) d x \geq 1-64 \delta^{1 / 2} n^{1 / 4}$.

Using this along with some technical tools, it can be shown that $\phi_{s}$ is large. Perhaps the main contribution of [35] is to move the smoothness conditions from requirements on the input (i.e., the algorithm) to tools for the proof.

In summary, analyzing the ball walk has led to many interesting developments: isoperimetric inequalities, more general methods of proving convergence $\left(\phi_{s}\right)$ and many tricks for sampling to get around the fact that it is not rapidly mixing from general starting points (or distributions). The analysis of the speedy walk shows that most points are good starting points. However, it is an open question as to whether the ball walk is rapidly mixing from a predetermined starting point, e.g., the centroid.

## 6 Mixing of hit-and-run

Hit-and-run, introduced by Smith [45], offers the attractive possibility of long steps. There is some evidence that it is fast in practice [5, 48].

### 6.1 Warm start

Lovász [30] showed that hit-and-run mixes rapidly from a warm start in a convex body $K$. If we start with an $M$-warm distribution, then in

$$
O\left(\frac{M^{2}}{\varepsilon^{2}} n^{2} D^{2} \ln \left(\frac{M}{\varepsilon}\right)\right)
$$

steps, the distance between the current distribution and the stationary is at $\operatorname{most} \varepsilon$. This is essentially the same bound as for the ball walk, and so hit-andrun is no worse. The proof is based on cross-ratio isoperimetry (Theorem 4.5) for convex bodies and a new lemma about the overlap of one-step distributions. For $x \in K$, let $y$ be a random step from $x$. Then the step-size $F(x)$ at $x$ is defined as

$$
\mathrm{P}(|x-y| \leq F(x))=\frac{1}{8}
$$

The lemma below asserts that if $u, v$ are close in Euclidean distance and crossratio distance then their one-step distributions overlap substantially. This is analogous to Lemma 5.1 for the ball walk.

Lemma 6.1 Let $u, v \in K$. Suppose that

$$
d_{K}(u, v)<\frac{1}{8} \text { and }|u-v|<\frac{2}{\sqrt{n}} \max \{F(u), F(v)\} .
$$

Then

$$
\left\|P_{u}-P_{v}\right\|_{t v}<1-\frac{1}{500}
$$

Hit-and-run generalizes naturally to sampling arbitrary functions. The isoperimetry, the one-step lemma and the bound on $\phi_{s}$ were all extended to arbitrary logconcave densities in [36]. Thus, hit-and-run is rapidly mixing for any logconcave
density from a warm start. While the analysis is along the lines of that in [30] and uses the tools developed in [35], it has to overcome substantial additional difficulties.

So hit-and-run is at least as fast as the ball walk. But is it faster? Can it get stuck in corners (points of small local conductance) like the ball walk?

### 6.2 Any start

Let us revisit the bad example for the ball walk: starting near the apex of a rotational cone. If we start hit-and-run at any interior point, then it exhibits a small, but inverse polynomial, drift towards the base of the cone. Thus, although the initial steps are tiny, they rapidly get larger and the current point moves away from the apex. This example shows two things. First, the "step-size" of hit-and-run can be arbitrarily small (near the apex), but hit-and-run manages to escape from such regions. This phenomenon is in fact general as shown by the following theorem, proved recently in [38].

Theorem 6.2 The conductance of hit-and-run in a convex body of diameter $D$ is $\Omega(1 / n D)$.

Unlike the ball walk, we can bound the conductance of hit-and-run (for arbitrarily small subsets). From this we get a bound on mixing time.

Theorem 6.3 Let $K$ be a convex body that contains a unit ball and has centroid $z_{K}$. Suppose that $\mathrm{E}_{K}\left(\left|x-z_{K}\right|^{2}\right) \leq R^{2}$ and $\left\|Q_{0} / Q\right\| \leq M$. Then after

$$
t \geq C n^{2} R^{2} \ln ^{3} \frac{M}{\varepsilon}
$$

steps, where $C$ is an absolute constant, we have $\left\|Q_{t}-Q\right\| \leq \varepsilon$.
The theorem improves on the bound for the ball walk (12) by reducing the dependence on $M$ and $\varepsilon$ from polynomial (which is unavoidable for the ball walk) to logarithmic, while maintaining the (optimal) dependence on $R$ and $n$. For a body in near-isotropic position, $R=O(\sqrt{n})$ and so the mixing time is $O^{*}\left(n^{3}\right)$. One also gets a polynomial bound starting from any single interior point. If $x$ is at distance $d$ from the boundary, then the distribution obtained after one step from $x$ has $\left\|Q_{1} / Q\right\| \leq(n / d)^{n}$ and so applying the above theorem, the mixing time is $O\left(n^{4} \ln ^{3}(n / d \varepsilon)\right)$.

The main tool in the proof is a new isoperimetric inequality based on "average" distance (Theorem 4.5). The proof of conductance is on the same lines as shown for Theorem 5.2 in the previous section. It uses Lemma 6.1 for comparing one-step distributions.

Theorems 6.2 and 6.3 have been extended in [38] to the case of sampling an exponential density function, i.e., $f(x)$ is proportional to $e^{a^{T} x}$ for some fixed vector $a$. It remains open to determine if hit-and-run has high conductance for arbitrary logconcave functions.

As in the ball walk analysis, it is not known (even in the convex body case) if starting at the centroid is as good as a warm start. Also, while the theorem is the best possible in terms of $R$, it is conceivable that for an isotropic body the mixing rate is $O\left(n^{2}\right)$.

## 7 Efficient sampling

Let $f$ be a density in $\mathbb{R}^{n}$ with corresponding measure $\pi_{f}$. Sampling $f$, i.e., generating independent random points distributed according to $\pi_{f}$ is a basic algorithmic problem with many applications. We have seen in previous sections that if $f$ is logconcave there are natural random walks in $\mathbb{R}^{n}$ that will converge to $\pi_{f}$. Does this yield an efficient sampling algorithm?

### 7.1 Rounding

Take the case when $f$ is uniform over a convex body $K$. The convergence depends on the diameter $D$ of $K$ (or the second moment). So the resulting algorithm to get a random point would take $\operatorname{poly}(n, D)$ steps. However, the input to the algorithm is only $D$ and an oracle. So we would like an algorithm whose dependence on $D$ is only logarithmic. How can this be done? The Ellipsoid algorithm can be used to find a transformation that achieves $D=$ $O\left(n^{1.5}\right)$ in poly $(n, \log D)$ steps.

Isotropic position provides a better solution. For a convex body in isotropic position $D \leq n$. For an isotropic logconcave distribution, $(1-\varepsilon)$ of its measure lies in a ball of radius $\sqrt{n} \ln (1 / \varepsilon)$. But how to make $f$ isotropic? One way is by sampling. We get $m$ random points from $f$ and compute an affine tranformation that makes this set of points isotropic. We then apply this transformation to $f$. It is shown in [43], that the resulting distribution is near-isotropic with $m=$ $O\left(n \log ^{2} n\right)$ points for convex bodies and and $m=O\left(n \log ^{3} n\right)$ for logconcave densities [35] with large probability.

Although this sounds cyclic (we need samples to make the sampling efficient) one can bootstrap and make larger and larger subsets of $f$ isotropic. For a convex body $K$ such an algorithm was given in [23]. Its complexity is $O^{*}\left(n^{5}\right)$. This has been improved to $O^{*}\left(n^{4}\right)$ in [37]. The basic approach in [23] is to define a series of bodies, $K_{i}=K \cap 2^{i / n} B_{n}$. Then $K_{0}=B_{n}$ is isotropic upto a radial scaling. Given that $K_{i}$ is 2 -isotropic, $K_{i+1}$ will be 6 -isotropic and so we can sample efficiently from it. We use these samples to compute a transformation that makes $K_{i+1} 2$-isotropic and continue. The number of samples required in each phase is $O^{*}(n)$ and the total number of phases is $O(n \log D)$. Since each sample is drawn from a near-isotropic convex body, the sample complexity is $O^{*}\left(n^{3}\right)$ on average $\left(O^{*}\left(n^{4}\right)\right.$ for the first point and $O^{*}\left(n^{3}\right)$ for subsequent points since we have a warm start). This gives an overall complexity of $O^{*}\left(n^{5}\right)$. The improvement to $O^{*}\left(n^{4}\right)$ uses ideas from the latest volume algorithm [37], including sampling from an exponential density and the pencil construction (see Section 9).

A similar method also works for making a logconcave density $f$ isotropic [35]. We consider a series of level sets

$$
L_{i}=\left\{x \in \mathbb{R}^{n}: f(x) \geq M_{f} / 2^{\left(1+\frac{1}{n}\right)^{i}}\right\}
$$

where $M_{f}$ is the maximum value of $f$. In phase $i$, we make the restriction of $f$ to $L_{i}$ isotropic. The complexity of this algorithm is $O^{*}\left(n^{5}\right)$. It is an open question to reduce this to $O^{*}\left(n^{4}\right)$.

### 7.2 Independence

The second important issue to be addressed is that of independence. If we examine the current point every $m$ steps for some $m$, then are these points independent? Unfortunately, they might not be independent even if $m$ is as large as the mixing time. Another problem is that the distribution might not be exactly $\pi_{f}$. The latter problem is easier to deal with. Suppose that $x$ is from some distribution $\pi$ so that $\left\|\pi-\pi_{f}\right\|_{t v} \leq \varepsilon$. Typically this affects the algorithm using the samples by some small function of $\varepsilon$. There is a general way to handle this (sometimes called divine intervention). We can pretend that $x$ is drawn from $\pi_{f}$ with probability $1-\varepsilon$ and from some other distribution with probability at most $\varepsilon$. If we draw $k$ samples, then the probability of success (i.e., each sample is drawn from the desired distribution) is at least $1-k \varepsilon$.

Although points spaced apart by $m$ steps might not be independent, they are "nearly" independent in the following sense. Two random variables $X, Y$ will be called $\mu$-independent $(0<\mu<1)$ if for any two sets $A, B$ in their ranges,

$$
|\mathrm{P}(X \in A, Y \in B)-\mathrm{P}(X \in A) \mathrm{P}(Y \in B)| \leq \mu
$$

The next lemma summarizes some properties of $\mu$-independence.
Lemma 7.1 1. Let $X$ and $Y$ be $\mu$-independent, and $f, g$ be two measurable functions. Then $f(X)$ and $g(Y)$ are also $\mu$-independent.
2. Let $X, Y$ be $\mu$-independent random variables such that $0 \leq X \leq a$ and $0 \leq Y \leq b$. Then

$$
|\mathrm{E}(X Y)-\mathrm{E}(X) \mathrm{E}(Y)| \leq \mu a b
$$

3. Let $X_{0}, X_{1}, \ldots$, be a Markov chain, and assume that for some $i>0, X_{i+1}$ is $\mu$-independent from $X_{i}$. Then $X_{i+1}$ is $\mu$-independent from $\left(X_{0}, \ldots, X_{i}\right)$.

The guarantee that $\pi$ is close to $\pi_{f}$ will imply the following.
Lemma 7.2 Let $Q$ be the stationary distribution of a Markov chain and $t$ be large enough so that for any starting distribution $Q_{0}$ with $\left\|Q_{0} / Q\right\| \leq 4 M$ we have $\left\|Q_{t}-Q\right\|_{t v} \leq \varepsilon$. Let $X$ be a random point from a starting distribution $Q_{0}$ such that $\left\|Q_{0} / Q\right\| \leq M$. Then the point $Y$ obtained by taking t steps of the chain starting at $X$ is $2 \varepsilon$-independent from $X$.

Proof. Let $A, B \subseteq \mathbb{R}^{n}$; we claim that

$$
\begin{aligned}
& |\mathrm{P}(X \in A, Y \in B)-\mathrm{P}(X \in A) \mathrm{P}(Y \in B)| \\
& =\mathrm{P}(X \in A)|\mathrm{P}(Y \in B \mid X \in A)-\mathrm{P}(Y \in B)| \\
& \leq 2 \varepsilon
\end{aligned}
$$

Since

$$
\begin{aligned}
& |\mathrm{P}(X \in A, Y \in B)-\mathrm{P}(X \in A) \mathrm{P}(Y \in B)| \\
& =|\mathrm{P}(X \in \bar{A}, Y \in B)-\mathrm{P}(X \in \bar{A}) \mathrm{P}(Y \in B)|
\end{aligned}
$$

we may assume that $Q_{0}(A) \geq 1 / 2$. Let $Q_{0}^{\prime}$ be the restriction of $Q_{0}$ to $A$, scaled to a probability measure. Then $Q_{0}^{\prime} \leq 2 Q_{0}$ and so $\left\|Q_{0}^{\prime} / Q\right\| \leq 4\left\|Q_{0} / Q\right\| \leq 4 M$. Imagine running the Markov chain with starting distribution $Q_{0}^{\prime}$. Then, by the assumption on $t$,

$$
\begin{aligned}
|\mathrm{P}(Y \in B \mid X \in A)-\mathrm{P}(Y \in B)| & =\left\|Q_{t}^{\prime}(B)-Q_{t}(B)\right\| \\
& \leq\left\|Q_{t}^{\prime}(B)-Q(B)\right\|+\left\|Q_{t}(B)-Q(B)\right\| \\
& \leq 2 \varepsilon
\end{aligned}
$$

and so the claim holds.
Various versions of this lemma, adapted to the mixing guarantee at hand, have been used in the literature. See $[34,23,37]$ for developments along this line.

## 8 Application I: Convex optimization

Let $S \subset \mathbb{R}^{n}$, and $f: S \rightarrow \mathbb{R}$ be a real-valued function. Optimization is the following basic problem: $\min f(x)$ s.t. $x \in S$, that is, find a point $x \in S$ which minimizes by $f$. We denote by $x^{*}$ a solution for the problem. When the set $S$ and the function $f$ are convex ${ }^{2}$, we obtain a class of problems which are solvable in poly $(n, \log (1 / \varepsilon))$ time where $\epsilon$ defines an optimality criterion. If $x$ is the point found, then $\left|x-x^{*}\right| \leq \epsilon$.

The problem of minimizing a convex function over a convex set in $\mathbb{R}^{n}$ is a common generalization of well-known geometric optimization problems such as linear programming as well as a variety of combinatorial optimization problems including matchings, flows and matroid intersection, all of which have polynomial-time algorithms [15]. It is shown in [15] that the Ellipsoid method [47, 25] solves this problem in polynomial time when $K$ is given by a separation oracle. A different, more efficient algorithm is given in [46]. Here, we discuss the recent algorithm of [6] which is based on random sampling.

Note that minimizing a quasi-convex function is easily reduced to the feasibility problem: to minimize a quasi-convex function $f(x)$, we simply add the constraint $f(x) \leq t$ and search (in a binary fashion) for the optimal $t$.

[^2]In the description below, we assume that the convex set $K$ is contained in the axis-aligned cube of width $R$ centered at the origin; further if $K$ is nonempty then it contains a cube of width $r$. It is easy to show that any algorithm with this input specification needs to make at least $n \log (R / r)$ oracle queries. The parameter $L$ is equal to $\log \frac{R}{r}$.

```
Algorithm.
Input: A separation oracle for a convex set \(K\) and \(L\).
Output: A point in \(K\) or a guarantee that \(K\) is empty.
    1. Let \(P\) be the axis-aligned cube of side length \(R\) and
    center \(z=0\).
    2. Check if \(z\) is in \(K\). If so, report \(z\) and stop. If not,
    set
        \(H=\left\{x \mid a^{T} x \leq a^{T} z\right\}\).
    where \(a^{T} x \leq b\) is the halfspace containing \(K\) reported by
    the oracle.
3. Set \(P=P \cap H\). Pick \(m\) random points \(y^{1}, y^{2}, \ldots, y^{m}\) from \(P\).
    Set \(z\) to be their average.
    4. Repeat steps 2 and 3 at most \(2 n L\) times. Report \(K\) is
    empty.
```

The number of samples required in each iteration, $m$, is $O(n)$. Roughly speaking, the algorithm is computing an approximate centroid in each iteration. The idea of an algorithm based on computing the exact centroid was suggested in 1965 by Y. Levin [28]. Indeed, if we could compute the centroid in each iteration, then by Lemma 2.3, the volume of $P$ falls by a constant factor $(1-1 / e)$ in each iteration. But, finding the centroid, is \#P-hard, i.e., computationally intractable.

The idea behind the algorithm is that an approximate centroid can be computed using $O(n)$ random points and the volume of $P$ is likely to drop by a constant factor in each iteration with this choice of $z$. This is formalized in the next lemma. Although we need it only for convex bodies, it holds for arbitrary logconcave densities.

Lemma 8.1 Let $g$ be a logconcave density in $\mathbb{R}^{n}$ and $z$ be the average of $m$ random points from $\pi_{g}$. If $H$ is a halfspace containing $z$,

$$
\mathrm{E}\left(\pi_{g}(H)\right) \geq\left(\frac{1}{e}-\sqrt{\frac{n}{m}}\right)
$$

Proof. First observe that we can assume $g$ is in isotropic position. This is because a linear transformation $A$ affects the volume of a set $S$ as $\operatorname{vol}(A S)=$


Figure 3: An illustration of the algorithm.
$|\operatorname{det}(A)| \operatorname{vol}(S)$ and so the ratio of the volumes of two subsets is unchanged by the transformation. Applying this to all the level sets of $g$, we get that the ratio of the measures of two subsets is unchanged.

Since $z=\frac{1}{m} \sum_{i=1}^{m} y^{i}$,

$$
\begin{aligned}
\mathrm{E}\left(|z|^{2}\right)=\frac{1}{m^{2}} \sum_{i=1}^{m} \mathrm{E}\left(\left|y^{i}\right|^{2}\right) & =\frac{1}{m} \mathrm{E}\left(\left|y^{i}\right|^{2}\right) \\
& =\frac{1}{m} \sum_{j=1}^{n} \mathrm{E}\left(\left(y_{j}^{i}\right)^{2}\right)=\frac{n}{m}
\end{aligned}
$$

where the first equality follows from the independence between $y^{i}$, s , and equalities of the second line follow from the isotropic position. Let $h$ be a unit vector normal to $H$. We can assume that $h=e_{1}=(1,0, \ldots, 0)$.

Next, let $f$ be the marginal of $g$ along $h$, i.e.,

$$
\begin{equation*}
f(y)=\int_{x: x_{1}=y} g(x) d x . \tag{13}
\end{equation*}
$$

It is easy to see that $f$ is isotropic. The next lemma (from [35]; see [6] for the case of $f$ arising from convex bodies) states that its maximum must be bounded.

Lemma 8.2 Let $f: \mathbb{R} \rightarrow R_{+}$be an isotropic logconcave density function. Then,

$$
\max _{y} f(y)<1
$$

Using Lemma 2.3,

$$
\begin{aligned}
\int_{z_{1}}^{\infty} f(y) d y & =\int_{0}^{\infty} f(y) d y-\int_{0}^{z_{1}} f(y) d y \\
& \geq \frac{1}{e}-\left|z_{1}\right| \max _{y} f(y) \\
& \geq \frac{1}{e}-|z|
\end{aligned}
$$

The lemma follows from the bound on $\mathrm{E}(|z|)$.
The guarantee on the algorithm follows immediately. This optimal guarantee is also obtained in [46]; the Ellipsoid algorithm needs $O\left(n^{2} L\right)$ oracle calls.

Theorem 8.3 With high probability, the algorithm works correctly using at most $2 n L$ oracle calls (and iterations).

The algorithm can also be modified for optimization given a membership oracle only and a point in $K$. It has a similar flavor: get random points from $K$; restrict $K$ using the function value at the average of the random points; repeat. The oracle complexity turns out to be $O\left(n^{5} L\right)$ which is an improvement on previous methods. This has been improved for linear objective functions using a variant of simulated annealing [19].

## 9 Application II: Volume computation

Finally, we come to perhaps the most important application and the principal motivation behind many developments in the theory of random walks: the problem of computing the volume of a convex body.

Let $K$ be a convex body in $\mathbb{R}^{n}$ of diameter $D$ such that $B_{n} \subset K$. The next theorem from [4], improving on [13], essentially says that a deterministic algorithm cannot estimate the volume efficiently.

Theorem 9.1 For every deterministic algorithm that runs in time $O\left(n^{a}\right)$ and outputs two numbers $A$ and $B$ such that $A \leq \operatorname{vol}(K) \leq B$ for any convex body $K$, there is some convex body for which the ratio $B / A$ is at least

$$
\left(\frac{c n}{a \log n}\right)^{n}
$$

where $c$ is an absolute constant.
So, in polynomial-time, the best possible approximation is exponential in $n$ and to get a factor 2 approximation (say), one needs exponential time. The basic idea of the proof is simple. Consider an oracle that answers "yes" for any point in a unit ball and "no" to any point outside. After $m$ "yes" answers, the convex body $K$ could be anything between the ball and the convex hull of the $m$ query points. The ratio of these volumes is exponential in $n$.

Given this lower bound, the following result of Dyer, Frieze and Kannan [12] is quite remarkable.

| Authors | Complexity | New ingredient(s) |
| :--- | :---: | :---: |
| Dyer-Frieze-Kannan [12] | $n^{23}$ | Everything |
| Lovász-Simonovits [32] | $n^{16}$ | Localization lemma |
| Applegate-Kannan [3] | $n^{10}$ | Logconcave sampling |
| Lovász [29] | $n^{10}$ | Ball walk |
| Dyer-Frieze [11] | $n^{8}$ | Better error analysis |
| Lovász-Simonovits [34] | $n^{7}$ | Many improvements |
| Kannan-Lovász-Simonovits [23] | $n^{5}$ | Isotropy, speedy walk |
| Lovász-Vempala [37] | $n^{4}$ | Annealing, hit-and-run |

Table 1: Complexity comparison

Theorem 9.2 For any convex body $K$ and any $0 \leq \varepsilon, \delta \leq 1$, there is a randomized algorithm which computes an estimate $V$ such that with probability at least $1-\delta$, we have $(1-\varepsilon) \operatorname{vol}(K) \leq V \leq(1+\varepsilon) \operatorname{vol}(K)$, and the number of oracle calls is $\operatorname{poly}(n, 1 / \varepsilon, \log (1 / \delta)$.

Using randomness, we can go from an exponential approximation to an arbitrarily small one!

The main tool used in [12] is sampling by a random walk. They actually used the grid walk and showed that by "fixing up" $K$ a bit without changing its volume by much, the grid walk can sample nearly random points in polynomial time. Even though the walk is discrete, its analysis relies on a continuous isoperimetric inequality, quite similar to the one used for the analysis of the ball walk. The original algorithm of Dyer, Frieze and Kannan had complexity $O^{*}\left(n^{23}\right)$. In the years since, there have been many interesting improvements. These are summarized in Table 9. In this section, we describe the latest algorithm from [37] whose complexity is $O^{*}\left(n^{4}\right)$.

Let us first review the common structure of previous volume algorithms. Assume that the diameter $D$ of $K$ is poly $(n)$. All these algorithms reduce volume computation to sampling from a convex body, using the "Multi-Phase MonteCarlo" technique. They construct a sequence of convex bodies $K_{0} \subseteq K_{1} \subseteq$ $\cdots \subseteq K_{m}=K$, where $K_{0}=B_{n}$ or some body whose volume is easily computed. They estimate the ratios $\operatorname{vol}\left(K_{i-1}\right) / \operatorname{vol}\left(K_{i}\right)$ by generating sufficiently many independent (nearly) uniformly distributed random points in $K_{i}$ and counting the fraction that lie in $K_{i-1}$. The product of these estimates is an estimate of $\operatorname{vol}\left(K_{0}\right) / \operatorname{vol}(K)$.

In order to get a sufficiently good estimate for the ratio $\operatorname{vol}\left(K_{i-1}\right) / \operatorname{vol}\left(K_{i}\right)$, one needs about $m \operatorname{vol}\left(K_{i}\right) / \operatorname{vol}\left(K_{i-1}\right)$ random points. So we would like to have the ratios $\operatorname{vol}\left(K_{i}\right) / \operatorname{vol}\left(K_{i-1}\right)$ be small. But, the ratio of $\operatorname{vol}(K)$ and $\operatorname{vol}\left(K_{0}\right)$ could be $n^{\Omega(n)}$ and so $m$ has to be $\Omega(n)$ just to keep the ratios $\operatorname{vol}\left(K_{i}\right) / \operatorname{vol}\left(K_{i-1}\right)$ polynomial. The best choice is to keep these ratios bounded; this can be achieved e.g., if $K_{0}=B_{n}$ and $K_{i}=K \cap\left(2^{i / n} B_{n}\right)$ for $i=1,2, \ldots, m=\Theta(n \log n)$.

Thus, the total number of random points used is $O\left(m^{2}\right)=O^{*}\left(n^{2}\right)$. Since $\operatorname{vol}\left(K_{i}\right) \leq 2 \operatorname{vol}\left(K_{i-1}\right)$ for this sequence, a random point in $K_{i-1}$ provides a warm start for sampling from $K_{i}$. So each sample takes $O^{*}\left(n^{3}\right)$ steps to generate, giving an $O^{*}\left(n^{5}\right)$ algorithm. In [3, 34], sampling uniformly from $K_{i}$ was replaced by sampling from a smooth logconcave function to avoid bad local conductance and related issues.

The number of phases, $m$, enters the running time as its square and one would like to make it as small as possible. But, due to the reasons described above, $m=\Theta(n \log n)$ is optimal for this type of algorithm and reducing $m$ any further (i.e., $o(n)$ ) seems to be impossible for this type of method.

The algorithm in [37] can be viewed as a variation of simulated annealing. Introduced by Kirkpatrick et al. [26], simulated annealing is a general-purpose randomized search method for optimization. It walks randomly in the space of possible solutions, gradually adjusting a parameter called "temperature". At high temperature, the random walk converges to the uniform distribution over the whole space; as the temperature drops, the stationary distribution becomes more and more biased towards the optimal solutions.

Instead of a sequence of bodies, the algorithm in [37] constructs a sequence of functions $f_{0} \leq f_{1} \leq \cdots \leq f_{m}$ that "connect" a function $f_{0}$ whose integral is easy to find to the characteristic function $f_{m}$ of $K$. The ratios $\left(\int f_{i-1}\right) /\left(\int f_{i}\right)$ can be estimated by sampling from the distribution whose density function is proportional to $f_{i}$, and averaging the function $f_{i-1} / f_{i}$ over the sample points. Previous algorithms can be viewed as the special case where the functions $f_{i}$ are characteristic functions of the convex bodies $K_{i}$. By choosing a different set of $f_{i}$, the algorithm uses only $m=O^{*}(\sqrt{n})$ phases, and $O^{*}(\sqrt{n})$ sample points in each phase. In fact, it uses exponential functions of the form $f(x)=e^{-a^{T} x / T}$ restricted to some convex body. The temperature $T$ will start out at a small value and increase gradually. This is the reverse of what happens in simulated annealing.

Besides annealing, the algorithm uses a pre-processing step called the pencil construction. We describe it next.

Let $K$ be the given body in $\mathbb{R}^{n}$ and $\varepsilon>0$. Let $C$ denote the cone in $\mathbb{R}^{n+1}$ defined by

$$
C=\left\{x \in \mathbb{R}^{n+1}: x_{0} \geq 0, \sum_{i=1}^{n} x_{i}^{2} \leq x_{0}^{2}\right\}
$$

where $x=\left(x_{0}, x_{1} \ldots, x_{n}\right)^{\top}$. We define a new convex body $K^{\prime} \in \mathbb{R}^{n+1}$ as follows:

$$
K^{\prime}=([0,2 D] \times K) \cap C
$$

In words, $K^{\prime}$ is a sharpened $(n+1)$-dimensional "pencil" whose cross-section is $K$ and its tip is at the origin. See Fig. 4 for an illustration. The idea of the algorithm is to start with a function that is concentrated near the tip of the pencil, and is thus quite close to an exponential over a cone, and gradually move to a nearly constant function over the whole pencil, which would give us the volume of the pencil. The integral of the starting function is easily estimated.


Figure 4: The pencil construction when $K$ is a square.

The sharpening takes less than half of the volume of the pencil away. Hence, if we know the volume of $K^{\prime}$, it is easy to estimate the volume of $K$ by generating $O\left(1 / \varepsilon^{2}\right)$ sample points from the uniform distribution on $[0,2 D] \times K$ and then counting how many of them fall into $K^{\prime}$.

We describe the annealing part of the algorithm in a bit more detail. For each real number $a>0$, let

$$
Z(a)=\int_{K^{\prime}} e^{-a x_{0}} d x
$$

where $x_{0}$ is the first coordinate of $x$. For $a \leq \varepsilon^{2} / D$, an easy computation shows that

$$
(1-\varepsilon) \operatorname{vol}\left(K^{\prime}\right) \leq Z(a) \leq \operatorname{vol}(K)
$$

On the other hand, for $a \geq 2 n$ the value of $Z(a)$ is essentially the same as the integral over the whole cone which is easy to compute. So, if we select a sequence $a_{0}>a_{1}>\cdots>a_{m}$ for which $a_{0} \geq 2 n$ and $a_{m} \leq \varepsilon^{2} / D$, then we can estimate $\operatorname{vol}\left(K^{\prime}\right)$ by

$$
Z\left(a_{m}\right)=Z\left(a_{0}\right) \prod_{i=0}^{m-1} \frac{Z\left(a_{i+1}\right)}{Z\left(a_{i}\right)}
$$

The algorithm estimates each ratio $R_{i}=Z\left(a_{i+1}\right) / Z\left(a_{i}\right)$ as follows. Let $\mu_{i}$ be the probability distribution over $K^{\prime}$ with density proportional to $e^{-a_{i} x_{0}}$, i.e., for $x \in K^{\prime}$,

$$
\frac{d \mu_{i}(x)}{d x}=\frac{e^{-a_{i} x_{0}}}{Z\left(a_{i}\right)}
$$

To estimate the ratio $R_{i}$, the algorithm draws random samples $X^{1}, \ldots, X^{k}$ from $\mu_{i}$, and computes

$$
W_{i}=\frac{1}{m} \sum_{j=1}^{m} e^{\left(a_{i}-a_{i+1}\right)\left(X^{j}\right)_{0}}
$$

It is easy to see that $\mathrm{E}\left(W_{i}\right)=R_{i}$. The main lemma in the analysis is that the second moment of $W_{i}$ is small.

Lemma 9.3 Let $X$ be a random sample from $d \mu_{i}$ and

$$
Y=e^{\left(a_{i}-a_{i+1}\right) X_{0}}
$$

Then,

$$
\frac{\mathrm{E}\left(Y^{2}\right)}{\mathrm{E}(Y)^{2}} \leq \frac{a_{i+1}^{2}}{a_{i}\left(2 a_{i+1}-a_{i}\right)}
$$

From the lemma, it follows that with

$$
a_{i+1}=a_{i}\left(1-\frac{1}{\sqrt{n}}\right)
$$

we get that $\mathrm{E}\left(Y^{2}\right) / \mathrm{E}(Y)^{2}$ is bounded by a constant. So with $k$ samples,

$$
\frac{\mathrm{E}\left(W_{i}^{2}\right)}{\mathrm{E}\left(W_{i}\right)^{2}} \leq\left(1+\frac{O(1)}{k}\right)
$$

Hence, the standard deviation of the estimate $Z=W_{1} W_{2} \ldots W_{m}$ is at most $\varepsilon$ times $\mathrm{E}\left(W_{1} W_{2} \ldots W_{m}\right)=\operatorname{vol}(K)$, for $k=O\left(m / \varepsilon^{2}\right)$. Further, the number of phases needed to go from $a_{0}=2 n$ to $a_{m} \leq \varepsilon^{2} / D$ is only $\sqrt{n} \log \left(D / \varepsilon^{2}\right)$. So the total number of sample points needed is only $O^{*}(n)$ (it would be interesting to show that this is a lower bound for any algorithm that uses a blackbox sampler).

As mentioned earlier, the samples obtained are not truly independent. This introduces technical complications. In previous algorithms, the random variable estimating the ratio was bounded (between 1 and 2 ) and so we could directly use Lemma 7.1. For the new algorithm, the individual ratios being estimated could be unbounded. To handle this further properties of $\mu$-independence are developed in [37]. We do not go into the details here.

How fast can we sample from $\mu_{i}$ ? Sampling from $\mu_{0}$ is easy. But each $\mu_{i-1}$ no longer provides a warm start for $\mu_{i}$, i.e., $d \mu_{i-1}(x) / d \mu_{i}(x)$ could be unbounded. However, as the next lemma asserts, the $L_{2}$ distance is bounded.
Lemma $9.4\left\|\mu_{i-1} / \mu_{i}\right\|<8$.
The proof of this lemma and that of Lemma 9.3 are both based on the following property of logconcave functions proved in [37].

Lemma 9.5 For $a>0$, any convex body $K$ and logconcave function $f$ in $\mathbb{R}^{n}$, the function

$$
Z(a)=a^{n} \int_{K} f(a x) d x
$$

is logconcave.
Finally, hit-and-run only needs bounded $L_{2}$ norm to sample efficiently, and by the version of Theorem 6.3 for the exponential density, we get each sample in $O^{*}\left(n^{3}\right)$ time. Along with the bound on the number of samples, this gives the complexity bound of $O *\left(n^{4}\right)$ for the volume algorithm.

It is apparent that any improvement in the mixing rate of random walks will directly affect the complexity of volume computation. Such improvements seem to consistently yield interesting new mathematics as well.

## References

[1] D. Aldous and J. Fill: Reversible Markov Chains and Random Walks on Graphs. Monograph in preparation. http://statwww.berkeley.edu/users/aldous/RWG/book.html
[2] N. Alon: Eigenvalues and expanders. Combinatorica, 6, (1986), 83-96.
[3] D. Applegate and R. Kannan: Sampling and integration of near log-concave functions. Proc. 23th ACM STOC (1990), 156-163.
[4] I. Bárány and Z. Füredi: Computing the Volume is Difficult. Discrete \& Computational Geometry 2, (1987), 319-326.
[5] H. C. Berbee, C.G. Boender, A. H. G. Rinnooy Kan, C. L. Scheffer, R. L. Smith and J. Telgen: Hit-and-run algorithms for the identification of nonredundant linear inequalities. Math. Prog., 37, (1987), 184-207.
[6] D. Bertsimas and S. Vempala: Solving convex programs by random walks. J. ACM 51(4), (2004), 540-556.
[7] B. Bollobás: Volume estimates and rapid mixing. In Flavors of geometry, 151-182, MSRI Publ. 31, Cambridge University Press, Cambridge 1997.
[8] P. Diaconis and L. Saloff-Coste: Logarithmic Sobolev inequalities for finite Markov chains. Ann. Appl. Prob., 6(3), (1996), 695-750.
[9] A. Dinghas: Über eine Klasse superadditiver Mengenfunktionale von Brunn-Minkowski-Lusternik-schem Typus. Math. Zeitschr. 68 (1957), 111-125.
[10] J. Dodziuk and W. S. Kendall: Combinatorial Laplacians and isoperimetric inequality. In From Local Times to Global Geometry, Control and Physics, (Ed. K.D. Ellworthy), Pitman Res. Notes in Math. 150, (1986), 68-74.
[11] M. Dyer and A. Frieze: Computing the volume of convex bodies: a case where randomness provably helps. Proc. Symp. Appl. Math. 44, (1991), 123-169.
[12] M. Dyer, A. Frieze and R. Kannan: A random polynomial-time algorithm for approximating the volume of convex bodies. J. Assoc. Comput. Mach. 38 (1991), 1-17.
[13] G. Elekes: A Geometric Inequality and the Complexity of Computing Volume. Discrete \& Computational Geometry 1, (1986), 289-292.
[14] A. Frieze, R. Kannan and N. Polson: Sampling from log-concave distributions: Ann. Appl. Probab. 4 (1994), 812-837; Correction: Sampling from log-concave distributions, ibid. 4 (1994), 1255.
[15] M. Grötschel, L. Lovász, and A. Schrijver: Geometric Algorithms and Combinatorial Optimization, Springer, 1988.
[16] B. Grünbaum: Partitions of mass-distributions and convex bodies by hyperplanes, Pacific J. Math. 10, (1960), 1257-1261.
[17] M. Jerrum and A. Sinclair: Approximating the permanent. Siam J. Computing 18, (1989), 1149-1178.
[18] A. Kalai and S. Vempala: Efficient algorithms for universal portfolios. J. Machine Learning Research 3, (2002), 423-440.
[19] A. Kalai and S. Vempala: Simulated Annealing for Convex Optimization. Manuscript available at http://wwwmath.mit.edu/~vempala/papers/anneal.ps
[20] R. Kannan, L. Lovász: Faster mixing via average conductance. Proc. of STOC, (1999), 282-287.
[21] R. Kannan, L. Lovász and R. Montenegro: Blocking conductance and mixing in random walks. Manuscript.
[22] R. Kannan, L. Lovász and M. Simonovits: Isoperimetric problems for convex bodies and a localization lemma. Discr. Comput. Geom. 13 (1995), 541-559.
[23] R. Kannan, L. Lovász and M. Simonovits: Random walks and an $O^{*}\left(n^{5}\right)$ volume algorithm for convex bodies. Random Structures and Algorithms 11 (1997), 1-50.
[24] R. Kannan, J. Mount, S. Tayur: A randomized algorithm to optimize over certain convex sets. Math. Oper. Res. 20(3), 529-549, 1995.
[25] L. G. Khachiyan: A polynomial algorithm in linear programming. (in Russian), Doklady Akedamii Nauk SSSR, 244, 1093-1096, 1979 (English translation: Soviet Mathematics Doklady, 20, 191-194, 1979).
[26] S. Kirkpatrick, C.D. Gelatt Jr, and M.P. Vecchi: Optimization by Simulated Annealing. Science 220 (1983), 671—680.
[27] L. Leindler: On a certain converse of Hölder's Inequality II. Acta Sci. Math. Szeged 33 (1972), 217-223.
[28] A. Yu. Levin: On an algorithm for the minimization of convex functions. (in Russian), Doklady Akedemii Nauk SSSR, 160, 1244-1247, 1965 (English translation: Soviet Mathematics Doklady, 6, 286-290, 1965).
[29] L. Lovász: How to compute the volume? Jber. d. Dt. Math.-Verein, Jubiläumstagung 1990, B. G. Teubner, Stuttgart, 138-151.
[30] L. Lovász: Hit-and-run mixes fast. Mathematical Programming 86, (1999), 443-461.
[31] L. Lovász: Random Walks on Graphs: A Survey. Combinatorics, Paul Erdős is Eighty, Vol. 2 (ed. D. Miklós, V. T. Sós, T. Szőnyi), János Bolyai Mathematical Society, Budapest, 1996, 353-397.
[32] L. Lovász and M. Simonovits: The Mixing rate of Markov chains, an isoperimetric inequality, and computing the volume. Proc. 31st IEEE Annual Symp. on Found. of Comp. Sci. (1990), 346-354.
[33] L. Lovász and M. Simonovits: On the randomized complexity of volume and diameter. Proc. 33rd IEEE Annual Symp. on Found. of Comp. Sci. (1992), 482-491.
[34] L. Lovász and M. Simonovits: Random walks in a convex body and an improved volume algorithm. Random Structures and Alg. 4 (1993), 359412.
[35] L. Lovász and S. Vempala: The geometry of logconcave functions and an $O^{*}\left(n^{3}\right)$ sampling algorithm. Microsoft Research Tech. Rep. MSR-TR-200304. Available at: http://www-math.mit.edu/~vempala/papers/logconball.ps (preliminary version in Proc. of FOCS, 2003).
[36] L. Lovász and S. Vempala: Hit-and-run is fast and fun. Microsoft Research Tech. Rep. MSR-TR-2003-05. Available at http://wwwmath.mit.edu/ ${ }^{\text {vempala/papers/logcon-hitrun.ps (preliminary version in }}$ Proc. of FOCS, 2003).
[37] L. Lovász and S. Vempala: Simulated Annealing in Convex Bodies an $O^{*}\left(n^{4}\right)$ volume algorithm. Proc. of FOCS (2003), 650-659. http://wwwmath.mit.edu/~vempala/papers/vol4.ps
[38] L. Lovász and S. Vempala: Hit-and-run from a corner. Proc. of STOC (2004), 310-314. http://www-math.mit.edu/~ vempala/papers/start.ps
[39] J. R. Norris: Markov Chains. Cambridge University Press, 1997.
[40] A. Prékopa: Logarithmic concave measures with applications to stochastic programming. Acta Sci. Math. Szeged 32 (1971), 301-316.
[41] A. Prékopa: On Logarithmic Concave Measures and Functions. Acta Sci. Math. Szeged 34 (1973), 335-343.
[42] D. Revuz: Markov Chains. North Holland, 1975.
[43] M. Rudelson: Random vectors in the isotropic position. J. Funct. Anal. 164 (1999), 60-72.
[44] M. Simonovits: How to compute the volume in high dimension? Math. Prog. (B), 97, (2003), 337-374.
[45] R.L. Smith: Efficient Monte-Carlo procedures for generating points uniformly distributed over bounded regions. Operations Res. 32, (1998), 12961308.
[46] P. M. Vaidya: A new algorithm for minimizing convex functions over convex sets. Mathematical Programming, 73, (1996), 291-341.
[47] D. B. Yudin and A. S. Nemirovski: Estimation of the information complexity of mathematical programming problems (in Russian). Ekonomika i Matematicheskie Metody 12, (1976), 128-142. (English translation: Matekon 13, 2, (1976), 3-45).
[48] Z.B. Zabinsky, R.L. Smith, J.F. McDonald, H.E. Romeijn and D.E. Kaufman: Improving hit-and-run for global optimization. J. Global Optim. $\mathbf{3}$ (1993), 171-192.


[^0]:    *To appear in the MSRI volume on the Foundations of Geometric Algorithms
    $\dagger$ Supported by NSF award CCR-0307536 and a Sloan foundation fellowship.

[^1]:    ${ }^{1}$ For Markov chains on discrete state spaces, the characterization is much simpler; see e.g., [39].

[^2]:    ${ }^{2}$ In fact, it is enough for $f$ to be quasi-convex.

