
Rapid Convergence of the Unadjusted Langevin Algorithm: Isoperimetry Suffices

Santosh S. Vempala
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332
vempala@gatech.edu

Andre Wibisono
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332
wibisono@gatech.edu

Abstract

We study the Unadjusted Langevin Algorithm (ULA) for sampling from a probability distribution $\nu = e^{-f}$ on \mathbb{R}^n . We prove a convergence guarantee in Kullback-Leibler (KL) divergence assuming ν satisfies log-Sobolev inequality and f has bounded Hessian. Notably, we do not assume convexity or bounds on higher derivatives. We also prove convergence guarantees in Rényi divergence of order $q > 1$ assuming the limit of ULA satisfies either log-Sobolev or Poincaré inequality.

1 Introduction

Sampling is a fundamental algorithmic task. Many applications require sampling from probability distributions in high-dimensional spaces, and in modern applications the probability distributions are complicated and non-logconcave. While the setting of logconcave functions is well-studied, it is important to have efficient sampling algorithms with good convergence guarantees beyond the logconcavity assumption. There is a close interplay between sampling and optimization, either via optimization as a limit of sampling (annealing) [34, 55], or via sampling as optimization in the space of distributions [36, 62]. Motivated by the widespread use of non-convex optimization and sampling, there is resurgent interest in understanding non-logconcave sampling.

In this paper we study a simple algorithm, the Unadjusted Langevin Algorithm (ULA), for sampling from a target probability distribution $\nu = e^{-f}$ on \mathbb{R}^n . ULA requires oracle access to the gradient ∇f of the log density $f = -\log \nu$. In particular, ULA does not require knowledge of f , which makes it applicable in practice where we often only know ν up to a normalizing constant.

As the step size $\epsilon \rightarrow 0$, ULA recovers the Langevin dynamics, which is a continuous-time stochastic process in \mathbb{R}^n that converges to ν . We recall the optimization interpretation of the Langevin dynamics for sampling as the gradient flow of the Kullback-Leibler (KL) divergence with respect to ν in the space of probability distributions with the Wasserstein metric [36]. When ν is strongly logconcave, the KL divergence is a strongly convex objective function, so the Langevin dynamics as gradient flow converges exponentially fast [6, 60]. From the classical theory of Markov chains and diffusion processes, there are several known conditions milder than logconcavity that are sufficient for rapid convergence *in continuous time*. These include isoperimetric inequalities such as Poincaré inequality or log-Sobolev inequality (LSI). Along the Langevin dynamics in continuous time, Poincaré inequality implies an exponential convergence rate in χ^2 -divergence, while LSI—which is stronger—implies an exponential convergence rate in KL divergence (as well as in Rényi divergence).

However, in discrete time, sampling under Poincaré inequality or LSI is a more challenging problem. ULA is an inexact discretization of the Langevin dynamics, and it converges to a biased limit $\nu_\epsilon \neq \nu$. When ν is strongly logconcave and smooth, it is known how to control the bias and prove a convergence guarantee on KL divergence along ULA [17, 21, 22, 24]. When ν is strongly

logconcave, there are many sampling algorithms with provable rapid convergence; these include the ball walk and hit-and-run [37, 43, 44, 42] (which give truly polynomial algorithms), various discretizations of the overdamped or underdamped Langevin dynamics [21, 22, 24, 8, 26] (which have polynomial dependence on smoothness parameters but low dependence on dimension), and the Hamiltonian Monte Carlo [47, 48, 25, 39, 16]. It is of great interest to extend these results to non-logconcave densities ν , where existing results require strong assumptions with bounds that grow exponentially with the dimension or other parameters [2, 18, 45, 49]. There are recent works that analyze convergence of sampling using various techniques such as reflection coupling [28], kernel methods [29], and higher-order integrators [40], albeit still under some strong conditions such as distant dissipativity, which is similar to strong logconcavity outside a bounded domain.

In this paper we study the convergence along ULA under minimal (and necessary) isoperimetric assumptions, namely, LSI and Poincaré inequality. These are sufficient for fast convergence in continuous time; moreover, in the case of logconcave distribution, the log-Sobolev and Poincaré constants can be bounded and lead to convergence guarantees for efficient sampling in discrete time. However, do they suffice on their own without the assumption of logconcavity?

We note that LSI and Poincaré inequality apply to a wider class of measures than logconcave distributions. In particular, LSI and Poincaré inequality are preserved under bounded perturbation and Lipschitz mapping, whereas logconcavity is destroyed. Given these properties, it is easy to exhibit examples of non-logconcave distributions satisfying LSI or Poincaré inequality. For example, we can take a small perturbation of a convex body to make it nonconvex but still satisfies isoperimetry; then the uniform probability distribution (or a smooth version of it) on the body is not logconcave but satisfies LSI and Poincaré inequality. Similarly, we can start with a strongly logconcave distribution and make bounded perturbations; then the resulting (normalized) probability distribution is not logconcave, but it satisfies LSI and Poincaré inequality. See Figure 1 for an illustration.




Figure 1: Illustrations of non-logconcave distributions satisfying isoperimetry: uniform distribution on a nonconvex set (left) and a perturbation of a logconcave distribution (right).

We measure the mode of convergence using KL divergence and Rényi divergence of order $q \geq 1$, which is stronger. Our first main result says the only further assumption we need is smoothness. We say $\nu = e^{-f}$ is L -smooth if ∇f is L -Lipschitz. Here $H_\nu(\rho)$ is the KL divergence between ρ and ν . See Theorem 2 in Section 3.1 for more detail.

Theorem 2. *Assume $\nu = e^{-f}$ satisfies log-Sobolev inequality with constant $\alpha > 0$ and is L -smooth. ULA with step size $0 < \epsilon \leq \frac{\alpha}{4L^2}$ satisfies*

$$H_\nu(\rho_k) \leq e^{-\alpha\epsilon k} H_\nu(\rho_0) + \frac{8\epsilon n L^2}{\alpha}.$$

For $0 < \delta < 4n$, ULA with $\epsilon \leq \frac{\alpha\delta}{16L^2n}$ reaches error $H_\nu(\rho_k) \leq \delta$ after $k \geq \frac{1}{\alpha\epsilon} \log \frac{2H_\nu(\rho_0)}{\delta}$ iterations.

For example, if we start with a Gaussian $\rho_0 = \mathcal{N}(x^*, \frac{1}{L}I)$ where x^* is a stationary point of f (which we can find, e.g., via gradient descent), then $H_\nu(\rho_0) = \tilde{O}(n)$ (see Lemma 1), and Theorem 2 gives an iteration complexity of $k = \tilde{\Theta}(\frac{L^2n}{\alpha^2\delta})$ to achieve $H_\nu(\rho_k) \leq \delta$ using ULA with step size $\epsilon = \Theta(\frac{\alpha\delta}{L^2n})$.

The result above matches previous known bounds for ULA when ν is strongly logconcave [17, 21, 22, 24]. Our result complements the recent work of Ma et al. [45] who study the underdamped version of the Langevin dynamics under LSI and show an iteration complexity for the discrete-time algorithm that has better dependence on the dimension ($\sqrt{\frac{n}{\delta}}$ in place of $\frac{n}{\delta}$ above for ULA), but under an additional smoothness assumption (f has bounded third derivatives) and with higher polynomial dependence on other parameters. Our result also complements the work of Mangoubi and Vishnoi [49] who study the Metropolis-adjusted version of ULA (MALA) for non-logconcave ν and show a $\log(\frac{1}{\delta})$ iteration complexity from a warm start, under the additional assumption that f has bounded third and fourth derivatives in an appropriate ∞ -norm.

We note that in general some isoperimetry condition is needed for rapid mixing of Markov chains (such as Langevin dynamics and ULA), otherwise there are bad regions in the state space from which the chains take arbitrarily long to escape. Smoothness or bounded Hessian is a common assumption needed for the analysis of discrete-time algorithms (such as gradient descent or ULA).

In the second part of this paper, we study the convergence of Rényi divergence of order $q > 1$ along ULA. Rényi divergence is a family of generalizations of KL divergence [56, 59, 11], which becomes stronger as the order q increases. There are physical and operational interpretations of Rényi divergence [31, 3]. Rényi divergence has been useful in many applications, including for the exponential mechanism in differential privacy [27, 1, 12, 52], lattice-based cryptography [4], information-theoretic encryption [35], variational inference [41], machine learning [32, 50], information theory and statistics [20, 53], and black hole physics [23].

Our second result proves a convergence bound for the Rényi divergence of order $q > 1$. While Rényi divergence is a stronger measure of convergence than KL divergence, the situation is more complicated. First, we can hope to converge to the biased limit ν_ϵ only for finite q for any step size ϵ (as we illustrate with an example). Second, it is unclear how to bound the Rényi divergence between ν_ϵ and ν . We first show the following convergence guarantees of Rényi divergence along the continuous-time Langevin dynamics under LSI or Poincaré inequality; see Theorem 3 and Theorem 5. Here $R_{q,\nu}(\rho)$ is the Rényi divergence of order q between ρ and ν .

Theorem 3. *Suppose ν satisfies LSI with constant $\alpha > 0$. Let $q \geq 1$. Along the Langevin dynamics,*

$$R_{q,\nu}(\rho_t) \leq e^{-\frac{2\alpha t}{q}} R_{q,\nu}(\rho_0).$$

Theorem 5. *Suppose ν satisfies Poincaré inequality with constant $\alpha > 0$. Let $q \geq 2$. Along the Langevin dynamics,*

$$R_{q,\nu}(\rho_t) \leq \begin{cases} R_{q,\nu}(\rho_0) - \frac{2\alpha t}{q} & \text{if } R_{q,\nu}(\rho_0) \geq 1 \text{ and as long as } R_{q,\nu}(\rho_t) \geq 1, \\ e^{-\frac{2\alpha t}{q}} R_{q,\nu}(\rho_0) & \text{if } R_{q,\nu}(\rho_0) \leq 1. \end{cases}$$

Notice that under Poincaré inequality, compared to LSI, the convergence is slower in the beginning before it becomes exponential. For a reasonable starting distribution (such as a Gaussian centered at a stationary point), this leads to an extra factor of n compared to the convergence under LSI. We then turn to discrete time and show the convergence of Rényi divergence along ULA to the biased limit ν_ϵ under the assumption that ν_ϵ itself satisfies either LSI or Poincaré inequality. We combine this with a decomposition result on Rényi divergence to derive a convergence guarantee for Rényi divergence to ν along ULA; see Theorem 4 and Theorem 6.

In what follows, we review KL divergence and its properties along the Langevin dynamics in Section 2 and prove a convergence guarantee for KL divergence along ULA under LSI in Section 3. We provide a review of Rényi divergence and its properties along the Langevin dynamics in Section 4. We then prove the convergence guarantee for Rényi divergence along ULA under LSI in Section 5 and under Poincaré inequality in Section 6. We conclude with a discussion in Section 7.

2 Review of KL divergence along Langevin dynamics

In this section we review the definition of Kullback-Leibler (KL) divergence, log-Sobolev inequality, and the convergence of KL divergence along the Langevin dynamics in continuous time under log-Sobolev inequality. See Appendix A.1 for a review on notation.

2.1 KL divergence

Let ρ, ν be probability distributions on \mathbb{R}^n , represented via their probability density functions with respect to the Lebesgue measure on \mathbb{R}^n . We assume ρ, ν have full support and smooth densities.

Recall the **Kullback-Leibler (KL) divergence** of ρ with respect to ν is

$$H_\nu(\rho) = \int_{\mathbb{R}^n} \rho(x) \log \frac{\rho(x)}{\nu(x)} dx. \quad (1)$$

KL divergence is the relative form of *Shannon entropy* $H(\rho) = -\int_{\mathbb{R}^n} \rho(x) \log \rho(x) dx$. Whereas Shannon entropy can be positive or negative, KL divergence is nonnegative and minimized at ν :

$H_\nu(\rho) \geq 0$ for all ρ , and $H_\nu(\rho) = 0$ if and only if $\rho = \nu$. Therefore, KL divergence serves as a measure of (albeit asymmetric) “distance” of a probability distribution ρ from a base distribution ν . KL divergence is a relatively strong measure of distance; for example, Pinsker’s inequality implies that KL divergence controls total variation distance. Furthermore, under log-Sobolev (or Talagrand) inequality, KL divergence also controls the quadratic Wasserstein W_2 distance, as we review below. We say $\nu = e^{-f}$ is *L-smooth* if f has bounded Hessian: $-LI \preceq \nabla^2 f(x) \preceq LI$ for all $x \in \mathbb{R}^n$. We provide the proof of Lemma 1 in Appendix B.1.1.

Lemma 1. *Suppose $\nu = e^{-f}$ is L-smooth. Let $\rho = \mathcal{N}(x^*, \frac{1}{L}I)$ where x^* is a stationary point of f . Then $H_\nu(\rho) \leq f(x^*) + \frac{n}{2} \log \frac{L}{2\pi}$.*

2.2 Log-Sobolev inequality

Recall we say ν satisfies the **log-Sobolev inequality (LSI)** with a constant $\alpha > 0$ if for all smooth function $g: \mathbb{R}^n \rightarrow \mathbb{R}$ with $\mathbb{E}_\nu[g^2] < \infty$,

$$\mathbb{E}_\nu[g^2 \log g^2] - \mathbb{E}_\nu[g^2] \log \mathbb{E}_\nu[g^2] \leq \frac{2}{\alpha} \mathbb{E}_\nu[\|\nabla g\|^2]. \quad (2)$$

Recall the **relative Fisher information** of ρ with respect to ν is

$$J_\nu(\rho) = \int_{\mathbb{R}^n} \rho(x) \left\| \nabla \log \frac{\rho(x)}{\nu(x)} \right\|^2 dx. \quad (3)$$

LSI is equivalent to the following relation between KL divergence and Fisher information for all ρ :

$$H_\nu(\rho) \leq \frac{1}{2\alpha} J_\nu(\rho). \quad (4)$$

Indeed, to obtain (4) we choose $g^2 = \frac{\rho}{\nu}$ in (2); conversely, to obtain (2) we choose $\rho = \frac{g^2 \nu}{\mathbb{E}_\nu[g^2]}$ in (4).

LSI is an isoperimetry condition and implies, among others, concentration of measure and sub-Gaussian tail property [38]. LSI was first shown by Gross [30] for the case of Gaussian ν . It was extended by Bakry and Émery [6] to strongly log-concave ν ; namely, when $f = -\log \nu$ is α -strongly convex, then ν satisfies LSI with constant α . However, LSI applies more generally. For example, the classical perturbation result by Holley and Stroock [33] states that LSI is stable under bounded perturbation. Furthermore, LSI is preserved under a Lipschitz mapping. In one dimension, there is an exact characterization of when a probability distribution on \mathbb{R} satisfies LSI [9]. Moreover, LSI satisfies a tensorization property [38]: If ν_1, ν_2 satisfy LSI with constants $\alpha_1, \alpha_2 > 0$, respectively, then $\nu_1 \otimes \nu_2$ satisfies LSI with constant $\min\{\alpha_1, \alpha_2\} > 0$. Thus, there are many examples of non-logconcave distributions ν on \mathbb{R}^n satisfying LSI (with a constant independent of dimension). There are also Lyapunov function criteria and exponential integrability conditions that can be used to verify when a probability distribution satisfies LSI; see for example [14, 15, 51, 61, 7].

2.2.1 Talagrand inequality

Recall the **Wasserstein distance** between ρ and ν is

$$W_2(\rho, \nu) = \inf_{\Pi} \mathbb{E}_\Pi[\|X - Y\|^2]^{\frac{1}{2}} \quad (5)$$

where the infimum is over joint distributions Π of (X, Y) with the correct marginals $X \sim \rho, Y \sim \nu$. Recall we say ν satisfies **Talagrand inequality** with a constant $\alpha > 0$ if for all ρ :

$$\frac{\alpha}{2} W_2(\rho, \nu)^2 \leq H_\nu(\rho). \quad (6)$$

Talagrand’s inequality implies concentration of measure of Gaussian type. It was first studied by Talagrand [58] for Gaussian ν , and extended by Otto and Villani [54] to all ν satisfying LSI; namely, if ν satisfies LSI with constant $\alpha > 0$, then ν also satisfies Talagrand’s inequality with the same constant [54, Theorem 1]. Therefore, under LSI, KL divergence controls the Wasserstein distance. Moreover, when ν is log-concave, LSI and Talagrand’s inequality are equivalent [54, Corollary 3.1]. We recall in Appendix A.2 the geometric interpretation of LSI and Talagrand’s inequality from [54].

2.3 Langevin dynamics

The **Langevin dynamics** for target distribution $\nu = e^{-f}$ is a continuous-time stochastic process $(X_t)_{t \geq 0}$ in \mathbb{R}^n that evolves following the stochastic differential equation:

$$dX_t = -\nabla f(X_t) dt + \sqrt{2} dW_t \quad (7)$$

where $(W_t)_{t \geq 0}$ is the standard Brownian motion in \mathbb{R}^n with $W_0 = 0$.

If $(X_t)_{t \geq 0}$ evolves following the Langevin dynamics (7), then their probability density function $(\rho_t)_{t \geq 0}$ evolves following the **Fokker-Planck equation**:

$$\frac{\partial \rho_t}{\partial t} = \nabla \cdot (\rho_t \nabla f) + \Delta \rho_t = \nabla \cdot \left(\rho_t \nabla \log \frac{\rho_t}{\nu} \right). \quad (8)$$

Here $\nabla \cdot$ is the divergence and Δ is the Laplacian operator. We provide a derivation in Appendix A.3. From (8), if $\rho_t = \nu$, then $\frac{\partial \rho_t}{\partial t} = 0$, so ν is the stationary distribution for the Langevin dynamics (7). Moreover, the Langevin dynamics brings any distribution $X_t \sim \rho_t$ closer to the target distribution ν , as the following lemma shows.

Lemma 2. *Along the Langevin dynamics (7) (or equivalently, the Fokker-Planck equation (8)),*

$$\frac{d}{dt} H_\nu(\rho_t) = -J_\nu(\rho_t). \quad (9)$$

We provide the proof of Lemma 2 in Appendix B.1.2. Since $J_\nu(\rho) \geq 0$, the identity (9) shows KL divergence is decreasing along the Langevin dynamics, so indeed the distribution ρ_t converges to ν .

2.3.1 Exponential convergence of KL divergence along Langevin dynamics under LSI

When ν satisfies LSI, KL divergence converges exponentially fast along the Langevin dynamics.

Theorem 1. *Suppose ν satisfies LSI with constant $\alpha > 0$. Along the Langevin dynamics (7),*

$$H_\nu(\rho_t) \leq e^{-2\alpha t} H_\nu(\rho_0). \quad (10)$$

Furthermore, $W_2(\rho_t, \nu) \leq \sqrt{\frac{2}{\alpha} H_\nu(\rho_0)} e^{-\alpha t}$.

We provide the proof of Theorem 1 in Appendix B.1.3. We also recall the optimization interpretation of Langevin dynamics as the gradient flow of KL divergence in the space of distributions with the Wasserstein metric [36, 60, 54]. Then the exponential convergence rate in Theorem 1 is a manifestation of the general fact that gradient flow converges exponentially fast under gradient domination condition. This provides a justification for using the Langevin dynamics for sampling from ν , as a natural steepest descent flow that minimizes the KL divergence H_ν .

3 Unadjusted Langevin Algorithm

Suppose we wish to sample from a smooth target probability distribution $\nu = e^{-f}$ in \mathbb{R}^n . The **Unadjusted Langevin Algorithm (ULA)** with step size $\epsilon > 0$ is the discrete-time algorithm

$$x_{k+1} = x_k - \epsilon \nabla f(x_k) + \sqrt{2\epsilon} z_k \quad (11)$$

where $z_k \sim \mathcal{N}(0, I)$ is an independent standard Gaussian random variable in \mathbb{R}^n . Let ρ_k denote the probability distribution of x_k that evolves following ULA.

As $\epsilon \rightarrow 0$, ULA recovers the Langevin dynamics (7) in continuous-time. However, for fixed $\epsilon > 0$, ULA converges to a biased limiting distribution $\nu_\epsilon \neq \nu$. Therefore, KL divergence $H_\nu(\rho_k)$ does not tend to 0 along ULA, as it has an asymptotic bias $H_\nu(\nu_\epsilon) > 0$.

Example 1. *Let $\nu = \mathcal{N}(0, \frac{1}{\alpha} I)$. The ULA iteration is $x_{k+1} = (1 - \epsilon\alpha)x_k + \sqrt{2\epsilon}z_k$. For $0 < \epsilon < \frac{2}{\alpha}$, the limit is $\nu_\epsilon = \mathcal{N}(0, \frac{1}{\alpha(1 - \frac{\epsilon\alpha}{2})})$ and the bias is $H_\nu(\nu_\epsilon) = \frac{n}{2} \left(\frac{\epsilon\alpha}{2(1 - \frac{\epsilon\alpha}{2})} + \log(1 - \frac{\epsilon\alpha}{2}) \right)$. In particular, $H_\nu(\nu_\epsilon) \leq \frac{n\epsilon^2\alpha^2}{16(1 - \frac{\epsilon\alpha}{2})^2} = O(\epsilon^2)$.*

3.1 Convergence of KL divergence along ULA under LSI

When ν satisfies LSI and a smoothness condition, we can prove a convergence guarantee in KL divergence along ULA. Recall we say $\nu = e^{-f}$ is L -smooth if $-LI \preceq \nabla^2 f(x) \preceq LI$ for all $x \in \mathbb{R}^n$. A key part in our analysis is the following lemma which bounds the decrease in KL divergence along one iteration of ULA. Here $x_{k+1} \sim \rho_{k+1}$ is the output of one step of ULA (11) from $x_k \sim \rho_k$.

Lemma 3. *Suppose ν satisfies LSI with constant $\alpha > 0$ and is L -smooth. If $0 < \epsilon \leq \frac{\alpha}{4L^2}$, then along each step of ULA (11),*

$$H_\nu(\rho_{k+1}) \leq e^{-\alpha\epsilon} H_\nu(\rho_k) + 6\epsilon^2 nL^2. \quad (12)$$

We provide the proof of Lemma 3 in Appendix B.2.1. The proof of Lemma 3 compares the evolution of KL divergence along one step of ULA with the evolution along Langevin dynamics in continuous time (which converges exponentially fast under LSI), and bounds the discretization error; see Figure 2 for an illustration. This comparison technique has been used in many papers. Our proof structure is similar to that of Cheng and Bartlett (17), whose analysis needs ν to be strongly logconcave.

With Lemma 3 we can prove our main result on the convergence rate of ULA under LSI. We provide the proof of Theorem 2 in Appendix B.2.2.

Theorem 2. *Assume $\nu = e^{-f}$ satisfies log-Sobolev inequality with constant $\alpha > 0$ and is L -smooth. ULA with step size $0 < \epsilon \leq \frac{\alpha}{4L^2}$ satisfies*

$$H_\nu(\rho_k) \leq e^{-\alpha\epsilon k} H_\nu(\rho_0) + \frac{8\epsilon nL^2}{\alpha}.$$

For $0 < \delta < 4n$, ULA with $\epsilon \leq \frac{\alpha\delta}{16L^2n}$ reaches error $H_\nu(\rho_k) \leq \delta$ after $k \geq \frac{1}{\alpha\epsilon} \log \frac{2H_\nu(\rho_0)}{\delta}$ iterations.

In particular, suppose $\delta < 4n$ and we choose the largest permissible step size $\epsilon = \Theta\left(\frac{\alpha\delta}{L^2n}\right)$. Suppose we start with a Gaussian $\rho_0 = \mathcal{N}(x^*, \frac{1}{L}I)$, where x^* is a stationary point of f (which we can find, e.g., via gradient descent), so $H_\nu(\rho_0) \leq f(x^*) + \frac{n}{2} \log \frac{L}{2\pi} = \tilde{O}(n)$ by Lemma 1. Theorem 2 states that to achieve $H_\nu(\rho_k) \leq \delta$, ULA has iteration complexity $k = \tilde{\Theta}\left(\frac{L^2n}{\alpha^2\delta}\right)$. Since LSI implies Talagrand's inequality, Theorem 2 also yields a convergence guarantee in Wasserstein distance. As $k \rightarrow \infty$, Theorem 2 implies the following bound on the bias of ULA under LSI. However, we note the bound $O(\epsilon)$ may be loose, since from Example 1 we see $H_\nu(\nu_\epsilon) = \Theta(\epsilon^2)$ in Gaussian case.

Corollary 1. *Suppose ν satisfies LSI with constant $\alpha > 0$ and is L -smooth. For $0 < \epsilon \leq \frac{\alpha}{4L^2}$, the biased limit ν_ϵ of ULA with step size ϵ satisfies $H_\nu(\nu_\epsilon) \leq \frac{8nL^2\epsilon}{\alpha}$ and $W_2(\nu, \nu_\epsilon)^2 \leq \frac{16nL^2\epsilon}{\alpha^2}$.*

4 Review of Rényi divergence along Langevin dynamics

4.1 Rényi divergence

Rényi divergence (56) is a family of generalizations of KL divergence. See (59, 11) for properties of Rényi divergence.

For $q > 0, q \neq 1$, the **Rényi divergence** of order q of a probability distribution ρ with respect to ν is

$$R_{q,\nu}(\rho) = \frac{1}{q-1} \log F_{q,\nu}(\rho) \quad (13)$$

where

$$F_{q,\nu}(\rho) = \mathbb{E}_\nu \left[\left(\frac{\rho}{\nu} \right)^q \right] = \int_{\mathbb{R}^n} \nu(x) \frac{\rho(x)^q}{\nu(x)^q} dx = \int_{\mathbb{R}^n} \frac{\rho(x)^q}{\nu(x)^{q-1}} dx. \quad (14)$$

Rényi divergence is the relative form of Rényi entropy (56): $H_q(\rho) = \frac{1}{q-1} \log \int \rho(x)^q dx$. The case $q = 1$ is defined via limit, and recovers the KL divergence (1): $R_{1,\nu}(\rho) = \lim_{q \rightarrow 1} R_{q,\nu}(\rho) = H_\nu(\rho)$. Rényi divergence has the property that $R_{q,\nu}(\rho) \geq 0$ for all ρ , and $R_{q,\nu}(\rho) = 0$ if and only if $\rho = \nu$. Furthermore, the map $q \mapsto R_{q,\nu}(\rho)$ is increasing (see Section B.3.1). Therefore, Rényi divergence provides an alternative measure of “distance” of ρ from ν , which becomes stronger as q increases. In particular, $R_{\infty,\nu}(\rho) = \log \left\| \frac{\rho}{\nu} \right\|_\infty = \log \sup_x \frac{\rho(x)}{\nu(x)}$ is finite if and only if ρ is *warm* relative to ν . It is possible that $R_{q,\nu}(\rho) = \infty$ for large enough q , as the following example shows.

Example 2. Let $\rho = \mathcal{N}(0, \sigma^2 I)$ and $\nu = \mathcal{N}(0, \lambda^2 I)$. If $\sigma^2 > \lambda^2$ and $q \geq \frac{\sigma^2}{\sigma^2 - \lambda^2}$, then $R_{q,\nu}(\rho) = \infty$. Otherwise, $R_{q,\nu}(\rho) = \frac{n}{2} \log \frac{\lambda^2}{\sigma^2} - \frac{n}{2(q-1)} \log(q - (q-1)\frac{\sigma^2}{\lambda^2})$.

The following is analogous to Lemma 1. We provide the proof of Lemma 4 in Appendix B.3.2.

Lemma 4. Suppose $\nu = e^{-f}$ is L -smooth. Let $\rho = \mathcal{N}(x^*, \frac{1}{L}I)$ where x^* is a stationary point of f . Then for all $q \geq 1$, $R_{q,\nu}(\rho) \leq f(x^*) + \frac{n}{2} \log \frac{L}{2\pi}$.

4.1.1 Log-Sobolev inequality

For $q > 0$, we define the **Rényi information** of order q of ρ with respect to ν as

$$G_{q,\nu}(\rho) = \mathbb{E}_\nu \left[\left(\frac{\rho}{\nu} \right)^q \left\| \nabla \log \frac{\rho}{\nu} \right\|^2 \right] = \mathbb{E}_\nu \left[\left(\frac{\rho}{\nu} \right)^{q-2} \left\| \nabla \frac{\rho}{\nu} \right\|^2 \right] = \frac{4}{q^2} \mathbb{E}_\nu \left[\left\| \nabla \left(\frac{\rho}{\nu} \right)^{\frac{q}{2}} \right\|^2 \right]. \quad (15)$$

The case $q = 1$ recovers relative Fisher information (3): $G_{1,\nu}(\rho) = \mathbb{E}_\nu \left[\frac{\rho}{\nu} \left\| \nabla \log \frac{\rho}{\nu} \right\|^2 \right] = J_\nu(\rho)$. We have the following relation under log-Sobolev inequality. Note the case $q = 1$ recovers LSI (4). We provide the proof of Lemma 5 in Appendix B.3.3.

Lemma 5. Suppose ν satisfies LSI with constant $\alpha > 0$. Let $q \geq 1$. For all ρ ,

$$\frac{G_{q,\nu}(\rho)}{F_{q,\nu}(\rho)} \geq \frac{2\alpha}{q^2} R_{q,\nu}(\rho). \quad (16)$$

4.2 Langevin dynamics

Along the Langevin dynamics (7) for ν , we can compute the rate of change of the Rényi divergence.

Lemma 6. For all $q > 0$, along the Langevin dynamics (7),

$$\frac{d}{dt} R_{q,\nu}(\rho_t) = -q \frac{G_{q,\nu}(\rho_t)}{F_{q,\nu}(\rho_t)}. \quad (17)$$

We provide the proof of Lemma 6 in Appendix B.3.4. In particular, $\frac{d}{dt} R_{q,\nu}(\rho_t) \leq 0$, so Rényi divergence is always decreasing along the Langevin dynamics. Furthermore, analogous to how the Langevin dynamics is the gradient flow of KL divergence under the Wasserstein metric, one can also show that the Langevin dynamics is the the gradient flow of Rényi divergence with respect to a suitably defined metric (which depends on ν) on the space of distributions; see [13].

4.2.1 Convergence of Rényi divergence along Langevin dynamics under LSI

When ν satisfies LSI, Rényi divergence converges exponentially fast along the Langevin dynamics. Note the case $q = 1$ recovers the exponential convergence rate of KL divergence from Theorem 1.

Theorem 3. Suppose ν satisfies LSI with constant $\alpha > 0$. Let $q \geq 1$. Along the Langevin dynamics,

$$R_{q,\nu}(\rho_t) \leq e^{-\frac{2\alpha t}{q}} R_{q,\nu}(\rho_0).$$

We provide the proof of Theorem 3 in Appendix B.3.5. Theorem 3 shows that if the initial Rényi divergence is finite, then it converges exponentially fast. However, even if initially the Rényi divergence is ∞ , it will be finite along the Langevin dynamics, after which time Theorem 3 applies. This is because when ν satisfies LSI, the Langevin dynamics satisfies a *hypercontractivity* property [30, 10, 60]; see Section B.3.6. Furthermore, as shown in [13], we can combine the exponential convergence rate above with the hypercontractivity property to improve the exponential rate to be 2α , independent of q , at the cost of some initial waiting time; here we leave the rate as above for simplicity.

5 Rényi divergence along ULA

In this section we prove a convergence guarantee for Rényi divergence along ULA under the assumption that the biased limit satisfies LSI. As before, let $\nu = e^{-f}$, and let ν_ϵ denote the biased limit of ULA (11) with step size $\epsilon > 0$. We note that the bias $R_{q,\nu}(\nu_\epsilon)$ may be ∞ for large enough q .

Example 3. As in Examples [1](#) and [2](#) let $\nu = \mathcal{N}(0, \frac{1}{\alpha}I)$, so $\nu_\epsilon = \mathcal{N}(0, \frac{1}{\alpha(1-\frac{\epsilon\alpha}{2})}I)$. The bias is

$$R_{q,\nu}(\nu_\epsilon) = \begin{cases} \frac{n}{2(q-1)} (q \log(1 - \frac{\epsilon\alpha}{2}) - \log(1 - \frac{q\epsilon\alpha}{2})) & \text{if } 1 < q < \frac{2}{\epsilon\alpha}, \\ \infty & \text{if } q \geq \frac{2}{\epsilon\alpha}. \end{cases}$$

Thus, for each fixed $q > 1$, there is an asymptotic bias $R_{q,\nu}(\nu_\epsilon)$ which is finite for small enough ϵ . In Example [3](#) we have $R_{q,\nu}(\nu_\epsilon) = O(\epsilon^2)$. In general, we assume for each $q > 1$ there is a **growth function** $g_q(\epsilon)$ that controls the bias: $R_{q,\nu}(\nu_\epsilon) \leq g_q(\epsilon)$ for small $\epsilon > 0$, and $\lim_{\epsilon \rightarrow 0} g_q(\epsilon) = 0$.

5.1 Decomposition of Rényi divergence

For order $q > 1$, we have the following decomposition of Rényi divergence.

Lemma 7. Let $q > 1$. For all probability distribution ρ ,

$$R_{q,\nu}(\rho) \leq \left(\frac{q - \frac{1}{2}}{q - 1}\right) R_{2q,\nu_\epsilon}(\rho) + R_{2q-1,\nu}(\nu_\epsilon). \quad (18)$$

We provide the proof of Lemma [7](#) in Appendix [B.4.1](#). The first term in the bound above is the Rényi divergence with respect to the biased limit, which converges exponentially fast under LSI (see Lemma [8](#)). The second term in [\(18\)](#) is the bias, which is controlled by the growth function $g_{2q-1}(\epsilon)$.

5.2 Rapid convergence of Rényi divergence with respect to ν_ϵ along ULA

We show Rényi divergence with respect to the biased limit ν_ϵ converges exponentially fast along ULA, assuming ν_ϵ itself satisfies LSI.

Assumption 1. The probability distribution ν_ϵ satisfies LSI with a constant $\beta \equiv \beta_\epsilon > 0$.

We can verify Assumption [1](#) in the Gaussian case. However, it is unclear how to verify Assumption [1](#) in general. One might hope to prove that if ν satisfies LSI, then Assumption [1](#) holds.

Example 4. Let $\nu = \mathcal{N}(0, \frac{1}{\alpha}I)$, so $\nu_\epsilon = \mathcal{N}(0, \frac{1}{\alpha(1-\frac{\epsilon\alpha}{2})}I)$ satisfies LSI with $\beta = \alpha(1 - \frac{\epsilon\alpha}{2})$.

Under Assumption [1](#) we can prove an exponential convergence rate to the biased limit ν_ϵ .

Lemma 8. Assume Assumption [1](#). Suppose $\nu = e^{-f}$ is L -smooth, and let $0 < \epsilon \leq \min\{\frac{1}{3L}, \frac{1}{9\beta}\}$. For $q \geq 1$, along ULA [\(11\)](#),

$$R_{q,\nu_\epsilon}(\rho_k) \leq e^{-\frac{\beta\epsilon k}{q}} R_{q,\nu_\epsilon}(\rho_0). \quad (19)$$

We provide the proof of Lemma [8](#) in Appendix [B.4.2](#). In the proof of Lemma [8](#) we decompose each step of ULA as a sequence of two operations; see Figure [3](#) for an illustration. In the first part we take a gradient step, which is a deterministic bijective map, so it preserves Rényi divergence. In the second part we add an independent Gaussian, which is evolution along the heat flow, and we derive a formula on the decrease in Rényi divergence (which is similar to [\(17\)](#) along the Langevin dynamics).

5.3 Convergence of Rényi divergence along ULA under LSI

We combine Lemma [7](#) and Lemma [8](#) to obtain the following characterization of the convergence of Rényi divergence along ULA under LSI. We provide the proof of Theorem [4](#) in Appendix [B.4.3](#).

Theorem 4. Assume Assumption [1](#). Suppose $\nu = e^{-f}$ is L -smooth, and let $0 < \epsilon \leq \min\{\frac{1}{3L}, \frac{1}{9\beta}\}$. Let $q > 1$, and suppose $R_{2q,\nu_\epsilon}(\rho_0) < \infty$. Then along ULA [\(11\)](#),

$$R_{q,\nu}(\rho_k) \leq \left(\frac{q - \frac{1}{2}}{q - 1}\right) R_{2q,\nu_\epsilon}(\rho_0) e^{-\frac{\beta\epsilon k}{2q}} + g_{2q-1}(\epsilon). \quad (20)$$

For $\delta > 0$, let $g_q^{-1}(\delta) = \sup\{\epsilon > 0: g_q(\epsilon) \leq \delta\}$. Theorem [4](#) states that to achieve $R_{q,\nu}(\rho_k) \leq \delta$, it suffices to run ULA with step size $\epsilon = \Theta(\min\{\frac{1}{L}, g_{2q-1}^{-1}(\frac{\delta}{2})\})$ for $k = O(\frac{1}{\beta\epsilon} \log \frac{R_{2q,\nu_\epsilon}(\rho_0)}{\delta})$ iterations. Suppose δ is small so $g_{2q-1}^{-1}(\frac{\delta}{2}) < \frac{1}{L}$. Note ν_ϵ is $\frac{1}{2\epsilon}$ -smooth, so if we choose ρ_0

to be a Gaussian with covariance $2\epsilon I$, we have $R_{2q, \nu_\epsilon}(\rho_0) = \tilde{O}(n)$ by Lemma 4. Therefore, Theorem 4 yields an iteration complexity of $k = \tilde{O}\left(\frac{1}{\beta g_{2q-1}^{-1}(\delta/2)}\right)$. For example, if $g_q(\epsilon) = O(\epsilon)$, then $g_q^{-1}(\delta) = \Omega(\delta)$, so the iteration complexity is $k = \tilde{O}\left(\frac{1}{\beta\delta}\right)$ with $\epsilon = \Theta(\delta)$. If $g_q(\epsilon) = O(\epsilon^2)$, as in Example 3, then $g_q^{-1}(\delta) = \Omega(\sqrt{\delta})$, so the iteration complexity is $k = \tilde{O}\left(\frac{1}{\beta\sqrt{\delta}}\right)$ with $\epsilon = \Theta(\sqrt{\delta})$.

6 Poincaré inequality

We recall ν satisfies **Poincaré inequality (PI)** with a constant $\alpha > 0$ if for all smooth $g: \mathbb{R}^n \rightarrow \mathbb{R}$,

$$\text{Var}_\nu(g) \leq \frac{1}{\alpha} \mathbb{E}_\nu[\|\nabla g\|^2] \quad (21)$$

where $\text{Var}_\nu(g) = \mathbb{E}_\nu[g^2] - \mathbb{E}_\nu[g]^2$ is the variance of g under ν . Poincaré inequality is an isoperimetry condition which is weaker than LSI. LSI implies PI with the same constant; in fact, PI is a linearization of LSI (4), i.e., when $\rho = (1+\eta g)\nu$ as $\eta \rightarrow 0$ [57, 60]. Furthermore, it is known Talagrand's inequality implies PI with the same constant, and PI is also a linearization of Talagrand's inequality [54]. Poincaré inequality is better behaved than LSI [15], and there are various Lyapunov criteria and integrability conditions to verify when a distribution satisfies Poincaré inequality [5, 51, 19].

6.1 Convergence of Rényi divergence along Langevin dynamics under Poincaré inequality

When ν satisfies Poincaré inequality, Rényi divergence converges along the Langevin dynamics. The convergence is initially linear, then becomes exponential once Rényi divergence falls below 1.

Theorem 5. *Suppose ν satisfies Poincaré inequality with constant $\alpha > 0$. Let $q \geq 2$. Along the Langevin dynamics,*

$$R_{q, \nu}(\rho_t) \leq \begin{cases} R_{q, \nu}(\rho_0) - \frac{2\alpha t}{q} & \text{if } R_{q, \nu}(\rho_0) \geq 1 \text{ and as long as } R_{q, \nu}(\rho_t) \geq 1, \\ e^{-\frac{2\alpha t}{q}} R_{q, \nu}(\rho_0) & \text{if } R_{q, \nu}(\rho_0) \leq 1. \end{cases}$$

We provide the proof of Theorem 5 in Appendix B.5.2. Theorem 5 states that starting from $R_{q, \nu}(\rho_0) \geq 1$, the Langevin dynamics reaches $R_{q, \nu}(\rho_t) \leq \delta$ in $t \leq O\left(\frac{q}{\alpha} (R_{q, \nu}(\rho_0) + \log \frac{1}{\delta})\right)$ time.

6.2 Rapid convergence of Rényi divergence with respect to ν_ϵ along ULA

We assume the biased limit ν_ϵ satisfies Poincaré inequality.

Assumption 2. *The distribution ν_ϵ satisfies Poincaré inequality with a constant $\beta \equiv \beta_\epsilon > 0$.*

Under Assumption 2 we can show Rényi divergence with respect to ν_ϵ converges at a rate similar to the Langevin dynamics; see Lemma 18 in Appendix B.5.3.

6.3 Convergence of Rényi divergence along ULA under Poincaré inequality

We combine Lemma 7 and Lemma 18 to obtain the following convergence of Rényi divergence along ULA under Poincaré inequality. We provide the proof of Theorem 6 in Appendix B.5.4.

Theorem 6. *Assume Assumption 2. Suppose $\nu = e^{-f}$ is L -smooth, and let $0 < \epsilon \leq \min\left\{\frac{1}{3L}, \frac{1}{9\beta}\right\}$. Let $q > 1$ and assume $1 \leq R_{2q, \nu_\epsilon}(\rho_0) < \infty$. Along ULA (11), for $k \geq k_0 := \frac{2q}{\beta_\epsilon} (R_{2q, \nu_\epsilon}(\rho_0) - 1)$,*

$$R_{q, \nu}(\rho_k) \leq \left(\frac{q - \frac{1}{2}}{q - 1}\right) e^{-\frac{\beta_\epsilon(k-k_0)}{2q}} + g_{2q-1}(\epsilon). \quad (22)$$

This yields an iteration complexity for ULA under Poincaré which is a factor of n larger than the complexity under LSI; see Appendix B.5.5.

7 Discussion

In this paper we proved convergence guarantees on KL and Rényi divergence along ULA under isoperimetry and bounded Hessian, without assuming convexity or bounds on higher derivatives. It would be interesting to verify when Assumptions 1 and 2 hold or whether they follow from isoperimetry and bounded Hessian of the target density. Another intriguing question is whether there is an affine-invariant version of the Langevin dynamics. This might lead to a sampling algorithm with logarithmic dependence on smoothness parameters, rather than the current polynomial dependence.

Acknowledgment

The first author was supported in part by NSF awards CCF-1563838 and CCF-1717349. The authors would like to thank Kunal Talwar for explaining the application of Rényi divergence to data privacy. The authors thank Yu Cao, Jianfeng Lu, and Yulong Lu for alerting us to their work [13] on Rényi divergence. The authors also thank Xiang Cheng and Peter Bartlett for helpful comments on an earlier version of this paper.

References

- [1] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 308–318. ACM, 2016.
- [2] David Applegate and Ravi Kannan. Sampling and integration of near log-concave functions. In *Proceedings of the Twenty-third Annual ACM Symposium on Theory of Computing*, STOC '91, pages 156–163, New York, NY, USA, 1991. ACM.
- [3] John C Baez. Rényi entropy and free energy. *arXiv preprint arXiv:1102.2098*, 2011.
- [4] Shi Bai, Tancrède Lepoint, Adeline Roux-Langlois, Amin Sakzad, Damien Stehlé, and Ron Steinfeld. Improved security proofs in lattice-based cryptography: using the rényi divergence rather than the statistical distance. *Journal of Cryptology*, 31(2):610–640, 2018.
- [5] Dominique Bakry, Franck Barthe, Patrick Cattiaux, Arnaud Guillin, et al. A simple proof of the Poincaré inequality for a large class of probability measures. *Electronic Communications in Probability*, 13:60–66, 2008.
- [6] Dominique Bakry and Michel Émery. Diffusions hypercontractives. In *Séminaire de Probabilités XIX 1983/84*, pages 177–206. Springer, 1985.
- [7] Jean-Baptiste Bardet, Nathaël Gozlan, Florent Malrieu, and Pierre-André Zitt. Functional inequalities for Gaussian convolutions of compactly supported measures: Explicit bounds and dimension dependence. *Bernoulli*, 24(1):333–353, 2018.
- [8] Espen Bernton. Langevin Monte Carlo and JKO splitting. In *Conference On Learning Theory, COLT 2018, Stockholm, Sweden, 6-9 July 2018*, pages 1777–1798, 2018.
- [9] Sergej G Bobkov and Friedrich Götze. Exponential integrability and transportation cost related to logarithmic Sobolev inequalities. *Journal of Functional Analysis*, 163(1):1–28, 1999.
- [10] Sergej G Bobkov, Ivan Gentil, and Michel Ledoux. Hypercontractivity of Hamilton–Jacobi equations. *Journal de Mathématiques Pures et Appliquées*, 80(7):669–696, 2001.
- [11] SG Bobkov, GP Chistyakov, and Friedrich Götze. Rényi divergence and the central limit theorem. *The Annals of Probability*, 47(1):270–323, 2019.
- [12] Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*, pages 635–658. Springer, 2016.
- [13] Yu Cao, Jianfeng Lu, and Yulong Lu. Exponential decay of Rényi divergence under Fokker–Planck equations. *Journal of Statistical Physics*, pages 1–13, 2018.
- [14] Djalil Chafaï. Entropies, convexity, and functional inequalities: On ϕ -entropies and ϕ -Sobolev inequalities. *Journal of Mathematics of Kyoto University*, 44(2):325–363, 2004.
- [15] Djalil Chafaï and Florent Malrieu. On fine properties of mixtures with respect to concentration of measure and Sobolev type inequalities. In *Annales de l’IHP Probabilités et statistiques*, volume 46, pages 72–96, 2010.
- [16] Zongchen Chen and Santosh S. Vempala. Optimal convergence rate of Hamiltonian Monte Carlo for strongly logconcave distributions. *arXiv preprint arXiv:1905.02313*, 2019.
- [17] Xiang Cheng and Peter Bartlett. Convergence of Langevin MCMC in KL-divergence. In Firdaus Janoos, Mehryar Mohri, and Karthik Sridharan, editors, *Proceedings of Algorithmic Learning Theory*, volume 83 of *Proceedings of Machine Learning Research*, pages 186–211. PMLR, 07–09 Apr 2018.

- [18] Xiang Cheng, Niladri S Chatterji, Yasin Abbasi-Yadkori, Peter L Bartlett, and Michael I Jordan. Sharp convergence rates for Langevin dynamics in the nonconvex setting. *arXiv preprint arXiv:1805.01648*, 2018.
- [19] Thomas A Courtade. Bounds on the Poincaré constant for convolution measures. *arXiv preprint arXiv:1807.00027*, 2018.
- [20] Imre Csiszár. Generalized cutoff rates and Rényi’s information measures. *IEEE Transactions on information theory*, 41(1):26–34, 1995.
- [21] Arnak Dalalyan. Further and stronger analogy between sampling and optimization: Langevin Monte Carlo and gradient descent. In *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pages 678–689. PMLR, 07–10 Jul 2017.
- [22] Arnak S Dalalyan and Avetik Karagulyan. User-friendly guarantees for the langevin monte carlo with inaccurate gradient. *Stochastic Processes and their Applications*, 2019.
- [23] Xi Dong. The gravity dual of Rényi entropy. *Nature communications*, 7:12472, 2016.
- [24] Alain Durmus, Szymon Majewski, and Błażej Miasojedow. Analysis of Langevin Monte Carlo via convex optimization. *arXiv preprint arXiv:1802.09188*, 2018.
- [25] Alain Durmus, Eric Moulines, and Eero Saksman. On the convergence of Hamiltonian Monte Carlo. *arXiv preprint arXiv:1705.00166*, 2017.
- [26] Raaz Dwivedi, Yuansi Chen, Martin J. Wainwright, and Bin Yu. Log-concave sampling: Metropolis-Hastings algorithms are fast! In *Conference On Learning Theory, COLT 2018, Stockholm, Sweden, 6-9 July 2018*, pages 793–797, 2018.
- [27] Cynthia Dwork and Guy N Rothblum. Concentrated differential privacy. *arXiv preprint arXiv:1603.01887*, 2016.
- [28] Andreas Eberle, Arnaud Guillin, and Raphael Zimmer. Couplings and quantitative contraction rates for Langevin dynamics. *The Annals of Probability*, 47(4):1982–2010, 2019.
- [29] Jackson Gorham and Lester Mackey. Measuring sample quality with kernels. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1292–1301, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.
- [30] Leonard Gross. Logarithmic Sobolev inequalities. *American Journal of Mathematics*, 97(4):1061–1083, 1975.
- [31] Peter Harremoës. Interpretations of Rényi entropies and divergences. *Physica A: Statistical Mechanics and its Applications*, 365(1):57–62, 2006.
- [32] Yun He, A Ben Hamza, and Hamid Krim. A generalized divergence measure for robust image registration. *IEEE Transactions on Signal Processing*, 51(5):1211–1220, 2003.
- [33] Richard Holley and Daniel Stroock. Logarithmic Sobolev inequalities and stochastic Ising models. *Journal of Statistical Physics*, 46(5):1159–1194, 1987.
- [34] Richard Holley and Daniel Stroock. Simulated annealing via Sobolev inequalities. *Communications in Mathematical Physics*, 115(4):553–569, 1988.
- [35] Mitsugu Iwamoto and Junji Shikata. Information theoretic security for encryption based on conditional Rényi entropies. In *International Conference on Information Theoretic Security*, pages 103–121. Springer, 2013.
- [36] Richard Jordan, David Kinderlehrer, and Felix Otto. The variational formulation of the Fokker–Planck equation. *SIAM Journal on Mathematical Analysis*, 29(1):1–17, January 1998.
- [37] R. Kannan, L. Lovász, and M. Simonovits. Random walks and an $O^*(n^5)$ volume algorithm for convex bodies. *Random Structures and Algorithms*, 11:1–50, 1997.
- [38] Michel Ledoux. Concentration of measure and logarithmic Sobolev inequalities. *Séminaire de probabilités de Strasbourg*, 33:120–216, 1999.

- [39] Yin Tat Lee and Santosh S Vempala. Convergence rate of Riemannian Hamiltonian Monte Carlo and faster polytope volume computation. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1115–1121. ACM, 2018.
- [40] Xuechen Li, Denny Wu, Lester Mackey, and Murat A Erdogdu. Stochastic Runge–Kutta accelerates Langevin Monte Carlo and beyond. *arXiv preprint arXiv:1906.07868*, 2019.
- [41] Yingzhen Li and Richard E Turner. Rényi divergence variational inference. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 1073–1081. Curran Associates, Inc., 2016.
- [42] L. Lovász and S. Vempala. Fast algorithms for logconcave functions: sampling, rounding, integration and optimization. In *FOCS*, pages 57–68, 2006.
- [43] L. Lovász and S. Vempala. The geometry of logconcave functions and sampling algorithms. *Random Struct. Algorithms*, 30(3):307–358, 2007.
- [44] László Lovász and Santosh S. Vempala. Hit-and-run from a corner. *SIAM Journal on Computing*, 35(4):985–1005, 2006.
- [45] Yi-An Ma, Niladri Chatterji, Xiang Cheng, Nicolas Flammarion, Peter Bartlett, and Michael I Jordan. Is there an analog of Nesterov acceleration for MCMC? *arXiv preprint arXiv:1902.00996*, 2019.
- [46] Michael C. Mackey. *Time’s Arrow: The Origins of Thermodynamics Behavior*. Springer-Verlag, 1992.
- [47] Oren Mangoubi and Aaron Smith. Rapid mixing of Hamiltonian Monte Carlo on strongly log-concave distributions. *arXiv preprint arXiv:1708.07114*, 2017.
- [48] Oren Mangoubi and Nisheeth Vishnoi. Dimensionally tight bounds for second-order Hamiltonian Monte Carlo. In *Advances in Neural Information Processing Systems 31*, pages 6027–6037. Curran Associates, Inc., 2018.
- [49] Oren Mangoubi and Nisheeth K Vishnoi. Nonconvex sampling with the Metropolis-adjusted Langevin algorithm. *arXiv preprint arXiv:1902.08452*, 2019.
- [50] Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. Multiple source adaptation and the Rényi divergence. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pages 367–374. AUAI Press, 2009.
- [51] Georg Menz and André Schlichting. Poincaré and logarithmic Sobolev inequalities by decomposition of the energy landscape. *Ann. Probab.*, 42(5):1809–1884, 09 2014.
- [52] Ilya Mironov. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 263–275. IEEE, 2017.
- [53] D Morales, L Pardo, and I Vajda. Rényi statistics in directed families of exponential experiments. *Statistics: A Journal of Theoretical and Applied Statistics*, 34(2):151–174, 2000.
- [54] Felix Otto and Cédric Villani. Generalization of an inequality by Talagrand and links with the logarithmic Sobolev inequality. *Journal of Functional Analysis*, 173(2):361–400, 2000.
- [55] Maxim Raginsky, Alexander Rakhlin, and Matus Telgarsky. Non-convex learning via stochastic gradient Langevin dynamics: A nonasymptotic analysis. In Satyen Kale and Ohad Shamir, editors, *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pages 1674–1703, Amsterdam, Netherlands, 07–10 Jul 2017. PMLR.
- [56] Alfréd Rényi et al. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*. The Regents of the University of California, 1961.
- [57] OS Rothaus. Diffusion on compact Riemannian manifolds and logarithmic Sobolev inequalities. *Journal of functional analysis*, 42(1):102–109, 1981.
- [58] M Talagrand. Transportation cost for Gaussian and other product measures. *Geometric and Functional Analysis*, 6:587–600, 01 1996.
- [59] Tim Van Erven and Peter Harremoos. Rényi divergence and Kullback-Leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014.

- [60] Cédric Villani. *Topics in optimal transportation*. Number 58 in Graduate Studies in Mathematics. American Mathematical Society, 2003.
- [61] Feng-Yu Wang and Jian Wang. Functional inequalities for convolution of probability measures. In *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, volume 52, pages 898–914. Institut Henri Poincaré, 2016.
- [62] Andre Wibisono. Sampling as optimization in the space of measures: The Langevin dynamics as a composite optimization problem. In *Conference On Learning Theory, COLT 2018, Stockholm, Sweden, 6-9 July 2018*, pages 2093–3027, 2018.