

Fast Convergence of Φ -Divergence Along the Unadjusted Langevin Algorithm and Proximal Sampler

Siddharth Mitra

Andre Wibisono

Department of Computer Science, Yale University, New Haven, CT, USA

SIDDHARTH.MITRA@YALE.EDU

ANDRE.WIBISONO@YALE.EDU

Editors: Gautam Kamath and Po-Ling Loh

Abstract

We study the mixing time of two popular discrete-time Markov chains in continuous space, the Unadjusted Langevin Algorithm and the Proximal Sampler, which are discretizations of the Langevin dynamics. We extend mixing time analyses for these Markov chains to hold in Φ -divergence. We show that any Φ -divergence arising from a twice-differentiable strictly convex function Φ converges to 0 exponentially fast along these Markov chains, under the assumption that their stationary distributions satisfy the corresponding Φ -Sobolev inequality, which holds for example when the target distribution of the Langevin dynamics is strongly log-concave. Our setting includes as special cases popular mixing time regimes, namely the mixing in chi-squared divergence under a Poincaré inequality, and the mixing in relative entropy under a log-Sobolev inequality. Our results follow by viewing the sampling algorithms as noisy channels and bounding the contraction coefficients arising in the appropriate strong data processing inequalities.

Keywords: Langevin dynamics, unadjusted Langevin algorithm, proximal sampler, Φ -divergence, Φ -Sobolev inequalities, strong data processing inequalities

1. Introduction

Sampling from a probability distribution is a fundamental task that appears in many fields, including machine learning, statistics, and Bayesian inference (Gelman et al., 1995; MacKay, 2003; Robert et al., 1999; Von Toussaint, 2011; Johannes and Polson, 2003). Suppose we wish to obtain samples from a continuous probability distribution $\nu \propto \exp(-f)$ on \mathbb{R}^d ; a common approach is to construct a Markov chain which admits ν as its invariant or stationary distribution, and then draw samples after some initial burn-in time. A rigorous study of burn-in time (Gilks et al., 1995; Geyer, 2011) leads one to analyze the *mixing time* of a Markov chain, which tracks how quickly the Markov chain converges to its stationary distribution (Levin et al., 2017). Different choices of distances¹ between probability distributions lead to different guarantees, which can be bounded in terms of each other, to yield interesting results and bounds.

A general, useful, and well-studied family of divergences are that of Φ -divergences (Csiszár, 1967)², which include many popular divergences such as Kullback-Leibler (KL) divergence, chi-squared divergence, Total Variation (TV) distance, and squared-Hellinger distance. For any convex function $\Phi : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ with $\Phi(1) = 0$, the Φ -divergence between probability distributions μ and ν such that $\mu \ll \nu$ is defined by

$$D_{\Phi}(\mu \parallel \nu) := \mathbb{E}_{\nu} \left[\Phi \left(\frac{\mu}{\nu} \right) \right], \quad (1)$$

1. We use “distance” loosely here to refer to metrics such as Wasserstein metric, or divergences such as KL divergence.

2. Also commonly known as f -divergences in the literature.

and it is $+\infty$ if $\mu \not\ll \nu$. For example, the KL divergence or relative entropy corresponds to $\Phi(x) = x \log x$, chi-squared divergence corresponds to $\Phi(x) = (x - 1)^2$, TV distance corresponds to $\Phi(x) = \frac{1}{2}|x - 1|$, and squared-Hellinger distance corresponds to $\Phi(x) = \frac{1}{2}(\sqrt{x} - 1)^2$. In general, Φ -divergences do not satisfy a triangle inequality, but some examples which do are the TV distance and Marton's divergence (Sason and Verdú, 2016, (72)). Further examples of Φ -divergences can be found in Table 1 in Appendix A. The family of Φ -divergences have found profound applications from hypothesis and distribution testing (Pensia et al., 2024; Györfi and Vajda, 2002; Gretton and Györfi, 2008), and neuroscience (Nemenman et al., 2004; Belitski et al., 2008), to reinforcement learning (Ho et al., 2022; Panaganti et al., 2024). We study Φ -divergences in the context of mixing time for Markov chains, and extend the analysis from KL divergence³ to Φ -divergence.

We study the mixing time of Markov chains in Φ -divergence under the assumption that their stationary distribution satisfy a Φ -Sobolev inequality. These inequalities include as special cases popular isoperimetric inequalities such as the log-Sobolev inequality (LSI) and Poincaré inequality, and are defined as follows.

Definition 1 *A probability distribution ν satisfies a Φ -Sobolev inequality (Φ SI) with constant $\alpha > 0$ if for all probability distributions $\mu \ll \nu$, we have*

$$2\alpha D_{\Phi}(\mu \parallel \nu) \leq \text{FI}_{\Phi}(\mu \parallel \nu), \quad (2)$$

where $D_{\Phi}(\mu \parallel \nu)$ is the Φ -divergence defined in (1), and $\text{FI}_{\Phi}(\mu \parallel \nu)$ is the Φ -Fisher information defined by

$$\text{FI}_{\Phi}(\mu \parallel \nu) := \mathbb{E}_{\nu} \left[\left\| \nabla \frac{\mu}{\nu} \right\|^2 \Phi'' \left(\frac{\mu}{\nu} \right) \right].$$

We define the Φ -Sobolev constant of ν to be the optimal (largest) constant α such that the above inequality holds:

$$\alpha_{\Phi\text{SI}}(\nu) := \inf_{\mu} \frac{\text{FI}_{\Phi}(\mu \parallel \nu)}{2D_{\Phi}(\mu \parallel \nu)} \quad (3)$$

where the infimum is taken over all probability distributions μ with $0 < D_{\Phi}(\mu \parallel \nu) < \infty$.

The inequality (2) is equivalent to saying that for all smooth functions $g : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ with $\mathbb{E}_{\nu}[g] = 1$,

$$2\alpha \text{Ent}_{\Phi}^{\nu}(g) \leq \mathcal{E}_{\Phi}^{\nu}(g), \quad (4)$$

where

$$\text{Ent}_{\Phi}^{\nu}(g) := \mathbb{E}_{\nu}[\Phi(g)] - \Phi(\mathbb{E}_{\nu}[g]) \quad \text{and} \quad \mathcal{E}_{\Phi}^{\nu}(g) := \mathbb{E}_{\nu} [\|\nabla g\|^2 \Phi''(g)].$$

To see this, take g to be the density function of μ with respect to ν .

For $\Phi(x) = x \log x$, the Φ -divergence is the KL divergence and the Φ -Sobolev inequality is the log-Sobolev inequality. For $\Phi(x) = (x - 1)^2$, the Φ -divergence is the chi-squared divergence and the Φ -Sobolev inequality is the Poincaré inequality. Further examples are mentioned in Table 1 in Appendix A. The Poincaré inequality is the weakest Φ -Sobolev inequality in that it is implied by any other Φ -Sobolev inequality (Chafai, 2004, Section 2.2). This extends the more well-known

3. And other specific analyses such as in chi-squared divergence.

statement that the log-Sobolev inequality implies the Poincaré inequality. Additionally, one can start from a Φ -Sobolev inequality and show that it satisfies a Φ -Sobolev inequality for $\Phi(x) = x^p - 1$ for $p \in (1, 2]$ but without a tight constant (Chafaï, 2004, Section 2.2). This class of Φ -Sobolev inequalities are related to Beckner inequalities (Bakry et al., 2014, Section 7.6.2).

We discuss how the Φ -Sobolev constant of a distribution evolves along various operations such as convolutions and pushforwards in Section 2.1. These properties will be crucial when analyzing mixing time. We now introduce the Markov chains we will study.

1.1. Langevin dynamics

The Langevin dynamics to sample from $\nu \propto \exp(-f)$ on \mathbb{R}^d is the following stochastic differential equation (SDE):

$$dX_t = -\nabla f(X_t) dt + \sqrt{2} dW_t, \quad (5)$$

where W_t is the standard Brownian motion on \mathbb{R}^d . The Langevin dynamics admits ν as the stationary or invariant distribution, and hence is a natural process to study for sampling (Bakry et al., 2014). It also has the natural optimization interpretation as the gradient flow to minimize KL divergence in the space of probability distributions over \mathbb{R}^d with the Wasserstein W_2 metric, see (Jordan et al., 1998; Wibisono, 2018).

The Langevin dynamics is a continuous-time Markov process, and it needs to be discretized in time in order to implement in practice. We will focus on two discretizations of the Langevin dynamics – the Unadjusted Langevin Algorithm (ULA), which has been well-studied in (Roberts and Tweedie, 1996; Roberts and Rosenthal, 1998; Dalalyan, 2017; Cheng and Bartlett, 2018; Vempala and Wibisono, 2019; Chewi et al., 2022a), and the Proximal Sampler, which has been studied in (Lee et al., 2021; Chen et al., 2022; Yuan et al., 2023; Kook et al., 2024; Kook and Zhang, 2025).

The convergence of Φ -divergence along Langevin dynamics is easy to establish (Corollary 9) and the properties of Φ -divergence along continuous-time dynamics are well-studied (Chafaï, 2004; Dolbeault and Li, 2018; Achleitner et al., 2015). We study the convergence of Φ -divergence for discrete-time samplers in Theorems 2 and 3.

1.2. Unadjusted Langevin algorithm

The Unadjusted Langevin Algorithm (ULA) is a simple discretization of the Langevin dynamics (5) and is given by the following update

$$X_{k+1} = X_k - \eta \nabla f(X_k) + \sqrt{2\eta} Z_k, \quad (6)$$

where $Z_k \sim \mathcal{N}(0, I)$ and $\eta > 0$ is the stepsize. It is well-known (Roberts and Tweedie, 1996) that the ULA is a biased discretization, which means that it admits as its stationary distribution $\nu^\eta \neq \nu$ for all $\eta > 0$; furthermore, as $\eta k \rightarrow t$ and $\eta \rightarrow 0$, $\nu^\eta \rightarrow \nu$ and ULA (6) recovers the Langevin dynamics (5). The ULA has been widely studied, and we now have guarantees on its mixing time in numerous settings: Dalalyan (2017) studies the mixing time to ν in Wasserstein W_2 distance under ν being strongly log-concave; Cheng and Bartlett (2018) study the mixing time to ν in KL divergence under ν being strongly log-concave; Vempala and Wibisono (2019) study mixing to ν in KL divergence under an LSI assumption on ν , and also study mixing to the biased limit ν^η in Rényi divergence under an LSI assumption on ν^η ; Altschuler and Talwar (2023) study mixing to ν^η under a strong log-concavity assumption on ν ; and Chewi et al. (2022a) study mixing to ν under modified

LSI assumptions on ν . We refer the reader to [Chewi \(2024\)](#) for a comprehensive overview of recent mixing time results.

A probability distribution $\nu \propto \exp(-f)$ is said to be $L > 0$ smooth if $-LI \preceq \nabla^2 f \preceq LI$, and all of the aforementioned works make smoothness assumptions on the target distribution ν . This is common for discrete-time analysis. In [Theorem 2](#), we show mixing guarantees for ULA in Φ -divergence for smooth target distributions ν and under the assumption that the stationary distribution ν^η of ULA satisfies the corresponding Φ -Sobolev inequality ([Definition 1](#)). To the best of our knowledge, we are the first to study mixing time guarantees in any Φ -divergence for ULA. We prove [Theorem 2](#) in [Section 3.2](#).

Theorem 2 *Suppose the stationary distribution ν^η of ULA satisfies a Φ -Sobolev inequality with optimal constant $\alpha > 0$, and ν is L -smooth for some $0 < \alpha \leq L < \infty$. Let $X_k \sim \rho_k$ evolve following ULA (6) with step size $0 < \eta \leq 1/L$ from $X_0 \sim \rho_0$. Then for all $k \geq 1$,*

$$D_\Phi(\rho_k \parallel \nu^\eta) \leq \left(1 + \frac{2\alpha\eta}{(1+\eta L)^2}\right)^{-k} D_\Phi(\rho_0 \parallel \nu^\eta). \quad (7)$$

As $\eta k \rightarrow t$ and $\eta \rightarrow 0$, [Theorem 2](#) recovers the correct $\exp(-2\alpha t)$ convergence rate for Langevin dynamics under a Φ -Sobolev inequality (see [Corollary 9](#) in [Section 2.1](#)). We show that the rate of [Theorem 2](#) is tight for KL divergence via an explicit calculation for the case when ν is a Gaussian, see [Proposition 17](#) in [Appendix G.1](#).

A common setting in which ν^η satisfies a Φ -Sobolev inequality is when ν is strongly log-concave (see [Lemma 13](#) in [Section 3.3](#)). To the best of our knowledge, properties of the limiting distribution ν^η under assumptions on ν which are weaker than strong log-concavity remain unknown. Similar to how log-Sobolev and Poincaré inequalities are stable to bounded perturbations ([Holley and Stroock, 1987](#)), Φ -Sobolev inequalities also enjoy similar properties ([Chafaï, 2004](#), [Section 3.3](#)). Therefore, an interesting question one can ask is if the stationary distribution ν^η continues to satisfy a Φ -Sobolev inequality when ν undergoes a suitable perturbation; we leave this for future work.

[Theorem 2](#) shows convergence to the stationary distribution ν^η . The bias $D_\Phi(\nu^\eta \parallel \nu)$ between the stationary distribution and the target distribution depends on the choice of Φ . For example, for the Ornstein-Uhlenbeck process where $\nu = \mathcal{N}(0, \frac{1}{\alpha}I)$, we have $\nu^\eta = \mathcal{N}(0, \frac{2}{\alpha(2-\eta\alpha)}I)$ (see [Appendix G.1](#)). In this case, the bias has linear dependence in the dimension d for KL divergence ($\Phi(x) = x \log x$) and exponential dependence in d for chi-squared divergence ($\Phi(x) = (x - 1)^2$).

1.3. Proximal Sampler

While the simplicity of ULA is appealing, its biased limiting distribution makes it less ideal for applications requiring sampling with high accuracy. The Proximal Sampler ([Lee et al., 2021](#); [Chen et al., 2022](#)) is an alternative discretization that addresses this shortcoming. Given the close connections between optimization and sampling, along with the modern interpretation of sampling as optimization in the space of distributions ([Jordan et al., 1998](#); [Wibisono, 2018](#)), the Proximal Sampler as introduced in [Lee et al. \(2021\)](#) can be seen to be a sampling analogue of the proximal point algorithm in optimization ([Martinet, 1970](#); [Rockafellar, 1976](#)). The Proximal Sampler was analyzed in great generality in [Chen et al. \(2022\)](#), who show mixing guarantees in KL divergence, Rényi divergence, chi-square divergence, and W_2 distance, under isoperimetry or log-concavity assumptions on ν . In [Theorem 3](#), we show mixing guarantees for the Proximal Sampler in Φ -divergence under

the assumption that ν satisfies the corresponding Φ -Sobolev inequality (Definition 1). Our results extend the mixing time of the Proximal Sampler to Φ -divergence.

The Proximal Sampler considers an augmented (X, Y) space $\mathbb{R}^d \times \mathbb{R}^d$ and performs Gibbs sampling on the joint space. We will use the appropriate superscripts to denote probability distributions on their respective spaces.⁴ Let the target distribution we wish to sample from be $\nu^X \propto \exp(-f)$ on \mathbb{R}^d . The joint target distribution on $\mathbb{R}^d \times \mathbb{R}^d$ is defined as:

$$\nu^{XY}(x, y) \propto \exp\left(-f(x) - \frac{\|x - y\|^2}{2\eta}\right), \quad (8)$$

for step size $\eta > 0$. The Proximal Sampler, initialized from $X_0 \sim \rho_0^X$, is the following two-step algorithm:

Step 1 (forward step): Sample $Y_k \mid X_k \sim \nu^{Y|X=X_k} = \mathcal{N}(X_k, \eta I)$ (9)

Step 2 (backward step): Sample $X_{k+1} \mid Y_k \sim \nu^{X|Y=Y_k}$

Note that ν^{XY} has the desired target distribution ν^X as the X marginal, and that the Y marginal ν^Y is a smoothed version of ν^X , i.e., $\nu^Y = \nu^X * \mathcal{N}(0, \eta I)$. The forward step of the algorithm is easy to implement as it corresponds to drawing a Gaussian random variable. For the backward step, implementation is possible given access to a *Restricted Gaussian Oracle* (RGO). A RGO is an oracle that, given any $y \in \mathbb{R}^d$, outputs a sample from $\nu^{X|Y=y}$, i.e. from

$$\nu^{X|Y}(x \mid y) \propto_x \exp\left(-f(x) - \frac{\|x - y\|^2}{2\eta}\right).$$

Similar to Lee et al. (2021); Chen et al. (2022), in our main result for the Proximal Sampler (Theorem 3), we consider an ideal implementation of the sampler where we have exact access to a RGO. Specific cases where the RGO can be implemented efficiently are discussed in Lee et al. (2021), and improved implementations of the RGO along with analysis of the Proximal Sampler with inexact RGO implementations have been an active area of research (Fan et al., 2023; Liang and Chen, 2022; Altschuler and Chewi, 2024). Fan et al. (2023) show an implementation via approximate rejection sampling and Altschuler and Chewi (2024) show an approach based on the Metropolis-adjusted Langevin algorithm.

We mention a basic approach to implement the RGO via rejection sampling in Appendix B. We also mention the oracle complexity (in this case, the expected number of calls to the first order oracle of f) when sampling using the Proximal Sampler with the rejection sampling based RGO implementation in Corollary 4.

When both steps are implemented exactly (i.e. given access to an exact RGO), the Proximal Sampler corresponds to Gibbs sampling from the stationary distribution ν^{XY} , and the algorithm is therefore unbiased. Step 1 is referred to as the forward step as it corresponds to evolving along the (forward) heat flow, and step 2 is referred to as the backward step as it corresponds to backward heat flow; this perspective is further made clear in Sections 4.1 and 4.2. We now state the main theorem describing the mixing time of the Proximal Sampler in Φ -divergence. We prove Theorem 3 in Section 4.3.

4. For example, ρ^{XY} denotes a distribution on the joint space, ρ^X refers to a distribution on the X space, and $\rho^{Y|X=x}$ denotes a conditional distribution supported on the Y space.

Theorem 3 *Suppose ν^X satisfies a Φ -Sobolev inequality with optimal constant α . Let $X_k \sim \rho_k^X$ evolve along the Proximal Sampler (9) with step size $\eta > 0$ from $X_0 \sim \rho_0^X$. Then for all $k \geq 1$,*

$$D_{\Phi}(\rho_k^X \parallel \nu^X) \leq \frac{D_{\Phi}(\rho_0^X \parallel \nu^X)}{(1 + \alpha\eta)^{2k}}.$$

We show the tightness of Theorem 3 in KL divergence by doing an explicit calculation for the case when ν is Gaussian, see Proposition 18 in Appendix G.2. The setting of Theorem 3 includes popular mixing time results as special cases. For example, it directly implies the mixing in KL divergence under a log-Sobolev inequality, and mixing in chi-squared divergence under a Poincaré inequality. Under a strong log-concavity assumption on ν , it implies mixing in all Φ -divergences. Φ -Sobolev inequalities and Φ -divergences for different choices of Φ are mentioned in Table 1 in Appendix A.

Theorem 3 considers the ideal implementation of the Proximal Sampler where we assume exact access to the RGO. Combined with an RGO implementation via rejection sampling, which requires an additional smoothness assumption on f , Theorem 3 yields the following corollary.

Corollary 4 *Suppose ν^X satisfies a Φ -Sobolev inequality with optimal constant α and is L -smooth. Then for any $\epsilon > 0$, the Proximal Sampler with $\eta \asymp \frac{1}{Ld}$ and with rejection sampling based RGO implementation (as described in Appendix B) outputs $X_k \sim \rho_k^X$ with $D_{\Phi}(\rho_k^X \parallel \nu^X) \leq \epsilon$ whenever $k \geq \frac{Ld}{2\alpha} \log \frac{D_{\Phi}(\rho_0^X \parallel \nu^X)}{\epsilon}$. The expected number of oracle calls to f is $\mathcal{O}\left(\frac{Ld}{\alpha} \log \frac{D_{\Phi}(\rho_0^X \parallel \nu^X)}{\epsilon}\right)$.*

1.4. Related work

Our main results, Theorems 2 and 3, study the mixing time in Φ -divergence of the ULA (6) and Proximal Sampler (9) respectively. We give an overview of the algorithms along with prior mixing time results in Sections 1.2 and 1.3, and refer the reader to those sections for the corresponding references. We discuss other related works here; see also Chewi (2024) for a comprehensive overview of Langevin-based samplers.

Both the ULA (6) and Proximal Sampler (9) have the Langevin dynamics (5) as the limiting dynamics as $\eta \rightarrow 0$, and the convergence of Φ -divergence along the Langevin dynamics is well-known (see Corollary 9, and see also Chafaï (2004); Dolbeault and Li (2018); Achleitner et al. (2015) for a general discussion for diffusions). To study the convergence for discrete-time algorithms, our main tool is that of Strong Data Processing Inequalities (SDPIs) (see Section 2.2). SDPIs are Markov chain-dependent strengthenings of the data processing inequality, and are a fundamental concept in information theory. For a comprehensive overview of SDPIs, we refer the reader to Polyanskiy and Wu (2024, Chapter 33). Just as data processing inequalities hold in many different metrics, so do SDPIs. Raginsky (2016) study SDPIs in Φ -divergence for discrete-space Markov chains, and du Pin Calmon et al. (2017); Polyanskiy and Wu (2017) study SDPIs in lesser generality than Φ -divergences, but for continuous-space chains and networks.

SDPI-inspired techniques have been used in prior works studying the mixing time of Langevin based algorithms, but most works use them implicitly, and to the best of our knowledge, none study Φ -divergences in general. Vempala and Wibisono (2019) use them to study the convergence of ULA to ν^η in Rényi divergence, Chen et al. (2022) use them for the Proximal Sampler, Yuan et al. (2023) mention SDPIs and use them for the Proximal Sampler on graphs, Kook et al. (2024); Kook and Zhang (2025) use them for the constrained Proximal Sampler. We make the SDPI-based approach

explicit (Section 2.2) and use it for Φ -divergence. It should also be noted that Rényi divergence is not a Φ -divergence, but it is a simple transformation of a Φ -divergence, so SDPI-based analyses have also been used to show mixing time in Rényi divergence: Vempala and Wibisono (2019) present mixing guarantees in Rényi divergence via SDPI for the ULA, and Chen et al. (2022) for the Proximal Sampler.

Bounding the contraction coefficient (Definition 10) is the key step in showing a Markov chain satisfies a SDPI, and the main idea used to bound the contraction coefficient for Langevin-based Markov chains is by taking the time derivative of the divergence along two simultaneous stochastic processes (Lemma 8). This idea traces back to de Bruijn’s identity (Stam, 1959), and similar methods have been used well beyond information theory, for example to study diffusions and diffusion models (Chafaï, 2004; Albergo et al., 2023; Vempala and Wibisono, 2019; Kook et al., 2024).

Organization We go over the necessary background material in Section 2, and then discuss the convergence of Φ -divergence along ULA in Section 3, and along the Proximal Sampler in Section 4. We summarize and discuss open questions in Section 5 to conclude.

2. Preliminaries

A distribution $\nu \propto \exp(-f)$ on \mathbb{R}^d with a twice-differentiable potential function f is α -strongly log-concave for some $\alpha > 0$ if $\alpha I \preceq \nabla^2 f$, and is L -smooth for some $L > 0$ if $-LI \preceq \nabla^2 f \preceq LI$. When $\alpha = 0$, we call ν (weakly) log-concave. Throughout, we take $\Phi: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ to be a twice-differentiable strictly convex function with $\Phi(1) = 0$. Whenever we refer to any probability distribution, we always take it to be a member of $\mathcal{P}_{2,ac}(\mathbb{R}^d)$, i.e. the set of probability distributions on \mathbb{R}^d which are absolutely continuous with respect to Lebesgue measure and have finite second moment. We also refer to distributions via their densities with respect to Lebesgue measure. We denote \mathcal{N}_t as shorthand for $\mathcal{N}(0, tI)$ where $t > 0$, and use $\mathcal{N}(\mu, \Sigma)$ to refer to a Gaussian distribution with mean $\mu \in \mathbb{R}^d$ and positive-definite covariance matrix $\Sigma \in \mathbb{R}^{d \times d}$.

2.1. Φ -divergences, Φ -Sobolev inequalities, and their properties

Recall the Φ -divergence between probability distributions is defined in (1). To study the mixing time of the ULA (Theorem 2) and Proximal Sampler (Theorem 3) in Φ -divergence, we assume a Φ -Sobolev inequality assumption on the stationary distribution (Definition 1).

Our proofs of Theorems 2 and 3 are based on SDPIs (Section 2.2), and core to this SDPI-based proof strategy will be analyzing the change of the Φ -Sobolev constant along various operations such as convolution and pushforward. These operations arise by interpreting the Markov chains (6) and (9) as updates in the space of distributions, and will be described in Sections 3 and 4 respectively. The evolution of the Φ -Sobolev constant along these operations is classical (Chafaï, 2004) and we discuss them now.

The following lemma tells us how the Φ SI constant evolves along a Lipschitz pushforward map.

Lemma 5 (Chafaï, 2004, Remark 7) *Assume ν satisfy Φ -Sobolev inequality with optimal constant $\alpha_{\Phi\text{SI}}(\nu)$. Let $T: \mathbb{R}^d \rightarrow \mathbb{R}^d$ be a γ -Lipschitz map. Then, the pushforward $\tilde{\nu} = T_{\#}\nu$ satisfies Φ -Sobolev inequality with optimal constant*

$$\alpha_{\Phi\text{SI}}(\tilde{\nu}) \geq \frac{\alpha_{\Phi\text{SI}}(\nu)}{\gamma^2}.$$

The next lemma describes the change of the Φ SI constant after convolution.

Lemma 6 (*Chafaï, 2004, Corollary 3.1*) *Assume μ and ν satisfy Φ SI with optimal constants $\alpha_{\Phi\text{SI}}(\mu)$ and $\alpha_{\Phi\text{SI}}(\nu)$, respectively. Then the convolution $\mu * \nu$ satisfies Φ SI with constant*

$$\frac{1}{\alpha_{\Phi\text{SI}}(\mu * \nu)} \leq \frac{1}{\alpha_{\Phi\text{SI}}(\mu)} + \frac{1}{\alpha_{\Phi\text{SI}}(\nu)}.$$

The following lemma tells us that when ν is α -strongly log-concave, it also satisfies Φ -Sobolev inequality with the same constant.

Lemma 7 (*Chafaï, 2004, Corollary 2.1*) *If ν is α -strongly log-concave for some $\alpha > 0$, then ν satisfies Φ SI with constant*

$$\alpha_{\Phi\text{SI}}(\nu) \geq \alpha.$$

We conclude this subsection by describing the rate of change of Φ -divergence along simultaneous evolutions of the same SDE. The following lemma will be crucial in the SDPI-based approach. The same lemma described in terms of Markov semigroup theory can be found in [Chewi \(2024, Theorem 8.3.1\)](#). We prove [Lemma 8](#) in [Appendix C](#).

Lemma 8 *Suppose $X_t \sim \mu_t$ and $X_t \sim \nu_t$ with initial conditions μ_0 and ν_0 are two solutions of the following SDE:*

$$dX_t = b_t(X_t) dt + \sqrt{2c} dW_t, \quad (10)$$

where $b_t : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a time-varying drift function, c is a positive constant, and W_t is the standard Brownian motion on \mathbb{R}^d . Then for all $t \geq 0$,

$$\frac{d}{dt} D_{\Phi}(\mu_t \parallel \nu_t) = -c \text{FI}_{\Phi}(\mu_t \parallel \nu_t).$$

As mentioned in [Section 1.4](#), when [\(10\)](#) is taken to be Brownian motion (i.e. the drift $b_t = 0$), ν_t is fixed to be the Lebesgue measure, and $\Phi(x) = x \log x$, [Lemma 8](#) corresponds to de Bruijn's identity ([Stam, 1959](#)). The SDE [\(10\)](#) is more general than the Langevin dynamics [\(5\)](#) as it includes a time-varying drift function; this is required to study the Proximal Sampler, since as we will see in [Section 4](#), the backward step of the Proximal Sampler corresponds to an SDE with time-varying drift.

As an easy consequence of [Lemma 8](#), we have the following exponential convergence of Φ -divergence for the Langevin dynamics when ν satisfies a Φ -Sobolev inequality.

Corollary 9 *Suppose $X_t \sim \rho_t$ evolves along the Langevin dynamics [\(5\)](#) to sample from $\nu \propto \exp(-f)$, and let ν satisfy a Φ -Sobolev inequality with optimal constant $\alpha > 0$. Then,*

$$D_{\Phi}(\rho_t \parallel \nu) \leq e^{-2\alpha t} D_{\Phi}(\rho_0 \parallel \nu).$$

Proof Applying [Lemma 8](#) (with $\mu_t = \rho_t$ and $\nu_t = \nu$) along with the Φ -Sobolev inequality of ν yields,

$$\frac{d}{dt} D_{\Phi}(\mu_t \parallel \nu) = -\text{FI}_{\Phi}(\mu_t \parallel \nu) \leq -2\alpha D_{\Phi}(\mu_t \parallel \nu).$$

Integrating the differential inequality from 0 to t using Grönwall's lemma completes the proof. ■

2.2. Strong data processing inequalities

We now discuss Strong Data Processing Inequalities (SDPIs), which will be our primary proof strategy. For any Φ -divergence, the data processing inequality states that for any two distributions μ and ν , and for any Markov kernel \mathbf{Q} , $D_\Phi(\mu\mathbf{Q} \parallel \nu\mathbf{Q}) \leq D_\Phi(\mu \parallel \nu)$, provided all quantities are well-defined.⁵ *Strong* data processing inequalities check if the inequality is strict, and quantify the decrease (Polyanskiy and Wu, 2024, 2017; Raginsky, 2016). They do so by fixing the second input distribution (i.e. ν) and then varying the first distribution (i.e. μ), to see the worst case ratio $\frac{D_\Phi(\mu\mathbf{Q} \parallel \nu\mathbf{Q})}{D_\Phi(\mu \parallel \nu)}$. The ratio is called the *contraction coefficient*, and when it is strictly less than 1 for all valid distributions μ , we say that (\mathbf{Q}, ν) satisfies a strong data processing inequality.

Next, we define the contraction coefficient in Definition 10 and define SDPIs in Φ -divergence in Definition 11.

Definition 10 *Let ν be a probability distribution, \mathbf{Q} be a Markov kernel, and D_Φ be a Φ -divergence. Then the contraction coefficient ε_{D_Φ} is defined as follows*

$$\varepsilon_{D_\Phi}(\mathbf{Q}, \nu) := \sup_{\rho: 0 < D_\Phi(\rho \parallel \nu) < \infty} \frac{D_\Phi(\rho\mathbf{Q} \parallel \nu\mathbf{Q})}{D_\Phi(\rho \parallel \nu)}. \quad (11)$$

Definition 11 *Let ν be a probability distribution, \mathbf{Q} be a Markov kernel, and D_Φ be a Φ -divergence. Further define $\varepsilon_{D_\Phi}(\mathbf{Q}, \nu)$ as in (11). Then we say that (\mathbf{Q}, ν) satisfies a strong data processing inequality in Φ -divergence when $\varepsilon_{D_\Phi}(\mathbf{Q}, \nu) < 1$. In particular, we have,*

$$D_\Phi(\mu\mathbf{Q} \parallel \nu\mathbf{Q}) \leq \varepsilon_{D_\Phi}(\mathbf{Q}, \nu) D_\Phi(\mu \parallel \nu), \quad (12)$$

where μ is any distribution such that $D_\Phi(\mu \parallel \nu) < \infty$.

As a direct consequence of Definition 11, we can see that if ν is invariant for \mathbf{Q} , then for any μ ,

$$D_\Phi(\mu\mathbf{Q}^k \parallel \nu) \leq \varepsilon_{D_\Phi}(\mathbf{Q}, \nu)^k D_\Phi(\mu \parallel \nu). \quad (13)$$

Hence, if we desire to show a mixing time result for Markov chain \mathbf{Q} with invariant distribution ν , then we require a bound on $\varepsilon_{D_\Phi}(\mathbf{Q}, \nu)$ which is strictly less than 1.

3. Convergence Along ULA

Recall the ULA update (6) for sampling from $\nu \propto \exp(-f)$. Defining $\rho_k := \text{law}(X_k)$, the update (6) can be seen as the following update in distribution:

$$\rho_{k+1} = (\text{id} - \eta\nabla f)_\# \rho_k * \mathcal{N}(0, 2\eta I). \quad (14)$$

Denote the Markov kernel of ULA as \mathbf{P} (i.e. $\rho_{k+1} = \rho_k \mathbf{P} = \rho_0 \mathbf{P}^{k+1}$), and define $F: \mathbb{R}^d \rightarrow \mathbb{R}^d$ by $F(x) = x - \eta\nabla f(x)$. Recall from Section 1.2 that ULA has a biased stationary distribution ν^η . Then as per (13), in order to obtain a mixing time result, our goal is to bound $\varepsilon_{D_\Phi}(\mathbf{P}, \nu^\eta)$. We now provide a proof outline for this.

5. For an example of the Markov kernel notation, and its behaviour as an operator on distributions, see (14) and the subsequent text.

3.1. Proof outline

We wish to bound $\varepsilon_{\mathbf{D}_\Phi}(\mathbf{P}, \nu^\eta)$. To that end, let μ be an arbitrary probability distribution and recall by Definition 10 that we want to control

$$\frac{\mathbf{D}_\Phi(\mu \mathbf{P} \parallel \nu^\eta \mathbf{P})}{\mathbf{D}_\Phi(\mu \parallel \nu^\eta)} = \frac{\mathbf{D}_\Phi(F_{\#}\mu * \mathcal{N}(0, 2\eta I) \parallel F_{\#}\nu^\eta * \mathcal{N}(0, 2\eta I))}{\mathbf{D}_\Phi(\mu \parallel \nu^\eta)}.$$

Suppose that F is a bijective map. Then, as Φ -divergence is invariant to simultaneous bijective deterministic maps, we have $\mathbf{D}_\Phi(\mu \parallel \nu^\eta) = \mathbf{D}_\Phi(F_{\#}\mu \parallel F_{\#}\nu^\eta)$.⁶ Therefore, under the assumption that F is bijective, the quantity we wish to bound is

$$\frac{\mathbf{D}_\Phi(F_{\#}\mu * \mathcal{N}(0, 2\eta I) \parallel F_{\#}\nu^\eta * \mathcal{N}(0, 2\eta I))}{\mathbf{D}_\Phi(F_{\#}\mu \parallel F_{\#}\nu^\eta)}.$$

Denoting $F_{\#}\mu * \mathcal{N}(0, tI) = \mu_t$ and $F_{\#}\nu^\eta * \mathcal{N}(0, tI) = \nu_t$, we want to bound

$$\frac{\mathbf{D}_\Phi(\mu_{2\eta} \parallel \nu_{2\eta})}{\mathbf{D}_\Phi(\mu_0 \parallel \nu_0)}.$$

As convolving with a Gaussian can be viewed as the solution to the heat equation, this quantity can be bounded as a consequence of Lemma 8 (for the Brownian motion SDE, i.e., for (10) with $b_t = 0$ and $c = \frac{1}{2}$) under the assumption that ν_t satisfies a Φ -Sobolev inequality for $t \in [0, 2\eta]$. The final expression obtained is independent of μ , and therefore this provides a bound on $\varepsilon_{\mathbf{D}_\Phi}(\mathbf{P}, \nu^\eta)$.

This is the outline we follow and the assumptions stated in Theorem 2 are to ensure that the assumptions mentioned in the proof sketch go through.

3.2. Proof of Theorem 2

In this section we prove Theorem 2. The key lemma for doing so is the following. The proof of Lemma 12 follows the outline mentioned in Section 3.1 and is in Appendix D.

Lemma 12 *Let \mathbf{P} denote the ULA Markov kernel with update (14). Suppose ν^η satisfies a Φ -Sobolev inequality with optimal constant α , ν is L -smooth, and $\eta \leq \frac{1}{L}$. Then,*

$$\varepsilon_{\mathbf{D}_\Phi}(\mathbf{P}, \nu^\eta) \leq \frac{(1 + \eta L)^2}{(1 + \eta L)^2 + 2\alpha\eta}.$$

Lemma 12 gives a bound on the contraction coefficient for ULA, with which we can prove Theorem 2.

Proof of Theorem 2 By an application of Definition 11 and by noting that ν^η is stationary for \mathbf{P} , we have that,

$$\mathbf{D}_\Phi(\rho_k \parallel \nu^\eta) \leq \varepsilon_{\mathbf{D}_\Phi}(\mathbf{P}, \nu^\eta)^k \mathbf{D}_\Phi(\rho_0 \parallel \nu^\eta).$$

Using Lemma 12, we get that,

$$\mathbf{D}_\Phi(\rho_k \parallel \nu^\eta) \leq \left(\frac{(1 + \eta L)^2}{(1 + \eta L)^2 + 2\alpha\eta} \right)^k \mathbf{D}_\Phi(\rho_0 \parallel \nu^\eta) = \left(1 + \frac{2\alpha\eta}{(1 + \eta L)^2} \right)^{-k} \mathbf{D}_\Phi(\rho_0 \parallel \nu^\eta).$$

6. This can be seen as a consequence of data processing inequality, i.e. for any μ and ρ and bijective map F , $\mathbf{D}_\Phi(\mu \parallel \rho) \geq \mathbf{D}_\Phi(F_{\#}\mu \parallel F_{\#}\rho) \geq \mathbf{D}_\Phi(F_{\#}^{-1}F_{\#}\mu \parallel F_{\#}^{-1}F_{\#}\rho) = \mathbf{D}_\Phi(\mu \parallel \rho)$, hence all the inequalities must be equalities.

■

An explicit calculation for the Ornstein-Uhlenbeck process implying the tightness of Theorem 2 for KL divergence can be found in Proposition 17 in Appendix G.1.

3.3. Property of the biased limit

We conclude this section by showing that the biased limit satisfies a Φ -Sobolev inequality under strong log-concavity of ν . We prove Lemma 13 in Appendix E.

Lemma 13 *Suppose $\nu \propto \exp(-f)$ is α -strongly log-concave and L -smooth. Consider the ULA (6) to sample from ν with step size $\eta \leq \frac{1}{L}$. Then the biased limit ν^η satisfies:*

$$\alpha_{\Phi\text{SI}}(\nu^\eta) \geq \frac{\alpha}{2}.$$

4. Convergence Along Proximal Sampler

Recall the Proximal Sampler from Section 1.3 with update given by (9), with the forward and backward steps. We will denote the Proximal Sampler as $\mathbf{P}_{\text{prox}} = \mathbf{P}_{\text{prox}}^+ \mathbf{P}_{\text{prox}}^-$ where $\mathbf{P}_{\text{prox}}^+$ corresponds to the forward step and $\mathbf{P}_{\text{prox}}^-$ to the backward step. Each step of the Proximal Sampler is an update on \mathbb{R}^d , and so this perspective via composition of Markov kernels is valid. In terms of notation, we have $\rho_k^X := \text{law}(X_k)$, $\rho_k^Y := \text{law}(Y_k)$, and therefore $\rho_k^X \mathbf{P}_{\text{prox}} = \rho_{k+1}^X$, $\rho_k^X \mathbf{P}_{\text{prox}}^+ = \rho_k^Y$, and $\rho_k^Y \mathbf{P}_{\text{prox}}^- = \rho_{k+1}^X$.

As mentioned in Section 1.3, we know that ν^X is stationary for the Proximal Sampler, and from (13), our goal is to then bound the contraction coefficient $\varepsilon_{\mathbf{D}_\Phi}(\mathbf{P}_{\text{prox}}, \nu^X)$. The following lemma shows that this can be bounded by the product of the contraction coefficients of the forward and backward steps.

Lemma 14 *Let $\mathbf{P}_{\text{prox}} = \mathbf{P}_{\text{prox}}^+ \mathbf{P}_{\text{prox}}^-$ denote the Proximal Sampler (9) with joint stationary distribution ν^{XY} . Then,*

$$\varepsilon_{\mathbf{D}_\Phi}(\mathbf{P}_{\text{prox}}, \nu^X) \leq \varepsilon_{\mathbf{D}_\Phi}(\mathbf{P}_{\text{prox}}^+, \nu^X) \varepsilon_{\mathbf{D}_\Phi}(\mathbf{P}_{\text{prox}}^-, \nu^Y).$$

Proof By Definition 10,

$$\begin{aligned} \varepsilon_{\mathbf{D}_\Phi}(\mathbf{P}_{\text{prox}}, \nu^X) &= \sup_{\mu} \frac{\mathbf{D}_\Phi(\mu \mathbf{P}_{\text{prox}} \parallel \nu^X \mathbf{P}_{\text{prox}})}{\mathbf{D}_\Phi(\mu \parallel \nu^X)} \\ &= \sup_{\mu} \left[\frac{\mathbf{D}_\Phi(\mu \mathbf{P}_{\text{prox}}^+ \parallel \nu^X \mathbf{P}_{\text{prox}}^+)}{\mathbf{D}_\Phi(\mu \parallel \nu^X)} \times \frac{\mathbf{D}_\Phi(\mu \mathbf{P}_{\text{prox}}^+ \mathbf{P}_{\text{prox}}^- \parallel \nu^X \mathbf{P}_{\text{prox}}^+ \mathbf{P}_{\text{prox}}^-)}{\mathbf{D}_\Phi(\mu \mathbf{P}_{\text{prox}}^+ \parallel \nu^X \mathbf{P}_{\text{prox}}^+)} \right] \\ &\leq \varepsilon_{\mathbf{D}_\Phi}(\mathbf{P}_{\text{prox}}^+, \nu^X) \varepsilon_{\mathbf{D}_\Phi}(\mathbf{P}_{\text{prox}}^-, \nu^X \mathbf{P}_{\text{prox}}^+). \end{aligned}$$

The claim follows by noting that $\nu^Y = \nu^X \mathbf{P}_{\text{prox}}^+$ (which is further made clear in Section 4.1). ■

In light of Lemma 14, we will bound each of the contraction coefficients separately and get our convergence result. As each contraction coefficient can at most equal 1 (as a consequence of data processing inequality), this perspective implies that showing that either coefficient is strictly less than 1 yields a convergence guarantee in Φ -divergence. We will in fact show both coefficients for the forward step and the backward step are strictly less than 1.

4.1. Forward step

Without loss of generality consider $k = 0$, and recall the forward step of the Proximal Sampler (9), where $X_0 \sim \rho_0^X$ and $Y_0 | X_0 \sim \mathcal{N}(X_0, \eta I)$. To understand the forward step, we wish to understand ρ_0^Y . For any $y \in \mathbb{R}^d$, $\rho_0^Y(y) = \int \nu^{Y|X}(y | x) \rho_0^X(x) dx = \int \frac{1}{(2\pi\eta)^{\frac{d}{2}}} \exp(-\frac{\|y-x\|^2}{2\eta}) \rho_0^X(x) dx$, and therefore $\rho_0^Y = \rho_0^X * \mathcal{N}(0, \eta I)$. This implies that $\mathbf{P}_{\text{prox}}^+$ corresponds to evolving along the heat flow (i.e. $dX_t = dW_t$) for time η . And so, to get a control on $\varepsilon_{\text{D}_\Phi}(\mathbf{P}_{\text{prox}}^+, \nu^X)$, we can use Lemma 8 with $b_t = 0$ and $c = \frac{1}{2}$, just like was done for ULA (discussed in Section 3). This leads to the following guarantee on the contraction coefficient. We prove Lemma 15 in Appendix F.1.

Lemma 15 *Let $\mathbf{P}_{\text{prox}} = \mathbf{P}_{\text{prox}}^+ \mathbf{P}_{\text{prox}}^-$ denote the Proximal Sampler (9) with step size $\eta > 0$, to sample from ν^X where ν^X satisfies a Φ -Sobolev inequality with optimal constant $\alpha > 0$. Then,*

$$\varepsilon_{\text{D}_\Phi}(\mathbf{P}_{\text{prox}}^+, \nu^X) \leq \frac{1}{1 + \alpha\eta}.$$

4.2. Backward step

We now focus on the backward step of the Proximal Sampler. For the forward step, we were able to relate it to the Brownian motion SDE ($dX_t = dW_t$) by observing how it acts on distributions (i.e. on ρ_0^X , and seeing that $\rho_0^Y = \rho_0^X * \mathcal{N}(0, \eta I)$). For the backward step, this approach is not as clear. Indeed, writing $\rho_1^X(x) = \int \nu^{X|Y}(x | y) \rho_0^Y(y) dy$ does not immediately yield an SDE interpretation. By step 2 of the Proximal Sampler (9), we want to find an SDE such that when initialized from a point mass δ_y at any $y \in \mathbb{R}^d$, the output has distribution $\nu^{X|Y=y}$ at time η ; and therefore in general, when initialized at ν^Y the SDE will output ν^X , and when initialized at ρ_0^Y the SDE will output ρ_1^X . It turns out we can obtain this by reversing the heat flow path from ν^X to ν^Y .

For the forward step, we have that $dX_t = dW_t$ where if $X_0 \sim \nu^X$, then $X_\eta \sim \nu^Y$. The time reversal of this SDE is called the *backward heat flow* and is given by

$$dY_t = \nabla \log(\nu^X * \mathcal{N}_{\eta-t})(Y_t) dt + dW_t. \quad (15)$$

By construction, if we start (15) from ν^Y , then for any $t \in [0, \eta]$, the marginal law along (15) is the same as $\nu^X * \mathcal{N}_{\eta-t}$. Such a reverse SDE construction is popular in diffusion models (Chen et al., 2023), and details regarding it can be found in Föllmer (2005); Cattiaux et al. (2023). Rigorous connections between (15) and the Proximal Sampler can be found in Chen et al. (2022, Appendix A.1.2); see also the exposition in Chewi (2024, Chapter 8.3) and Kook et al. (2024, Appendix B.2).

The following lemma describes the contraction coefficient for the backward step. We provide the proof of Lemma 16 in Appendix F.2.

Lemma 16 *Let $\mathbf{P}_{\text{prox}} = \mathbf{P}_{\text{prox}}^+ \mathbf{P}_{\text{prox}}^-$ denote the Proximal Sampler (9) with step size $\eta > 0$, to sample from ν^X where ν^X satisfies a Φ -Sobolev inequality with optimal constant $\alpha > 0$. Then,*

$$\varepsilon_{\text{D}_\Phi}(\mathbf{P}_{\text{prox}}^-, \nu^Y) \leq \frac{1}{1 + \alpha\eta}.$$

The reason that both Lemmas 15 and 16 have the same result is for two reasons. First, it is because Lemma 8 does not depend on the drift b_t . Second, it is because the backward heat flow is the

time reversal of the forward heat flow SDE, and therefore by construction, there is a correspondence between the marginal distributions of the forward and backward processes in the sense that if we start (15) from ν^Y , then for any $t \in [0, \eta]$, the marginal law along (15) is the same as $\nu^X * \mathcal{N}_{\eta-t}$.

4.3. Proof of Theorem 3

Equipped with Lemmas 15 and 16, the proof of Theorem 3 follows easily.

Proof of Theorem 3 By repeated application of Definition 11 and by the fact that ν^X is stationary for the Proximal Sampler, we have that,

$$D_{\Phi}(\rho_k^X \parallel \nu^X) \leq \varepsilon_{D_{\Phi}}(\mathbf{P}_{\text{prox}}, \nu^X)^k D_{\Phi}(\rho_0^X \parallel \nu^X).$$

Further using Lemmas 14, 15, and 16, we get,

$$D_{\Phi}(\rho_k^X \parallel \nu^X) \leq \frac{D_{\Phi}(\rho_0^X \parallel \nu^X)}{(1 + \alpha\eta)^{2k}}.$$

■

An explicit calculation for the Ornstein-Uhlenbeck process implying the tightness of Theorem 3 for KL divergence can be found in Proposition 18 in Appendix G.2.

5. Discussion

We study the mixing time in Φ -divergence for two popular discretizations of the Langevin dynamics, namely the Unadjusted Langevin Algorithm (ULA) and the Proximal Sampler. Our results show mixing to the stationary distributions of these Markov chains under the stationary distributions satisfying Φ -Sobolev inequalities. As the Proximal Sampler is unbiased, this implies mixing to the target distribution ν of interest. However, this is not the case for the ULA, where our mixing guarantees are to the biased limit ν^η of ULA. While ν^η can be shown to satisfy a Φ -Sobolev inequality under strong log-concavity of ν (Lemma 13), it is interesting to study what can hold under weaker assumptions on ν . For example one can ask if the fact that ν^η satisfies a Φ -Sobolev inequality is closed under certain perturbations of ν . Alternatively, one can also ask for biased convergence guarantees of ULA to ν in Φ -divergence directly. One can also pose all of these questions for other Markov chains, such as Hamiltonian Monte Carlo and the underdamped Langevin algorithm.

Our results are based on strong data processing inequalities in Φ -divergence, and proceed by bounding the respective contraction coefficients. However, different forms of (strong) data processing inequalities exist, namely in terms of mutual information and Φ -mutual information. It is interesting to ask if our approaches can extend to these other forms of strong data processing inequalities, to yield for example, the convergence of the mutual information functional. Additionally, we require the function Φ to be twice differentiable, which prohibits many popular Φ -divergences such as total variation distance and “hockey-stick” divergences (Sason and Verdú, 2016). These divergences are popular in differential privacy applications (Asodeh et al., 2020) and it would therefore be instructive to study extensions to such divergences directly.

Finally, our results imply mixing in all (twice-differentiable) Φ -divergences for strongly log-concave target distributions. It is interesting to ask if mixing bounds can be obtained in a variety of Φ -divergences under weaker absolute isoperimetric assumptions on ν (such as a log-Sobolev inequality, or in other words, a Φ -Sobolev inequality with $\Phi(x) = x \log x$).

Acknowledgments

We thank Varun Jog, Sinho Chewi, and Yuliang Wang for helpful comments and references. S.M. and A.W. were supported by NSF Award CCF #2403391.

References

- Franz Achleitner, Anton Arnold, and Dominik Stürzer. Large-time behavior in non-symmetric Fokker-Planck equations. *Rivista di Matematica della Università di Parma*, 6(1):1–68, 2015.
- Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023.
- Jason Altschuler and Kunal Talwar. Resolving the Mixing Time of the Langevin Algorithm to its Stationary Distribution for Log-Concave Sampling. In *Proceedings of Thirty Sixth Conference on Learning Theory*, Proceedings of Machine Learning Research. PMLR, 2023.
- Jason M Altschuler and Sinho Chewi. Faster high-accuracy log-concave sampling via algorithmic warm starts. *Journal of the ACM*, 71(3):1–55, 2024.
- Shahab Asoodeh, Mario Diaz, and Flavio P Calmon. Contraction of E_γ -Divergence and Its Applications to Privacy. *arXiv preprint arXiv:2012.11035*, 2020.
- Dominique Bakry, Ivan Gentil, and Michel Ledoux. *Analysis and geometry of Markov diffusion operators*, volume 103. Springer, 2014.
- A Belitski, A Gretton, C Magri, Y Murayama, M Montemurro, N Logothetis, and S Panzeri. Local field potentials and spiking activity in primary visual cortex convey independent information about natural stimuli. *Journal of Neuroscience*, 28(22):5696–5709, 2008.
- Patrick Cattiaux, Giovanni Conforti, Ivan Gentil, and Christian Léonard. Time reversal of diffusion processes under a finite entropy condition. In *Annales de l’Institut Henri Poincaré (B) Probabilités et Statistiques*, volume 59, pages 1844–1881. Institut Henri Poincaré, 2023.
- Djalil Chafaï. Entropies, convexity, and functional inequalities: On Φ -entropies and Φ -Sobolev inequalities. *Journal of Mathematics of Kyoto University*, 2004.
- Sitan Chen, Sinho Chewi, Jerry Li, Yuanzhi Li, Adil Salim, and Anru R. Zhang. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. In *The Eleventh International Conference on Learning Representations*, 2023.
- Yongxin Chen, Sinho Chewi, Adil Salim, and Andre Wibisono. Improved analysis for a proximal algorithm for sampling. In *Conference on Learning Theory*, pages 2984–3014. PMLR, 2022.
- Xiang Cheng and Peter Bartlett. Convergence of Langevin MCMC in KL-divergence. In *Algorithmic Learning Theory*, pages 186–211. PMLR, 2018.
- Sinho Chewi. Log-concave sampling. 2024. Draft available at: <https://chewisinho.github.io>.

- Sinho Chewi, Murat A. Erdogdu, Mufan (Bill) Li, Ruoqi Shen, and Matthew Zhang. Analysis of Langevin Monte Carlo from Poincaré to log-Sobolev. In *Proceedings of Thirty Fifth Conference on Learning Theory*, Proceedings of Machine Learning Research, pages 1–2. PMLR, 2022a.
- Sinho Chewi, Patrik R Gerber, Chen Lu, Thibaut Le Gouic, and Philippe Rigollet. The query complexity of sampling from strongly log-concave distributions in one dimension. In *Conference on Learning Theory*, pages 2041–2059. PMLR, 2022b.
- I. Csiszár. Information-type measures of difference of probability distributions and indirect observations. *Studia Sci. Math. Hungar.*, 2:299–318, 1967.
- Arnak Dalalyan. Further and stronger analogy between sampling and optimization: Langevin Monte Carlo and gradient descent. In *Conference on Learning Theory*, pages 678–689. PMLR, 2017.
- Jean Dolbeault and Xingyu Li. φ -entropies: Convexity, coercivity and hypocoercivity for Fokker–Planck and kinetic Fokker–Planck equations. *Mathematical Models and Methods in Applied Sciences*, 28(13):2637–2666, 2018.
- Flavio du Pin Calmon, Yury Polyanskiy, and Yihong Wu. Strong data processing inequalities for input constrained additive noise channels. *IEEE Transactions on Information Theory*, 64(3):1879–1892, 2017.
- Jiaojiao Fan, Bo Yuan, and Yongxin Chen. Improved dimension dependence of a proximal algorithm for sampling. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 1473–1521. PMLR, 2023.
- Hans Föllmer. An entropy approach to the time reversal of diffusion processes. In *Stochastic Differential Systems Filtering and Control: Proceedings of the IFIP-WG 7/1 Working Conference Marseille-Luminy, France, March 12–17, 1984*, pages 156–163. Springer, 2005.
- Andrew Gelman, John B Carlin, Hal S Stern, and Donald B Rubin. *Bayesian data analysis*. Chapman & Hall/CRC, 1995.
- Charles J Geyer. Introduction to Markov chain Monte Carlo. *Handbook of Markov chain Monte Carlo*, 20116022(45):22, 2011.
- Walter R Gilks, Sylvia Richardson, and David J Spiegelhalter. Introducing Markov chain Monte Carlo. *Markov chain Monte Carlo in practice*, 1, 1995.
- Arthur Gretton and László Györfi. Nonparametric independence tests: Space partitioning and kernel approaches. In *Algorithmic Learning Theory: 19th International Conference, ALT 2008, Budapest, Hungary, October 13-16, 2008. Proceedings 19*, pages 183–198. Springer, 2008.
- László Györfi and I Vajda. Asymptotic distributions for goodness-of-fit statistics in a sequence of multinomial models. *Statistics & probability letters*, 56(1):57–67, 2002.
- Chin Pang Ho, Marek Petrik, and Wolfram Wiesemann. Robust Φ -Divergence MDPs. *Advances in Neural Information Processing Systems*, 35:32680–32693, 2022.
- Richard Holley and Daniel W. Stroock. Logarithmic Sobolev inequalities and stochastic Ising models. *Journal of Statistical Physics*, 46:1159–1194, 1987.

- Michael Johannes and Nicholas Polson. MCMC Methods for Continuous-Time Financial Econometric. *Handbook of Financial Econometrics*, 2, 2003.
- Richard Jordan, David Kinderlehrer, and Felix Otto. The variational formulation of the Fokker-Planck equation. *SIAM journal on mathematical analysis*, 29(1):1–17, 1998.
- Yunbum Kook and Matthew S Zhang. Rényi-infinity constrained sampling with d^3 membership queries. In *Symposium on Discrete Algorithms (SODA)*, 2025.
- Yunbum Kook, Santosh S. Vempala, and Matthew S. Zhang. In-and-Out: Algorithmic Diffusion for Sampling Convex Bodies. *Advances in Neural Information Processing Systems*, 2024.
- Yin Tat Lee, Ruoqi Shen, and Kevin Tian. Structured log-concave sampling with a restricted Gaussian oracle. In *Conference on Learning Theory*, pages 2993–3050. PMLR, 2021.
- David A Levin, Yuval Peres, and Elizabeth Wilmer. *Markov chains and mixing times*, volume 107. American Mathematical Society, 2017.
- Jiaming Liang and Yongxin Chen. A proximal algorithm for sampling from non-smooth potentials. In *2022 Winter Simulation Conference (WSC)*, pages 3229–3240. IEEE, 2022.
- David JC MacKay. *Information theory, inference and learning algorithms*. Cambridge University Press, 2003.
- Bernard Martinet. Brève communication. régularisation d’inéquations variationnelles par approximations successives. *Revue française d’informatique et de recherche opérationnelle. Série rouge*, 4(R3):154–158, 1970.
- Ilya Nemenman, William Bialek, and Rob de Ruyter van Steveninck. Entropy and information in neural spike trains: Progress on the sampling problem. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 69(5):056111, 2004.
- Kishan Panaganti, Adam Wierman, and Eric Mazumdar. Model-Free Robust Φ -Divergence Reinforcement Learning Using Both Offline and Online Data. *arXiv preprint arXiv:2405.05468*, 2024.
- Ankit Pensia, Varun Jog, and Po-Ling Loh. The sample complexity of simple binary hypothesis testing. In *Proceedings of Thirty Seventh Conference on Learning Theory*, Proceedings of Machine Learning Research. PMLR, 30 Jun–03 Jul 2024.
- Yury Polyanskiy and Yihong Wu. Strong data-processing inequalities for channels and Bayesian networks. In *Convexity and Concentration*, pages 211–249. Springer, 2017.
- Yury Polyanskiy and Yihong Wu. *Information Theory: From Coding to Learning*. Cambridge University Press, 2024. Draft available at: <http://www.stat.yale.edu/~yw562/teaching/itbook-export.pdf>.
- Maxim Raginsky. Strong data processing inequalities and Φ -Sobolev inequalities for discrete channels. *IEEE Transactions on Information Theory*, 62(6):3355–3389, 2016.

- Christian P Robert, George Casella, and George Casella. *Monte Carlo statistical methods*, volume 2. Springer, 1999.
- Gareth O Roberts and Jeffrey S Rosenthal. Optimal scaling of discrete approximations to Langevin diffusions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60(1): 255–268, 1998.
- Gareth O Roberts and Richard L Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, 2(4):341–363, 1996.
- R Tyrrell Rockafellar. Monotone operators and the proximal point algorithm. *SIAM journal on control and optimization*, 14(5):877–898, 1976.
- Igal Sason and Sergio Verdú. f -divergence inequalities. *IEEE Transactions on Information Theory*, 62(11):5973–6006, 2016.
- Aart J Stam. Some inequalities satisfied by the quantities of information of Fisher and Shannon. *Information and Control*, 2(2):101–112, 1959.
- Santosh Vempala and Andre Wibisono. Rapid convergence of the Unadjusted Langevin Algorithm: Isoperimetry suffices. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Udo Von Toussaint. Bayesian inference in physics. *Reviews of Modern Physics*, 83(3):943, 2011.
- Andre Wibisono. Sampling as optimization in the space of measures: The Langevin dynamics as a composite optimization problem. In *Conference on Learning Theory*, pages 2093–3027. PMLR, 2018.
- Bo Yuan, Jiaojiao Fan, Jiaming Liang, Andre Wibisono, and Yongxin Chen. On a Class of Gibbs Sampling over Networks. In *Proceedings of Thirty Sixth Conference on Learning Theory*, Proceedings of Machine Learning Research, 2023.

Appendix A. Examples of Φ -divergences

$\Phi(x)$	$\mathbf{D}_\Phi(\mu \parallel \nu)$	$\mathbf{D}_\Phi(\mu \parallel \nu)$ name	Φ -Sobolev inequality
$x \log x$	$\int d\mu \log \frac{d\mu}{d\nu}$	KL divergence	log-Sobolev inequality
$(x - 1)^2$	$\int \frac{(d\mu - d\nu)^2}{d\nu}$	chi-squared divergence	Poincaré inequality
$\frac{1}{2}(\sqrt{x} - 1)^2$	$\frac{1}{2} \int (\sqrt{d\mu} - \sqrt{d\nu})^2$	squared Hellinger distance	–
$\frac{1}{2} x - 1 $	$\frac{1}{2} \int d\mu - d\nu $	TV distance	–
$-\log x$	$\int d\nu \log \frac{d\nu}{d\mu}$	reverse KL divergence	–
$\frac{1}{x} - x$	$2 + \int \frac{(d\mu - d\nu)^2}{d\mu}$	reverse chi-squared divergence	–

Table 1: Common Φ functions along with corresponding Φ -divergences (1) and Φ -Sobolev inequalities (Definition 1).

Appendix B. Restricted Gaussian Oracle

Here we show a basic implementation of the RGO via rejection sampling. As mentioned in Section 1.3, enhanced implementations of the RGO under weaker assumptions have been an active area of research and we refer the reader to Fan et al. (2023) and the references therein.

Recall from Section 1.3 the conditional distribution the RGO seeks to sample from:

$$\nu^{X|Y}(x | y) \propto_x \exp\left(-f(x) - \frac{\|x - y\|^2}{2\eta}\right).$$

Define $g_y(x) := f(x) + \frac{\|x - y\|^2}{2\eta}$ so that for any fixed $y \in \mathbb{R}^d$, the target distribution for the RGO is $\tilde{\nu}_y(x) \propto \exp(-g_y(x))$. Suppose the potential function f is L -smooth and that $\eta < \frac{1}{L}$. In this case, $\tilde{\nu}_y$ is strongly log-concave with condition number $\frac{1+L\eta}{1-L\eta}$. and the RGO can be implemented efficiently via rejection sampling.

Suppose $\pi \propto \exp(-V)$ is β -strongly log-concave and M -smooth. The rejection sampling method to sample from π is the following:

1. Compute the minimizer x^* of V , so that for any $z \in \mathbb{R}^d$, $V(z) \geq V(x^*) + \frac{\beta}{2}\|z - x^*\|^2$.
2. Draw $Z \sim \mathcal{N}(x^*, \frac{1}{\beta}I)$ and accept it with probability

$$\exp\left(-V(Z) + V(x^*) + \frac{\beta}{2}\|Z - x^*\|^2\right).$$

Repeat this until acceptance.

The output of this method is distributed according to π and the expected number of iterations is $(\frac{M}{\beta})^{d/2}$ (Chewi et al., 2022b, Theorem 7).

Applying this to sample from $\tilde{\nu}_y$ with $\eta \asymp \frac{1}{Ld}$ gives a valid implementation of the RGO under smoothness of f with $\mathcal{O}(1)$ many iterations in expectation. Specifically, $M = L + \frac{1}{\eta}$ and $\beta = -L + \frac{1}{\eta}$ and therefore, $\frac{M}{\beta} = \frac{1+L\eta}{1-L\eta}$. So if $\eta = \frac{1}{Ld}$, $(\frac{M}{\beta})^{d/2} = (1 + \frac{2}{d-1})^{d/2} = \mathcal{O}(1)$.

Appendix C. Proof of Lemma 8

Proof Begin by recalling that if $X_t \sim \rho_t$ where $dX_t = b_t(X_t) dt + \sqrt{2c} dW_t$, then $\rho_t : \mathbb{R}^d \rightarrow \mathbb{R}$ satisfies the Fokker-Planck equation, given by:

$$\partial_t \rho_t = -\nabla \cdot (\rho_t b_t) + c \Delta \rho_t.$$

Also note that using the identity $\Delta \rho = \nabla \cdot (\rho \nabla \log \rho)$, the above can be written as:

$$\partial_t \rho_t = -\nabla \cdot (\rho_t b_t) + c \nabla \cdot (\rho_t \nabla \log \rho_t). \quad (16)$$

We identify μ_t and ν_t with their densities with respect to Lebesgue measure, and further denote their relative density as $h_t = \frac{\mu_t}{\nu_t}$. We also assume enough regularity to take the differential under the integral sign and use $\int fg$ as a shorthand for $\int f(x)g(x) dx$. Throughout the proof, we use integration by parts in various steps, denoted by (IBP). With all of this in mind, we have the following:

$$\begin{aligned} \partial_t D_\Phi(\mu_t \parallel \nu_t) &= \partial_t \int \nu_t \Phi(h_t) \\ &= \int (\partial_t \nu_t) \Phi(h_t) + \int \nu_t (\partial_t \Phi(h_t)) \\ &= \int (\partial_t \nu_t) \Phi(h_t) + \int \nu_t \Phi'(h_t) \frac{\nu_t \partial_t \mu_t - \mu_t \partial_t \nu_t}{\nu_t^2} \\ &= \underbrace{\int (\partial_t \nu_t) \Phi(h_t)}_{T_1} + \underbrace{\int \Phi'(h_t) (\partial_t \mu_t)}_{T_2} - \underbrace{\int \Phi'(h_t) \frac{\mu_t}{\nu_t} (\partial_t \nu_t)}_{T_3} \end{aligned}$$

We will now handle each of these terms separately. We have:

$$\begin{aligned} T_1 &= \int (\partial_t \nu_t) \Phi(h_t) \\ &\stackrel{(16)}{=} \int (-\nabla \cdot (\nu_t b_t) + c \nabla \cdot (\nu_t \nabla \log \nu_t)) \Phi(h_t) \\ &= - \int \nabla \cdot (\nu_t b_t) \Phi(h_t) + c \int \nabla \cdot (\nu_t \nabla \log \nu_t) \Phi(h_t) \\ &\stackrel{(IBP)}{=} \int \langle \nu_t b_t, \nabla(\Phi(h_t)) \rangle - c \int \langle \nu_t \nabla \log \nu_t, \nabla(\Phi(h_t)) \rangle \\ &= \int \langle \nu_t b_t, \Phi'(h_t) \nabla \frac{\mu_t}{\nu_t} \rangle - c \int \langle \nu_t \nabla \log \nu_t, \Phi'(h_t) \nabla \frac{\mu_t}{\nu_t} \rangle \end{aligned}$$

We also have:

$$\begin{aligned}
 T_2 &= \int \Phi'(h_t)(\partial_t \mu_t) \\
 &\stackrel{(16)}{=} \int \Phi'(h_t) (-\nabla \cdot (\mu_t b_t) + c \nabla \cdot (\mu_t \nabla \log \mu_t)) \\
 &\stackrel{(\text{IBP})}{=} \int \langle \mu_t b_t, \nabla(\Phi'(h_t)) \rangle - c \int \langle \mu_t \nabla \log \mu_t, \nabla(\Phi'(h_t)) \rangle \\
 &= \int \langle \mu_t b_t, \Phi''(h_t) \nabla \frac{\mu_t}{\nu_t} \rangle - c \int \langle \mu_t \nabla \log \mu_t, \Phi''(h_t) \nabla \frac{\mu_t}{\nu_t} \rangle
 \end{aligned}$$

We also have:

$$\begin{aligned}
 T_3 &= \int \Phi'(h_t) \frac{\mu_t}{\nu_t} (\partial_t \nu_t) \\
 &\stackrel{(16)}{=} \int \Phi'(h_t) \frac{\mu_t}{\nu_t} (-\nabla \cdot (\nu_t b_t) + c \nabla \cdot (\nu_t \nabla \log \nu_t)) \\
 &\stackrel{(\text{IBP})}{=} \int \left\langle \nu_t b_t, \nabla \left(\Phi'(h_t) \frac{\mu_t}{\nu_t} \right) \right\rangle - c \int \left\langle \nu_t \nabla \log \nu_t, \nabla \left(\Phi'(h_t) \frac{\mu_t}{\nu_t} \right) \right\rangle \\
 &= \int \langle \mu_t b_t, \Phi''(h_t) \nabla \frac{\mu_t}{\nu_t} \rangle + \int \langle \nu_t b_t, \Phi'(h_t) \nabla \frac{\mu_t}{\nu_t} \rangle \\
 &\quad - c \int \left\langle \mu_t \nabla \log \nu_t, \Phi''(h_t) \nabla \frac{\mu_t}{\nu_t} \right\rangle - c \int \left\langle \nu_t \nabla \log \nu_t, \Phi'(h_t) \nabla \frac{\mu_t}{\nu_t} \right\rangle
 \end{aligned}$$

Therefore, combining the above, we see many terms cancel and we have the following:

$$\begin{aligned}
 \partial_t \mathsf{D}_\Phi(\mu_t \parallel \nu_t) &= T_1 + T_2 - T_3 \\
 &= c \int \left\langle \mu_t \nabla \log \nu_t, \Phi''(h_t) \nabla \frac{\mu_t}{\nu_t} \right\rangle - c \int \left\langle \mu_t \nabla \log \mu_t, \Phi''(h_t) \nabla \frac{\mu_t}{\nu_t} \right\rangle \\
 &= -c \int \left\langle \nabla \log \frac{\mu_t}{\nu_t}, \Phi''(h_t) \nabla \frac{\mu_t}{\nu_t} \right\rangle \mu_t \\
 &= -c \mathbb{E}_{\mu_t} \left[\left\langle \nabla \log \frac{\mu_t}{\nu_t}, \Phi''(h_t) \nabla \frac{\mu_t}{\nu_t} \right\rangle \right] \\
 &= -c \mathbb{E}_{\nu_t} \left[\frac{\mu_t}{\nu_t} \left\langle \nabla \log \frac{\mu_t}{\nu_t}, \Phi''(h_t) \nabla \frac{\mu_t}{\nu_t} \right\rangle \right] \\
 &= -c \mathbb{E}_{\nu_t} \left[\left\langle \nabla \frac{\mu_t}{\nu_t}, \Phi''(h_t) \nabla \frac{\mu_t}{\nu_t} \right\rangle \right] \\
 &= -c \mathsf{FI}_\Phi(\mu_t \parallel \nu_t),
 \end{aligned}$$

which proves the desired statement. ■

Appendix D. Proof Of Lemma 12

Proof Let μ be an arbitrary probability distribution such that $\mathsf{D}_\Phi(\mu \parallel \nu^\eta) < \infty$. As ν is L -smooth (i.e. $-LI \preceq \nabla^2 f \preceq LI$) and $\eta \leq \frac{1}{L}$, $F(x) = x - \eta \nabla f(x)$ is a bijective map. Furthermore, it also holds that $\|F\|_{\text{Lip}} \leq 1 + \eta L$.

Therefore, from Lemma 5, we see that $\alpha_{\Phi\text{SI}}(F_{\#}\nu^\eta) \geq \frac{\alpha}{(1+\eta L)^2}$. Now for $t \geq 0$, let $\mu_t := F_{\#}\mu * \mathcal{N}(0, tI)$ and $\nu_t := F_{\#}\nu^\eta * \mathcal{N}(0, tI)$. For $t \geq 0$, further denote $\alpha_{\Phi\text{SI}}(\nu_t) = \alpha_t$ as shorthand. Also note that $\mathcal{N}(0, tI)$ is $1/t$ -strongly log-concave and therefore from Lemma 7, $\alpha_{\Phi\text{SI}}(\mathcal{N}(0, tI)) \geq 1/t$. Hence, Lemma 6 implies that

$$\alpha_t \geq \frac{\alpha}{(1 + \eta L)^2 + \alpha t}.$$

The rate of change of Φ -divergence between μ_t and ν_t is given by Lemma 8. Applying Lemma 8 for (10) with $b_t = 0$ and $c = \frac{1}{2}$, which is what the (Gaussian convolution) evolution of μ_t and ν_t corresponds to, implies,

$$\begin{aligned} \frac{d}{dt} D_\Phi(\mu_t \parallel \nu_t) &= -\frac{1}{2} \text{Fl}_\Phi(\mu_t \parallel \nu_t) \\ &\leq -\alpha_t D_\Phi(\mu_t \parallel \nu_t), \end{aligned}$$

where the inequality follows from ν_t satisfying a Φ -Sobolev inequality. Applying Grönwall's lemma and integrating the differential inequality from 0 to 2η yields,

$$\frac{D_\Phi(\mu_{2\eta} \parallel \nu_{2\eta})}{D_\Phi(\mu_0 \parallel \nu_0)} \leq \exp\left(-\int_0^{2\eta} \alpha_t dt\right) \leq \frac{(1 + \eta L)^2}{(1 + \eta L)^2 + 2\alpha\eta}.$$

Finally, note that $D_\Phi(\mu_{2\eta} \parallel \nu_{2\eta}) = D_\Phi(\mu\mathbf{P} \parallel \nu^\eta\mathbf{P})$ and $D_\Phi(\mu_0 \parallel \nu_0) = D_\Phi(\mu \parallel \nu^\eta)$, where the latter holds as F is bijective. This, along with the fact that μ was arbitrary completes the proof. \blacksquare

Appendix E. Proof of Lemma 13

Proof Recall the ULA update in law (14). Under the assumptions of α -strong log-concavity, L -smoothness, and $\eta \leq \frac{1}{L}$, it holds that $F(x) = x - \eta \nabla f(x)$ is a bijective map with $\|F\|_{\text{Lip}} \leq 1 - \alpha\eta$. Also note that $\mathcal{N}_{2\eta}$ is $\frac{1}{2\eta}$ -SLC and therefore satisfies a Φ -Sobolev inequality with the same constant (Lemma 7). Suppose ρ_k satisfies α_k - ΦSI . Then Lemmas 5 and 6 imply that

$$\frac{1}{\alpha_{k+1}} \leq \frac{(1 - \alpha\eta)^2}{\alpha_k} + 2\eta.$$

Therefore, suppose we start from ρ_0 such that $\alpha_0 \geq \frac{\alpha}{2}$. Then by induction $\alpha_k \geq \frac{\alpha}{2}$ for all $k \geq 0$. Indeed, let us show that if $\alpha_k \geq \frac{\alpha}{2}$, then $\alpha_{k+1} \geq \frac{\alpha}{2}$. This follows as:

$$\frac{1}{\alpha_{k+1}} \leq \frac{2(1 - \alpha\eta)^2}{\alpha} + 2\eta = \frac{2}{\alpha}(1 - \alpha\eta(1 - \alpha\eta)) \leq \frac{2}{\alpha}.$$

Therefore, taking $k \rightarrow \infty$, we get that for the limiting distribution, $\alpha_{\Phi\text{SI}}(\nu^\eta) \geq \frac{\alpha}{2}$. \blacksquare

Appendix F. Deferred Proofs From Section 4

F.1. Proof of Lemma 15

Proof Let μ be an arbitrary distribution such that $D_\Phi(\mu \parallel \nu^X) < \infty$. Define $\mu_t := \mu * \mathcal{N}_t$ and $\nu_t := \nu^X * \mathcal{N}_t$. Further denote α_t as shorthand for $\alpha_{\Phi\text{SI}}(\nu_t)$. By assumption, we know that ν^X

satisfies a Φ -Sobolev inequality with optimal constant α . Also note that \mathcal{N}_t is $\frac{1}{t}$ -SLC and therefore satisfies a Φ -Sobolev inequality with the same constant (Lemma 7). Therefore by Lemma 6 we have that

$$\alpha_t \geq \frac{\alpha}{1 + \alpha t}. \quad (17)$$

Hence, applying Lemma 8 on the SDE $dX_t = dW_t$ gives us,

$$\frac{d}{dt} D_\Phi(\mu_t \parallel \nu_t) = -\frac{1}{2} \text{Fl}_\Phi(\mu_t \parallel \nu_t).$$

We can then apply the Φ -Sobolev inequality of ν_t and integrate the differential inequality from 0 to η to yield

$$\frac{D_\Phi(\mu_\eta \parallel \nu_\eta)}{D_\Phi(\mu_0 \parallel \nu_0)} \leq \exp\left(-\int_0^\eta \alpha_t dt\right) \leq \frac{1}{1 + \alpha\eta}.$$

Observe that $D_\Phi(\mu_\eta \parallel \nu_\eta) = D_\Phi(\mu \mathbf{P}^+ \parallel \nu^X \mathbf{P}^+)$ and that $D_\Phi(\mu_0 \parallel \nu_0) = D_\Phi(\mu \parallel \nu^X)$. As μ was arbitrary, this gives a valid bound on the contraction coefficient and concludes the proof. \blacksquare

E.2. Proof of Lemma 16

Proof The proof follows similarly to that of Lemma 15. Let μ be an arbitrary distribution such that $D_\Phi(\mu \parallel \nu^Y) < \infty$. Define μ_t to be the marginal law at time t when starting the SDE (15) from μ , and let ν_t be the marginal law at time t when starting the SDE (15) from ν^Y . Further denote α_t as shorthand for $\alpha_{\Phi\text{SI}}(\nu_t)$. As by construction of the SDE (15), $\nu_t = \nu^X * \mathcal{N}_{\eta-t}$, the same argument used to derive (17) yields that

$$\alpha_t \geq \frac{\alpha}{1 + \alpha(\eta - t)}.$$

Hence, applying Lemma 8 on the SDE (15) gives us,

$$\frac{d}{dt} D_\Phi(\mu_t \parallel \nu_t) = -\frac{1}{2} \text{Fl}_\Phi(\mu_t \parallel \nu_t).$$

We can then apply the Φ -Sobolev inequality of ν_t and integrate the differential inequality from 0 to η to yield

$$\frac{D_\Phi(\mu_\eta \parallel \nu_\eta)}{D_\Phi(\mu_0 \parallel \nu_0)} \leq \exp\left(-\int_0^\eta \alpha_t dt\right) \leq \frac{1}{1 + \alpha\eta}.$$

Observe that $D_\Phi(\mu_\eta \parallel \nu_\eta) = D_\Phi(\mu \mathbf{P}^- \parallel \nu^Y \mathbf{P}^-)$ and that $D_\Phi(\mu_0 \parallel \nu_0) = D_\Phi(\mu \parallel \nu^Y)$. As μ was arbitrary, this gives a valid bound on the contraction coefficient and concludes the proof. \blacksquare

Appendix G. Ornstein-Uhlenbeck Process

Recall the Langevin dynamics (5) to sample from $\nu \propto \exp(-f)$. When ν is a Gaussian, the Langevin dynamics is known as the Ornstein-Uhlenbeck process. For example, when $\nu = \mathcal{N}(0, \frac{1}{\alpha}I)$, the Langevin dynamics (5) becomes:

$$dX_t = -\alpha X_t dt + \sqrt{2} dW_t.$$

When the initial distribution ρ_0 is a Gaussian, the marginal distributions ρ_t admit convenient Gaussian forms. We will leverage this to show that Theorems 2 and 3 are tight for KL divergence.

Throughout, let $\nu = \mathcal{N}(0, \frac{1}{\alpha}I)$ for some $\alpha > 0$. Recall that this is α -strongly log-concave, and hence also satisfies a Φ -Sobolev inequality with the same constant (Lemma 7).

G.1. ULA evolution

The ULA update (6) for $\nu = \mathcal{N}(0, \frac{1}{\alpha}I)$ takes the following form:

$$X_{k+1} = (1 - \alpha\eta)X_k + \sqrt{2\eta}Z_k, \quad (18)$$

and the solution to (18) is:

$$X_k = (1 - \alpha\eta)^k X_0 + \sqrt{\frac{2(1 - (1 - \alpha\eta)^{2k})}{\alpha(2 - \alpha\eta)}} Z, \quad (19)$$

where $Z \sim \mathcal{N}(0, I)$. In this case, the biased limit ν^η is:

$$\nu^\eta = \mathcal{N}\left(0, \frac{2}{\alpha(2 - \alpha\eta)}I\right). \quad (20)$$

To show a tightness result, we simply take $\rho_0 = \mathcal{N}(\mathbf{1}, I)$ where $\mathbf{1}$ is the d -dimensional all-ones vector.

Proposition 17 *Consider the ULA (6) for $\nu = \mathcal{N}(0, \frac{1}{\alpha}I)$ with $\rho_0 = \mathcal{N}(\mathbf{1}, I)$ and $\eta \leq \frac{1}{\alpha}$. Then, $\text{KL}(\rho_k \parallel \nu^\eta) = \mathcal{O}(d\alpha(1 - \eta\alpha)^{2k})$.*

Proof From (19), we can see that $X_k \sim \rho_k$ where

$$\rho_k = \mathcal{N}\left((1 - \eta\alpha)^k \mathbf{1}, \frac{2 + (1 - \eta\alpha)^{2k}(2\alpha - \eta\alpha^2 - 2)}{\alpha(2 - \eta\alpha)}I\right). \quad (21)$$

Taking $k \rightarrow \infty$ in (21), we see that the biased limit is indeed given by (20), i.e.,

$$\nu^\eta = \mathcal{N}\left(0, \frac{2}{\alpha(2 - \alpha\eta)}I\right).$$

Therefore, using the formula for KL divergence between two multivariate Gaussians, we get

$$\text{KL}(\rho_k \parallel \nu^\eta) = \frac{d}{2} \left[\frac{(1 - \eta\alpha)^{2k}\alpha(2 - \eta\alpha)}{2} + \frac{(1 - \eta\alpha)^{2k}(2\alpha - \eta\alpha^2 - 2)}{2} - \log\left(1 + \frac{(1 - \eta\alpha)^{2k}(2\alpha - \eta\alpha^2 - 2)}{2}\right) \right].$$

This simplifies to reveal that

$$\text{KL}(\rho_k \parallel \nu^\eta) = \mathcal{O}(d\alpha(1 - \eta\alpha)^{2k}).$$

■

G.2. Proximal Sampler evolution

Similar to the ULA setting in Appendix G.1, we will take the target distribution to be $\nu^X = \mathcal{N}(0, \frac{1}{\alpha}I)$ and the starting distribution to be $\rho_0^X = \mathcal{N}(\mathbf{1}, I)$ where $\mathbf{1}$ is the d -dimensional all-ones vector.

The Proximal Sampler update (9) reveals that $\rho_0^Y = \mathcal{N}(\mathbf{1}, (1 + \eta)I)$. Moreover, we also have that for any $y \in \mathbb{R}^d$,

$$\nu^{X|Y}(\cdot | y) = \mathcal{N}\left(\frac{y}{1 + \eta\alpha}, \frac{\eta}{1 + \eta\alpha}I\right),$$

and consequently, as $\rho_1^X(x) = \int \nu^{X|Y}(x | y)\rho_0^Y(y) dy$, ρ_1^X is Gaussian. This reasoning extends to show that all of the iterates $\{\rho_k^X\}$ and $\{\rho_k^Y\}$ for $k \geq 0$ are Gaussian. Denoting $\rho_k^X = \mathcal{N}(m_k, c_k I)$, where for all $k \geq 0$, $m_k \in \mathbb{R}^d$ and $c_k > 0$, we have that,

$$m_{k+1} = \frac{m_k}{1 + \eta\alpha} \quad \text{and} \quad c_{k+1} = \frac{1}{(1 + \eta\alpha)^2} \left(c_k - \frac{1}{\alpha}\right) + \frac{1}{\alpha}.$$

Therefore, as $m_0 = \mathbf{1}$ and $c_0 = 1$,

$$m_k = \frac{1}{(1 + \eta\alpha)^k} \mathbf{1} \quad \text{and} \quad c_k = \frac{1}{(1 + \eta\alpha)^{2k}} \left(1 - \frac{1}{\alpha}\right) + \frac{1}{\alpha}. \quad (22)$$

We have the following proposition.

Proposition 18 *Consider the Proximal Sampler (9) for $\nu^X = \mathcal{N}(0, \frac{1}{\alpha}I)$ with $\rho_0^X = \mathcal{N}(\mathbf{1}, I)$ and $\eta > 0$. Then, $\text{KL}(\rho_k \parallel \nu^\eta) = \mathcal{O}(d\alpha(1 - \eta\alpha)^{-2k})$.*

Proof Recall from (22) that $\rho_k^X = \mathcal{N}(m_k, c_k I)$ with

$$m_k = \frac{1}{(1 + \eta\alpha)^k} \mathbf{1} \quad \text{and} \quad c_k = \frac{1}{(1 + \eta\alpha)^{2k}} \left(1 - \frac{1}{\alpha}\right) + \frac{1}{\alpha}.$$

Therefore, using the formula for KL divergence between two multivariate Gaussians, we get

$$\text{KL}(\rho_k^X \parallel \nu^X) = \frac{d}{2} \left[\frac{\alpha}{(1 + \eta\alpha)^{2k}} + \frac{\alpha - 1}{(1 + \eta\alpha)^{2k}} - \log \left(1 + \frac{\alpha - 1}{(1 + \eta\alpha)^{2k}}\right) \right].$$

This simplifies to reveal that

$$\text{KL}(\rho_k \parallel \nu^\eta) = \mathcal{O}(d\alpha(1 - \eta\alpha)^{-2k}).$$

■

Conclusion From Corollary 9, Proposition 17, and Proposition 18, we can see that the rates of convergence of KL divergence along the continuous time Langevin dynamics, ULA, and the Proximal Sampler are $(d\alpha \exp(-2\alpha t))$, $(d\alpha(1 - \alpha\eta)^{2k})$, and $(d\alpha(1 + \alpha\eta)^{-2k})$ respectively.