

---

# SVGD as a kernelized Wasserstein gradient flow of the chi-squared divergence

---

**Sinho Chewi**  
MIT  
schewi@mit.edu

**Thibaut Le Gouic**  
MIT  
tlegouic@mit.edu

**Chen Lu**  
MIT  
chenl819@mit.edu

**Tyler Maunu**  
MIT  
maunut@mit.edu

**Philippe Rigollet**  
MIT  
rigollet@mit.edu

## Abstract

Stein Variational Gradient Descent (SVGD), a popular sampling algorithm, is often described as the kernelized gradient flow for the Kullback-Leibler divergence in the geometry of optimal transport. We introduce a new perspective on SVGD that instead views SVGD as the (kernelized) gradient flow of the chi-squared divergence which, we show, exhibits a strong form of uniform exponential ergodicity under conditions as weak as a Poincaré inequality. This perspective leads us to propose an alternative to SVGD, called Laplacian Adjusted Wasserstein Gradient Descent (LAWGD), that can be implemented from the spectral decomposition of the Laplacian operator associated with the target density. We show that LAWGD exhibits strong convergence guarantees and good practical performance.

## 1 Introduction

The seminal paper of Jordan, Kinderlehrer, and Otto [JKO98] has profoundly reshaped our understanding of sampling algorithms. What is now commonly known as the *JKO scheme* interprets the evolution of marginal distributions of a Langevin diffusion as a gradient flow of a Kullback-Leibler (KL) divergence over the Wasserstein space of probability measures. This optimization perspective on Markov Chain Monte Carlo (MCMC) has not only renewed our understanding of algorithms based on Langevin diffusions [Dal17a; Ber18; CB18; Wib18; DMM19; VW19], but has also fueled the discovery of new MCMC algorithms inspired by the diverse and powerful optimization toolbox [Mar+12; Sim+16; Che+18; Ber18; Hsi+18; Wib18; Ma+19; Wib19; Che+20a; DR20; Zha+20b].

In order to arrive at a practical sampling algorithm, one must discretize the Wasserstein gradient flow. The most common discretization is to discretize the Langevin diffusion, resulting in the Unadjusted Langevin Algorithm (ULA) [Dal17b; DM17]. However, it is unclear whether this diffusion based discretization is the most effective one. In fact, ULA is asymptotically biased, which results in slow convergence and often requires ad-hoc adjustments [Dwi+19]. To overcome this limitation, various methods that track the Wasserstein gradient flow more closely have been recently developed [Ber18; Wib18; SKL20].

An alternative sampling approach that avoids diffusions is to construct a sequence of *deterministic* mappings that approximately pushes forward an initial distribution to the target distribution. Let  $F$  denote a functional over the Wasserstein space of distributions. The Wasserstein gradient flow of  $F$  may be described as the deterministic and time-inhomogeneous Markov process  $(X_t)_{t \geq 0}$  started at a random variable  $X_0 \sim \mu_0$  and evolving according to  $\dot{X}_t = -[\nabla_{W_2} F(\mu_t)](X_t)$ , where  $\mu_t$  denotes the distribution of  $X_t$ . Here  $[\nabla_{W_2} F(\mu)](\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^d$  is the Wasserstein gradient of  $F$  at  $\mu$ . If

$F(\mu) = D_{\text{KL}}(\mu \parallel \pi)$ , where  $\pi \propto e^{-V}$  is a given target distribution on  $\mathbb{R}^d$ , it is known [AGS08; Vil09; San17] that  $\nabla_{W_2} F(\mu) = \nabla \ln(d\mu/d\pi)$ . Therefore, a natural discretization of the Wasserstein gradient flow with step size  $h > 0$ , albeit one that cannot be implemented since it depends on the distribution  $\mu_t$  of  $X_t$ , is:

$$X_{t+1} = X_t - h \nabla \ln \left( \frac{d\mu_t}{d\pi}(X_t) \right), \quad t = 0, 1, 2, \dots$$

While  $\mu_t$  can, in principle, be estimated by evolving a large number of particles  $X_t^{[1]}, \dots, X_t^{[N]}$ , estimation of  $\mu_t$  is hindered by the curse of dimensionality and this approach still faces significant computational challenges despite attempts to improve the original JKO scheme [SKL20; WL20].

A major advance in this direction was achieved by allowing for *approximate* Wasserstein gradients, which makes the push forward maps tractable. More specifically, Stein Variational Gradient Descent (SVGD), recently proposed by [LW16] (see Section 2 for more details), consists in replacing  $\nabla_{W_2} F(\mu)$  by its image  $\mathcal{K}_\mu \nabla_{W_2} F(\mu)$  under the integral operator  $\mathcal{K}_\mu : L^2(\mu) \rightarrow L^2(\mu)$  associated to a chosen kernel  $K : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$  and defined by  $\mathcal{K}_\mu f(x) := \int K(x, y) f(y) d\mu(y)$  for  $f \in L^2(\mu)$ . This leads to the following process:

$$X_{t+1} = X_t - h[\mathcal{K}_{\mu_t} \nabla_{W_2} F(\mu_t)](X_t), \quad t = 0, 1, 2, \dots \quad (\text{SVGD}_p)$$

where we apply the integral operator  $\mathcal{K}_{\mu_t}$  individually to each coordinate of the Wasserstein gradient. In turn, this *kernelization trick* overcomes most of the above computational bottleneck. Building on this perspective, [DNS19] introduced a new geometry, different from the Wasserstein geometry and which they call the *Stein geometry*, in which the continuous limit of (SVGD<sub>p</sub>) becomes the gradient flow of the KL divergence.

However, despite this recent advance, the theoretical properties of SVGD are still largely unexplored, resulting in little understanding of SVGD’s known problems, such as mode collapse or a lack of guidance on how to choose an appropriate kernel  $K$ . Consequently diffusion-based algorithms remain the dominant choice for applications. In this work, we provide a new and stronger theoretical footing for the development of such deterministic mappings.

**Our contributions.** We introduce, in Section 2.3, a new perspective on SVGD by viewing it as kernelized gradient flow of the chi-squared divergence rather than the KL divergence. This perspective is fruitful in two ways. First, it uses a single integral operator  $\mathcal{K}_\pi$ —as opposed to (SVGD<sub>p</sub>), which requires a family of integral operators  $\mathcal{K}_\mu, \mu \ll \pi$ —providing a conceptually clear guideline for choosing  $K$ , namely:  $K$  should be chosen to make  $\mathcal{K}_\pi$  approximately equal to the identity operator. Second, under the idealized choice  $\mathcal{K}_\pi = \text{id}$ , we show that this gradient flow converges exponentially fast in KL divergence as soon as the target distribution  $\pi$  satisfies a Poincaré inequality. In fact, our results are stronger than exponential convergence and they highlight *strong uniform ergodicity*: the gradient flow forgets the initial distribution after a finite time that is at most half of the Poincaré constant. To establish this exponential convergence under a relatively weak condition (Poincaré inequality), we employ the following technique. While the gradient flow aims at minimizing the chi-squared divergence by following the curve in Wasserstein space with steepest descent, we do not track its progress with the objective function itself, the chi-squared divergence, but instead we track it with the KL divergence. This is in a sense dual to argument employed in [Che+20a], where the chi-squared divergence is used to track the progress of a gradient flow on the KL divergence. A more standard analysis relying on Łojasiewicz inequalities also yields rates of convergence on the chi-squared divergence under stronger assumptions such as a log-Sobolev inequality, and log-concavity. These results establish the first finite-time theoretical guarantees for SVGD in an idealized setting.




Figure 1: Sampling from a mixture of two 2D Gaussians with LAWGD. See Appendix C.

Beyond providing a better understanding of SVGD, our novel perspective is instrumental in the development of a new sampling algorithm, which we call Laplacian Adjusted Wasserstein Gradient Descent (LAWGD) and present in Section 4. We show that it possesses a striking theoretical property:

assuming that the target distribution  $\pi$  satisfies a Poincaré inequality, LAWGD converges exponentially fast, with *no* dependence on the Poincaré constant. This scale invariance has been recently demonstrated for the Newton-Langevin diffusion [Che+20a], but under the additional assumption that  $\pi$  is log-concave. A successful implementation of LAWGD hinges on the spectral decomposition of a certain differential operator which is within reach of modern PDE solvers. As a proof of concept, we show that LAWGD, implemented using a naïve finite differences method, performs well on mixtures of Gaussians in one or two dimensions, whereas SVGD fails. This is an indication that our novel perspective could be the correct one to further advance the state-of-the-art for sampling via deterministic mappings. Implementing LAWGD in high dimensions is challenging, and we are not advocating for it as the definitive solution of the sampling problem. Instead, LAWGD serves as the start of a family of interacting particle systems with an interacting potential that depends strongly and non-trivially on the target distribution and furthermore comes with strong theoretical guarantees. We hope this work can encourage further research in the application of numerical PDEs for sampling.

**Related work.** Since its introduction in [LW16], a number of variants of SVGD have been considered. They include a stochastic version [Li+20], a version that approximates the Newton direction in Wasserstein space [Det+18], a version that uses matrix kernels [Wan+19], an accelerated version [Liu+19], and a hybrid with Langevin [Zha+20a]. Several works have studied theoretical properties of SVGD, including its interpretation as a gradient flow under a modified geometry [Liu17; DNS19], and its asymptotic convergence [LLN19].

**Notation.** In this paper, all probability measures are assumed to have densities w.r.t. Lebesgue measure; therefore, we frequently abuse notation by identifying a probability measure with its Lebesgue density. For a differentiable kernel  $K : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ , we denote by  $\nabla_1 K : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d$  (resp.  $\nabla_2 K$ ) the gradient of the kernel w.r.t. the first (resp. second) argument. When describing particle algorithms, we use a subscript to denote the time index and brackets to denote the particle index, i.e.,  $X_t^{[i]}$  refers to the  $i$ th particle at time (or iteration number)  $t$ .

## 2 SVGD as a kernelized Wasserstein gradient flow

### 2.1 Wasserstein gradient flows

In this section, we review the theory of gradient flows on the space  $\mathcal{P}_{2,\text{ac}}(\mathbb{R}^d)$  of probability measures absolutely continuous w.r.t. Lebesgue measure and possessing a finite second moment, equipped with the 2-Wasserstein metric  $W_2$ . We refer readers to [Vil03; San15; San17] for introductory treatments of optimal transport, and to [AGS08; Vil09] for detailed treatments of Wasserstein gradient flows.

Let  $F : \mathcal{P}_{2,\text{ac}}(\mathbb{R}^d) \rightarrow \mathbb{R} \cup \{\infty\}$  be a functional defined on Wasserstein space. We say that a curve  $(\mu_t)_{t \geq 0}$  of probability measures is a *Wasserstein gradient flow* for the functional  $F$  if it satisfies

$$\partial_t \mu_t = \operatorname{div}(\mu_t \nabla_{W_2} F(\mu_t)) \quad (1)$$

in a weak sense. Here,  $\nabla_{W_2} F(\mu) := \nabla \delta F(\mu)$  is the Wasserstein gradient of the functional  $F$  at  $\mu$ , where  $\delta F(\mu) : \mathbb{R}^d \rightarrow \mathbb{R}$  is the *first variation* of  $F$  at  $\mu$ , defined by

$$\lim_{\varepsilon \rightarrow 0} \frac{F(\mu + \varepsilon \xi) - F(\mu)}{\varepsilon} = \int \delta F(\mu) d\xi, \quad \text{for all } \xi \text{ with } \int d\xi = 0,$$

and  $\nabla$  denotes the usual (Euclidean) gradient. Hence, the Wasserstein gradient, at each  $\mu \in \mathcal{P}_{2,\text{ac}}(\mathbb{R}^d)$ , is a map from  $\mathbb{R}^d$  to  $\mathbb{R}^d$ .

Using the *continuity equation*, we can give an Eulerian interpretation to the evolution equation (1) (see [San15, §4] and [AGS08, §8]). Given a family of vector fields  $(v_t)_{t \geq 0}$ , let  $(X_t)_{t \geq 0}$  be a curve in  $\mathbb{R}^d$  with random initial point  $X_0 \sim \mu_0$ , and such that  $(X_t)_{t \geq 0}$  is an integral curve of the vector fields  $(v_t)_{t \geq 0}$ , that is,  $\dot{X}_t = v_t(X_t)$ . If we let  $\mu_t$  denote the law of  $X_t$ , then  $(\mu_t)_{t \geq 0}$  evolves according to the *continuity equation*

$$\partial_t \mu_t = -\operatorname{div}(\mu_t v_t). \quad (2)$$

Comparing (1) and (2), we see that (1) describes the evolution of the marginal law  $(\mu_t)_{t \geq 0}$  of the curve  $(X_t)_{t \geq 0}$  with  $X_0 \sim \mu_0$  and  $\dot{X}_t = -[\nabla_{W_2} F(\mu_t)](X_t)$ .

Wasserstein calculus provides the following (formal) calculation rule: the Wasserstein gradient flow  $(\mu_t)_{t \geq 0}$  for the functional  $F$  dissipates  $F$  at the rate  $\partial_t F(\mu_t) = -\mathbb{E}_{\mu_t}[\|\nabla_{W_2} F(\mu_t)\|^2]$ . More generally, for a curve  $(\mu_t)_{t \geq 0}$  evolving according to the continuity equation (2), the time-derivative of  $F$  is given by  $\partial_t F(\mu_t) = \mathbb{E}_{\mu_t} \langle \nabla_{W_2} F(\mu_t), v_t \rangle$ .

In this paper, we are primarily concerned with two functionals: the Kullback-Leibler (KL) divergence  $D_{\text{KL}}(\cdot \parallel \pi)$ , and the chi-squared divergence  $\chi^2(\cdot \parallel \pi)$  (see, e.g., [Tsy09]). It is a standard exercise [AGS08; San15] to check that Wasserstein gradients of these functionals are, respectively,

$$(\nabla_{W_2} D_{\text{KL}}(\cdot \parallel \pi))(\mu) = \nabla \ln \frac{d\mu}{d\pi}, \quad (\nabla_{W_2} \chi^2(\cdot \parallel \pi))(\mu) = 2\nabla \frac{d\mu}{d\pi}. \quad (3)$$

## 2.2 SVGD as a kernelized gradient flow of the KL divergence

SVG $D^1$  is achieved by replacing the Wasserstein gradient  $\nabla \ln(d\mu_t/d\pi)$  of the KL divergence with  $\mathcal{K}_{\mu_t} \nabla \ln(d\mu_t/d\pi)$ , leading to the particle evolution equation (SVG $D_p$ ).

Recalling that  $\pi \propto e^{-V}$ , we get

$$\mathcal{K}_{\mu_t} \nabla \ln \frac{d\mu_t}{d\pi}(x) := \int K(x, \cdot) \nabla \ln \frac{d\mu_t}{d\pi} d\mu_t = \int K(x, \cdot) \nabla V d\mu_t - \int \nabla_2 K(x, \cdot) d\mu_t, \quad (4)$$

where, in the second identity, we used integration by parts. This expression shows that rather than having to estimate the distribution  $\mu_t$ , it is sufficient to estimate the expectation  $\int \nabla_2 K(x, \cdot) d\mu_t$ . This is the key to the computational tractability of SVG $D$ . Indeed, the kernelized gradient flow can be implemented by drawing  $N$  particles  $X_0^{[1]}, \dots, X_0^{[N]} \stackrel{\text{i.i.d.}}{\sim} \mu_0$  and following the coupled dynamics

$$\dot{X}_t^{[i]} = -\mathcal{K}_{\mu_t} \nabla \ln \frac{d\mu_t}{d\pi}(X_t^{[i]}) = - \int K(X_t^{[i]}, \cdot) \nabla V d\mu_t + \int \nabla_2 K(X_t^{[i]}, \cdot) d\mu_t, \quad i \in [N].$$

With this, we can simply estimate the expectation with respect to  $\mu_t$  with an average over all particles, which yields the SVG $D$  algorithm:

$$X_{t+1}^{[i]} = X_t^{[i]} - \frac{h}{N} \sum_{j=1}^N K(X_t^{[i]}, X_t^{[j]}) \nabla V(X_t^{[j]}) + \frac{h}{N} \sum_{j=1}^N \nabla_2 K(X_t^{[i]}, X_t^{[j]}), \quad i \in [N]. \quad (5)$$

## 2.3 SVGD as a kernelized gradient flow of the chi-squared divergence

Recall from Section 2.1 that by the continuity equation, the continuous limit of the particle evolution equation (SVG $D_p$ ) translates into the following PDE that describes the evolution of the distribution  $\mu_t$  of  $X_t$ :

$$\partial_t \mu_t = \text{div}(\mu_t \mathcal{K}_{\mu_t} \nabla \ln \frac{d\mu_t}{d\pi}). \quad (\text{SVG}D_d)$$

We make the simple observation that

$$\mathcal{K}_{\mu_t} \nabla \ln \frac{d\mu_t}{d\pi}(x) = \int K(x, y) \nabla \ln \frac{d\mu_t}{d\pi}(y) d\mu_t(y) = \int K(x, y) \nabla \frac{d\mu_t}{d\pi}(y) d\pi(y) = \mathcal{K}_{\pi} \nabla \frac{d\mu_t}{d\pi}(x).$$

Thus, the continuous-dynamics of SVG $D$ , as given in (SVG $D_d$ ), can equivalently be expressed as

$$\partial_t \mu_t = \text{div}(\mu_t \mathcal{K}_{\pi} \nabla \frac{d\mu_t}{d\pi}). \quad (\text{SVG}D)$$

To interpret this equation, we recall that the Wasserstein gradient of the chi-squared divergence  $\chi^2(\cdot \parallel \pi)$  at  $\mu$  is  $2\nabla(d\mu/d\pi)$  (by (3)), so the gradient flow for the chi-squared divergence is

$$\partial_t \mu_t = 2 \text{div}(\mu_t \nabla \frac{d\mu_t}{d\pi}). \quad (\text{CSF})$$

Comparing (SVG $D$ ) and (CSF), we see that (up to a factor of 2), SVG $D$  can be understood as the flow obtained by replacing the gradient of the chi-squared divergence,  $\nabla(d\mu/d\pi)$ , by  $\mathcal{K}_{\pi} \nabla(d\mu/d\pi)$ .

Although (SVG $D_d$ ) and (SVG $D$ ) are equivalent ways of expressing the same dynamics, the formulation of (SVG $D$ ) presents a significant advantage: it involves a kernel integral operator  $\mathcal{K}_{\pi}$  that does not change with time and depends only on the target distribution  $\pi$ .

<sup>1</sup>Throughout this paper, we call SVG $D$  the generalization of the original method of [LW16; Liu17] that was introduced in [Wan+19].

### 3 Chi-squared gradient flow

In this section, study the idealized case where  $\mathcal{K}_\pi$  taken to be the identity operator. In this case, (SVGD) reduces to the gradient flow CSF. The existence, uniqueness, and regularity of this flow are studied in [OT11; OT13] and [AGS08, Theorem 11.2.1].

The rate of convergence of the gradient flow of the KL divergence is closely related to two functional inequalities: the Poincaré inequality controls the rate of exponential convergence in chi-squared divergence ([Pav14, Theorem 4.4], [Che+20a]) while a log-Sobolev inequality characterizes the rate of exponential convergence of the KL divergence [BGL14, Theorem 5.2.1]. In this section, we show that these inequalities also guarantee exponential rates of convergence of CSF.

Recall that  $\pi$  satisfies a *Poincaré inequality* with constant  $C_P$  if

$$\mathrm{var}_\pi f \leq C_P \mathbb{E}_\pi[\|\nabla f\|^2], \quad \text{for all locally Lipschitz } f \in L^2(\pi), \quad (\text{P})$$

while  $\pi$  satisfies a *log-Sobolev inequality* with constant  $C_{\text{LSI}}$

$$\mathrm{ent}_\pi(f^2) := \mathbb{E}_\pi[f^2 \ln(f^2)] - \mathbb{E}_\pi[f^2] \ln \mathbb{E}_\pi[f^2] \leq 2C_{\text{LSI}} \mathbb{E}_\pi[\|\nabla f\|^2] \quad (\text{LSI})$$

for all locally Lipschitz  $f$  for which  $\mathrm{ent}_\pi(f^2) < \infty$ .

We briefly review some facts regarding the strength of these assumptions. It is standard that the log-Sobolev inequality is stronger than the Poincaré inequality: (LSI) implies (P) with constant  $C_P \leq C_{\text{LSI}}$ . In turn, if  $\pi$  is  $\alpha$ -strongly log-concave, i.e.  $\nabla^2 V \succeq \alpha I_d$ , then it implies the validity of (LSI) with  $C_{\text{LSI}} \leq 1/\alpha$ , and thus a Poincaré inequality holds too. However, a Poincaré inequality is in general much weaker than strong log-concavity. For instance, if  $\lambda_\pi^2$  denotes the largest eigenvalue of the covariance matrix of  $\pi$ , then it is currently known that  $\pi$  satisfies a Poincaré inequality as soon as it is log-concave, with  $C_P \leq C(d)\lambda_\pi^2$ , where  $C(d)$  is a dimensional constant [Bob99; AB15; LV17], and the well-known *Kannan-Lovász-Simonovitz* (KLS) conjecture [KLS95] asserts that  $C(d)$  does not actually depend on the dimension.

Our first result shows that a Poincaré inequality suffices to establish exponential decay of the KL divergence along CSF. In fact, we establish a remarkable property, which we call *strong uniform ergodicity*: under a Poincaré inequality, CSF forgets its initial distribution after a time of no more than  $C_P/2$ . Uniform ergodicity is central in the theory of Markov processes [MT09, Ch. 16] but is often limited to compact state spaces. Moreover, this theory largely focuses on total variation, so the distance from the initial distribution to the target distribution is trivially bounded by 1.

**Theorem 1.** *Assume that  $\pi$  satisfies a Poincaré inequality (P) with constant  $C_P > 0$  and let  $(\mu_t)_{t \geq 0}$  denote the law of CSF. Assume that  $\chi^2(\mu_0 \parallel \pi) < \infty$ . Then,*

$$D_{\text{KL}}(\mu_t \parallel \pi) \leq D_{\text{KL}}(\mu_0 \parallel \pi) e^{-\frac{2t}{C_P}}, \quad \forall t \geq 0. \quad (6)$$

*In fact, a stronger convergence result holds:*

$$D_{\text{KL}}(\mu_t \parallel \pi) \leq (D_{\text{KL}}(\mu_0 \parallel \pi) \wedge 2) e^{-\frac{2t}{C_P}}, \quad \forall t \geq \frac{C_P}{2}. \quad (7)$$

*Proof.* Given the Wasserstein gradients (3) in Section 2.1, we get that  $(\mu_t)_{t \geq 0}$  satisfies

$$\partial_t D_{\text{KL}}(\mu_t \parallel \pi) = -2 \mathbb{E}_{\mu_t} \left\langle \nabla \ln \frac{d\mu_t}{d\pi}, \nabla \frac{d\mu_t}{d\pi} \right\rangle = -2 \mathbb{E}_\pi \left[ \left\| \nabla \frac{d\mu_t}{d\pi} \right\|^2 \right].$$

Applying the Poincaré inequality (P) with  $f = d\mu_t/d\pi - 1$ , we get

$$\partial_t D_{\text{KL}}(\mu_t \parallel \pi) \leq -\frac{2}{C_P} \chi^2(\mu_t \parallel \pi) \leq -\frac{2}{C_P} D_{\text{KL}}(\mu_t \parallel \pi),$$

where, in the last inequality, we use the fact that  $D_{\text{KL}}(\cdot \parallel \pi) \leq \chi^2(\cdot \parallel \pi)$  (see [Tsy09, §2.4]). The bound (6) follows by applying Grönwall's inequality.

To prove (7), we use the stronger inequality  $D_{\text{KL}}(\cdot \parallel \pi) \leq \ln[1 + \chi^2(\cdot \parallel \pi)]$  (see [Tsy09, §2.4]). Our differential inequality now reads:

$$\partial_t D_{\text{KL}}(\mu_t \parallel \pi) \leq -\frac{2}{C_P} (e^{D_{\text{KL}}(\mu_t \parallel \pi)} - 1) \iff \partial_t \psi(D_{\text{KL}}(\mu_t \parallel \pi)) \leq -\frac{2}{C_P} \psi(D_{\text{KL}}(\mu_t \parallel \pi)),$$

where  $\psi(x) = 1 - e^{-x} \leq 1$ . Grönwall's inequality now yields

$$\psi(D_{\text{KL}}(\mu_t \parallel \pi)) \leq e^{-\frac{2t}{C_P}} \psi(D_{\text{KL}}(\mu_0 \parallel \pi)) \leq e^{-\frac{2t}{C_P}}.$$

Note that  $x \leq 2\psi(x)$  whenever  $\psi(x) \leq 1/e$ . Thus, if  $t \geq C_P/2$ , we get  $\psi(D_{\text{KL}}(\mu_t \parallel \pi)) \leq e^{-1}$  so

$$D_{\text{KL}}(\mu_t \parallel \pi) \leq 2\psi(D_{\text{KL}}(\mu_t \parallel \pi)) \leq e^{-\frac{2t}{C_P}},$$

which, together with (6), completes the proof of (7).  $\square$

*Remark 1.* In [Che+20a], it was observed that the chi-squared divergence decays exponentially fast along the gradient flow  $(\mu_t)_{t \geq 0}$  for the KL divergence, provided that  $\pi$  satisfies a Poincaré inequality. This observation is made precise and more general in [MMS09] where it is noted that the *gradient flow of a functional  $\mathfrak{L}$  dissipates a different functional  $\mathfrak{V}$  at the same rate that the gradient flow of  $\mathfrak{V}$  dissipates the functional  $\mathfrak{L}$* . A similar method is used to study the thin film equation in [CT02] and [Car11, §5].

Since we are studying the gradient flow of the chi-squared divergence, it is natural to ask whether CSF converges to  $\pi$  in chi-squared divergence as well. In the next results, we show quantitative decay of the chi-squared divergence along the gradient flow under a Poincaré inequality (P), but we obtain only a polynomial rate of decay. However, if we additionally assume either that  $\pi$  is log-concave or that it satisfies a log-Sobolev inequality (LSI), then we obtain exponential decay of the chi-squared divergence along CSF.

**Theorem 2.** *Suppose that  $\pi$  satisfies a Poincaré inequality (P). Then, provided  $\chi^2(\mu_0 \parallel \pi) < \infty$ , the law  $(\mu_t)_{t \geq 0}$  of CSF satisfies*

$$\chi^2(\mu_t \parallel \pi) \leq \chi^2(\mu_0 \parallel \pi) \wedge \left(\frac{9C_P}{8t}\right)^2.$$

*If we further assume that  $\pi$  is log-concave, then*

$$\chi^2(\mu_t \parallel \pi) \leq \chi^2(\mu_0 \parallel \pi) e^{-\frac{t}{2C_P}}.$$

*Proof.* The proof is deferred to Appendix B.  $\square$

Under the stronger assumption (LSI), we can show strong uniform ergodicity as in Theorem 1.

**Theorem 3.** *Assume that  $\pi$  satisfies a log-Sobolev inequality (LSI). Let  $(\mu_t)_{t \geq 0}$  denote the law of CSF, and assume that  $\chi^2(\mu_0 \parallel \pi) < \infty$ . Then, for all  $t \geq 7C_{\text{LSI}}$ ,*

$$\chi^2(\mu_t \parallel \pi) \leq (\chi^2(\mu_0 \parallel \pi) \wedge 2) e^{-\frac{t}{9C_{\text{LSI}}}}.$$

*Proof.* The proof is deferred to Appendix B.  $\square$

Convergence in chi-squared divergence was studied in recent works such as [CLL19; VW19; Che+20a]. From standard comparisons between information divergences (see [Tsy09, §2.4]), it implies convergence in total variation distance, Hellinger distance, and KL divergence. Moreover, recent works have shown that the Poincaré inequality (P) yields transportation-cost inequalities which bound the 2-Wasserstein distance by powers of the chi-squared divergence [Din15; Led18; Che+20a; Liu20], so we obtain convergence in the 2-Wasserstein distance as well. In particular, we mention that [Che+20a] uses the chi-squared gradient flow (CSF) to prove a transportation-cost inequality.

## 4 Laplacian Adjusted Wasserstein Gradient Descent (LAWGD)

While the previous section leads to a better understanding of the convergence properties of SVGD in the case that  $\mathcal{K}_\pi$  is the identity operator, it is still unclear how to choose the kernel  $K$  to approach this idealized setup. For SVGD with a general kernel  $K$ , the calculation rules of Section 2.1 together with the method of the previous section yield the formula

$$\partial_t D_{\text{KL}}(\mu_t \parallel \pi) = -\mathbb{E}_\pi \left\langle \nabla \frac{d\mu_t}{d\pi}, \mathcal{K}_\pi \nabla \frac{d\mu_t}{d\pi} \right\rangle,$$

for the dissipation of the KL divergence along SVGD. From this, a natural way to proceed is to seek an inequality of the form

$$\mathbb{E}_\pi \langle f, \mathcal{K}_\pi f \rangle \gtrsim \mathbb{E}_\pi [f^2], \quad \text{for all locally Lipschitz } f \in L^2(\pi). \quad (8)$$

Applying this inequality to each coordinate of  $\nabla(\mathrm{d}\mu_t/\mathrm{d}\pi)$  separately and using a Poincaré inequality would then allow us to conclude as in the proof of Theorem 1. The inequality (8) can be interpreted as a positive lower bound on the smallest eigenvalue of the operator  $\mathcal{K}_\pi$ . However, this approach is doomed to fail; under mild conditions on the kernel  $K$ , it is a standard fact that the eigenvalues of  $\mathcal{K}_\pi$  form a sequence converging to 0, so no such spectral gap can hold.<sup>2</sup>

This suggests that any approach which seeks to prove finite-time convergence results for SVGD in the spirit of Theorem 1 must exploit finer properties of the eigenspaces of the operator  $\mathcal{K}_\pi$ . Motivated by this observation, we develop a new algorithm called Laplacian Adjusted Wasserstein Gradient Descent (LAWGD) in which the kernel  $K$  is chosen carefully so that  $\mathcal{K}_\pi = \mathcal{L}^{-1}$  is the inverse of the generator of the Langevin diffusion that has  $\pi$  as invariant measure.

More precisely, the starting point for our approach is the following integration-by-parts formula, which is a crucial component of the theory of Markov semigroups [BGL14]:

$$\mathbb{E}_\pi \langle \nabla f, \nabla g \rangle = \mathbb{E}_\pi [f \mathcal{L} g], \quad \text{for all locally Lipschitz } f, g \in L^2(\pi), \quad (9)$$

where  $\mathcal{L} := -\Delta + \langle \nabla V, \nabla \cdot \rangle$ . The operator  $\mathcal{L}$  is the (negative) generator of the standard Langevin diffusion with stationary distribution  $\pi$  [Pav14, §4.5]. We refer readers to Appendix A for background on the spectral theory of  $\mathcal{L}$ .

In order to use (9), we replace the vector field  $-\mathcal{K}_\pi \nabla(\mathrm{d}\mu_t/\mathrm{d}\pi)$  by the vector field  $-\nabla \mathcal{K}_\pi(\mathrm{d}\mu_t/\mathrm{d}\pi)$ . The new dynamics follow the evolution equation

$$\partial_t \mu_t = \operatorname{div} \left( \mu_t \nabla \mathcal{K}_\pi \frac{\mathrm{d}\mu_t}{\mathrm{d}\pi} \right). \quad (\text{LAWGD})$$

The vector field in the above continuity equation may also be written

$$-\nabla \mathcal{K}_\pi \frac{\mathrm{d}\mu_t}{\mathrm{d}\pi}(x) = - \int \nabla_1 K(x, \cdot) \frac{\mathrm{d}\mu_t}{\mathrm{d}\pi} \mathrm{d}\pi = - \int \nabla_1 K(x, \cdot) \mathrm{d}\mu_t.$$

Replacing  $\mu_t$  by an empirical average over particles and discretizing the process in time, we again obtain an implementable algorithm, which we give as Algorithm 1.

---

#### Algorithm 1 LAWGD

---

```

1: procedure LAWGD( $K, \mathcal{L}, \mu_0$ )
2:   draw  $N$  particles  $X_0^{[1]}, \dots, X_0^{[N]} \stackrel{\text{i.i.d.}}{\sim} \mu_0$ 
3:   for  $t = 1, \dots, T - 1$  do
4:     for  $i = 1, \dots, N$  do
5:        $X_{t+1}^{[i]} \leftarrow X_t^{[i]} - \frac{h}{N} \sum_{j=1}^N \nabla_1 K_{\mathcal{L}}(X_t^{[i]}, X_t^{[j]})$ 
6:     end for
7:   end for
8:   return  $X_T^{[1]}, \dots, X_T^{[N]}$ 
9: end procedure

```

---

A careful inspection of Algorithm 1 reveals that the update equation for the particles in Algorithm 1 does not involve the potential  $V$  directly, unlike the SVGD algorithm (5); thus, the kernel for LAWGD must contain all the information about  $V$ .

Our choice for the kernel  $K$  is guided by the following observation (based on (9)):

$$\partial_t D_{\text{KL}}(\mu_t \parallel \pi) = - \mathbb{E}_\pi \left\langle \nabla \frac{\mathrm{d}\mu_t}{\mathrm{d}\pi}, \nabla \mathcal{K}_\pi \frac{\mathrm{d}\mu_t}{\mathrm{d}\pi} \right\rangle = - \mathbb{E}_\pi \left[ \frac{\mathrm{d}\mu_t}{\mathrm{d}\pi} \mathcal{L} \mathcal{K}_\pi \frac{\mathrm{d}\mu_t}{\mathrm{d}\pi} \right].$$

<sup>2</sup>It is enough that  $K$  is a symmetric kernel with  $K \in L^2(\pi \otimes \pi)$ , and that  $\pi$  is not discrete (so that  $L^2(\pi)$  is infinite-dimensional); see [BGL14, Appendix A.6].

As a result, we choose  $K$  to ensure that  $\mathcal{K}_\pi = \mathcal{L}^{-1}$ . This choice yields

$$\partial_t D_{\text{KL}}(\mu_t \parallel \pi) = -\mathbb{E}_\pi \left[ \left( \frac{d\mu_t}{d\pi} - 1 \right)^2 \right] = -\chi^2(\mu_t \parallel \pi). \quad (10)$$

It remains to see which kernel  $K$  implements  $\mathcal{K}_\pi = \mathcal{L}^{-1}$ . To that end, assume that  $\mathcal{L}$  has a discrete spectrum and let  $(\lambda_i, \phi_i)$ ,  $i = 0, 1, 2, \dots$  be its eigenvalue-eigenfunction pairs where  $\lambda_i$ s are arranged in nondecreasing order. Assume further that  $\lambda_1 > 0$  (which amounts to a Poincaré inequality; see Appendix A) and define the following *spectral kernel*:

$$\mathcal{K}_\mathcal{L}(x, y) = \sum_{i=1}^{\infty} \frac{\phi_i(x)\phi_i(y)}{\lambda_i} \quad (11)$$

We now show that this choice of kernel endows LAWGD with a remarkable property: it converges to the target distribution exponentially fast, with a rate which has no dependence on the Poincaré constant. Moreover, akin to CSF—see (7)—it also exhibits strong uniform ergodicity.

**Theorem 4.** *Assume that  $\mathcal{L}$  has a discrete spectrum and that  $\pi$  satisfies a Poincaré inequality (P) with some finite constant. Let  $(\mu_t)_{t \geq 0}$  be the law of LAWGD with the kernel described above. Then,*

$$D_{\text{KL}}(\mu_t \parallel \pi) \leq (D_{\text{KL}}(\mu_0 \parallel \pi) \wedge 2) e^{-t}, \quad \forall t \geq 1.$$

*Proof.* In light of (10), the proof is identical to that of Theorem 1. □

The convergence rate in Theorem 4 has no dependence on the target measure. This *scale-invariant convergence* also appears in [Che+20a], where it is shown for the Newton-Langevin diffusion with a strictly log-concave target measure  $\pi$ . In Theorem 4, we obtain similar guarantees under the much weaker assumption of a Poincaré inequality; indeed, there are many examples of non-log-concave distributions which satisfy a Poincaré inequality [VW19].

## 5 Experiments

To implement Algorithm 1, we numerically approximate the kernel  $K = \mathcal{K}_\mathcal{L}$  given in (11). When  $\pi$  is the standard Gaussian distribution on  $\mathbb{R}$ , the eigendecomposition of the operator  $\mathcal{L}$  in (9) is known explicitly in terms of the Hermite polynomials [BGL14, §2.7.1], and we approximate the kernel via a truncated sum:  $\hat{K}(x, y) = \sum_{i=1}^k \lambda_i^{-1} \phi_i(x)\phi_i(y)$  (Figure 2) involving the smallest eigenvalues of  $\mathcal{L}$ .

In the general case, we implement a basic finite difference (FD) method to approximate the eigenvalues and eigenfunctions of  $\mathcal{L}$ . We obtain better numerical results by first transforming the operator  $\mathcal{L}$  into the Schrödinger operator  $\mathcal{L}_S := -\Delta + V_S$ , where  $V_S := \frac{1}{4}\|\nabla V\|^2 - \frac{1}{2}\Delta V$ . If  $\phi_S$  is an eigenfunction of  $\mathcal{L}_S$  with eigenvalue  $\lambda$  (normalized such that  $\int \phi_S^2 = 1$ ), then  $\phi := e^{V/2} \phi_S$  is an eigenfunction of  $\mathcal{L}$  also with eigenvalue  $\lambda$  (and normalized such that  $\int \phi^2 d\pi = 1$ ); see [BGL14, §1.15.7].

On a grid of points (with spacing  $\varepsilon$ ), if we replace the Laplacian with the FD operator  $\Delta_\varepsilon f(x) := \{f(x - \varepsilon) + f(x + \varepsilon) - 2f(x)\}/\varepsilon^2$  (in 1D), then the FD Schrödinger operator  $\mathcal{L}_{S,\varepsilon} := -\Delta_\varepsilon + V_S$  can be represented as a sparse matrix, and its eigenvalues and (unit) eigenvectors are found with standard linear algebra solvers.

When the potential  $V$  is known only up to an additive constant, then the approximate eigenfunctions produced by this method are not normalized correctly; instead, they satisfy  $\|\phi\|_{L^2(\pi)} = C$  for some constant  $C$  (which is the same for each eigenfunction). In turn, this causes the kernel  $K$  in LAWGD to be off by a multiplicative constant. For implementation purposes, however, this constant is absorbed in the step size of Algorithm 1. We also note that the eigenfunctions are differentiated using a FD approximation.

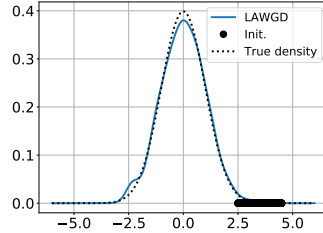


Figure 2: Samples from the standard Gaussian distribution generated by LAWGD, with kernel approximated by Hermite polynomials. For details, see Appendix C.



To demonstrate, we sample from a mixture of three Gaussians:  $\frac{2}{5}\mathcal{N}(-3, 1) + \frac{1}{5}\mathcal{N}(0, 1) + \frac{2}{5}\mathcal{N}(4, 2)$ . We compare LAWGD with SVGD using the RBF kernel and median-based bandwidth as in [LW16]. We approximate the eigenfunctions and eigenvalues using a finite difference scheme, on 256 grid points evenly spaced between  $-14$  and  $14$ . Constant step sizes for LAWGD and SVGD are tuned and the algorithms are run for 5000 iterations, and the samples are initialized to be uniform on  $[1, 4]$ . The results are displayed in Figure 3. All 256 discrete eigenfunctions and eigenvalues are used.




Figure 3: LAWGD and SVGD run with constant step size for a mixture of three Gaussians. Both kernel density estimators use the same bandwidth.

## 6 Open questions

We conclude this paper with some interesting open questions. The introduction of the chi-squared divergence as an objective function allows us to obtain both theoretical insights about SVGD and a new algorithm, LAWGD. This perspective opens the possibility of identifying other functionals defined over Wasserstein space and that yield gradient flows which are amenable to mathematical analysis and efficient computation. Towards this goal, an intriguing direction is to develop alternative methods, besides kernelization, which provide effective implementations of Wasserstein gradient flows. Finally, we note that LAWGD provides a hitherto unexplored connection between sampling and computing the spectral decomposition of the Schrödinger operator, the latter of which has been intensively studied in numerical PDEs. We hope our work further stimulates research at the intersection of these communities.

## A Review of spectral theory

In this paper, we consider elliptic differential operators of the form  $\mathcal{L} = -\Delta + \langle \nabla V, \nabla \cdot \rangle$ , where  $V$  is a continuously differentiable potential. In this section, we provide a brief review of the spectral theory of these operators, and we refer to [Eva10, §6.5] for a standard treatment.

The operator  $\mathcal{L}$  (when suitably interpreted) is a linear operator defined on a domain  $\mathcal{D} \subset L^2(\pi)$ . For any locally Lipschitz function  $f \in L^2(\pi)$ , integration by parts shows that

$$\mathbb{E}_\pi[f \mathcal{L} f] = \mathbb{E}_\pi[\|\nabla f\|^2].$$

Therefore,  $\mathcal{L}$  has a non-negative spectrum. Also, we have  $\mathcal{L}1 = 0$ , so that 0 is always an eigenvalue of  $\mathcal{L}$ . We say that  $\mathcal{L}$  has a *discrete spectrum* if it has a countable sequence of eigenvalues  $0 = \lambda_0 \leq \lambda_1 \leq \lambda_2 \leq \lambda_3 \leq \dots$  and corresponding eigenfunctions  $(\phi_i)_{i=1}^\infty$  which form a basis of  $\mathcal{D}$ . The eigenfunctions can be chosen to be orthogonal and normalized such that  $\|\phi_i\|_{L^2(\pi)} = 1$ ; we always assume this is the case. Then,  $\mathcal{L}$  can be expressed as

$$\mathcal{L} = \sum_{i=1}^{\infty} \lambda_i \langle \phi_i, \cdot \rangle_{L^2(\pi)} \phi_i.$$

The operator  $\mathcal{L}$  has a discrete spectrum under the following condition ([Fri34], [RS78, Theorem XIII.67], [BGL14, Corollary 4.10.9]):

$$V_S \in L^1_{\text{loc}}(\mathbb{R}^d), \quad \inf V_S > -\infty, \quad \text{and} \quad \lim_{\|x\| \rightarrow \infty} V_S(x) = +\infty,$$

where  $V_S := -\Delta V + \frac{1}{2}\|\nabla V\|^2$ . Moreover, under this condition we also have  $\lambda_i \rightarrow \infty$  as  $i \rightarrow +\infty$ . For example, this condition is satisfied for  $V(x) = \|x\|^\alpha$  for  $\alpha > 1$ , but not for  $\alpha = 1$ . In fact, for  $\alpha = 1$ , the spectrum of  $\mathcal{L}$  is not discrete [BGL14, §4.1.1].

The *Poincaré inequality* (P) is interpreted as a *spectral gap inequality*, since it asserts that  $\lambda_1 = 1/C_P > 0$ . Thus, under a Poincaré inequality,  $\mathcal{L} : \mathcal{D} \cap \{f \in L^2(\pi) \mid \mathbb{E}_\pi f = 0\} \rightarrow L^2(\pi)$  is bijective. Moreover, if it has a discrete spectrum, its inverse satisfies

$$\mathcal{L}^{-1} = \sum_{i=1}^{\infty} \lambda_i^{-1} \langle \phi_i, \cdot \rangle_{L^2(\pi)} \phi_i.$$

## B Proofs of the convergence guarantees for the chi-squared flow

In this section, we give proofs of the convergence results we stated in Section 3.

*Proof of Theorem 2 (non-log-concave case).* According to [Che+20a, Proposition 1], the Poincaré inequality implies the following inequality for the chi-squared divergence:

$$\chi^2(\mu \parallel \pi)^{3/2} \leq \frac{9C_P}{4} \mathbb{E}_\mu \left[ \left\| \nabla \frac{d\mu}{d\pi} \right\|^2 \right], \quad \forall \mu \ll \pi. \quad (12)$$

Since the Wasserstein gradient of  $\chi^2(\cdot \parallel \pi)$  at  $\mu$  is given by  $2\nabla(d\mu/d\pi)$  (see Section 2.1), it yields

$$\partial_t \chi^2(\mu_t \parallel \pi) = -4 \mathbb{E}_{\mu_t} \left[ \left\| \nabla \frac{d\mu_t}{d\pi} \right\|^2 \right] \leq -\frac{16}{9C_P} \chi^2(\mu_t \parallel \pi)^{3/2}$$

Solving the above differential inequality yields

$$\chi^2(\mu_t \parallel \pi) \leq \frac{\chi^2(\mu_0 \parallel \pi)}{\{1 + 8t\sqrt{\chi^2(\mu_0 \parallel \pi)}/(9C_P)\}^2},$$

which implies the desired result.  $\square$

We now prepare for the proof of exponentially fast convergence in chi-squared divergence for log-concave measures. The key to proving such results lies in differential inequalities of the form

$$\chi^2(\mu \parallel \pi) \leq C_{PL} \mathbb{E}_\mu \left[ \left\| \nabla \frac{d\mu}{d\pi} \right\|^2 \right], \quad \forall \mu \ll \pi, \quad (13)$$

which may be interpreted as a *Polyak-Łojasiewicz (PL) inequality* [KNS16] for the functional  $\chi^2(\cdot \parallel \pi)$ . PL inequalities are well-known in the optimization literature, and can be even used when the objective is not convex [Che+20b]. In contrast, the preceding proof uses the weaker inequality (12), which may be interpreted as a *Łojasiewicz inequality* [Loj63].

To see that a PL inequality readily yields exponential convergence, observe that

$$\partial_t \chi^2(\mu_t \parallel \pi) = -4 \mathbb{E}_{\mu_t} \left[ \left\| \nabla \frac{d\mu_t}{d\pi} \right\|^2 \right] \leq -\frac{4}{C_{PL}} \chi^2(\mu_t \parallel \pi).$$

Together with Grönwall's inequality, the differential inequality yields  $\chi^2(\mu_t \parallel \pi) \leq \chi^2(\mu_0 \parallel \pi) e^{-\frac{4t}{C_{PL}}}$ .

In order to prove a PL inequality of the type (13), we require two ingredients. The first one is a transportation-cost inequality for the chi-squared divergence proven in [Liu20], building on the works [Din15; Led18]. It asserts that if  $\pi$  satisfies a Poincaré inequality (P), then the following inequality holds:

$$W_2^2(\mu, \pi) \leq 2C_P \chi^2(\mu \parallel \pi), \quad \forall \mu \ll \pi. \quad (14)$$

For the second ingredient, we use an argument of [OV00] to show that if  $\pi$  satisfies a chi-squared transportation-cost inequality such as (14), and in addition is log-concave, then it satisfies an inequality of the type (13). We remark that the converse statement, that is, if  $\pi$  satisfies a PL inequality (13) then it satisfies an appropriate chi-squared transportation-cost inequality, was proven in [Che+20a] without the additional assumption of log-concavity. It implies that for log-concave distributions, the PL inequality (13) and the chi-squared transportation-cost inequality (14) are, in fact, equivalent.

**Theorem 5.** *Let  $\pi$  be log-concave, and assume that for some  $q \in (1, \infty)$  and a constant  $C > 0$ ,*

$$W_2^2(\mu, \pi) \leq C \chi^2(\mu \parallel \pi)^{2/q}, \quad \forall \mu \ll \pi.$$

*Then,*

$$\chi^2(\mu \parallel \pi)^{2/p} \leq 4C \mathbb{E}_\mu \left[ \left\| \nabla \frac{d\mu}{d\pi} \right\|^2 \right], \quad \forall \mu \ll \pi, \quad (15)$$

*where  $p$  satisfies  $1/p + 1/q = 1$ .*

*Proof.* Following [OV00], let  $T$  be the optimal transport map from  $\mu$  to  $\pi$ . Since  $\chi^2(\cdot \parallel \pi)$  is displacement convex [OT11; OT13] and has Wasserstein gradient  $2\nabla(d\mu/d\pi)$  at  $\mu$  (c.f. Section 2.1), the “above-tangent” formulation of displacement convexity ([Vil03, Proposition 5.29]) yields

$$0 = \chi^2(\pi \parallel \pi) \geq \chi^2(\mu \parallel \pi) + 2 \mathbb{E}_\mu \langle \nabla \frac{d\mu}{d\pi}, T - \text{id} \rangle \geq \chi^2(\mu \parallel \pi) - 2W_2(\mu, \pi) \sqrt{\mathbb{E}_\mu [\|\nabla \frac{d\mu}{d\pi}\|^2]},$$

where we used the Cauchy-Schwarz inequality for the last inequality. Rearranging the above display and using the transportation-cost inequality assumed in the statement of theorem, we get

$$\chi^2(\mu \parallel \pi) \leq 2W_2(\mu, \pi) \sqrt{\mathbb{E}_\mu [\|\nabla \frac{d\mu}{d\pi}\|^2]} \leq 2\sqrt{C \mathbb{E}_\mu [\|\nabla \frac{d\mu}{d\pi}\|^2]} \chi^2(\mu \parallel \pi)^{1/q}.$$

The result follows by rearranging the terms.  $\square$

*Proof of Theorem 2 (log-concave case).* From the transportation-cost inequality (14) and Theorem 5 with  $p = q = 2$ , we obtain

$$\chi^2(\mu \parallel \pi) \leq 8C_P \mathbb{E}_\mu [\|\nabla \frac{d\mu}{d\pi}\|^2].$$

This PL inequality together with Grönwall’s inequality readily yields the result.  $\square$

We conclude this section with the proof of Theorem 3, which shows exponential convergence of CSF in chi-squared divergence under the assumption of a log-Sobolev inequality (LSI) (but without the assumption of log-concavity).

*Proof of Theorem 3.* We first claim that

$$\partial_t \chi^2(\mu_t \parallel \pi) \leq -\frac{4}{9C_{\text{LSI}}} [\chi^2(\mu_t \parallel \pi) + 1]^{3/2} \ln[\chi^2(\mu_t \parallel \pi) + 1]. \quad (16)$$

Indeed, applying (LSI), we obtain

$$\partial_t \chi^2(\mu_t \parallel \pi) = -4 \int \|\nabla \frac{d\mu_t}{d\pi}\|^2 d\mu_t = -\frac{16}{9} \int \|\nabla |\frac{d\mu_t}{d\pi}|^{3/2}\|^2 d\pi \leq -\frac{8}{9C_{\text{LSI}}} \text{ent}_\pi(|\frac{d\mu_t}{d\pi}|^3).$$

Next, the variational formula for the entropy gives

$$\text{ent}_\pi f = \sup\{\mathbb{E}_\pi(fg) : g \text{ satisfies } \mathbb{E}_\pi \exp g = 1\},$$

see [Han16, Lemma 3.15] or [BLM13, Theorem 4.13]. Choosing  $g = \ln(d\mu_t/d\pi)$  yields

$$\begin{aligned} \text{ent}_\pi(|\frac{d\mu_t}{d\pi}|^3) &\geq \mathbb{E}_\pi [|\frac{d\mu_t}{d\pi}|^3 \ln \frac{d\mu_t}{d\pi}] = \frac{1}{3} \mathbb{E}_\pi [|\frac{d\mu_t}{d\pi}|^3 \ln(|\frac{d\mu_t}{d\pi}|^3)] \\ &\geq \frac{1}{3} \mathbb{E}_\pi [|\frac{d\mu_t}{d\pi}|^3] \ln \mathbb{E}_\pi [|\frac{d\mu_t}{d\pi}|^3] \\ &\geq \frac{1}{2} \mathbb{E}_\pi [|\frac{d\mu_t}{d\pi}|^2]^{3/2} \ln \mathbb{E}_\pi [|\frac{d\mu_t}{d\pi}|^2] \\ &= \frac{1}{2} [\chi^2(\mu_t \parallel \pi) + 1]^{3/2} \ln[\chi^2(\mu_t \parallel \pi) + 1], \end{aligned}$$

where in the second inequality, we used that  $x \mapsto x \ln x$  is convex on  $\mathbb{R}_+$  and in the third, we used that it increasing when  $x \geq 1$  together with

$$\mathbb{E}_\pi [|\frac{d\mu_t}{d\pi}|^2] = 1 + \chi^2(\mu_t \parallel \pi) \geq 1.$$

This proves (16).

To simplify the inequality (16), we use the crude bounds

$$\ln[\chi^2(\mu_t \parallel \pi) + 1] \geq \begin{cases} 1, & \text{if } \chi^2(\mu_t \parallel \pi) \geq e - 1 \\ \chi^2(\mu_t \parallel \pi)/2, & \text{otherwise.} \end{cases}$$

It yields respectively

$$\partial_t \chi^2(\mu_t \parallel \pi) \leq -\frac{2}{9C_{\text{LSI}}} \begin{cases} 2\chi^2(\mu_t \parallel \pi)^{3/2}, & \text{if } \chi^2(\mu_t \parallel \pi) \geq e-1, \\ \chi^2(\mu_t \parallel \pi), & \text{otherwise.} \end{cases} \quad (17)$$

Solving the differential inequality in the first case yields

$$e-1 \leq \chi^2(\mu_t \parallel \pi) \leq \left[ \frac{9C_{\text{LSI}} \sqrt{\chi^2(\mu_0 \parallel \pi)}}{9C_{\text{LSI}} + 2t \sqrt{\chi^2(\mu_0 \parallel \pi)}} \right]^2 \leq \left[ \frac{9C_{\text{LSI}}}{2t} \right]^2,$$

so that in this first case, it must holds that

$$t \leq \frac{9C_{\text{LSI}}}{2\sqrt{e-1}} < 3.5C_{\text{LSI}} =: t_0.$$

Therefore, if  $t \geq t_0$ , we are in the second case. In particular,  $\chi^2(\mu_{t_0} \parallel \pi) \leq e-1 \leq 2$  and integrating the differential inequality between  $t_0$  and  $t$  we get

$$\chi^2(\mu_t \parallel \pi) \leq \chi^2(\mu_{t_0} \parallel \pi) e^{-\frac{2(t-t_0)}{9C_{\text{LSI}}}} \leq (\chi^2(\mu_0 \parallel \pi) \wedge 2) e^{-\frac{2(t-t_0)}{9C_{\text{LSI}}}},$$

where in the last inequality, we used the fact that  $t \mapsto \chi^2(\mu_t \parallel \pi)$  is decreasing so that it also holds  $\chi^2(\mu_{t_0} \parallel \pi) \leq \chi^2(\mu_0 \parallel \pi)$ . In particular, taking  $t \geq 2t_0 = 7C_{\text{LSI}}$  yields the desired result.  $\square$

## C Details for the experiments

We give additional details for the experiments presented in this paper. All methods were implemented in Python. Since the Schrödinger operator requires the Laplacian and gradient of the potential  $V$ , we employ automatic differentiation to avoid laborious calculations of these derivatives.

The *probabilists'* Hermite polynomials are well-known to be eigenfunctions of the 1D Ornstein-Uhlenbeck operator  $\mathcal{L}$  given by  $\mathcal{L}f(x) := -f''(x) + xf'(x)$ , and they satisfy the recursive relationship  $H_{n+1}(x) = xH_n(x) - nH_{n-1}(x)$ , with  $H_0(x) = 1$  and  $H_1(x) = x$ . It also holds that  $H'_n(x) = nH_{n-1}(x)$ . With these equations, it is easy to check that the eigenvalue corresponding to  $H_n$  is  $\lambda_n = n$ . These are used as the eigenfunctions and eigenvalues in the standard normal example given in Figure 2. In the simulation, we use the first 150 Hermite polynomials. We run **LAWGD** for 2000 iterations with a constant step size, with initial points drawn uniformly from the interval  $[2.5, 4.5]$ .

In Figure 1, we display an example of sampling 50 particles from a mixture of two 2-dimensional Gaussian distributions given by  $\pi = \frac{1}{2}\mathcal{N}((-1, -1)^\top, I_2) + \frac{1}{2}\mathcal{N}((1, 1)^\top, I_2)$ . To run this experiment, we use a 2-dimensional FD method, which approximates the Laplacian as

$$\Delta_\varepsilon f(x, y) := \frac{f(x - \varepsilon, y) + f(x + \varepsilon, y) + f(x, y - \varepsilon) + f(x, y + \varepsilon) - 4f(x)}{\varepsilon^2}.$$

We again use the Schrödinger operator for stability and use FD again to compute the gradients of the eigenfunctions. We use a  $128 \times 128$  grid of evenly spaced  $x$  and  $y$  values between  $-6$  and  $6$ . We calculate only the bottom 100 eigenvalues and eigenfunctions, since the other eigenfunctions incur additional computational cost without noticeably changing the result. Any negative eigenvalues (which arise from numerical errors) are discarded.

Additionally, we display the results from running SVGD with the RBF kernel and median-based bandwidth on this example with a less favorable initialization. True samples from  $\pi$  are displayed for comparison. Both **LAWGD** and SVGD are run for 20000 iterations with a constant step size. The samples from **LAWGD** tend to move very fast from their initial positions and then tend to settle into their final positions as seen in Figure 4. On the other hand, with constant step size, the samples of SVGD do not seem to converge, and one must use a decreasing step size scheme in order for the particles to stabilize. We also note that many of the samples generated by SVGD tend to blow up with a constant step size.

In Figure 5, we plot the particles of **LAWGD** and SVGD at iterations 100, 200, 1000, and 2000 to compare the speed of convergence.

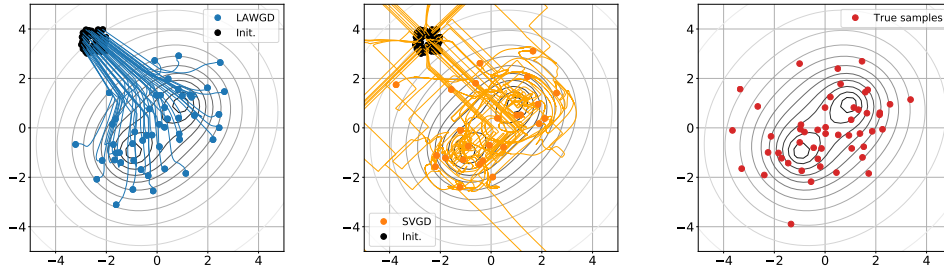


Figure 4: Left: 50 particles and trajectories generated from  $\frac{1}{2}\mathcal{N}((-1, -1)^\top, I_2) + \frac{1}{2}\mathcal{N}((1, 1)^\top, I_2)$  with **LAWGD**. Middle: 50 particles and trajectories generated by **SVGD**. Right: true samples from the distribution.

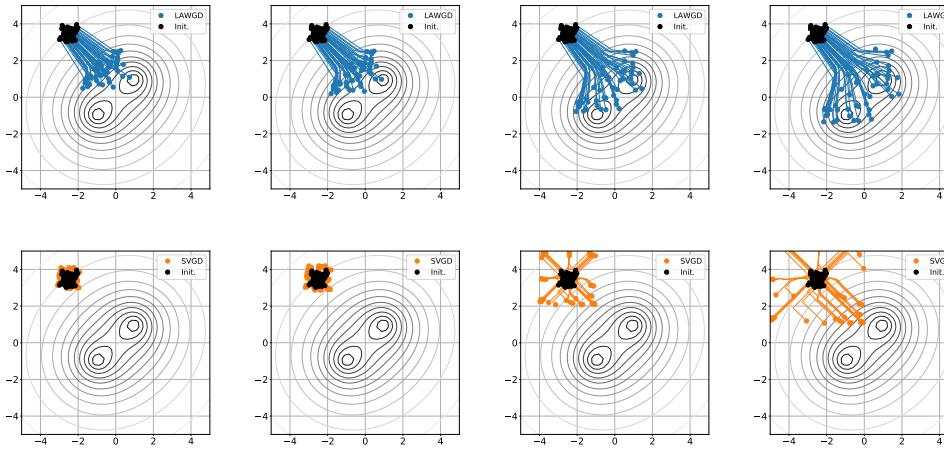


Figure 5: Top: **LAWGD** after 100, 200, 1000, and 2000 iterations. Bottom: **SVGD** after 100, 200, 1000, and 2000 iterations.

## D Broader impact

The sampling algorithms designed in this paper have the potential to improve a wide variety of Bayesian methods and therefore have an indirect impact on various domains such as health and medicine where such methods are pervasive. Sampling algorithms are also used for the generation of automated spam messages, which have potentially negative effects on society. Since this paper is primarily focused on theory, these questions are not addressed here.

## E Acknowledgments

Philippe Rigollet was supported by NSF awards IIS-1838071, DMS-1712596, DMS-TRIPODS-1740751. Sinho Chewi was supported by the Department of Defense (DoD) through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program. Thibaut Le Gouic was supported by ONR grant N00014-17-1-2147 and NSF IIS-1838071.

We thank the reviewers for very helpful suggestions regarding the presentation of the paper.

## References

- [AB15] D. Alonso-Gutiérrez and J. Bastero. *Approaching the Kannan-Lovász-Simonovits and variance conjectures*. Vol. 2131. Lecture Notes in Mathematics. Springer, Cham, 2015, pp. x+148.
- [AGS08] L. Ambrosio, N. Gigli, and G. Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2008.
- [Ber18] E. Bernton. “Langevin Monte Carlo and JKO splitting”. In: *Proceedings of the 31st Conference On Learning Theory*. Ed. by S. Bubeck, V. Perchet, and P. Rigollet. Vol. 75. Proceedings of Machine Learning Research. PMLR, 2018, pp. 1777–1798.
- [BGL14] D. Bakry, I. Gentil, and M. Ledoux. *Analysis and geometry of Markov diffusion operators*. Vol. 348. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Springer, Cham, 2014, pp. xx+552.
- [BLM13] S. Boucheron, G. Lugosi, and P. Massart. *Concentration inequalities*. A nonasymptotic theory of independence, With a foreword by Michel Ledoux. Oxford University Press, Oxford, 2013, pp. x+481.
- [Bob99] S. G. Bobkov. “Isoperimetric and analytic inequalities for log-concave probability measures”. In: *The Annals of Probability* 27.4 (1999), pp. 1903–1921.
- [Car11] E. A. Carlen. “Functional inequalities and dynamics”. In: *Nonlinear PDE’s and Applications*. Springer, 2011, pp. 17–85.
- [CB18] X. Cheng and P. Bartlett. “Convergence of Langevin MCMC in KL-divergence”. In: *Algorithmic Learning Theory 2018*. Vol. 83. Proc. Mach. Learn. Res. (PMLR). Proceedings of Machine Learning Research PMLR, 2018, p. 26.
- [Che+18] X. Cheng, N. S. Chatterji, P. L. Bartlett, and M. I. Jordan. “Underdamped Langevin MCMC: A non-asymptotic analysis”. In: *Proceedings of the 31st Conference On Learning Theory*. Ed. by S. Bubeck, V. Perchet, and P. Rigollet. Vol. 75. Proceedings of Machine Learning Research. PMLR, July 2018, pp. 300–323.
- [Che+20a] S. Chewi, T. Le Gouic, C. Lu, T. Maunu, P. Rigollet, and A. J. Stromme. “Exponential ergodicity of mirror-Langevin diffusions”. In: *arXiv e-prints*, arXiv:2005.09669 (May 2020).
- [Che+20b] S. Chewi, T. Maunu, P. Rigollet, and A. Stromme. “Gradient descent algorithms for Bures-Wasserstein barycenters”. In: *Proceedings of Thirty Third Conference on Learning Theory*. Ed. by J. Abernethy and S. Agarwal. Vol. 125. Proceedings of Machine Learning Research. PMLR, Sept. 2020, pp. 1276–1304.
- [CLL19] Y. Cao, J. Lu, and Y. Lu. “Exponential decay of Rényi divergence under Fokker-Planck equations”. In: *J. Stat. Phys.* 176.5 (2019), pp. 1172–1184.
- [CT02] J. A. Carrillo and G. Toscani. “Long-time asymptotics for strong solutions of the thin film equation”. In: *Comm. Math. Phys.* 225.3 (2002), pp. 551–571.
- [Dal17a] A. Dalalyan. “Further and stronger analogy between sampling and optimization: Langevin Monte Carlo and gradient descent”. In: *Proceedings of the 2017 Conference on Learning Theory*. Ed. by S. Kale and O. Shamir. Vol. 65. Proceedings of Machine Learning Research. Amsterdam, Netherlands: PMLR, 2017, pp. 678–689.
- [Dal17b] A. S. Dalalyan. “Theoretical guarantees for approximate sampling from smooth and log-concave densities”. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 79.3 (2017), pp. 651–676.
- [Det+18] G. Detommaso, T. Cui, Y. Marzouk, A. Spantini, and R. Scheichl. “A Stein variational Newton method”. In: *Advances in Neural Information Processing Systems*. 2018, pp. 9169–9179.
- [Din15] Y. Ding. “A note on quadratic transportation and divergence inequality”. In: *Statist. Probab. Lett.* 100 (2015), pp. 115–123.
- [DM17] A. Durmus and É. Moulines. “Nonasymptotic convergence analysis for the unadjusted Langevin algorithm”. In: *Ann. Appl. Probab.* 27.3 (2017), pp. 1551–1587.

- [DMM19] A. Durmus, S. Majewski, and B. Miasojedow. “Analysis of Langevin Monte Carlo via convex optimization”. In: *J. Mach. Learn. Res.* 20 (2019), Paper No. 73, 46.
- [DNS19] A. Duncan, N. Nuesken, and L. Szpruch. “On the geometry of Stein variational gradient descent”. In: *arXiv e-prints*, arXiv:1912.00894 (Dec. 2019).
- [DR20] A. S. Dalalyan and L. Riou-Durand. “On sampling from a log-concave density using kinetic Langevin diffusions”. In: *Bernoulli* 26.3 (2020), pp. 1956–1988.
- [Dwi+19] R. Dwivedi, Y. Chen, M. J. Wainwright, and B. Yu. “Log-concave sampling: Metropolis-Hastings algorithms are fast”. In: *Journal of Machine Learning Research* 20.183 (2019), pp. 1–42.
- [Eva10] L. C. Evans. *Partial differential equations*. Second. Vol. 19. Graduate Studies in Mathematics. American Mathematical Society, Providence, RI, 2010, pp. xxii+749.
- [Fri34] K. Friedrichs. “Spektraltheorie Halbbeschränkter Operatoren Und Anwendung Auf Die Spektralzerlegung von Differentialoperatoren”. In: *Mathematische Annalen* 109.1 (1934), pp. 465–487.
- [Han16] R. van Handel. *Probability in high dimension*. Lecture Notes (Princeton University). 2016.
- [Hsi+18] Y.-P. Hsieh, A. Kavis, P. Rolland, and V. Cevher. “Mirrored Langevin dynamics”. In: *Advances in Neural Information Processing Systems*. 2018, pp. 2878–2887.
- [JKO98] R. Jordan, D. Kinderlehrer, and F. Otto. “The variational formulation of the Fokker-Planck equation”. In: *SIAM Journal on Mathematical Analysis* 29.1 (1998), pp. 1–17.
- [KLS95] R. Kannan, L. Lovász, and M. Simonovits. “Isoperimetric problems for convex bodies and a localization lemma”. In: *Discrete Comput. Geom.* 13.3-4 (1995), pp. 541–559.
- [KNS16] H. Karimi, J. Nutini, and M. Schmidt. “Linear convergence of gradient and proximal-gradient methods under the Polyak-Lojasiewicz condition”. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer. 2016, pp. 795–811.
- [Led18] M. Ledoux. *Remarks on some transportation cost inequalities*. 2018.
- [Li+20] L. Li, Y. Li, J.-G. Liu, Z. Liu, and J. Lu. “A stochastic version of Stein variational gradient descent for efficient sampling”. In: *Commun. Appl. Math. Comput. Sci.* 15.1 (2020), pp. 37–63.
- [Liu+19] C. Liu, J. Zhuo, P. Cheng, R. Zhang, and J. Zhu. “Understanding and accelerating particle-based variational inference”. In: *International Conference on Machine Learning*. 2019, pp. 4082–4092.
- [Liu17] Q. Liu. “Stein variational gradient descent as gradient flow”. In: *Advances in Neural Information Processing Systems 30*. Ed. by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett. Curran Associates, Inc., 2017, pp. 3115–3123.
- [Liu20] Y. Liu. “The Poincaré inequality and quadratic transportation-variance inequalities”. In: *Electron. J. Probab.* 25 (2020), Paper No. 1, 16.
- [LLN19] J. Lu, Y. Lu, and J. Nolen. “Scaling limit of the Stein variational gradient descent: The mean field regime”. In: *SIAM Journal on Mathematical Analysis* 51.2 (2019), pp. 648–671.
- [Loj63] S. Lojasiewicz. “Une propriété topologique des sous-ensembles analytiques réels”. In: *Les équations aux dérivées partielles* 117 (1963), pp. 87–89.
- [LV17] Y. T. Lee and S. S. Vempala. “Eldan’s stochastic localization and the KLS hyperplane conjecture: an improved lower bound for expansion”. In: *58th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2017*. IEEE Computer Soc., Los Alamitos, CA, 2017, pp. 998–1007.

- [LW16] Q. Liu and D. Wang. “Stein variational gradient descent: A general purpose Bayesian inference algorithm”. In: *Advances in Neural Information Processing Systems*. 2016, pp. 2378–2386.
- [Ma+19] Y.-A. Ma, N. Chatterji, X. Cheng, N. Flammarion, P. Bartlett, and M. I. Jordan. “Is there an analog of Nesterov acceleration for MCMC?” In: *arXiv e-prints*, arXiv:1902.00996 (Feb. 2019).
- [Mar+12] J. Martin, L. C. Wilcox, C. Burstedde, and O. Ghattas. “A stochastic Newton MCMC method for large-scale statistical inverse problems with application to seismic inversion”. In: *SIAM J. Sci. Comput.* 34.3 (2012), A1460–A1487.
- [MMS09] D. Matthes, R. J. McCann, and G. Savaré. “A family of nonlinear fourth order equations of gradient flow type”. In: *Comm. Partial Differential Equations* 34.10-12 (2009), pp. 1352–1397.
- [MT09] S. Meyn and R. L. Tweedie. *Markov chains and stochastic stability*. 2nd. USA: Cambridge University Press, 2009.
- [OT11] S.-i. Ohta and A. Takatsu. “Displacement convexity of generalized relative entropies”. In: *Adv. Math.* 228.3 (2011), pp. 1742–1787.
- [OT13] S.-i. Ohta and A. Takatsu. “Displacement convexity of generalized relative entropies. II”. In: *Comm. Anal. Geom.* 21.4 (2013), pp. 687–785.
- [OV00] F. Otto and C. Villani. “Generalization of an inequality by Talagrand and links with the logarithmic Sobolev inequality”. In: *Journal of Functional Analysis* 173.2 (2000), pp. 361–400.
- [Pav14] G. A. Pavliotis. *Stochastic processes and applications: diffusion processes, the Fokker-Planck and Langevin equations*. Vol. 60. Springer, 2014.
- [RS78] M. Reed and B. Simon. *Methods of modern mathematical physics. IV. Analysis of operators*. Academic Press [Harcourt Brace Jovanovich, Publishers], New York-London, 1978, pp. xv+396.
- [San15] F. Santambrogio. *Optimal transport for applied mathematicians*. Vol. 87. Progress in Nonlinear Differential Equations and their Applications. Calculus of variations, PDEs, and modeling. Birkhäuser/Springer, Cham, 2015, pp. xxvii+353.
- [San17] F. Santambrogio. “{Euclidean, metric, and Wasserstein} gradient flows: an overview”. In: *Bulletin of Mathematical Sciences* 7.1 (2017), pp. 87–154.
- [Sim+16] U. Simsekli, R. Badeau, A. T. Cemgil, and G. Richard. “Stochastic quasi-Newton Langevin Monte Carlo”. In: *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*. ICML’16. New York, NY, USA: JMLR.org, 2016, pp. 642–651.
- [SKL20] A. Salim, A. Korba, and G. Luise. “Wasserstein proximal gradient”. In: *arXiv e-prints*, arXiv:2002.03035 (Feb. 2020).
- [Tsy09] A. B. Tsybakov. *Introduction to nonparametric estimation*. Springer Series in Statistics. Revised and extended from the 2004 French original, Translated by Vladimir Zaiats. Springer, New York, 2009, pp. xii+214.
- [Vil03] C. Villani. *Topics in optimal transportation*. Vol. 58. Graduate Studies in Mathematics. American Mathematical Society, Providence, RI, 2003, pp. xvi+370.
- [Vil09] C. Villani. *Optimal transport*. Vol. 338. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Old and new. Springer-Verlag, Berlin, 2009, pp. xxii+973.
- [VW19] S. Vempala and A. Wibisono. “Rapid convergence of the unadjusted Langevin algorithm: isoperimetry suffices”. In: *Advances in Neural Information Processing Systems* 32. Ed. by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett. Curran Associates, Inc., 2019, pp. 8094–8106.
- [Wan+19] D. Wang, Z. Tang, C. Bajaj, and Q. Liu. “Stein variational gradient descent with matrix-valued kernels”. In: *Advances in Neural Information Processing Systems* 32. Ed. by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett. Curran Associates, Inc., 2019, pp. 7836–7846.



- [Wib18] A. 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*. Ed. by S. Bubeck, V. Perchet, and P. Rigollet. Vol. 75. Proceedings of Machine Learning Research. PMLR, 2018, pp. 2093–3027.
- [Wib19] A. Wibisono. “Proximal Langevin algorithm: rapid convergence under isoperimetry”. In: *arXiv e-prints*, arXiv:1911.01469 (Nov. 2019).
- [WL20] Y. Wang and W. Li. “Information Newton’s flow: second-order optimization method in probability space”. In: *arXiv e-prints*, arXiv:2001.04341 (Jan. 2020).
- [Zha+20a] J. Zhang, R. Zhang, L. Carin, and C. Chen. “Stochastic particle-optimization sampling and the non-asymptotic convergence theory”. In: ed. by S. Chiappa and R. Calandra. Vol. 108. Proceedings of Machine Learning Research. Online: PMLR, 26–28 Aug 2020, pp. 1877–1887.
- [Zha+20b] K. S. Zhang, G. Peyré, J. Fadili, and M. Pereyra. “Wasserstein control of mirror Langevin Monte Carlo”. In: *arXiv e-prints*, arXiv:2002.04363 (Feb. 2020).