

---

# Learning Rate Free Bayesian Inference in Constrained Domains

---

**Louis Sharrock**

Department of Mathematics and Statistics  
Lancaster University, UK  
l.sharrock@lancaster.ac.uk

**Lester Mackey**

Microsoft Research New England  
Cambridge, MA  
lmackey@microsoft.com

**Christopher Nemeth**

Department of Mathematics and Statistics  
Lancaster University, UK  
c.nemeth@lancaster.ac.uk

## Abstract

We introduce a suite of new particle-based algorithms for sampling on constrained domains which are entirely learning rate free. Our approach leverages coin betting ideas from convex optimisation, and the viewpoint of constrained sampling as a mirrored optimisation problem on the space of probability measures. Based on this viewpoint, we also introduce a unifying framework for several existing constrained sampling algorithms, including mirrored Langevin dynamics and mirrored Stein variational gradient descent. We demonstrate the performance of our algorithms on a range of numerical examples, including sampling from targets on the simplex, sampling with fairness constraints, and constrained sampling problems in post-selection inference. Our results indicate that our algorithms achieve competitive performance with existing constrained sampling methods, without the need to tune any hyperparameters.

## 1 Introduction

The problem of sampling from unnormalised probability distributions is of central importance to computational statistics and machine learning. Standard approaches to this problem include Markov chain Monte Carlo (MCMC) [11, 73] and variational inference (VI) [8, 40, 92]. More recently, there has been growing interest in particle-based variational inference (ParVI) methods, which deterministically evolve a collection of particles towards the target distribution and combine favourable aspects of both MCMC and VI. Perhaps the most well-known of these methods is Stein variational gradient descent (SVGD) [61], which iteratively updates particles according to a form of gradient descent on the Kullback-Leibler (KL) divergence. This approach has since given rise to several extensions [16, 26, 93, 94, 101] and found success in a wide range of applications [26, 64, 71].

While such methods have enjoyed great success in sampling from unconstrained distributions, they typically break down when applied to constrained targets. These include targets that are not integrable over the entire Euclidean space (e.g., the uniform distribution), when the target density is undefined outside a particular domain (e.g., the Dirichlet distribution), or when the target must satisfy certain additional constraints (e.g., fairness constraints in Bayesian inference [63]). Notable examples include latent Dirichlet allocation [9], ordinal data models [44], regularised regression [15], survival analysis [48], and post-selection inference [53, 85]. In recent years, there has been increased interest in this topic, with several new methodological and theoretical developments [1, 19, 42, 43, 43, 57, 58, 79, 83, 99, 100].

One limitation shared by all of these methods is that their performance depends, often significantly, on a suitable choice of learning rate. In principle, convergence rates for existing constrained sampling schemes (e.g., Theorems 1 - 3 in [1] or Theorems 1 - 2 in [83]) allow one to compute the optimal learning rate for a given problem. In practice, however, this optimal learning rate is a function of the unknown target and thus cannot be computed. Motivated by this problem, in this paper we introduce a suite of new sampling algorithms for constrained domains which are entirely learning rate free.

**Our contributions** We first propose a unifying perspective of several existing constrained sampling algorithms, based on the viewpoint of sampling in constrained domains as a ‘mirrored’ optimisation problem on the space of probability measures. Based on this perspective, we derive a general class of solutions to the constrained sampling problem via the mirrored Wasserstein gradient flow (MWGF). The MWGF includes, as special cases, several existing constrained sampling methods, e.g., the mirrored Langevin dynamics in [42] and the mirrored SVGD (MSVGD) algorithm in [79]. Using this formulation, we also introduce mirrored versions of several other particle-based sampling algorithms which are suitable for sampling in constrained domains, and study their convergence properties.

By leveraging this perspective, and extending the coin sampling methodology recently introduced in [78], we then obtain a suite of learning rate free algorithms suitable for constrained domains. In particular, we introduce the mirrored coin Wasserstein gradient descent algorithm, which includes as a special case Coin MSVGD, as well as several other tuning-free constrained sampling algorithms. We also outline an alternative approach which does not rely on a mirror mapping, namely, coin mollified interaction energy descent (Coin MIED). This algorithm can be viewed as the natural coin sampling analogue of the method recently introduced in [57]. Finally, we demonstrate the performance of our methods on a wide range of examples, including sampling on the simplex, sampling with fairness constraints, and constrained sampling problems in post-selection inference. Our empirical results indicate that our methods achieve competitive performance with the optimal performance of existing approaches, with no need to tune a learning rate.

## 2 Preliminaries

**The Wasserstein Space.** We will make use of the following notation. For any  $\mathcal{X} \subset \mathbb{R}^d$ , we write  $\mathcal{P}_2(\mathcal{X})$  for the set of probability measures  $\mu$  on  $\mathcal{X}$  with finite 2<sup>nd</sup> moment:  $\int_{\mathcal{X}} \|x\|^2 \mu(dx) < \infty$  for  $\|\cdot\|$  the Euclidean norm. For  $\mu \in \mathcal{P}_2(\mathcal{X})$ , we write  $L^2(\mu)$  for the set of measurable  $f : \mathcal{X} \rightarrow \mathbb{R}$  with finite  $\|f\|_{L^2(\mu)}^2 = \int_{\mathcal{X}} |f(x)|^2 \mu(dx)$ . For  $f, g \in L^2(\mu)$ , we also define  $\langle f, g \rangle_{L^2(\mu)} = \int_{\mathcal{X}} f(x)^\top g(x) \mu(dx)$ . Given  $\mu \in \mathcal{P}_2(\mathcal{X})$  and  $T \in L^2(\mu)$ , we write  $T_{\#}\mu$  for the push-forward of  $\mu$  under  $T$  (the distribution of  $T(X)$  when  $X \sim \mu$ ). For each  $\mu, \nu \in \mathcal{P}_2(\mathcal{X})$ , the Wasserstein 2-distance between  $\mu$  and  $\nu$  is defined by  $W_2^2(\mu, \nu) = \inf_{\gamma \in \Gamma(\mu, \nu)} \int_{\mathcal{X} \times \mathcal{X}} \|x - y\|^2 \gamma(dx, dy)$ , where  $\Gamma(\mu, \nu)$  denotes the set of couplings between  $\mu$  and  $\nu$ . The Wasserstein distance  $W_2$  is indeed a distance over  $\mathcal{P}_2(\mathcal{X})$  [2, Chapter 7.1], with the metric space  $(\mathcal{P}_2(\mathcal{X}), W_2)$  known as the Wasserstein space. Finally, for a functional  $\mathcal{F} : \mathcal{P}(\mathcal{X}) \rightarrow (-\infty, \infty]$ , we write  $\nabla_{W_2} \mathcal{F}(\mu)$  for the Wasserstein gradient of  $\mathcal{F}$  at  $\mu$ , which exists under appropriate regularity conditions [2, Lemma 10.4.1].

**Wasserstein Gradient Flows.** Suppose that  $\pi \in \mathcal{P}_2(\mathcal{X})$  and that  $\pi$  admits a density  $\pi \propto e^{-V}$  with respect to Lebesgue measure,<sup>1</sup> where  $V : \mathcal{X} \rightarrow \mathbb{R}$  is a smooth potential function. The problem of sampling from  $\pi$  can be recast as an optimisation problem over  $\mathcal{P}_2(\mathcal{X})$  [e.g., 96]. In particular, one can view the target  $\pi$  as the solution of

$$\pi = \arg \min_{\mu \in \mathcal{P}_2(\mathcal{X})} \mathcal{F}(\mu) \quad (1)$$

where  $\mathcal{F} : \mathcal{P}(\mathcal{X}) \rightarrow (-\infty, \infty]$  is a functional uniquely minimised at  $\pi$ . In the unconstrained case,  $\mathcal{X} = \mathbb{R}^d$ , a now rather classical method for solving this problem is to identify a continuous process which pushes samples from an initial distribution  $\mu_0$  to the target  $\pi$ . This corresponds to finding a family of vector fields  $(v_t)_{t \geq 0}$ , with  $v_t \in L^2(\mu_t)$ , which transports  $\mu_t$  to  $\pi$  via the continuity equation

$$\frac{\partial \mu_t}{\partial t} + \nabla \cdot (v_t \mu_t) = 0. \quad (2)$$

A standard choice is  $v_t = -\nabla_{W_2} \mathcal{F}(\mu_t)$ , the negative Wasserstein gradient, in which case the solution  $(\mu_t)_{t \geq 0}$  of (2) is referred to as the Wasserstein gradient flow (WGF) of  $\mathcal{F}$ . For different choices

<sup>1</sup>In a slight abuse of notation, throughout this paper we will use the same letter to refer to a distribution and its density with respect to Lebesgue measure.

of  $(v_t)_{t \geq 0}$  in (2), one can obtain the continuous-time limit of several popular sampling algorithms [3, 20, 39, 45, 50, 61, 96]. Unfortunately, it is not straightforward to extend these approaches to the case in which  $\mathcal{X} \subset \mathbb{R}^d$  is a constrained domain. For example, the vector fields  $v_t$  may push the random variable  $x_t$  outside of its support  $\mathcal{X}$ , thus rendering all future updates undefined.

### 3 Mirrored Wasserstein Gradient Flows

We now outline one way to extend the WGF approach to the constrained case, based on the use of a mirror map. We will first require some basic definitions from convex optimisation [e.g., 74]. Let  $\mathcal{X}$  be a convex domain in  $\mathbb{R}^d$ . Let  $\phi : \mathcal{X} \rightarrow \mathbb{R} \cup \{\infty\}$  be a proper, lower semicontinuous, convex function of Legendre type [74]. This implies, in particular, that  $\nabla\phi(\mathcal{X}) = \mathbb{R}^d$  and  $\nabla\phi : \mathcal{X} \rightarrow \mathbb{R}^d$  is bijective [74, Theorem 26.5]. Moreover, its inverse  $(\nabla\phi)^{-1} : \mathbb{R}^d \rightarrow \mathcal{X}$  satisfies  $(\nabla\phi)^{-1} = \nabla\phi^*$ , where  $\phi^* : \mathbb{R}^d \rightarrow \mathbb{R}$  denotes the Fenchel conjugate of  $\phi$ , defined as  $\phi^*(y) = \sup_{x \in \mathcal{X}} \langle x, y \rangle - \phi(x)$ . We will refer to  $\nabla\phi : \mathcal{X} \rightarrow \mathbb{R}^d$  as the mirror map and  $\nabla\phi(\mathcal{X}) = \mathbb{R}^d$  as the dual space.

#### 3.1 Constrained Sampling as Optimisation

Using the mirror map  $\nabla\phi : \mathcal{X} \rightarrow \mathbb{R}^d$ , we can now reformulate the constrained sampling problem as the solution of a ‘mirrored’ version of the optimisation problem in (1). Let us define  $\nu = (\nabla\phi)_\# \pi$ , so  $\pi = (\nabla\phi^*)_\# \nu$ . We can then view the target  $\pi$  as the solution of

$$\pi = (\nabla\phi^*)_\# \nu, \quad \nu = \arg \min_{\eta \in \mathcal{P}_2(\mathbb{R}^d)} \mathcal{F}(\eta), \quad (3)$$

where now  $\mathcal{F} : \mathcal{P}_2(\mathbb{R}^d) \rightarrow (-\infty, \infty]$  is a functional uniquely minimised by  $\nu = (\nabla\phi)_\# \pi$ . The motivation for this formulation is clear. Rather than directly solving a constrained optimisation problem for the target  $\pi \in \mathcal{P}_2(\mathcal{X})$  as in (1), we could instead solve an unconstrained optimisation problem for the mirrored target  $\nu \in \mathcal{P}_2(\mathbb{R}^d)$  and then recover the target via  $\pi = (\nabla\phi^*)_\# \nu$ . This is a rather natural extension of the formulation of sampling as optimisation to the constrained setting.

#### 3.2 Mirrored Wasserstein Gradient Flows

One way to solve the mirrored optimisation problem in (3) is to find a continuous process which transports samples from  $\eta_0 := (\nabla\phi)_\# \mu_0 \in \mathcal{P}_2(\mathbb{R}^d)$  to the mirrored target  $\nu$ . According to the mirror map, this process will implicitly also push samples from  $\mu_0$  to the target  $\pi$ . More precisely, we would like to find  $(w_t)_{t \geq 0}$ , where  $w_t : \mathbb{R}^d \rightarrow \mathbb{R}^d$ , which evolves  $\eta_t$  on  $\mathbb{R}^d$  and thus  $\mu_t$  on  $\mathcal{X}$ , through

$$\frac{\partial \eta_t}{\partial t} + \nabla \cdot (w_t \eta_t) = 0 \quad , \quad \mu_t = (\nabla\phi^*)_\# \eta_t \quad (4)$$

where  $(w_t)_{t \geq 0}$  should ensure convergence of  $\eta_t$  to  $\nu$ . In the case that  $w_t = -\nabla_{W_2} \mathcal{F}(\eta_t)$ , we will refer to (4) as the mirrored Wasserstein gradient flow (MWGF) of  $\mathcal{F}$ . The idea of (4) is to simulate the WGF in the (unconstrained) dual space and then map back to the primal (constrained) space using a carefully chosen map. This approach bypasses the inherent difficulties associated with simulating a gradient flow on the constrained space itself. While a similar approach has previously been used to derive certain mirrored sampling schemes [42, 79, 83], these works focus on specific choices of  $\mathcal{F}$  and  $(w_t)_{t \geq 0}$ , while we consider the general case. Our general perspective provides a unifying framework for these approaches, whilst also allowing us to obtain new constrained sampling algorithms and to study their convergence properties.

**Algorithms and Theoretical Guarantees.** In App. B, we outline several algorithms which arise as special cases of the MWGF and study in detail their properties in continuous time. These include existing algorithms, namely mirrored Langevin dynamics (MLD) [42] and mirrored Stein variational gradient descent (MSVGD) [79] and two new algorithms which correspond to mirrored versions of Laplacian adjusted Wasserstein gradient descent (LAWGD) [20] and kernel Stein discrepancy descent (KSDD) [50]. We now consider one of these algorithms, MSVGD, in some more detail.

##### 3.2.1 Mirrored Stein Variational Gradient Descent

MSVGD [79] is a mirrored version of the popular SVGD algorithm. In [79], MSVGD is derived using a mirrored Stein operator in a manner analogous to the original SVGD derivation [61]. Here, we take

---

**Algorithm 1** Mirrored SVGD
 

---

**input:** target density  $\pi$ , kernel  $k$ , mirror function  $\phi$ , particles  $(x_0^i)_{i=1}^N \sim \mu$ , step size  $\gamma$ .  
**initialise:**  $y_0^i = \nabla \phi(x_0^i)$  for  $i \in [N]$   
**for**  $t = 0, \dots, T - 1$  **do**  
   For  $i \in [N]$ ,  $y_{t+1}^i = y_t^i - \gamma P_{\hat{\eta}_t, k_\phi} \nabla \log \left( \frac{\hat{\eta}_t}{\nu} \right) (y_t^i)$ , where  $\hat{\eta}_t = \frac{1}{N} \sum_{j=1}^N \delta_{y_t^j}$ .  
   For  $i \in [N]$ ,  $x_{t+1}^i = \nabla \phi^*(y_{t+1}^i)$ .  
**return**  $(x_T^i)_{i=1}^N$ .

---

an alternative perspective, which originates in the MWGF. This formulation is more closely aligned with other more recent papers analysing the non-asymptotic convergence of SVGD [50, 75, 82, 83].

We will require the following additional notation. Let  $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$  denote a positive semi-definite kernel and  $\mathcal{H}_k$  the corresponding reproducing kernel Hilbert space (RKHS) of real-valued functions  $f : \mathcal{X} \rightarrow \mathbb{R}$ , with inner product  $\langle \cdot, \cdot \rangle_{\mathcal{H}_k}$  and norm  $\| \cdot \|_{\mathcal{H}_k}$ . Let  $\mathcal{H} := \mathcal{H}_k^d$  denote the Cartesian product RKHS, consisting of elements  $f = (f_1, \dots, f_d)$  for  $f_i \in \mathcal{H}_k$ . We write  $S_{\mu, k} : L^2(\mu) \rightarrow \mathcal{H}$  for the operator  $S_{\mu, k} f = \int k(x, \cdot) f(x) d\mu(x)$ . We assume  $\int k(x, x) d\mu(x) < \infty$  for  $\mu \in \mathcal{P}_2(\mathcal{X})$ , in which case  $\mathcal{H} \subset L^2(\mu)$  for all  $\mu \in \mathcal{P}_2(\mathcal{X})$  [51, Section 3.1]. We also define the identity embedding  $i : \mathcal{H} \rightarrow L^2(\mu)$  with adjoint  $i^* = S_{\mu, k}$ . Finally, we define  $P_{\mu, k} : L^2(\mu) \rightarrow L^2(\mu)$  as  $P_{\mu, k} = i S_{\mu, k}$ .

We are now ready to introduce our perspective on MSVGD. We begin by introducing the continuous-time MSVGD dynamics, which are defined by the continuity equation

$$\frac{\partial \eta_t}{\partial t} + \nabla \cdot (w_t \eta_t) = 0, \quad w_t = -P_{\eta_t, k_\phi} \nabla \log \left( \frac{\eta_t}{\nu} \right) \quad (5)$$

with  $\mu_t = (\nabla \phi^*)_{\#} \eta_t$  and  $k_\phi(y, y') = k(\nabla \phi^*(y), \nabla \phi^*(y'))$ , for some base kernel  $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ . In App. B.2, we derive convergence results for (5). The key point here is that, given samples of  $\eta$ , estimating  $P_{\eta, k_\phi} \nabla \log \left( \frac{\eta}{\nu} \right)$  is straightforward. In particular, if  $\lim_{\|y'\| \rightarrow \infty} k_\phi(y', \cdot) \nu(y') = 0$ , then  $P_{\eta, k_\phi} \nabla \log \left( \frac{\eta}{\nu} \right)(y) = - \int [\nabla \log \nu(y') k_\phi(y', y) + \nabla_{y'} k_\phi(y', y)] d\eta(y')$ . To obtain an implementable algorithm, we must discretise in time. By applying an explicit Euler discretisation to (5), we arrive at

$$\eta_{t+1} = \left( \text{id} - \gamma P_{\eta_t, k_\phi} \nabla \log \left( \frac{\eta_t}{\nu} \right) \right)_{\#} \eta_t \quad (6)$$

where  $\mu_{t+1} = (\nabla \phi^*)_{\#} \eta_{t+1}$  and  $\gamma > 0$  is a step size. This scheme represents the population limit of the MSVGD algorithm. Suppose  $\eta_0 \in \mathcal{P}_2(\mathbb{R}^d)$  admits a density with respect to the Lebesgue measure. Then, for each  $t \in \mathbb{N}_0$ ,  $\eta_t$  is the density of  $y_t$ , where

$$y_{t+1} = y_t - \gamma P_{\eta_t, k_\phi} \nabla \log \left( \frac{\eta_t}{\nu} \right) (y_t) \quad \text{and} \quad y_0 \sim \eta_0 = (\nabla \phi)_{\#} \mu_0, \quad (7)$$

and  $\mu_t$  is the density of  $x_t$ , where  $x_{t+1} = \nabla \phi^*(y_{t+1})$ . In practice, the densities  $(\eta_t)_{t \in \mathbb{N}_0}$  are unknown and must be estimated using samples. In particular, to approximate (7), we initialise a collection of particles  $y_0^i \stackrel{i.i.d.}{\sim} \eta_0 = (\nabla \phi)_{\#} \mu_0$  for  $i \in [N]$  and then update

$$y_{t+1}^i = y_t^i - \gamma P_{\hat{\eta}_t, k_\phi} \nabla \log \left( \frac{\hat{\eta}_t}{\nu} \right) (y_t^i) \quad \text{for} \quad \hat{\eta}_t = \frac{1}{N} \sum_{j=1}^N \delta_{y_t^j} \quad (8)$$

where, as per usual,  $x_{t+1}^i = \nabla \phi^*(y_{t+1}^i)$ . This is precisely the MSVGD algorithm introduced in [79] and since also studied in [83].

## 4 Mirrored Coin Sampling

By construction, any algorithm derived as a discretisation of the MWGF, including MSVGD, will depend on an appropriate choice of learning rate  $\gamma$ . In this section, we consider an alternative approach, leading to new algorithms which are entirely learning rate free. Our approach can be viewed as an extension of the coin sampling framework introduced in [78] to the constrained setting, based on the reformulation of the constrained sampling problem as a mirrored optimisation problem.

---

**Algorithm 2** Coin MSVGD
 

---

**input:** initial density  $\mu$ , target density  $\pi$ , kernel  $k$ , mirror function  $\phi$ , particles  $(x_0^i)_{i=1}^N \sim \mu$   
**initialise:**  $y_0^i = \nabla \phi(x_0^i)$  for  $i \in [N]$   
**for**  $t = 1, \dots, T$  **do**  
   For  $i \in [N]$ ,  $y_t^i = y_0^i - \frac{\sum_{s=1}^{t-1} P_{\hat{\eta}_s, k_\phi} \nabla \log(\frac{\hat{\eta}_s}{\nu})(y_s^i)}{t} \left( 1 - \sum_{s=1}^{t-1} \langle P_{\hat{\eta}_s, k_\phi} \nabla \log(\frac{\hat{\eta}_s}{\nu})(y_s^i), y_s^i - y_0^i \rangle \right)$ .  
   For  $i \in [N]$ ,  $x_t^i = \nabla \phi^*(y_t^i)$ .  
**return**  $(x_T^i)_{i=1}^N$ .

---

#### 4.1 Coin Sampling

We begin by reviewing the general coin betting framework [22, 68, 69]. Consider a gambler with initial wealth  $w_0 = 1$  who makes bets on a series of adversarial coin flips  $c_t \in \{-1, 1\}$ , where  $+1$  denotes heads and  $-1$  denotes tails. We encode the gambler's bet by  $x_t \in \mathbb{R}$ . In particular,  $\text{sign}(x_t) \in \{-1, 1\}$  denotes whether the bet is on heads or tails, and  $|x_t| \in \mathbb{R}$  denotes the size of the bet. Thus, in the  $t^{\text{th}}$  round, the gambler wins  $|x_t c_t|$  if  $\text{sign}(c_t) = \text{sign}(x_t)$  and loses  $|x_t c_t|$  otherwise. Finally, we write  $w_t$  for the wealth of the gambler at the end of the  $t^{\text{th}}$  round. The gambler's wealth thus accumulates as  $w_t = 1 + \sum_{s=1}^t c_s x_s$ . We assume that the gambler's bets satisfy  $x_t = \beta_t w_{t-1}$ , for a betting fraction  $\beta_t \in [-1, 1]$ , which means the gambler does not borrow any money.

In [68], the authors show how this approach can be used to solve convex optimisation problems on  $\mathbb{R}^d$ , i.e.,  $x^* = \arg \min_{x \in \mathbb{R}^d} f(x)$  for convex  $f : \mathbb{R}^d \rightarrow \mathbb{R}$ . In particular, [68] consider a betting game with 'outcomes'  $c_t = -\nabla f(x_t)$ , replacing scalar multiplications  $c_t x_t$  by scalar products  $\langle c_t, x_t \rangle$ . In this case, under certain assumptions, they show that  $f(\frac{1}{T} \sum_{t=1}^T x_t) \rightarrow f(x^*)$  at a rate determined by the betting strategy. There are many suitable choices for the betting strategy. Perhaps the simplest of these is  $\beta_t = t^{-1} \sum_{s=1}^{t-1} c_s$ , known as the Krichevsky-Trofimov (KT) betting strategy [52, 68], which yields the sequence of bets  $x_t = \beta_t w_{t-1} = t^{-1} \sum_{s=1}^{t-1} c_s (1 + \sum_{s=1}^{t-1} \langle c_s, x_s \rangle)$ .

In [78], this approach was extended to solve optimisation problems on the space of probability measures. In this case, several further modifications are necessary. First, in each round, one now bets  $x_t - x_0$  rather than  $x_t$ , where  $x_0$  is distributed according to some initial betting distribution  $\mu_0 \in \mathcal{P}_2(\mathbb{R}^d)$ . In this case, viewing  $x_t : \mathbb{R}^d \rightarrow \mathbb{R}^d$  as a function that maps  $x_0 \mapsto x_t(x_0)$ , one can define a betting distribution  $\mu_t \in \mathcal{P}_2(\mathbb{R}^d)$  as the push-forward of  $\mu_0$  under the function  $x_t$ . Under this definition, if  $x_0 \sim \mu_0$ , then  $x_t := x_t(x_0)$  is distributed according to the betting distribution  $\mu_t$ .

Using these modifications in the original coin betting framework, it is now possible to solve optimisation problems of the form  $\mu^* = \arg \min_{\mu \in \mathcal{P}_2(\mathcal{X})} \mathcal{F}(\mu)$ . In particular, by considering a betting game with  $c_t = -\nabla_{W_2} \mathcal{F}(\mu_t)(x_t)$ , [78] obtain the so-called coin Wasserstein gradient descent algorithm, which corresponds to a learning-rate free analogue of the more standard Wasserstein gradient descent. Alternatively, setting  $c_t = -P_{\mu_t, k} \nabla_{W_2} \mathcal{F}(\mu_t)$ , one can obtain a learning-rate free analogue of the population limit of SVGD, in which the updates are given by

$$x_t = x_0 - \frac{\sum_{s=1}^{t-1} P_{\mu_s, k} \nabla_{W_2} \mathcal{F}(\mu_s)(x_s)}{t} \left( 1 - \sum_{s=1}^{t-1} \langle P_{\mu_s, k} \nabla_{W_2} \mathcal{F}(\mu_s)(x_s), x_s - x_0 \rangle \right) \quad (9)$$

where, as defined above,  $\mu_t = (x_t)_{\#} \mu_0$ . In practice, the densities  $(\mu_t)_{t \in \mathbb{N}_0}$  and thus  $(c_t)_{t \in \mathbb{N}}$  are almost always intractable, and we will approximate them using a set of interacting particles.

#### 4.2 Mirrored Coin Sampling

By leveraging the formulation of constrained sampling as a 'mirrored' optimisation problem over the space of probability measures, we can extend the coin sampling approach to constrained domains. In particular, as an alternative to discretising the MWGF to solve this optimisation problem (Sec. 3.2), we propose instead to use a mirrored version of coin Wasserstein gradient descent (Sec. 4.1). The idea is to use coin Wasserstein gradient descent to solve the optimisation problem in the dual space, and then map back to the primal space using the mirror map. This suggests the following scheme.

---

**Algorithm 3** Coin MIED

---

**Input:** target density  $\pi$ , mollifier  $\phi_\epsilon$ , particles  $\omega_0^n = (x_0^i)_{i=1}^N$ .

**for**  $t = 1, \dots, T$  **do**

$$\text{For } i \in [N], x_t^i = x_0^i - \frac{\sum_{s=1}^{t-1} \nabla_{x_s^i} \log E_\epsilon(\omega_s^n)}{t} \left( 1 - \sum_{s=1}^{t-1} \langle \nabla_{x_s^i} \log E_\epsilon(\omega_s^n), x_s^i - x_0^i \rangle \right).$$

**return**  $(x_T^i)_{i=1}^N$ .

---

Initialise at  $y_0 = \nabla \phi(x_0) \in \mathbb{R}^d$ , where  $x_0 \in \mathcal{X}$ . Then, update

$$y_t = y_0 + \frac{\sum_{s=1}^{t-1} c_{\eta_s}(y_s)}{t} \left( 1 + \sum_{s=1}^{t-1} \langle c_{\eta_s}(y_s), y_s - y_0 \rangle \right), \quad x_t = \nabla \phi^*(x_t) \quad (10)$$

with  $\eta_t = (y_t)_\# \eta_0$  and  $\mu_t = (\nabla \phi^*)_\# \eta_t$ , and where  $c_{\eta_s} = -\nabla_{W_2} \mathcal{F}(\eta_s)$ , or some variant thereof. We will refer to this approach as mirrored coin sampling, or mirrored coin Wasserstein gradient descent. In order to obtain an implementable algorithm, we proceed as follows. First, decide on a suitable  $\mathcal{F}$  and a suitable  $(c_{\eta_t})_{t \in \mathbb{N}_0}$ , e.g.,  $\mathcal{F}(\eta) = \text{KL}(\eta|\nu)$ , and  $(c_{\eta_t})_{t \in \mathbb{N}_0} = (-\mathcal{P}_{\eta_t, k_\phi} \nabla_{W_2} \mathcal{F}(\eta_t))_{t \in \mathbb{N}_0}$ . Then, substitute these into (10), and approximate the updates using a set of interacting particles, as in existing ParVIs algorithms. Following these steps, we obtain a learning rate free analogue of MSVGD [79], which we will refer to as Coin MSVGD (Alg. 2). For different choices of  $(c_{\eta_t})_{t \in \mathbb{N}_0}$  in (10), we also obtain learning rate free analogues of two other mirrored ParVI algorithms we introduced in App. B, namely mirrored LAWGD (App. B.3), and mirrored KSDD (App. B.4). These algorithms, which we term Coin MLAWGD and Coin MKSDD, are given in Alg. 6 and Alg. 7 in App. D.

Even in the unconstrained setting, establishing the convergence of coin Wasserstein gradient descent remains an open problem. In [78], the authors provide a technical sufficient condition under which it is possible to establish convergence to the target measure in the population limit; however, it is difficult to verify this condition in general. In the interest of completeness, in App. C, we provide an analogous convergence result, adapted appropriately to the constrained setting.

### 4.3 Alternative Approaches

In this section, we outline an alternative approach, based on a ‘coinification’ of the mollified interaction energy descent (MIED) method recently introduced in [57]. MIED is based on minimising a function known as the mollified interaction energy (MIE)  $\mathcal{E}_\epsilon : \mathcal{P}(\mathbb{R}^d) \rightarrow [0, \infty]$ . In order to obtain a practical sampling algorithm, [57] propose to minimise a discrete version of the logarithmic MIE, using, e.g., Adagrad [30] or Adam [47]. That is, minimising the function

$$\log E_\epsilon(\omega^n) := \log \frac{1}{N^2} \sum_{i=1}^n \sum_{j=1}^n \phi_\epsilon(x^i - x^j) (\pi(x^i) \pi(x^j))^{-\frac{1}{2}}, \quad (11)$$

where  $\omega^n = \{x^1, \dots, x^n\}$ , and where  $(\phi_\epsilon)_{\epsilon > 0}$  is a family of mollifiers (see Definition 3.1 in [57]). This approach can be adapted to handle the constrained case, as outlined in [57]. In particular, if there exists a differentiable  $f : \mathbb{R}^d \rightarrow \mathcal{X}$  such that  $\mathcal{X} \setminus f(\mathbb{R}^d)$  has Lebesgue measure zero, then one can reduce the constrained problem to an unconstrained one by minimising  $\log E_\epsilon(f(\omega^n))$ , where  $f(\omega^n) = \{f(x^1), \dots, f(x^n)\}$ . Alternatively, if  $\mathcal{X} = \{x \in \mathbb{R}^d : g(x) \leq 0\}$  for some differentiable  $g : \mathbb{R}^d \rightarrow \mathbb{R}$ , then one can use a variant of the dynamic barrier method in [36].

As an alternative to the optimisation methods considered in [57], we instead propose to use a learning-rate free algorithm to minimise (11), based on coin betting. We term the resulting method, summarised in Alg. 3, Coin MIED. In comparison to our mirrored coin sampling algorithm, this approach has the advantage that it only requires a differentiable (not necessarily bijective) map.

## 5 Related Work

**Sampling as Optimisation.** The connection between sampling and optimisation has a long history, dating back to the work of [45], who established that the evolution of the law of the Langevin diffusion corresponds to the WGF of the KL divergence. In recent years, there has been renewed interest in

this perspective [96]. In particular, the viewpoint of existing algorithms such as LMC [23, 24, 32, 33] and SVGD [31, 61] as discretisations of WGFs has proved fruitful in deriving sharp convergence rates [5, 33, 51, 75, 80, 82]. It has also resulted in the rapid development of new sampling algorithms, inspired by ideas from the optimisation literature such as proximal methods [97], coordinate descent [28, 29], Nesterov acceleration [18, 25, 65], Newton methods [66, 81, 95], and coin betting [77, 78].

**Sampling in Constrained Domains.** Recent interest in constrained sampling problems has resulted in a range of general-purpose algorithms [42, 57, 79, 99, 100]. In cases where it is possible to explicitly parameterise the target domain in a lower dimensional space, one can use variants of classical methods, including LMC [95], rejection sampling [27], Hamiltonian Monte Carlo (HMC) [13], and Riemannian manifold HMC [54, 70]. Some other approaches are based on the use of mirror mapping. In particular, [42] proposed mirrored Langevin dynamics and its first-order discretisation. More recently, [99] introduced the mirrored Langevin diffusion and the mirror Langevin algorithm, which has since also been studied in [1, 19, 43, 58]. Along the same lines, [79] introduced two variants of SVGD suitable for constrained domains based on a mirror map; see also [83]. In certain cases (e.g., [55, 56]) one cannot explicitly parameterise the manifold of interest. In this setting, different approaches are required. [12] introduced a constrained version of HMC for this case. More recently, two ParVI algorithms were introduced for this setting, namely, O-SVGD [100] and MIED [57].

## 6 Numerical Results

In this section, we perform an empirical evaluation of Coin MSVGD (Sec. 6.1 - 6.2) and Coin MIED (Sec. 6.3). We compare our methods to their learning-rate-dependent analogues, namely, MSVGD and MIED. We also include a comparison with Stein variational mirror descent (SVMD) [79], and projected version of Coin SVGD and SVGD that ensure the iterates remain within the domain. We provide additional experimental details in App. F and additional numerical results in App. G.

### 6.1 Simplex Targets

Following [79], we first test the performance of our algorithms on two 20-dimensional targets defined on the simplex: the sparse Dirichlet posterior of [70] and the quadratic simplex target of [1]. We employ the IMQ kernel and the entropic mirror map [6]; and use  $N = 50$  particles,  $T = 500$  iterations, and repeat each experiment 5 times with different random seeds. In Fig. 1 and Fig. 5 - 6 (App. F.1), we plot the energy distance [84] to a set of surrogate ground truth samples, either obtained i.i.d. (sparse Dirichlet target) or using the No-U-Turn Sampler (NUTS) [41] (quadratic target). For the sparse Dirichlet posterior, Coin MSVGD converges more rapidly than MSVGD and SVMD, even for well-chosen values of the learning rate (Fig. 5a in App. F.1). For the quadratic target, Coin MSVGD generally converges more rapidly than MSVGD but not as fast as SVMD [79], which takes advantage of the log-concavity of the target in the primal space (Fig. 5a in App. F.1). After 500 iterations, Coin MSVGD has comparable performance to MSVGD and SVMD with optimal learning rates but significantly outperforms both algorithms for sub-optimal learning rates (Fig. 1).

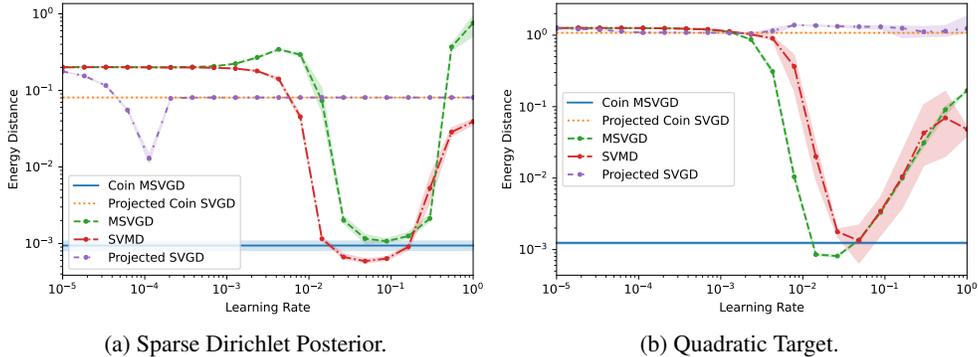


Figure 1: **Results for the simplex targets in [70] and [1].** Posterior approximation quality for Coin MSVGD, MSVGD, projected Coin SVGD, projected SVGD, and SVMD.

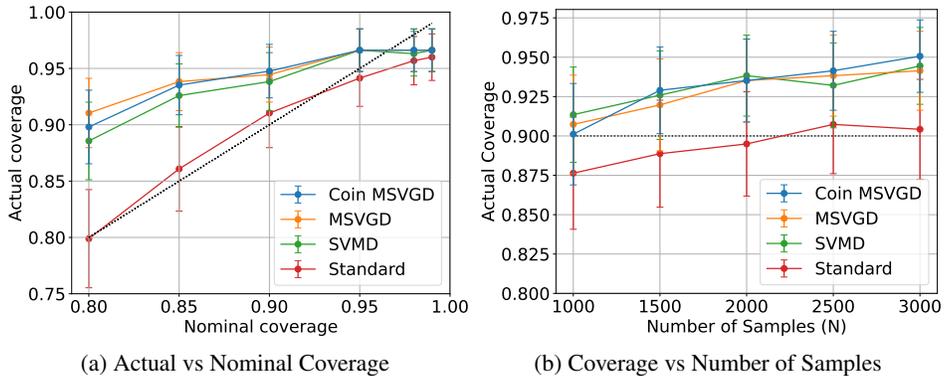


Figure 2: **Coverage results for the post-selection inference.** Coverage of post-selection confidence intervals for Coin MSVGD, MSVGD, and SVMD.

## 6.2 Confidence Intervals for Post Selection Inference

We next consider a constrained sampling problem arising in post-selection inference [53, 85]. Suppose we are given data  $(X, y) \in \mathbb{R}^{n \times p} \times \mathbb{R}^n$ . We are interested in obtaining valid confidence intervals (CIs) for regression coefficients obtained via the randomised Lasso [86], defined as the solution of  $\hat{\beta} = \arg \min_{\beta \in \mathbb{R}^p} [\frac{1}{2} \|y - X\beta\|_2^2 + \lambda \|\beta\|_1 - \omega^\top \beta + \frac{\varepsilon}{2} \|\beta\|_2^2]$ , where  $\lambda, \varepsilon \in \mathbb{R}_+$  are penalty parameters and  $\omega \in \mathbb{R}^p$  is an auxiliary random vector. Let  $\hat{\beta}_E \in \mathbb{R}^q$  represent the non-zero coefficients of  $\hat{\beta}$  and  $\hat{z}_E = \text{sign}(\hat{\beta}_E)$  for their signs. If the support  $E$  were known a priori, we could determine CIs for  $\beta_E$  using the asymptotic normality of  $\beta_E = (X_E^\top X_E)^{-1} X_E^\top y$ . However, when  $E$  is based on data, these ‘classical’ CIs will no longer be valid. In this case, one approach is to condition on the result of the initial model selection, i.e., knowledge of  $E$  and  $\hat{z}_E$ . In practice, this means sampling from the so-called selective distribution, which has support  $\mathcal{D} = \{(\hat{\beta}_E, \hat{z}_{-E}) : \text{sign}(\hat{\beta}_E) = \hat{z}_E, \|\hat{z}_{-E}\|_\infty \leq 1\}$ , and density proportional to [87]

$$\hat{g}(\hat{\beta}_E, \hat{z}_{-E}) \propto g\left(\varepsilon(\hat{\beta}_E, 0)^\top - X^\top(y - X_E \hat{\beta}_E) + \lambda(\hat{z}_E, \hat{z}_{-E})^\top\right). \quad (12)$$

We first consider the model setup described in [76, Section 5.3]; see App. F.2 for full details. In Fig. 8 (App. G.3), we plot the energy distance between Coin MSVGD and projected Coin SVGD to a set of 1000 samples obtained using NUTS [41] for the selective distribution of the randomized Lasso in which two features are selected. Similar to our previous examples, the performance of MSVGD is very sensitive to the choice of learning rate. Meanwhile, Coin MSVGD converges rapidly to the true target, with no need to tune a learning rate. Similar to the examples considered in Sec. 6.1, the projected methods fail to converge to the true target, highlighting the need for the mirrored approach.

In Fig. 2, we plot the coverage of the CIs obtained using Coin MSVGD, MSVGD, SVMD, and the norejection algorithm in `selectiveInference` [88], as we vary either the nominal coverage or the total number of samples. For MSVGD and SVMD, we use RMSProp [89] to adapt the learning rate, with an initial learning rate of  $\gamma = 0.1$ , following [79]. As the nominal coverage varies, Coin MSVGD achieves similar coverage to MSVGD and SVMD; and significantly higher actual coverage than norejection (Fig. 2a). Meanwhile, as the number of samples varies, Coin MSVGD, MSVGD, and SVMD consistently all obtain a higher coverage than the fixed nominal coverage of 90% (Fig. 2b). This is important for small sample sizes, where the standard approach undercovers. The performance of MSVGD and SVMD, of course, is highly dependent on an appropriate learning rate. In Fig. 9 (App. G.3), we provide additional plots illustrating how the coverage of the CIs obtained using MSVGD and SVMD can deteriorate when the learning-rate is chosen sub-optimally.

We next consider a post-selection inference problem involving the HIV-1 drug resistance dataset studied in [7, 72]; see App. F.2 for full details. The goal is to identify statistically significant mutations associated with the response to the drug Lamivudine (3TC). The randomised Lasso selects a subset of mutations, for which we compute 90% CIs using 5000 samples and five methods: ‘unadjusted’ (the unadjusted CIs), ‘standard’ (the method in `selectiveInference`), MSVGD, SVMD, and Coin MSVGD. Our results indicate that the CIs obtained by Coin MSVGD are similar to those obtained

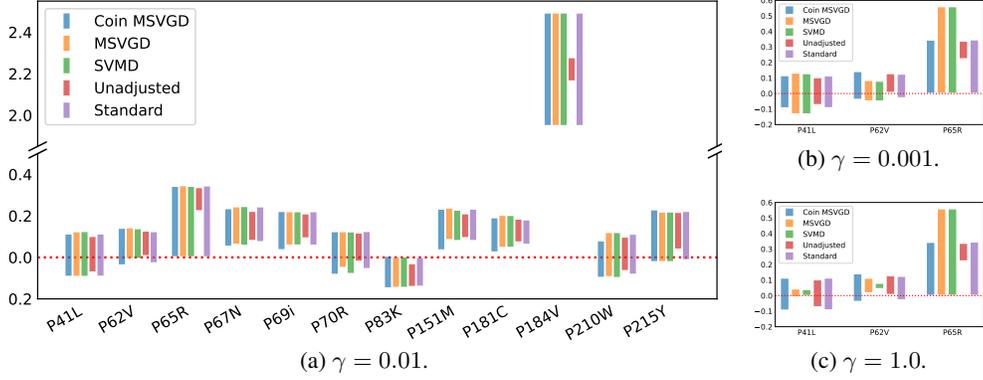


Figure 3: **Real data results for post-selection inference.** Confidence intervals for the mutations selected by the randomised Lasso as candidates for HIV-1 drug resistance. We report (a) all mutations when MSVGD and SVMD use a well chosen learning rate ( $\gamma = 0.01$ ) and (b) - (c) a subset of mutations when MSVGD and SVMD use a smaller ( $\gamma = 0.001$ ) or larger ( $\gamma = 1.0$ ) learning rate.

via the ‘standard’ approach, which we view as a benchmark, as well as by MSVGD and SVMD (with a well chosen learning rate). Importantly, they differ from the CIs obtained using the unadjusted approach (see, e.g., mutation P65R or P184V in Fig. 3). Meanwhile, for sub-optimal choices of the learning rate (Fig. 3b and Fig. 3c, Fig. 10 in App. G.3), the CIs obtained using MSVGD and SVMD differ substantially from those obtained using the baseline ‘standard’ approach and by Coin MSVGD, once more highlighting the sensitivity of these methods to the choice of an appropriate learning rate.

### 6.3 Fairness Bayesian Neural Network

Finally, following [57, 62, 63, 67], we train a fairness Bayesian neural network to predict whether an individual’s income is greater than \$50,000, with gender as a protected characteristic. We use the Adult Income dataset [49]. The dataset is of the form  $\mathcal{D} = \{x_i, y_i, z_i\}_{i=1}^n$ , where  $x_i$  denote the feature vectors,  $y_i$  denote the labels (i.e., whether the income is greater than \$50,000), and  $z_i$  denote the protected attribute (i.e., the gender). We train a two-layer Bayesian neural network  $\hat{y}(x; w)$  with weights  $w$  and place a standard Gaussian prior on each weight independently. The fairness constraint is given by  $g(\theta) = \text{Cov}_{(x,y,z) \sim \mathcal{D}}[z, \hat{y}(x; \theta)]^2 - t \leq 0$  for some  $t > 0$ . In testing, we evaluate each method using a Calder-Verwer (CV) score [14], a standard measure of disparate impact. We run each method for  $T = 2000$  iterations and use  $N = 50$  particles. In Fig. 4a, we plot the trade-off curve between test accuracy and CV score, for  $t \in \{10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}, 2 \times 10^{-3}, 5 \times 10^{-3}, 10^{-2}\}$ . We compare the results for Coin MIED, MIED, and the methods in [63], using the default implementations. In this experiment, Coin MIED is clearly preferable to the other methods, achieving a much larger Pareto front. In Fig. 4b and Fig. 4c, we plot the constraint and the MIE versus training iterations for both coin MIED and MIED. Once again, coin MIED tends to outperform MIED, achieving both lower values of the constraint and lower values of the energy.

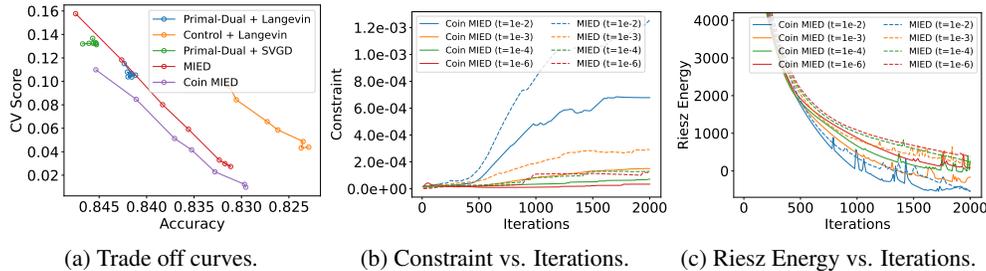


Figure 4: **Results for the fairness Bayesian neural network.** (a) Trade-off curves showing disparate impact versus test accuracy for MIED, coin MIED, and the methods in [63]. (b) - (c) The constraint  $\text{Cov}_{(x,y,z) \sim \mathcal{D}}[z, \hat{y}(x; \theta)]^2$  and the MIE versus iterations, for MIED and Coin MIED.

## 7 Discussion

**Summary.** In this paper, we introduce several new particle-based algorithms for constrained sampling which are entirely learning rate free. Our first approach is based on the coin sampling framework introduced in [78], and a perspective of constrained sampling as a mirrored optimisation problem on the space of probability measures, which we also use to unify several existing constrained sampling algorithms, and study their theoretical properties. Our second approach can be viewed as the coin sampling analogue of the recently proposed MIED algorithm [57]. Empirically, our algorithms achieve comparable or superior performance to other particle-based constrained sampling algorithms, with no need to tune a learning rate.

**Limitations.** We highlight three limitations. First, like any mirrored sampling algorithm, mirrored coin sampling (e.g., Coin MSVGD), necessarily depends on the availability of a mirror map that can appropriately capture the constraints of the problem at hand. Second, Coin MSVGD has a cost of  $O(N^2)$  per update. Finally, establishing the convergence of mirrored coin sampling under more standard conditions, e.g., a mirrored log-Sobolev inequality, remains an open problem. We leave this as an interesting direction for future work.

## References

- [1] Kwangjun Ahn and Sinho Chewi. Efficient constrained sampling via the mirror-Langevin algorithm. In *Proceedings of the 35th International Conference on Neural Information Processing Systems (NeurIPS 2021)*, Online, 2021. 1, 2, 7, 19, 20, 30, 31, 32
- [2] Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient Flows: In Metric Spaces and in the Space of Probability Measures*. Birkhäuser, Basel, 2008. 2, 16, 22, 27
- [3] Michael Arbel, Anna Korba, Adil Salim, and Arthur Gretton. Maximum Mean Discrepancy Gradient Flow. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems (NeurIPS 2019)*, Vancouver, Canada, 2019. 3
- [4] Dominique Bakry, Ivan Gentil, and Michel Ledoux. *Analysis and Geometry of Markov Diffusion Operators*. Springer Cham, 2014. 24
- [5] Krishnakumar Balasubramanian, Sinho Chewi, Murat A. Erdogdu, Adil Salim, and Matthew Zhang. Towards a Theory of Non-Log-Concave Sampling: First-Order Stationarity Guarantees for Langevin Monte Carlo. In *Proceedings of the 35th Annual Conference on Learning Theory (COLT 2022)*, London, UK, 2022. 7
- [6] Amir Beck and Marc Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, 2003. 7
- [7] Nan Bi, Jelena Markovic, Lucy Xia, and Jonathan Taylor. Inferactive data analysis. *Scandinavian Journal of Statistics*, 47(1):212–249, 2020. 8
- [8] David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: a review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017. 1
- [9] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. Latent Dirichlet Allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003. 1, 30
- [10] Sergey G Bobkov. Isoperimetric and Analytic Inequalities for Log-Concave Probability Measures. *The Annals of Probability*, 27(4):1903–1921, 1999. 19
- [11] Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng. *Handbook of Markov Chain Monte Carlo*. Chapman and Hall/CRC, 1st edition, 2011. 1
- [12] Marcus Brubaker, Mathieu Salzmann, and Raquel Urtasun. A Family of MCMC Methods on Implicitly Defined Manifolds. In *Proceedings of the 15th International Conference on Artificial Intelligence and Statistics (AISTATS)*, La Palma, Canary Islands, 2012. 7
- [13] Simon Byrne and Mark Girolami. Geodesic Monte Carlo on Embedded Manifolds. *Scandinavian Journal of Statistics*, 40(4):825–845, 2013. 7

- [14] Toon Calders and Sicco Verwer. Three Naive Bayes Approaches for Discrimination-Free Classification. *Data Min. Knowl. Discov.*, 21(2):277–292, 2010. 9
- [15] Gilles Celeux, Mohammed El Anbari, Jean-Michel Marin, and Christian P Robert. Regularization in Regression: Comparing Bayesian and Frequentist Methods in a Poorly Informative Situation. *Bayesian Analysis*, 7(2):477–502, 2012. 1
- [16] Peng Chen and Omar Ghattas. Projected Stein Variational Gradient Descent. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (NeurIPS 2020)*, Vancouver, Canada, 2020. 1
- [17] Yuansi Chen. An Almost Constant Lower Bound of the Isoperimetric Coefficient in the KLS Conjecture. *Geometric and Functional Analysis*, 31(1):34–61, 2021. 19
- [18] Xiang Cheng, Niladri S. Chatterji, Peter L. Bartlett, and Michael I. Jordan. Underdamped Langevin MCMC: A non-asymptotic analysis. In *Proceedings of the 31st Annual Conference on Learning Theory (COLT 2018)*, volume 75, pages 1–24, Stockholm, Sweden, 2018. 7
- [19] Sinho Chewi, Thibaut Le Gouic, Chen Lu, Tyler Maunu, Philippe Rigollet, and Austin Stromme. Exponential ergodicity of mirror-Langevin diffusions. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (NeurIPS 2020)*, Online, 2020. 1, 7, 18, 19, 20
- [20] Sinho Chewi, Thibaut Le Gouic, Chen Lu, Tyler Maunu, and Philippe Rigollet. SVGD as a kernelized Wasserstein gradient flow of the chi-squared divergence. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (NeurIPS 2020)*, pages 2098–2109, 2020. 3, 22, 23, 24
- [21] Kacper Chwialkowski, Heiko Strathmann, and Arthur Gretton. A Kernel Test of Goodness of Fit. In *Proceedings of the 33rd International Conference on Machine Learning (ICML 2016)*, New York, NY, 2016. 24
- [22] Ashok Cutkosky and Francesco Orabona. Black-Box Reductions for Parameter-free Online Learning in Banach Spaces. In *Proceedings of the 31st Annual Conference on Learning Theory (COLT 2018)*, Stockholm, Sweden, 2018. 5
- [23] Arnak S Dalalyan. Theoretical guarantees for approximate sampling from smooth and log-concave densities. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 79(3):651–676, 2017. 7
- [24] Arnak S Dalalyan and Avetik Karagulyan. User-friendly guarantees for the Langevin Monte Carlo with inaccurate gradient. *Stochastic Processes and their Applications*, 129(12):5278–5311, 2019. 7
- [25] Arnak S Dalalyan and Lionel Riou-Durand. On sampling from a log-concave density using kinetic Langevin diffusions. *Bernoulli*, 26(3):1956–1988, 2020. 7
- [26] Gianluca Detommaso, Tiangang Cui, Alessio Spantini, Youssef Marzouk, and Robert Scheichl. A Stein variational Newton method. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems (NIPS 2018)*, Montreal, Canada, 2018. 1
- [27] Persi Diaconis, Susan Holmes, and Mehrdad Shahshahani. Sampling from a manifold. *Advances in modern statistical theory and applications: a Festschrift in honor of Morris L. Eaton*, 10:102–125, 2013. 7
- [28] Zhiyan Ding, Qin Li, Jianfeng Lu, and Stephen J. Wright. Random Coordinate Langevin Monte Carlo. In *Proceedings of the 34th Annual Conference on Learning Theory (COLT 2021)*, volume 134, pages 1–28, Boulder, CO, 2021. 7
- [29] Zhiyan Ding, Qin Li, Jianfeng Lu, and Stephen J. Wright. Random Coordinate Underdamped Langevin Monte Carlo. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics (AISTATS 2021)*, San Diego, CA, 2021. 7

- [30] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(61):2121–2159, 2011. [6](#), [29](#)
- [31] Andrew Duncan, Nikolas Nüsken, and Lukasz Szpruch. On the geometry of Stein variational gradient descent. *Journal of Machine Learning Research*, 24:1–40, 2023. [7](#), [20](#), [21](#)
- [32] Alain Durmus, Szymon Majewski, and Blazej Miasojedow. Analysis of Langevin Monte Carlo via Convex Optimization. *Journal of Machine Learning Research*, 20(1):2666–2711, 2019. [7](#)
- [33] Alain Durmus and Éric Moulines. High-dimensional Bayesian inference via the unadjusted Langevin algorithm. *Bernoulli*, 25(4A):2854–2882, 2019. [7](#)
- [34] Murat A. Erdogdu and Rasa Hosseinzadeh. On the Convergence of Langevin Monte Carlo: The Interplay between Tail Growth and Smoothness. In *Proceedings of the 34th Annual Conference on Learning Theory (COLT 2021)*, Boulder, CO, 2021. [19](#)
- [35] Damien Garreau, Wittawat Jitkrittum, and Motonobu Kanagawa. Large sample analysis of the median heuristic. *arXiv preprint*, 2017. [29](#)
- [36] Chengyue Gong, Xingchao Liu, and Qiang Liu. Bi-objective trade-off with dynamic barrier gradient descent. In *Proceedings of the 35th International Conference on Neural Information Processing Systems (NeurIPS 2021)*, Online, 2021. [6](#)
- [37] Jackson Gorham and Lester Mackey. Measuring Sample Quality with Kernels. In *Proceedings of the 34th International Conference on Machine Learning (ICML 2017)*, Sydney, Australia, 2017. [24](#), [29](#)
- [38] T H Gronwall. Note on the Derivatives with Respect to a Parameter of the Solutions of a System of Differential Equations. *Annals of Mathematics*, 20(4):292–296, 1919. [18](#)
- [39] Richard Grumitt, Biwei Dai, and Uros Seljak. Deterministic Langevin Monte Carlo with Normalizing Flows for Bayesian Inference. In *Proceedings of the 36th International Conference on Neural Information Processing Systems (NeurIPS 2022)*, New Orleans, LA, 2022. [3](#), [19](#)
- [40] Matthew D. Hoffman, David M. Blei, Chong Wang, and John Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 14(1):1303–1347, 2013. [1](#)
- [41] Matthew D Hoffman and Andrew Gelman. The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(1):1593–1623, 2014. [7](#), [8](#), [31](#), [33](#)
- [42] Ya-Ping Hsieh, Ali Kavis, Paul Rolland, and Volkan Cevher. Mirrored Langevin Dynamics. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems (NIPS 2018)*, Montreal, Canada, 2018. [1](#), [2](#), [3](#), [7](#), [16](#), [17](#), [19](#), [20](#), [22](#), [24](#), [27](#)
- [43] Qijia Jiang. Mirror Langevin Monte Carlo: the Case Under Isoperimetry. In *Proceedings of the 35th International Conference on Neural Information Processing Systems (NeurIPS 2021)*, Online, 2021. [1](#), [7](#), [19](#), [20](#)
- [44] Valen E. Johnson and James H. Albert. *Ordinal Data Modeling*. Springer, New York, NY, 1999. [1](#)
- [45] 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. [3](#), [6](#)
- [46] Ravi Kannan, László Lovász, and Miklós Simonovits. Isoperimetric problems for convex bodies and a localization lemma. *Discrete & Computational Geometry*, 13(3):541–559, 1995. [19](#)
- [47] Diederik P. Kingma and Jimmy Ba. Adam: a method for stochastic optimisation. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR '15)*, pages 1–13, San Diego, CA, 2015. [6](#), [29](#)

- [48] John P. Klein and Melvin L. Moeschberger. *Survival Analysis: Techniques for Censored and Truncated Data*. Springer, New York, NY, 2003. 1
- [49] Ron Kohavi. Scaling up the Accuracy of Naive-Bayes Classifiers: A Decision-Tree Hybrid. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, KDD'96, pages 202–207. AAAI Press, 1996. 9
- [50] Anna Korba, Pierre-Cyril Aubin-Frankowski, Szymon Majewski, and Pierre Ablin. Kernel Stein Discrepancy Descent. In *Proceedings of the 38th International Conference on Machine Learning (ICML 2021)*, Online, 2021. 3, 4, 20, 24, 25, 26
- [51] Anna Korba, Adil Salim, Michael Arbel, Giulia Luise, and Arthur Gretton. A Non-Asymptotic Analysis for Stein Variational Gradient Descent. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (NeurIPS 2020)*, Vancouver, Canada, 2020. 4, 7, 21
- [52] Raphael E Krichevsky and Victor K Trofimov. The performance of universal encoding. *IEEE Transactions on Information Theory*, 27(2):199–207, 1981. 5
- [53] Jason D Lee, Dennis L Sun, Yuekai Sun, and Jonathan E Taylor. Exact post-selection inference, with application to the lasso. *The Annals of Statistics*, 44(3):907–927, 2016. 1, 8
- [54] Yin Tat Lee and Santosh S Vempala. Convergence Rate of Riemannian Hamiltonian Monte Carlo and Faster Polytope Volume Computation. In *Proceedings of the 50th Annual ACM Symposium on the Theory of Computing (STOC 2018)*, pages 1115–1121, Los Angeles, CA, 2018. 7
- [55] Tony Lelièvre, Mathias Rousset, and Gabriel Stoltz. *Free Energy Computations*. Imperial College Press, jun 2010. 7
- [56] Tony Lelièvre and Gabriel Stoltz. Partial differential equations and stochastic methods in molecular dynamics. *Acta Numerica*, 25:681–880, 2016. 7
- [57] Lingxiao Li, Qiang Liu, Anna Korba, Mikhail Yurochkin, and Justin Solomon. Sampling with Mollified Interaction Energy Descent. In *Proceedings of the 11th International Conference on Learning Representations*, Kigali, Rwanda, 2023. 1, 2, 6, 7, 9, 10, 29, 31
- [58] Ruilin Li, Molei Tao, Santosh S. Vempala, and Andre Wibisono. The Mirror Langevin Algorithm Converges with Vanishing Bias. In *Proceedings of the 33rd International Conference on Algorithmic Learning Theory (ALT 2022)*, Paris, France, 2022. 1, 7, 19
- [59] Qiang Liu. Stein variational gradient descent as gradient flow. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS 2017)*, pages 3118–3126, Red Hook, NY, 2017. 20
- [60] Qiang Liu, Jason D. Lee, and Michael Jordan. A Kernelized Stein Discrepancy for Goodness-of-fit Tests. In *Proceedings of the 33rd International Conference on Machine Learning (ICML 2016)*, New York, NY, 2016. 24, 25
- [61] Qiang Liu and Dilin Wang. Stein Variational Gradient Descent: A General Purpose Bayesian Inference Algorithm. In *Proceedings of the 30th Conference on Neural Information Processing Systems (NIPS 2016)*, Barcelona, Spain, 2016. 1, 3, 7, 29
- [62] Suyun Liu and Luis Nunes Vicente. Accuracy and fairness trade-offs in machine learning: a stochastic multi-objective approach. *Computational Management Science*, 19(3):513–537, 2022. 9, 31
- [63] Xingchao Liu, Xin T. Tong, and Qiang Liu. Sampling with Trustworthy Constraints: A Variational Gradient Framework. In *Proceedings of the 35th International Conference on Neural Information Processing Systems (NeurIPS 2021)*, Online, 2021. 1, 9, 31
- [64] Yang Liu, Prajit Ramachandran, Qiang Liu, and Jian Peng. Stein Variational Policy Gradient. In *Proceedings of the Conference on Uncertainty In Artificial Intelligence (UAI 2017)*, Sydney, Australia, 2017. 1

- [65] Yi-An Ma, Niladri S Chatterji, Xiang Cheng, Nicolas Flammarion, Peter L Bartlett, and Michael I Jordan. Is there an analog of Nesterov acceleration for gradient-based MCMC? *Bernoulli*, 27(3):1942–1992, may 2021. 7
- [66] James Martin, Lucas C Wilcox, Carsten Burstedde, and Omar Ghattas. A Stochastic Newton MCMC Method for Large-Scale Statistical Inverse Problems with Application to Seismic Inversion. *SIAM Journal on Scientific Computing*, 34(3):1460–1487, 2012. 7
- [67] Natalia Martinez, Martin Bertran, and Guillermo Sapiro. Minimax Pareto Fairness: A Multi Objective Perspective. In *Proceedings of the 37th International Conference on Machine Learning (ICML 2020)*, Online, 2020. 9
- [68] Francesco Orabona and David Pal. Coin Betting and Parameter-Free Online Learning. In *Proceedings of the 30th Conference on Neural Information Processing Systems (NIPS 2016)*, Barcelona, Spain, 2016. 5
- [69] Francesco Orabona and Tatiana Tommasi. Training Deep Networks without Learning Rates Through Coin Betting. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS 2017)*, Long Beach, CA, 2017. 5, 28
- [70] Sam Patterson and Yee Whye Teh. Stochastic gradient Riemannian Langevin dynamics on the probability simplex. In *Proceedings of the 27th International Conference on Neural Information Processing Systems (NIPS 2013)*, volume 26, Lake Tahoe, NV, 2013. 7, 30, 31
- [71] Yunchen Pu, Zhe Gan, Ricardo Henao, Chunyuan Li, Shaobo Han, and Lawrence Carin. VAE Learning via Stein Variational Gradient Descent. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS 2017)*, Long Beach, CA, 2017. 1
- [72] Soo-Yon Rhee, Jonathan Taylor, Gauhar Wadhera, Asa Ben-Hur, Douglas L Brutlag, and Robert W Shafer. Genotypic predictors of human immunodeficiency virus type 1 drug resistance. *Proceedings of the National Academy of Sciences*, 103(46):17355–17360, 2006. 8
- [73] Christian P Robert and George Casella. *Monte Carlo Statistical Methods*. Springer-Verlag, New York, 2 edition, 2004. 1
- [74] Ralph Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, NJ, 1970. 3
- [75] Adil Salim, Lukang Sun, and Peter Richtárik. A Convergence Theory for SVGD in the Population Limit under Talagrand’s Inequality T1. In *Proceedings of the 39th International Conference on Machine Learning (ICML 2022)*, Online, 2022. 4, 7
- [76] Amir Sepehri and Jelena Markovic. Non-reversible, tuning- and rejection-free Markov chain Monte Carlo via iterated random functions. *arXiv preprint*, 2017. 8, 30
- [77] Louis Sharrock, Daniel Dodd, and Christopher Nemeth. CoinEM: Tuning-Free Particle-Based Variational Inference for Latent Variable Models. *arXiv preprint*, 2023. 7
- [78] Louis Sharrock and Christopher Nemeth. Coin Sampling: Gradient-Based Bayesian Inference without Learning Rates. *To appear in Proceedings of the 40th International Conference on Machine Learning (ICML 2023)*, 2023. 2, 4, 5, 6, 7, 10, 27, 28
- [79] Jiaxin Shi, Chang Liu, and Lester Mackey. Sampling with Mirrored Stein Operators. In *Proceedings of the 10th International Conference on Learning Representations (ICLR 2022)*, Online, 2022. 1, 2, 3, 4, 6, 7, 8, 20, 21, 23, 29, 30, 31
- [80] Jiaxin Shi and Lester Mackey. A Finite-Particle Convergence Rate for Stein Variational Gradient Descent. *arXiv preprint*, 2022. 7
- [81] Umut Simsekli, Roland Badeau, A. Taylan Cemgil, and Gaël Richard. Stochastic Quasi-Newton Langevin Monte Carlo. In *Proceedings of the 33rd International Conference on Machine Learning (ICML 2016)*, New York, NY, 2016. 7
- [82] Lukang Sun, Avetik Karagulyan, and Peter Richtarik. Convergence of Stein Variational Gradient Descent under a Weaker Smoothness Condition. *arXiv preprint*, 2022. 4, 7, 21

- [83] Lukang Sun and Peter Richtárik. A Note on the Convergence of Mirrored Stein Variational Gradient Descent under (L0,L1)-Smoothness Condition. *arXiv preprint*, 2022. 1, 2, 3, 4, 7, 20, 21, 22
- [84] Gábor J. Székely and Maria L. Rizzo. Energy statistics: A class of statistics based on distances. *Journal of Statistical Planning and Inference*, 143(8):1249–1272, 2013. 7
- [85] Jonathan Taylor and Robert J Tibshirani. Statistical learning and selective inference. *Proceedings of the National Academy of Sciences*, 112(25):7629–7634, 2015. 1, 8
- [86] Xiaoying Tian, Nan Bi, and Jonathan Taylor. MAGIC: a general, powerful and tractable method for selective inference. *arXiv preprint*, 2016. 8, 30
- [87] Xiaoying Tian, Snigdha Panigrahi, Jelena Markovic, Nan Bi, and Jonathan Taylor. Selective sampling after solving a convex problem. *arXiv preprint*, 2016. 8, 30
- [88] Ryan Tibshirani, Rob Tibshirani, Jonatha Taylor, Joshua Loftus, Stephen Reid, and Jelena Markovic. selectiveInference: Tools for Post-Selection Inference, 2019. 8, 33
- [89] Tijmen Tieleman and Geoffrey E Hinton. Lecture 6.5-rmsprop: divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2):26–31, 2012. 8, 29, 31, 32, 33
- [90] Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer, New York, NY, 2009. 19, 23
- [91] Santosh S. Vempala and Andre Wibisono. Rapid Convergence of the Unadjusted Langevin Algorithm: Isoperimetry Suffices. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems (NeurIPS 2019)*, Vancouver, Canada, 2019. 19
- [92] Martin J Wainwright and Michael I Jordan. Graphical Models, Exponential Families, and Variational Inference. *Foundations and Trends in Machine Learning*, 1(1-2):1–305, 2008. 1
- [93] Dilin Wang, Ziyang Tang, Chandrajit Bajaj, and Qiang Liu. Stein Variational Gradient Descent with Matrix-Valued Kernels. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems (NeurIPS 2019)*, Vancouver, Canada, 2019. 1
- [94] Dilin Wang, Zhe Zeng, and Qiang Liu. Stein Variational Message Passing for Continuous Graphical Models. In *Proceedings of the 35th International Conference on Machine Learning (ICML 2018)*, Stockholm, Sweden, 2018. 1
- [95] Xiao Wang, Qi Lei, and Ioannis Panageas. Fast Convergence of Langevin Dynamics on Manifold: Geodesics meet Log-Sobolev. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (NeurIPS 2020)*, Vancouver, Canada, 2020. 7
- [96] Andre Wibisono. Sampling as optimization in the space of measures: The Langevin dynamics as a composite optimization problem. In *Proceedings of the 31st Annual Conference on Learning Theory (COLT 2018)*, Stockholm, Sweden, 2018. 2, 3, 7, 19
- [97] Andre Wibisono. Proximal Langevin Algorithm: Rapid Convergence Under Isoperimetry. *arXiv preprint*, 2019. 7
- [98] Edwin B Wilson. Probable Inference, the Law of Succession, and Statistical Inference. *Journal of the American Statistical Association*, 22(158):209–212, 1927. 30
- [99] Kelvin Shuangjian Zhang, Gabriel Peyré, Jalal Fadili, and Marcelo Pereyra. Wasserstein Control of Mirror Langevin Monte Carlo. In *Proceedings of the 33rd Annual Conference on Learning Theory (COLT 2020)*, Graz, Austria, 2020. 1, 7, 19
- [100] Ruqi Zhang, Qiang Liu, and Xin T. Tong. Sampling in Constrained Domains with Orthogonal-Space Variational Gradient Descent. In *Proceedings of the 36th International Conference on Neural Information Processing Systems (NeurIPS 2022)*, New Orleans, LA, 2022. 1, 7
- [101] Jingwei Zhuo, Chang Liu, Jiaxin Shi, Jun Zhu, Ning Chen, and Bo Zhang. Message Passing Stein Variational Gradient Descent. In *Proceedings of the 35th International Conference on Machine Learning (ICML 2018)*, Stockholm, Sweden, 2018. 1

## A Calculus in the Wasserstein Space

We recall the following from [2, Chapter 10]. Suppose  $\mu \in \mathcal{P}_2(\mathbb{R}^d)$  and  $\xi \in L^2(\mu)$ . Let  $\mathcal{F}$  be a proper and lower semi-continuous functional on  $\mathcal{P}_2(\mathbb{R}^d)$ . We say that  $\xi \in L^2(\mu)$  belongs to the Fréchet subdifferential of  $\mathcal{F}$  at  $\mu$  and write  $\xi \in \partial\mathcal{F}(\mu)$  if, for any  $\nu \in \mathcal{P}_2(\mathbb{R}^d)$ ,

$$\liminf_{\nu \rightarrow \mu} \frac{\mathcal{F}(\nu) - \mathcal{F}(\mu) - \int_{\mathbb{R}^d} \langle \xi(x), \mathbf{t}_\mu^\nu(x) - x \rangle \mu(dx)}{W_2(\nu, \mu)} \geq 0, \quad (13)$$

where  $\mathbf{t}_\mu^\nu : \mathbb{R}^d \rightarrow \mathbb{R}^d$  denotes the optimal transport map from  $\mu$  to  $\nu$  [2, Chapter 7.1]. Under mild conditions [2, Lemma 10.4.1], the subdifferential  $\partial\mathcal{F}$  is single-valued,  $\partial\mathcal{F}(\mu) = \{\nabla_{W_2}\mathcal{F}(\mu)\}$ , and  $\nabla_{W_2}\mathcal{F}(\mu) \in L^2(\mu)$  is given by

$$\nabla_{W_2}\mathcal{F}(\mu) = \nabla \frac{\partial\mathcal{F}(\mu)}{\partial\mu}(x) \quad \text{for } \mu\text{-a.e. } x \in \mathbb{R}^d, \quad (14)$$

where  $\frac{\partial\mathcal{F}(\mu)}{\partial\mu} : \mathbb{R}^d \rightarrow \mathbb{R}$  denotes the first variation of  $\mathcal{F}$  at  $\mu$ , that is, the unique function such that

$$\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} (\mathcal{F}(\mu + \varepsilon\zeta) - \mathcal{F}(\mu)) = \int_{\mathbb{R}^d} \frac{\partial\mathcal{F}(\mu)}{\partial\mu}(x) \zeta(dx), \quad (15)$$

where  $\zeta = \nu - \mu$ , and  $\nu \in \mathcal{P}_2(\mathbb{R}^d)$ . We will refer to  $\nabla_{W_2}\mathcal{F}(\mu)$  as the Wasserstein gradient of  $\mathcal{F}$  at  $\mu$ .

## B Examples of Mirrored Wasserstein Gradient Flows

In this section, we outline several algorithms which arise as special cases of the continuity equation introduced in Sec 3.2, namely

$$\frac{\partial\eta_t}{\partial t} + \nabla \cdot (w_t \eta_t) = 0 \quad , \quad \mu_t = (\nabla\phi^*)_{\#}\eta_t, \quad (16)$$

where  $(w_t)_{t \geq 0}$  are some sequence of vector fields chosen to ensure convergence of  $\eta_t$  to  $\nu = (\nabla\phi)_{\#}\pi$ .

### B.1 Mirrored Langevin Dynamics

Suppose that  $\mathcal{F}(\eta) = \text{KL}(\eta|\nu)$ , where  $\nu = (\nabla\phi)_{\#}\pi$  is the mirrored (or dual) target as defined in Sec. 3.1. In addition, let  $w_t = -\nabla_{W_2}\mathcal{F}(\eta_t)$ . In this case, we have that  $\nabla_{W_2}\mathcal{F}(\eta) = \nabla \log\left(\frac{\eta}{\nu}\right) = \nabla \log \eta - \nabla \log \nu = \nabla \log \eta + \nabla W$  [e.g. 2, Chapter 10], where in the final equality we have introduced the notation  $W$  for the potential of  $\nu$ .

**Remark 1.** *Under the assumption that the target  $\pi \propto e^{-V}$  and the mirrored target  $\nu = (\nabla\phi)_{\#}\pi \propto e^{-W}$ , the Monge-Ampere equation determines the relationship between  $W$  and  $V$ . Namely, we have that [e.g., 42]*

$$e^{-V} = e^{-W \circ \nabla\phi} \det \nabla^2 \phi \quad , \quad e^{-W} = e^{-V \circ \nabla\phi^*} \det \nabla^2 \phi^*. \quad (17)$$

Thus, in particular, the potential  $W(y)$  of the mirrored target  $\nu$  evaluated at  $y = \nabla\phi(x)$  can be expressed in terms of the potential  $V(x)$  of the original target  $\pi$  evaluated at  $x$ , as

$$W(y) = V(x) + \log \det \nabla^2 \phi(x). \quad (18)$$

For the choice of  $\mathcal{F}$  and  $(w_t)_{t \geq 0}$  described above, the mirrored Wasserstein gradient flow in (16) now reads

$$\frac{\partial\eta_t}{\partial t} = \nabla \cdot \left( \eta_t \nabla \log \frac{\eta_t}{\nu} \right) = \nabla \cdot (\eta_t \nabla W + \nabla \eta_t) \quad , \quad \mu_t = (\nabla\phi^*)_{\#}\eta_t. \quad (19)$$

The PDE in (19) is nothing more than the Fokker-Planck equation describing the evolution of the law of the overdamped Langevin stochastic differential equation (SDE) with respect to the mirrored target  $\nu \propto e^{-W}$ . That is to say, if  $y_0 \sim \eta_0$ , and  $(y_t)_{t \geq 0}$  is described by the SDE

$$dy_t = -\nabla W(y_t) dt + \sqrt{2} db_t \quad , \quad x_t = \nabla\phi^*(y_t), \quad (20)$$

where  $b = (b_t)_{t \geq 0}$  is a standard Brownian motion, then, if we define  $(\eta_t)_{t \geq 0}$  as the law of  $(y_t)_{t \geq 0}$ ,  $(\eta_t)_{t \geq 0}$  satisfies (19). The equation in (20) is precisely the mirrored Langevin dynamics (MLD) introduced in [42]. We now study the properties of these dynamics in continuous time, starting with the dissipation of  $\text{KL}(\mu_t|\pi)$  and  $\chi^2(\mu_t|\pi)$  along the MLD. These results are a natural extension of existing results in the unconstrained case to our setting.

**Proposition 1.** *The dissipation of  $\text{KL}(\cdot|\pi)$  along the mirrored Wasserstein gradient flow in (19) is given by*

$$\frac{d\text{KL}(\mu_t|\pi)}{dt} = - \left\| [\nabla^2 \phi]^{-1} \nabla \log \left( \frac{\mu_t}{\pi} \right) \right\|_{L^2(\mu_t)}^2. \quad (21)$$

*Proof.* Using differential calculus in the Wasserstein space and the chain rule, we have that

$$\frac{d\text{KL}(\eta_t|\nu)}{dt} = \left\langle w_t, \nabla \log \left( \frac{\eta_t}{\nu} \right) \right\rangle_{L^2(\eta_t)} = - \left\| \nabla \log \left( \frac{\eta_t}{\nu} \right) \right\|_{L^2(\eta_t)}^2. \quad (22)$$

By [42, Theorem 2], we have that  $\text{KL}(\eta|\nu) = \text{KL}(\mu|\pi)$ . To deal with the term on the RHS, first note that, using the formula for the change of variable in a probability density, we have that

$$\eta(y) = \mu(\nabla \phi^*(y)) |\det \nabla^2 \phi(\nabla \phi^*(y))|^{-1}, \quad \nu(y) = \pi(\nabla \phi^*(y)) |\det \nabla^2 \phi(\nabla \phi^*(y))|^{-1}. \quad (23)$$

Thus, in particular, if  $y = \nabla \phi(x)$ , then we have that

$$\log \left( \frac{\eta}{\nu} \right) (y) = \log \left( \frac{\mu}{\pi} \right) (\nabla \phi^*(y)) = \log \left( \frac{\mu}{\pi} \right) (x). \quad (24)$$

Using this, and the fact that  $\eta = (\nabla \phi)_\# \mu$ , we can now compute

$$\left\| \nabla \log \left( \frac{\eta_t}{\nu} \right) \right\|_{L^2(\eta_t)}^2 = \int \left\| \nabla_y \log \left( \frac{\eta_t}{\nu} \right) (y) \right\|^2 d\eta_t(y) \quad (25)$$

$$= \int \left\| \nabla_y \log \left( \frac{\eta_t}{\nu} \right) (y) \right\|^2 d((\nabla \phi)_\# \mu_t)(y) \quad (26)$$

$$= \int \left\| \nabla_y \log \left( \frac{\eta_t}{\nu} \right) (\nabla \phi(x)) \right\|^2 d\mu_t(x) \quad (27)$$

$$= \int \left\| \nabla_y \log \left( \frac{\mu_t}{\pi} \right) (x) \right\|^2 d\mu_t(x) \quad (28)$$

$$= \int \left\| [\nabla^2 \phi(x)]^{-1} \nabla_x \log \left( \frac{\mu_t}{\pi} \right) (y) \right\|^2 d\mu_t(y) \quad (29)$$

$$= \left\| [\nabla^2 \phi(y)]^{-1} \nabla_y \log \left( \frac{\mu_t}{\pi} \right) (y) \right\|_{L^2(\mu_t)}^2. \quad (30)$$

The conclusion now follows.  $\square$

**Proposition 2.** *The dissipation of  $\chi^2(\cdot|\pi)$  along the mirrored Wasserstein gradient flow in (19) is given by*

$$\frac{d\chi^2(\mu_t|\pi)}{dt} = -2 \left\| [\nabla^2 \phi]^{-1} \nabla \frac{\mu_t}{\pi} \right\|_{L^2(\pi)}^2. \quad (31)$$

*Proof.* The proof is rather similar to the previous one. In this case, using the fact that  $\nabla_{W_2} \chi^2(\eta|\nu) = 2\nabla \frac{\eta}{\nu}$ , we have

$$\frac{d\chi^2(\eta_t|\nu)}{dt} = \left\langle -\nabla \log \left( \frac{\eta_t}{\nu} \right), 2\nabla \frac{\eta_t}{\nu} \right\rangle_{L^2(\eta_t)} = -2 \left\| \nabla \frac{\eta_t}{\nu} \right\|_{L^2(\nu)}^2.$$

We next observe, using now the fact that  $\frac{\eta}{\nu}(y) = \frac{\mu}{\pi}(x)$  which follows from (23), that

$$\chi^2(\eta_t|\nu) = \int \left( \frac{\eta_t}{\nu}(y) - 1 \right)^2 d\nu(y) = \int \left( \frac{\eta_t}{\nu}(y) - 1 \right)^2 d((\nabla \phi)_\# \pi)(y) \quad (32)$$

$$= \int \left( \frac{\eta_t}{\nu}(\nabla \phi(x)) - 1 \right)^2 d\pi(x) = \int \left( \frac{\mu_t}{\pi}(x) - 1 \right)^2 d\pi(x) = \chi^2(\mu_t|\pi). \quad (33)$$

In addition, based on a very similar argument, we have that

$$\left\| \nabla \left( \frac{\eta_t}{\nu} \right) \right\|_{L^2(\nu)}^2 = \int \left\| \nabla_y \left( \frac{\eta_t}{\nu} \right) (y) \right\|^2 d\nu(y) = \int \left\| \nabla_y \left( \frac{\eta_t}{\nu} \right) (y) \right\|^2 d((\nabla \phi)_\# \pi)(y) \quad (34)$$

$$= \int \left\| \nabla_y \left( \frac{\eta_t}{\nu} \right) (\nabla \phi(x)) \right\|^2 d\pi(x) = \int \left\| \nabla_y \left( \frac{\mu_t}{\pi} \right) (x) \right\|^2 d\pi(x) \quad (35)$$

$$= \int \left\| [\nabla^2 \phi(y)]^{-1} \nabla_y \left( \frac{\mu_t}{\pi} \right) (y) \right\|^2 d\pi(y) = \left\| [\nabla^2 \phi(y)]^{-1} \nabla_y \left( \frac{\mu_t}{\pi} \right) (y) \right\|_{L^2(\pi)}^2. \quad (36)$$

This completes the proof.  $\square$

Since the RHS of both (21) and (31) are non-positive, these results show that  $\text{KL}(\mu_t|\pi)$  and  $\chi^2(\mu_t|\pi)$  decrease along the MLD. As an immediate corollary of these results, we have the following continuous time convergence rate for the time-average of  $\|[\nabla^2\phi]^{-1}\nabla\log(\frac{\mu_t}{\pi})\|_{L^2(\mu_t)}^2$  and  $\|[\nabla^2\phi]^{-1}\nabla(\frac{\mu_t}{\pi})\|_{L^2(\pi)}^2$ .

**Proposition 3.** *For any  $t \geq 0$ , it holds that*

$$\min_{0 \leq s \leq t} \left\| [\nabla^2\phi]^{-1}\nabla\log\left(\frac{\mu_t}{\pi}\right) \right\|_{L^2(\eta_t)}^2 \leq \frac{1}{t} \int_0^t \left\| [\nabla^2\phi]^{-1}\nabla\log\left(\frac{\mu_s}{\pi}\right) \right\|_{L^2(\mu_s)}^2 ds \leq \frac{\text{KL}(\mu_0|\pi)}{t}. \quad (37)$$

**Proposition 4.** *For any  $t \geq 0$ , it holds that*

$$\min_{0 \leq s \leq t} \left\| [\nabla^2\phi]^{-1}\nabla\left(\frac{\mu_t}{\pi}\right) \right\|_{L^2(\eta_t)}^2 \leq \frac{1}{t} \int_0^t \left\| [\nabla^2\phi]^{-1}\nabla\left(\frac{\mu_s}{\pi}\right) \right\|_{L^2(\mu_s)}^2 ds \leq \frac{\chi^2(\mu_0|\pi)}{t}. \quad (38)$$

*Proof.* These results follow straightforwardly from integration of (21) or (31) and the non-negativity of the KL or  $\chi^2$  divergence.  $\square$

To obtain stronger convergence results, we will impose additional assumptions on the target measure. In the first case, we can assume that the target satisfies a *mirrored log-Sobolev inequality*.

**Proposition 5.** *Assume that  $\pi$  satisfies the following mirrored log-Sobolev inequality for some  $\lambda > 0$ , viz*

$$\text{KL}(\mu|\pi) \leq \frac{1}{2\lambda} \left\| [\nabla^2\phi]^{-1}\nabla\log\left(\frac{\mu}{\pi}\right) \right\|_{L^2(\mu)}^2. \quad (39)$$

*Then the KL divergence converges exponentially along the mirrored Wasserstein gradient flow in (19):*

$$\text{KL}(\mu_t|\pi) \leq e^{-2\lambda t} \text{KL}(\mu_0|\pi). \quad (40)$$

*Proof.* Substituting the log-Sobolev inequality (39) into (21), we have that

$$\frac{d\text{KL}(\mu_t|\pi)}{dt} = - \left\| [\nabla^2\phi]^{-1}\nabla\log\left(\frac{\mu_t}{\pi}\right) \right\|_{L^2(\mu_t)}^2 \leq -2\lambda \text{KL}(\mu_t|\pi). \quad (41)$$

The conclusion now follows straightforwardly from Grönwall's inequality [38].  $\square$

**Remark 2.** *The assumption in Proposition 5 is precisely equivalent to the assumption that the mirrored target  $\nu = (\nabla\phi)_{\#}\pi$  satisfies the classical log-Sobolev inequality. This assumption holds, for example, for all strongly log-concave distributions, with constant equal to the inverse of the strong convexity constant of the potential. It is worth emphasising that this is distinct from the so-called mirror log-Sobolev inequality, used in, e.g., [19] to establish convergence of the mirror Langevin diffusion, which is distinct from the mirrored Langevin diffusion that we study here. We provide some additional details on this approach at the end of this section.*

Alternatively, under the assumption that the target satisfies a *mirrored Poincaré inequality*, we can obtain exponential convergence in a number of metrics.

**Proposition 6.** *Assume that  $\pi$  satisfies the following mirrored Poincaré inequality for some  $\kappa > 0$ , viz*

$$\text{Var}_{\pi} \left[ \left( \frac{d\mu_t}{d\pi} - 1 \right) \right] \leq \frac{1}{\kappa} \left\| [\nabla^2\phi]^{-1}\nabla \left( \frac{d\mu_t}{d\pi} - 1 \right) \right\|_{L^2(\pi)}^2. \quad (42)$$

*Then the total variation distance, Hellinger distance, KL divergence, and the  $\chi^2$  divergence all converge exponentially along the mirrored Wasserstein gradient flow in (19):*

$$\max(2\|\mu_t - \pi\|_{\text{TV}}^2, H^2(\mu_t, \pi), \frac{\lambda}{2}W_2^2(\mu_t, \pi), \text{KL}(\mu_t|\pi), \chi^2(\mu_t|\pi)) \leq e^{-2\kappa t} \chi^2(\mu_0|\pi). \quad (43)$$

*Proof.* The proof proceeds similarly to above. In particular, this time substituting the Poincaré inequality (42) into (22), we have

$$\frac{d\chi^2(\mu_t|\pi)}{dt} = -2 \left\| [\nabla^2\phi]^{-1}\nabla\frac{\mu_t}{\pi} \right\|_{L^2(\pi)}^2 \leq -2\kappa\chi^2(\mu_t|\pi).$$

The result for the  $\chi^2$  divergence follows via an application of Grönwall's inequality. The remaining bounds follow from standard comparison inequalities (e.g., Section 2.4 in [90]).  $\square$

**Remark 3.** *The assumption used in Proposition 6 is equivalent to the assumption that the mirrored target  $\nu = (\nabla\phi)_{\#}\pi$  satisfies the classical Poincaré inequality. The class of measures satisfying this inequality is rather large, including all strongly log-concave measures (due to the Bakry-Émery criterion) and, more generally, all log-concave measures [10, 17, 46].*

In practice, of course, it is necessary to discretise the dynamics in (20) in time. Applying a standard Euler-Maruyama discretisation to (20), one arrives at the *mirrored Langevin algorithm* (MLA) from [42], viz

$$y_{t+1} = y_t - h\nabla W(y_t) + \sqrt{2h}\xi_t \quad (44)$$

$$x_{t+1} = \nabla\phi^*(y_{t+1}), \quad (45)$$

where  $h > 0$  is the step size or learning rate and where  $(\xi_t)_{t \in \mathbb{N}}$  are a sequence of i.i.d. standard normal random variables. Alternatively, one also arrives at the same algorithm by considering the so-called forward-flow discretisation of the mirrored WGF in (19); see, e.g., [96] in the unconstrained setting. In particular, this corresponds to

$$\eta_{t+\frac{1}{2}} = (\text{id} - \gamma\nabla W)_{\#}\eta_t \quad , \quad \eta_{t+1} = \mathcal{N}(0, 2\gamma\text{id}) * \eta_{t+\frac{1}{2}} \quad (46)$$

where  $*$  denotes the convolution and where, as elsewhere,  $\mu_{t+1} = (\nabla\phi^*)_{\#}\eta_{t+1}$ . From an Eulerian perspective, the discretisation in (46) is precisely equivalent to

$$y_{t+\frac{1}{2}} = y_t - \gamma\nabla W(y_t) \quad , \quad y_{t+1} = y_{t+\frac{1}{2}} + \sqrt{2\gamma}\xi_t, \quad (47)$$

with  $x_{t+1} = \nabla\phi^*(y_{t+1})$ . Rewriting (47) as a single update equation, we arrive back at the mirrored Langevin algorithm in (44) - (45). In this case, we have the following convergence rate.

**Proposition 7.** *Suppose that  $\pi$  satisfies the mirrored log-Sobolev inequality with constant  $\lambda > 0$ . In addition, suppose that  $\pi$  is mirror- $L$ -smooth.<sup>2</sup> Then, if  $0 < \gamma < \frac{\lambda}{4L^2}$ , the iterates  $(x_t)_{t \geq 0}$  of the mirrored Langevin algorithm satisfy*

$$\text{KL}(\mu_t || \pi) \leq e^{-\lambda\gamma t} \text{KL}(\mu_0 || \pi) + \frac{8\gamma dL^2}{\lambda}. \quad (48)$$

*Proof.* By [42, Theorem 2], and (25) - (30) in the proof of Proposition 1, if  $\pi$  satisfies a mirrored-log-Sobolev inequality with constant  $\lambda > 0$ , then  $\nu = (\nabla\phi)_{\#}\pi$  satisfies a logarithmic-Sobolev inequality with the same constant  $\lambda$ . In addition, if  $\pi$  is mirror- $L$ -smooth, then  $\nu$  is  $L$ -smooth. Thus, it follows from [91, Theorem 1] that

$$\text{KL}(\eta_k || \nu) \leq e^{-\lambda\gamma t} \text{KL}(\eta_0 || \nu) + \frac{8\gamma dL^2}{\lambda}. \quad (49)$$

Finally, using the invariance of the KL divergence under the mirror map  $\nabla\phi$  [42, Theorem 2], the result follows.  $\square$

**Remark 4.** *Using a similar approach, one can also obtain bounds on  $\chi^2(\mu_k || \pi)$  under a log-Sobolev inequality [34, Theorem 4] or for the Renyi divergence  $\mathcal{R}_q(\mu_t || \pi)$  under a log-Sobolev [91, Theorem 2] or a Poincaré [91, Theorem 3] inequality.*

**Remark 5.** *One could also consider a forward Euler discretisation of the dynamics in (19), which leads to a mirrored version of the so-called deterministic Langevin Monte Carlo algorithm [39].*

**Remark 6.** *It is worth distinguishing between the mirrored Langevin dynamics in (20) and the so-called mirror Langevin dynamics or mirror Langevin diffusion, which refers to the solution of*

$$dy_t = -\nabla V(x_t)dt + \sqrt{2} [\nabla^2\phi(x_t)]^{\frac{1}{2}} db_t \quad (50)$$

$$x_t = \nabla\phi^*(y_t). \quad (51)$$

*These dynamics and the corresponding Euler-Maruyama discretisation known as the mirror Langevin algorithm were first studied in [99], and further analysed in [1, 43, 58]. An alternative time discretisation has also been studied in [19].*

<sup>2</sup>We say that  $\pi \propto e^{-V}$  is mirror- $L$ -smooth if  $\nu \propto e^{-W}$  is  $L$ -smooth. This holds, in particular, if  $W = -V \circ \nabla\phi^* + \log \det \nabla^2\phi^*$  has Hessian bounded by  $L$ .

The law  $(\mu_t)_{t \geq 0}$  of the solution of this SDE satisfies the Fokker-Planck equation [43, App. A.2]

$$\frac{\partial \mu_t}{\partial t} = \nabla \cdot (\mu_t v_t), \quad v_t = [\nabla^2 \phi]^{-1} \nabla \log \frac{\mu_t}{\pi}. \quad (52)$$

Recalling that  $\nabla_{W_2} \text{KL}(\mu_t | \pi) = \nabla \log \frac{\mu_t}{\pi}$ , one can view (52) as a special case of the so-called Wasserstein mirror flow (WMF), which is defined by [e.g. 1, App. C]

$$\frac{\partial \mu_t}{\partial t} + \nabla \cdot (v_t \mu_t) = 0, \quad v_t = -(\nabla^2 \phi)^{-1} \nabla_{W_2} \mathcal{F}(\mu_t). \quad (53)$$

The WMF is rather different from the MWGF defined in Sec. 3.2. In particular, (53) describes the evolution of  $\mu_t$  according to the Wasserstein tangent vector  $-(\nabla^2 \phi)^{-1} \nabla_{W_2} \mathcal{F}(\mu_t)$ , rather than the evolution of the mirrored  $\eta_t := (\nabla \phi)_{\#} \mu_t$  according to the tangent vector  $\nabla_{W_2} \mathcal{F}(\eta_t)$ . We refer to [19] for a detailed discussion of the mirrored Langevin dynamics from this perspective.

## B.2 Mirrored Stein Variational Gradient Descent

We now turn our attention to the MSVGD dynamics given in Sec. 3.2.1, which are defined by the continuity equation

$$\frac{\partial \eta_t}{\partial t} + \nabla \cdot (w_t \eta_t) = 0, \quad w_t = -P_{\eta_t, k_\phi} \nabla \log \left( \frac{\eta_t}{\nu} \right) \quad (54)$$

with  $\mu_t = (\nabla \phi^*)_{\#} \eta_t$  and where  $k_\phi$  is the kernel defined according to  $k_\phi(y, y') = k(\nabla \phi^*(y), \nabla \phi^*(y'))$ , given some kernel  $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ . This is nothing more than the SVGD gradient flow [31, 50, 59] on the mirror (i.e., the dual) space. In this case, we have the following results. We note that the results can also found in [79, App. H]; we provide them here to highlight the analogues with the convergence results obtained for the MLD in Sec. B.1.

**Proposition 8.** *The dissipation of the KL along the mirrored SVGD gradient flow is given by*

$$\frac{d\text{KL}(\mu_t | \pi)}{dt} = - \left\| S_{\mu_t, k} [\nabla^2 \phi]^{-1} \nabla \log \left( \frac{\mu_t}{\pi} \right) \right\|_{\mathcal{H}_k^d}^2. \quad (55)$$

*Proof.* The proof is similar to the proof of Proposition 1. In particular, using the chain rule, we have

$$\frac{d\text{KL}(\eta_t | \nu)}{dt} = \left\langle w_t, \nabla \log \left( \frac{\eta_t}{\nu} \right) \right\rangle_{L^2(\eta_t)} = - \left\| S_{\eta_t, k_\phi} \nabla \log \left( \frac{\eta_t}{\nu} \right) \right\|_{\mathcal{H}_{k_\phi}^d}^2. \quad (56)$$

By Theorem 2 in [42], we have that  $\text{KL}(\eta_t | \nu) = \text{KL}(\mu_t | \pi)$ . In addition, once more appealing to (24), we can compute

$$\left\| S_{\eta_t, k_\phi} \nabla \log \left( \frac{\eta_t}{\nu} \right) \right\|_{\mathcal{H}_{k_\phi}^d}^2 \quad (57)$$

$$= \int \int k(\nabla \phi^*(y), \nabla \phi^*(y')) \left\langle \nabla_y \log \left( \frac{\eta_t}{\nu} \right) (y), \nabla_{y'} \log \left( \frac{\eta_t}{\nu} \right) (y') \right\rangle d\eta_t(y) d\eta_t(y') \quad (58)$$

$$= \int \int k(x, x') \left\langle \nabla_y \log \left( \frac{\mu_t}{\pi} \right) (x), \nabla_{y'} \log \left( \frac{\mu_t}{\pi} \right) (x') \right\rangle d\mu_t(x) d\mu_t(x') \quad (59)$$

$$= \int \int k(x, x') \left\langle [\nabla^2 \phi(x)]^{-1} \nabla_x \log \left( \frac{\mu_t}{\pi} \right) (x), [\nabla^2 \phi(x')]^{-1} \nabla_{x'} \log \left( \frac{\mu_t}{\pi} \right) (x') \right\rangle d\mu_t(x) d\mu_t(x') \quad (60)$$

$$= \left\| S_{\mu_t, k} [\nabla^2 \phi]^{-1} \nabla \log \left( \frac{\mu_t}{\pi} \right) \right\|_{\mathcal{H}_k^d}^2,$$

from which the result then follows.  $\square$

**Remark 7.** *The quantity on the RHS of (55) is referred to in [83] as the mirrored Stein Fisher information of  $\mu$  with respect to  $\pi$ , with mirror map  $\nabla \phi$ . As noted in [83], it is quite natural to interpret this quantity as the Stein Fisher information [31] between the mirrored target  $(\nabla \phi)_{\#} \mu$  and  $(\nabla \phi)_{\#} \pi$ , given the kernel  $k_\phi(\cdot, \cdot) = k(\nabla \phi^*(\cdot), \nabla \phi^*(\cdot))$ .*

Similar to before, a straightforward consequence of this result is the following continuous time convergence rate for the average of the mirrored stein discrepancy along the continuous-time MSVGD dynamics.

**Proposition 9.** For all  $t \geq 0$ , it holds that

$$\min_{0 \leq s \leq t} \left\| S_{\mu_s, k} [\nabla^2 \phi]^{-1} \nabla \log \left( \frac{\mu_s}{\pi} \right) \right\|_{\mathcal{H}_k^d}^2 \leq \frac{1}{t} \int_0^t \left\| S_{\mu_s, k} [\nabla^2 \phi]^{-1} \nabla \log \left( \frac{\mu_s}{\pi} \right) \right\|_{\mathcal{H}_k^d}^2 ds \quad (61)$$

$$\leq \frac{\text{KL}(\mu_0 | \pi)}{t}. \quad (62)$$

*Proof.* The result follows upon integrating (55), using also the non-negativity of the KL divergence.  $\square$

To obtain faster convergence guarantees for the continuous-time mirrored SVGD dynamics, we must once again assume some additional properties on the target. A natural choice in this case is the *mirrored Stein log-Sobolev inequality* (see also [79, Definition 2]), which represents the natural analogue of the Stein log-Sobolev inequality introduced in [31] (see also [51]) in the mirrored setting. As we now show, this assumption guarantees that the KL divergence converges exponentially fast along the SVGD dynamics.

**Proposition 10.** Assume that  $\pi$  satisfies the mirrored Stein log-Sobolev inequality for some  $\lambda > 0$ , viz

$$\text{KL}(\mu | \pi) \leq \frac{1}{2\lambda} \left\| S_{\mu, k} [\nabla^2 \phi]^{-1} \nabla \log \left( \frac{\mu}{\pi} \right) \right\|_{\mathcal{H}_k^d}^2. \quad (63)$$

Then

$$\text{KL}(\mu_t | \pi) \leq e^{-2\lambda t} \text{KL}(\mu_0 | \pi). \quad (64)$$

*Proof.* The result is a straightforward consequence of (55) and (63). In particular, substituting (63) into (55), we have

$$\frac{d\text{KL}(\mu_t | \pi)}{dt} = - \left\| S_{\mu_t, k} [\nabla^2 \phi]^{-1} \nabla \log \left( \frac{\mu_t}{\pi} \right) \right\|_{\mathcal{H}_k^d}^2 \leq 2\lambda \text{KL}(\mu_t | \pi). \quad (65)$$

We conclude by applying Grönwall's inequality.  $\square$

**Remark 8.** The assumption of a mirrored Stein log-Sobolev inequality is equivalent to the assumption that the mirrored target  $\nu = (\nabla \phi)_{\#} \pi$  satisfies the standard Stein log-Sobolev inequality. The Stein log-Sobolev inequality is rather less well known than the classical log-Sobolev inequality. It is now well known that this condition fails to hold if the kernel  $k$  is too regular with respect to the target  $\pi$ , e.g., if  $\pi$  has exponential tails and the derivatives of the kernel  $k$  and the potential  $V$  grow at most polynomially [31, 51]. It remains an open problem to establish conditions for which this conditions holds.

In discrete time, [83] recently established the following complexity bound for MSVGD in the population limit, based on an extension of the arguments in [82]. In particular, this result is obtained by applying the arguments in [82] in the mirror space, after imposing suitable assumptions on the mirrored target. For completeness, we recall this result here.

**Theorem 1** (Theorem 1 in [83]). Suppose that the following assumptions hold:

- (1) The mirror function  $\phi : \mathcal{X} \rightarrow (-\infty, +\infty)$  is strongly  $K$ -convex.
- (2) There exist  $B_1, B_2 > 0$  such that for all  $x \in \mathcal{X}$ ,  $k(x, x') \leq B_1^2$  and  $\partial_{x_i} \partial_{x_j} k(x, x')|_{x=x'} \leq B_2^2$  for all  $i \in [d], j \in [d]$ .
- (3) There exist  $L_0, L_1 \geq 0$  such that  $\|\nabla^2 W(y)\| \leq L_0 + L_1 \|\nabla W(y)\|$ .
- (4) There exists  $p \geq 1$ ,  $y_0 \in \mathbb{R}^d$ , and  $s > 0$  such that

$$\int \exp [s \|y - y_0\|^p] d\nu(y) < \infty. \quad (66)$$

- (5) There exists  $p \geq 1$  and a constant  $C_p$  such that  $\|\nabla W(y)\| \leq C_p (\|y\|^p + 1)$  for any  $y \in \mathbb{R}^d$ .

Then, provided the learning rate satisfies is sufficiently small - see (21) in [83] - then

$$\frac{1}{T} \sum_{t=0}^{T-1} \left\| S_{\mu_t, k} [\nabla^2 \phi]^{-1} \nabla \log \left( \frac{\mu_t}{\pi} \right) \right\|_{\mathcal{H}_k^d}^2 \leq \frac{2\text{KL}(\mu_0 || \pi)}{T\gamma}. \quad (67)$$

*Proof.* See Appendices A - D in [83].  $\square$

### B.3 Mirrored Laplacian Adjusted Wasserstein Gradient Descent

Based on the perspective in Sec. 3.2, we can also obtain new, mirrored versions of several other sampling algorithms and establish their asymptotic properties. We begin by introducing the MWGF of the chi-squared divergence; see also [20, Section 3] or [2, Theorem 11.2.1] in the unconstrained case. In particular, suppose that  $\mathcal{F}(\eta) = \chi^2(\eta || \nu)$  with  $\nabla_{W_2} \mathcal{F}(\eta) = \nabla \left( \frac{\eta}{\nu} \right)$ . Setting  $w_t = -\nabla_{W_2} \mathcal{F}(\eta_t)$  in (4), we then have

$$\frac{\partial \eta_t}{\partial t} + \nabla \cdot (\eta_t w_t) = 0 \quad , \quad w_t = -2\nabla \left( \frac{\eta_t}{\nu} \right) \quad (68)$$

with  $\mu_t = (\nabla \phi)_{\#} \eta_t$ . We then have the following results.

**Proposition 11.** *The dissipation of  $\text{KL}(\cdot || \pi)$  along the mirrored Wasserstein gradient flow in (68) is given by*

$$\frac{d\text{KL}(\mu_t || \pi)}{dt} = -2 \left\| [\nabla^2 \phi]^{-1} \nabla \frac{\mu_t}{\pi} \right\|_{L^2(\pi)}^2. \quad (69)$$

*Proof.* The proof is essentially identical to the proof of Proposition 2, although the roles of the Lyapunov functional and the Wasserstein gradient are now reversed. In particular, we now have

$$\frac{d\text{KL}(\eta_t || \nu)}{dt} = \left\langle 2\nabla \frac{\eta_t}{\nu}, -\nabla \log \left( \frac{\eta_t}{\nu} \right) \right\rangle_{L^2(\eta_t)} = -2 \left\| \nabla \frac{\eta_t}{\nu} \right\|_{L^2(\nu)}^2.$$

By Theorem 2 in [42] and (34) - (36) in the proof of Proposition 2, we have that  $\text{KL}(\eta_t || \nu) = \text{KL}(\mu_t || \pi)$  and  $\left\| \nabla \left( \frac{\eta_t}{\nu} \right) \right\|_{L^2(\nu)}^2 = \left\| [\nabla^2 \phi]^{-1} \nabla \left( \frac{\mu_t}{\pi} \right) \right\|_{L^2(\pi)}^2$ , from which the conclusion follows.  $\square$

**Proposition 12.** *The dissipation of  $\chi^2(\cdot || \pi)$  along the mirrored Wasserstein gradient flow in (68) is given by*

$$\frac{d\chi^2(\mu_t || \pi)}{dt} = -4 \left\| [\nabla^2 \phi]^{-1} \nabla \left( \frac{\mu_t}{\pi} \right) \right\|_{L^2(\mu_t)}^2. \quad (70)$$

*Proof.* Similar to the last result, the proof follows very closely the proof of Proposition 1. On this occasion, we have

$$\frac{d\chi^2(\eta_t || \nu)}{dt} = \left\langle w_t, 2\nabla \left( \frac{\eta_t}{\nu} \right) \right\rangle_{L^2(\eta_t)} = -4 \left\| \nabla \left( \frac{\eta_t}{\nu} \right) \right\|_{L^2(\eta_t)}^2. \quad (71)$$

Arguing as in (32) - (33) and (34) - (36) in the proof of Proposition 2, we have that  $\chi^2(\eta || \nu) = \chi^2(\mu || \pi)$  and  $\left\| \nabla \left( \frac{\eta_t}{\nu} \right) \right\|_{L^2(\eta_t)}^2 = \left\| [\nabla^2 \phi]^{-1} \nabla \left( \frac{\mu_t}{\pi} \right) \right\|_{L^2(\mu_t)}^2$ , from which the conclusion follows.  $\square$

Similar to before, we can also obtain exponential convergence to the target under a mirrored functional inequality. Here, we just consider a mirrored version of the Poincaré inequality.

**Proposition 13.** *Assume that  $\pi$  satisfies the following mirrored Poincaré inequality for some  $\kappa > 0$ , viz*

$$\text{Var}_{\pi} \left[ \left( \frac{\mu_t}{\pi} - 1 \right) \right] \leq \frac{1}{\kappa} \left\| [\nabla^2 \phi]^{-1} \nabla \left( \frac{\mu_t}{\pi} - 1 \right) \right\|_{L^2(\pi)}^2. \quad (72)$$

*Then the KL divergence converges exponentially along the mirrored Wasserstein gradient flow in (68):*

$$\text{KL}(\mu_t || \pi) \leq \text{KL}(\mu_0 || \pi) e^{-2\lambda t}. \quad (73)$$

*Furthermore, for  $t \geq \frac{1}{2\lambda}$ , the following convergence rate holds:*

$$\text{KL}(\mu_t || \pi) \leq (\text{KL}(\mu_0 || \pi) \wedge 2) e^{-2\lambda t}. \quad (74)$$

*Proof.* The proof follows closely the proof of [20, Theorem 1]. Substituting the mirrored Poincaré inequality (72) into (69), we have that

$$\frac{d\text{KL}(\mu_t|\pi)}{dt} = - \left\| [\nabla^2 \phi]^{-1} \nabla \left( \frac{\mu_t}{\pi} \right) \right\|_{L^2(\mu_t)}^2 \leq -2\lambda \chi^2(\mu_t|\pi) \leq -2\lambda \text{KL}(\mu_t|\pi). \quad (75)$$

where in the final inequality we have used the fact that  $\text{KL}(\cdot|\pi) \leq \chi^2(\cdot|\pi)$  (e.g., [90, Section 2.4]). The first bound now follows straightforwardly from Grönwall's inequality.

For the second bound, we will instead use the inequality  $\text{KL}(\cdot|\pi) \leq \log(1 + \chi^2(\cdot|\pi))$  (e.g., [90, Section 2.4]), which implies that  $e^{\text{KL}(\cdot|\pi)} - 1 \leq \chi^2(\cdot|\pi)$ . Using this in (75), it follows that  $\frac{d\text{KL}(\mu_t|\pi)}{dt} \leq -2\lambda(e^{\text{KL}(\mu_t|\pi)} - 1)$  or, equivalently, that

$$\frac{d}{dt} \left( 1 - e^{-\text{KL}(\mu_t|\pi)} \right) \leq -2\lambda \left( 1 - e^{-\text{KL}(\mu_t|\pi)} \right). \quad (76)$$

Applying Grönwall's lemma, we then have that

$$\left( 1 - e^{-\text{KL}(\mu_t|\pi)} \right) \leq e^{-2\lambda t} \left( 1 - e^{-\text{KL}(\mu_t|\pi)} \right) \leq e^{-2\lambda t}, \quad (77)$$

where in the final line we have used the elementary inequality  $g(x) = 1 - e^{-x} \leq 1$ , for  $x \geq 0$ . Finally, observe that if  $t \geq \frac{1}{2\lambda}$ , then  $(1 - e^{-\text{KL}(\mu_t|\pi)}) \leq e^{-2\lambda t} \leq e^{-1}$ . Combining this with the fact that  $x \leq 2g(x)$  whenever  $g(x) \leq e^{-1}$ , we have that

$$\text{KL}(\mu_t|\pi) \leq 2 \left( 1 - e^{-\text{KL}(\mu_t|\pi)} \right) \leq 2e^{-2\lambda t}. \quad (78)$$

for  $t \geq \frac{1}{2\lambda}$ . Alongside the bound in (73), this completes the proof of (74).  $\square$

In [20], the authors provide an interpretation of the continuous-time SVGD dynamics as a kernelised gradient flow of the  $\chi^2$  divergence. Similarly, we can interpret the continuous-time MSVGD dynamics as a kernelised, mirrored gradient flow of the  $\chi^2$  divergence; see also [79, App. H]. In particular, based on the simple observation that [20, Section 2.3]

$$\mathcal{P}_{\eta_t, k_\phi} \nabla \log \left( \frac{\eta_t}{\nu} \right) = \int k(x, \cdot) \nabla \log \left( \frac{\eta_t}{\nu}(x) \right) d\eta_t(x) \quad (79)$$

$$= \int k(x, \cdot) \nabla \left( \frac{\eta_t}{\nu}(x) \right) d\nu(x) = \mathcal{P}_{\nu, k_\phi} \nabla \left( \frac{\eta_t}{\nu} \right). \quad (80)$$

the continuous-time MSVGD dynamics in (5) can equivalently be written as

$$\frac{\partial \eta_t}{\partial t} + \nabla \cdot (w_t \eta_t) = 0, \quad w_t = -\mathcal{P}_{\nu, k_\phi} \nabla \left( \frac{\eta_t}{\nu} \right), \quad (81)$$

with  $\mu_t = (\nabla \phi^*)_{\#} \eta_t$ . On comparison of (81) and (68), it is clear that, up to a factor of two, MSVGD can be interpreted as the mirrored Wasserstein gradient flow obtained by replacing the gradient of chi-squared divergence,  $\nabla \left( \frac{\eta_t}{\nu} \right)$ , by  $\mathcal{P}_{\nu, k_\phi} \nabla \left( \frac{\eta_t}{\nu} \right)$ . In [79, App. H], the authors obtain continuous-time convergence results for this scheme, under a mirrored Stein Poincaré inequality. Here, we proceed differently.

Following [20], suppose now that we replace  $\mathcal{P}_{\nu, k_\phi}$  by the vector field  $\nabla \mathcal{P}_{\nu, k_\phi}$ . The new dynamics, which we will refer to as the mirrored Laplacian adjusted Wasserstein gradient descent (MLAWGD) dynamics, are then given by

$$\frac{\partial \eta_t}{\partial t} + \nabla \cdot (w_t \eta_t) = 0, \quad w_t = -\nabla \mathcal{P}_{\nu, k_\phi} \frac{\eta_t}{\pi} \quad (82)$$

where, naturally,  $\mu_t = (\nabla \phi^*)_{\#} \eta_t$ . The dynamics in (82) can essentially be viewed as a change of measure of the standard LAWGD dynamics in [20], designed to converge to the dual target  $\nu = (\nabla \phi)_{\#} \pi$ .

Suppose, in addition, that the kernel  $k_\phi$  were chosen such that  $\mathcal{P}_{\nu, k_\phi} = -\mathcal{L}_\nu^{-1}$ , where  $\mathcal{L}_\nu$  denotes the infinitesimal generator of the (overdamped) Langevin diffusion with stationary distribution  $\nu$ . Thus, in particular,

$$\mathcal{L}_\nu f_\phi(y) = -\langle \nabla W(y), \nabla f_\phi(y) \rangle + \Delta f_\phi(y). \quad (83)$$

where  $f_\phi : \mathbb{R}^d \rightarrow \mathbb{R}$  denotes the function  $f_\phi(\cdot) = f(\nabla\phi^*(\cdot))$ , given some base function  $f : \mathcal{X} \rightarrow \mathbb{R}$ . Letting  $y = \nabla\phi(x)$ , using (18), and then the chain rule, we can rewrite this as

$$\mathcal{L}_\nu f(x) = -\langle [\nabla^2\phi(x)]^{-1} [\nabla V(x) + \nabla \log \det \nabla^2\phi(x)], [\nabla^2\phi(x)]^{-1} \nabla f(x) \rangle \quad (84)$$

$$+ \text{Tr} [-[\nabla^2\phi(x)]^{-1} \nabla^3\phi(x) [\nabla^2\phi(x)]^{-2} \nabla f(x) + [\nabla^2\phi(x)]^{-2} \nabla^2 f(x)]. \quad (85)$$

We will denote this kernel  $k_{\mathcal{L}_\nu}$ . In practice, computing  $k_{\mathcal{L}_\nu}$  requires computing the spectral decomposition of the operator  $\mathcal{L}_\nu$ , which is challenging even in the moderate dimensions. Nonetheless, following [20, Theorem 4], for this choice of kernel we can show that the MLAWGD dynamics converge exponentially fast to the target distribution, at a rate independent of the Poincaré constant.

**Proposition 14.** *Assume that  $\pi$  satisfies a mirrored Poincaré inequality with some constant  $\kappa > 0$ . In addition, suppose that  $\mathcal{L}_\nu$  has a discrete spectrum. Then the KL divergence converges exponentially along the LAWGD dynamics, with a rate independent of  $\kappa$ :*

$$\text{KL}(\mu_t|\pi) \leq (\text{KL}(\mu_0|\pi) \wedge 2) e^{-t}. \quad (86)$$

*Proof.* Using standard rules for calculus on the Wasserstein space, and the integration by parts formula  $\mathbb{E}_\nu \langle \nabla f, \nabla g \rangle = \mathbb{E}_\nu [f \mathcal{L}_\nu g]$  [e.g., 4], we have

$$\frac{d\text{KL}(\eta_t|\nu)}{dt} = -\left\langle \nabla \left( \frac{\eta_t}{\nu} \right), \nabla P_{\nu, k_{\mathcal{L}_\nu}} \left( \frac{\eta_t}{\nu} \right) \right\rangle_{L^2(\pi)} = \int \frac{\eta_t}{\nu} \mathcal{L}_\nu P_{\nu, k_{\mathcal{L}_\nu}} \frac{\eta_t}{\nu} d\nu(x) = -\left\| \frac{\eta_t}{\nu} \right\|_{L^2(\nu)}^2.$$

By [42, Theorem 2] and (34) - (36) in the proof of Proposition 2, we have that  $\text{KL}(\eta_t|\nu) = \text{KL}(\mu_t|\pi)$  and  $\left\| \nabla \left( \frac{\eta_t}{\nu} \right) \right\|_{L^2(\nu)}^2 = \left\| [\nabla^2\phi]^{-1} \left( \frac{\mu_t}{\pi} \right) \right\|_{L^2(\pi)}^2$ , which implies that

$$\frac{d\text{KL}(\mu_t|\pi)}{dt} = -\left\| \frac{\mu_t}{\pi} \right\|_{L^2(\pi)}^2 = -\left\| \left( \frac{\mu_t}{\pi} - 1 \right) \right\|_{L^2(\pi)}^2 = -\mathcal{X}^2(\mu_t|\pi) \leq -\text{KL}(\mu_t|\pi), \quad (87)$$

where in the final line we have once again used the fact that  $\text{KL}(\cdot|\pi) \leq \mathcal{X}^2(\cdot|\pi)$ . The conclusion now follows via Grönwall's inequality.  $\square$

To obtain an implementable MLAWGD algorithm, we use the observation that the vector field in the continuity equation can be rewritten as [20]

$$-\nabla P_{\nu, k_{\mathcal{L}_\nu}} \frac{\eta_t}{\nu}(\cdot) = -\int \nabla_1 k_{\mathcal{L}_\nu}(\cdot, y) \frac{\eta_t}{\nu} d\nu(y) = -\int \nabla_1 k_{\mathcal{L}_\nu}(\cdot, y) d\eta(y). \quad (88)$$

Next, using a forward Euler discretisation in time, we arrive at

$$\eta_{t+1} = \left( \text{id} - \gamma \nabla P_{\nu, k_{\mathcal{L}_\nu}} \frac{\eta_t}{\nu} \right)_{\#} \eta_t, \quad (89)$$

where  $\mu_{t+1} = (\nabla\phi)_{\#} \eta_{t+1}$  and where, as elsewhere,  $\gamma > 0$  is a step size and  $\text{id}$  is the identity map. Finally, we obtain the MLAWGD algorithm by approximating the expectations in (89) with samples. In particular, after initialising a collection of particles  $y_0^i \stackrel{\text{i.i.d.}}{\sim} \eta_0 = (\nabla\phi)_{\#} \mu_0$  for  $i \in [N]$ , we update

$$y_{t+1}^i = y_t^i - \gamma \frac{1}{N} \sum_{j=1}^N \nabla_{y^i} k_{\mathcal{L}_\nu}(y_t^i, y_t^j) \quad (90)$$

$$x_{t+1}^i = \nabla\phi^*(y_{t+1}^i). \quad (91)$$

#### B.4 Mirrored Kernel Stein Discrepancy Descent

Finally, we consider the MWGF of the kernel Stein discrepancy (KSD) [21, 37, 60]. In so doing, we will arrive at the mirrored version of the kernel Stein discrepancy descent (KSDD) algorithm introduced in [50]. We begin by defining  $\mathcal{F}(\eta) = \frac{1}{2} \text{KSD}_\phi^2(\eta|\nu)$ , where  $\text{KSD}_\phi(\eta|\nu)$  is the *mirrored KSD*, which we define as

$$\text{KSD}_\phi(\eta|\nu) = \sqrt{\int_{\mathbb{R}^d} \int_{\mathbb{R}^d} k_{\nu, \phi}(y, y') d\eta(y) d\eta(y')}, \quad (92)$$

---

**Algorithm 4** Mirrored LAWGD
 

---

**input:** target density  $\pi$ , mirror function  $\phi$ , particles  $(x_0^i)_{i=1}^N \sim \mu$ , step size  $\gamma$ .  
**initialise:** for  $i \in [N]$ ,  $y_0^i = \nabla\phi(x_0^i)$ .  
**for**  $t = 0, \dots, T-1$  **do**  
   For  $i \in [N]$ ,  $y_{t+1}^i = y_t^i - \gamma \frac{1}{N} \sum_{j=1}^N \nabla_{y^i} k_{\mathcal{L}_\nu}(y_t^i, y_t^j)$ .  
   For  $i \in [N]$ ,  $x_{t+1}^i = \nabla\phi^*(y_{t+1}^i)$ .  
**return**  $(x_T^i)_{i=1}^N$ .

---

where  $k_{\nu, \eta}$  is the mirrored Stein kernel, defined in terms of the score  $s_\nu = \nabla \log \nu$  and the positive semi-definite kernel  $k_\phi$  as

$$\begin{aligned}
 k_{\nu, \phi}(y, y') &= \nabla_y \log \nu(y)^\top k_\phi(y, y') \nabla_{y'} \log \nu(y') + \nabla_y \log \nu(y)^\top \nabla_{y'} k_\phi(y, y') \\
 &\quad + \nabla_{y'} k_\phi^\top(y, y') \nabla \log \nu(y') + \langle \nabla_y k_\phi(y, \cdot), \nabla_{y'} k_\phi(\cdot, y') \rangle_{\mathcal{H}_{k_\phi}^d}, \tag{93}
 \end{aligned}$$

and where, as before,  $k_\phi$  is the kernel defined according to  $k_\phi(y, y') = k(\nabla\phi^*(y), \nabla\phi^*(y'))$ , for some base kernel  $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ . The MKSD is nothing more than the KSD between  $\eta = (\nabla\phi)_\# \mu$  and  $\nu = (\nabla\phi)_\# \pi$ , with respect to the kernel  $k_\phi$ .

In this case, one can show that  $\nabla_{W_2} \mathcal{F}(\eta) = \mathbb{E}_{y \sim \eta} [\nabla_2 k_{\nu, \phi}(y, \cdot)]$  [50, Proposition 2]. We can thus define the MKSD gradient flow according to

$$\frac{\partial \eta_t}{\partial t} + \nabla \cdot (w_t \eta_t) = 0, \quad w_t = -\mathbb{E}_{y \sim \eta_t} [\nabla_2 k_{\nu, \phi}(y, \cdot)] \tag{94}$$

where, as elsewhere,  $\mu_t = (\nabla\phi^*)_\# \eta_t$ . Similar to the previous sections, we will now outline some of the properties of this scheme in continuous time.

**Proposition 15.** *The dissipation of  $\frac{1}{2} \text{KSD}_\phi^2(\eta_t | \nu)$  along the mirrored Wasserstein gradient flow in (94) is given by*

$$\frac{d \frac{1}{2} \text{KSD}_\phi^2(\eta_t | \nu)}{dt} = -\mathbb{E}_{y' \sim \eta_t} \left[ \|\mathbb{E}_{y \sim \eta_t} [\nabla_{y'} k_{\nu, \phi}(y, y')]\|^2 \right] \tag{95}$$

*Proof.* The result follows straightforwardly using the chain rule. In particular, we have

$$\frac{d \frac{1}{2} \text{KSD}_\phi^2(\eta_t | \nu)}{dt} = \langle w_t, \nabla_{W_2} \frac{1}{2} \text{KSD}^2(\eta_t | \nu) \rangle_{L^2(\eta_t)} \tag{96}$$

$$= -\mathbb{E}_{y' \sim \eta_t} \left[ \|\mathbb{E}_{y \sim \eta_t} [\nabla_{y'} k_{\nu, \phi}(y, y')]\|^2 \right]. \tag{97}$$

□

**Remark 9.** *Unlike elsewhere, here the dissipation result is given in terms of the mirrored densities  $(\eta_t)_{t \geq 0}$  and the mirrored target  $\nu = (\nabla\phi)_\# \pi$ , rather than densities  $(\mu_t)_{t \geq 0}$  and the target  $\pi$ . The crucial difference in this case is that the objective functional  $\frac{1}{2} \text{KSD}^2(\eta_t | \nu)$  is not equal to  $\frac{1}{2} \text{KSD}^2(\mu_t | \pi)$ , whereas previously it was true that, e.g.,  $\text{KL}(\eta_t | \nu) = \text{KL}(\mu_t | \pi)$  or  $\chi^2(\eta_t | \nu) = \chi^2(\mu_t | \pi)$ .*

*Nonetheless, it is possible to give a dissipation result in terms of the target  $\pi$ . We begin by rewriting the LHS of (95). Recall that the squared KSD between  $\mu$  and  $\pi$  is identical to the Stein Fisher information between  $\mu$  and  $\pi$  [e.g., 60, Theorem 3.6]. This also holds for  $\eta = (\nabla\phi)_\# \mu$  and  $\nu = (\nabla\phi)_\# \pi$  in the mirrored space, and thus we have*

$$\text{KSD}_\phi^2(\eta | \nu) = \left\| \mathcal{S}_{\mu, k_\phi} \nabla \log \left( \frac{\eta}{\nu} \right) \right\|_{\mathcal{H}_{k_\phi}^d}^2 \tag{98}$$

$$= \left\| \mathcal{S}_{\mu, k} [\nabla^2 \phi]^{-1} \nabla \log \left( \frac{\mu}{\pi} \right) \right\|_{\mathcal{H}_k^d}^2, \tag{99}$$

*where in the second line we have used the result obtained in (57) - (60); see the proof of Proposition 8. Thus, in particular,  $\mathcal{F}(\eta)$  is equal to (half of) the mirrored Stein Fisher information.*

---

**Algorithm 5** Mirrored KSDD
 

---

**input:** target density  $\pi$ , kernel  $k$ , mirror function  $\phi$ , particles  $(x_0^i)_{i=1}^N \sim \mu$ , step size  $\gamma$ .  
**initialise:** for  $i \in [N]$ ,  $y_0^i = \nabla\phi(x_0^i)$ .  
**for**  $t = 0, \dots, T-1$  **do**  
   For  $i \in [N]$ ,  $y_{t+1}^i = y_t^i - \gamma \frac{1}{N^2} \sum_{j=1}^N \nabla_{y^i} k_{\nu, \phi}(y_t^j, y_t^i)$ .  
   For  $i \in [N]$ ,  $x_{t+1}^i = \nabla\phi^*(y_{t+1}^i)$ .  
**return**  $(x_T^i)_{i=1}^N$ .

---

We now turn our attention to the RHS of (95). We would like to simplify  $\nabla_{y'} k_{\nu, \phi}(y, y')$ , as defined in (93). Let  $y = \nabla\phi(x)$  and  $y' = \nabla\phi(x')$ . Using the definition of  $k_\phi$ , we have  $k_\phi(y, y') = k(x, x')$ . Using also the chain rule, we have

$$\nabla_y k_\phi(y, y') = [\nabla^2\phi(x)]^{-1} \nabla_x k(x, x') \quad , \quad \nabla_{y'} k_\phi(y, y') = [\nabla^2\phi(x')]^{-1} \nabla_{x'} k(x, x') \quad (100)$$

Similarly, it holds that

$$\langle \nabla_y k_\phi(y, \cdot), \nabla_{y'} k_\phi(\cdot, y') \rangle_{\mathcal{H}_{k_\phi}^d} = \langle [\nabla^2\phi(x)]^{-1} \nabla_x k(x, \cdot), [\nabla^2\phi(x')]^{-1} \nabla_{x'} k(\cdot, x') \rangle_{\mathcal{H}_k^d}. \quad (101)$$

Finally, due to (18), we have  $\nabla_y \log \nu(y) = [\nabla^2\phi(x)]^{-1} [\nabla_x \log \pi(x) - \nabla_x \log \det \nabla^2\phi(x)]$ . Putting everything together, we thus have that

$$\begin{aligned} k_{\nu, \phi}(y, y') &= ([\nabla^2\phi(x)]^{-1} [\nabla_x \log \pi(x) - \nabla_x \log \det \nabla^2\phi(x)])^\top k(x, x') \\ &\quad \cdot ([\nabla^2\phi(x)]^{-1} [\nabla_x \log \pi(x) - \nabla_x \log \det \nabla^2\phi(x)]) \\ &\quad + ([\nabla^2\phi(x)]^{-1} [\nabla_x \log \pi(x) - \nabla_x \log \det \nabla^2\phi(x)])^\top ([\nabla^2\phi(x)]^{-1} \nabla_x k(x, x')) \\ &\quad + ([\nabla^2\phi(x')]^{-1} \nabla_{x'} k(x, x'))^\top ([\nabla^2\phi(x)]^{-1} [\nabla_x \log \pi(x) - \nabla_x \log \det \nabla^2\phi(x)]) \\ &\quad + \langle [\nabla^2\phi(x)]^{-1} \nabla_x k(x, \cdot), [\nabla^2\phi(x')]^{-1} \nabla_{x'} k(\cdot, x') \rangle \\ &:= \tilde{k}_{\pi, \phi}(x, x'). \end{aligned} \quad (102)$$

$$:= \tilde{k}_{\pi, \phi}(x, x'). \quad (103)$$

Using (99) and (103) we thus have the following dissipation result for the mirrored Stein Fisher information in terms of the target  $\pi$ :

$$\frac{d}{dt} \left\| S_{\mu_t, k} [\nabla^2\phi]^{-1} \nabla \log \left( \frac{\mu_t}{\pi} \right) \right\|_{\mathcal{H}_k^d}^2 = -2 \mathbb{E}_{x' \sim \mu_t} \left[ \left\| \mathbb{E}_{x \sim \mu_t} \left[ \tilde{k}_{\pi, \phi}(x, x') \right] \right\|^2 \right]. \quad (104)$$

To obtain an implementable MKSDD algorithm, it is, of course, necessary to discretise in time and in space. In this case, following [50], we will first discretise in space, replacing the measure  $\eta$  by a discrete approximation  $\hat{\eta} = \frac{1}{N} \sum_{j=1}^N \delta_{y^j}$ . This yields a tractable loss function

$$F(\omega^N) := \mathcal{F}(\hat{\eta}) = \frac{1}{2N^2} \sum_{i, j=1}^N k_{\nu, \phi}(y^i, y^j), \quad (105)$$

where  $\omega^N = \{y^1, \dots, y^N\}$ , with partial gradients given by

$$\nabla_{y^i} F(\omega^N) := \mathcal{F}(\hat{\eta}) = \frac{1}{N^2} \sum_{j=1}^N \nabla_2 k_{\nu, \phi}(y^j, y^i). \quad (106)$$

The mirrored KSDD algorithm thus takes the following form. Initialise a set of particles  $y_0^i \sim \eta_0 = (\nabla\phi)_{\#} \mu_0$  for  $i \in [N]$  and then update

$$y_{t+1}^i = y_t^i - \gamma \frac{1}{N^2} \sum_{j=1}^N \nabla_2 k_{\nu, \phi}(y_t^j, y_t^i) \quad , \quad x_{t+1}^i = \nabla\phi^*(y_{t+1}^i). \quad (107)$$

## C Convergence of Mirrored Coin Wasserstein Gradient Descent

**Proposition 16.** *Suppose that  $\pi$  is strongly log-concave and that  $\|\nabla_{W_2} \mathcal{F}(\eta_t)(y_t)\| \leq 1$  for all  $t \in [T]$ . In addition, suppose that  $(\eta_t)_{t \in \mathbb{N}_0}$  and  $\nu$  satisfy [78, Assumption B.1 or B.2]. Then there exists a mirror map  $\phi$  such that*

$$\begin{aligned} \text{KL} \left( \frac{1}{T} \sum_{t=1}^T \mu_t \middle| \middle| \pi \right) &\leq \frac{L}{T} \left[ 1 + \int_{\mathcal{X}} \|\nabla \phi(x)\| \sqrt{T \ln \left( 1 + \frac{96K_1^2 T^2 \|\nabla \phi(x)\|^2}{\varepsilon^2} \right)} \pi(\mathrm{d}x) \right. \\ &\quad \left. + \int_{\mathcal{X}} \|\nabla \phi(x)\| \sqrt{T \ln \left( 1 + \frac{96T^2 \|\nabla \phi(x)\|^2}{\varepsilon^2} \right)} \mu_0(\mathrm{d}x) \right]. \end{aligned} \quad (108)$$

*Proof.* First note that, using Jensen's inequality, and [42, Theorem 2], we have that

$$\text{KL} \left( \frac{1}{T} \sum_{t=1}^T \mu_t \middle| \middle| \pi \right) \leq \frac{1}{T} \sum_{t=1}^T \text{KL}(\mu_t \middle| \pi) \quad (109)$$

$$= \frac{1}{T} \sum_{t=1}^T [\text{KL}(\eta_t \middle| \nu) - \text{KL}(\nu \middle| \nu)] \quad (110)$$

$$:= \frac{1}{T} \sum_{t=1}^T \mathcal{F}(\eta_t) - \mathcal{F}(\nu) \quad (111)$$

where in the final line we have defined  $\mathcal{F}(\eta) := \text{KL}(\eta \middle| \pi)$ . Now, under the assumption that  $\pi$  is strongly log-concave, [42, Theorem 3] guarantees the existence of a mirror map  $\phi : \mathcal{X} \rightarrow \mathbb{R} \cup \{\infty\}$  such that the mirrored target  $\nu = (\nabla \phi)_{\#} \pi$  is also strongly log-concave. It follows, in particular, that  $\mathcal{F}(\eta) := \text{KL}(\eta \middle| \pi)$  is geodesically convex [e.g., 2, Chapter 9]. Thus, arguing as in the proof of Proposition 3.3 in [78], setting  $w_0(x) = 1$ , it follows that

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathcal{F}(\eta_t) - \mathcal{F}(\nu) &\leq \frac{L}{T} \left[ 1 + \int_{\mathbb{R}^d} \|y\| \sqrt{T \ln \left( 1 + \frac{96K_1^2 T^2 \|y\|^2}{\varepsilon^2} \right)} \nu(\mathrm{d}y) \right. \\ &\quad \left. + \int_{\mathbb{R}^d} \|y\| \sqrt{T \ln \left( 1 + \frac{96T^2 \|y\|^2}{\varepsilon^2} \right)} \eta_0(\mathrm{d}y) \right]. \end{aligned} \quad (112)$$

Finally, once more using the definition of the mirror map and in particular the fact that  $\nu = (\nabla \phi)_{\#} \pi$  and  $\eta_0 = (\nabla \phi)_{\#} \eta_0$ , we can rewrite the RHS in (112) as the RHS in (108). This completes the proof.  $\square$

## D Other Mirrored Coin ParVI Algorithms

Recall, from Sec. 4.2, the general form of the mirrored coin sampling algorithm. Initialise at  $y_0 = \nabla \phi(x_0) \in \mathbb{R}^d$ , where  $x_0 \in \mathcal{X}$ . Then, update

$$y_t = y_0 + \frac{\sum_{s=1}^{t-1} c_{\eta_s}(y_s)}{t} \left( 1 + \sum_{s=1}^{t-1} \langle c_{\eta_s}(y_s), y_s - y_0 \rangle \right), \quad x_t = \nabla \phi^*(x_t), \quad (113)$$

where  $c_{\eta_t} = -\nabla_{W_2} \mathcal{F}(\eta_t)$  or some variant thereof. For different choices of  $\mathcal{F}$  and  $(c_{\eta_t})_{t \in \mathbb{N}_0}$ , one obtains learning-rate free analogues of different mirrored ParVI algorithms.

In Sec. 4.2, we provided the Coin MSVGD algorithm (Alg. 2), which corresponds to  $\mathcal{F} = \text{KL}(\cdot \middle| \nu)$  and  $(c_{\eta_t})_{t \in \mathbb{N}_0} = (-P_{\eta_t, k_{\phi}} \nabla \log(\frac{\eta_t}{\nu}))_{t \in \mathbb{N}_0}$ . This can be viewed as the mirrored coin sampling analogue of the MSVGD algorithm (Sec. 3.2.1 and App. B.2). We now provide two more mirrored coin sampling algorithms, which correspond to the coin sampling analogues of MLAWGD (App. B.3) and MKSDD (App. B.4).

## D.1 Coin MLAWGD

---

### Algorithm 6 Coin MLAWGD

---

**input:** initial density  $\mu_0$ , target density  $\pi$ , mirror map  $\phi$ , initial particles  $(x_0^i)_{i=1}^N \sim \mu_0$ .

**initialise:** for  $i \in [N]$ ,  $y_0^i = \nabla\phi(x_0^i)$ .

**for**  $t = 1, \dots, T$  **do**

For  $i \in [N]$ ,

$$y_t^i = y_0^i - \frac{\sum_{s=1}^{t-1} \frac{1}{N} \sum_{j=1}^N \nabla_{y^i} k_{\mathcal{L}_\nu}(y_s^i, y_s^j)}{t} \left( 1 - \sum_{s=1}^{t-1} \left\langle \frac{1}{N} \sum_{j=1}^N \nabla_{y^i} k_{\mathcal{L}_\nu}(y_s^i, y_s^j), y_s^i - y_0^i \right\rangle \right)$$

$$x_t^i = \nabla\phi^*(y_t^i).$$

where  $k_{\mathcal{L}_\nu}$  is the mirrored LAWGD kernel defined in (83).

**return**  $(x_T^i)_{i=1}^N$ .

---

## D.2 Coin MKSDD

---

### Algorithm 7 Coin MKSDD

---

**input:** initial density  $\mu_0$ , target density  $\pi$ , kernel  $k$ , mirror map  $\phi$ , initial particles  $(x_0^i)_{i=1}^N \sim \mu_0$ .

**initialise:** for  $i \in [N]$ ,  $y_0^i = \nabla\phi(x_0^i)$ .

**for**  $t = 1, \dots, T$  **do**

For  $i \in [N]$

$$y_t^i = y_0^i - \frac{\sum_{s=1}^{t-1} \frac{1}{N^2} \sum_{j=1}^N \nabla_{y^i} k_{\nu, \phi}(y_s^j, y_s^i)}{t} \left( 1 - \sum_{s=1}^{t-1} \left\langle \frac{1}{N^2} \sum_{j=1}^N \nabla_{y^i} k_{\nu, \phi}(y_s^j, y_s^i), y_s^i - y_0^i \right\rangle \right)$$

$$x_t^i = \nabla\phi^*(y_t^i).$$

where  $k_{\nu, \phi}$  is the mirrored Stein kernel defined in (93).

**return**  $(x_T^i)_{i=1}^N$ .

---

## E Adaptive Coin MSVGD

In Sec. 4.1 - 4.2, we described a betting strategy which is valid under the assumption that the sequence of outcomes  $(c_{\eta_t})_{t \in \mathbb{N}}$  are bounded by one, in the sense that  $\|c_{\eta_t}(y_t^i)\| \leq 1$  for all  $t \in [T]$  and for all  $i \in [N]$  [see also 78, Section 3.4]. In practice, this may not be the case. If, instead,  $\|c_{\eta_t}(y_t^i)\| \leq L$  for some constant  $L$ , then one can simply replace  $c_{\eta_t}(y)$  by its normalised version, namely  $\hat{c}_{\eta_t}(y) = c_{\eta_t}(y)/L$  in Alg. 2. If, on the other hand, the constant  $L$  is unknown in advance, then one can use a coordinate-wise empirical estimate  $L_{t,j}^i$  which is adaptively updated as the algorithm runs, following [69] and later [78, App. D].<sup>3</sup> We provide full details of this approach, as it applies in the case of Coin MSVGD, in Alg. 8. Naturally, one can also obtain an analogous version of Coin MIED; we omit the full details here in the interest of brevity.

Finally, when we use a coin sampling algorithm to perform inference for a (fairness) Bayesian neural network (see Sec. 6.3), we alter the denominator of the betting fraction in Alg. 8 such that it is at least  $100L_{t,j}^i$ . Thus, the update equation in Alg. 8 (or, analogously, the update equation in the adaptive version of Coin MIED) now reads as

$$y_{t,j}^i = y_{0,j}^i + \frac{\sum_{s=1}^{t-1} c_{s,j}^i}{\max(G_{t,j}^i + L_{t,j}^i, 100L_{t,j}^i)} \left( 1 + \frac{R_{t,j}^i}{L_{t,j}^i} \right). \quad (114)$$

This modification, recommended in [69, Sec. 6], has the effect of restricting the value of the particles in the first iterations of the algorithm. We should emphasise that these adaptive algorithms, which in practice we use in all experiments, are still free of any tuning parameters.

<sup>3</sup>Here,  $i \in [N]$  indexes the particle, and  $j \in [d]$  indexes the dimension (see Alg. 8).

---

**Algorithm 8** Adaptive Coin MSVGD

---

**input:** initial measure  $\mu_0$ , target density  $\pi$ , kernel  $k$ , mirror map  $\phi$ , initial particles  $(x_0^i)_{i=1}^N \sim \mu_0$ .  
**initialise:** for  $i \in [N]$ ,  $y_0^i = \nabla\phi(x_0^i)$ ; for  $i \in [N]$  and  $j \in [d]$ ,  $L_{0,j}^i = 0$ ,  $G_{0,j}^i = 0$ ,  $R_{0,j}^i = 0$ .

**for**  $t = 1$  **to**  $T$  **do**

For  $i \in [N]$ ,

$$c_{t-1}^i = -P_{\hat{\eta}_t, k_\phi} \nabla \log \left( \frac{\hat{\eta}_{t-1}}{\nu} \right) (y_{t-1}^i), \quad \hat{\eta}_{t-1} = \frac{1}{N} \sum_{k=1}^N \delta_{y_{t-1}^k} \quad (\text{compute SVGD gradient})$$

For  $i \in [N]$ ,  $j \in [d]$ ,

$$L_{t,j}^i = \max(L_{t-1,j}^i, |c_{t-1,j}^i|) \quad (\text{update max. observed scale})$$

$$G_{t,j}^i = G_{t-1,j}^i + |c_{t-1,j}^i| \quad (\text{update sum of abs. value of gradients})$$

$$R_{t,j}^i = \max(R_{t-1,j}^i + \langle c_{t-1,j}^i, y_{t-1,j}^i - y_{0,j}^i \rangle, 0) \quad (\text{update total reward})$$

$$y_{t,j}^i = y_{0,j}^i + \frac{\sum_{s=1}^{t-1} c_{s,j}^i}{G_{t,j}^i + L_{t,j}^i} \left( 1 + \frac{R_{t,j}^i}{L_{t,j}^i} \right). \quad (\text{update the particles})$$

For  $i \in [N]$ ,  $x_t^i = \nabla\phi^*(y_t^i)$ .

**output:**  $(x_T^i)_{i=1}^N$ .

---

It is worth noting that, in terms of computational and memory cost, the adaptive Coin MSVGD algorithm is no worse than MSVGD when the latter is paired, as is common, with a method such as Adagrad [30], RMSProp [89], or Adam [47]. The computational cost per iteration is  $O(N^2)$ , due to the kernelised gradient. Meanwhile, in terms of memory requirements, it is necessary to keep track of the maximum observed absolute value of the gradients (component-wise), the sum of the absolute value of the gradients (component-wise), and the total reward (component-wise). This is similar to, e.g., Adam, which keeps track of exponential moving averages of the gradient and the (component-wise) squared gradient [47].

## F Additional Experimental Details

In this section, we provide additional details relevant to the experiments carried out in Sec. 6. We implement all methods using Python 3, PyTorch, and TensorFlow. We use the existing implementations of MSVGD and SVMD [79] and MIED [57] to run these methods. We perform all experiments using a MacBook Pro 16" (2021) laptop with Apple M1 Pro chip and 16GB of RAM.

For the experiments in Sec. 6.1 and 6.2, we use the inverse multiquadric (IMQ) kernel  $k(x, x') = (1 + \{\|x - x'\|_2^2/h^2\})^{-\frac{1}{2}}$  due to its favourable convergence control properties (e.g., [37]). Unless otherwise stated, the bandwidth  $h$  is determined using the median heuristic [35, 61]. For the experiment in Sec. 6.3, following [57], we use the  $s$ -Riesz family of mollifiers  $\{\phi_\epsilon^s\}$  with  $s = d + 10^{-4}$  and  $\epsilon = 10^{-8}$ .

### F.1 Simplex Targets

For both experiments in Sec. 6.1, we initialise all methods with i.i.d samples from a Dirichlet(5) distribution and use the entropic mirror map

$$\phi(x) = \sum_{k=1}^d x_k \log x_k + \left( 1 - \sum_{k=1}^d x_k \right) \log \left( 1 - \sum_{k=1}^d x_k \right). \quad (115)$$

We run all methods for  $T = 500$  iterations using  $N = 50$  particles. In the interest of a fair comparison, we run all of the learning-rate dependent methods (MSVGD, SVMD, projected SVGD) using RMSProp [89], an extension of the Adagrad algorithm [30], to adapt the learning rates.

**Sparse Dirichlet Posterior.** We first consider the sparse Dirichlet posterior introduced in [70], which takes the form

$$\pi(x) = \frac{1}{Z} \prod_{k=1}^{d+1} x_k^{n_k + \alpha_k - 1}, \quad x \in \text{int}(\Delta), \quad (116)$$

for parameters  $\alpha_1, \dots, \alpha_{d+1} > 0$  and observations  $n_1, \dots, n_{d+1}$ , where  $n_k$  is the number of observations of category  $k$ . We consider the sparse regime, in which most of the  $n_k$  are zero and the remaining  $n_k$  are large, e.g., of order  $O(d)$ . In particular, as in [79], we set  $\alpha_k = 0.1$  for  $k \in \{1, \dots, 20\}$  and use count data  $n_1 = 90$ ,  $n_2 = n_3 = 5$ , and  $n_j = 0$  for  $j > 3$ . This target is used to replicate the multimodal sparse conditionals that arise in Latent Dirichlet Allocation (LDA) [9]. While this distribution is not log-concave, the dual-target is strictly log-concave under the entropic mirror map.

**Quadratic Target** We consider the quadratic simplex target introduced in [1]. In this case, the target takes the form

$$\pi(x) = \frac{1}{Z} \exp\left(-\frac{1}{2\sigma^2} x^\top A x\right), \quad (117)$$

where  $A \in \mathbb{R}^{d \times d}$  is a symmetric positive semidefinite matrix with all entries bounded in magnitude by 1. Following [79], we set  $\sigma = 0.01$  and generate  $A$  by normalising products of random matrices with i.i.d. elements drawn from  $\text{Unif}[-1, 1]$ .

## F.2 Confidence Intervals for Post Selection Inference

Given data  $(X, y) \in \mathbb{R}^{n \times p} \times \mathbb{R}^n$ , the randomised Lasso is defined as the solution of (e.g., [86, Section 3.1] or [87, Section 1.1])

$$\hat{\beta} = \arg \min_{\beta \in \mathbb{R}^p} \left[ \frac{1}{2} \|y - X\beta\|_2^2 + \lambda \|\beta\|_1 - \omega^\top \beta + \frac{\varepsilon}{2} \|\beta\|_2^2 \right], \quad (118)$$

where  $\lambda$  is the  $\ell_1$  penalty parameter,  $\varepsilon \in \mathbb{R}_+$  is a ridge term, and  $\omega \in \mathbb{R}^p$  is an auxiliary random vector drawn from some distribution  $\mathbb{G}$  with density  $g$ . We choose  $\mathbb{G}$  to be a zero-mean independent Gaussian distribution. The scale of this distribution, and the parameter  $\varepsilon$ , are both determined automatically by the `randomizedLasso` function in the `selectiveInference` package. The distribution of interest has density proportional to (e.g., [87, Section 4.2])

$$\hat{g}(\hat{\beta}_E, \hat{z}_{-E}) \propto g\left(\varepsilon(\hat{\beta}_E, 0)^\top - X^\top(y - X_E \hat{\beta}_E) + \lambda(\hat{z}_E, \hat{z}_{-E})^\top\right), \quad (119)$$

with support given by  $\mathcal{D} = \{(\hat{\beta}_E, \hat{z}_{-E}) : \text{sign}(\hat{\beta}_E) = \hat{z}_E, \|\hat{z}_{-E}\|_\infty \leq 1\}$ . In all experiments, following [86], we integrate out  $\hat{z}_{-E}$  analytically and reparameterise  $\hat{\beta}_E$  as  $\hat{z}_E \odot |\hat{\beta}_E|$ . This yields a log-concave density for  $|\hat{\beta}_E|$ , with domain given by the non-negative orthant in  $\mathbb{R}^q$ . We then use the mirror mapping

$$\phi(x) = \sum_{k=1}^d (x_k \log x_k - x_k). \quad (120)$$

**Synthetic Example.** We consider the simulation described in Section 5.3 in [76], with  $n = 100$  and  $p = 40$ . In particular, the design matrix  $X$  is generated from an equi-correlated model. Namely, for  $i = 1, \dots, n$ , the rows  $x_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Sigma)$ , where  $\Sigma_{ij} = \delta_{i,j} + \rho(1 - \delta_{i,j})$ , with  $\rho = 0.3$ . We also normalise the columns of  $X$  to have norm 1. Meanwhile, the response is distributed as  $y \sim \mathcal{N}(0, I_n)$ , independent of  $X$ . That is the true parameter  $\beta = 0$ . We set  $\varepsilon = \frac{\hat{\sigma}^2}{\sqrt{n}}$ , where  $\hat{\sigma}$  denotes the empirical standard deviation of the response. We set  $\lambda$  using the default value returned by `theoretical.lambda` in the `selectiveInference` package, multiplied by a factor of  $0.7n$ . This adjustment reduces the  $\ell_1$  regularisation, ensuring that a reasonable subset of the  $p = 40$  features are selected. Following [79], we use RMSProp for MSVDG and SVMMD, with an initial learning rate of 0.01. We use  $N = 50$  particles and run each algorithm for  $T = 1000$  iterations. Finally, in Fig. 2, the error bars for the actual coverage levels correspond to 95% Wilson intervals [98].

**HIV-1 Drug Resistance.** We use the vitro measurement of log-fold change under 3TC as the response and include mutations that appear more than 10 times in the dataset as predictors. The resulting dataset consists of  $n = 633$  cases and  $p = 91$  predictors. In the selection stage, we apply the randomised Lasso, setting the regularisation parameter to the theoretical value in `selectiveInference`,

multiplied by a factor of  $n$ . We implement SVMD and MSVGD using RMSProp [89], with an initial learning rate of  $\gamma = 0.01$  following [79]. We implement each method for  $T = 2000$  iterations, using  $N = 50$  particles.

### F.3 Fairness Bayesian Neural Network

We use a train-test split of 80% / 20%, as in [57, 63]. We implement the methods in [57, 63] using their source code, using the default hyperparameters. We run all algorithms for  $T = 2000$  iterations, and using  $N = 50$  particles. Similar to [57], we found that one of the algorithms (Control + SVGD) in [62] got stuck after initialisation, with accuracy close to 0.75 and did not improve during training. We thus omit the results for this method in Fig. 4.

## G Additional Numerical Results

### G.1 Simplex Targets

In Fig. 5 and Fig. 6, we provide additional results for the sparse Dirichlet posterior [70] and the quadratic simplex target [1] considered in Sec. 6.1. In particular, in Fig. 5a and Fig. 6a, we plot the energy distance to a set of ground truth samples either obtained i.i.d. (sparse Dirichlet target) or using the NUTS [41] (quadratic target) versus the number of iterations for Coin MSVGD, MSVGD, SVMD, and projected versions of Coin SVGD and SVGD. Meanwhile, in Fig. 5b and Fig. 6b, we compare the performance of Coin MSVGD with the performance of MSVGD, for multiple choices of the learning rate.

As observed in Sec. 6.1, for the sparse Dirichlet posterior, Coin MSVGD converges more rapidly than MSVGD and SVMD, even for well-chosen values of the learning rate (Fig. 5). Meanwhile, for the quadratic target, Coin MSVGD generally converges more rapidly than MSVGD but not as fast as SVMD [79], which leverages the log-concavity of the target in the primal space (Fig. 6). In both cases, for poorly chosen values of the learning rate, MSVGD (e.g.,  $\gamma = 0.5$  in Fig. 5b and  $\gamma = 0.0001, \gamma = 0.001, \text{ or } \gamma = 0.5$  in Fig. 6b) entirely fails to converge to the true target.

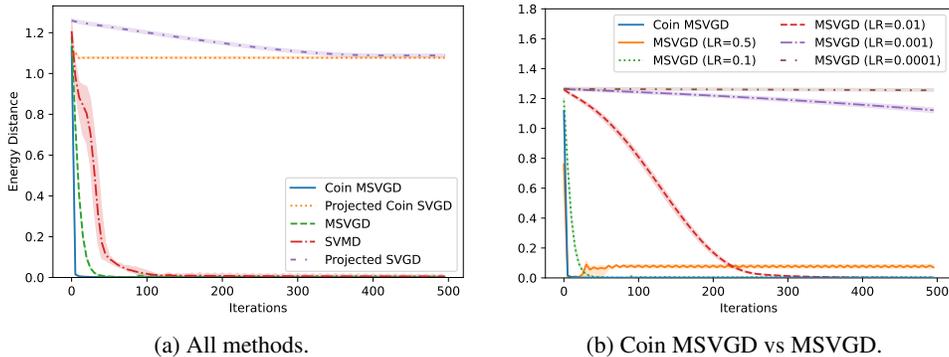


Figure 5: **Additional results for the sparse dirichlet posterior in [70].** (a) Energy distance vs. iterations for Coin MSVGD, MSVGD, SVMD, projected Coin SVGD, and projected SVGD; and (b) energy distance vs iterations for Coin MSVGD and MSVGD, using five different learning rates.

### G.2 Sampling from a Uniform Distribution

In this section, we consider the task of sampling from the uniform distribution on the unit square, namely  $\text{Unif}[-1, 1]^2$ . For MSVGD and Coin MSVGD, we once again use the entropic mirror map. We also now use the RBF kernel, with a fixed bandwidth of  $h = 0.01$ .<sup>4</sup> For MIED and Coin MIED, following [57], we reparameterise the particles using  $\tanh$  to eliminate the constraint. We show results for several choices of mollifiers, in each casing choosing the smallest  $\epsilon$  which guarantees the

<sup>4</sup>In this example, the use of an adaptive kernel (e.g., with the bandwidth chosen according to the median rule), can cause the samples from MSVGD to collapse [57].

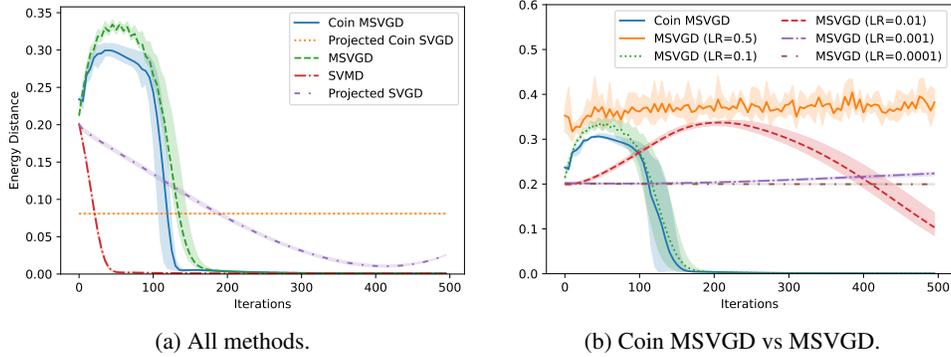


Figure 6: **Additional results for the quadratic simplex target in [1].** (a) Energy distance vs. iterations for Coin MSVGD, MSVGD, SVM, projected Coin SVGD, and projected SVGD; and (b) energy distance vs iterations for Coin MSVGD and MSVGD, using five different learning rates.

optimisation remains stable. We initialise all methods with particles drawn i.i.d. from  $\mathcal{U}[-0.5, 0.5]^2$ . For both methods, in the interest of a fair comparison, we adapt the learning rates using RMSProp [89]. We use  $N = 100$  particles and run each experiment for  $T = 250$  iterations. We then assess the quality of the posterior approximations by computing both the energy distance to a set of 1000 samples drawn i.i.d. from the true target distribution.

The results for this experiment are shown in Fig. 7. Similar to our previous synthetic examples, Coin MSVGD is competitive with the optimal performance of MSVGD. Meanwhile, for sub-optimal choices of the learning rate, the coin sampling algorithm is clearly preferable. The same is also true for Coin MIED which, depending on the choice of mollifier, obtains comparable or superior performance to MIED with the optimal choice of learning rate.

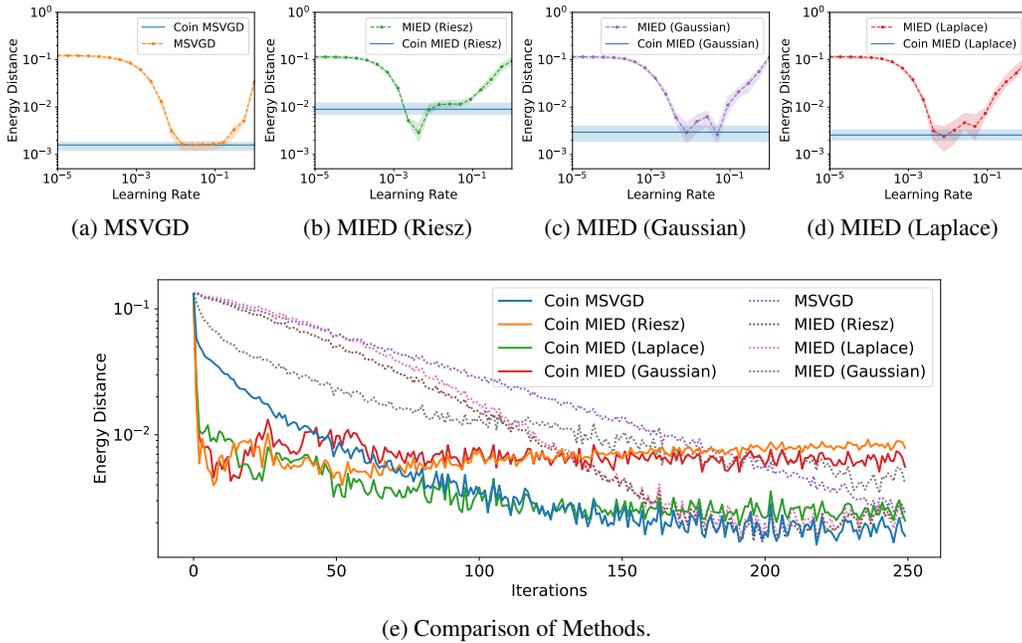


Figure 7: **Results for the uniform target:**  $\text{Unif}[-1, 1]^2$ . Energy distance vs learning rate after 250 iterations for (a) Coin MSVGD and MSVGD; and (b) - (d) Coin MIED and MIED, for three different choices of mollifier.

### G.3 Post-Selection Inference

In this section, we provide additional results for the post-selection inference examples considered in Sec. 6.2. In Fig. 8, we plot the energy distance between (Coin) MSVGD and projected (Coin) SVGD to a set of 1000 samples obtained using NUTS [41] for a two-dimensional selective distribution obtained via the randomized Lasso (i.e., when only two features are selected by the Lasso). We run each method 5 times, using different random seeds for each run.

Similar to our previous examples, the performance of MSVGD is very sensitive to the choice of learning rate. In particular, if the learning rate is too small, the particles converge very slowly. On the other hand, if the learning rate is too big, the particles may fail to converge entirely (e.g.,  $\gamma = 0.1$  in Fig 8a). Meanwhile, Coin MSVGD converges rapidly to the true target, with no need to tune a learning rate. Once again, the projected versions of SVGD and Coin SVGD both fail to converge to the true target (Fig. 8b), highlighting the need for the mirrored approach.

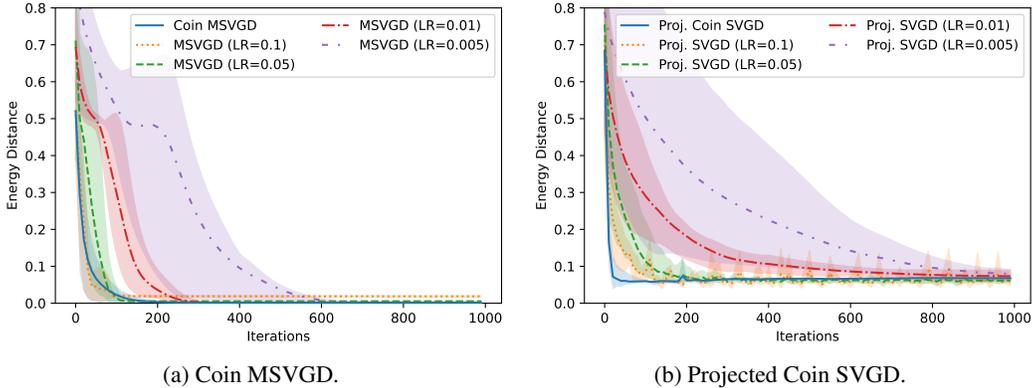


Figure 8: **Additional results for post-selection inference.** Posterior approximation quality for (a) mirror Coin SVGD and MSVGD and (b) projected Coin SVGD and projected SVGD, for the two-dimensional selective distribution described in Sec. 6.2.

In Fig. 9, we once again plot the coverage of the CIs obtained using Coin MSVGD, MSVGD, SVMD, and the norejection algorithm in `selectiveInference` [88], as we vary either the nominal coverage or the total number of samples. Now, however, we consider two sub-optimal choices of the learning rate for MSVGD and SVMD. In particular, we now use an initial learning rate of  $\gamma = 0.001$  (Fig. 9a - 9b) or  $\gamma = 0.1$  (Fig. 9c - 9d) for RMSProp [89], rather than  $\gamma = 0.01$  as in Fig. 2 (see Sec. 6.2). In this case, we see that the performance of both MSVGD and SVMD deteriorates. In particular, they both now obtain significantly lower coverage than Coin MSVGD. As noted in the main text, this increased coverage is of particular importance for smaller sample sizes (see Fig. 9b and Fig. 9c), and for greater nominal coverage levels (see Fig. 9a and 9c), where the standard CIs generally undercover.

Finally, in Fig. 10, we perform a similar experiment for the real-data example involving the HIV-1 drug resistance dataset. In particular, we compute 90% CIs for the subset of mutations selected by the randomized Lasso, using  $N = 5000$  samples and the same five methods as before: ‘unadjusted’ (the unadjusted CIs), ‘standard’ (the method in `selectiveInference`), MSVGD, SVMD, and Coin MSVGD. Now, however, we use sub-optimal choices of the initial learning rate in RMSProp, namely,  $\gamma = 0.001$  (Fig. 10a) and  $\gamma = 1.0$  (Fig. 10b), compared to  $\gamma = 0.01$  in the main paper.

Similar to the synthetic example, the impact of using a sub-optimal choice of learning rate is clear. In particular, CIs obtained using MSVGD and SVMD now differ substantially from those obtained by Coin MSVGD and the standard approach. This means, in particular, that several mutations are now incorrectly flagged as significant by MSVGD and SVMD (e.g., mutation P41L, P62V, P70R, or P215Y in Fig. 10b), when in reality there is insufficient evidence after conditioning on the selection event.

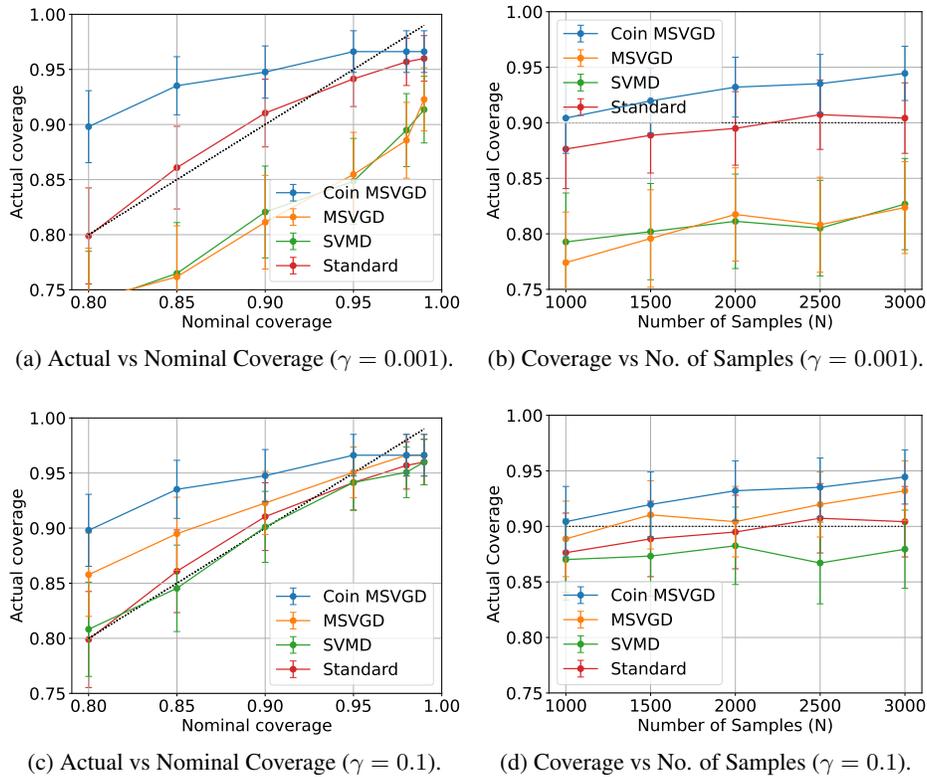


Figure 9: **Additional coverage results for the post-selection inference.** Coverage of post-selection confidence intervals for Coin MSVGD, MSVGD, and SVMD, using sub-optimal learning rates.

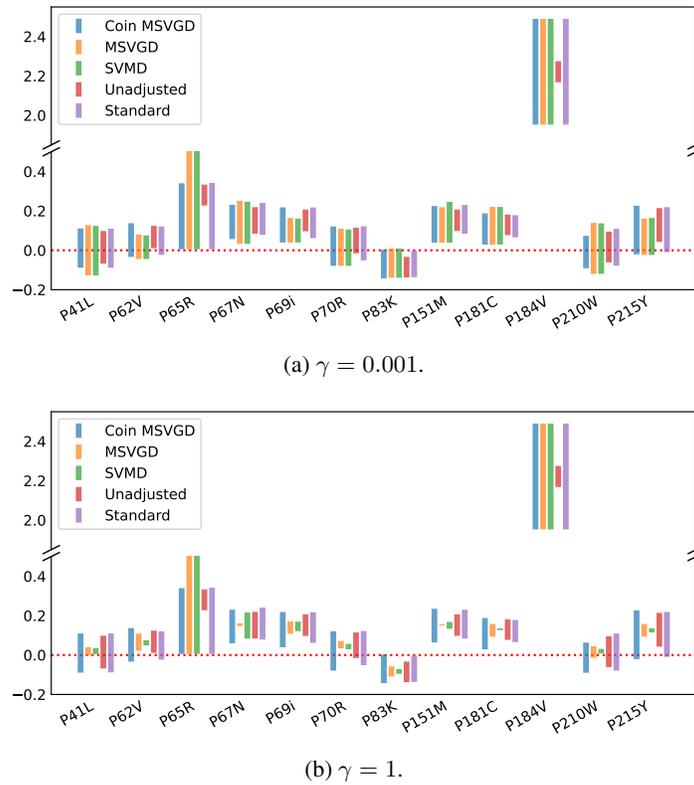


Figure 10: **Additional real data results for post-selection inference.** Unadjusted and post-selection confidence intervals for the mutations selected by the randomised Lasso as candidates for HIV-1 drug resistance, for sub-optimal choices of the learning rate.