Anna Korba¹ Pierre-Cyril Aubin-Frankowski² Szymon Majewski³ Pierre Ablin⁴

Abstract

Among dissimilarities between probability distributions, the Kernel Stein Discrepancy (KSD) has received much interest recently. We investigate the properties of its Wasserstein gradient flow to approximate a target probability distribution π on \mathbb{R}^d , known up to a normalization constant. This leads to a straightforwardly implementable, deterministic score-based method to sample from π , named KSD Descent, which uses a set of particles to approximate π . Remarkably, owing to a tractable loss function, KSD Descent can leverage robust parameter-free optimization schemes such as L-BFGS; this contrasts with other popular particle-based schemes such as the Stein Variational Gradient Descent algorithm. We study the convergence properties of KSD Descent and demonstrate its practical relevance. However, we also highlight failure cases by showing that the algorithm can get stuck in spurious local minima.

1. Introduction

An important problem in machine learning and computational statistics is to sample from an intractable target distribution π . In Bayesian inference for instance, π corresponds to the posterior probability of the parameters, which is required to compute the posterior predictive distribution. It is known only up to an intractable normalization constant. In Generative Adversarial Networks (GANs, Goodfellow et al., 2014), the goal is to generate data which distribution is similar to the training set defined by samples of π . In the first setting, one has access to the score of π (the gradient of its log density), while in the second, one has access to samples of π . Assessing how different the target π and a given approximation μ are can be performed through a dissimilarity function $D(\mu|\pi)$. As summarized by Simon-Gabriel (2018), classical dissimilarities include f-divergences such as the KL (Kullback-Leibler) or the χ^2 (Chi-squared), the Wasserstein distances in Optimal Transport (OT), and Integral Probability Metrics (IPMs), such as the Maximum Mean Discrepancy (MMD, Gretton et al., 2012). Dissimilarity functions can hence be used to characterize π since, under mild assumptions, they only vanish at $\mu = \pi$. Setting $\mathcal{F}(\mu) = D(\mu | \pi)$, assuming also in our case that $\pi \in \mathcal{P}_2(\mathbb{R}^d)$, the set of probability measures μ with finite second moment $(\int ||x||^2 d\mu(x) < \infty)$, the sampling task can then be recast as an optimization problem over $\mathcal{P}_2(\mathbb{R}^d)$

$$\min_{\mu \in \mathcal{P}_2(\mathbb{R}^d)} \mathcal{F}(\mu). \tag{1}$$

Starting from an initial distribution μ_0 , one can then apply a descent scheme to (1) to converge to π . In particular, one can consider the Wasserstein gradient flow of \mathcal{F} over $\mathcal{P}_2(\mathbb{R}^d)$. This operation can be interpreted as a vector field continuously displacing the particles constituting μ .

Among dissimilarities, the Kernel Stein Discrepancy (KSD, introduced independently by Liu et al., 2016; Chwialkowski et al., 2016; Gorham & Mackey, 2017) writes as follows

$$\mathrm{KSD}(\mu|\pi) = \sqrt{\iint k_{\pi}(x, y) \mathrm{d}\mu(x) \mathrm{d}\mu(y)}, \qquad (2)$$

where k_{π} is the Stein kernel, defined through the score of π , $s(x) = \nabla \log \pi(x)$, and through a positive semi-definite kernel k (see Section 2.1 for the meaning of $\nabla \cdot_1$ or of ∇_2)

$$k_{\pi}(x,y) = s(x)^{T} s(y) k(x,y) + s(x)^{T} \nabla_{2} k(x,y) + \nabla_{1} k(x,y)^{T} s(y) + \mathbf{\nabla} \cdot_{1} \nabla_{2} k(x,y).$$
(3)

The great advantage of the KSD is that it can be readily computed when one has access to the score of π and uses a discrete measure $\hat{\mu}$, since (2) writes as a finite double sum of k_{π} in this case. Furthermore the definition of the KSD was inspired by Stein's method (see Anastasiou et al., 2021, for a review) and no sampling over π is required in (2). This motivated the use of the KSD in a growing number of problems. The KSD has been widely used in nonparametric statistical tests for goodness-of-fit (e.g. Xu & Matsuda, 2020; Kanagawa et al., 2020). It was also used for sampling tasks: to select a suitable set of static points to approximate π , adding a new one at each iteration (Chen et al., 2018; 2019); to compress (Riabiz et al., 2020) or reweight (Hodgkinson et al.,

¹CREST, ENSAE, Institut Polytechnique de Paris ²CAS, MINES ParisTech, Paris, France ³CMAP, Ecole Polytechnique, Institut Polytechnique de Paris ⁴CNRS and DMA, Ecole Normale Supérieure, Paris, France. Correspondence to: Anna Korba <anna.korba@ensae.fr>.

Proceedings of the 38th International Conference on Machine Learning, PMLR 139, 2021. Copyright 2021 by the author(s).

2020) Markov Chain Monte Carlo (MCMC) outputs; and to learn a static transport map from μ_0 to π (Fisher et al., 2021). In this paper, we consider $\mathcal{F}(\mu) = \frac{1}{2} \text{KSD}^2(\mu | \pi)$ and its Wasserstein gradient $\nabla_{W_2} \mathcal{F}$ to define a flow over particles to approximate π .

Related works. Minimizing a dissimilarity D is a popular approach to fit an unnormalized density model in the machine learning and computational statistics literature. For instance, Hyvärinen (2005) proposed to minimize the Fisher divergence. Alternatively, D is often taken as the KL divergence. Indeed, since the seminal paper by Jordan et al. (1998), the Wasserstein gradient flow of the KL has been extensively studied and related to the Langevin Monte Carlo (LMC) algorithm (e.g. Wibisono, 2018; Durmus et al., 2019). However, an unbiased time-discretization of the KL flow is hard to implement (Salim et al., 2020). To tackle this point, a recent successful kernel-based approximation of the KL flow was introduced as the Stein Variational Gradient Descent (SVGD, Liu & Wang, 2016). Several variants were considered (see Chewi et al., 2020, and references therein). Another line of work considers D as the MMD (Mrouch et al., 2019; Arbel et al., 2019) with either regularized or exact Wasserstein gradient flow of the MMD, especially for GAN training. However, these two approaches require samples of π to evaluate the gradient of the MMD. As the KSD is a specific case of the MMD with the Stein kernel (Chen et al., 2018), our approach is similar to Arbel et al. (2019) but better suited for a score-based sampling task, owing to the properties of the Stein kernel.

Contributions. In this paper, in contrast with the aforementioned approaches, we choose the dissimilarity D in (1) to be the KSD. As in SVGD, our approach, KSD Descent, optimizes the positions of a finite set of particles to approximate π , but through a descent scheme of the KSD (in contrast to the KL for SVGD) in the space of probability measures. KSD Descent comes with several advantages. First, it benefits from a closed-form cost function which can be optimized with a fast and hyperparameter-free algorithm such as L-BFGS (Liu & Nocedal, 1989). Second, our analysis comes with several theoretical guarantees (namely, existence of the flow and a descent lemma in discrete time) under a Lipschitz assumption on the gradient of k_{π} . We also provide negative results highlighting some weaknesses of the convergence of the KSD gradient flow, such as the absence of exponential decay near equilibrium. Moreover, stationary points of the KSD flow may differ from the target π and even be local minima of the flow, which implies that some particles are stuck far from high-probability regions of π . Sometimes a simple annealing strategy mitigates this convergence issue. On practical machine learning problems, the performance of KSD Descent highly depends on the local minimas of $\log \pi$. KSD Descent achieves comparable performance to SVGD on convex (i.e., π log-concave) toy

examples and Bayesian inference tasks, while it is outperformed on non-convex tasks with several saddle-points like independent component analysis.

This paper is organized as follows. Section 2 introduces the necessary background on optimal transport and on the Kernel Stein Discrepancy. Section 3 presents our approach and discusses its connections with related works. Section 4 is devoted to the theoretical analysis of KSD Descent. Our numerical results are to be found in Section 5.

2. Background

This section introduces the high-level idea of the gradient flow approach to sampling. It also summarizes the known properties of the KSD.

2.1. Notations

The space of l continuously differentiable functions on \mathbb{R}^d is $C^l(\mathbb{R}^d)$. The space of smooth functions with compact support is $C_c^{\infty}(\mathbb{R}^d)$. If $\psi : \mathbb{R}^d \to \mathbb{R}^p$ is differentiable, we denote by $J\psi : \mathbb{R}^d \to \mathbb{R}^{p \times d}$ its Jacobian. If p = 1, we denote by $\nabla \psi$ the gradient of ψ . Moreover, if $\nabla \psi$ is differentiable, the Jacobian of $\nabla \psi$ is the Hessian of ψ denoted $H\psi$. If p = d, $\nabla \cdot \psi$ denotes the divergence of ψ , i.e. the trace of the Jacobian. We also denote by $\Delta \psi$ the Laplacian of ψ , where $\Delta \psi = \nabla \cdot \nabla \psi$. For a differentiable kernel $k : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}, \nabla_1 k$ (resp. $\nabla_2 k$) is the gradient of the kernel w.r.t. the first (resp. second) variable, while $H_1 k$ denotes its Hessian w.r.t. the first variable.

Consider the set $\mathcal{P}_2(\mathbb{R}^d)$ of probability measures μ on \mathbb{R}^d with finite second moment and $\mathcal{P}_{c}(\mathbb{R}^{d})$ the set of probability measures with compact support. For $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, we denote by $d\mu/d\pi$ its Radon-Nikodym density if μ is absolutely continuous w.r.t. π . For any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, $L^2(\mu)$ is the space of functions $f : \mathbb{R}^d \to \mathbb{R}$ such that $\int \|f\|^2 d\mu < \infty$. We denote by $\|\cdot\|_{L^2(\mu)}$ and $\langle\cdot,\cdot\rangle_{L^2(\mu)}$ respectively the norm and the inner product of the Hilbert space $L^2(\mu)$. Given a measurable map $T : \mathbb{R}^d \to \mathbb{R}^d$ and $\mu \in \mathcal{P}_2(\mathbb{R}^d), T_{\#}\mu$ is the pushforward measure of μ by T. We consider, for $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$, the 2-Wasserstein distance $W_2(\mu,\nu)$, and we refer to the metric space $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ as the Wasserstein space. In a Riemannian interpretation of $(\mathcal{P}_2(\mathbb{R}^d), W_2)$, the tangent space of $\mathcal{P}_2(\mathbb{R}^d)$ at μ is denoted $\mathcal{T}_{\mu}\mathcal{P}_{2}(\mathbb{R}^{d})$ and is a subset of $L^{2}(\mu)$ (Otto, 2001). We refer to Appendix A.2 for more details on the Wasserstein distance and related flows.

2.2. Lyapunov analysis and gradient flows

To sample from a target distribution π , a now classical approach consists in identifying a continuous process which moves particles from an initial probability distribution μ_0 toward samples of π . This can be expressed as searching for

vector fields $v_t : \mathbb{R}^d \to \mathbb{R}^d$ transporting the distribution μ_t through the continuity equation (see Appendix A.1)

$$\frac{\partial \mu_t}{\partial t} + \operatorname{div}(\mu_t v_t) = 0 \tag{4}$$

where v_t should ensure the convergence of μ_t to π , for some topology over measures, in finite or infinite time. Due to v_t , eq. (4) is nonlinear over μ_t . Cauchy-Lipschitz-style assumptions for existence and uniqueness of the solution of (4) are provided in Appendix A.3. The continuity equation ensures that the mass is conserved and that it is not teleported as for a mixture $\mu_t = (1 - t)\mu_0 + t\pi$. In order to adjust the position of particles only depending on the present distribution μ_t and to have an automated choice of v_t at any given time, it is favorable to have v_t as a function of μ_t , written as v_{μ_t} .

A principled way to select such a $v_{\mu t}$ is to define it based on a Lyapunov functional $\mathcal{F}(\mu)$ over measures, decreasing along the Wasserstein gradient flow (see Appendix A.2)

$$\dot{\mathcal{F}}(\mu_t) := \frac{\mathrm{d}\mathcal{F}(\mu_t)}{\mathrm{d}t} = \langle \nabla_{W_2}\mathcal{F}(\mu_t), v_{\mu_t} \rangle_{L^2(\mu_t)} \le 0.$$
(5)

Any dissimilarity $\mathcal{F}(\cdot) = D(\cdot|\pi)$ is a valid Lyapunov candidate since it is non-negative and vanishes at π . Hence, (4) can be seen as a continuous descent scheme of (1) or, conversely, (1)-(5) can be interpreted as a way to choose v_{μ_t} in (4) to steer μ_0 to π . In short, any Lyapunov-based approach rests upon three quantities $(\mathcal{F}(\mu_t), v_{\mu_t}, \dot{\mathcal{F}}(\mu_t))$, related by (5). A natural choice of v_{μ_t} satisfying (5) and realizing the steepest descent is the Wasserstein gradient itself, $v_{\mu_t} = -\nabla_{W_2} \mathcal{F}(\mu_t)$. Depending on the choice of \mathcal{F} , this v_{μ_t} may be hard to implement, or require specific analysis of the resulting dissipation function $\dot{\mathcal{F}}(\mu_t)$. Otherwise, to ensure that $\dot{\mathcal{F}}(\mu_t)$ only vanishes at π , one can choose v_{μ_t} so that both $\mathcal{F}(\mu_t)$ and $\dot{\mathcal{F}}(\mu_t)$ are known dissimilarities. As a matter of fact, if there exists a dissimilarity \tilde{D} separating measures such that $-\dot{\mathcal{F}}(\mu) \geq \tilde{D}(\mu|\pi)$, then π is asymptotically stable for the flow and, if $D(\mu|\pi) > \mathcal{F}(\mu)$, then π is exponentially stable (by Gronwall's lemma). This relates the Lyapunov analysis to functional inequalities (Villani, 2003), expressing domination w.r.t. \mathcal{F} of the Wasserstein gradient of \mathcal{F} under specific assumptions on π , e.g. log-Sobolev for the KL, or Poincaré for the χ^2 (Chewi et al., 2020).

2.3. Kernel Stein Discrepancy

Consider a positive semi-definite kernel $k : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ and its corresponding RKHS \mathcal{H}_k of real-valued functions on \mathbb{R}^d . The space \mathcal{H}_k is a Hilbert space with inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}_k}$ and norm $\|\cdot\|_{\mathcal{H}_k}$. Moreover, k satisfies the reproducing property: $\forall f \in \mathcal{H}_k, f(x) = \langle f, k(x, \cdot) \rangle_{\mathcal{H}_k}$; which for smooth kernels also holds for derivatives, e.g. $\partial_i f(x) =$ $\langle f, (\nabla_1 k(x, \cdot))_i \rangle_{\mathcal{H}_k}$ (see Saitoh & Sawano, 2016). Let $\mu \in \mathcal{P}_2(\mathbb{R}^d)$. If $\int k(x, x) d\mu(x) < \infty$, then the integral operator associated to the kernel k and measure μ , denoted by $S_{\mu,k}: L^2(\mu) \to \mathcal{H}_k$ and defined as

$$S_{\mu,k}f = \int k(x,\cdot)f(x)\mathrm{d}\mu(x),\tag{6}$$

is a Hilbert-Schmidt operator and $\mathcal{H}_k \subset L^2(\mu)$. In this case, the identity embedding $\iota : \mathcal{H}_k \to L^2(\mu)$ is a bounded operator and it is the adjoint of $S_{\mu,k}$ (i.e., $\iota^* = S_{\mu,k}$ (Steinwart & Christmann, 2008, Theorems 4.26 and 4.27). Hence, for any $(f,g) \in \mathcal{H}_k \times L^2(\mu), \langle \iota f, g \rangle_{L^2(\mu)} = \langle f, S_{\mu,k}g \rangle_{\mathcal{H}_k}$. We denote by \mathcal{H}_k^d the Cartesian product RKHS consisting of elements $f = (f_1, \ldots, f_d)$ with $f_i \in \mathcal{H}_k$, and with inner product $\langle f, g \rangle_{\mathcal{H}_k^d} = \sum_{i=1}^d \langle f_i, g_i \rangle_{\mathcal{H}_k}$. For vector-inputs, we extend $S_{\mu,k}$, applying it component-wise.

The Stein kernel k_{π} (3) is a reproducing kernel and satisfies a Stein identity ($\int_{\mathbb{R}^d} k_{\pi}(x, \cdot) d\pi(x) = 0$) under mild regularity assumptions on k and π .¹ It allows for several interpretations of the KSD (2) as already discussed by Liu et al. (2016). It can be introduced as an IPM in the specific case of a Stein operator applied to \mathcal{H}_k (e.g. Gorham & Mackey, 2017). It can then be identified as an asymmetric MMD in $\mathcal{H}_{k_{\pi}}$ (see Section 3.3). Alternatively, the squared KSD can be seen as a kernelized Fisher divergence, where the Fisher information $\|\nabla \log \left(\frac{d\mu}{d\pi}\right)\|_{L^2(\mu)}^2$ is smoothed through the kernel integral operator, i.e. $\mathrm{KSD}^2(\mu|\pi) = \|S_{\mu,k}\nabla \log \left(\frac{d\mu}{d\pi}\right)\|_{\mathcal{H}_k^d}^2$. In this sense, the squared KSD has also been referred to as the Stein Fisher information (Duncan et al., 2019). Hence, minimizing the KSD can be thought as a kernelized version of score-matching (Hyvärinen, 2005).

The metrization of weak convergence by the KSD, i.e. that $\lim_{t\to\infty} \text{KSD}(\mu_t | \pi) = 0$ implies the weak convergence of μ_t to π , depends on the choice of the kernel relatively to the target. This question has been considered by Gorham & Mackey (2017), who show this is the case under assumptions akin to strong log-concavity of π at infinity (namely distant dissipativity, Eberle, 2016), and for a kernel k with a slow decay rate. This includes finite Gaussian mixtures with common covariance and kernels that are translation-invariant with heavy-tails and non-vanishing Fourier transform, such as the inverse multi-quadratic (IMQ) kernel defined by $k(x,y) = (c^2 + ||x - y||_2^2)^{\beta}$ for c > 0 and $\beta \in (-1,0)$, or its variants considered in Chen et al. (2018).

3. Sampling as optimization of the KSD

This section defines the KSD Descent and relates it to other gradient flows, especially the MMD gradient flow, of which the KSD Descent is a special case. In all the following, we assume that $k \in C^{3,3}(\mathbb{R}^d \times \mathbb{R}^d, \mathbb{R})$, and that π is such that $s = \nabla \log \pi \in C^2(\mathbb{R}^d)$.

¹e.g., k is a Gaussian kernel and π is a smooth density fully supported on \mathbb{R}^d , see Liu et al. (2016, Theorem 3.7).

3.1. Continuous time dynamics

Consider the functional $\mathcal{F} : \mathcal{P}_2(\mathbb{R}^d) \to [0, +\infty), \mu \mapsto \frac{1}{2} \text{KSD}^2(\mu|\pi)$ defined over the Wasserstein space. If $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ satisfies some mild regularity conditions (i.e., it has a C^1 density w.r.t. Lebesgue measure, and it is in the domain of \mathcal{F} , see Appendix A.2), the gradient of \mathcal{F} at μ is well-defined and denoted by $\nabla_{W_2} \mathcal{F}(\mu) \in L^2(\mu)$. We shall consider the following assumptions on the Stein kernel:

- (A₁) There exists a map $L(\cdot) \in C^0(\mathbb{R}^d, \mathbb{R}_+)$, which is μ integrable for any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, such that, for any $y \in \mathbb{R}^d$, the maps $x \mapsto \nabla_1 k_\pi(x, y)$ and $x \mapsto \nabla_2 k_\pi(x, y)$ are L(y)-Lipschitz.
- (A₂) There exists m > 0 such that for any $\mu \in \mathcal{P}_c(\mathbb{R}^d)$, for all $y \in \mathbb{R}^d$, we have $\|\int \nabla_2 k_{\pi}(x, y) d\mu(x)\| \le m (1 + \|y\| + \int \|x\| d\mu(x))$.
- (A₃) The map $(x, y) \mapsto ||H_1k_{\pi}(x, y)||_{op}$ is $\mu \otimes \nu$ -integrable for every $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$.
- (A₄) For all $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, $\int k_{\pi}(x, x) d\mu(x) < \infty$.

The KSD gradient flow is defined as the flow induced by the continuity equation:

$$\frac{\partial \mu_t}{\partial t} + \operatorname{div}(\mu_t v_{\mu_t}) = 0 \text{ for } v_{\mu_t} := -\nabla_{W_2} \mathcal{F}(\mu_t).$$
 (7)

Assumptions (\mathbf{A}_1) and (\mathbf{A}_2) ensure that the KSD gradient flow exists and is unique, they are further discussed in Appendix A.3. Assumptions (\mathbf{A}_1) and (\mathbf{A}_3) are needed so that the Hessian of \mathcal{F} is well defined (see Section 4). Assumption (\mathbf{A}_4) guarantees that the integral operator $S_{\mu,k_{\pi}}$ (6) is well-defined and that $\mathcal{F}(\mu) < \infty$ for all $\mu \in \mathcal{P}_2(\mathbb{R}^d)$.

Lemma 1. Assume that k, its derivatives up to order 3, and their product by ||x - y|| are uniformly bounded over \mathbb{R}^d ; and that s is Lipschitz and has a bounded Hessian over \mathbb{R}^d . Then Assumptions (A₁), (A₃) and (A₄) hold. If, furthermore there exists M > 0 and M_0 such that, for all $x \in \mathbb{R}^d$, $||s(x)|| \leq M\sqrt{||x||} + M_0$, then Assumption (A₂) also holds.

See the proof in Appendix B.1. Smoothed Laplace distributions π paired with Gaussian k satisfy the assumptions of Lemma 1. For Gaussian π , s is linear, so Assumptions (\mathbf{A}_1), (\mathbf{A}_3) and (\mathbf{A}_4) hold for smooth kernels, but Assumption (\mathbf{A}_2) does not hold in general because of the $s(x)^{\top}s(y)$ term in k_{π} . Notice that most of our results are stated without Assumption (\mathbf{A}_2), which is only required to establish the global existence of KSD flow, in the sense that the particle trajectories are well-defined and do not explode in finite-time.

Proposition 2. Under Assumptions (A₁) and (A₂), the W_2 gradient of \mathcal{F} evaluated at μ and its dissipation (5) along (7) are

$$\nabla_{W_2} \mathcal{F}(\mu) = \mathbb{E}_{x \sim \mu} [\nabla_2 k_\pi(x, \cdot)], \tag{8}$$

$$\dot{\mathcal{F}}(\mu_t) = -\mathbb{E}_{y \sim \mu_t} \left[\|\mathbb{E}_{x \sim \mu_t} [\nabla_2 k_\pi(x, y)]\|^2 \right].$$
(9)

Since the r.h.s. of (9) is negative, Proposition 2 shows that the squared KSD w.r.t. π decreases along the KSD gradient flow dynamics. In other words, \mathcal{F} is indeed a Lyapunov functional for the dynamics (7) as discussed in Section 2.2.

3.2. Discrete time and discrete measures

A straightforward time-discretization of (7) is a gradient descent in the Wasserstein space applied to $\mathcal{F}(\mu) = \frac{1}{2} \text{KSD}^2(\mu | \pi)$. Starting from an initial distribution $\mu_0 \in \mathcal{P}_2(\mathbb{R}^d)$, it writes as follows at iteration $n \in \mathbb{N}$,

$$\mu_{n+1} = (I - \gamma \nabla_{W_2} \mathcal{F}(\mu_n))_{\#} \mu_n, \qquad (10)$$

for a step-size $\gamma > 0$. However for discrete measures $\hat{\mu} = \frac{1}{N} \sum_{j=1}^{N} \delta_{x^j}$, we can make the problem more explicit setting a loss function

$$F([x^j]_{j=1}^N) := \mathcal{F}(\hat{\mu}) = \frac{1}{2N^2} \sum_{i,j=1}^N k_\pi(x^i, x^j).$$
(11)

Problem (1) then corresponds to a standard non-convex optimization problem over the finite-dimensional, Euclidean, space of particle positions. The gradient of F is readily obtained as

$$\nabla_{x^i} F([x^j]_{j=1}^N) = \frac{1}{N^2} \sum_{j=1}^N \nabla_2 k_\pi(x^j, x^i).$$

since, by symmetry of k_{π} , $\nabla_1 k_{\pi}(x, y) = \nabla_2 k_{\pi}(y, x)$. As both *F* and $\nabla_{x^i} F$ can be explicitly computed, one can implement the KSD Descent either using a gradient descent (Algorithm 1) or through a quasi-Newton algorithm such as L-BFGS (Algorithm 2). As a matter of fact, L-BFGS (Liu & Nocedal, 1989) is often faster and more robust than the conventional gradient descent. It also does not require choosing critical hyper-parameters, such as a learning rate, since L-BFGS performs a line-search to find suitable step-sizes. It only requires a tolerance parameter on the norm of the gradient, which is in practice set to machine precision. A techni-

Algorithm 1 KSD Descent GD

Input: initial particles $(x_0^i)_{i=1}^N \sim \mu_0$, number of iterations M, step-size γ

for
$$n = 1$$
 to M do
 $[x_{n+1}^i]_{i=1}^N = [x_n^i]_{i=1}^N - \frac{\gamma}{N^2} \sum_{j=1}^N [\nabla_2 k_\pi(x_n^j, x_n^i)]_{i=1}^N,$
end for
Return: $[x_M^i]_{i=1}^N.$
(12)

cal descent lemma for (10) (Proposition 14) showing that \mathcal{F} decreases at each iteration (10) is provided in Appendix A.6. It requires the boundedness of $(\|L(\cdot)\|_{L^2(\mu_n)})_{n\geq 0}$, the L^2 -norm of the Lipschitz constants of Assumption (**A**₁) along the flow, as well as the convexity of $L(\cdot)$ and a compactly-supported initialization.

Algorithm 2 KSD Descent L-BFGS				
Input: initial particles $(x_0^i)_{i=1}^N \sim \mu_0$, tolerance tol				
Return: $[x_*^i]_{i=1}^N = L$ -BFGS $(F, \nabla F, [x_0^i]_{i=1}^N, \text{tol})$.				

Remark 1. As L-BFGS requires access to exact gradients, Algorithm 2 cannot be used in a stochastic setting. However this can be done for Algorithm 1 by subsampling over particles in the double sum in Equation (11). Moreover, in some settings like Bayesian inference, the score itself writes as a sum over observations. In this case, the loss F writes as a *double* sum over observations, and a stochastic variant of Algorithm 1 tailored for this problem could be devised, in the spirit of Clémençon et al. (2016).

3.3. Related work

Several recent works fall within the framework sketched in Section 2.2. In SVGD, Liu et al. (2016) take \mathcal{F} as the KL, and set $v_{\mu_t} = -S_{\mu,k} \nabla \ln \left(\frac{d\mu_t}{d\pi} \right)$ to obtain $\dot{\mathcal{F}}$ as the (squared) KSD. Integrating this inequality w.r.t. time yields a 1/T convergence rate for the average KSD between μ_t and π for $t \in [0,T]$. This enabled Korba et al. (2020, Proposition 5, Corollary 6) to obtain a discrete-time descent lemma for bounded kernels, as well as rates of convergence for the averaged KSD. In contrast, since the dissipation (9) of the KSD along its W_2 gradient flow does not correspond to any dissimilarity, our descent lemma for (10) (Proposition 14) does not yield similar rates of convergence. Alternatively, in the LAWGD algorithm recently proposed by Chewi et al. (2020), \mathcal{F} is the KL, and $v_{\mu_t} = -\nabla S_{\pi,k_{\mathcal{L}_{\pi}}}\left(\frac{d\mu_t}{d\pi}\right)$ with $k_{\mathcal{L}_{\pi}}$ chosen such that $\dot{\mathcal{F}}$ is the χ^2 , by taking $S_{\pi,k_{\mathcal{L}_{\pi}}}$ as the inverse of the diffusion operator:

$$\mathcal{L}_{\pi}: f \mapsto -\Delta f - \langle \nabla \log \pi, \nabla f \rangle. \tag{13}$$

Their elegant approach results in a linear convergence of the KL along their flow, but implementing LAWGD in practice requires to compute the spectrum of \mathcal{L}_{π} . It is in general as difficult as solving a linear PDE, and Chewi et al. (2020) admit it is unlikely to scale in high dimensions.²

Beyond studies on the KL, Mroueh et al. (2019) considered \mathcal{F} as the MMD and pick v_{μ_t} based on a kernelized Sobolev norm so that $\dot{\mathcal{F}}$ resembles the MMD, but without proving convergence of their scheme. Arbel et al. (2019) also analyzed \mathcal{F} as the MMD, but for $v_{\mu_t} = -\nabla_{W_2} \frac{1}{2} \text{ MMD}^2(\mu_t, \pi)$ with similar $\dot{\mathcal{F}}$ as ours and with a dedicated analysis of their MMD-GD flow. We recall that $\text{MMD}^2(\mu, \pi) = \|\int k(x, \cdot)d\mu(x) - \int k(x, \cdot)d\pi(x)\|_{\mathcal{H}_k}^2$. Since the Stein kernel satisfies the Stein's identity $\int k_{\pi}(x, \cdot)d\pi(x) = 0$, the

KSD (2) can be identified to an MMD with the Stein kernel (Chen et al., 2018). However, the assumptions of Arbel et al. (2019) $-\nabla k$ is *L*-Lipschitz for $L \in \mathbb{R}^+$ - do not hold in general for unbounded Stein kernels. Here, we provide the right set of assumptions (**A**₁)-(**A**₂) on k_{π} for the flow to exist and for a descent lemma to hold. Also, as noted on Figure 1, the sample-based MMD flow, defined through

$$\nabla_{W_2} \frac{1}{2} \operatorname{MMD}^2(\mu, \pi) = \int \nabla_2 k(x, \cdot) \mathrm{d}(\mu - \pi)(x) \quad (14)$$

can fail dramatically while KSD flow succeeds. This suggests that the geometrical properties of the KSD flow are more favorable than the ones of the regular MMD flow. In other words, choosing an appropriate (target-dependent) kernel, as in our method or in LAWGD, appears more propitious than taking a kernel k unrelated to π .

Related to the optimization of the KSD, Stein points (Chen et al., 2018) also propose to use the KSD loss for sampling, but the loss is minimized using very different tools. While KSD descent uses a first order information by following the gradient (or L-BFGS) direction with a fixed number of particles, Stein points use a Frank-Wolfe scheme, adding particles one by one in a greedy fashion. Given N particles x^1, \ldots, x^N , in Stein points, the next particle is set as

$$x^{N+1} \in \operatorname*{arg\,min}_{x} \frac{1}{2} k_{\pi}(x, x) + \sum_{i=1}^{N} k_{\pi}(x, x^{i}).$$

This problem is solved using derivative-free (a.k.a. zerothorder) algorithms like grid-search or random sampling. The main drawback of such an approach is that it scales poorly with the dimension d when compared to first-order methods. In the same spirit, (Futami et al., 2019) have recently proposed a similar algorithm to optimize the MMD.

4. Theoretical properties of the KSD flow

In this section, we provide a theoretical study of the convergence of the KSD Wasserstein gradient flow, assessing the convexity of \mathcal{F} and discussing the stationary points of its gradient flow. Remarkably, we encounter pitfalls similar to other deterministic flows derived from IPMs. This issue arises because IPMs are always mixture convex, but seldom geodesically convex. We first investigate the convexity properties of \mathcal{F} along W_2 geodesics and show that exponential convergence near equilibrium cannot hold in general. Then, we examine some stationary points of its W_2 gradient flow, which explain the failure cases met in Section 5, where $\hat{\mu}_n$ converges to a degenerate measure.

4.1. Convexity properties of the KSD flow

As is well-known, decay along W_2 gradient flows can be obtained from convexity properties along geodesics. A natural object of interest is then the Hessian of the objective

²The update rules of SVGD, LAWGD and MMD-GD can be found in Appendix A.4.

 \mathcal{F} . We define below this object, in a similar way as Duncan et al. (2019). We recall that $\{\nabla \psi, \psi \in C_c^{\infty}(\mathbb{R}^d)\}$ is by definition dense in $\mathcal{T}_{\mu}\mathcal{P}_2(\mathbb{R}^d) \subset L^2(\mu)$ for any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ (Ambrosio et al., 2008, Definition 8.4.1).

Definition 1. Consider $\psi \in C_c^{\infty}(\mathbb{R}^d)$ and the path ρ_t from μ to $(I + \nabla \psi)_{\#}\mu$ given by: $\rho_t = (I + t\nabla \psi)_{\#}\mu$, for all $t \in [0, 1]$. The Hessian of \mathcal{F} at μ , $H\mathcal{F}_{|\mu}$, is defined as a symmetric bilinear form on $C_c^{\infty}(\mathbb{R}^d)$ associated with the quadratic form $\operatorname{Hess}_{\mu} \mathcal{F}(\psi, \psi) := \frac{d^2}{dt^2}\Big|_{t=0} \mathcal{F}(\rho_t)$.

Definition 1 can be straightforwardly related to the usual symmetric bilinear form defined on $\mathcal{T}_{\mu}\mathcal{P}_{2}(\mathbb{R}^{d}) \times \mathcal{T}_{\mu}\mathcal{P}_{2}(\mathbb{R}^{d})$ (Otto & Villani, 2000, Section 3)³.

Proposition 3. Under Assumptions (A₁) and (A₃), the Hessian of \mathcal{F} at μ is given, for any $\psi \in C_c^{\infty}(\mathbb{R}^d)$, by

$$\operatorname{Hess}_{\mu} \mathcal{F}(\psi, \psi) = \mathbb{E}_{x, y \sim \mu} \left[\nabla \psi(x)^{T} \nabla_{1} \nabla_{2} k_{\pi}(x, y) \nabla \psi(y) \right] \\ + \mathbb{E}_{x, y \sim \mu} \left[\nabla \psi(x)^{T} H_{1} k_{\pi}(x, y) \nabla \psi(x) \right].$$
(15)

A proof of Proposition 3 is provided in Appendix B.3. Our computations are similar to the ones in (Arbel et al., 2019, Lemma 23) with some terms getting simpler owing to the Stein's identity satisfied by the Stein kernel. As for the squared MMD (Arbel et al., 2019, Proposition 5), the squared KSD is unlikely to be geodesically convex. Indeed, while the first term is always positive, the second term in (15) can in general take negative values, unless $H_1k_{\pi}(x, y)$ is positive for all $x, y \in \text{supp}(\mu)$. Nevertheless, at $\mu = \pi$, this second term vanishes, again owing to the Stein's property of k_{π} .

Corollary 4. Under Assumptions (A₁), (A₃) and (A₄), the Hessian of \mathcal{F} at π is given, for any $\psi \in C_c^{\infty}(\mathbb{R}^d)$, by

$$\operatorname{Hess}_{\pi} \mathcal{F}(\psi, \psi) = \|S_{\pi, k_{\pi}} \mathcal{L}_{\pi} \psi\|_{\mathcal{H}_{k_{\pi}}}^2$$

where $S_{\pi,k_{\pi}}$ and \mathcal{L}_{π} are defined in Equations (6) and (13).

A proof of Corollary 4 is provided in Appendix B.4. We now study the curvature properties near equilibrium, characterized by $\text{Hess}_{\pi} \mathcal{F}$. In particular, inspired by the methodology described in Villani (2003) and recently applied by Duncan et al. (2019), we expect exponential convergence of solutions initialized near π whenever the Hessian is bounded from below by a quadratic form on the tangent space of $\mathcal{P}_2(\mathbb{R}^d)$ at π , included in $L^2(\pi)$.

Definition 2. We say that *exponential decay near equilibrium* holds if there exists $\lambda > 0$ such that for any $\psi \in C_c^{\infty}(\mathbb{R}^d)$,

$$\operatorname{Hess}_{\pi} \mathcal{F}(\psi, \psi) \ge \lambda \|\nabla \psi\|_{L^{2}(\pi)}^{2}.$$
 (16)

³The W_2 Hessian of \mathcal{F} , denoted $H\mathcal{F}_{|\mu}$ is an operator over $\mathcal{T}_{\mu}\mathcal{P}_2(\mathbb{R}^d)$ verifying $\langle H\mathcal{F}_{|\mu}v_t, v_t \rangle_{L^2(\mu)} = \frac{d^2}{dt^2}\Big|_{t=0} \mathcal{F}(\rho_t)$ if $t \mapsto \rho_t$ is a geodesic starting at μ with vector field $t \mapsto v_t$.

According to Corollary 4, (16) can be seen as a kernelized version of the following form of the Poincaré inequality for π (Bakry et al., 2013, Chapter 5)

$$\|\mathcal{L}_{\pi}\psi\|_{L_{2}(\pi)}^{2} \geq \lambda_{\pi} \|\nabla\psi\|_{L_{2}(\pi)}^{2}.$$
 (17)

Condition (16) is similar to the Stein-Poincaré inequality (Duncan et al., 2019, Lemma 32). We will now argue that (16) is hardly ever satisfied, thus obtaining an impossibility result reminiscent of the one for SVGD in (Duncan et al., 2019, Lemma 36), which states that exponential convergence (of the KL) for the SVGD gradient flow does not hold whenever π has exponential tails and the derivatives of $\nabla \log \pi$ and k grow at most at a polynomial rate. We start with the following characterization of exponential decay near equilibrium:

Proposition 5. Let $T_{\pi,k_{\pi}} = S^*_{\pi,k_{\pi}} \circ S_{\pi,k_{\pi}}$ and $L^2_0(\pi) = \{\phi \in L^2(\pi), \int \phi d\pi = 0\}$. The exponential decay near equilibrium (16) holds if and only if $\mathcal{L}^{-1}_{\pi} : L^2_0(\pi) \to L^2_0(\pi)$, the inverse of $\mathcal{L}_{\pi}|_{L^2_0(\pi)}$, is well-defined, bounded, and for all $\phi \in L^2_0(\pi)$ we have

$$\langle \phi, T_{\pi,k_{\pi}}\phi \rangle_{L_2(\pi)} \ge \lambda \langle \phi, \mathcal{L}_{\pi}^{-1}\phi \rangle_{L_2(\pi)}.$$
 (18)

See Appendix B.5 for a proof. By the spectral theorem for compact, self-adjoint operators (Kreyszig, 1978, Section 8.3), $T_{\pi,k_{\pi}}$ has a discrete spectrum $(l_n)_{n\in\mathbb{N}^*}$ which satisfies $l_n \geq 0$ and $l_n \rightarrow 0$. Under mild assumptions on π , the operator \mathcal{L}_{π} also has a discrete, positive spectrum (Chewi et al., 2020, Appendix A). Proposition 5 implies the following necessary condition on the spectrum of \mathcal{L}_{π}^{-1} and $T_{\pi,k_{\pi}}$ for the exponential decay near equilibrium (16) to hold.

Corollary 6. If \mathcal{L}_{π}^{-1} has a discrete spectrum $(\lambda_n)_{n \in \mathbb{N}^*}$ and (16) holds, then $\lambda_n = \mathcal{O}(l_n)$, i.e. the eigenvalue decay of \mathcal{L}_{π}^{-1} is at least as fast as the one of $T_{\pi,k_{\pi}}$.

We also show that if $\mathcal{H}_{k_{\pi}}$ is infinite dimensional and exponential convergence near equilibrium holds, then \mathcal{L}_{π}^{-1} has a discrete spectrum (Lemma 16). We now present our impossibility result on the exponential decay near equilibrium.

Theorem 7. Let $\pi \propto e^{-V}$. Assume that $V \in C^2(\mathbb{R}^d)$, ∇V is Lipschitz and \mathcal{L}_{π} has a discrete spectrum. Then exponential decay near equilibium does not hold.

The main idea behind the proof of Theorem 7 (Appendix B.7) is that $T_{\pi,k_{\pi}}$ is nuclear (Steinwart & Christmann, 2008, Theorem 4.27), which implies that its eigenvalues $(l_n)_{n \in \mathbb{N}^*}$ are summable. On the other hand the eigenvalue decay of \mathcal{L}_{π}^{-1} when π is a Gaussian can be seen to be of order $O(1/n^{1/d})$ (Lemma 17 in Appendix B.7), which is not summable. The general case is obtained by comparison with a Gaussian. Remark that by Lemma 16, the assumption that \mathcal{L}_{π} has a discrete spectrum in Theorem 7 can be exchanged for an assumption that $\mathcal{H}_{k_{\pi}}$ is infinite dimensional, which is discussed in the proof of Theorem 7. Despite the

lack of (strong) geodesic convexity near equilibrium, we still observe empirically good convergence properties of the KSD flow for discrete measures to a stationary measure. Hence, we now investigate these stationary measures.

4.2. Stationary measures of the KSD flow

The KSD gradient flow leads to a deterministic algorithm, as for SVGD and LAWGD. To study the convergence of these algorithms in continuous time, it is relevant to characterize the stationary measures, i.e. the ones which cancel the dissipation $\dot{\mathcal{F}}$ (5) of the objective functional \mathcal{F} along the relative gradient flow dynamics. Unfortunately, unlike for the SVGD and LAWGD algorithms, the dissipation related to the KSD flow (9) does not yield a dissimilarity between measures. Consequently, the study of the stationary measures of the KSD is more involved. We discuss below when failure cases may happen.

Lemma 8. Assume Assumption (\mathbf{A}_4) holds. Then $\mathcal{H}_{k_{\pi}}$ does not contain non-zero constant functions.

A proof of Lemma 8 is provided in Appendix B.8. This result has the immediate consequence of considerably restricting the number of candidate fully-supported measures that are stationary for the KSD gradient flow. Consider one such measure μ_{∞} . At equilibrium, $\dot{\mathcal{F}}(\mu_{\infty}) = 0$; which implies that $\int k_{\pi}(x, .)d\mu_{\infty}(x)$ is μ_{∞} -a.e. equal to a constant function c. Since $x \mapsto c$ is then also an element of $\mathcal{H}_{k_{\pi}}$, the previous lemma implies that c = 0. Hence, if μ_{∞} and π are full-support, $\mathcal{F}(\mu_{\infty}) = 0$. Provided k_{π} is characteristic (Sriperumbudur et al., 2011), then $\mu_{\infty} = \pi$.

However, as Algorithms 1 and 2 rely on discrete measures, the dissipation $\dot{\mathcal{F}}$ (9) can vanish even for $\mu \neq \pi$ because μ is not full-support. Depending on the properties of π and k, this may happen even for trivial measures such as a single Dirac mass, as stated in the following Lemma.

Lemma 9. Let x_0 such that $s(x_0) = 0$ and $Js(x_0)$ is invertible, and consider a translation-invariant kernel $k(x, y) = \phi(x - y)$, for $\phi \in C^3(\mathbb{R}^d)$. Then δ_{x_0} is a stable fixed measure of (7), i.e. it is stationary and any small push-forward of δ_{x_0} is attracted back by the flow.

Proof: For $\varepsilon > 0$ and $\psi \in C_c^{\infty}(\mathbb{R}^d)$, set $\mu_{\varepsilon} = (I + \varepsilon \psi)_{\#} \delta_{x_0}$. We then have $\mathcal{F}(\mu_{\varepsilon}) = \frac{1}{2} k_{\pi}(x_0 + \varepsilon \psi(x_0), x_0 + \varepsilon \psi(x_0))$. Expanding $k_{\pi}(x, x)$ at the first order around $x = x_0$ gives $2\mathcal{F}(\mu_{\varepsilon}) = \varepsilon^2 ||[Js(x_0)]\psi(x_0)||^2\phi(0) - \Delta\phi(0) + o(\varepsilon^2)$. This quantity is minimized for $\psi(x_0) = 0$, which shows that δ_{x_0} is indeed a local minimum for \mathcal{F} .

Importantly, this result applies whenever the score *s* vanishes at x_0 , not only when x_0 is a local stable minimum of the potential $\log(\pi)$. This means that for a single particle, KSD descent is attracted to any stationary point of $\log(\pi)$, whereas SVGD converges only to local maxima of $\log(\pi)$

(Liu & Wang, 2016). Nonetheless, if π is log-concave, there is no spurious stationary point.

For cases more general than Lemma 9, we are interested in the sets that are kept invariant by the gradient flow. For these sets, an erroneous initialization may prevent the particles from reaching the support of π . We provide below a general result holding for any deterministic flow, beyond our specific choice (7) of v_{μ_t} , and thus holding also for SVGD.

Definition 3. Let $\mathcal{M} \subset \mathbb{R}^d$ be a closed nonempty set. We say that \mathcal{M} is a flow-invariant support set for the flow $(\mu_t)_{t\geq 0}$ of (4) if for any μ_0 s.t. $\operatorname{supp}(\mu_0) \subset \mathcal{M}$, we have that the flow verifies $\operatorname{supp}(\mu_t) \subset \mathcal{M}$ for all $t \geq 0$.

Proposition 10. (Informal) Let $\mathcal{M} \subset \mathbb{R}^d$ be a smooth nonempty submanifold and $\mu_0 \in \mathcal{P}_c(\mathbb{R}^d)$ with $\operatorname{supp}(\mu_0) \subset \mathcal{M}$. Assume that, for a deterministic $(v_{\mu_t})_{t\geq 0}$ satisfying classical Caratheodory-Lipschitz assumptions (Appendix A.3), we have $v_{\mu_t}(x) \in T_{\mathcal{M}}(x)$ where $T_{\mathcal{M}}(x)$ is the tangent space to \mathcal{M} at x. Then \mathcal{M} is flow-invariant for (4).

The formal statement, Proposition 20, stated and proved in Appendix B.9, can be in particular applied to the ubiquitous radial kernels and to planes of symmetry of π , i.e. affine subspaces $\mathcal{M} \subset \mathbb{R}^d$ such that the density of π is symmetric w.r.t. \mathcal{M} . Lemma 11 is illustrated in Section 5 for a mixture of two Gaussians with the same variance (Figure 3).

Lemma 11. Let \mathcal{M} be a plane of symmetry of π and consider a radial kernel $k(x, y) = \phi(||x - y||^2/2)$ with $\phi \in C^3(\mathbb{R})$. Then, for all $(x, y) \in \mathcal{M}^2$, $\nabla_2 k_{\pi}(x, y) \in T_{\mathcal{M}}(x)$ and \mathcal{M} is flow-invariant for (7).

Proof idea: We show that all the terms in $\nabla_2 k_{\pi}(x, y)$ belong to $T_{\mathcal{M}}(x)$. This implies that the convex combination $\nabla_{W_2} \mathcal{F}(\mu)(y) = \mathbb{E}_{x \sim \mu}[\nabla_2 k_{\pi}(x, y)] \in T_{\mathcal{M}}(x)$. We then apply Proposition 10 to conclude.

5. Experiments

In this section, we discuss the performance of KSD Descent to sample from π in practice, on toy examples and realworld problems. The code to reproduce the experiments and a package to use KSD Descent are available at https: //github.com/pierreablin/ksddescent. For all our experiments, we use a Gaussian kernel, as we did not notice any difference in practice w.r.t. the IMQ kernel. Its bandwith is selected by cross-validation. Implementation details and additional experiments can be found in Appendix D.

5.1. Toy examples

In the first example, we choose π to be a standard 2D Gaussian, and a Gaussian k with unit bandwidth. We initialize with 50 particles drawn from a Gaussian with mean (1, 1). Figure 1 displays the trajectories of several different meth-



Figure 1. Toy example with 2D standard Gaussian. The green points represent the initial positions of the particles. The light grey curves correspond to their trajectories under the different v_{μ_t} .



Figure 2. Convergence speed of KSD and SVGD on a Gaussian problem in 1D, with 30 particles.

ods: SVGD, KSD Descent implemented using gradient descent (Algorithm 1) and L-BFGS (Algorithm 2), and the MMD flow (Arbel et al., 2019). To assess the convergence of the algorithms, for SVGD we monitored the norm of the displacement, while for the KSD and MMD gradient flows we used the tolerance parameter of L-BFGS. KSD Descent successfully pushes the particles towards the target distribution, with a final result that is well-distributed around the mode. While KSD performs similarly to SVGD, we can notice that the trajectories of the particles behave very differently. Indeed, SVGD trajectories appear to be at first driven by the score term in the update, while the repulsive term acts later to spread the particles around the mode. In contrast, trajectories of the particles updated by KSD Descent are first influenced by the last repulsive term of the update, which seems to determine their global layout, and are then translated and contracted towards the mode under the action of the driving terms. Finally, for the MMD descent, some particles collapse around the mode, while others stay far from the target. This behavior was documented in Arbel et al. (2019), and can be partly circumvented by injecting some noise in the updates.

We then compare the convergence speed of KSD Descent and SVGD in terms of the KL or KSD objective (see Figure 2). With a fine-tuned step-size, SVGD is the fastest method. However, taking a step-size too large leads to nonconvergence, while one too small leads to slow convergence. It should be stressed that it is hard to select a good step-size, or to implement a line-search strategy, since SVGD does not minimize a simple function. In contrast, the empirical KSD (11) can be minimized using L-BFGS, which does not have any critical hyper-parameter.



Figure 3. KSD Descent applied on a balanced mixture of Gaussian with small variance (0.1) in 2D. The centroids are at (-1, 0) and (1, 0). The green crosses indicate the initial particle positions, while the blue ones are the final positions. The light red arrows correspond to the score directions. Left: the initial positions are all chosen close to the line x = 0, which corresponds to an axis of symmetry of π . Right: even when the initialization is more spread, some particles are still caught in this spurious local minimum.

In our second example, we apply KSD Descent for π taken as a symmetric mixture of two Gaussians with the same variance. This highlights the results of Section 4.2. If initialized on the axis of symmetry, the particles are indeed stuck on it, as stated in Lemma 11. We noticed that, for a large variance of π (e.g. in [0.2, 1]), this axis is unstable. However, when the variance is too small (e.g. set to 0.1 as in Figure 3), the axis can even become a locally stable set. We also observed that, for a distribution initialized exclusively on one side of the axis, a single component of the mixture can be favored. This is a classical behavior of score-based methods, depending typically on the variance of π (Wenliang, 2020).

To fix this issue, we consider an annealing strategy as suggested by Wenliang (2020). It consists in adding an inverse temperature variable β to the log density of the model, i.e. $\pi^{\beta}(x) \propto \exp(-\beta V(x))$ for $\pi(x) \propto \exp(-V(x))$, with $\beta \in (0, 1]$. This is easily implemented with score-based methods, since it simply corresponds to multiplying s(x) by β . When β is small, annealing smoothes the target distribution and the last term of the Stein kernel, repulsive at short distance, becomes dominant; on the other hand, for β close to 1, we recover the true log density. To implement this method, we start with $\beta = 0.1$, and run the KSD Descent to obtain particles at 'high temperature'. KSD Descent is then re-run starting from these particles, setting now $\beta = 1$. One can see that this strategy successfully solves the issues encountered when the KSD flow was failing to converge to the target π (Figure 4). This correction differs from the noiseinjection strategy proposed in Arbel et al. (2019) for the MMD flow, which is rather related to randomized smoothing (Duchi et al., 2012). Noise-injection would prevent the use of L-BFGS in our case, as it requires exact values of the gradients of previous iterations. Annealing on the other hand is compatible with Algorithm 2.



Figure 4. Effect of the annealing strategy on a mixture of Gaussians. Left: without annealing, some particles fall into a spurious minimum. Middle: with a higher temperature ($\beta = 0.1$), the particles are more spread out. **Right**: starting from the particles in the middle figure and setting $\beta = 1$ we converge to a distribution which minimizes the KSD, and has no spurious particles.



Figure 5. Bayesian ICA results. Left: p = 2. Middle: p = 4. Right: p = 8. Each dot corresponds to the Amari distance between an estimated matrix and the true unmixing matrix.

5.2. Bayesian Independent Component Analysis

Independent Component Analysis (ICA, Comon, 1994) is the generative model $x = W^{-1}s$, where x is an observed sample in \mathbb{R}^p , $W \in \mathbb{R}^{p \times p}$ is the unknown square unmixing matrix, and $s \in \mathbb{R}^p$ are the independent sources. We assume that each component has the same density $s_i \sim p_s$. The likelihood of the model is p(x|W) = $\log |W| + \sum_{i=1}^{p} p_s([Wx]_i)$. For our prior, we assume that W has i.i.d. entries, of law $\mathcal{N}(0, 1)$. The posterior is $p(W|x) \propto p(x|W)p(W)$, and the score is given by $s(W) = W^{-\top} - \psi(Wx)x^{\top} - W$, where $\psi = -\frac{p'_s}{p_s}$. In practice, we choose p_s such that $\psi(\cdot) = \tanh(\cdot)$. We then use the presented algorithms to draw particles $W \sim p(W|x)$. We use N = 10 particles, and take 1000 samples x from the ICA model for $p \in \{2, 4, 8\}$. Each method outputs N estimated unmixing matrices, $[\tilde{W}_i]_{i=1}^N$. We compute the Amari distance (Amari et al., 1996) between each \tilde{W}_i and W: the Amari distance vanishes if and only if the two matrices are the same up to scale and permutation, which are the natural indeterminacies of ICA. We repeat the experiment 50 times, resulting in 500 values for each algorithm (Figure 5). We also add the results of a random output, where the estimated matrices are obtained with i.i.d. $\mathcal{N}(0,1)$ entries. We see that for this experiment, KSD performs barely better than random, while SVGD finds matrices with lower Amari distance. One explanation is that the ICA likelihood is highly non-convex (Cardoso, 1998). This is easily seen with the invariances of the problem: permuting the rows of W does not change p(x|W). As a consequence, the posterior has

many saddle points, in which particles might get trapped. Unfortunately, the annealing strategy proposed above did not improve the achieved performance for this problem.

5.3. Real-world data

We compare KSD Descent and SVGD in the Bayesian logistic regression setting described in Gershman et al. (2012); Liu & Wang (2016). Given datapoints $d_1, \ldots, d_q \in \mathbb{R}^p$, and labels $y_1, \ldots, y_q \in \{\pm 1\}$, the labels y_i are modelled as $p(y_i = 1 | d_i, w) = (1 + \exp(-w^{\top} d_i))^{-1}$ for some $w \in \mathbb{R}^p$. The parameters w follow the law $p(w|\alpha) =$ $\mathcal{N}(0, \alpha^{-1}I_p)$, and $\alpha > 0$ is drawn from an exponential law $p(\alpha) = \exp(0.01)$. The parameter vector is then $x = [w, \log(\alpha)] \in \mathbb{R}^{p+1}$, and we use Algorithm 2 to obtain samples from $p(x|(d_i, y_i)_{i=1}^q)$ for 13 datasets, with N = 10particles for each.

The learning rate for SVGD and the bandwidth of the kernel for both methods are chosen through grid-search, and for each problem we select the hyper-parameters yielding the best test accuracy. For all problems, the running times of SVGD with the best step-size and of KSD Descent were similar, while KSD Descent has the advantage of having one less hyper-parameter. We present on Figure 6 the accuracy of each method on each dataset, where KSD Descent was applied without annealing since it did not change the numerical results. Our results show we match the SVGD performance without having to fine-tune the step-size, owing to Algorithm 2. We posit that KSD succeeds on this task because the posterior $p(x|(d_i, y_i)_{i=1}^q)$ is log-concave, and does not have saddle points.



Discussion. KSD Descent benefits from a tractable loss and can be straightforwardly implemented with L-BFGS, achieving performance on par with SVGD on convex problems. However its dissipation has non-trivial stationary points, which prevents its use for non-convex problems with saddlepoints such as ICA. Convergence of kernel-based sampling schemes is known to be difficult, and we provided some intuitions on the reasons for it. This leaves the door open to a more in-depth analysis of kernel-based gradient flows, especially for unbounded kernels.

Acknowledgments. A.K. thanks the GENES, and S.M. the A.N.R ABSint (Projet-ANR-18-CE40-0034) for the financial support. P.A. thanks A.N.R. ANR19-P3IA-0001.

References

- Amari, S.-I., Cichocki, A., and Yang, H. H. A new learning algorithm for blind signal separation. In Advances in Neural Information Processing Systems, pp. 757–763, 1996.
- Ambrosio, L. and Crippa, G. Continuity equations and ODE flows with non-smooth velocity. *Proceedings of the Royal Society of Edinburgh: Section A Mathematics*, 144 (6):1191–1244, 2014.
- Ambrosio, L., Gigli, N., and Savaré, G. Gradient flows: in metric spaces and in the space of probability measures. Springer Science & Business Media, 2008.
- Anastasiou, A., Barp, A., Briol, F.-X., Ebner, B., Gaunt, R. E., Ghaderinezhad, F., Gorham, J., Gretton, A., Ley, C., Liu, Q., Mackey, L., Oates, C. J., Reinert, G., and Swan, Y. Stein's method meets statistics: A review of some recent developments. arXiv preprint arXiv:2105.03481, 2021.
- Arbel, M., Korba, A., Salim, A., and Gretton, A. Maximum mean discrepancy gradient flow. *Advances in Neural Information Processing Systems*, 2019.
- Aubin, J.-P. and Frankowska, H. *Set-Valued Analysis*. Birkhäuser Boston, 1990.
- Bakry, D., Gentil, I., and Ledoux, M. Analysis and geometry of Markov diffusion operators, volume 348. Springer Science & Business Media, 2013.
- Bonnet, B. and Frankowska, H. Differential inclusions in Wasserstein spaces: The Cauchy-Lipschitz framework. *Journal of Differential Equations*, 271:594–637, January 2021.
- Brezis, H. Functional analysis, Sobolev spaces and partial differential equations. Springer Science & Business Media, 2010.
- Cardoso, J.-F. Blind signal separation: statistical principles. *Proceedings of the IEEE*, 86(10):2009–2025, 1998.
- Chen, W. Y., Mackey, L., Gorham, J., Briol, F.-X., and Oates, C. J. Stein points. *International Conference on Machine Learning (ICML)*, 2018.
- Chen, W. Y., Barp, A., Briol, F.-X., Gorham, J., Girolami, M., Mackey, L., Oates, C., et al. Stein point Markov Chain Monte Carlo. *International Conference on Machine Learning (ICML)*, 2019.
- Chewi, S., Gouic, T. L., Lu, C., Maunu, T., and Rigollet, P. SVGD as a kernelized Wasserstein gradient flow of the chi-squared divergence. *Advances in Neural Information Processing Systems*, 2020.

- Chwialkowski, K., Strathmann, H., and Gretton, A. A kernel test of goodness of fit. In *International Conference on Machine Learning (ICML)*, pp. 2606–2615, 2016.
- Clémençon, S., Colin, I., and Bellet, A. Scaling-up empirical risk minimization: optimization of incomplete u-statistics. *Journal of Machine Learning Research*, 17 (1):2682–2717, 2016.
- Comon, P. Independent component analysis, a new concept? *Signal processing*, 36(3):287–314, 1994.
- Duchi, J. C., Bartlett, P. L., and Wainwright, M. J. Randomized smoothing for stochastic optimization. SIAM Journal on Optimization, 22(2):674–701, 2012.
- Duncan, A., Nüsken, N., and Szpruch, L. On the geometry of Stein variational gradient descent. arXiv preprint arXiv:1912.00894, 2019.
- Durmus, A., Majewski, S., and Miasojedow, B. Analysis of Langevin Monte Carlo via convex optimization. *Journal* of Machine Learning Research, 20(73):1–46, 2019.
- Eberle, A. Reflection couplings and contraction rates for diffusions. *Probability theory and related fields*, 166(3): 851–886, 2016.
- Fisher, M. A., Nolan, T., Graham, M. M., Prangle, D., and Oates, C. J. Measure transport with kernel Stein discrepancy. arXiv preprint arXiv:2010.11779, 130:1054–1062, 2021.
- Futami, F., Cui, Z., Sato, I., and Sugiyama, M. Bayesian posterior approximation via greedy particle optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 3606–3613, 2019.
- Gershman, S., Hoffman, M., and Blei, D. Nonparametric variational inference. In *International Conference on Machine Learning (ICML)*, pp. 235–242, 2012.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, volume 27, pp. 2672–2680, 2014.
- Gorham, J. and Mackey, L. Measuring sample quality with kernels. In *International Conference on Machine Learning (ICML)*, volume 70, pp. 1292–1301. JMLR. org, 2017.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. A kernel two-sample test. *Journal of Machine Learning Research*, 13:723–773, 2012.

- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., del R'10, J. F., Wiebe, M., Peterson, P., G'erard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., and Oliphant, T. E. Array programming with NumPy. *Nature*, 585(7825):357–362, September 2020. doi: 10. 1038/s41586-020-2649-2. URL https://doi.org/ 10.1038/s41586-020-2649-2.
- Hodgkinson, L., Salomone, R., and Roosta, F. The reproducing Stein kernel approach for post-hoc corrected sampling. arXiv preprint arXiv:2001.09266, 2020.
- Hunter, J. D. Matplotlib: A 2d graphics environment. Computing in Science & Engineering, 9(3):90–95, 2007. doi: 10.1109/MCSE.2007.55.
- Hyvärinen, A. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(24):695–709, 2005.
- Jordan, R., Kinderlehrer, D., and Otto, F. The variational formulation of the Fokker-Planck equation. *SIAM Journal on Mathematical Analysis*, 29(1):1–17, 1998.
- Kanagawa, H., Jitkrittum, W., Mackey, L., Fukumizu, K., and Gretton, A. A kernel Stein test for comparing latent variable models. arXiv preprint arXiv:1907.00586, 2020.
- Klenke, A. *Probability theory: a comprehensive course*. Springer Science & Business Media, 2013.
- Korba, A., Salim, A., Arbel, M., Luise, G., and Gretton, A. A non-asymptotic analysis for Stein variational gradient descent. *Advances in Neural Information Processing Systems*, 33, 2020.
- Kreyszig, E. Introductory functional analysis with applications, volume 1. Wiley New York, 1978.
- Liu, D. C. and Nocedal, J. On the limited memory BFGS method for large scale optimization. *Mathematical programming*, 45(1-3):503–528, 1989.
- Liu, Q. and Wang, D. Stein variational gradient descent: A general purpose Bayesian inference algorithm. In Advances in Neural Information Processing Systems, pp. 2378–2386, 2016.
- Liu, Q., Lee, J., and Jordan, M. A kernelized Stein discrepancy for goodness-of-fit tests. In *International Conference on Machine Learning (ICML)*, pp. 276–284, 2016.
- Mroueh, Y., Sercu, T., and Raj, A. Sobolev descent. In International Conference on Artificial Intelligence and Statistics (AISTATS), pp. 2976–2985, 2019.

- Otto, F. The geometry of dissipative evolution equations: the porous medium equation. *Communications in Partial Differential Equations*, 26(1-2):101–174, 2001.
- Otto, F. and Villani, C. Generalization of an inequality by Talagrand and links with the logarithmic Sobolev inequality. *Journal of Functional Analysis*, 173(2):361–400, 2000.
- Pankrashkin, K. Introduction to the spectral theory. *Lecture* notes, Université Paris-Sud, 2014.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703*, 2019.
- Pavliotis, G. A. Stochastic processes and applications: diffusion processes, the Fokker-Planck and Langevin equations, volume 60. Springer, 2014.
- Renardy, M. and Rogers, R. C. *An introduction to partial differential equations*, volume 13. Springer Science & Business Media, 2006.
- Riabiz, M., Chen, W., Cockayne, J., Swietach, P., Niederer, S. A., Mackey, L., Oates, C., et al. Optimal thinning of MCMC output. arXiv preprint arXiv:2005.03952, 2020.
- Saitoh, S. and Sawano, Y. *Theory of Reproducing Kernels* and Applications. Springer Singapore, 2016.
- Salim, A., Korba, A., and Luise, G. Wasserstein proximal gradient. In *Advances in Neural Information Processing Systems*, 2020.
- Simon-Gabriel, C. J. Distribution-Dissimilarities in Machine Learning. PhD thesis, Eberhard Karls Universität Tübingen, Germany, 2018.
- Sriperumbudur, B. K., Fukumizu, K., and Lanckriet, G. R. Universality, characteristic kernels and RKHS embedding of measures. *Journal of Machine Learning Research*, 12 (70):2389–2410, 2011.
- Steinwart, I. and Christmann, A. *Support vector machines*. Springer Science & Business Media, 2008.
- Villani, C. Optimal transportation, dissipative PDE's and functional inequalities. In *Optimal transportation and applications*, pp. 53–89. Springer, 2003.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A.,

Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.

- Wenliang, L. K. Blindness of score-based methods to isolated components and mixing proportions. arXiv preprint arXiv:2008.10087, 2020.
- Wibisono, A. Sampling as optimization in the space of measures: The Langevin dynamics as a composite optimization problem. *Conference on Learning Theory*, 2018.
- Xu, W. and Matsuda, T. A Stein goodness-of-fit test for directional distributions. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 108, pp. 320–330, 2020.

A. Background

A.1. The continuity equation

Let T > 0. Consider a weakly continuous family of probability measures on \mathbb{R}^d , $\mu : (0,T) \to \mathcal{P}_2(\mathbb{R}^d)$, $t \mapsto \mu_t$. It satisfies a continuity equation (Ambrosio et al., 2008, Section 8.1) if there exists $(v_t)_{t \in (0,T)}$ such that $v_t \in L^2(\mu_t)$ and :

$$\frac{\partial \mu_t}{\partial t} + \operatorname{div}(\mu_t v_t) = 0 \qquad \text{in } \mathbb{R}^d \times (0, T)$$
(19)

holds in the sense of distributions, i.e. for any $\phi \in C_c^{\infty}(\mathbb{R}^d)$, the identity

$$\frac{\mathrm{d}}{\mathrm{d}t} \int_{\mathbb{R}^d} \phi(x) d\mu_t(x) = \int_{\mathbb{R}^d} \langle \nabla \phi(x), v_t(x) \rangle d\mu_t(x)$$
(20)

holds for any $t \in (0, T)$. By an integration by parts, the r.h.s. of the previous identity can be written:

$$\int_{\mathbb{R}^d} \langle \nabla \phi(x), v_t(x) \rangle d\mu_t(x) = -\int_{\mathbb{R}^d} \phi(x) \operatorname{div}(\mu_t(x)v_t(x)) \mathrm{d}x.$$
(21)

Hence, the identity (20) can be rewritten as

$$\int_{\mathbb{R}^d} \phi(x) \frac{\partial \mu_t(x)}{\partial t} dx + \int_{\mathbb{R}^d} \phi(x) \operatorname{div}(\mu_t(x)v_t(x)) dx = 0.$$
(22)

This equation expresses the law of conservation of mass for any volume of a moving fluid. Assume d = 3 and denote by (v_t^x, v_t^y, v_t^z) the projections of the vector field on the axes $(\overrightarrow{x}, \overrightarrow{y}, \overrightarrow{z})$. Then, the continuity equation has the form

$$\frac{\partial \mu_t}{\partial t} + \operatorname{div}(\mu_t v_t) = \frac{\partial \mu_t}{\partial t} + \frac{\partial (\mu_t v_t^x)}{\partial x} + \frac{\partial (\mu_t v_t^y)}{\partial y} + \frac{\partial (\mu_t v_t^z)}{\partial z} = 0.$$
(23)

A.2. Differentiability and convexity on the Wasserstein space

Let $(\mu, \nu) \in \mathcal{P}_2(\mathbb{R}^d)$. The Wasserstein 2 distance is defined as :

$$W_2^2(\mu,\nu) = \inf_{q \in \mathcal{Q}(\mu,\nu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^2 dq(x,y)$$
(24)

where $\mathcal{Q}(\mu, \nu)$ is the set of couplings between μ and ν , *i.e.* the set of nonnegative measures q over $\mathbb{R}^d \times \mathbb{R}^d$ such that $P_{1\#}q = \mu$ (resp. $P_{2\#}q = \nu$) where $P_1 : (x, y) \mapsto x$ (resp. $P_2 : (x, y) \mapsto y$).

The Wasserstein space $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ is not a Riemannian manifold but it can be equipped with a Riemannian structure and interpretation (Otto, 2001). In this geometric interpretation, the tangent space to $\mathcal{P}_2(\mathbb{R}^d)$ at μ is included in $L_2(\mu)$, and is equipped with the scalar product, defined for $f, g \in L_2(\mu)$ by:

$$\langle f,g\rangle_{L_2(\mu)} = \int_{\mathbb{R}^d} f(x)g(x)\mathrm{d}\mu(x).$$
(25)

Let $\mathcal{F} : \mathcal{P}_2(\mathbb{R}^d) \to \mathbb{R}$ be a functional on the Wasserstein space. We clarify in this section the notions of differentiability of \mathcal{F} that we consider in this setting. The notion of Fréchet subdifferentiability and its properties have been extended to the Wasserstein framework in (Ambrosio et al., 2008, Chapter 10). We first recall that if it exists, the *first variation of* \mathcal{F} *evaluated at* $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ is the unique function $\frac{\partial \mathcal{F}(\mu)}{\partial \mu} : \mathbb{R}^d \to \mathbb{R}$ s.t.

$$\lim_{\epsilon \to 0} \frac{1}{\epsilon} (\mathcal{F}(\mu + \epsilon \xi) - \mathcal{F}(\mu)) = \int_{\mathbb{R}^d} \frac{\partial \mathcal{F}(\mu)}{\partial \mu} (x) d\xi(x)$$
(26)

for all $\xi = \nu - \mu$, where $\nu \in \mathcal{P}_2(\mathbb{R}^d)$. Under mild regularity assumptions, the W_2 gradient of \mathcal{F} corresponds to the gradient of the first variation of \mathcal{F} , as stated below.

Definition 4. (Ambrosio et al., 2008, Lemma 10.4.1). Let $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, absolutely continuous with respect to the Lebesgue measure, with density in $C^1(\mathbb{R}^d)$ and such that $\mathcal{F}(\mu) < \infty$. The *subdifferential* of \mathcal{F} at μ is the map $\nabla_{W_2} \mathcal{F}(\mu)$ defined by:

$$\nabla_{W_2} \mathcal{F}(\mu)(x) = \nabla \frac{\partial \mathcal{F}(\mu)}{\partial \mu}(x) \text{ for } \mu\text{-a.e. } x \in \mathbb{R}^d,$$
(27)

and for every vector field $\xi \in C_c^{\infty}(\mathbb{R}^d; \mathbb{R}^d)$,

$$\int_{\mathbb{R}^d} \langle \nabla_{W_2} \mathcal{F}(\mu)(x), \xi(x) \rangle \mathrm{d}\mu(x) = -\int_{\mathbb{R}^d} \frac{\partial \mathcal{F}(\mu)}{\partial \mu}(x) \operatorname{div}(\mu(x)\xi(x)) \mathrm{d}x.$$
(28)

Moreover, $\nabla_{W_2} \mathcal{F}(\mu)$ belongs to the tangent space of $\mathcal{P}_2(\mathbb{R}^d)$ at μ , which is included in $L^2(\mu)$.

A.3. Cauchy-Lipschitz assumptions for the existence and uniqueness of the Wasserstein gradient flow

Let T > 0, and denote by \mathcal{L}^1 the standard one-dimensional Lebesgue measure on [0, T] and L^1 the space of measurable and integrable functions w.r.t. Lebesgue measure. The Cauchy-Lipschitz assumptions below for existence and uniqueness of the flow on [0, T] are adapted to our flows from their general differential inclusion version by Bonnet & Frankowska (2021). They hold for an initial distribution $\mu_0 \in \mathcal{P}_c(\mathbb{R}^d)$, the space of probability measures with compact support. If v only depends on time and not on μ_t , they then write as

- (**B**₁) (Carathéodory vector fields) $v : [0,T] \times \mathbb{R}^d \to \mathbb{R}^d$ is such that $t \mapsto v(t,x)$ is \mathcal{L}^1 -measurable for all $x \in \mathbb{R}^d$ and $x \mapsto v(t,x)$ is continuous for \mathcal{L}^1 -almost every $t \in [0,T]$.
- (**B**₂) (Sublinear growth) There exists a map $m(\cdot) \in L^1([0,T]; \mathbb{R}_+)$ such that $|v(t,x)| \leq m(t)(1+|x|)$ for \mathcal{L}^1 -almost every $t \in [0,T]$ and all $x \in \mathbb{R}^d$.
- (**B**₃) (Lipschitz vector field) For any compact set $\mathcal{K} \subset \mathbb{R}^d$, there exists a map $l_{\mathcal{K}}(\cdot) \in L^1([0,T],\mathbb{R}_+)$ such that $\operatorname{Lip}(v(t,\cdot);\mathcal{K}) \leq l_{\mathcal{K}}(t)$ for \mathcal{L}^1 -almost every $t \in [0,T]$.

For R > 0, we denote by $\mathcal{K} := B(0, R)$ the closed ball of radius R in \mathbb{R}^d ; and for any function f, we denote by $||f(\cdot)||_{\infty,\mathcal{K}} := \sup_{x \in \mathcal{K}} |f(x)|$ its supremum over \mathcal{K} and by $\operatorname{Lip}(f(\cdot); \mathcal{K})$ the Lipschitz constant of the restriction of f on \mathcal{K} . If v only depends on μ_t , then the following assumptions should hold for every R > 0:

- (**C**₁) For any $\mu \in \mathcal{P}_c(\mathbb{R}^d)$, $v_{|\mathcal{K}}(\mu)(\cdot) \in C^0(\mathcal{K}; \mathbb{R}^d)$.
- (C₂) There exists m > 0 such that for any $\mu \in \mathcal{P}_c(\mathbb{R}^d)$, for all $y \in \mathbb{R}^d$, we have $\|v(\mu)(y)\| \le m \left(1 + \|y\| + \int \|x\| d\mu(x)\right)$
- (C₃) There exists $l_{\mathcal{K}} > 0$ such that for any $\mu \in \mathcal{P}_c(\mathcal{K})$, we have $\operatorname{Lip}(v(\mu)(\cdot); \mathcal{K}) \leq l_{\mathcal{K}}$,
- (C₄) There exists $L_{\mathcal{K}} > 0$ such that for any $\mu, \nu \in \mathcal{P}_c(\mathcal{K})$, we have $\|v_{|\mathcal{K}}(\mu)(\cdot) v_{|\mathcal{K}}(\nu)(\cdot)\|_{\infty,\mathcal{K}} \le L_{\mathcal{K}}W_2(\mu, v)$.

Relation with Assumption (A₁). Consider $\mu \in \mathcal{P}_c(\mathbb{R}^d)$ and take R > 0 such that $\operatorname{supp}(\mu) \subset \mathcal{K} := B(0, R)$. Our Assumption (A₁) implies that $y \mapsto v_{\mu}(y) = \int \nabla_2 k_{\pi}(x, y) d\mu(x)$ is Lipschitz with constant $l_{\mathcal{K}} = \sup_{y \in \mathcal{K}} L(y)$, so Assumptions (C₁) and (C₃) hold. Assumption (C₂) corresponds to Assumption (A₂). Finally, for $\nu \in \mathcal{P}_c(\mathcal{K})$, since $\nabla_2 k_{\pi}(\cdot, y)$ is $l_{\mathcal{K}}$ -Lipschitz for $y \in \mathcal{K}$, we have that $x \mapsto ||\nabla_2 k_{\pi}(x, y)||$ is also $l_{\mathcal{K}}$ -Lipschitz, hence

$$\|v_{\mu}(y) - v_{\nu}(y)\| \le \sup\left\{\int_{M} f(x) \mathrm{d}(\mu - \nu)(x) \mid \operatorname{Lipschitz} f : \mathcal{K} \to \mathbb{R}, \operatorname{Lip}(f) \le l_{\mathcal{K}}\right\} = l_{\mathcal{K}} W_{1}(\mu, \nu) \le l_{\mathcal{K}} W_{2}(\mu, \nu),$$

where the last inequality between 1-Wasserstein distance W_1 and W_2 is a consequence of Jensen's inequality, so Assumption (C_4) is satisfied.

For other kernel-based updates as presented in Appendix A.4, the kernel is at most multiplied once by the score s, so for uniformly Lipschitz s and kernels with bounded derivatives, the Lipschitz constant is uniform for $v(\mu)(\cdot)$. This was for instance assumed in Arbel et al. (2019). However for updates involving the Stein kernel such as in KSD Descent, the analysis is more intricate. As a matter of fact, when π is a standard Gaussian distribution, we have s(x) = x. Hence the Stein kernel verifies, for a smooth translation-invariant kernel with k(x, x) = 1, that $k_{\pi}(x, x) = C + ||x||^2$ for $x \in \mathbb{R}^d$, with *C* a constant determined by the last term in k_{π} . So, for a Gaussian π , the Lipschitz constant of $v(\delta_y)(\cdot)$ is larger than 2||y||, explaining the reason why we only required for a *y*-dependent Lipschitz constant in Assumption (**A**₁) and did not assume it to be uniform. This is related to the fact that $k_{\pi}(x, x)$ is in general unbounded for the Stein kernel, which precludes most of the classical assumptions made in the kernel literature.

Comment on Assumption (C₂). The main role of Assumption (C₂) is to impede the well-known phenomenon of finite-time explosion of a trajectory as for the system $z'(t) = z(t)^2$ in \mathbb{R} . It could in principle be replaced by any assumption preventing particles from escaping to infinity in finite time. One such lighter assumption would write as follows:

"For any r > 0, T > 0 and $\mu_0 \in \mathcal{P}_c(\mathbb{R}^d)$ with $\operatorname{supp}(\mu_0) \subset rB(0,1)$ there exists $R(r,T) < \infty$ such that $\operatorname{supp}(\mu_t) \subset R(r,T)B(0,1)$ for $(\mu_t)_{t \in [0,T]}$ solving (7)."

A.4. Descent updates of the algorithms presented in Section 3.3

From a computational viewpoint, both the SVGD, LAWGD and MMD flow algorithms, presented in Section 3.3, only propose gradient descent schemes with the following respective updates $x_n^i \leftarrow x_n^i - \gamma D^i$ with

$$\begin{aligned} \mathcal{D}_{SVGD}^{i} &= \frac{1}{N} \sum_{j=1}^{N} \left[k(x_{n}^{j}, x_{n}^{i}) s(x_{n}^{j}) + \nabla_{1} k(x_{n}^{j}, x_{n}^{i}) \right], \\ \mathcal{D}_{LAWGD}^{i} &= \frac{1}{N} \sum_{j=1}^{N} \nabla_{2} k_{\mathcal{L}_{\pi}}(x_{n}^{j}, x_{n}^{i}), \\ \mathcal{D}_{MMD-GD}^{i} &= \frac{1}{N} \sum_{j=1}^{N} \left[\nabla_{2} k(x_{n}^{j}, x_{n}^{i}) - \nabla_{2} k(y^{j}, x_{n}^{i}) \right]. \end{aligned}$$

where the MMD-GD update requires extra samples $(y^j)_{j=1}^N \sim \pi$, since it is sample-based rather than score-based. Notice that all these updates have the same iteration complexity of order $\mathcal{O}(N^2)$. Intriguingly LAWGD has the same update rule as ours (12) but for their kernel $k_{\mathcal{L}_{\pi}}$ which, as discussed in Section 3.3, does not incorporate an off-the-shelf kernel k, unlike the Stein kernel k_{π} (3).

A.5. Background on diffusion operator \mathcal{L}_{π}

In this section, for the convenience of the reader who is not familiar with the spectral theory of diffusion operators, we formulate Lemma 12 which provides a formal construction of the diffusion operator $\mathcal{L}_{\pi} = -\langle \nabla \log \pi, \nabla \rangle - \Delta$ on the space $L_2(\pi)$ and gathers a number of technical facts about it. Those facts form a background for the proofs of results related to lack of exponential convergence near equilibrium of KSD flow, which are presented in Appendix B. We provide the proof of lemma 12 in Appendix B.11.

Lemma 12. Let $\pi \propto e^{-V}$ be a probability measure on \mathbb{R}^d and assume that $V : \mathbb{R}^d \to \mathbb{R}$ is in $C^1(\mathbb{R}^d)$. Let $\hat{\mathcal{L}}_{\pi} = \langle \nabla V, \nabla \rangle - \Delta$ on $C_c^{\infty}(\mathbb{R}^d)$. This operator can be extended to a positive self-adjoint operator on $L^2(\pi)$ with dense domain $\mathcal{D}(\mathcal{L}_{\pi}) \subset L^2(\pi)$ which we denote by \mathcal{L}_{π} . Moreover, $C_c^{\infty}(\mathbb{R}^d)$ is dense in $\mathcal{D}(\mathcal{L}_{\pi})$, for the norm:

$$\|\phi\|_{\mathcal{L}_{\pi}} = \left(\langle\phi, \mathcal{L}_{\pi}\phi\rangle + \|\phi\|_{L_{2}(\pi)}^{2}\right)^{1/2}.$$
(29)

From that it follows, that $\mathcal{D}(\mathcal{L}_{\pi})$ is the subset of the weighted Sobolev space $W_0^{1,2}(\pi)^4$ (that is, the closure of $C_c^{\infty}(\mathbb{R}^d)$ in $W^{1,2}(\pi)$) and for all $f \in \mathcal{D}(\mathcal{L}_{\pi})$ we have

$$\|\nabla f\|_{L^2(\pi)}^2 = \langle f, \mathcal{L}_\pi f \rangle_{L^2(\pi)},\tag{30}$$

where ∇f is the weak derivative of f. This implies that the kernel of \mathcal{L}_{π} consists of π -almost everywhere constant functions. Furthermore, for any $f \in \mathcal{D}(\mathcal{L}_{\pi})$ we can find a sequence $\phi_n \in C_c^{\infty}(\mathbb{R}^d)$, such that $\lim_{n\to\infty} \|\phi_n - f\|_{L_2(\pi)} = 0$ and $\lim_{n\to\infty} \|\mathcal{L}_{\pi}\phi_n - \mathcal{L}_{\pi}f\|_{L_2(\pi)} = 0$.

⁴Note that the meaning of 0 in the notation $W_0^{1,2}(\pi)$ differs from the meaning of 0 in $L_0^2(\pi) = \{\phi \in L^2(\pi), \int \phi d\pi = 0\}$.

A.6. A Descent lemma for KSD Descent

A descent lemma for (10) is a proposition stating that \mathcal{F} decreases at each iteration of the time-discretized flow. It should hold for both continuous or discrete initializations. Arbel et al. (2019) proved a similar result for $\mathcal{F} = \frac{1}{2} \text{ MMD}^2$ which we recall below for completeness.

Proposition 13. (Arbel et al., 2019, Proposition 4) Let μ_n be defined by (10) for $\mathcal{F} = \frac{1}{2} \text{ MMD}^2$ and assume that $k \in C^{1,1}(\mathbb{R}^d \times \mathbb{R}^d)$ with *l*-Lipschitz gradient: $\|\nabla k(x, x') - \nabla k(y, y')\| \le l(\|x - y\| + \|x' - y'\|)$ for all $x, x', y, y' \in \mathbb{R}^d$. Then, for $\gamma \le \frac{2}{l}$, the sequence $(\mathcal{F}(\mu_n))_{n \ge 0}$ is decreasing and for any $n \ge 0$:

$$\mathcal{F}(\mu_{n+1}) - \mathcal{F}(\mu_n) \le -\gamma \left(1 - \frac{3\gamma l}{2}\right) \|\nabla_{W_2} \mathcal{F}(\mu_n)\|_{L^2(\mu_n)}^2.$$
(31)

Although KSD Descent is a special case of MMD Descent, assuming a uniform Lipschitz constant for k_{π} is excessive, even for targets as simple as a single Gaussian distribution, as discussed in Appendix A.3. In contrast, we consider the weaker Assumption (**A**₁) on the Stein kernel: there exists a map $L(\cdot) \in C^0(\mathbb{R}^d; \mathbb{R}_+)$ such that, for any $y \in \mathbb{R}^d$, the maps $x \mapsto \nabla_1 k_{\pi}(x, y)$ and $x \mapsto \nabla_1 k_{\pi}(y, x)$ are L(y)-Lipschitz.

We shall assume below the convexity of $L(\cdot)$ and the boundedness of its L^2 -norm along the trajectory, i.e. $(||L||_{L^2(\mu_n)})_{n\geq 0}$ is bounded. An example of more explicit, though tighter, assumptions are that L satisfies a subpolynomial growth at infinity, i.e. there exists R > 0, c > 0 and $m \in \mathbb{N}$ such that $L(x) \leq \tilde{L}(x) = c||x||^m$ for $x \in \mathbb{R}^d$ with $||x|| \geq R$, combined with boundedness of the corresponding moment along the flow, i.e. there exists M_m such that $\mathbb{E}_{x\sim\mu_n}||x||^{2m} \leq M_m$. In this case, instead of the continuous L, one can consider the convex and continuous function defined by $\tilde{L}(x) = c||x||^m + \sup_{y\in B(0,R)} L(y)$.

Proposition 14. Suppose Assumption (A₁) holds and that $\mu_0 \in \mathcal{P}_c(\mathbb{R}^d)$. Assume additionally that $L(\cdot)$ is convex, belongs to $L^2(\mu_n)$, and that $\|L(\cdot)\|_{L^2(\mu_n)} \leq M$ for any $n \geq 0$, where μ_n is defined by (10). Then, for any $\gamma \leq \frac{1}{M}$:

$$\mathcal{F}(\mu_{n+1}) - \mathcal{F}(\mu_n) \le -\gamma \left(1 - \gamma M\right) \|\nabla_{W_2} \mathcal{F}(\mu_n)\|_{L^2(\mu_n)}^2 \le 0.$$
(32)

See the proof in Appendix B.12.

B. Proofs

We shall often use that, for smooth and symmetric k, $\nabla_2 k(x, y) = \nabla_1 k(y, x)$, for any $x, y \in \mathbb{R}^d$. Moreover if k is translation-invariant, $\nabla_1 k(x, y) = -\nabla_2 k(x, y)$.

Remark 2. For any $k \in C^2(\mathbb{R}^d \times \mathbb{R}^d, \mathbb{R})$ and π such that $s \in C^1(\mathbb{R}^d)$, k_{π} and its gradient are written

$$k_{\pi}(x,y) = s(x)^{\top} s(y) k(x,y) + s(x)^{\top} \nabla_2 k(x,y) + s(y)^{\top} \nabla_1 k(x,y) + \boldsymbol{\nabla} \cdot_1 \nabla_2 k(x,y).$$
(33)
$$\nabla_2 k_{\pi}(x,y) = s(x)^{\top} s(y) \nabla_2 k(x,y) + J s(y)^{\top} s(x) k(x,y) + H_2 k(x,y) s(x)$$

$$+ Js(y)^{\top} \nabla_1 k(x, y) + \nabla_2 (\boldsymbol{\nabla} \cdot_1 \nabla_2 k(x, y)).$$
(34)

If k is translation-invariant, since $\nabla_1 k(x, y) = -\nabla_2 k(x, y)$, we have:

$$k_{\pi}(x,y) = s(x)^{\top} s(y) k(x,y) + (s(x) - s(y))^{\top} \nabla_2 k(x,y) + \nabla \cdot_1 \nabla_2 k(x,y),$$
(35)

$$\nabla_{2}k_{\pi}(x,y) = s(x)^{\top}s(y)\nabla_{2}k(x,y) + Js(y)^{\top}s(x)k(x,y) + H_{1}k(x,y)s(x) - Js(y)^{\top}\nabla_{2}k(x,y) - \nabla_{2}(\nabla \cdot_{2} \nabla_{2}k(x,y)).$$
(36)

Notice that, in this case, odd derivatives of k vanish for x = y.

B.1. Proof of Lemma 1

To prove that Assumption (A₁) holds, we want to show that $x \mapsto \nabla_1 k_\pi(x, y)$ and $x \mapsto \nabla_2 k_\pi(x, y)$ are L(y)-Lipschitz for some $L(\cdot) \in C^0(\mathbb{R}^d, \mathbb{R}_+)$ that is μ -integrable for any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, so L should have a quadratic growth at most. We will leverage the regularity of k and s to write this problem as that of upper-bounding over \mathbb{R}^d the gradients of these functions (which are related to the Hessian of k_{π}). To prove that Assumption (A₄) holds, we want on the other hand to identify an integrable function that upper bounds $||H_1k_{\pi}(x, y)||_{op}$.

We denote by Js(x) the Jacobian of s at x, which is also the Hessian of $\log(\pi)$, and by Hs(x) the Hessian of s at x, which is also the tensor of third derivatives of $\log(\pi)$, i.e. $Hs(x)_{ijk} = \frac{\partial^3 \log(\pi)}{\partial x_i \partial x_j \partial x_k}(x)$.

We compute the gradient of the first term in (33):

$$\nabla_x \left(s(x)^\top s(y) k(x,y) \right) = J s(x)^\top s(y) k(x,y) + s(x)^\top s(y) \nabla_1 k(x,y).$$
(37)

Differentiating a second time, we get that

$$H_1(s(x)^{\top}s(y)k(x,y)) = H_s(x)s(y)k(x,y) + \nabla_1 k(x,y)s(y)^{\top} J_s(x) + (J_s(x)^{\top}s(y)) \nabla_1 k(x,y)^{\top} + s(x)^{\top}s(y)H_1k(x,y)$$

where Hs(x)s(y) is a $d \times d$ matrix of entries $[Hs(x)s(y)]_{jk} = \sum_{i=1}^{n} Hs(x)_{ijk}s(y)_i$.

Then, the gradient of the second term in (33) is

$$\nabla_x \left(s(x)^\top \nabla_2 k(x, y) \right) = J s(x)^\top \nabla_2 k(x, y) + \nabla_{1,2} k(x, y) s(x), \tag{38}$$

where $\nabla_{1,2}k(x,y) = [\partial_{x_i}\partial_{y_j}k(x,y)]_{i,j}$, and its Hessian is given by:

$$H_1\left(s(x)^{\top} \nabla_2 k(x,y)\right) = H_s(x) \nabla_2 k(x,y) + J_s(x)^{\top} \nabla_{1,2} k(x,y) + \nabla_{1,2} k(x,y) J_s(x) + \nabla_{1,1,2} k(x,y) s(x), \quad (39)$$

where $\nabla_{1,1,2}k(x,y) = [\partial_{x_i,x_j}\partial_{y_l}k(x,y)]_{i,j,l}$ is a tensor of third derivatives of k.

The Hessian of the last two terms in (33) is straightforward to compute. Hence, collecting all the terms together, we derive that

$$H_{1}k_{\pi}(x,y) = Hs(x)s(y)k(x,y) + \nabla_{1}k(x,y)s(y)^{\top}Js(x) + (Js(x)^{\top}s(y))\nabla_{1}k(x,y)^{\top} + s(x)^{\top}s(y)H_{1}k(x,y)$$
(40)
$$Hs(x)\nabla_{2}k(x,y) + Js(x)^{\top}\nabla_{1,2}k(x,y) + \nabla_{1,2}k(x,y)Js(x) + \nabla_{1,1,2}k(x,y)s(x)$$
$$+ \nabla_{1,1,1}k(x,y)s(y) + H_{1}\operatorname{Tr}(\boldsymbol{\nabla} \cdot_{2} \nabla_{1}k(x,y)).$$

In this expression, the only problematic terms to upper bound H_1k_{π} uniformly w.r.t. x are the ones where s(x) appears (since Js(x) and Hs(x) are bounded by assumption). However, since s is C_1 -Lipschitz, we have that $||s(x)|| \le ||s(y)|| + C_1 ||x-y||$. We then use that the kernel, its derivatives, and the derivatives up to order 3 multiplied by ||x - y|| are bounded for all x and y by a given $B \ge 0$. Hence, for instance,

$$\|s(x)^{\top}s(y)H_1k(x,y)\| \le \|s(y)\| \cdot (\|s(y)\| + C_1\|x - y\|) \cdot \|H_1k(x,y)\| \le B\|s(y)\|^2 + C_1B\|s(y)\|.$$

As $||Js(x)||_{op} \leq C_1$ and $||Hs(x)||_{op} \leq C_2$ for all $x \in \mathbb{R}^d$, we have that $||s(x)|| \leq C_1 ||x|| + ||s(0)||$ and a similar computation gives that

$$||H_1k_{\pi}(x,y)||_{op} \le C_2 B||s(y)|| + 2C_1 B||s(y)|| + B||s(y)||^2 + C_1 B||s(y)|| + C_2 B + 2C_1 B + B||s(y)|| + C_1 B + B||s(y)|| + B$$
(41)

$$\leq \tilde{C}_2 \|y\|^2 + \tilde{C}_0,\tag{42}$$

where $\tilde{C}_2 \ge 0$ and $\tilde{C}_0 \ge 0$ only depend on B, C_2, C_1 and ||s(0)||. Consequently, the function $(x, y) \mapsto ||H_1k_{\pi}(x, y)||_{op}$ is integrable for any $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$, so Assumption (A₃) holds.⁵

Similarly, one can show based on (34) that

$$\nabla_{1,2}k_{\pi}(x,y) = (Js(x)^{\top}s(y))\nabla_{2}k(x,y)^{\top} + s(x)^{\top}s(y)\nabla_{1,2}k(x,y) + Js(y)^{\top}Js(x)k(x,y) + \nabla_{1}k(x,y)s(x)^{\top}Js(y) + H_{2}k(x,y)Js(x) + \nabla_{1,2,2}k(x,y)s(x) + Js(y)^{\top}H_{1}k(x,y) + \nabla_{1,2}\operatorname{Tr}(\nabla \cdot_{2} \nabla_{1}k(x,y)).$$

⁵We could do without the assumption of Lipschitzianity of s to prove that Assumption (**A**₃) holds. We could also merely require that the derivatives of k are bounded, without considering the multiplication by ||x - y||. As a matter of fact, assuming that $||Hs(x)||_{op} \leq C_2$ for all $x \in \mathbb{R}^d$ implies that $||Js(x)||_{op} \leq C_2 ||x|| + C_1$ and $||s(x)|| \leq C_2 ||x||^2 + C_0$ for some $C_0 \geq 0$ and $C_1 \geq 0$. Hence (40) yields an upper bound that is quadratic in ||x|| and ||y||, so $\mu \otimes \nu$ -integrable as in Assumption (**A**₃).

Using the same inequalities that led to (42), we can find two constants $\hat{C}_2 \ge 0$ and $\hat{C}_0 \ge 0$ such that

$$|\nabla_{1,2}k_{\pi}(x,y)||_{op} \leq \hat{C}_{2}||y||^{2} + \hat{C}_{0}$$

Setting $L(y) = \max(\tilde{C}_2, \hat{C}_2) \|y\|^2 + \max(\tilde{C}_0, \hat{C}_0)$, the functions $x \mapsto \nabla_1 k_\pi(x, y)$ and $x \mapsto \nabla_2 k_\pi(x, y)$ are L(y)-Lipschitz since L(y) is an upper bound of the norm of their Jacobians. Furthermore, $L(\cdot) \in C^0(\mathbb{R}^d, \mathbb{R}_+)$, and is μ -integrable for any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, so Assumption (A₁) holds.

To prove that Assumption (\mathbf{A}_4) holds, notice that, based on (33),

$$|k_{\pi}(x,x)|_{op} = \left| \|s(x)\|^{2} k(x,x) + s(x)^{\top} \nabla_{2} k(x,x) + s(x)^{\top} \nabla_{1} k(x,x) + \nabla \cdot_{1} \nabla_{2} k(x,x) \right|$$

$$\leq B \left(\|s(x)\|^{2} + 2\|s(x)\| + 1 \right).$$

Since, for all $x \in \mathbb{R}^d$, $||s(x)|| \le C_1 ||x|| + ||s(0)||$, the function $x \mapsto k_\pi(x, x)$ is bounded by a quadratic function and is thus μ -integrable for any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$. This shows that Assumption (A₄) holds.

We now prove that Assumption (A₂) is also satisfied if there exists M > 0 and M_0 such that, for all $x \in \mathbb{R}^d$, $||s(x)|| \le M\sqrt{||x||} + M_0$. Then, based on (34) and as $\sqrt{||x||} \le ||x|| + 1$,

$$\begin{aligned} \|\nabla_2 k_{\pi}(x,y)\| &\leq \|s(x)\| \cdot \|s(y)\| \cdot \|\nabla_2 k(x,y)\| + BC_1 \|s(x)\| + B\|s(x)\| + BC_1 + B \\ &\leq (\|s(y)\| + C_1 \|x - y\|) \cdot \|s(y)\| \cdot \|\nabla_2 k(x,y)\| + (BC_1 + B) \cdot (M\sqrt{\|x\|} + M_0 + 1) \\ &\leq B\|s(y)\|^2 + BC_1 \|s(y)\| + (BC_1 + B) \cdot (M\|x\| + M + M_0 + 1) \\ &\leq B(M\sqrt{\|y\|} + M_0)^2 + BC_1 (M\|y\| + M + M_0) + (BC_1 + B) \cdot (M\|x\| + M + M_0 + 1) \\ &\leq m(1 + \|y\| + \|x\|), \end{aligned}$$

for some m > 0 depending on B, M, C_1 and M_0 . Integrating over μ for any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ gives that $\|\int \nabla_2 k_{\pi}(x, y) d\mu(x)\| \le m (1 + \|y\| + \int \|x\| d\mu(x))$, so Assumption (A₂) holds, which concludes the proof.

B.2. Proof of Proposition 2

Lemma 15. Suppose Assumption (A₁) holds. The first variation of \mathcal{F} evaluated at $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ is then the function defined for any $y \in \mathbb{R}^d$ by:

$$\frac{\partial \mathcal{F}(\mu)}{\partial \mu}(y) = \int_{\mathbb{R}^d} k_{\pi}(x, y) \mathrm{d}\mu(x).$$

Proof. Let $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$, and $\xi = \mu - \nu$. Using the symmetry of k_{π} , we get

$$\frac{1}{\epsilon} \left[\mathcal{F}(\mu + \epsilon\xi) - \mathcal{F}(\mu) \right] = \frac{1}{2\epsilon} \left[\iint_X k_\pi(x, y) d(\mu + \epsilon\xi)(x) d(\mu + \epsilon\xi)(y) - \iint_X k_\pi(x, y) d\mu(x) d\mu(y) \right]$$
$$= \iint_{\mathbb{R}^d} k_\pi(x, y) d\mu(x) d\xi(y) + \frac{\epsilon}{2} \iint_X k_\pi(x, y) d\xi(x) d\xi(y)$$

Hence,

$$\lim_{\epsilon \to 0} \frac{1}{\epsilon} (\mathcal{F}(\mu + \epsilon\xi) - \mathcal{F}(\mu)) = \iint_{\mathbb{R}^d} k_{\pi}(x, y) d\mu(x) d\xi(y).$$

Assume μ satisfies the assumptions of Definition 4. To obtain the expression for the W_2 gradient of \mathcal{F} , we first need to exchange the integration and the gradient with respect to y. Since $y \mapsto \nabla_2 k_{\pi}(x, y)$ is L(x)-Lipschitz, we have for all $v \in \mathbb{R}^d, 0 < \epsilon \leq 1$:

$$\left|\frac{k_{\pi}(x, y + \epsilon v) - k_{\pi}(x, y)}{\epsilon} - \langle \nabla_2 k_{\pi}(x, y), v \rangle\right| \le \frac{L(x)\epsilon \|v\|_2^2}{2} \le \frac{L(x)\|v\|_2^2}{2}$$

where the right hand side is integrable. Therefore by the Lebesgue dominated convergence theorem, we have the following interchange of the gradient and the integral when computing the W_2 -gradient:

$$\nabla_{W_2} \mathcal{F}(\mu)(y) = \nabla \frac{\partial \mathcal{F}(\mu)}{\partial \mu}(y) = \int \nabla_2 k_\pi(x, y) \mathrm{d}\mu(x).$$
(43)

Then, using the continuity equation (7) and an integration by parts,⁶ the dissipation of \mathcal{F} along its W_2 gradient flow is obtained as follows:

$$\frac{\mathrm{d}\mathcal{F}(\mu_t)}{\mathrm{d}t} = \int \frac{\partial\mathcal{F}(\mu_t)}{\partial\mu_t} \frac{\partial\mu_t}{\partial t} = \int \frac{\partial\mathcal{F}(\mu_t)}{\partial\mu} \operatorname{div}\left(\mu_t \nabla \frac{\partial\mathcal{F}(\mu_t)}{\partial\mu}\right) = -\int \left\|\nabla \frac{\partial\mathcal{F}(\mu_t)}{\partial\mu}\right\|^2 \mathrm{d}\mu_t.$$

Plugging (43) in the r.h.s. of this equality leads to the final formula.

B.3. Proof of Proposition 3

We compute the second time derivative $\ddot{\mathcal{F}}(\rho_t)$ where ρ_t is a path from μ to $(I + \nabla \psi)_{\#}\mu$ given by: $\rho_t = (I + t\nabla \psi)_{\#}\mu$, for all $t \in [0, 1]$. By Lemma 22, $\ddot{\mathcal{F}}(\rho_t)$ is well-defined and $\ddot{\mathcal{F}}(\mu) = \ddot{\mathcal{F}}(\rho_t)|_{t=0}$ is given by:

$$\ddot{\mathcal{F}}(\mu) = \int \left[\nabla\psi^{\top}(z)\nabla_1\nabla_2k_{\pi}(z,z')\nabla\psi(z')\right]d\mu(z')d\mu(z) + \int \left[\nabla\psi^{\top}(z)H_1k_{\pi}(z,z')\nabla\psi(z)\right]d\mu(z')d\mu(z).$$
(44)

B.4. Proof of Corollary 4

Under Assumption (A_3) , by the dominated convergence theorem, we can exchange the order of the integral and derivative in the second term of (44):

$$\iint \left[\nabla\psi^{\top}(z)H_1k_{\pi}(z,z')\nabla\psi(z)\right]d\rho_t(z')d\rho_t(z) = \int \left[\nabla\psi^{\top}(z)H_1\left(\int k_{\pi}(z,z')d\rho_t(z')\right)\nabla\psi(z)\right]d\rho_t(z).$$
 (45)

The latter (45) vanishes when $\rho_t = \pi$, thanks to the property of the Stein kernel : $\int k_{\pi}(z, z') d\pi(z') = 0$.

Hence, by considering $\psi \in C_c^{\infty}(\mathbb{R}^d)$ and a path ρ_t from π to $(I + \nabla \psi)_{\#}\pi$ given by: $\rho_t = (I + t\nabla \psi)_{\#}\pi$, for all $t \in [0, 1]$.

$$\mathcal{F}(\pi) = \operatorname{Hess}_{\pi}(\psi, \psi),$$

where

$$\operatorname{Hess}_{\pi}(\psi,\psi) = \iint \left[\nabla \psi^{\top}(z) \nabla_{1} \nabla_{2} k_{\pi}(z,z') \nabla \psi(z') \right] \mathrm{d}\pi(z') \mathrm{d}\pi(z).$$

By an integration by parts⁷ w.r.t. x, we have, for $\mathcal{L}_{\pi} : f \mapsto -\Delta f + \langle -\nabla \log \pi, \nabla f \rangle$,

$$\begin{aligned} \operatorname{Hess}_{\pi}(\psi,\psi) &= \sum_{i,j=1}^{d} \iint \frac{\partial \psi(x)}{\partial x_{i}} \cdot \frac{\partial^{2} k_{\pi}(x,y)}{\partial x_{i} \partial y_{j}} \cdot \frac{\partial \psi(y)}{\partial y_{j}} \mathrm{d}\pi(x) \mathrm{d}\pi(y) \\ &= -\sum_{i,j=1}^{d} \iint \frac{\partial^{2} \psi(x)}{\partial x_{i}^{2}} \cdot \frac{\partial k_{\pi}(x,y)}{\partial y_{j}} \cdot \frac{\partial \psi(y)}{\partial y_{j}} \mathrm{d}\pi(x) \mathrm{d}\pi(y) - \sum_{i,j=1}^{d} \iint \frac{\partial \psi(x)}{\partial x_{i}} \cdot \frac{\partial \log \pi(x)}{\partial x_{i}} \cdot \frac{\partial k_{\pi}(x,y)}{\partial y_{j}} \cdot \frac{\partial \psi(y)}{\partial y_{j}} \mathrm{d}\pi(x) \mathrm{d}\pi(y) \\ &= \sum_{j=1}^{d} \iint \mathcal{L}_{\pi} \psi(x) \cdot \frac{\partial k_{\pi}(x,y)}{\partial y_{j}} \cdot \frac{\partial \psi(y)}{\partial y_{j}} \mathrm{d}\pi(x) \mathrm{d}\pi(y). \end{aligned}$$

We then repeat the same steps, performing an integration by parts w.r.t. y_i and using Lemma 21 in the last equality,

$$\operatorname{Hess}_{\pi}(\psi,\psi) = \int \mathcal{L}_{\pi}\psi(x)k_{\pi}(x,y)\mathcal{L}_{\pi}\psi(y)\mathrm{d}\pi(x)\mathrm{d}\pi(y) = \|S_{\pi,k_{\pi}}(\mathcal{L}_{\pi}\psi)\|_{\mathcal{H}_{k_{\pi}}}^{2}.$$

Notice that under Assumption (A₄), we have $\mathcal{H}_{k_{\pi}} \subset L^2(\pi)$, so $S_{\pi,k_{\pi}}$ and its adjoint are well-defined (see Section 2.3).

⁶By the regularity of μ , from the assumptions of Definition 4, there are no boundary terms in the integration.

⁷First, differentiate the product $\frac{\partial \psi(x)}{\partial x_i}\pi(x)$, then integrate $\frac{\partial^2 k_\pi(x,y)}{\partial x_i \partial y_j}$ w.r.t. x_i . Since ψ is compactly supported, it vanishes at infinity. Since π has a C^1 -density w.r.t. the Lebesgue measure, any integral over $d\pi$ bearing on the boundary of the support of π vanishes. We thus do not have any extra integral on the boundary when performing the integration by parts.

B.5. Proof of Proposition 5

Denote by $L_0^2(\pi)$ the closed subspace of $L^2(\pi)$ consisting of functions ϕ such that $\int \phi(x) d\pi(x) = 0$. For any $f \in \mathcal{D}(\mathcal{L}_{\pi})$, $\int \mathcal{L}_{\pi} f(x) d\pi(x) = 0$; hence $\operatorname{Im}(\mathcal{L}_{\pi})$ is a subset of $L_0^2(\pi)$, which is itself a subset of $L^2(\pi)$.

Let $T_{\pi,k_{\pi}} = S^*_{\pi,k_{\pi}} \circ S_{\pi,k_{\pi}}$, and assume that $V = -\log(\pi)$ is a $C^1(\mathbb{R}^d)$ function. We want to show that exponential decay near equilibrium (16) holds, if and only if $\mathcal{L}^{-1}_{\pi} : L^2_0(\pi) \to L^2_0(\pi)$, the inverse of $\mathcal{L}_{\pi}|_{L^2_0(\pi)}$, is a well-defined linear operator, and

$$\langle \phi, T_{\pi,k_{\pi}} \phi \rangle_{L_2(\pi)} \ge \lambda \langle \phi, \mathcal{L}_{\pi}^{-1} \phi \rangle_{L_2(\pi)}$$
(46)

holds for all $\phi \in L^2_0(\pi)$.

Proof. First, assume that the exponential convergence near equilibrium holds. We will show that \mathcal{L}_{π}^{-1} is a well-defined bounded linear operator on $L_0^2(\pi)$, and then that the inequality 46 holds. Let $\psi \in C_c^{\infty}(\mathbb{R}^d)$. By Corollary 4, the Hessian of \mathcal{F} at π can be written as

$$\operatorname{Hess}_{\pi}(\psi,\psi) = \langle S_{\pi,k_{\pi}}(\mathcal{L}_{\pi}\psi), S_{\pi,k_{\pi}}\mathcal{L}_{\pi}\psi \rangle_{\mathcal{H}_{k_{\pi}}} = \langle T_{\pi,k_{\pi}}(\mathcal{L}_{\pi}\psi), \mathcal{L}_{\pi}\psi \rangle_{L^{2}(\pi)}$$

where $T_{\pi,k_{\pi}} = S^*_{\pi,k_{\pi}} \circ S_{\pi,k_{\pi}}$. By Lemma 12, we have that $\|\nabla\psi\|^2_{L^2(\pi)} = \langle\psi, \mathcal{L}\psi\rangle_{L^2(\pi)}$, and exponential convergence near equilibrium can be written:

$$\langle T_{\pi,k_{\pi}}(\mathcal{L}_{\pi}\psi), \mathcal{L}_{\pi}\psi\rangle_{L^{2}(\pi)} \geq \lambda \langle \psi, \mathcal{L}_{\pi}\psi\rangle_{L^{2}(\pi)}.$$
(47)

Note, that by Assumption (A₄), $T_{k_{\pi},\pi}$ is a bounded operator (Steinwart & Christmann, 2008, Theorem 4.27), and it follows by an application of the Cauchy-Schwartz inequality that (47) implies:

$$\|\mathcal{L}_{\pi}\psi\|_{L^{2}(\pi)}^{2} \geq \frac{\lambda}{\|T_{k_{\pi},\pi}\|_{op}} \langle \psi, \mathcal{L}_{\pi}\psi \rangle_{L^{2}(\pi)},$$
(48)

where $\|\cdot\|_{op}$ denotes the operator norm. Now, let ψ be an arbitrary element of $\mathcal{D}(\mathcal{L}_{\pi})$, the domain of \mathcal{L}_{π} . By Lemma 12, there exists a sequence $(\psi_n)_{n=1}^{\infty} \subset C_c^{\infty}(\mathbb{R}^d)$ converging strongly to ψ , such that $\mathcal{L}_{\pi}\psi_n$ converges strongly to $\mathcal{L}_{\pi}\psi$ as well. Hence, (48) holds for all $\psi \in \mathcal{D}(\mathcal{L}_{\pi})$.

We will now show, that the spectrum of \mathcal{L}_{π} , $\sigma(\mathcal{L}_{\pi})$ is contained in $\{0\} \cup [\frac{\lambda}{\|T_{k_{\pi},\pi}\|_{op}},\infty)$. Suppose that there exists a $\sigma \in (0, \frac{\lambda}{\|T_{k_{\pi},\pi}\|_{op}}) \cap \sigma(\mathcal{L}_{\pi})$. If σ is in the point spectrum of \mathcal{L}_{π} , then by definition there would exist a vector $v \in \mathcal{D}(\mathcal{L}_{\pi})$ such that $\mathcal{L}_{\pi}v = \sigma v$, which would contradict the inequality (48). On the other hand, if σ is not in the point spectrum of \mathcal{L}_{π} , then by Weyl's criterion (Pankrashkin, 2014, Theorem 7.22) we can find an orthonormal sequence $(u_n)_{n=1}^{\infty} \in \mathcal{D}(\mathcal{L}_{\pi})$ such that $(\mathcal{L}_{\pi} - \sigma)u_n$ converges to 0 in $L_2(\pi)$. An obvious calculation shows that this would contradict (48). Hence, $\sigma(\mathcal{L}_{\pi}) \subset \{0\} \cup [\frac{\lambda}{\|T_{k_{\pi},\pi}\|},\infty)$.

We note that $L_0^2(\pi)$ is itself a Hilbert space with the inner product inherited from $L^2(\pi)$. The image of \mathcal{L}_{π} is contained in $L_0^2(\pi)$ (recall that for any $f \in \mathcal{D}(\mathcal{L}_{\pi})$, $\int \mathcal{L}_{\pi}f(x)d\pi(x) = 0$), and it is dense in $L_0^2(\pi)$ since $\operatorname{Im}(T) = (\operatorname{Ker}(T))^{\perp}$ for self-adjoint operators T (Brezis, 2010, Corollary 2.18). Furthermore, since $\sigma(\mathcal{L}_{\pi}) \subset \{0\} \cup [\frac{\lambda}{\|T_{k\pi,\pi}\|_{op}}, \infty)$ and the kernel of \mathcal{L}_{π} consists of constants (Appendix A.5), the operator $\tilde{\mathcal{L}}_{\pi} = \mathcal{L}_{\pi}|_{L_0^2(\pi)}$ (the restriction of \mathcal{L}_{π} to $L_0^2(\pi)$) is strictly positive and self-adjoint on $L_0^2(\pi)$. Since consequently $\sigma(\mathcal{L}_{\pi}^{-1}) \in (0, \frac{\|T_{k\pi,\pi}\|_{op}}{\lambda}]$, $\mathcal{L}_{\pi}^{-1} := \tilde{\mathcal{L}}_{\pi}^{-1}$ is a well-defined, bounded, self-adjoint and positive operator from $L_0^2(\pi)$ to $L_0^2(\pi)$. We now take an arbitrary $\psi \in \mathcal{D}(\mathcal{L}_{\pi})$ and denote $\phi = \mathcal{L}_{\pi}\psi$. We can write $\psi = \psi' + C$, where $C = \int \psi(x)d\pi(x)$ is a constant and $\psi' \in L_0^2(\pi) \cap \mathcal{D}(\mathcal{L}_{\pi})$. We then have $\mathcal{L}_{\pi}\psi = \mathcal{L}_{\pi}\psi'$ and $\mathcal{L}_{\pi}^{-1}\mathcal{L}_{\pi}\psi = \psi' = \psi - C$. We also note, that since $\phi \in \operatorname{Im}(\mathcal{L}_{\pi}) \subset L_0^2(\pi)$ we have $\langle C, \phi \rangle_{L^2(\pi)} = 0$. Now by a direct substitution of $\phi = \mathcal{L}_{\pi}\psi$ in (47) we obtain:

$$\langle T_{\pi,k_{\pi}}\phi,\phi\rangle_{L^{2}(\pi)} \geq \lambda \langle \psi'+C,\phi\rangle_{L^{2}(\pi)} = \lambda \langle \mathcal{L}_{\pi}^{-1}\phi,\phi\rangle_{L^{2}(\pi)}.$$
(49)

for all ϕ in the image of \mathcal{L}_{π} . Given that this image is dense in $L_0^2(\pi)$, and that $T_{k_{\pi},\pi}$ and \mathcal{L}_{π} are continuous, this is equivalent to (49) holding for all $\phi \in L_0^2(\pi)$.

Now we prove the reverse implication, that is that if \mathcal{L}_{π}^{-1} is well-defined, bounded, and (49) holds for all $\phi \in L_0^2(\pi)$, then (47) holds for all $\psi \in C_c^{\infty}(\mathbb{R}^d)$. This follows trivially from the fact, that for $\psi \in C_c^{\infty}(\mathbb{R}^d)$ we have $\mathcal{L}_{\pi}\psi \in L_0^2(\pi)$ and $\mathcal{L}_{\pi}^{-1}\mathcal{L}_{\pi}\psi = \psi - C$ where $C = \int \psi(x)d\pi(x)$ as previously. Again, using the fact that $\langle C, \mathcal{L}_{\pi}\psi \rangle_{L^2(\pi)} = 0$ and a direct substitution $\phi = \mathcal{L}_{\pi}\psi$ into (49), we obtain (47) for all $\psi \in C_c^{\infty}(\mathbb{R}^d)$.

B.6. Proof of Corollary 6

We begin this section by an additional result, that shows that the assumption that the spectrum of \mathcal{L}_{π}^{-1} is necessarily discrete if exponential decay near equilibrium (16) holds and the RKHS of k_{π} is infinite dimensional.

Lemma 16. If exponential decay near equilibrium (16) holds, and the RKHS for k_{π} is infinite dimensional, then \mathcal{L}_{π}^{-1} has a discrete spectrum.

Proof. By Proposition 5, exponential convergence near equilibrium implies that \mathcal{L}_{π}^{-1} is a well-defined bounded linear operator on $L_0^2(\pi)$ and the inequality (18) holds for all $\phi \in L_2^0(\pi)$. Let $(l_n)_{n \in \mathbb{N}}$ be the eigenvalues of T_{k,π_k} in descending order. Using the max-min variational formula for the eigenvalues of a compact operator, and applying Proposition 5, we get:

$$l_n = \sup_{\substack{E \subset L_0^2(\pi) \\ \dim(E)=n}} \inf_{\substack{\phi \in E \\ \phi \neq 0}} \frac{\langle \phi, T_{k,\pi_k} \phi \rangle_{L^2(\pi)}}{\langle \phi, \phi \rangle_{L^2(\pi)}} \ge \lambda \sup_{\substack{E \subset L_0^2(\pi) \\ \dim(E)=n}} \inf_{\substack{\phi \in E \\ \phi \neq 0}} \frac{\langle \phi, \mathcal{L}_{\pi}^{-1} \phi \rangle_{L^2(\pi)}}{\langle \phi, \phi \rangle_{L^2(\pi)}}.$$
(50)

We will now show, that this implies that the spectrum of \mathcal{L}_{π}^{-1} is discrete. Let $\mathcal{L}_{\pi}^{1/2}$ be the square root of $\mathcal{L}_{\pi}|_{L_{0}^{2}(\pi)}$, the restriction of \mathcal{L}_{π} to $L_{0}^{2}(\pi)$. As a consequence of the analysis in Appendix B.5, the operator $\mathcal{L}_{\pi}^{1/2}$ is well-defined, strictly positive and self-adjoint from $\mathcal{D}(\mathcal{L}_{\pi}) \cap L_{0}^{2}(\pi)$ to $L_{0}^{2}(\pi)$. By retricting the supremum to ϕ of the form $\phi = \mathcal{L}_{\pi}^{1/2} \psi$, where $\psi \in \mathcal{D}(\mathcal{L}_{\pi}) \cap L_{0}^{2}(\pi)$, we obtain the following lower bound:

$$\sup_{\substack{E \subset L_{0}^{2}(\pi) \\ \dim(E)=n}} \inf_{\substack{\phi \in E \\ \phi \neq 0}} \frac{\langle \phi, \mathcal{L}_{\pi}^{-1} \phi \rangle_{L^{2}(\pi)}}{\langle \phi, \phi \rangle_{L^{2}(\pi)}} \geq \sup_{\substack{E \subset \mathcal{L}_{\pi}^{1/2}(L_{0}^{2}(\pi) \cap \mathcal{D}(\mathcal{L}_{\pi})) \\ \dim(E)=n}} \inf_{\substack{\phi \in E \\ \phi \neq 0}} \frac{\langle \phi, \mathcal{L}_{\pi}^{-1} \phi \rangle_{L^{2}(\pi)}}{\langle \phi, \phi \rangle_{L^{2}(\pi)}} = \left(\inf_{\substack{E \subset L_{0}^{2}(\pi) \cap \mathcal{D}(\mathcal{L}_{\pi}) \\ \dim(E)=n}} \sup_{\substack{\psi \in E \\ \dim(E)=n}} \frac{\langle \psi, \psi \rangle_{L^{2}(\pi)}}{\langle \psi, \mathcal{L}_{\pi} \psi \rangle_{L^{2}(\pi)}} = \left(\inf_{\substack{E \subset L_{0}^{2}(\pi) \cap \mathcal{D}(\mathcal{L}_{\pi}) \\ \dim(E)=n}} \sup_{\substack{\psi \in E \\ \dim(E)=n}} \frac{\langle \psi, \psi \rangle_{L^{2}(\pi)}}{\langle \psi, \psi \rangle_{L^{2}(\pi)}} = \left(\inf_{\substack{E \subset L_{0}^{2}(\pi) \cap \mathcal{D}(\mathcal{L}_{\pi}) \\ \dim(E)=n}} \sup_{\substack{\psi \in E \\ \psi \neq 0}} \frac{\langle \psi, \psi \rangle_{L^{2}(\pi)}}{\langle \psi, \psi \rangle_{L^{2}(\pi)}} \right)^{-1}$$
(51)

It now follows that $\mathcal{L}_{\pi}|_{L_{0}^{2}(\pi)}$ has discrete spectrum. Indeed, if it was not the case, by (Pankrashkin, 2014, Theorem 8.1) the quotients within the parenthesis on the right hand side of (51), would be upper bounded by $\inf \sigma_{\text{ess}}(\mathcal{L}_{\pi}|_{L_{0}^{2}(\pi)}) > 0$, where σ_{ess} is the essential spectrum, and thus the right hand side of (51) would be lower bounded by a constant equal to $\left(\inf \sigma_{\text{ess}}(\mathcal{L}_{\pi}|_{L_{0}^{2}(\pi)})\right)^{-1} > 0$. This is not possible since the eigenvalues $(l_{n})_{n \in \mathbb{N}}$ converge to zero. Since \mathcal{L}_{π} is positive and unbounded on $L_{0}^{2}(\pi)$ and its spectrum is purely discrete, it follows that \mathcal{L}_{π}^{-1} is compact and, by extension, $(\mathcal{L}_{\pi} + I)^{-1}$: $L^{2}(\pi) \to L^{2}(\pi)$ is compact, which means that \mathcal{L}_{π} has compact resolvent.

The proof of Corollary 6 now follows readily:

Proof of Corollary 6. We denote by $(l_n)_{n\in\mathbb{N}}$ the eigenvalues of T_{k,π_k} in descending order and by $(\lambda_n)_{n\in\mathbb{N}}$ the eigenvalues of compact operator $\mathcal{L}_{\pi}^{-1}: L_0^2(\pi) \to L_0^2(\pi)$ also in descending order. By max-min variational characterization of eigenvalues for compact operators, it follows from (50) that $l_n \ge \lambda . \lambda_n$, hence $\lambda_n = \mathcal{O}(l_n)$.

B.7. Bound on eigenvalue decay and proof of Theorem 7

Lemma 17. Let $\gamma_d \sim \mathcal{N}(0, I_d)$ be the standard *d*-dimensional Gaussian measure, and let $\mathcal{L}_{\gamma_d} = -\Delta + \langle x, \nabla \rangle$ be the Ornstein-Uhlenbeck operator on $L^2(\gamma_d)$. If we denote by ρ_n the *n*-th smallest eigenvalue of \mathcal{L}_{γ_d} , then we have:

$$\rho_n = \mathcal{O}(n^{1/d}).$$

Proof. For a multi-index $\alpha = (k_1, \ldots, k_d)$, the multivariate Hermite polynomial is for any $x = (x_1, \ldots, x_d) \in \mathbb{R}^d$:

$$H_{\alpha}(x) = \prod_{i=1}^{d} H_{k_i}(x_i)$$

where for $k_i \in \mathbb{N}$, H_{k_i} denotes the usual one-dimensional k_i th-order Hermite polynomial. It is well known that multivariate Hermite polynomials form an orthogonal basis of $L^2(\gamma_d)$ and that we have

$$\mathcal{L}_{\gamma_d} H_\alpha = |\alpha| H_\alpha$$

where $|\alpha| = \sum_{i=1}^{d} k_i$ (Bakry et al., 2013, Section 2.7.1). Therefore any $k \in \mathbb{N}$ is an eigenvalue of \mathcal{L}_{γ_d} , with multiplicity equal to

$$S_k = \binom{k+d-1}{d-1}$$

which is the number of solutions of the equation $k = \sum_{i=1}^{d} k_i$, where $\{k_i\}_{i=1}^{d}$ takes its values in the set of nonnegative integers. This means, that

$$o_n = k \iff \sum_{i=1}^{k-1} S_i < n \le \sum_{i=1}^k S_i.$$

Since S_i is a polynomial in *i* of degree d - 1, then $\sum_{i=1}^k S_i$ is a polynomial in *k* of degree *d*. Therefore we have $\rho_n = \mathcal{O}(n^{1/d})$.

Corollary 18. The conclusion of Lemma 17 holds for any Schrödinger operator on $L_2(\mathbb{R}^d)$, defined for L > 0 as follows,

$$\mathcal{H}_{\nu_L} := -\Delta + \frac{1}{4}L^2 \|x\|^2 - \frac{1}{2}dL.$$
(52)

Proof. Let $\nu_L \sim \mathcal{N}(0, \frac{1}{L}I_d)$ be a normal measure, and let $\mathcal{L}_{\nu_L} = -\Delta + L\langle x, \nabla \rangle$ be the associated Ornstein-Uhlenbeck operator on $L_2(\nu_L)$. It is easy to see that the map $R_L : L_2(\gamma_d) \to L_2(\nu_L)$ given by:

$$(R_L\phi)(x) = \phi(\sqrt{Lx})$$

is unitary, and that $\mathcal{L}_{\gamma_d} = R_L^* \mathcal{L}_{\nu_L} R_L$. Furthermore, it is standard that \mathcal{L}_{ν_L} is unitarily equivalent to \mathcal{H}_{ν_L} , see (Pavliotis, 2014, Proposition 4.7). Since unitary equivalence preserves the spectrum, the thesis follows.

Lemma 19. Suppose that $\pi \propto e^{-V}$ where V is a $C^2(\mathbb{R}^d)$ potential such that ∇V is L-Lipschitz and assume that \mathcal{L}_{π} has discrete spectrum. If we denote by $\tilde{\lambda}_n$ the n-th smallest eigenvalue of the operator \mathcal{L}_{π} (counting the multiplicities), then

$$\tilde{\lambda}_n \le \mathcal{O}(n^{1/d})$$

Proof. It is easy to show, that if ∇V is Lipschitz and $\int e^{-V(x)} dx$ is finite, then $\lim_{|x|\to\infty} V(x) = \infty$. This is a consequence of the fact that $\lim_{R\to\infty} \int_{|x|>R} e^{-V(x)} dx = 0$ and that on the set $\{y : ||x-y|| \leq 1, \langle \nabla V(x), y-x \rangle \leq 0\}$ we have $-V(y) \geq -V(x) - \frac{L}{2} ||x-y||_2^2$. Therefore V(x) has to diverge to ∞ as $|x| \to \infty$. It follows, that V attains a minimum on \mathbb{R}^d . Assume, without loss of generality, that V attains its (not necessarily unique) minimum at 0. The operator \mathcal{L}_{π} is unitarily equivalent to the Schrödinger operator (Equation 4.130, Pavliotis (2014))

$$\mathcal{H}_{\pi} = -\Delta + \frac{1}{4} \|\nabla V\|^2 - \frac{1}{2} \Delta V$$

on $L_2(\mathbb{R}^d)$. Since unitary equivalence preserves the spectrum, this operator has by assumption a discrete spectrum. We will show that the eigenvalues of $\mathcal{H}_{\nu_L} + dLI$ dominate those of \mathcal{H}_{π} where \mathcal{H}_{ν_L} is defined in (52). For any $\phi \in D(\mathcal{H}_L)$, we have

$$+\infty > \langle \phi, \mathcal{H}_{\nu_L} + dLI\phi \rangle_{L_2(\mathbb{R}^d)} = \langle \phi, \mathcal{H}_{\pi}\phi \rangle_{L_2(\mathbb{R}^d)} + \frac{1}{4} \int \left(L^2 \|x\|^2 - \|\nabla V(x)\|^2 + 2dL + 2\Delta V(x) \right) \phi(x)^2 \mathrm{d}x$$
$$\geq \langle \phi, \mathcal{H}_{\pi}\phi \rangle_{L_2(\mathbb{R}^d)},$$

where the last inequality follows from the *L*-Lipschitz regularity of ∇V . It follows from the above calculation that the domain of the quadratic form $\langle \phi, \mathcal{H}_{\nu_L} + dLI \rangle_{L_2(\mathbb{R}^d)}$, denoted $Q(\mathcal{H}_{\nu_L} + dLI)$ is contained in the domain of the quadratic form $\langle \phi, \mathcal{H}_{\pi} \phi \rangle_{L_2(\mathbb{R}^d)}$, denoted $Q(\mathcal{H}_{\pi})$. It is now a simple consequence of the the Rayleigh-Ritz variational formula that the eigenvalues of $\mathcal{H}_{\nu_L} + dLI$ dominate those of \mathcal{H}_{π} (Pankrashkin, 2014, Corollary 8.6), that is, $\tilde{\lambda}_n \leq \rho_n + dL$ where ρ_n is the *n*-th smallest eigenvalue of the operator \mathcal{H}_{ν_L} . By Corollary 18, we have that $\rho_n = \mathcal{O}(n^{1/d})$, which concludes the proof. \Box

Proof of Theorem 7. Let $(l_n)_{n=1}^{\infty}$ be the eigenvalues of $T_{k_{\pi},\pi}$ and $(\lambda_n)_{n=1}^{\infty}$ be the eigenvalues of \mathcal{L}_{π}^{-1} both in descending order. Since $T_{k_{\pi},\pi}$ is a trace class operator (Steinwart & Christmann, 2008, Theorem 4.27), we know that $\sum_{n=1}^{\infty} l_n < +\infty$. By Corollary 18 exponential convergence near equilibrium would imply that $\sum_{n=1}^{\infty} \lambda_n < +\infty$ as well. On the other hand we have that $\lambda_n = \tilde{\lambda_n}^{-1}$, where $\tilde{\lambda}_n$ are the positive eigenvalues of \mathcal{L}_{π} in a nondecreasing order. By Lemma 19 we have $\tilde{\lambda}_n = \mathcal{O}(n^{1/d})$, which implies that there exists an $N \in \mathbb{N}$ and a constant C > 0, such that for all $n \geq N$ we have $\tilde{\lambda}_n \leq Cn^{1/d}$ and thus $\lambda_n \geq Cn^{-1/d}$. This contradicts the summability of λ_n , and concludes the proof.

B.8. Proof of Lemma 8

Let $\mathcal{H}_0 = \left\{ \sum_{i=1}^m \alpha_i k_{\pi}(\cdot, x_i); m \in \mathbb{N}; \alpha_1, \dots, \alpha_m \in \mathbb{R}; x_1, \dots, x_m \in \mathbb{R}^d \right\}$. Recall that $\mathcal{H}_{k_{\pi}}$ is the set of functions on \mathbb{R}^d for which there exists an \mathcal{H}_0 -Cauchy sequence $(f_n)_n \in \mathcal{H}_0^{\mathbb{N}}$ converging pointwise to f. Let $c \in \mathbb{R}$. Let $f_0 = \sum_{i=1}^m \alpha_i k_{\pi}(\cdot, x_i) \in \mathcal{H}_0$ and assume that $f_0 = c$. Integrating f_0 w.r.t. π yields

$$c = \int f_0(x) d\pi(x) = \sum_{i=1}^m \alpha_i \int k_\pi(x_i, x) d\pi(x) = 0.$$

Hence c = 0. Similarly, let $f \in \mathcal{H}_{k_{\pi}}$ such that f = c, fix $(f_n)_n \in \mathcal{H}_0^{\mathbb{N}}$ such that $||f_n - f||_{\mathcal{H}_{k_{\pi}}} \to 0$.

Under Assumption (A₄), the injection $\iota : \mathcal{H}_{k_{\pi}} \to L^2(\pi)$ is continuous, linear and bounded. Indeed, for any $g \in \mathcal{H}_{k_{\pi}}$

$$\|\iota g\|_{L^{2}(\pi)}^{2} = \int g(x)^{2} d\pi(x) = \int \langle g, k_{\pi}(x, \cdot) \rangle_{\mathcal{H}_{k_{\pi}}}^{2} d\pi(x) \leq \|g\|_{\mathcal{H}_{k_{\pi}}}^{2} \int k_{\pi}(x, x) d\pi(x) =: c_{\pi}^{2} \|g\|_{\mathcal{H}_{k_{\pi}}}^{2},$$

where $c_{\pi} < +\infty$ exists by Assumption (A₄). So $||f_n - f||_{\mathcal{H}_{k_{\pi}}} \to 0$ implies that $||f_n - f||_{L^2(\pi)} \to 0$. Since $\int f_n(x) d\pi(x) = 0$, by the reproducing property and Cauchy-Schwartz inequality, we have:

$$|c| = \left| \int f(x) d\pi(x) \right| = \left| \int (f_n - f)(x) d\pi(x) \right| \le \int |\langle f_n - f, k_\pi(x, \cdot) \rangle_{\mathcal{H}_{k_\pi}} |d\pi(x)| \le \|f_n - f\|_{\mathcal{H}_{k_\pi}} c_\pi^2.$$

which shows that c = 0 since $||f_n - f||_{\mathcal{H}_{k_{\pi}}} \to 0$.

B.9. Proof of Proposition 10

Rather than providing a proof only for the more restrictive smooth submanifolds considered in Proposition 10, we express the result below for general closed nonempty sets. This formulation involves contingent cones which are crucial quantities for studying the invariance of non-smooth sets. They can be informally described as the collection of directions at xthat point either inward or that are tangent to the set \mathcal{M} . More formally, for $d_{\mathcal{M}}(y)$ the distance of $y \in \mathbb{R}^d$ to \mathcal{M} , $T_{\mathcal{M}}(x) := \{v | \liminf_{h \to 0^+} d_{\mathcal{M}}(x + hv)/h = 0\}$. Non-smooth sets were experimentally met whenever we considered Gaussian mixtures with more than three components for which there was no simple axis of symmetry (as on Figure 8). These are cases more intricate than the ones considered in Lemma 11.

Proposition 20. Let $\mathcal{M} \subset \mathbb{R}^d$ be a closed nonempty set and $\mu_0 \in P_c(\mathbb{R}^d)$ with $\operatorname{supp}(\mu_0) \subset \mathcal{M}$. Assume that, for a deterministic $(v_{\mu_t})_{t\geq 0}$ satisfying the Caratheodory-Lipschitz Assumptions (\mathbb{C}_1) - (\mathbb{C}_4) , we have $v_{\mu_t}(x) \in T_{\mathcal{M}}(x)$ where $T_{\mathcal{M}}(x)$ is the contingent cone of \mathcal{M} at $x \in \mathcal{M}$. Then \mathcal{M} is flow-invariant for (4).

Proof. Consider any $x_0 \in \text{supp}(\mu_0)$. By Assumption (\mathbb{C}_2) and Gronwall's lemma, $x'(t) = v_{\mu_t}(x(t))$ can only generate trajectories that do not explode in finite time, i.e. there is no $T < \infty$ such that $\limsup_{t \to T^-} ||x(t)|| = +\infty$. Since $v(\mu)(\cdot)$ is continuous by Assumption (\mathbb{C}_1), and $v_{\mu_t}(x) \in T_{\mathcal{M}}(x)$, we can apply an invariance result (Aubin & Frankowska, 1990, Theorem 10.4.1) stating that all the generated trajectories $x(\cdot)$ stay within \mathcal{M} at all times. This can be informally understood as using directions that are always tangent or pointing within \mathcal{M} cannot push $x(\cdot)$ outside of \mathcal{M} .

The superposition principle (Ambrosio & Crippa, 2014, Theorem 3.4) states that $\operatorname{supp}(\mu_t)$ is contained in the set of positions x(t) reached by all the trajectories satisfying $x'(t) = v_{\mu_t}(x(t))$ for some $x_0 = x(0) \in \operatorname{supp}(\mu_0)$. Consequently $\operatorname{supp}(\mu_t) \subset \mathcal{M}$, namely \mathcal{M} is flow-invariant for (4).

B.10. Proof of Lemma 11

Assume that $d\pi(x) \propto e^{-V(x)} d\lambda(x)$, then $s(x) = -\nabla V(x)$. Translating and reindexing, w.l.o.g. we can take $\mathcal{M} = \operatorname{span}(e_I, \ldots, e_d)$. Thus by symmetry of π w.r.t. \mathcal{M} , we have that $s_i(x) = (\nabla \log \pi(x))_i = 0$ for all i < I and $x \in \mathcal{M}$,

so $s(x) \in \mathcal{M}$. Hence, for every $x \in \mathcal{M}$, $\partial_j s_i(x) = 0$, for i < I and $I \leq j \leq d$. Therefore, for any $u \in \mathcal{M}$ $[J(s)(x)]u = [J(s)(x)]_{|\mathcal{M}}u$, i.e. \mathcal{M} is left invariant under the action of any Jacobian matrix [J(s)(x)] for $x \in \mathcal{M}$.

Consider a radial kernel $k(x,y) = \phi(||x-y||^2/2)$ with $\phi \in C^3(\mathbb{R})$. Then $\nabla_2 k(x,y) = (x-y)\phi'(||x-y||^2/2)$ and $H_1k(x,y) = \phi'(||x-y||^2/2)I_d + (x-y)\otimes(x-y)\phi''(||x-y||^2/2)$. As, for all $x, y \in \mathcal{M}, x-y \in \mathcal{M}, \nabla_2 k(x,y) \in \mathcal{M}$ and \mathcal{M} is left invariant under the action of $H_1k(x,y)$. Since $\nabla \cdot 1 \nabla_2 k(x,y)$ is radial as well, $\nabla_2 (\nabla \cdot 1 \nabla_2 k(x,y)) \in \text{span}(x-y)$. We have thus shown that all the terms in (36) belong to \mathcal{M} if $x, y \in \mathcal{M}$, so $\nabla_2 k_\pi(x,y) \in \mathcal{M}$.

Consequently $\nabla_{W_2} \mathcal{F}(\mu)(y) = \mathbb{E}_{x \sim \mu}[\nabla_2 k_{\pi}(x, y)] \in \mathcal{M}$ for any μ such that $\operatorname{supp}(\mu) \subset \mathcal{M}$. Since the contingent cone of a subspace at any point is equal to the subspace itself, we conclude by applying Proposition 20 to \mathcal{M} .

B.11. Proof of Lemma 12

We start by noting, that $\tilde{\mathcal{L}}_{\pi}(\phi) = \nabla V \cdot \nabla \phi$ is well-defined on $C_c^{\infty}(\mathbb{R}^d)$. Denote $Z = \int_{\mathbb{R}^d} e^{-V(x)} dx$. Using integration by parts, we have for $\phi, \psi \in C_c^{\infty}(\mathbb{R}^d)$:

$$\begin{split} \langle \tilde{\mathcal{L}}_{\pi} \phi, \psi \rangle_{L_{2}(\pi)} &= \frac{1}{Z} \int_{\mathbb{R}^{d}} (\nabla V \cdot \nabla \phi(x) - \Delta \phi(x)) \psi(x) e^{-V(x)} \mathrm{d}x \\ &= \frac{1}{Z} \int_{\mathbb{R}^{d}} (\nabla V \cdot \nabla \phi(x)) \psi(x) e^{-V(x)} + \nabla \phi(x) \cdot (\nabla \psi(x) e^{-V(x)}) \mathrm{d}x \\ &= \frac{1}{Z} \int \langle \nabla \phi(x), \nabla \psi(x) \rangle e^{-V(x)} \mathrm{d}x = \langle \nabla \phi, \nabla \psi \rangle_{L_{2}(\pi)}. \end{split}$$

It follows that $\tilde{\mathcal{L}}_{\pi}$ is symmetric, and since $\langle \tilde{\mathcal{L}}_{\pi}f, f \rangle = \|\nabla f\|_{L_{2}(\pi)}^{2}$, it is positive as well. We can now define \mathcal{L}_{π} as the Friedriechs extension of $\tilde{\mathcal{L}}_{\pi}$ over $L_{2}(\pi)$ (Pankrashkin, 2014, Definition 2.17). This means that, when we consider the bilinear form $F(\phi, \psi) = \langle \tilde{\mathcal{L}}_{\pi}\phi, \psi \rangle$, there exists the smallest closed bilinear form \overline{F} on $L_{2}(\pi)$ which extends it (Pankrashkin, 2014, Proposition 2.16). The operator associated with \overline{F} is a self-adjoint extension of $\tilde{\mathcal{L}}_{\pi}$ on $L_{2}(\pi)$, which is also positive. We denote this extension by \mathcal{L}_{π} . Furthermore, the domain of \mathcal{L}_{π} is by definition contained in the domain of \overline{F} , which is the closure of C_{c}^{∞} for the norm (29) (Pankrashkin, 2014, Proposition 2.8). The claim of density of $C_{c}^{\infty}(\mathbb{R}^{d})$ in $\mathcal{D}(\mathcal{L}_{\pi})$ for the norm (29) now follows.

Recall, that we have shown that for $\phi \in C_c^{\infty}(\mathbb{R}^d)$ we have $\langle \phi, \mathcal{L}_{\pi}\phi \rangle = \|\nabla \phi\|_{L^2(\pi)}^2$. Therefore

$$\|\phi\|_{\mathcal{L}_{\pi}}^{2} = \left(\|\nabla\phi\|_{L^{2}(\pi)}^{2} + \|\phi\|_{L^{2}(\pi)}^{2}\right)$$

which is the $W^{1,2}(\pi)$ Sobolev norm. This means, that the domain of the closure \overline{F} is equal to $W_0^{1,2}(\pi)$, and hence $\mathcal{D}(\mathcal{L}_{\pi}) \subset W_0^{1,2}(\pi)$. In fact, one can establish that $\mathcal{D}(\mathcal{L}_{\pi}) = W_0^{1,2}(\pi) = W^{1,2}(\pi)$, though we will not need this. It now follows that for any $f \in \mathcal{D}(\mathcal{L}_{\pi})$ there exists a weak derivative $\nabla f \in L_2(\pi)$, and it is easy to establish the equality $\|\nabla f\|_{L_2(\pi)}^2 = \langle f, \mathcal{L}_{\pi}f \rangle_{L_2(\pi)}$ by approximating f in the norm $\|\cdot\|_{W^{1,2}(\pi)}$.

It is now easy to show, that the kernel of \mathcal{L}_{π} consists of constant functions. If $\mathcal{L}_{\pi}f = 0$, we have $\|\nabla f\|_{L^{2}(\pi)}^{2} = \langle f, \mathcal{L}_{\pi}f \rangle_{L^{2}(\pi)} = 0$. One can deduce, that if the weak derivative is zero then the function is constant by a standard argument using mollifiers that we sketch below:

Define the convolution in $L^2(\pi)$ by $f \star g(x) = \int f(x-y)g(y)d\pi(y)$. By standard properties of the convolution, for any $\phi \in C_c^{\infty}$, we have $f \star \phi \in C^{\infty} \cap L^2(\pi)$ and $\nabla(f \star \phi) = f \star (\nabla \phi)$, and if f has a weak derivative then by definition $f \star (\nabla \phi) = (\nabla f) \star \phi$. Let ϕ_n be now a mollifier then again, by standard arguments, $f \star \phi_n$ converges to f in $L_2(\pi)$, and also $f \star \phi_n$ is smooth with $\nabla(f \star \phi_n) = (\nabla f) \star \phi_n = 0$, hence $f \star \phi_n$ is constant. It follows that f is constant.

We are left to show that the set $\{(\phi, \mathcal{L}_{\pi}\phi) \in L^2(\pi) \times L^2(\pi) : \phi \in C_c^{\infty}(\mathbb{R}^d)\}$ is dense in the graph of $L^2(\pi)$ for the topology inherited from the normed space $L^2(\pi) \times L^2(\pi)$. The domain of the closed form \overline{F} with the norm (29) is a Hilbert space (Pankrashkin, 2014, Definition 2.5), with an inner product denoted by $\langle \cdot, \cdot \rangle_{\mathcal{L}_{\pi}}$. We have that for any $f \in \mathcal{D}(\mathcal{L}_{\pi})$ there exists a sequence $(\phi_n)_{n=1}^{\infty}$ such that $\lim_{n\to\infty} \|\phi_n - f\|_{\mathcal{L}_{\pi}} = 0$. By the Cauchy-Schwarz inequality, we have that for any $\psi \in C_c^{\infty}$:

$$|\langle \psi, \mathcal{L}_{\pi}(f - \phi_n) \rangle_{L^2(\pi)} + \langle \psi, f - \phi_n \rangle_{L^2(\pi)}| = |\langle \psi, f - \phi_n \rangle_{\mathcal{L}_{\pi}}| \le \|\psi\|_{\mathcal{L}_{\pi}} \|f - \phi_n\|_{\mathcal{L}_{\pi}},$$

and since ϕ_n converges to f in $L^2(\pi)$ strongly, we obtain $\lim_{n\to\infty} \langle \psi, \mathcal{L}_{\pi}(f-\phi_n) \rangle = 0$ for all $\psi \in C_c^{\infty}(\mathbb{R}^d)$. Since $C_c^{\infty}(\mathbb{R}^d)$ is dense in $L^2(\pi)$, it follows that $\mathcal{L}_{\pi}\phi_n$ converges weakly to $\mathcal{L}_{\pi}f$. By a version of Mazur's lemma on weakly and

strongly closed convex sets (Renardy & Rogers, 2006, Lemma 10.19), it follows that there exists a sequence of finite convex combinations of $\mathcal{L}_{\pi}\phi_n$ converging strongly to $\mathcal{L}_{\pi}f$. More precisely, there exists a function $N : \mathbb{N} \to \mathbb{N}$ and a sequence of sets of real numbers $\{\alpha_{k,m}\}_{m=k}^{N(k)}$ such that $\alpha_{k,m} \ge 0$, $\sum_{m=k}^{N(k)} \alpha_{k,m} = 1$, and

$$\lim_{k \to \infty} \sum_{n=k}^{N(k)} \alpha_{k,n} \mathcal{L}_{\pi} \phi_n = \mathcal{L}_{\pi} f$$

in the strong topology on $L^2(\pi)$. It follows by linearity of \mathcal{L}_{π} , that the functions $\psi_k = \sum_{n=k}^{N(k)} \alpha_{k,n} \phi_n$ are C_c^{∞} functions that converge to f in $L^2(\pi)$ strongly, and $\mathcal{L}_{\pi}\psi_k$ converges to $\mathcal{L}_{\pi}f$ strongly in $L^2(\pi)$ as well.

B.12. Proof of Proposition 14

To perform the computations related to Lemma 22 for the induction formula (10), we need a compactly-supported pushforward, however $\nabla_{W_2} \mathcal{F}(\mu_n)(\cdot)$ is not compactly supported in general. We consequently leverage the compactness of the measure it is applied on to perform our analysis. We thus first show by induction that, owing to (10), $\mu_n \in \mathcal{P}_c(\mathbb{R}^d)$ for all $n \in \mathbb{N}$, since $\mu_0 \in \mathcal{P}_c(\mathbb{R}^d)$. Assume that $\mu_n \in \mathcal{P}_c(\mathbb{R}^d)$, then, by definition of the push-forward operation, since $\nabla_{W_2} \mathcal{F}(\mu_n) \in C^1(\mathbb{R}^d; \mathbb{R}^d)$,

$$\operatorname{supp}(\mu_{n+1}) \subset (I - \gamma \nabla_{W_2} \mathcal{F}(\mu_n)) \left(\operatorname{supp}(\mu_n)\right) \subset \operatorname{supp}(\mu_n) + \gamma \left(\max_{x \in \operatorname{supp}(\mu_n)} \|\nabla_{W_2} \mathcal{F}(\mu_n)(x)\|_2\right) B(0, 1) =: S_n.$$
(53)

Hence $\mu_{n+1} \in \mathcal{P}_c(\mathbb{R}^d)$ as claimed.

Consider a path between μ_n and μ_{n+1} of the form $\rho_t = (I - \gamma t \nabla_{W_2} \mathcal{F}(\mu_n))_{\#} \mu_n$. Set $\phi(x) = -\gamma \nabla_{W_2} \mathcal{F}(\mu_n)(x)$, and $s_t(x) = x + t\phi(x)$ which is distributed according to ρ_t for x distributed as μ_n . In general the function $(I - \gamma \nabla_{W_2} \mathcal{F}(\mu_n))$ is not compactly supported so we cannot apply Lemma 22 outright. But since the push-forward is only applied to the compactly supported μ_n in the definition of ρ_t , we can find, through a mollifier, a function $f_{n,t} \in C^1(\mathbb{R}^d; \mathbb{R}^d)$ such that it coincides with $(I - \gamma t \nabla_{W_2} \mathcal{F}(\mu_n))$ on S_n and is supported on $S_n + B(0, 1)$. So $\rho_t = (f_{n,t})_{\#} \mu_n$ and we can apply Lemma 22. Hence $t \mapsto \mathcal{F}(\rho_t)$ is differentiable and absolutely continuous, and consequently

$$\mathcal{F}(\rho_1) - \mathcal{F}(\rho_0) = \dot{\mathcal{F}}(\rho_0) + \int_0^1 \left[\dot{\mathcal{F}}(\rho_t) - \dot{\mathcal{F}}(\rho_0) \right] dt,$$
(54)

where

$$\dot{\mathcal{F}}(\rho_t) = \mathbb{E}_{\substack{x \sim \mu_n \\ x' \sim \mu_n}} \left[\nabla_1 k_\pi \left(s_t(x), s_t(x') \right)^\top \left(\phi(x) \right) \right].$$

The first term in the r.h.s. of (54) is

$$\dot{\mathcal{F}}(\rho_0) = -\gamma \mathbb{E}_{x \sim \mu_n} \left[\|\nabla_{W_2} \mathcal{F}(\mu_n)(x)\|^2 \right]$$

Since $s_t(x) - s_{t'}(x) = (t - t') \phi(x)$, by Assumption (A₁), we derive through Jensen's and Cauchy-Schwarz inequalities that

$$\begin{split} \left| \dot{\mathcal{F}}(\rho_{t}) - \dot{\mathcal{F}}(\rho_{t'}) \right| &= \left| \mathbb{E}_{\substack{x \sim \mu_{n} \\ x' \sim \mu_{n}}} \left[\left(\nabla_{1}k_{\pi}\left(s_{t}(x), s_{t}(x')\right) - \nabla_{1}k_{\pi}\left(s_{t'}(x), s_{t'}(x')\right) \right)^{\top} \phi(x) \right] \right| \\ &\leq \mathbb{E}_{\substack{x \sim \mu_{n} \\ x' \sim \mu_{n}}} \left[\left(\| \nabla_{1}k_{\pi}\left(s_{t}(x), s_{t}(x')\right) - \nabla_{1}k_{\pi}\left(s_{t'}(x), s_{t}(x')\right) \| + \| \nabla_{1}k_{\pi}\left(s_{t'}(x), s_{t}(x')\right) - \nabla_{1}k_{\pi}\left(s_{t'}(x), s_{t'}(x')\right) \| \right) \| \phi(x) \| \right] \\ &\leq |t - t'| \mathbb{E}_{\substack{x \sim \mu_{n} \\ x' \sim \mu_{n}}} \left[\left(L(s_{t}(x')) \| \phi(x) \| + L(s_{t'}(x)) \| \phi(x') \| \right) \| \phi(x) \| \right] \\ &\leq |t - t'| \left(\mathbb{E}_{x' \sim \mu_{n}} \left[L(s_{t}(x')) \right] \mathbb{E}_{x \sim \mu_{n}} \left[\| \phi(x) \|^{2} \right] + \mathbb{E}_{x \sim \mu_{n}} \left[L(s_{t'}(x)) \| \phi(x) \| \right] \mathbb{E}_{x' \sim \mu_{n}} \left[\| \phi(x') \|^{2} \right]^{\frac{1}{2}} \\ &\leq |t - t'| \left(\mathbb{E}_{x \sim \mu_{n}} \left[L(s_{t}(x)) \right] \mathbb{E}_{x \sim \mu_{n}} \left[\| \phi(x) \|^{2} \right] + \mathbb{E}_{x \sim \mu_{n}} \left[L(s_{t'}(x))^{2} \right]^{\frac{1}{2}} \mathbb{E}_{x \sim \mu_{n}} \left[\| \phi(x) \|^{2} \right]^{\frac{1}{2}} \right]^{\frac{1}{2}} \\ &\leq \gamma^{2} |t - t'| \left(\| L \|_{L^{1}(\rho_{t})} + \| L \|_{L^{2}(\rho_{t'})} \right) \mathbb{E}_{x \sim \mu_{n}} \left[\| \nabla_{W_{2}} \mathcal{F}(\mu_{n})(x) \|^{2} \right] \text{ since } \phi(x) = -\gamma \nabla_{W_{2}} \mathcal{F}(\mu_{n})(x). \end{split}$$

Hence, for t' = 0, we have

$$\left| \dot{\mathcal{F}}(\rho_t) - \dot{\mathcal{F}}(\rho_0) \right| \le t \gamma^2 \left(\|L\|_{L^1(\rho_t)} + \|L\|_{L^2(\mu_n)} \right) \mathbb{E}_{x \sim \mu_n} \left[\|\nabla_{W_2} \mathcal{F}(\mu_n)(x)\|^2 \right].$$

Then, since $\rho_t = ((1-t)I + t(I+\phi))_{\#}\mu_n$, by convexity of L, we can write :

$$\begin{split} \|L\|_{L^{1}(\rho_{t})} &= \int |L((1-t)x + t(x+\phi(x))| \mathrm{d}\mu_{n}(x) \\ &\leq (1-t) \|L\|_{L^{1}(\mu_{n})} + t \|L\|_{L^{1}(\mu_{n+1})} \leq (1-t) \|L\|_{L^{2}(\mu_{n})} + t \|L\|_{L^{2}(\mu_{n+1})} \leq M. \end{split}$$

We can thus upper bound the second term in the r.h.s. of (54),

$$\int_{0}^{1} |\dot{\mathcal{F}}(\rho_{t}) - \dot{\mathcal{F}}(\rho_{0})| dt \leq \int_{0}^{1} \left(t 2M\gamma^{2} \mathbb{E}_{x \sim \mu_{n}} \left[\|\nabla_{W_{2}} \mathcal{F}(\mu_{n})(x)\|^{2} \right] \right) dt = \gamma^{2} M \mathbb{E}_{x \sim \mu_{n}} \left[\|\nabla_{W_{2}} \mathcal{F}(\mu_{n})(x)\|^{2} \right].$$

Finally, since $\gamma M \leq 1$, (54) leads to

$$\mathcal{F}(\mu_{n+1}) - \mathcal{F}(\mu_n) \le -\gamma(1 - \gamma M) \mathbb{E}_{x \sim \mu_n} \left[\|\nabla_{W_2} \mathcal{F}(\mu_n)(x)\|^2 \right] \le 0.$$

C. Additional results

C.1. On the kernel integral operator

Lemma 21. For any $f, g \in L^2(\pi)$, we have that

$$\langle S_{\pi,k_{\pi}}f, S_{\pi,k_{\pi}}g\rangle_{\mathcal{H}_{k_{\pi}}} = \iint f(x)^{\top}g(y)k_{\pi}(x,y)\mathrm{d}\pi(x)\mathrm{d}\pi(y).$$

Proof. By using the reproducing property, we deduce that

$$\begin{split} \langle S_{\pi,k_{\pi}}f, S_{\pi,k_{\pi}}g \rangle_{\mathcal{H}^{d}_{k_{\pi}}} &= \langle \int k_{\pi}(x,\cdot)f(x)d\pi(x), \int k_{\pi}(y,\cdot)g(y)d\pi(y) \rangle_{\mathcal{H}^{d}_{k_{\pi}}} \\ &= \sum_{i=1}^{d} \langle \int k_{\pi}(x,\cdot)f_{i}(x)d\pi(x), \int k_{\pi}(y,\cdot)g_{i}(y)d\pi(y) \rangle_{\mathcal{H}_{k_{\pi}}} \\ &= \sum_{i=1}^{d} \iint f_{i}(x)k_{\pi}(x,y)g_{i}(y)d\pi(x)d\pi(y) = \iint f(x)^{\top}g(y)k_{\pi}(x,y)d\pi(x)d\pi(y). \end{split}$$

г	_	-	
L			
L			

C.2. On the differentiation of the squared KSD

For Lemma 22 below, our computations are similar to the ones in Arbel et al. (2019, Lemma 22 and 23) with some terms getting simpler owing to the Stein's property of the Stein kernel, but under a weaker assumption than a uniform Lipschitz constant for the kernel (see the discussion in Appendix A.6).

Lemma 22. Let $q \in \mathcal{P}_2(\mathbb{R}^d)$ and $\phi \in C_c^1(\mathbb{R}^d)$. Consider the path ρ_t from q to $(I + \nabla \phi)_{\#}q$ given by: $\rho_t = (I + t\nabla \phi)_{\#}q$, for all $t \in [0, 1]$. Suppose Assumptions (A₁) and (A₃) hold. Then, $\mathcal{F}(\rho_t)$ is twice differentiable in t with

$$\begin{split} \dot{\mathcal{F}}(\rho_t) = & \mathbb{E}_{\substack{(x) \sim q \\ (x') \sim q}} \left[\nabla_1 k_\pi (x + t \nabla \phi(x), x' + t \nabla \phi(x'))^\top \nabla \phi(x) \right], \\ \ddot{\mathcal{F}}(\rho_t) = & \mathbb{E}_{\substack{(x) \sim q \\ (x') \sim q}} \left[\nabla \phi(x')^\top \nabla_1 \nabla_2 k_\pi (x + t \nabla \phi(x), x' + t \nabla \phi(x')) \nabla \phi(x) \right] \\ & + & \mathbb{E}_{\substack{(x) \sim q^* \\ (x') \sim q^*}} \left[\nabla \phi(x)^\top H_1 k_\pi (x + t \nabla \phi(x), x' + t \nabla \phi(x')) \nabla \phi(x) \right]. \end{split}$$

Proof. The function $f: t \mapsto k_{\pi}(x + t\nabla\phi(x), x' + t\nabla\phi(x'))$ is differentiable for all $x, x' \in \mathbb{R}^d$, and its time derivative is :

$$\dot{f} = \nabla_1 k_\pi (x + t \nabla \phi(x), x' + t \nabla \phi(x'))^\top \nabla \phi(x) + \nabla_2 k_\pi (x + t \nabla \phi(x), x' + t \nabla \phi(x'))^\top \nabla \phi(x')$$

= $\nabla_1 k_\pi (x + t \nabla \phi(x), x' + t \nabla \phi(x'))^\top \nabla \phi(x) + \nabla_1 k_\pi (x' + t \nabla \phi(x'), x + t \nabla \phi(x))^\top \nabla \phi(x')$ (55)

using the symmetry of k_{π} . The two terms on the r.h.s. of the former equation are symmetric w.r.t. x and x', so we will focus on the first one. By the Cauchy-Schwartz inequality and Assumption (A₁),

$$\begin{aligned} |\nabla_1 k_\pi(x+t\nabla\phi(x),x'+t\nabla\phi(x'))^\top\nabla\phi(x)| &\leq \|\nabla_1 k_\pi(x+t\nabla\phi(x),x'+t\nabla\phi(x'))\|\|\nabla\phi(x)\| \\ &\leq (L(x+t\nabla\phi(x))\|x'+t\nabla\phi(x')\|+\|\nabla_1 k_\pi(x+\nabla\phi(x),0)\|)\|\nabla\phi(x)\| \end{aligned}$$

The r.h.s. of the above inequality is integrable in x because $\nabla \phi$ is compactly supported, L is continuous, and because $x' \mapsto ||x' + t\nabla \phi(x')||$ is integrable since $q \in \mathcal{P}_2(\mathbb{R}^d)$. Therefore, by the differentiation lemma (Klenke, 2013, Theorem 6.28), $\mathcal{F}(\rho_t)$ is differentiable and $\dot{\mathcal{F}}(\rho_t) = \mathbb{E}_{(x) \sim q, (x') \sim q}[\dot{f}]$, i.e.

$$\dot{\mathcal{F}}(\rho_t) = \mathbb{E}_{\substack{(x) \sim q \\ (x') \sim q}} \left[\nabla_1 k_\pi (x + t \nabla \phi(x), x' + t \nabla \phi(x'))^\top \nabla \phi(x) \right].$$

Now define the function $g: t \mapsto \nabla_1 k_\pi (x + t \nabla \phi(x), x' + t \nabla \phi(x'))^\top \nabla \phi(x)$. Its time derivative writes as

$$\dot{g} = \nabla \phi(x)^{\top} \nabla_2 \nabla_1 k_\pi (x + t \nabla \phi(x), x' + t \nabla \phi(x')) \nabla \phi(x') + \nabla \phi(x)^{\top} H_1 k_\pi (x + t \nabla \phi(x), x' + t \nabla \phi(x')) \nabla \phi(x).$$

The first term on the r.h.s. is integrable in (x, x') because it is compactly supported and continuous. We now deal with the second term. Under Assumption (A₃), $H_1k_{\pi}(x, y)$ is dominated by a q-integrable function. Then,

$$\|\nabla\phi(x)^{\top}H_1k_{\pi}(x+t\nabla\phi(x),x'+t\nabla\phi(x'))\nabla\phi(x)\| \le \|H_1k_{\pi}(x+t\nabla\phi(x),x'+t\nabla\phi(x'))\|_{op}\|\nabla\phi(x)\|^2$$

The r.h.s. of the above inequality is integrable because $(x, y) \mapsto H_1 k_\pi(x, y)$ is integrable and $\nabla \phi$ is bounded.

D. Additional experiments

Implementation details. The code for KSD Descent is written in Python using Pytorch (Paszke et al., 2019). We use Matplotlib (Hunter, 2007) for figures, Scipy (Virtanen et al., 2020) for the L-BFGS implementation, as well as Numpy (Harris et al., 2020). It is available at https://github.com/pierreablin/ksddescent.

Different initializations and variance for Gaussian mixtures. To discuss in greater detail the results of our second toy example presented in Section 5.1 and illustrated on Figure 3, we performed some more experiments for several choices of initialization and variance, for a mixture of two Gaussians with equal variance (see Figure 7). We also investigated the support of the stationary points of KSD Descent for a mixture of three Gaussians with different variances (see Figure 8). As discussed in Appendix B.9, the support is not necessarily a smooth submanifold or an axis of symmetry.

A comparison between KSD Descent and Stein points. We finally compare KSD Descent and Stein points (Chen et al., 2018). We choose a low dimensional problem, since the approach of Stein points cannot scale to large dimensions. We use the classical 2-D "banana" density. Figure 9 shows the behavior of the two algorithms. Interestingly, KSD Descent succeeds and does not fall into spurious local minima, even tough the density is not log-concave. We posit that this happens because here the potential $\log(\pi)$ does not have saddle points.



Figure 7. Results of KSD Descent for a mixture of two gaussians, depending on their variance and on the initialization of the algorithm. The green crosses indicate the initial particle positions, while the blue ones are the final positions.



Figure 8. Using KSD Descent to sample from an unbalanced mixture of Gaussians. Some particles get stuck in spurious zones, which are not a straight line nor a manifold. The green crosses indicate the initial particle positions, while the blue ones are the final positions.



Figure 9. Comparison of KSD Descent and Stein points on the "banana" dataset. Green points are the initial points for KSD Descent. Both methods work successfully here, even though it is not a log-concave distribution. We posit that KSD Descent succeeds because there is no saddle point in the potential.