

(De)-regularized Maximum Mean Discrepancy Gradient Flow

Zonghao Chen

ZONGHAO.CHEN.22@UCL.AC.UK

Department of Computer Science, University College London, London, WC1V 6LJ, UK

Aratrika Mustafi

ABM6733@PSU.EDU

Department of Statistics, Pennsylvania State University, University Park, PA, 16802 USA

Pierre Glaser

PIERREGLASER@MSN.COM

Gatsby Computational Neuroscience Unit, University College London, London, WC1V 6LJ, UK

Anna Korba

ANNA.KORBA@ENSAE.FR

ENSAE, CREST, Institut Polytechnique de Paris, Palaiseau, France

Arthur Gretton

ARTHUR.GRETTON@GMAIL.COM

Gatsby Computational Neuroscience Unit, University College London, London, WC1V 6LJ, UK

Bharath K. Sriperumbudur

BKS18@PSU.EDU

Department of Statistics, Pennsylvania State University, University Park, PA, 16802 USA

Editor: Quentin Berthet

Abstract

We introduce a (de)-regularization of the Maximum Mean Discrepancy (DrMMD) and its Wasserstein gradient flow. Existing gradient flows that transport samples from source distribution to target distribution with only target samples, either lack tractable numerical implementation (f -divergence flows) or require strong assumptions and modifications, such as noise injection, to ensure convergence (Maximum Mean Discrepancy flows). In contrast, DrMMD flow can simultaneously (i) guarantee near-global convergence for a broad class of targets in both continuous and discrete time, and (ii) be implemented in closed form using only samples. The former is achieved by leveraging the connection between the DrMMD and the χ^2 -divergence, while the latter comes by treating DrMMD as MMD with a de-regularized kernel. Our numerical scheme employs an adaptive de-regularization schedule throughout the flow to optimally balance the trade-off between discretization errors and deviations from the χ^2 regime. The potential application of the DrMMD flow is demonstrated across several numerical experiments, including a large-scale setting of training student/teacher networks.

Keywords: Wasserstein gradient flow, reproducing kernel Hilbert space, maximum mean discrepancy, f -divergences, spectral regularization

1. Introduction

Many applications in computational statistics and machine learning involve approximating a probability distribution π on \mathbb{R}^d (in terms of samples) when only partial information on π is accessible. For example, in Bayesian inference, π is known up to an intractable normalizing constant for complex models. The setting of interest in this work is the so-called generative modeling setting (Brock et al., 2019; Ho et al., 2020; Song et al., 2021; Franceschi et al., 2024) where one assumes access to a set of samples from the target distribution π , with the goal being to generate new samples from π . Recently, a popular framework to perform this

task involves solving a minimization problem in $\mathcal{P}(\mathbb{R}^d)$, the space of probability distributions over \mathbb{R}^d , by choosing the objective function to be a *dissimilarity* function $\mathcal{D}(\cdot\|\pi)$ (a distance or divergence) between probability distributions that satisfies: $\mathcal{D}(\nu\|\pi) = 0$ if and only if $\nu = \pi$. Since only samples from π are available, this problem is solved approximately—yet, the approximate minimizers may converge to π as the number of available samples increases. In particular, in the space of probability distributions with bounded second moment $\mathcal{P}_2(\mathbb{R}^d)$, a common approach is to solve this optimization problem by running a (sample-based) approximation of the Wasserstein gradient flow of the functional $\mathcal{F} = \mathcal{D}(\cdot\|\pi)$, which defines a path of distributions with steepest descent for \mathcal{F} with respect to the Wasserstein-2 distance.

In generative modeling, the choice of $\mathcal{D}(\cdot\|\pi)$ depends on two crucial aspects: First, its flow should admit consistent and preferably tractable numerical implementations using only samples from π , and second, under reasonable assumptions, it should guarantee convergence of its flow to π , the unique *global* minimizer. Combined, these two properties ensure that, in the large sample limit, this algorithm generates new samples from the target. Unfortunately, verifying these two properties simultaneously has proved to be a surprisingly challenging task. For instance, recent approaches based on Maximum Mean Discrepancy (MMD) (Arbel et al., 2019; Hertrich et al., 2024b), the sliced-Wasserstein distance (Liutkus et al., 2019) and the Sinkhorn divergence (Genevay et al., 2018) typically admit consistent finite-sample implementations, but their global convergence guarantees—when they exist—do not apply to most practical targets π of interest. To guarantee global convergence, one could instead choose \mathcal{D} as an f -divergence. For instance, the (reverse) KL divergence and χ^2 -divergence are geodesically convex (Villani et al., 2009, Definition 16.5) when the target is log-concave (i.e. $\pi \propto e^{-V}$ with V convex) (Ohta and Takatsu, 2011), and hence their flows enjoy better convergence behaviour. However, while the population Wasserstein gradient flows of the χ^2 and KL divergences are well-defined (Jordan et al., 1998; Chewi et al., 2020), they do not come naturally with consistent and tractable sample-based implementations. Multiple approaches propose to solve a surrogate optimization problem with samples at each iteration of the flow (Gao et al., 2019; Ansari et al., 2021; Simons et al., 2022; Birrell et al., 2022; Gu et al., 2024; Liu et al., 2024a); however, it remains to be formally established whether these surrogate problems preserve the desirable convergence guarantees of f -divergence flows.

In the face of the trade-offs present in the current approaches, a natural question arises: *Does there exist a divergence functional $\mathcal{D}(\cdot\|\pi)$ whose gradient flow both globally converges, and admits a tractable, consistent sample-based implementation?* In this work, we take a step towards a positive answer by constructing a “de-regularized” variant of the Maximum Mean Discrepancy (DrMMD) and its associated Wasserstein gradient flow. We prove that the DrMMD gradient flow converges exponentially to the global minimum up to a controllable barrier term for targets π that satisfy a Poincaré inequality, in both continuous and discrete time regimes. To do so, we establish and leverage a connection between the DrMMD and the χ^2 divergence, an f -divergence whose gradient flow benefits from strong convergence guarantees. By alternatively viewing DrMMD as MMD with a regularized kernel, DrMMD flow comes with a consistent and tractable implementation when only samples from the target π are available. In addition, given the empirical success of using adaptive kernels in MMD-based generative models (Galashov et al., 2025; Li et al., 2017; Arbel et al., 2018), our paper shows theoretically that using adaptive kernels through adaptive regularization indeed improves the convergence of MMD gradient flow.

This paper is organized as follows. Section 2 introduces the necessary background on reproducing kernel Hilbert spaces (RKHS), the MMD, χ^2 -divergence, and Wasserstein gradient flows. Section 3 introduces DrMMD and shows that DrMMD is a valid probability divergence that metrizes weak convergence. Section 4 uses DrMMD as the optimization objective to define a Wasserstein gradient flow in $\mathcal{P}_2(\mathbb{R}^d)$, and analyzes the convergence of DrMMD flow in continuous time. Sections 5 and 6 define an implementable DrMMD particle descent scheme with both time and space discretization and analyze its convergence. Section 7 discusses other Wasserstein gradient flows related to our DrMMD flow. Section 8 shows experiments that confirm our theoretical results. The proofs of all results are provided in Section 10, with the technical results being relegated to an appendix.

2. Background

In this section, we present the definitions and notation used throughout the paper.

2.1 Notations

Let \mathcal{L}^d be the Lebesgue measure on \mathbb{R}^d . $\mathcal{P}_2(\mathbb{R}^d)$ denotes the set of all Borel probability measures μ on \mathbb{R}^d with finite second moment. For $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, $\mu \ll \pi$ denotes that μ is absolutely continuous with respect to π . We use $\frac{d\mu}{d\pi}$ to denote the Radon-Nikodym derivative. We recall the standard definition of the Kullback-Leibler divergence, $\text{KL}(\mu\|\pi) = \int \log(\frac{d\mu}{d\pi})d\mu$ if $\mu \ll \pi$, $+\infty$ else.

For a continuous mapping $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$, $T_{\#}\mu$ denotes the push-forward measure of μ by T . For any $\pi \in \mathcal{P}_2(\mathbb{R}^d)$, $L^2(\pi)$ is the Hilbert space of (equivalence class of) functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ such that $\int |f|^2 d\pi < \infty$. We denote by $\|\cdot\|_{L^2(\pi)}$ and $\langle \cdot, \cdot \rangle_{L^2(\pi)}$ the norm and the inner product of $L^2(\pi)$. We denote by $C_c^\infty(\mathbb{R}^d)$ the space of infinitely differentiable functions from \mathbb{R}^d to \mathbb{R} with compact support. For a vector valued functions $g : \mathbb{R}^d \rightarrow \mathbb{R}^p$, we abuse the notation of $L^2(\pi)$ and claim $g \in L^2(\pi)$ if $g_i \in L^2(\pi)$ for all $i = 1, \dots, p$ along with $\|g\|_{L^2(\pi)}^2 := \sum_{i=1}^p \|g_i\|_{L^2(\pi)}^2$.

If $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is differentiable, we denote by ∇f the gradient of f and $\mathbf{H}f$ its Hessian. f is α -strongly convex if $\mathbf{H}f \succeq \alpha I$, i.e, $\mathbf{H}f(x) - \alpha I$ is positive semi-definite for any x , where I is the identity matrix (also denotes an identity operator depending on the context). For a vector valued function $g : \mathbb{R}^d \rightarrow \mathbb{R}^d$, if g_i is differentiable for all $i = 1, \dots, d$, $\nabla \cdot g$ denotes the divergence of g . We also denote by Δg the Laplacian of g , where $\Delta g = \nabla \cdot \nabla g$. We use $\|\cdot\|_F$ to denote the matrix Frobenius norm. $a \wedge b$ and $a \vee b$ denote the minimum and maximum of a and b , respectively.

2.2 Reproducing kernel Hilbert spaces

For a positive semi-definite kernel $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$, its corresponding reproducing kernel Hilbert space (RKHS) \mathcal{H} is a Hilbert space with inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ and norm $\|\cdot\|_{\mathcal{H}}$ (Aronszajn, 1950), such that (i) $k(x, \cdot) \in \mathcal{H}$ for all $x \in \mathbb{R}^d$, and (ii) the reproducing property holds, e.g. for all $f \in \mathcal{H}$, $x \in \mathbb{R}^d$, $f(x) = \langle f, k(x, \cdot) \rangle_{\mathcal{H}}$. We denote by \mathcal{H}^d the Cartesian product RKHS consisting of elements $f = (f_1, \dots, f_d)$ with $f_i \in \mathcal{H}$ with inner product $\langle f, g \rangle_{\mathcal{H}^d} = \sum_{i=1}^d \langle f_i, g_i \rangle_{\mathcal{H}}$.

When $\int k(x, x) d\pi(x) < \infty$, \mathcal{H} can be canonically injected into $L^2(\pi)$ using the operator $\iota_\pi : \mathcal{H} \rightarrow L^2(\pi)$, $f \mapsto f$ with adjoint $\iota_\pi^* : L^2(\pi) \rightarrow \mathcal{H}$ given by

$$\iota_\pi^* f(\cdot) = \int k(x, \cdot) f(x) d\pi(x).$$

The operator ι_π and its adjoint can be composed to form an $L^2(\pi)$ endomorphism $\mathcal{T}_\pi := \iota_\pi \iota_\pi^*$ called the *integral operator*, and a \mathcal{H} endomorphism $\Sigma_\pi := \iota_\pi^* \iota_\pi = \int k(\cdot, x) \otimes k(\cdot, x) d\pi(x)$ (where $(a \otimes b)c := \langle b, c \rangle_{\mathcal{H}} a$ for $a, b, c \in \mathcal{H}$) called the *covariance operator*. \mathcal{T}_π is compact, positive, self-adjoint, and can thus be diagonalized into an orthonormal system in $\{e_i\}_{i \geq 1}$ of $L^2(\pi)$ with associated eigenvalues $\varrho_1 \geq \dots \varrho_i \geq \dots \geq 0$.

In this paper, we make the following assumption on our kernel.

Assumption 1 $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is a continuous and c_0 -universal kernel, and there exists $K > 0$ such that $\sup_x k(x, x) \leq K$.

We refer the reader to Carmeli et al. (2010) for the definition of c_0 -universal kernel. The implication of Assumption 1 is that the RKHS \mathcal{H} is compactly embedded into $L^2(\pi)$ (Steinwart and Scovel, 2012, Lemma 2.3), and hence $k(x, x')$ has a absolute, uniform and pointwise convergent Mercer representation (Steinwart and Scovel, 2012, Corollary 3.5),

$$k(x, x') = \sum_{i \geq 1} \varrho_i e_i(x) e_i(x'), \quad (1)$$

for any x and x' in the support of π . Since the kernel is c_0 -universal, the RKHS \mathcal{H} is dense in $L^2(\pi)$ for all Borel probability measures π (Sriperumbudur et al., 2011, Section 3.1) and $\{e_i\}_{i \geq 1}$ becomes an orthonormal basis of $L^2(\pi)$ (Steinwart and Scovel, 2012, Theorem 3.1).

The power of the integral operator \mathcal{T}_π^r is defined as $\mathcal{T}_\pi^r f := \sum_{i \geq 1} \varrho_i^r \langle f, e_i \rangle_{L^2(\pi)} e_i$, $f \in L^2(\pi)$. For $f \in \text{Ran}(\mathcal{T}_\pi^r)$, there exists $q \in L^2(\pi)$ such that $f = \mathcal{T}_\pi^r q$. The exponent r quantifies the smoothness of the range space relative to the original RKHS \mathcal{H} with $0 < r < \frac{1}{2}$ (*resp.* $r > \frac{1}{2}$) yields spaces that are less (*resp.* more) smooth than \mathcal{H} with $\text{Ran}(\mathcal{T}_\pi^{1/2})$ being isometrically isomorphic to \mathcal{H} (Cucker and Zhou, 2007; Fischer and Steinwart, 2020).

We make an additional assumption—commonly employed in the kernel-based gradient flow literature (Glaser et al., 2021; He et al., 2024; Korba et al., 2020; Arbel et al., 2019)—on the regularity of the kernel that will be employed in studying the DrMMD gradient flow.

Assumption 2 $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is twice differentiable in the sense of (Steinwart and Christmann, 2008, Definition 4.35), i.e., for $i, j \in \{1, \dots, d\}$, both $\partial_i \partial_{i+d} k$ and $\partial_i \partial_j \partial_{i+d} \partial_{j+d} k$ exist and are continuous. There exist constants $K_{1d}, K_{2d} > 0$ such that $\|\nabla_1 k(x, \cdot)\|_{\mathcal{H}^d} := \sum_{i=1}^d \|\partial_i k(x, \cdot)\|_{\mathcal{H}} \leq \sqrt{K_{1d}}$ and $\|\mathbf{H}_1 k(x, \cdot)\|_{\mathcal{H}^{d \times d}} := \sum_{i,j=1}^d \|\partial_i \partial_j k(x, \cdot)\|_{\mathcal{H}} \leq \sqrt{K_{2d}}$ for all $x \in \mathbb{R}^d$.

Many kernels satisfy both Assumption 1 and 2, including the class of bounded, continuous, and translation invariant kernels on \mathbb{R}^d whose Fourier transforms have finite second and fourth moments. This is easy to verify by employing the Fourier transform representation of the RKHS (Wendland, 2004, Theorem 10.12) and noting that the finiteness of the RKHS norm of $\nabla_1 k(\cdot, x)$ and $\mathbf{H}_1 k(\cdot, x)$ for all x corresponds to the existence and finiteness of the second and fourth moments of the Fourier transform of the kernel, respectively. This condition is satisfied by the Gaussian kernel, Matérn kernels of order ν with $\nu + \frac{d}{2} \geq 2$ and the inverse multiquadratic kernel.

2.3 Maximum mean discrepancy and χ^2 -divergence

The Maximum Mean Discrepancy (MMD) (Gretton et al., 2012) between μ and π is defined as the RKHS norm of the difference between the mean embeddings¹ $m_\mu := \int k(x, \cdot) d\mu(x)$ and $m_\pi := \int k(x, \cdot) d\pi(x)$.

$$\text{MMD}(\mu \parallel \pi) := \left\| \int k(x, \cdot) d\mu(x) - \int k(x, \cdot) d\pi(x) \right\|_{\mathcal{H}} = \|m_\mu - m_\pi\|_{\mathcal{H}}.$$

The function $m_\mu - m_\pi$ is often referred to as the “witness function”. When the kernel k is c_0 -universal, $\text{MMD}(\mu \parallel \pi) = 0$ if and only if $\mu = \pi$, and the MMD metrizes the weak topology between probability measures (Sriperumbudur et al., 2010; Sriperumbudur, 2016). Given samples (x_1, \dots, x_n) and (y_1, \dots, y_m) from μ and π respectively, the MMD can be consistently estimated in multiple ways (Gretton et al., 2012). For instance, one can compute its “plug-in” estimator, e.g. $\text{MMD}(\hat{\mu} \parallel \hat{\pi})$, where $\hat{\mu} := \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$ and $\hat{\pi} := \frac{1}{m} \sum_{i=1}^m \delta_{y_i}$.

The χ^2 -divergence — a member of the family of f -divergences (Rényi, 1961) — is defined as the variance of the Radon-Nikodym derivative $\frac{d\mu}{d\pi}$ under π :

$$\chi^2(\mu \parallel \pi) := \int \left(\frac{d\mu}{d\pi} - 1 \right)^2 d\pi,$$

when $\mu \ll \pi$, and $+\infty$ otherwise. The χ^2 -divergence has a variational form (Nowozin et al., 2016; Nguyen et al., 2010):

$$\chi^2(\mu \parallel \pi) = \sup_{h \in \mathcal{M}(\mathbb{R}^d)} \int h d\mu - \int \left(h + \frac{1}{4} h^2 \right) d\pi,$$

where $\mathcal{M}(\mathbb{R}^d)$ denote the set of all measurable functions from \mathbb{R}^d to \mathbb{R} . When $\mu \ll \pi$, we prove in Lemma B.1 that the optimal $h^* = \frac{d\mu}{d\pi} - 1 \in L^2(\pi)$ so that it is sufficient to restrict the variational set to $L^2(\pi)$ in contrast to $\mathcal{M}(\mathbb{R}^d)$ for general f -divergences. Since in most cases, $\chi^2(\hat{\mu} \parallel \hat{\pi}) = +\infty$, the χ^2 -divergence does not admit plug-in estimators, and estimating it consistently involves more complicated strategies (Nguyen et al., 2010).

2.4 Wasserstein gradient flows

Gradient flows are dynamics that use local (e.g. *differential*) information about a given functional in order to minimize it as fast as possible. Their exact definition depends on the nature of the input space; in the familiar case of Euclidean space \mathbb{R}^d , the gradient flow of a sufficiently regular $F : \mathbb{R}^d \rightarrow \mathbb{R}$ given some initial condition x_0 is given by the solution $(x_t)_{t \geq 0}$ of $\partial_t x_t = -v(x_t)$, where v is the Fréchet subdifferential of F , a generalization of the notion of derivative to non-smooth functions (Kruger, 2003).

Gradient flows can be extended from Euclidean spaces to the more general class of *metric spaces* (Ambrosio et al., 2005). When the metric space in question is $\mathcal{P}_2(\mathbb{R}^d)$ endowed with the Wasserstein-2 distance, this gradient flow is called the Wasserstein gradient flow $(\mu_t)_{t \geq 0}$. The Wasserstein gradient flow of $\mathcal{F} : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ takes the particular form (Ambrosio et al., 2005, Lemma 10.4.1):

$$\partial_t \mu_t + \nabla \cdot (\mu_t v_t) = 0, \tag{2}$$

1. Such mean embeddings are well-defined under Assumption 1.

where v_t is the Fréchet subdifferential of $\mathcal{F} : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ evaluated at μ_t (Ambrosio et al., 2005, Definition 11.1.1). (2) is an instance of the *continuity equation* with velocity field v_t : under these dynamics, the mass of μ_t is transported in the direction v_t that decreases \mathcal{F} at the fastest rate at each time t . While (2) can be time-discretized in various ways (Santambrogio, 2017; Ambrosio et al., 2005), in this work, we will focus on the *forward Euler* scheme, defined as $\mu_{n+1} := (\text{I} - \gamma v_n)_\# \mu_n$ where $\gamma > 0$ is a step size parameter. Such a scheme is also known as the *Wasserstein Gradient Descent* of \mathcal{F} .

Just as gradient descent in Euclidean spaces, an instrumental property to characterize the convergence of the Wasserstein gradient descent of a functional \mathcal{F} is given by its *geodesic* convexity and smoothness. Among various ways, one can consider to characterize convexity and smoothness through lower and upper bounds on the Wasserstein Hessian of the functional \mathcal{F} (Villani et al., 2009, Proposition 16.2). Given any $\phi \in C_c^\infty(\mathbb{R}^d)$ that defines a constant speed geodesic² starting at μ : $\rho_t = (\text{I} + t\nabla\phi)_\# \mu$ for $0 \leq t \leq 1$, the Wasserstein Hessian of a functional $\mathcal{F} : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ at μ , denoted as $\text{Hess } \mathcal{F}|_\mu$, is an operator from $L^2(\mu)$ to $L^2(\mu)$:³

$$\langle \text{Hess } \mathcal{F}|_\mu \nabla\phi, \nabla\phi \rangle_{L^2(\mu)} = \left. \frac{d^2}{dt^2} \right|_{t=0} \mathcal{F}(\rho_t).$$

A functional \mathcal{F} is said to be *geodesically* M -smooth at μ if $\langle \text{Hess } \mathcal{F}|_\mu \nabla\phi, \nabla\phi \rangle_{L^2(\mu)} \leq M \|\nabla\phi\|_{L^2(\mu)}$, and is said to be *geodesically* Λ -convex at μ if $\langle \text{Hess } \mathcal{F}|_\mu \nabla\phi, \nabla\phi \rangle_{L^2(\mu)} \geq \Lambda \|\nabla\phi\|_{L^2(\mu)}$. Additionally, \mathcal{F} is *geodesically semiconvex* if $-\infty < \Lambda < 0$ and *geodesically strongly convex* if $\Lambda > 0$. Generally, a functional \mathcal{F} that is both smooth and strongly convex is preferred, because its Wasserstein gradient descent has an exponential rate of convergence under a small enough step size γ (Boyd and Vandenberghe, 2004, Section 9.3.1)(Bonet et al., 2024).

Given some probability measure $\pi \in \mathcal{P}_2(\mathbb{R}^d)$, the MMD flow (*resp.* χ^2 flow) is the Wasserstein gradient flow of the functional $\mathcal{F}_{\text{MMD}}(\cdot) = \text{MMD}(\cdot \| \pi)$ (*resp.* $\mathcal{F}_{\chi^2}(\cdot) = \chi^2(\cdot \| \pi)$). As with the MMD, the MMD flow has an analytic finite sample implementation and may be used to construct generative modeling algorithms (Hertrich et al., 2024b,a). The Wasserstein Hessian of \mathcal{F}_{MMD} for smooth kernels is not positively lower bounded (Arbel et al., 2019, Proposition 5), however, so MMD flow only converges up to an unknown barrier (Arbel et al., 2019, Theorem 6), with global convergence only under a strong (and unverifiable) assumption (Arbel et al., 2019, Proposition 7). More recent works (Boufadène and Vialard, 2024) have demonstrated the global convergence of the MMD flow when using the *Coulomb kernel*. This kernel is non-smooth, however, which complicates numerical implementations. In contrast, the Wasserstein Hessian of \mathcal{F}_{χ^2} is positively lower bounded (Ohta and Takatsu, 2011) for log-concave targets π , so \mathcal{F}_{χ^2} is *geodesically strongly convex* and χ^2 flow enjoys exponential rate of convergence. The exponential convergence of χ^2 flow towards the global minimum in fact holds for a broader class of targets π that satisfy a Poincaré inequality (Chewi et al., 2020). The χ^2 flow has so far lacked a tractable sample-based implementation, however, so it has not been widely used in practice.

2. See Appendix A for the definition.

3. Strictly speaking, $\text{Hess } \mathcal{F}|_\mu$ is an operator over the tangent space $\mathcal{T}_\mu \mathcal{P}_2(\mathbb{R}^d)$ which is a subset of $L^2(\mu)$ (Villani et al., 2009).

In the following sections, we will introduce a new Wasserstein gradient flow that combines the computational advantages of the MMD flow with the convergence properties of the χ^2 flow.

3. (De)-regularized Maximum Mean Discrepancy (DrMMD)

In this section, we introduce a (de)-regularized version of maximum mean discrepancy, or DrMMD in short. The DrMMD is rooted in a unified representation of the MMD and the χ^2 -divergence, given in the following proposition, which is proved in Section 10.1.

Proposition 3.1 (MMD and χ^2 -divergence) *Suppose $\frac{d\mu}{d\pi} - 1 \in L^2(\pi)$ for $\mu, \pi \in \mathcal{P}_2(\mathbb{R}^d)$. Then*

$$\text{MMD}^2(\mu||\pi) = \left\| \mathcal{T}_\pi^{1/2} \left(\frac{d\mu}{d\pi} - 1 \right) \right\|_{L^2(\pi)}^2 \quad \text{and} \quad \chi^2(\mu||\pi) = \left\| \mathbf{I} \left(\frac{d\mu}{d\pi} - 1 \right) \right\|_{L^2(\pi)}^2.$$

Remark 3.1 *The χ^2 identity follows from the definition but is provided for comparison purposes. Together, these identities express both the MMD and the χ^2 -divergence as functionals of the (centered) density ratio $\frac{d\mu}{d\pi} - 1$. While the χ^2 -divergence directly computes the $L^2(\pi)$ norm of the centred ratio, the MMD^2 first computes the image by the operator $\mathcal{T}_\pi^{1/2}$ before taking the $L^2(\pi)$ norm. The smoothing effect of the compact operator $\mathcal{T}_\pi^{1/2}$ —note that \mathcal{T}_π is compact if k is bounded as assumed in Assumption 1—has both positive and negative consequences: $\mathcal{F}_{\text{MMD}}(\cdot) = \text{MMD}^2(\cdot||\pi)$ admits finite sample estimators but is not geodesically convex, making the first-order optimization of MMD objective (as done in generative modeling) challenging (Arbel et al., 2019). In contrast, $\mathcal{F}_{\chi^2}(\cdot) = \chi^2(\cdot||\pi)$ is geodesically convex for log-concave targets π (Ohta and Takatsu, 2011) but is hard to estimate with samples.*

With these facts in mind, we introduce a divergence whose purpose is to combine the beneficial properties of both the χ^2 -divergence and MMD. To do so, this divergence computes the $L^2(\pi)$ norm of the image of $\frac{d\mu}{d\pi} - 1$ by an alternative operator which interpolates between \mathbf{I} and $\mathcal{T}_\pi^{1/2}$. We set this operator to be $((\mathcal{T}_\pi + \lambda \mathbf{I})^{-1} \mathcal{T}_\pi)^{1/2}$, where $\lambda > 0$ is a regularization parameter. The operator $((\mathcal{T}_\pi + \lambda \mathbf{I})^{-1} \mathcal{T}_\pi)^{1/2}$ can be seen as a (de)-regularization of the operator $\mathcal{T}_\pi^{1/2}$ used by the MMD—a similar idea has been used in kernel Fisher discriminant analysis (Mika et al., 1999), goodness-of-fit testing (Balasubramanian et al., 2021; Hagrass et al., 2024b), and two-sample testing (Harchaoui et al., 2007; Hagrass et al., 2024a). We call the resulting divergence the (De)-regularized Maximum Mean Discrepancy (DrMMD).

Definition 1 (DrMMD) *Suppose $\frac{d\mu}{d\pi} - 1 \in L^2(\pi)$ where $\mu, \pi \in \mathcal{P}_2(\mathbb{R}^d)$. Then the (de)-regularized maximum mean discrepancy (DrMMD) between $\mu, \pi \in \mathcal{P}_2(\mathbb{R}^d)$ is defined as*

$$\text{DrMMD}(\mu||\pi) = (1 + \lambda) \left\| ((\mathcal{T}_\pi + \lambda \mathbf{I})^{-1} \mathcal{T}_\pi)^{1/2} \left(\frac{d\mu}{d\pi} - 1 \right) \right\|_{L^2(\pi)}^2, \quad (3)$$

where $\lambda > 0$.

While all three operators, I , \mathcal{T}_π and $(\mathcal{T}_\pi + \lambda I)^{-1}\mathcal{T}_\pi$ are diagonalizable in the same eigenbasis of $L^2(\pi)$, the key difference between them lies in the behavior of their eigenvalues. The identity operator I has all eigenvalues 1, the integral operator \mathcal{T}_π has eigenvalues $(\varrho_i)_{i \geq 1}$ which decay to zero as $i \rightarrow \infty$ and the (de)-regularized integral operator $(\mathcal{T}_\pi + \lambda I)^{-1}\mathcal{T}_\pi$ has eigenvalues $(\varrho_i/(\varrho_i + \lambda))_{i \geq 1}$ which either decay to zero or converge to 1 depending on the choice of λ as $i \rightarrow \infty$. $(\mathcal{T}_\pi + \lambda I)^{-1}\mathcal{T}_\pi$ is known in the statistical estimation literature as Tikhonov regularization; alternative definitions of DrMMD could be obtained by using other regularizing operators, such as Showalter regularization in Engl et al. (1996). In this paper, we primarily focus on Tikhonov regularization and leave other types of regularization for future work.

One of the stated purposes of the DrMMD is to retain the computational benefits of the MMD, which are crucial for its use in particle algorithms for generative modeling. To this end, we provide an alternative representation of DrMMD which does not involve the density ratio $\frac{d\mu}{d\pi}$ directly, but only kernel expectations.

Proposition 3.2 (Density ratio-free and variational formulations) *DrMMD can be alternately represented as*

$$\text{DrMMD}(\mu||\pi) = (1 + \lambda) \left\| (\Sigma_\pi + \lambda I)^{-\frac{1}{2}} (m_\mu - m_\pi) \right\|_{\mathcal{H}}^2 \quad (4)$$

$$= (1 + \lambda) \sup_{h \in \mathcal{H}} \left\{ \int h \, d\mu - \int \left(\frac{h^2}{4} + h \right) d\pi - \frac{\lambda}{4} \|h\|_{\mathcal{H}}^2 \right\} \quad (5)$$

with $h_{\mu,\pi}^* = 2(\Sigma_\pi + \lambda I)^{-1}(m_\mu - m_\pi)$ being the witness function.

The proof is in Section 10.2. The density ratio-free representation (4) contains three expectations under μ and π : the mean embeddings m_μ and m_π , and the covariance operator Σ_π . Given samples $\{x_i\}_{i=1}^M \sim \mu$ and $\{y_i\}_{i=1}^N \sim \pi$, we can construct a plug-in finite sample estimator in (71) by replacing μ and π with their empirical counterparts: $\hat{\mu} = \frac{1}{M} \sum_{i=1}^M \delta_{x_i}$ and $\hat{\pi} = \frac{1}{N} \sum_{i=1}^N \delta_{y_i}$.

The density ratio-free representation (4) frames DrMMD as acting on the difference of m_μ and m_π similarly to MMD. In fact, up to a multiplicative factor $(1 + \lambda)$, DrMMD is MMD computed with respect to another kernel \tilde{k} defined as

$$\tilde{k}(x, x') = \left\langle (\Sigma_\pi + \lambda I)^{-\frac{1}{2}} k(x, \cdot), (\Sigma_\pi + \lambda I)^{-\frac{1}{2}} k(x', \cdot) \right\rangle_{\mathcal{H}}. \quad (6)$$

The derivations are provided in Section 10.3. The kernel \tilde{k} is symmetric and positive semi-definite by construction, and its associated reproducing kernel Hilbert space is $\tilde{\mathcal{H}}$. The density ratio-free representation of DrMMD in Proposition 3.2 is already known in (Balasubramanian et al., 2021; Harchaoui et al., 2007; Hagrass et al., 2024a,b) in the context of non-parametric hypothesis testing.

3.1 Properties of DrMMD

In this section, we establish various properties of DrMMD. As discussed earlier, DrMMD is constructed to interpolate between χ^2 -divergence and MMD to exploit the advantages associated with each. The following result formalizes the interpolation property.

Proposition 3.3 (Interpolation property) *Let $\mu, \pi \in \mathcal{P}_2(\mathbb{R}^d)$. If Assumption 1 holds and $\frac{d\mu}{d\pi} - 1 \in L^2(\pi)$, then*

$$\lim_{\lambda \rightarrow 0} \text{DrMMD}(\mu \parallel \pi) = \chi^2(\mu \parallel \pi), \quad \text{and} \quad \lim_{\lambda \rightarrow \infty} \text{DrMMD}(\mu \parallel \pi) = \text{MMD}^2(\mu \parallel \pi).$$

Proposition 3.3, whose proof can be found in Section 10.4, shows that DrMMD asymptotically becomes a probability divergence in the small and large λ regimes. We seek to use DrMMD as a minimizing objective in generative modeling algorithms; however, we need to ensure that DrMMD is a probability divergence for *any* fixed value of λ . This result holds, as shown next.

Proposition 3.4 (DrMMD is a probability divergence) *Under Assumption 1, for any $\lambda \in (0, \infty)$, DrMMD is a probability divergence, i.e., $\text{DrMMD}(\mu \parallel \pi) \geq 0$, with equality iff $\mu = \pi$. Moreover, DrMMD metrizes the weak topology between probability measures, i.e., $\text{DrMMD}(\mu_n \parallel \pi) \rightarrow 0$ iff μ_n converges weakly to $\pi \in \mathcal{P}_2(\mathbb{R}^d)$ as $n \rightarrow \infty$.*

As MMD with c_0 -universal kernels metrizes the weak convergence of distributions (Sriperumbudur, 2016; Simon-Gabriel et al., 2023), Proposition 3.4, whose proof can be found in Section 10.5, shows that for any $\lambda > 0$, DrMMD is “MMD-like” topologically speaking, and is different from the χ^2 -divergence which induces a strong topology (Agrawal and Horel, 2021).

Remark 3.2 *DrMMD is a specific case of a so-called “Moreau envelopes of f -divergences in reproducing kernel Hilbert spaces” introduced in Stein et al. (2025), when the f -divergence is taken as the χ^2 -divergence. This connection is uncovered by the variational formulation of DrMMD in (5). In contrast to general f -divergences, the Moreau envelope of the χ^2 -divergence enjoys a closed-form expression, as we highlight in this paper with various analytical formulas for DrMMD. The interpolation property (Proposition 3.3) and the metrization of weak convergence (Proposition 3.4) are proved concurrently in Corollaries 12 and 13 of Stein et al. (2025), relying on formulation via Moreau envelopes. In our case, we use direct computations thanks to the closed form of DrMMD.⁴*

4. Wasserstein Gradient Flow of DrMMD

Having introduced the DrMMD in the previous section, we now construct and analyze its *Wasserstein Gradient Flow* (WGF). As discussed in Section 2, WGFs define dynamics $(\mu_t)_{t \geq 0}$ in Wasserstein-2 space that minimize a given functional \mathcal{F} by transporting μ_t in the direction of steepest descent, given by the *Fréchet subdifferential* of \mathcal{F} evaluated at μ_t . Given some target distribution π from which we wish to sample, the WGF of $\mathcal{F}_{\text{DrMMD}}(\cdot) = \text{DrMMD}(\cdot \parallel \pi)$, called DrMMD flow, has the potential to form the basis of a generative modeling algorithm, since μ_t progressively minimizes its distance (in the DrMMD sense) to the target π .

To fulfill this potential, two additional ingredients are necessary. The first is to formally establish that μ_t reaches the global minimizer π , and the second is to design a tractable finite-sample algorithm that inherits the convergence properties of the original DrMMD flow.

4. We would like to clarify that Stein et al. (2025) appeared on arxiv when our paper was already under review at ICML 2024.

We defer the second point to Section 5 and focus in this section on showing how DrMMD benefits from its interpolation towards χ^2 -divergence such that the DrMMD flow achieves near-global convergence for a large class of target distributions.

4.1 DrMMD flow: Definition, existence, and uniqueness

To prove that the DrMMD flow is well-defined and admits solutions, the key is to show that the $\mathcal{F}_{\text{DrMMD}}$ admits Fréchet subdifferentials, as formalized in the following proposition.

Proposition 4.1 (DrMMD gradient flow) *Let $\lambda > 0$, and $\mu_0, \pi \in \mathcal{P}_2(\mathbb{R}^d)$. Under Assumption 1 and 2, the functional $\mathcal{F}_{\text{DrMMD}}$ admits Fréchet subdifferential of the form $(1 + \lambda)\nabla h_{\mu_t, \pi}^*$, where $h_{\mu_t, \pi}^*$ is the witness function defined in Proposition 3.2. Consequently, the DrMMD flow is well-defined and is the solution to the following equation*

$$\partial_t \mu_t - \nabla \cdot (\mu_t (1 + \lambda) \nabla h_{\mu_t, \pi}^*) = 0. \quad (7)$$

In addition, the DrMMD flow starting at μ_0 is unique because $\mathcal{F}_{\text{DrMMD}}$ is semiconvex, i.e., for any $\phi \in C_c^\infty(\mathbb{R}^d)$ and $\mu \in \mathcal{P}_2(\mathbb{R}^d)$,

$$\left| \langle \text{Hess } \mathcal{F}_{\text{DrMMD}}|_\mu \nabla \phi, \nabla \phi \rangle_{L^2(\mu)} \right| \leq 2(1 + \lambda) \frac{2\sqrt{K K_{2d}} + K_{1d}}{\lambda} \|\nabla \phi\|_{L^2(\mu)}^2. \quad (8)$$

The proof is in Section 10.6. By recalling the discussion of Wasserstein Hessian in Section 2.4, (8) indicates that $\mathcal{F}_{\text{DrMMD}}$ is both *geodesically smooth* and *geodesically semiconvex*, which is expected because DrMMD is equivalent to MMD with a regularized kernel \tilde{k} defined in (6) and \mathcal{F}_{MMD} is both *geodesically smooth* and *geodesically semiconvex* (Arbel et al., 2019, Proposition 5). Proposition 4.1 is proved concurrently in Corollaries 14 and 20 of Stein et al. (2025) relying on formulation via Moreau envelopes, while our proof uses the closed-form expression for DrMMD.

4.2 Near-global convergence of DrMMD flow

Having defined the DrMMD flow, we are now concerned with its convergence to the target π . Since DrMMD is constructed to interpolate between the MMD and the χ^2 -divergence, DrMMD flow is expected to recover the convergence properties of the χ^2 flow. With this goal in mind, we first study the Wasserstein Hessian of DrMMD and prove that it becomes *asymptotically* positive as $\lambda \rightarrow 0$ for strongly log-concave targets π . Next, to obtain a *non-asymptotic* convergence rate, we take another route and show that the DrMMD flow converges to π exponentially fast in KL divergence up to a barrier term that vanishes in the small λ regime, provided that π satisfies a Poincaré inequality.

4.2.1 NEAR-GEODESIC CONVEXITY OF $\mathcal{F}_{\text{DrMMD}}$

One popular approach to proving that the DrMMD flow $(\mu_t)_{t \geq 0}$ converges to the target π in terms of DrMMD is to show that $\mathcal{F}_{\text{DrMMD}}$ is *geodesically convex* (Ambrosio et al., 2005, Theorem 4.0.4), or, equivalently in our definition in Section 2.4, its Wasserstein Hessian is positive definite. In the next proposition, we show that the Wasserstein Hessian of $\mathcal{F}_{\text{DrMMD}}$ is indeed positive definite for small enough λ ; however, as we will see below (remark 4.2), the form of the result will not allow us to show convergence besides in the limit $\lambda \rightarrow 0$, which leads us to take a different approach in subsequent sections.

Proposition 4.2 (Near-geodesic convexity of $\mathcal{F}_{\text{DrMMD}}$) *Let $\mu, \pi \in \mathcal{P}_2(\mathbb{R}^d)$, $\mu, \pi \ll \mathcal{L}^d$ and $\phi \in C_c^\infty(\mathbb{R}^d)$. Under Assumption 1 and 2, let π be α strongly log-concave, i.e., $\pi \propto \exp(-V)$, $\mathbf{H}V \succeq \alpha \mathbf{I}$, and assume additionally that $x \mapsto \mathbf{H}V(x)$ is continuous. Then for all μ such that $x \mapsto \nabla \log \mu(x)$ is continuous and $\frac{d\mu}{d\pi} - 1 \in \mathcal{H}$,*

$$\langle \text{Hess } \mathcal{F}_{\text{DrMMD}}|_\mu \nabla \phi, \nabla \phi \rangle_{L^2(\mu)} \geq \alpha(1 + \lambda) \int \frac{d\mu}{d\pi}(x) \|\nabla \phi(x)\|^2 d\mu(x) - R(\lambda, \mu, \nabla \phi), \quad (9)$$

where $\lim_{\lambda \rightarrow 0} R(\lambda, \mu, \nabla \phi) = 0$.

The proof can be found in Section 10.7. To obtain this result, we relate the Wasserstein Hessian of $\mathcal{F}_{\text{DrMMD}}$ with that of \mathcal{F}_{χ^2} . When $\frac{d\mu}{d\pi} - 1 \in \mathcal{H}$, they coincide asymptotically as $\lambda \rightarrow 0$, showing that the interpolation properties of DrMMD to the χ^2 -divergence hold at the level of Wasserstein derivatives. Together with the fact that the Wasserstein Hessian of \mathcal{F}_{χ^2} is positive definite for α -strongly log-concave π (Ohta and Takatsu, 2011), we obtain the lower bound in (9). It is noteworthy that although DrMMD can be viewed as squared MMD with a regularized kernel \tilde{k} , the near-geodesic convexity in Proposition 4.2 is not observed for standard MMD with a fixed kernel k because the latter does not interpolate towards χ^2 -divergence.

Remark 4.1 (Geodesic convexity/smoothness trade-off in $\mathcal{F}_{\text{DrMMD}}$) *The geodesic smoothness and near-convexity of the functional $\mathcal{F}_{\text{DrMMD}}$ are characterized by (8) and (9) respectively via upper and lower bounds on the Wasserstein Hessian of $\mathcal{F}_{\text{DrMMD}}$. However, (8) and (9) impose contradictory conditions on the (de)-regularization parameter λ : (8) indicates that DrMMD is smoother if λ is larger while (9) indicates that DrMMD is more convex if λ is small enough. Consequently, there is a trade-off between the geodesic convexity and smoothness of $\mathcal{F}_{\text{DrMMD}}$, which will play an important role in Section 5.*

Remark 4.2 *Proposition 4.2 shows that for fixed μ and $\nabla \phi$, there exists λ small enough yet positive such that $\langle \text{Hess } \mathcal{F}_{\text{DrMMD}}|_\mu \nabla \phi, \nabla \phi \rangle_{L^2(\mu)} > 0$ at μ . The remainder term $R(\lambda, \mu, \nabla \phi)$ is only controlled in the limit as $\lambda \rightarrow 0$, however, which complicates the use of Proposition 4.2 to show global convergence of the DrMMD flow. In the next section, we employ a different set of techniques that rely on the Poincaré condition on π , a condition which, as we show, will ensure a sufficient dissipation of KL divergence along the DrMMD flow to obtain non-asymptotic near-global convergence.*

4.2.2 NEAR-GLOBAL CONVERGENCE OF DrMMD FLOW VIA POINCARÉ INEQUALITY

Even when a functional \mathcal{F} is not geodesically convex, convergence guarantees for its Wasserstein gradient flow $(\mu_t)_{t \geq 0}$ can still be obtained if the target π satisfies certain functional inequalities. Consider the χ^2 flow $(\nu_t)_{t \geq 0}$, for example: if π satisfies the Poincaré inequality, then $(\nu_t)_{t \geq 0}$ converges exponentially fast to π in terms of KL divergence (Chewi et al., 2020, Theorem 1), i.e.,

$$\text{KL}(\nu_T \| \pi) \leq \exp\left(-\frac{2T}{C_P}\right) \text{KL}(\nu_0 \| \pi). \quad (10)$$

Recall that π satisfies a Poincaré inequality (Pillaud-Vivien et al., 2020, Definition 1) if for all functions $f: \mathbb{R}^d \rightarrow \mathbb{R}$ such that $f, \nabla f \in L^2(\pi)$, there exists a constant $C_P > 0$ such that

$$\int f(x)^2 d\pi(x) - \left(\int f(x) d\pi(x) \right)^2 \leq C_P \|\nabla f\|_{L^2(\pi)}^2. \quad (11)$$

The smallest constant C_P for which (11) holds is called the Poincaré constant. The Poincaré inequality is widely used for studying the convergence of Langevin diffusions (Chewi et al., 2024) and χ^2 flow (Chewi et al., 2020; Trillos and Sanz-Alonso, 2020). The Poincaré condition is implied by the strong log-concavity of π , and is weaker than strong log-concavity because it also allows for nonconvex potentials. The set of probability measures satisfying a Poincaré inequality includes distributions with sub-gaussian tails or with exponential tails. This set is also closed under bounded perturbations and finite mixtures (see Vempala and Wibisono (2019) and Chewi et al. (2024) for a more detailed discussion).

Given the interpolation property of DrMMD to χ^2 -divergence, (10) suggests investigating the convergence of the DrMMD flow $(\mu_t)_{t \geq 0}$ in KL divergence under a Poincaré inequality. To this end, we first derive an upper bound on $\text{KL}(\mu_t \|\pi)$ along the DrMMD flow.

Theorem 4.1 (KL control of the DrMMD flow) *Suppose k satisfies Assumptions 1 and 2, and the target π and DrMMD gradient flow $(\mu_t)_{t \geq 0}$ satisfy the following conditions:*

1. π satisfies a Poincaré inequality with constant C_P .
2. $\mu_t, \pi \ll \mathcal{L}^d$.
3. $\frac{d\mu_t}{d\pi} - 1 \in \text{Ran}(\mathcal{T}_\pi^r)$ with $r > 0$, i.e., there exists $q_t \in L^2(\pi)$ such that $\frac{d\mu_t}{d\pi} - 1 = \mathcal{T}_\pi^r q_t$.
4. $\left\| \nabla (\log \pi)^\top \nabla \left(\frac{d\mu_t}{d\pi} \right) \right\|_{L^2(\pi)} \leq \mathcal{J}_t$ and $\left\| \Delta \left(\frac{d\mu_t}{d\pi} \right) \right\|_{L^2(\pi)} \leq \mathcal{I}_t$.
5. For all $i = 1, \dots, d$, $\lim_{x \rightarrow \infty} \left(h_{\mu_t, \pi}^*(x) - 2 \frac{d\mu_t}{d\pi}(x) \right) \left(\partial_i \frac{d\mu_t}{d\pi}(x) \right) \pi(x) \rightarrow 0$.

Then, for any $T \geq 0$,

$$\begin{aligned} \text{KL}(\mu_T \|\pi) &\leq \exp \left(-\frac{2(1+\lambda)}{C_P} T \right) \text{KL}(\mu_0 \|\pi) \\ &\quad + 4(1+\lambda) \lambda^r \int_0^T \exp \left(-\frac{2(1+\lambda)}{C_P} (T-t) \right) \|q_t\|_{L^2(\pi)} (\mathcal{J}_t + \mathcal{I}_t) dt. \end{aligned} \quad (12)$$

The proof, which can be found in Section 10.8, leverages the fact that the DrMMD can approximate not only the χ^2 -divergence, but also its *Wasserstein gradient*. The DrMMD's approximation properties can be combined with functional inequalities to obtain an upper-bound for the continuous-time dissipation of KL divergence along the flow, given by:

$$\frac{d}{dt} \text{KL}(\mu_t \|\pi) \leq -\frac{2(1+\lambda)}{C_P} \text{KL}(\mu_t \|\pi) + \underbrace{4(1+\lambda) \lambda^r \|q_t\|_{L^2(\pi)} (\mathcal{J}_t + \mathcal{I}_t)}_{\text{Approximation error}}, \quad (13)$$

from which Theorem 4.1 follows upon applying the Growall's lemma (Gronwall, 1919). The first term is strictly negative, while the second term is an approximation error term arising

from DrMMD not perfectly matching the χ^2 -divergence for $\lambda > 0$. When $\lambda = 0$, Theorem 4.1 recovers the exponential decay of KL divergence along χ^2 flow in (10).

Remark 4.3 (i) The second condition assumes that μ_t and π have densities.

(ii) The third condition that $\frac{d\mu_t}{d\pi} - 1 \in \text{Ran}(\mathcal{T}_\pi^r)$ is a regularity condition on the density ratio so that it can be well approximated by the witness function $h_{\mu_t, \pi}^*$. This assumption is known as the range assumption in the literature of kernel ridge regression (Cucker and Zhou, 2007; Fischer and Steinwart, 2020). We posit that this assumption can be relaxed to $\frac{d\mu_t}{d\pi} - 1 \in L^2(\pi)$ as in Proposition 3.3 if only asymptotic convergence is needed with no explicit rate.

(iii) The fourth condition is another regularity condition on the density ratio along the flow. This condition is automatically satisfied under a stronger range condition ($r = \frac{1}{2}$) in the third condition, i.e., $\frac{d\mu_t}{d\pi} - 1 \in \mathcal{H}$, along with a moment condition on the score function $\nabla \log \pi$. To see this, notice that we can further write (derivations are provided in Section 10.8)

$$\begin{aligned} \left\| \nabla (\log \pi)^\top \nabla \left(\frac{d\mu_t}{d\pi} \right) \right\|_{L^2(\pi)} &\leq \sqrt{K_{1d}} \|q_t\|_{L^2(\pi)} \|\nabla \log \pi\|_{L^2(\pi)} \\ \left\| \Delta \left(\frac{d\mu_t}{d\pi} \right) \right\|_{L^2(\pi)} &\leq \sqrt{K_{2d}} \|q_t\|_{L^2(\pi)}. \end{aligned} \tag{14}$$

As an illustration, we examine the DrMMD flow under a Gaussian approximation in Appendix C, following Lambert et al. (2022); Liu et al. (2024b), with a Gaussian kernel and Gaussian target π , for which explicit upper bounds \mathcal{I}_t and \mathcal{J}_t can be derived.

(iv) The fifth condition is a boundary condition that allows integration by parts equality in the proof. We highlight that many works (e.g., Theorem 1 of He et al. (2024), Theorem 2 of Nitanda et al. (2022), Lemma 6 of Vempala and Wibisono (2019)) on Wasserstein gradient flow apply integration by parts without explicitly stating this condition.

Additionally, if $\|q_t\|_{L^2(\pi)} \leq Q$, $\mathcal{J}_t \leq \mathcal{J}$, $\mathcal{I}_t \leq \mathcal{I}$ for all $0 \leq t \leq T$, then Theorem 4.1 will ensure KL convergence of the DrMMD flow up to a controllable barrier term, e.g. *near global convergence*.

Corollary 4.1 (Near global convergence of the DrMMD flow) *In addition to the assumptions of Theorem 4.1, if $\|q_t\|_{L^2(\pi)} \leq Q$, $\mathcal{J}_t \leq \mathcal{J}$, $\mathcal{I}_t \leq \mathcal{I}$ for all $0 \leq t \leq T$, where Q, \mathcal{J} , and \mathcal{I} are universal constants independent of λ , then for any $T \geq 0$,*

$$\text{KL}(\mu_T \| \pi) \leq \exp \left(-\frac{2(1+\lambda)}{C_P} T \right) \text{KL}(\mu_0 \| \pi) + 2\lambda^r C_P Q (\mathcal{J} + \mathcal{I}).$$

The proof of Corollary 4.1 follows directly from upper-bounding the second term of (12) with universal constants Q , \mathcal{J} , and \mathcal{I} , and using the closed-form expression of the resulting integral. Corollary 4.1 provides a condition under which the DrMMD flow will exhibit an exponential rate of convergence (linear convergence) in terms of KL divergence up to an extra approximation error term which vanishes as $\lambda \rightarrow 0$. If $\|q_t\|_{L^2(\pi)} \leq Q$, $\mathcal{J}_t \leq \mathcal{J}$, $\mathcal{I}_t \leq \mathcal{I}$ for all $0 \leq t \leq T$ as shown in Corollary 4.1, the approximation error is of explicit order $\mathcal{O}(\lambda^r)$. Therefore, in the continuous time regime, to have a smaller approximation error, it is beneficial to use a small (de)-regularization parameter λ , so that DrMMD flow operates

closer to the regime of χ^2 flow. However, as we will see in the next section, when it comes to time-discretized DrMMD flow, i.e., DrMMD gradient descent, there is a trade-off between the approximation error and the time discretization error such that the selection of λ would require more careful analysis to strike a good balance. Finally, a smaller Poincaré constant C_P results in both a faster rate of convergence and a smaller barrier.

Previously, Arbel et al. (2019) established (sublinear) global convergence of the MMD flow in terms of MMD distance by assuming that a Łojasiewicz inequality (or a variant of it if additionally performing noise injection, see Arbel et al., 2018, Proposition 8) holds along the flow. Our result thus complements that of Arbel et al. (2019) by showing that MMD-type functionals can achieve near-global convergence for targets satisfying a Poincaré inequality regardless of whether such inequalities hold, by studying their behavior in the f -divergence interpolation regime.

Remark 4.4 *Since DrMMD is asymmetric in its arguments, the reader may wonder why we focus on the gradient flow of $\text{DrMMD}(\cdot||\pi)$ instead of $\text{DrMMD}(\pi||\cdot)$. From the convergence standpoint, Theorem 4.1 shows that $\text{DrMMD}(\cdot||\pi)$ converges globally with an exponential rate up to a small barrier when π satisfies Poincaré inequality along with an extra regularity condition on the density ratio. This favorable convergence property is no longer true for $\text{DrMMD}(\pi||\cdot)$. Practically speaking, the Wasserstein gradient of DrMMD requires inverting a kernel integral operator. It is more efficient to do this just once with a kernel integral operator with respect to π as in (3), rather than with respect to μ_t , which would happen at each time (as done by He et al. 2024).*

Computing the DrMMD flow is intractable since the dynamics are in continuous-time, and in practice, we do not have access to π , but only samples from it. In the next two sections, we build a tractable approximation of the DrMMD flow which provably achieves near-global convergence under similar assumptions as the ones in Section 4. Compared to the DrMMD flow, this approximation combines a discretization in time (introduced in Section 5) with a particle-based space-discretization (introduced in Section 6). While the space-discretization techniques that we used are well-known in the Wasserstein gradient flow literature, our time-discretized scheme deviates from standard approaches, which are insufficient to guarantee near-global convergence in our case.

5. Time-discretized DrMMD flow

In this section, we first construct and analyze the forward Euler scheme of (7), a simple time-discretization of the DrMMD flow which we call DrMMD Descent. Compared to the DrMMD flow, the convergence of DrMMD descent is affected by an additional smoothness-related time-discretization error that blows up as λ approaches 0, thus preventing near-global convergence. To address this issue, we propose in Section 5.2 an alternative discrete-time scheme that adapts the value of the regularization coefficient λ across the descent iterates, which we call *Adaptive DrMMD Descent*. We show that this scheme converges in the KL divergence up to a barrier term that vanishes as the discretization step size goes to zero.

5.1 DrMMD Descent

The forward Euler discretization of the DrMMD flow (or DrMMD *Descent* in short) with

step size $\gamma > 0$ consists of sequence of probabilities $(\mu_n)_{n \in \mathbb{N}}$ defined by the recursion

$$\mu_{n+1} = (\mathbf{I} - \gamma(1 + \lambda)\nabla h_{\mu_n, \pi}^*)_{\#} \mu_n, \quad \mu_{n=0} = \mu_0, \quad (15)$$

where $h_{\mu_n, \pi}^* = 2(\Sigma_{\pi} + \lambda \mathbf{I})^{-1}(m_{\mu_n} - m_{\pi})$ is the DrMMD witness function. This scheme was previously considered in Arbel et al. (2019); in particular, Proposition 4 of Arbel et al. (2019) shows that the discrete-time MMD dissipation rate along the MMD Descent iterates follows the (continuous time) rate of the MMD flow up to an error term proportional to the step size γ and the smoothness parameters of the problem⁵.

Next, we turn to study the convergence of DrMMD descent. One way is to treat DrMMD as MMD with a regularized kernel \tilde{k} and follow Proposition 4 of Arbel et al. (2019), however this does not quantitatively take into account the role of (de)-regularization parameter λ that balances the trade-off between geodesic convexity and smoothness of $\mathcal{F}_{\text{DrMMD}}$. Instead, we adopt the same strategy of Section 4 that exploits the interpolation property of DrMMD towards χ^2 -divergence: in the following proposition, we study the dissipation of KL divergence along the DrMMD Descent when the target π satisfies a Poincaré inequality, in which the role of (de)-regularization parameter λ is highlighted in the approximation and discretization errors.

Proposition 5.1 (Descent lemma in KL) *Suppose k satisfies Assumption 1 and 2, and suppose the target π and DrMMD gradient descent iterates $(\mu_n)_{n \in \mathbb{N}}$ satisfy the following:*

1. π satisfies a Poincaré inequality with constant C_P and its potential is β -smooth, i.e., $\pi \propto \exp(-V)$ with $\mathbf{H}V \preceq \beta \mathbf{I}$.
2. $\mu_n, \pi \ll \mathcal{L}^d$.
3. $\frac{d\mu_n}{d\pi} - 1 \in \text{Ran}(\mathcal{T}_{\pi}^r)$ with $r > 0$, i.e., there exists $q_n \in L^2(\pi)$ such that $\frac{d\mu_n}{d\pi} - 1 = \mathcal{T}_{\pi}^r q_n$.
4. $\|q_n\|_{L^2(\pi)} \leq Q$, $\|\nabla V^{\top} \nabla \left(\frac{d\mu_n}{d\pi}\right)\|_{L^2(\pi)} \leq \mathcal{J}$, $\|\Delta \left(\frac{d\mu_n}{d\pi}\right)\|_{L^2(\pi)} \leq \mathcal{I}$ for all $n = 1, \dots, n_{\max}$.
5. For all $i = 1, \dots, d$, $\lim_{x \rightarrow \infty} \left(h_{\mu_n, \pi}^*(x) - 2\frac{d\mu_n}{d\pi}(x)\right) \left(\partial_i \frac{d\mu_n}{d\pi}(x)\right) \pi(x) \rightarrow 0$.
6. There exists a constant $1 < \zeta < 2$ such that for all $n = 1, \dots, n_{\max}$, the step size γ satisfies

$$\gamma \leq \frac{\zeta - 1}{2\zeta(1 + \lambda)\sqrt{\chi^2(\mu_n \|\pi) \frac{K_{2d}}{\lambda}}}. \quad (16)$$

Then for all $0 \leq n \leq n_{\max}$ and $0 < \lambda \leq 1$,

$$\begin{aligned} \text{KL}(\mu_{n+1} \|\pi) - \text{KL}(\mu_n \|\pi) &\leq -\frac{2}{C_P} \chi^2(\mu_n \|\pi) \gamma \\ &\quad + \underbrace{4\gamma\lambda^r Q (\mathcal{J} + \mathcal{I})}_{\text{Approximation error}} + \underbrace{8\gamma^2(\beta + \zeta^2)\chi^2(\mu_n \|\pi) \frac{K_{1d} + K_{2d}}{\lambda}}_{\text{Discretization error}}. \end{aligned} \quad (17)$$

5. In the MMD flow case, the main smoothness parameters is the Lipschitz constant of the kernel.

The proof is provided in Section 10.9. Conditions 1-5 of Proposition 5.1 are similar to conditions 1-5 of Theorem 4.1, which assumes a functional inequality on the target π and regularity on the density ratio $\frac{d\mu_n}{d\pi} - 1$. In a similar spirit to Theorem 4.1, the fourth regularity condition is automatically satisfied under a stronger range assumption on $\frac{d\mu_n}{d\pi} - 1$. For the sake of brevity, we directly assume uniform upper bounds in the fourth condition rather than writing it out in a separate corollary like Corollary 4.1. Compared to the continuous time regime, two extra conditions are necessary. The first one is a smoothness condition on the potential, $\mathbf{H}V \preceq \beta \mathbf{I}$, which is commonly used in the convergence analysis of discrete-time Langevin-based samplers (Dalalyan, 2017; Dalalyan and Karagulyan, 2019; Durmus et al., 2019; Dalalyan et al., 2022; Vempala and Wibisono, 2019). It can be relaxed to ∇V being Hölder-continuous with exponent $s \in [0, 1]$ (Chatterji et al., 2020). The second one, (16), is an upper bound on the step size γ , aligning with the principle that step size should be small enough for the discrete-time scheme to inherit the properties of its continuous analog. This condition will be more thoroughly discussed in Remark 5.4 when all the conditions on γ in Theorem 5.1 and Theorem 5.2 are presented.

Remark 5.1 (Approximation-discretization trade-off of DrMMD Descent)

If we compare the discrete-time KL dissipation of (17) with its continuous-time counterpart in (13), we see that the first two terms on the RHS of (17) admit continuous-time analogs present in (13). The discrete-time KL dissipation contains an additional (positive) term representing the time discretization error: unlike the approximation error term that vanishes as λ approaches 0, this term actually diverges as λ approaches 0. Therefore, replicating the arguments of the continuous-time result of Theorem 4.1 in the discrete-time regime would yield a barrier that does not vanish as $\lambda \rightarrow 0$, hinting at a trade-off between approximation and discretization similar to that of Remark 4.1. In the next section, we propose a refined adaptive discrete-time descent scheme that addresses the convergence issues.

5.2 Adaptive DrMMD Descent

The DrMMD flow and descent dynamics are defined for a value of λ that remains fixed throughout time. The KL dissipation provided in Proposition 5.1 is a function of λ , however; thus, to obtain a sequence of measures with better convergence guarantees than the DrMMD descent, we now construct and analyze a sequence of iterates obtained by selecting, at each iteration, the value of λ minimizing the sum of the approximation error and time-discretization error presented in (17). This sequence, which we term *Adaptive DrMMD descent*, is given by

$$\mu_{n+1} = (\mathbf{I} - \gamma(1 + \lambda)\nabla h_{\mu_n, \pi}^*)_{\#} \mu_n, \quad \lambda_n = \left(2\gamma\chi^2(\mu_n \parallel \pi) \frac{(\beta + \zeta^2)(K_{1d} + K_{2d})}{Q(\mathcal{J} + \mathcal{I})} \right)^{\frac{1}{r+1}}. \quad (18)$$

The particular choice of λ_n above minimizes the sum of the approximation error and time-discretization error presented in (17). This optimal choice indicates that λ_n should shrink towards 0 as $\chi^2(\mu_n \parallel \pi)$ decreases along DrMMD gradient descent: at the early stages of the scheme, it is desirable to have a larger λ_n , corresponding to a smoother objective functional and enabling larger step sizes; then as DrMMD gradient descent iterates μ_n get closer to π , a smaller λ_n enables the scheme to operate closer to the χ^2 flow regime, which metrizes a stronger topology and can better witness the difference between μ_n and π . Additionally,

as the potential V becomes less smooth, i.e., β gets larger, then λ_n should also increase to account for the loss of smoothness from V . Unlike the DrMMD descent in Section 5.1, since the Wasserstein gradient updating μ_n comes from a different DrMMD at each iteration, the adaptive scheme of (18) constitutes a significant departure from related works in Wasserstein gradient descent (Glaser et al., 2021; Arbel et al., 2019; Korba et al., 2021; Hertrich et al., 2024b, 2023; Chewi et al., 2020).

Remark 5.2 (Adaptive kernel) *Recent applications of MMD-based generative modeling algorithms with adaptive kernels (in particular, time-dependent kernel hyperparameters) demonstrate improved empirical performance over fixed kernels in both Wasserstein gradient flow on MMD (Galashov et al., 2025) and generative adversarial networks with an MMD critic (Li et al., 2017; Arbel et al., 2018): the latter can be related to gradient flow on the critic where μ_n is restricted to the output of a generator network (see e.g. Franceschi et al., 2024). As the DrMMD is an MMD with a regularized kernel \tilde{k} that depends on λ_n , the Adaptive DrMMD Descent thus falls into the former category. Galashov et al. (2025, Proposition 3.1) demonstrates faster convergence for MMD gradient flow with an adaptive kernel, for the parametric setting of Gaussian distributions π and μ_t . Our analysis is the first to prove theoretically that adaptive kernels can result in improved convergence for more general nonparametric settings. We believe that the theoretical analysis of adaptivity by varying other hyperparameters (such as the kernel bandwidth for RBF kernels) remains an interesting avenue for future work.*

By leveraging the quasi-descent lemma in KL divergence in Proposition 5.1, we are able to establish the following theorem, which provides a near-global convergence result of the Adaptive DrMMD gradient descent iterates in KL divergence.

Theorem 5.1 (Near-global convergence of adaptive DrMMD gradient descent)

Suppose k satisfies Assumption 1 and 2 and $K \leq 1$, and suppose the target π and adaptive DrMMD gradient descent iterates $(\mu_n)_{n \in \mathbb{N}}$ satisfy the following conditions:

1. π satisfies a Poincaré inequality with constant C_P and its potential is β -smooth, i.e. $\pi \propto \exp(-V)$ with $\mathbf{H}V \preceq \beta \mathbf{I}$.
2. $\mu_n, \pi \ll \mathcal{L}^d$.
3. $\frac{d\mu_n}{d\pi} - 1 \in \text{Ran}(\mathcal{T}_\pi^r)$ with $r > 0$, i.e., there exists $q_n \in L^2(\pi)$ such that $\frac{d\mu_n}{d\pi} - 1 = \mathcal{T}_\pi^r q_n$.
4. $\|q_n\|_{L^2(\pi)} \leq Q$, $\|\nabla V^\top \nabla \left(\frac{d\mu_n}{d\pi}\right)\|_{L^2(\pi)} \leq \mathcal{J}$, $\left\|\Delta \left(\frac{d\mu_n}{d\pi}\right)\right\|_{L^2(\pi)} \leq \mathcal{I}$ for all $n = 1, \dots, n_{\max}$.
5. For all $i = 1, \dots, d$, $\lim_{x \rightarrow \infty} \left(h_{\mu_n, \pi}^*(x) - 2\frac{d\mu_n}{d\pi}(x)\right) \left(\partial_i \frac{d\mu_n}{d\pi}(x)\right) \pi(x) \rightarrow 0$.
6. There exists a constant $1 < \zeta < 2$ such that the step size γ satisfies

$$\gamma \leq \frac{1}{8} \left(\frac{\zeta - 1}{\zeta}\right)^{\frac{2r+2}{2r+1}} \left(\frac{1}{Q}\right) \left(\frac{1}{\mathcal{J} + \mathcal{I}}\right)^{\frac{1}{2r+1}} \left(\frac{1}{K_{2d} \beta + \zeta^2}\right)^{\frac{r}{2r+1}} \wedge \frac{1}{4} \frac{\zeta - 1}{\zeta} \frac{1}{Q^2 K_{2d}} \wedge \frac{C_P}{2} \wedge 1. \quad (19)$$

Then by taking (de)-regularization parameter $\lambda_n = \left(2\gamma\chi^2(\mu_n\|\pi)\frac{(\beta+\zeta^2)(K_{1d}+K_{2d})}{Q(\mathcal{J}+\mathcal{I})}\right)^{\frac{1}{r+1}} \wedge 1$, we have

$$\begin{aligned} \text{KL}(\mu_{n_{\max}}\|\pi) &\leq \exp\left(-\frac{2n_{\max}\gamma}{C_P}\right) \text{KL}(\mu_0\|\pi) \\ &\quad + 4\gamma^{\frac{r}{r+1}}C_PQ^{\frac{2r+1}{r+1}}\left((K_{1d}+K_{2d})(\beta+\zeta^2)\right)^{\frac{r}{r+1}}(\mathcal{J}+\mathcal{I})^{\frac{1}{r+1}}. \end{aligned} \quad (20)$$

The proof can be found in Section 10.10. The conditions 1-5 of Theorem 5.1 are the same as conditions 1-5 of Proposition 5.1. Since the RHS of (19) are all constants, condition 6 is satisfied when the step size γ is small enough.

The implication of Theorem 5.1 is that DrMMD gradient descent exhibits an exponential rate of convergence (linear convergence) in terms of KL divergence up to an extra barrier of order $\mathcal{O}(\gamma^{\frac{r}{r+1}})$. The barrier term shows up as the result of picking the optimal regularization parameter λ_n that best trades off the approximation error and discretization error in Proposition 5.1. Theorem 5.1 is reminiscent of the convergence result of the Langevin Monte Carlo sampling algorithm, whose KL divergence also decreases exponentially up to an extra barrier, but of order $\mathcal{O}(\gamma)$ (Vempala and Wibisono, 2019, Theorem 2). Unlike the continuous-time result of Theorem 4.1, in which the barrier can be made arbitrarily small by taking small enough regularization, taking the step size γ in Theorem 5.1 to be arbitrarily small will significantly impact the rate of convergence, even though it is exponential in terms of n_{\max} . By making the step sizes adaptive with the number of iterations and imposing an extra condition, the barrier term actually vanishes, as demonstrated in the following theorem.

Theorem 5.2 (Global convergence of DrMMD gradient descent) *Suppose that k satisfies Assumptions 1 and 2, and that the conditions in Theorem 5.1 on DrMMD gradient descent iterates $(\mu_n)_{n\in\mathbb{N}}$, target distribution π , regularization coefficient λ_n and step size γ_n are satisfied. If additionally, the step size γ_n satisfies*

$$\gamma_n \leq \frac{1}{(K_{1d}+K_{2d})(\beta+\zeta^2)}\left(\frac{1}{Q(\mathcal{J}+\mathcal{I})}\right)^{\frac{1}{r}}\left(\frac{1}{8C_P}\right)^{\frac{r+1}{r}}\chi^2(\mu_n\|\pi)^{\frac{1}{r}}, \quad (21)$$

for all $n = 1, \dots, n_{\max}$, then

$$\text{KL}(\mu_{n_{\max}}\|\pi) \leq \prod_{n=1}^{n_{\max}} \left(1 - \frac{1}{C_P}\gamma_n\right) \text{KL}(\mu_0\|\pi). \quad (22)$$

The proof can be found in Section 10.11. Compared with (20), (22) provides a cleaner upper bound without the barrier term and leads to global convergence.

Remark 5.3 (Iteration complexity) *We now turn to analyze the iteration complexity of DrMMD gradient descent from Theorems 5.1 and 5.2.*

(i) *From Theorem 5.1, given an error threshold $\delta > 0$, DrMMD descent would reach $\text{KL}(\mu_{n_{\max}}\|\pi) \leq \delta$ after $n_{\max} \geq \frac{C_P}{2\gamma} \log \frac{\text{KL}(\mu_0\|\pi)}{\delta} = \mathcal{O}\left(\left(\frac{1}{\delta}\right)^{\frac{r+1}{r}} \log \frac{1}{\delta}\right)$ iterations. By comparison, when π satisfies a Poincaré inequality, Langevin Monte Carlo (LMC) has an iteration*

complexity of $\mathcal{O}(\frac{1}{\delta})$ up to logarithmic terms (Chewi et al., 2024, Theorem 7). As an approximation to the χ^2 flow, the iteration complexity of DrMMD flow is worse than LMC, because at each iterate of DrMMD flow, there is an extra approximation error in addition to time-discretization error (see our Proposition 5.1). With the optimal choice of regularization λ that balances these two errors, the overall per-step error is of order $\mathcal{O}(\gamma^{1+\frac{r}{r+1}})$. In contrast, at each iteration, LMC only incurs the time-discretization error, which is of order $\mathcal{O}(\gamma^2)$, smaller than that of DrMMD flow. As a result, LMC exhibits better iteration complexity than our DrMMD flow; however, LMC requires knowledge of the score of π , while DrMMD flow only requires samples from π .

(ii) From Theorem 5.2, we consider two cases. On the one hand, if there exists a threshold N_0 such that $\chi^2(\mu_n\|\pi) \geq n^{-r}$ holds for all $n \geq N_0$, then we select step size $\gamma_n \asymp C_P n^{-1}$ for all $n \geq N_0$ such that both (19) and (21) are satisfied, and consequently $\prod_{n=N_0}^{n_{\max}} \left(1 - \frac{1}{C_P} \gamma_n\right) = \mathcal{O}\left(\frac{1}{n_{\max}}\right) \rightarrow 0$ so the iteration complexity of DrMMD gradient descent is $\mathcal{O}(\frac{1}{\delta})$. On the other hand, if such a threshold N_0 does not exist, then there exists a subsequence $n_1, n_2, \dots, n_S, \dots$ such that $\chi^2(\mu_{n_s}\|\pi) \leq n_s^{-r}$ for all $s \geq 1$. Since KL divergence is smaller than χ^2 -divergence (Van Erven and Harremoës, 2014), we have $\text{KL}(\mu_{n_s}\|\pi) \leq n_s^{-r}$ for all $s \geq 1$. Notice that KL divergence is monotonically decreasing based on (63), we have $\text{KL}(\mu_n\|\pi) \leq n_s^{-r}$ for all $n_s \leq n \leq n_{s+1}$ so that $\lim_{n \rightarrow \infty} \text{KL}(\mu_n\|\pi) = 0$. Unfortunately, we are not able to derive iteration complexity in this case because the growth rate of $\{n_s\}_{s \geq 1}$ is unknown.

Remark 5.4 (Step size γ) Theorem 5.1 imposes a condition on the step size γ in (19) and Theorem 5.2 imposes an additional condition in (21). These conditions subsume the condition (16) on step size in Proposition 5.1. (See derivations in Section 10.10.) The conditions (19) and (21) become more stringent as the potential V becomes less smooth, i.e., when β gets larger, similar to the analysis in Langevin Monte Carlo (Balasubramanian et al., 2022; Vempala and Wibisono, 2019) and Stein Variational Gradient Descent (Korba et al., 2020). The condition also becomes more stringent as the density ratio becomes less regular, i.e., when r gets closer to 0 and $Q, \mathcal{J}, \mathcal{I}$ get larger, similar to He et al. (2024).

The adaptive DrMMD descent schemes of (15) and (18) defined via push-forward operations can be equivalently expressed by the following update scheme that defines a trajectory of samples $(y_n)_{n \in \mathbb{N}}$ whose distributions are precisely the Adaptive DrMMD descent iterates $(\mu_n)_{n \in \mathbb{N}}$,

$$y_{n+1} = y_n - \gamma(1 + \lambda_n) \nabla h_{\mu_n, \pi}^*(y_n), \quad y_0 \sim \mu_0. \quad (23)$$

Unfortunately, (23) is still intractable in practice because $h_{\mu_n, \pi}^*$ depends on the unknown distribution μ_n . Therefore, an additional discretization in space is needed to approximate (23) within a tractable algorithm. We propose to do so in the next section through a system of interacting particles, i.e., DrMMD particle descent.

6. DrMMD particle descent

Suppose we have M samples from the target distribution $\{x^{(i)}\}_{i=1}^M \sim \pi$ and N samples from the initial distribution $\{y_0^{(i)}\}_{i=1}^N \sim \mu_0$. The DrMMD particle descent is defined as:

$$y_{n+1}^{(i)} = y_n^{(i)} - \gamma(1 + \lambda_n) \nabla h_{\hat{\mu}_n, \hat{\pi}}^*(y_n^{(i)}), \quad (24)$$

where $\hat{\mu}_n = \frac{1}{N} \sum_{i=1}^N \delta_{y_n^{(i)}}$ and $\hat{\pi} = \frac{1}{M} \sum_{i=1}^M \delta_{x^{(i)}}$ denote respectively the empirical distribution of the particles at time step n and the target, and where $h_{\hat{\mu}_n, \hat{\pi}}^* = 2(\Sigma_{\hat{\pi}} + \lambda_n \mathbf{I})^{-1} (m_{\hat{\mu}_n} - m_{\hat{\pi}})$. Unfortunately, the KL divergence is ill-defined on empirical distributions $\hat{\mu}_n$, which means the analysis of Theorem 5.1 is no longer applicable to study the convergence of DrMMD particle descent. Therefore, in the next theorem, we instead resort to the Wasserstein-2 distance to analyze the convergence of DrMMD particle descent. For simplicity of presentation below, we assume $K \leq 1$.

Theorem 6.1 *Suppose that k satisfies Assumptions 1 and 2 with $K \leq 1$, and that all the conditions in Theorem 5.1 on DrMMD gradient descent iterates $(\mu_n)_{n \in \mathbb{N}}$, target distribution π , regularization coefficient λ_n and step size γ are satisfied. In addition, suppose $(\mu_n)_{n \in \mathbb{N}}$ has bounded fourth moment and the target π satisfies a Talagrand-2 inequality with constant C_T . Let the number of samples M, N satisfy*

$$\begin{aligned} M &\gtrsim \left(\frac{1}{\gamma}\right)^2 \left(\frac{1}{\min_{i=1, \dots, n_{\max}} \text{KL}(\mu_i \| \pi) Z \wedge 1} \right)^{\frac{2}{r+1}} \exp \left(\frac{8n_{\max} \gamma^{\frac{r}{r+1}} R}{\left(\min_{i=1, \dots, n_{\max}} \text{KL}(\mu_i \| \pi) Z \right)^{\frac{1}{r+1}} \wedge 1} \right), \\ N &\gtrsim \left(\frac{1}{\gamma}\right)^{\frac{2r}{r+1}} \exp \left(\frac{8n_{\max} \gamma^{\frac{r}{r+1}} R}{\left(\min_{i=1, \dots, n_{\max}} \text{KL}(\mu_i \| \pi) Z \right)^{\frac{1}{r+1}} \wedge 1} \right) \vee \left(\frac{1}{\gamma}\right)^{\frac{r(d \vee 4)}{2r+2}}, \end{aligned} \quad (25)$$

where \gtrsim means \geq up to constants, $R = K_{1d} + \sqrt{KK_{2d}}$ is a constant that only depends on the kernel, and Z is a constant that only depends on $\beta, \zeta, K_{1d}, K_{2d}, Q, \mathcal{J}, \mathcal{I}$. Then we have

$$\mathbb{E}[W_2(\hat{\mu}_{n_{\max}}, \pi)] \leq \sqrt{2C_T} \exp \left(-\frac{n_{\max} \gamma}{C_P} \right) \sqrt{\text{KL}(\mu_0 \| \pi)} + \mathcal{O} \left(\gamma^{\frac{r}{2r+2}} \right),$$

where the expectation is taken over initial samples $\{y_0^{(i)}\}_{i=1}^N$ drawn from μ_0 .

The proof is provided in Section 10.12. Theorem 6.1 shows that for sufficiently large sample size M and N , DrMMD particle descent exhibits an exponential rate of convergence (linear convergence) in terms of Wasserstein-2 distance up to an extra barrier of order $\mathcal{O} \left(\gamma^{\frac{r}{2r+2}} \right)$.

Remark 6.1 (Talagrand-2 inequality) *We say that the target distribution π satisfies a Talagrand-2 inequality with constant C_T if for any $\nu \in \mathcal{P}_2(\mathbb{R}^d)$,*

$$W_2(\mu, \pi) \leq \sqrt{2C_T \text{KL}(\mu \| \pi)}.$$

A Talagrand-2 inequality implies the Poincaré inequality in (11) with constant $C_P \leq C_T$, so the condition that π satisfies a Talagrand-2 inequality is stronger than the condition in

Theorems 4.1 and 5.1. A Talagrand-2 inequality allows linking of the two key components of Theorem 6.1: the population convergence in terms of KL divergence proved in Theorem 5.1, and the finite-particle propagation of chaos bound in terms of Wasserstein-2 distance proved in Proposition 10.1. Talagrand inequality is widely used in finite-particle convergence analysis of Wasserstein gradient flows (Shi and Mackey, 2024).

Remark 6.2 (Iteration and sample complexity) The choice of $\gamma = \mathcal{O}\left(\delta^{\frac{2r+2}{r}}\right)$ yields $\mathbb{E}[W_2(\hat{\mu}_{n_{\max}}, \pi)] \leq \delta$ with an iteration complexity of $n_{\max} = \mathcal{O}\left(\left(\frac{1}{\delta}\right)^{\frac{2r+2}{r}} \log \frac{1}{\delta}\right)$, which equals the square of the iteration complexity in Theorem 5.1 because Wasserstein-2 distance is of the same order as the square root of KL divergence. From Theorem 5.1, we have

$$\min_{i=1, \dots, n_{\max}} \text{KL}(\mu_i \| \pi) \leq \text{KL}(\mu_{n_{\max}} \| \pi) = \mathcal{O}(\delta^2).$$

Therefore, from (25), the sample complexity is at least

$$M = \text{poly exp}\left(\delta^{-\frac{2}{r} - \frac{2}{r+1}}\right), \quad N = \text{poly exp}\left(\delta^{-\frac{2}{r} - \frac{2}{r+1}}\right) \mathcal{O}\left(\delta^{-dV^4}\right).$$

The poly-exponential sample complexity originates from the propagation of chaos in interacting particle systems (Kac, 1956). Although $\mathcal{O}(\delta^{-dV^4})$ is subsumed by the poly-exponential term, we still make it explicit to show that N suffers from the curse of dimensionality as expected from the Wasserstein-2 distance between an empirical distribution and a continuous distribution (Kloekner, 2012; Lei, 2020). For comparison, the sample complexity of SVGD in Theorem 3 of Shi and Mackey (2024) is $\mathcal{O}(\exp \exp(\delta^{-2}))$. Similar results have also been established for mean-field Langevin dynamics (Suzuki et al., 2023; Chen et al., 2024).

Recently, Banerjee et al. (2025) proposed a refined finite-particle analysis of Stein Variational Gradient Descent (SVGD), which gives a uniform-in-time convergence bound, i.e., the bound does not blow up exponentially fast as the number of iterations n_{\max} grows. The analysis of Banerjee et al. (2025), however, heavily relies on the relation that the time derivative of the KL divergence in the course of SVGD equals the squared kernel Stein discrepancy (KSD), a fact which does not hold for our DrMMD flow. Recently, Chen et al. (2025) extended that analysis to finite-particle MMD gradient descent, but their analysis requires noise injection. We leave a more refined analysis of our finite-particle convergence result to future work.

Having established the convergence of DrMMD gradient flow/descent, we next show that DrMMD particle descent admits a closed-form implementation. Proposition 6.1 shows that $h_{\hat{\mu}_n, \hat{\pi}}^*$ in (24), defined through the inverse of covariance operators, is computable using Gram matrices.

Proposition 6.1 Given empirical distributions $\hat{\mu}_n = \frac{1}{N} \sum_{i=1}^N y_n^{(i)}$, $\hat{\pi} = \frac{1}{M} \sum_{i=1}^M x^{(i)}$ and Gram matrices $K_{xx} = k(x^{1:M}, x^{1:M}) \in \mathbb{R}^{M \times M}$ and $K_{xy} = k(x^{1:M}, y_n^{1:N}) \in \mathbb{R}^{M \times N}$, the witness function $h_{\hat{\mu}_n, \hat{\pi}}^*$ can be computed as:

$$\begin{aligned} h_{\hat{\mu}_n, \hat{\pi}}^*(\cdot) &= \frac{2}{N\lambda_n} k(\cdot, y^{1:N}) \mathbb{1}_N - \frac{2}{M\lambda_n} k(\cdot, x^{1:M}) \mathbb{1}_M - \frac{2}{N\lambda_n} k(\cdot, x^{1:M}) (M\lambda_n \mathbf{I} + K_{xx})^{-1} K_{xy} \mathbb{1}_N \\ &\quad + \frac{2}{M\lambda_n} k(\cdot, x^{1:M}) (M\lambda_n \mathbf{I} + K_{xx})^{-1} K_{xx} \mathbb{1}_M, \end{aligned} \quad (26)$$

where $\mathbb{1}_M \in \mathbb{R}^M$, $\mathbb{1}_N \in \mathbb{R}^N$ are column vectors of ones.

Algorithm 1 DrMMD particle descent

Input: Target samples $\{x^{(i)}\}_{i=1}^M \sim \pi$ and initial source samples $\{y_0^{(i)}\}_{i=1}^N \sim \mu_0$. Hyperparameters: step size γ , initial (de)-regularization coefficient λ_0 , maximum number of iterations n_{max} and regularity r .

For $n = 0$ to n_{max} :

1. Compute witness function $h_{\hat{\mu}_n, \hat{\pi}}$ from (26).
2. Compute $\text{DrMMD}(\hat{\mu}_n, \hat{\pi})$ with $h_{\hat{\mu}_n, \hat{\pi}}$ from (71).
3. Rescale regularization coefficient $\lambda_n \propto \text{DrMMD}(\hat{\mu}_n, \hat{\pi})^{\frac{1}{r+1}}$.
4. Update particles using (24):

$$y_{n+1}^{(i)} = y_n^{(i)} - \gamma(1 + \lambda_n) \nabla h_{\hat{\mu}_n, \hat{\pi}}^*(y_n^{(i)})$$

EndFor

Output: $\{y_{n_{max}}^{(i)}\}_{i=1}^N$.

The proof of Proposition 6.1 can be found in Section 10.14. The gradient of $h_{\hat{\mu}_n, \hat{\pi}}^*$ can be obtained using automatic differentiation libraries such as JAX (Bradbury et al., 2018).

As indicated in (18), the regularization parameter λ_n should be chosen to be proportional to $\chi^2(\mu_n \parallel \pi)^{\frac{1}{r+1}}$. In practice, however, both $\chi^2(\mu_n \parallel \pi)$ and r are not accessible and $\chi^2(\mu_n \parallel \pi)$ is not even well-defined for the particle descent algorithm. To address this, we use $\text{DrMMD}(\hat{\mu}_n \parallel \hat{\pi})$ as a proxy for $\chi^2(\mu_n \parallel \pi)$ (see Algorithm 1), which admits a closed-form expression with particles. The parameter r is picked via a search over a pre-defined set $\{0.1, 0.5, 1.0\}$. The step size γ should be chosen to satisfy the upper bound (21), which contains several constants that cannot generally be computed. In practice, our approach has been to select γ to be sufficiently small for the flow to converge empirically. The final algorithm is summarized in Algorithm 1.

At every iteration, computing $h_{\hat{\mu}_n, \hat{\pi}}^*$ with adaptive regularization λ_n has a time complexity of $\mathcal{O}(M^3 + NM + N^2)$ due to matrix inversion and multiplication. For DrMMD particle descent with fixed λ , however, the total computational cost can be reduced to $\mathcal{O}(NM + N^2)$, which is exactly the same as MMD flow, because inversion of the $M \times M$ Gram matrix is only required once, and so it can be pre-computed at initialization (see Remark 4.4). In contrast, when $N = M$, the complexity of Sinkhorn flow is $\mathcal{O}(N^2/\epsilon^3)$ (Feydy et al., 2019) with ϵ being the hyperparameter in Sinkhorn divergence, and the complexity of KALE flow is $\mathcal{O}(N^3)$ (Glaser et al., 2021).

7. Related Work

In this section, we discuss the works in the literature that are related to our proposed DrMMD flow and spectral (de)-regularization.

7.1 Gradient flows

Stein Variational Gradient Descent (SVGD) is a popular algorithm for sampling from distributions using only an unnormalized density. It can be written as either a gradient flow of the Kullback-Leibler (KL) divergence where the Wasserstein gradient of the KL is

preconditioned by \mathcal{T}_{μ_t} (Liu and Wang, 2016; Liu, 2017; Korba et al., 2020), or as a gradient flow of the χ^2 -divergence whose Wasserstein gradient is preconditioned by \mathcal{T}_π (Chewi et al., 2020),

$$\frac{\partial \mu_t}{\partial t} = \nabla \cdot \left(\mu_t \mathcal{T}_{\mu_t} \nabla \log \frac{d\mu_t}{d\pi} \right) = \nabla \cdot \left(\mu_t \mathcal{T}_\pi \nabla \frac{d\mu_t}{d\pi} \right).$$

Since SVGD may smooth the trajectory too much, He et al. (2024) considered a (de)-regularized SVGD flow,

$$\frac{\partial \mu_t}{\partial t} = \nabla \cdot \left(\mu_t (\mathcal{T}_{\mu_t} + \lambda \text{I})^{-1} \mathcal{T}_{\mu_t} \nabla \log \frac{d\mu_t}{d\pi} \right), \quad (27)$$

which approaches the KL gradient flow (Langevin diffusion) as $\lambda \rightarrow 0$, demonstrating faster convergence than SVGD. A key difference between (de)-regularization in (27) of He et al. (2024) and our DrMMD flow is that the flow in (27) is driven by the regularized version of the Wasserstein gradient of KL divergence while DrMMD flow is driven by the Wasserstein gradient of the regularized χ^2 -divergence. An alternative interpretation of this difference is that the flow in (27) is the gradient flow of KL divergence w.r.t. the regularized Stein geometry (Duncan et al., 2023), whereas the DrMMD flow is the gradient flow of regularized χ^2 -divergence w.r.t. the Wasserstein geometry.

In addition to sampling from unnormalized distributions, Wasserstein gradient flows (particularly MMD flows) are widely used in the field of generative modelling (Birrell et al., 2022; Gu et al., 2024; Hertrich et al., 2023, 2024b,a; Galashov et al., 2025). The MMD flow (with a smooth kernel) (Arbel et al., 2019) can be written as

$$\frac{\partial \mu_t}{\partial t} = \nabla \cdot \left(\mu_t \nabla \mathcal{T}_\pi \left(\frac{d\mu_t}{d\pi} - 1 \right) \right). \quad (28)$$

The MMD flow is known to get trapped in local minima, and several modifications have been proposed to avoid this in practice, such as noise injection (see Arbel et al., 2019, Proposition 8) or non-smooth kernels, e.g., based on negative distances (Sejdinovic et al., 2013). MMD gradient flows with non-smooth kernels have better empirical performance (Hertrich et al., 2024a,b), but they do not preserve discrete measure and rely on approximating implicit time discretizations (Hertrich et al., 2024a) or slicing (Hertrich et al., 2024b); and they have no local minima apart from the global one (Boufaden and Vialard, 2024).

Recall that our DrMMD flow takes the form

$$\frac{\partial \mu_t}{\partial t} = \nabla \cdot \left(\mu_t \nabla (\mathcal{T}_\pi + \lambda \text{I})^{-1} \mathcal{T}_\pi \left(\frac{d\mu_t}{d\pi} - 1 \right) \right),$$

which (de)-regularizes the MMD flow similarly to how (27) (de)-regularizes SVGD. It is both proved theoretically in Theorem 4.1, 5.1, 6.1, and verified empirically in Section 8, that (de)-regularization results in faster convergence than MMD flow.

Another closely related flow called LAWGD is considered in Chewi et al. (2020), which swaps the gradient and integral operators of SVGD, leading to the following flow:

$$\frac{\partial \mu_t}{\partial t} = \nabla \cdot \left(\mu_t \nabla \mathcal{T}_\pi \frac{d\mu_t}{d\pi} \right).$$

LAWGD closely resembles the MMD flow in (28), but Chewi et al. (2020) proposes to replace \mathcal{T}_π with an inverse diffusion operator, which requires computing the eigenspectrum of the latter and is unlikely to scale in high dimensions.

KALE (Glaser et al., 2021) kernelizes the variational formulation of the KL divergence in a similar way as DrMMD kernelizes the χ^2 -divergence in (5), but the KALE witness function does not have a closed form expression, so it requires solving a convex optimization problem, which makes the simulation of KALE gradient flow with particles computationally more expensive. Recently, Stein et al. (2025) studied kernelized variational formulation of f -divergences (referred to as Moreau envelopes of f -divergences in RKHS), which subsume both KALE (Glaser et al., 2021) and DrMMD. They prove that these functionals are lower semi-continuous and that their Wasserstein gradient flows are well-defined for smooth kernels. They do not study the convergence properties of their proposed flows, however.

The proposed DrMMD interpolates between MMD and χ^2 -divergence by using a specific spectral regularization known as Tikhnov regularization $(1+\lambda)(\mathcal{T}+\lambda\mathbf{I})^{-1}\mathcal{T}$, which interpolates between the identity operator \mathbf{I} as $\lambda \rightarrow 0$ and \mathcal{T} as $\lambda \rightarrow \infty$. DrMMD and its associated Wasserstein gradient flow can be easily extended to other spectral regularization strategies, such as the Showalter regularization, Landweber iteration, or cutoff regularization (Engl et al., 1996) using the techniques in (Bauer et al., 2007; Hagrass et al., 2024a). Alternative approximations to χ^2 -divergence have been proposed in the literature based on the idea of mollifiers, whose Wasserstein gradient flows have been constructed and for which convergence of the flows has been analyzed (Li et al., 2023; Craig et al., 2023, 2025). Compared to DrMMD flow, these gradient flows rely on additional approximations, such as the use of log-sum-exp in Li et al. (2023) and the use of numerical integration to estimate convolution in Craig et al. (2023, 2025)—and are not directly applicable in generative modeling settings where only samples are available.

7.2 Comparison with diffusion-based generative models

Diffusion-based generative models have been widely adopted in practice and are closely related to Wasserstein gradient flows (Song et al., 2021; Ho et al., 2020). These models generate high-quality samples by reversing a pre-defined forward diffusion process, which gradually corrupts data with noise. To implement the reverse process, the established practice is to estimate the score function via denoising score matching (Song et al., 2021). In contrast, Wasserstein gradient flows directly construct a trajectory by descending the objective in the steepest direction with respect to the Wasserstein metric. In particular, our proposed DrMMD gradient flow offers a tractable velocity field with a consistent finite-sample estimator without solving an additional optimization problem like score matching.

We emphasize that the main contribution of our paper is to establish convergence of the DrMMD flow, and that diffusion models for image generation require additional implementation details—most notably, the inductive biases introduced by deep neural networks. One possible avenue for future work is to simulate DrMMD gradient flows with kernels induced by learned deep neural network features on the data. In the case of MMD, this idea has been explored by Galashov et al. (2025), who demonstrate generation performance comparable to established diffusion models. Another promising direction, proposed by Hertrich et al. (2024b), involves first using MMD gradient flow with Riesz

kernels to generate particle trajectories, and then distilling these trajectories into a neural network-based generator. Both approaches would be of interest to extend our DrMMD gradient flow to domains such as image generation, as a topic for future work.

7.3 (De)-regularization for supervised learning and hypothesis testing

The idea of (de)-regularization is not new, and has been used in kernel Fisher discriminant analysis (Mika et al., 1999) and kernel ridge regression (Caponnetto and De Vito, 2007; Schölkopf and Smola, 2002). Subsequently, Harchaoui et al. (2007) employed this statistic in two-sample testing, where they constructed a test statistic that (de)-regularizes $\text{MMD}(\mu \parallel \pi)$ with both covariance operators Σ_μ, Σ_π . This work has been recently generalized in Hagrass et al. (2024a) to more general spectral regularizations. A (de)-regularized statistic is also employed by Balasubramanian et al. (2021); Hagrass et al. (2024b) in the context of a goodness-of-fit test. Balasubramanian et al. (2021) refers to (de)-regularized MMD as ‘Moderated MMD’. To the best of our knowledge, the present work represents the first instance of the (de)-regularized MMD being used as a distance functional in Wasserstein gradient flow. By only (de)-regularizing with Σ_π , DrMMD approaches the χ^2 -divergence in the limit, a crucial property that is exploited in the proofs of the convergence results of Theorems 4.1 and 5.1.

8. Experiments

In this section, we demonstrate the superior empirical performance of the proposed DrMMD descent in various experimental settings.

8.1 Three ring experiment

We follow the experimental set-up in Glaser et al. (2021) in which the target distribution π (•) is defined on a manifold in \mathbb{R}^2 consisting of three non-overlapping rings. The initial source distribution μ_0 (•) is a Gaussian distribution close to the vicinity of the first ring. In this setting, all f -divergence gradient flows, including Langevin diffusion, are ill-defined because the target π is not absolutely continuous with respect to the initial source μ_0 . Nevertheless, we will simulate χ^2 flow with an existing implementation of Liu et al. (2024a) that estimates the velocity field with a local linear estimator as one of the baseline methods. In contrast, kernel-based gradient flows like MMD, KALE, and DrMMD gradient flows are well-defined in this setting, and are also used as baseline methods for comparison.

We sample $N = M = 300$ samples from the initial source and the target distributions and run DrMMD descent with adaptive λ for $n_{max} = 100,000$ iterations, at which point all methods have converged. As in Glaser et al. (2021), we use a Gaussian kernel $k(x, x') = \exp(-0.5\|x - x'\|^2/l^2)$ with bandwidth $l = 0.3$. The step size for MMD descent is $\gamma = 10^{-2}$ and the step size for KALE and DrMMD descent is $\gamma = 10^{-3}$. We enforce a positive lower bound $\tilde{\lambda} = 10^{-3}$ for numerical stability and the regularity hyperparameter r is optimized over the set of $\{0.1, 0.5, 1.0\}$.

From Figure 2 **Left** and **Middle**, we can see that DrMMD descent outperforms MMD, KALE, and χ^2 descent in terms of all dissimilarity metrics with respect to the target π : MMD and Wasserstein-2 distance. Figure 1 is an animation plot visualizing the evolution of




Figure 1: Animation of MMD, KALE, χ^2 and DrMMD gradient descent on the Three-ring dataset.

particles under these descent schemes, which demonstrates that both KALE and DrMMD descent are sensitive to the mismatch of support and stay concentrated in the support of the target π , while particles of MMD descent can diffuse outside the support of π . Note that " χ^2 " denotes an alternate estimate of the χ^2 divergence due to Liu et al. (2024a): being an f -divergence, we would expect " χ^2 descent" to match the support of the target (as in KALE and DrMMD). This is not the case due to bias in the velocity field being learned from samples. Compared to KALE descent, DrMMD descent does not suffer from the numerical approximation error of the optimization routine when solving the velocity field of KALE, which explains its improved performance.

8.2 Gradient flow for training student/teacher networks

Next, we consider a large-scale setting following Arbel et al. (2019), where a student network is trained to imitate the outputs of a teacher network. We consider a two-layer neural network of the form

$$\psi(z, x) = G(b^1 + W^1 \sigma(W^0 z + b^0)),$$

where σ is the ReLU non-linearity and x is the concatenation of all network parameters $(b^1, W^1, b^0, W^0) \in \mathbb{R}^d$. G is an element-wise non-linear function $G: \mathbb{R} \rightarrow \mathbb{R}, x \mapsto \exp(-\frac{1}{4}x^2)$. The teacher network is of the form: $\Psi^T(z, \pi) = \int \psi(z, x) d\pi(x)$ where π denotes the teacher distribution, and the student network is $\Psi^S(z, \mu) = \int \psi(z, x) d\mu(x)$ where μ denotes the student distribution. Here we consider Gaussian distributed μ and π for simplicity. The




Figure 2: Left & Middle: Comparison of DrMMD descent with adaptive λ , MMD and KALE descent on three-ring synthetic data in terms of MMD and Wasserstein-2 distance with respect to the target π . **Right:** Comparison of MMD descent with and without noise injection, DrMMD descent with and without noise injection on training student/teacher networks in terms of validation MMD^2 distance.

student network can imitate the behavior of the teacher network by minimizing the objective⁶

$$\min_{\mu \in \mathcal{P}_2(\mathbb{R}^d)} \mathbb{E}_{z \sim \mathbb{P}_{\text{data}}} (\Psi^T(z, \pi) - \Psi^S(z, \mu))^2, \quad (29)$$

where \mathbb{P}_{data} is the distribution of the input data. If we define the kernel as the inner product of the neural network feature maps,

$$k(x, x') = \mathbb{E}_{z \sim \mathbb{P}_{\text{data}}} [\psi(z, x)^\top \psi(z, x')],$$

then the objective of (29) can be equivalently expressed as

$$\min_{\mu \in \mathcal{P}_2(\mathbb{R}^d)} \iint_{\mathcal{X}} k(x, x') d(\pi - \mu)(x) d(\pi - \mu)(x'),$$

which is precisely the $\text{MMD}(\mu \parallel \pi)^2$ under the kernel k . Since $G(x) = \exp(-\frac{1}{4}x^2)$, the kernel is bounded and so the MMD is well-defined. Also, since the kernel k has bounded first and second-order derivatives, it satisfies the Assumption 2.

Therefore, the training of the student network with objective (29) can be treated as an optimization problem of MMD^2 distance in $\mathcal{P}_2(\mathbb{R}^d)$, i.e., as MMD gradient flow. It is shown in Arbel et al. (2019) that MMD flow and its descent scheme will generally get stuck in local optima because MMD is not *geodesically* convex; therefore, noise injection has been proposed to escape these local optima.

With support from Theorems 4.1 and Theorem 5.1 on the convergence of DrMMD flow and its descent scheme, we propose to minimize $\text{DrMMD}(\mu \parallel \pi)$ and use DrMMD descent rather than minimizing $\text{MMD}(\mu \parallel \pi)^2$ directly. Although this does not directly minimize the objective in (29), the favorable convergence performance of DrMMD descent should result in a smaller $\text{MMD}(\mu \parallel \pi)^2$ at convergence.

In our experimental setting, we are given $M = 10$ particles $x^{(1)}, \dots, x^{(M)}$ from the teacher distribution $\pi = \mathcal{N}(0, \text{I})$ and $N = 1000$ particle $y_0^{(1)}, \dots, y_0^{(N)}$ from the initial student

6. Note that our setting is slightly different from Chizat and Bach (2018) in which μ, π are measures over the hidden neurons, while our setting follows Arbel et al. (2019) in which μ, π are measures over all the network parameters.

distribution $\mu_0 = \mathcal{N}(0, 10^{-3}\mathbf{I})$. The teacher particles are fixed while the student particles are updated according to Algorithm 1 at each time step. The initial (de)-regularization parameter is $\lambda_0 = 0.1$, the step size is $\gamma = 0.1$, we apply a lower bound $\tilde{\lambda} = 10^{-3}$, and the regularity r is optimized over the set of $\{0.1, 0.5, 1.0\}$. For the architecture of the neural network, there are 3 neurons in the hidden layer, and the output dimension is 1. The data distribution \mathbb{P}_{data} is a uniform distribution on the sphere in \mathbb{R}^p with $p = 50$. 2000 data are sampled from \mathbb{P}_{data} with 1000 as training dataset and another 1000 as validation dataset. The kernel $k(x, x') = \mathbb{E}_{z \sim \mathbb{P}_{\text{data}}} [\psi(z, x)^\top \psi(z, x')]$ is estimated by the average over 100 randomly selected samples from the training dataset at each iteration. DrMMD and MMD descent stop after $n_{\text{max}} = 15,000$ iterations when both converge. The final performance is evaluated in terms of $\text{MMD}^2(\mu_{n_{\text{max}}} \parallel \pi)$ with kernel k estimated by the average of 1000 samples in the validation dataset.

In Figure 2 **Right**, we report the performance of MMD descent (with and without noise injection) along with the DrMMD descent (with and without noise injection) in terms of MMD distance on the validation dataset. We can see that the DrMMD descent does not get stuck in a local optimum, and leads to much lower validation $\text{MMD}(\mu_{n_{\text{max}}} \parallel \pi)^2$ even without noise injection. We also run DrMMD descent with the noise injection scheme and find that noise injection can further improve the performance of DrMMD descent and outperforms MMD descent with noise injection. Although it is unclear whether the density ratio has enough regularity to meet the condition of Theorem 5.1, the kernel k satisfies the boundedness and smoothness conditions of Assumption 2 and the target π satisfies the Poincaré inequality since it is Gaussian. The DrMMD descent benefits from more favorable convergence properties, which explains its superior performance.

The code to reproduce all the experiments can be found in the following GitHub repository. <https://github.com/hudsonchen/DrMMD>.

9. Discussion

In this paper, we introduced (de)-regularization of the MMD (called DrMMD) and its associated Wasserstein gradient flow. As an interpolation between the MMD and χ^2 -divergence, the DrMMD gradient flow inherits strengths from both sides: it is easy to simulate in closed form with particles, and it has an exponential rate of convergence towards the global minimum up to a controllable barrier term when the target π satisfies a Poincaré inequality. Additionally, we provide the optimal adaptive selection of a regularization coefficient that best balances the approximation and time discretization errors in DrMMD gradient descent. Our work is the first to prove theoretically that an adaptive kernel through adaptive regularization can result in improved convergence of MMD gradient flow. The theoretical results are consistent with the empirical evidence in several numerical experiments.

Following our work, there remain a number of interesting open problems. For example, (i) Since the kernel bandwidth has been known to play an important role in the performance of kernel-based algorithms, it is of interest to study the adaptive choice of kernel bandwidth in the context of DrMMD gradient flow. (ii) To generalize our convergence analysis to the Wasserstein gradient flow of all Moreau envelopes of f -divergences in reproducing kernel Hilbert space, even when they do not have a closed-form expression as DrMMD. (iii) While the current work proposes an approximation to the χ^2 -squared flow in the generative

modeling setting, i.e., where the target distribution π is known only through samples, it will be interesting to construct approximations to χ^2 -flow in the sampling setting, i.e., where π is known in closed form (at least up to normalization).

10. Proofs

In this section, we present the proofs of the results presented in Sections 3–6.

10.1 Proof of Proposition 3.1

Note that

$$\begin{aligned} \text{MMD}(\mu||\pi)^2 &= \|m_\mu - m_\pi\|_{\mathcal{H}}^2 = \left\| \int k(\cdot, x) d\mu(x) - \int k(\cdot, x) d\pi(x) \right\|_{\mathcal{H}}^2 \\ &= \left\| \int k(\cdot, x) \left(\frac{d\mu}{d\pi}(x) - 1 \right) d\pi(x) \right\|_{\mathcal{H}}^2 = \left\| \iota_\pi^* \left(\frac{d\mu}{d\pi} - 1 \right) \right\|_{\mathcal{H}}^2 \\ &= \left\langle \mathcal{T}_\pi \left(\frac{d\mu}{d\pi} - 1 \right), \frac{d\mu}{d\pi} - 1 \right\rangle_{L^2(\pi)} = \left\| \mathcal{T}_\pi^{\frac{1}{2}} \left(\frac{d\mu}{d\pi} - 1 \right) \right\|_{L^2(\pi)}^2. \end{aligned}$$

Also recall that the χ^2 -divergence between μ and π is

$$\chi^2(\mu||\pi) = \int \left(\frac{d\mu}{d\pi} - 1 \right)^2 d\pi = \left\| \text{I} \left(\frac{d\mu}{d\pi} - 1 \right) \right\|_{L^2(\pi)}^2.$$

10.2 Proof of Proposition 3.2

Let $\mu \ll \pi$. In order to prove the alternative form of DrMMD in (4), we start from (4) and show that it recovers (3).

$$\begin{aligned} \text{DrMMD}(\mu||\pi) &= (1 + \lambda) \left\| (\Sigma_\pi + \lambda \text{I})^{-\frac{1}{2}} (m_\mu - m_\pi) \right\|_{\mathcal{H}}^2 \\ &= (1 + \lambda) \left\| (\iota_\pi^* \iota_\pi + \lambda \text{I})^{-\frac{1}{2}} \iota_\pi^* \left(\frac{d\mu}{d\pi} - 1 \right) \right\|_{\mathcal{H}}^2 \\ &= (1 + \lambda) \left\langle (\iota_\pi^* \iota_\pi + \lambda \text{I})^{-\frac{1}{2}} \iota_\pi^* \left(\frac{d\mu}{d\pi} - 1 \right), (\iota_\pi^* \iota_\pi + \lambda \text{I})^{-\frac{1}{2}} \iota_\pi^* \left(\frac{d\mu}{d\pi} - 1 \right) \right\rangle_{\mathcal{H}} \\ &= (1 + \lambda) \left\langle \iota_\pi (\iota_\pi^* \iota_\pi + \lambda \text{I})^{-1} \iota_\pi^* \left(\frac{d\mu}{d\pi} - 1 \right), \frac{d\mu}{d\pi} - 1 \right\rangle_{L^2(\pi)} \\ &= (1 + \lambda) \left\langle \iota_\pi \iota_\pi^* (\iota_\pi \iota_\pi^* + \lambda \text{I})^{-1} \left(\frac{d\mu}{d\pi} - 1 \right), \frac{d\mu}{d\pi} - 1 \right\rangle_{L^2(\pi)}, \end{aligned} \tag{30}$$

where the last equality follows by noticing $\iota_\pi (\iota_\pi^* \iota_\pi + \lambda \text{I})^{-1} \iota_\pi^* = \iota_\pi \iota_\pi^* (\iota_\pi \iota_\pi^* + \lambda \text{I})^{-1}$. Therefore,

$$\begin{aligned} \text{DrMMD}(\mu||\pi) &= (1 + \lambda) \left\langle \mathcal{T}_\pi (\mathcal{T}_\pi + \lambda \text{I})^{-1} \left(\frac{d\mu}{d\pi} - 1 \right), \frac{d\mu}{d\pi} - 1 \right\rangle_{L^2(\pi)} \\ &= (1 + \lambda) \left\| \left((\mathcal{T}_\pi + \lambda \text{I})^{-1} \mathcal{T}_\pi \right)^{1/2} \left(\frac{d\mu}{d\pi} - 1 \right) \right\|_{L^2(\pi)}^2, \end{aligned}$$

which follows from the positivity and self-adjointness of $\mathcal{T}_\pi (\mathcal{T}_\pi + \lambda \mathbf{I})^{-1}$. So (4) is proved. Next, we are going to prove the variational formulation in (5). Similarly, we start from (5) and show it recovers (3). Consider

$$\begin{aligned}
& (1 + \lambda) \sup_{h \in \mathcal{H}} \left\{ \int h d\mu - \int \left(\frac{h^2}{4} + h \right) d\pi - \frac{\lambda}{4} \|h\|_{\mathcal{H}}^2 \right\} \\
&= -(1 + \lambda) \inf_{h \in \mathcal{H}} \left\{ \int \frac{h^2}{4} d\pi - \int h(d\mu - d\pi) + \frac{\lambda}{4} \|h\|_{\mathcal{H}}^2 \right\} \\
&= -(1 + \lambda) \inf_{h \in \mathcal{H}} \left\{ \frac{1}{4} \langle h, \Sigma_\pi h \rangle_{\mathcal{H}} - \langle h, m_\mu - m_\pi \rangle_{\mathcal{H}} + \frac{\lambda}{4} \|h\|_{\mathcal{H}}^2 \right\} \\
&= -(1 + \lambda) \inf_{h \in \mathcal{H}} \left\{ \left\| \left(\frac{1}{4} \Sigma_\pi + \frac{\lambda}{4} \mathbf{I} \right)^{1/2} h - \frac{1}{2} \left(\frac{1}{4} \Sigma_\pi + \frac{\lambda}{4} \mathbf{I} \right)^{-1/2} (m_\mu - m_\pi) \right\|_{\mathcal{H}}^2 \right\} \\
&\quad + \frac{(1 + \lambda)}{4} \left\| \left(\frac{1}{4} \Sigma_\pi + \frac{\lambda}{4} \mathbf{I} \right)^{-1/2} (m_\mu - m_\pi) \right\|_{\mathcal{H}}^2. \tag{31}
\end{aligned}$$

The last equality follows from completing the squares, based on which it is easy to see that the infimum is achieved at $h_{\mu, \pi}^* = 2 (\Sigma_\pi + \lambda \mathbf{I})^{-1} (m_\mu - m_\pi)$. For $\mu \ll \pi$, following the same derivations in (30), $h_{\mu, \pi}^*$ can be alternatively expressed as

$$h_{\mu, \pi}^* = 2 (\mathcal{T}_\pi + \lambda \mathbf{I})^{-1} \mathcal{T}_\pi \left(\frac{d\mu}{d\pi} - 1 \right) \tag{32}$$

Plugging $h_{\mu, \pi}^*$ back into (31) recovers (3), so (5) is proved.

10.3 DrMMD is MMD with a regularized kernel \tilde{k}

Given the definition of $\tilde{k}(x, x') = \left\langle (\Sigma_\pi + \lambda \mathbf{I})^{-\frac{1}{2}} k(\cdot, x), (\Sigma_\pi + \lambda \mathbf{I})^{-\frac{1}{2}} k(\cdot, x') \right\rangle_{\mathcal{H}}$, it is clear that \tilde{k} is symmetric and positive definite so it has a unique associated reproducing kernel Hilbert space $\tilde{\mathcal{H}}$ (Steinwart and Christmann, 2008, Theorem 4.21) with canonical feature map $\tilde{k}(x, \cdot)$. Therefore,

$$\begin{aligned}
& \text{DrMMD}(\mu \| \pi) \\
&= (1 + \lambda) \left\| (\Sigma_\pi + \lambda \mathbf{I})^{-\frac{1}{2}} (m_\mu - m_\pi) \right\|_{\mathcal{H}}^2 \\
&= (1 + \lambda) \left\langle (\Sigma_\pi + \lambda \mathbf{I})^{-\frac{1}{2}} \int k(x, \cdot) d(\pi - \mu)(x), (\Sigma_\pi + \lambda \mathbf{I})^{-\frac{1}{2}} \int k(x', \cdot) d(\pi - \mu)(x') \right\rangle_{\mathcal{H}} \\
&= (1 + \lambda) \left\langle \int (\Sigma_\pi + \lambda \mathbf{I})^{-\frac{1}{2}} k(x, \cdot) d(\pi - \mu)(x), \int (\Sigma_\pi + \lambda \mathbf{I})^{-\frac{1}{2}} k(x', \cdot) d(\pi - \mu)(x') \right\rangle_{\mathcal{H}} \\
&= (1 + \lambda) \iint \left\langle (\Sigma_\pi + \lambda \mathbf{I})^{-\frac{1}{2}} k(x, \cdot), (\Sigma_\pi + \lambda \mathbf{I})^{-\frac{1}{2}} k(x', \cdot) \right\rangle_{\mathcal{H}} d(\pi - \mu)(x) d(\pi - \mu)(x') \\
&= (1 + \lambda) \iint \tilde{k}(x, x') d(\pi - \mu)(x) d(\pi - \mu)(x') \\
&= (1 + \lambda) \left\| \int \tilde{k}(x, \cdot) d(\mu - \pi)(x) \right\|_{\tilde{\mathcal{H}}}^2.
\end{aligned}$$

In the third and fourth equality above, we are using the fact that $(\Sigma_\pi + \lambda I)^{-\frac{1}{2}} k(x, \cdot) \in \mathcal{H}$ is Bochner integrable and the Bochner integral preserves inner product structure. So DrMMD is essentially MMD^2 with a different kernel \tilde{k} up to a multiplicative factor of $1 + \lambda$.

Next, we present the Mercer decomposition of \tilde{k} . Notice that $\{e_i\}_{i \geq 1}$ are the eigenfunctions of \mathcal{T}_π , so $\{\sqrt{\varrho_i} e_i\}_{i \geq 1}$ are the eigenfunctions of Σ_π . For x and x' in the support of π , \tilde{k} also enjoys a pointwise convergent Mercer decomposition

$$\begin{aligned} \tilde{k}(x, x') &= \left\langle (\Sigma_\pi + \lambda I)^{-\frac{1}{2}} k(\cdot, x), (\Sigma_\pi + \lambda I)^{-\frac{1}{2}} k(\cdot, x') \right\rangle_{\mathcal{H}} \\ &= \left\langle (\Sigma_\pi + \lambda I)^{-\frac{1}{2}} \left(\sum_{i \geq 1} \varrho_i e_i(x) e_i \right), (\Sigma_\pi + \lambda I)^{-\frac{1}{2}} \left(\sum_{i \geq 1} \varrho_i e_i(x') e_i \right) \right\rangle_{\mathcal{H}} \\ &= \left\langle \sum_{i \geq 1} \frac{\varrho_i}{\sqrt{\varrho_i + \lambda}} e_i(x) e_i, \sum_{i \geq 1} \frac{\varrho_i}{\sqrt{\varrho_i + \lambda}} e_i(x') e_i \right\rangle_{\mathcal{H}} \\ &= \sum_{i \geq 1} \frac{\varrho_i}{\varrho_i + \lambda} e_i(x) e_i(x'). \end{aligned} \tag{33}$$

More properties of the regularized kernel \tilde{k} are provided in Lemma B.3.

10.4 Proof of Proposition 3.3

Given that $\frac{d\mu}{d\pi} - 1 \in L^2(\pi)$, so

$$\text{DrMMD}(\mu \| \pi) = (1 + \lambda) \left\| \left((\mathcal{T}_\pi + \lambda I)^{-1} \mathcal{T}_\pi \right)^{1/2} \left(\frac{d\mu}{d\pi} - 1 \right) \right\|_{L^2(\pi)}^2 \leq (1 + \lambda) \left\| \frac{d\mu}{d\pi} - 1 \right\|_{L^2(\pi)}^2,$$

which is finite for any $\lambda \geq 0$. We are allowed to interchange the limit and integration according to the dominated convergence theorem (Rudin, 1976) to achieve,

$$\begin{aligned} \lim_{\lambda \rightarrow 0} \text{DrMMD}(\mu \| \pi) &= \left\| \lim_{\lambda \rightarrow 0} \left((\mathcal{T}_\pi + \lambda I)^{-1} \mathcal{T}_\pi \right)^{1/2} \left(\frac{d\mu}{d\pi} - 1 \right) \right\|_{L^2(\pi)}^2 \\ &= \left\| \frac{d\mu}{d\pi} - 1 \right\|_{L^2(\pi)}^2 = \chi^2(\mu \| \pi). \end{aligned}$$

From (4), we have that,

$$\begin{aligned} \sqrt{\text{DrMMD}(\mu \| \pi)} &= \sqrt{1 + \lambda} \left\| (\Sigma_\pi + \lambda I)^{-\frac{1}{2}} (m_\mu - m_\pi) \right\|_{\mathcal{H}} \\ &\leq \sqrt{1 + \lambda} \left\| (\Sigma_\pi + \lambda I)^{-\frac{1}{2}} \right\|_{op} \|m_\mu - m_\pi\|_{\mathcal{H}} \leq \sqrt{\frac{1 + \lambda}{\lambda}} \times \text{MMD}(\mu, \pi), \end{aligned} \tag{34}$$

where the last inequality follows by noticing that $\Sigma_\pi = \iota_\pi^* \iota_\pi$ shares the same eigenvalues as $\mathcal{T}_\pi = \iota_\pi \iota_\pi^*$, and hence the eigenvalues of $(\Sigma_\pi + \lambda I)^{-1}$ are $(\varrho_i + \lambda)^{-1}$ which all smaller than $\frac{1}{\lambda}$. Therefore, $\text{DrMMD}(\mu \| \pi) \leq \frac{1 + \lambda}{\lambda} \text{MMD}^2(\mu \| \pi)$.

On the other hand, using Lemma A.10 from (Hagrass et al., 2024a), we have

$$\text{MMD}(\mu \| \pi) = \|m_\mu - m_\pi\|_{\mathcal{H}} \leq \|(\Sigma_\pi + \lambda I)^{\frac{1}{2}}\|_{op} \left\| (\Sigma_\pi + \lambda I)^{-\frac{1}{2}} (m_\mu - m_\pi) \right\|_{\mathcal{H}}$$

$$= \sqrt{\frac{1}{1+\lambda}} \|(\Sigma_\pi + \lambda I)^{\frac{1}{2}}\|_{op} \sqrt{\text{DrMMD}(\mu\|\pi)} \leq \sqrt{\frac{K+\lambda}{1+\lambda}} \sqrt{\text{DrMMD}(\mu\|\pi)}, \quad (35)$$

where K is the upper bound on kernel k in Assumption 2 and hence an upper bound on the operator norm of Σ_π . Combining (34) and (35), we have

$$\frac{1+\lambda}{K+\lambda} \text{MMD}^2(\mu\|\pi) \leq \text{DrMMD}(\mu\|\pi) \leq \frac{1+\lambda}{\lambda} \text{MMD}^2(\mu\|\pi).$$

Therefore, $\lim_{\lambda \rightarrow \infty} \text{DrMMD}(\mu\|\pi) = \text{MMD}^2(\mu\|\pi)$, and the proposition is proved.

10.5 Proof of Proposition 3.4

To show that DrMMD is a probability divergence, we need to show that DrMMD enjoys non-negativity and definiteness. It is easy to see that $\text{DrMMD}(\mu\|\pi)$ is non-negative from its definition in Definition 1. Then, we prove definiteness, i.e., $\text{DrMMD}(\mu\|\pi) = 0$ if and only if $\mu = \pi$. For the first direction, assume $\text{DrMMD}(\mu\|\pi) = 0$, so $\|(\Sigma_\pi + \lambda I)^{-1/2} (m_\mu - m_\pi)\|_{\mathcal{H}}^2 = 0$. Since $(\Sigma_\pi + \lambda I)^{-1/2}$ is a non-singular operator, we must have that $m_\mu = m_\pi$ which implies $\mu = \pi$ as k is c_0 -universal and hence characteristic (Sriperumbudur et al., 2011). For the other direction, when $\mu = \pi$, immediately we can see $\text{DrMMD}(\mu\|\pi) = 0$.

Then we prove that DrMMD metrizes weak convergence. For the first direction, we know from (34) that $\text{DrMMD}(\mu_n\|\pi) \leq \frac{1+\lambda}{\lambda} \text{MMD}^2(\mu_n\|\pi)$ and $\text{MMD}^2(\mu_n\|\pi) \rightarrow 0$ as μ_n converges weakly to π (Simon-Gabriel et al., 2023). For the converse direction, we assume that $\text{DrMMD}(\mu_n\|\pi) \rightarrow 0$. From (35), we know that $\text{MMD}^2(\mu\|\pi) \leq \frac{K+\lambda}{1+\lambda} \text{DrMMD}(\mu\|\pi)$, therefore $\text{DrMMD}(\mu_n\|\pi) \rightarrow 0$ implies $\text{MMD}(\mu_n\|\pi) \rightarrow 0$, implying the weak convergence of μ_n to π , if k is characteristic (Simon-Gabriel et al., 2023).

10.6 Proof of Proposition 4.1

In order to show that $\mathcal{F}_{\text{DrMMD}}(\cdot) = \text{DrMMD}(\cdot\|\pi)$ admits a well-defined gradient flow, we follow the same techniques in Proposition 7 of Glaser et al. (2021) and Lemma B.2 of Chizat and Bach (2018), where the key is to show that $(1+\lambda)\nabla h_{\mu,\pi}^*$ is the Fréchet subdifferential of $\mathcal{F}_{\text{DrMMD}}$ evaluated at μ .⁷ According to Definition 10.1.1 of Ambrosio et al. (2005), it is equivalent to prove that, for any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ and $\phi \in C_c^\infty(\mathbb{R}^d)$,

$$\begin{aligned} \text{DrMMD}((I + \nabla\phi)_\# \mu\|\pi) - \text{DrMMD}(\mu\|\pi) &\geq (1+\lambda) \int \nabla\phi(x)^\top \nabla h_{\mu,\pi}^*(x) d\mu(x) \\ &\quad + o(\|\nabla\phi\|_{L^2(\mu)}). \end{aligned} \quad (36)$$

Define $\rho_t = (I + t\nabla\phi)_\# \mu$, $\varphi_t : \mathbb{R}^d \rightarrow \mathbb{R}^d, x \mapsto x + t\nabla\phi(x)$ and $g(t) = \text{DrMMD}(\rho_t\|\pi)$. Then from Lemma B.6 we know that $g(t)$ is continuous and differentiable with respect to t and

$$\left. \frac{d}{dt} \right|_{t=0} g(t) = (1+\lambda) \int \nabla\phi(x)^\top \nabla h_{\mu,\pi}^*(x) d\mu(x).$$

7. Although DrMMD can be viewed as squared MMD with a regularized kernel \tilde{k} , we are not using the technique in Arbel et al. (2019) because it relies on Lemma 10.4.1 of Ambrosio et al. (2005) which only provides the Fréchet subdifferential on probability measures μ that admit density functions. To construct the Wasserstein gradient flow of $\mathcal{F}_{\text{DrMMD}}$ up to full generality, we resort to the techniques of Glaser et al. (2021) and Chizat and Bach (2018) instead.

Since $t \mapsto g(t)$ is differentiable, using Taylor's theorem and mean value theorem (Rudin, 1976), we know that there exists $0 < \kappa < 1$ such that

$$\text{DrMMD}((I + \nabla\phi)_{\#}\mu|\pi) - \text{DrMMD}(\mu|\pi) = g(1) - g(0) = \frac{d}{dt}\Big|_{t=0} g(t) + \frac{d^2}{dt^2}\Big|_{t=\kappa} g(t).$$

Therefore, to prove (36), the goal is to prove that $\frac{d^2}{dt^2}\Big|_{t=\kappa} g(t) \geq o(\|\nabla\phi\|_{L^2(\mu)})$. To this end, since we know from Lemma B.6 that $t \mapsto \frac{d}{dt}g(t)$ is continuous and differentiable, we have

$$\begin{aligned} \frac{d^2}{dt^2} \text{DrMMD}(\rho_t|\pi) &= 2(1 + \lambda) \iint \nabla\phi(x)^\top \nabla_1 \nabla_2 \tilde{k}(\varphi_t(x), \varphi_t(y)) \nabla\phi(y) d\mu(x) d\mu(y) \\ &+ 2(1 + \lambda) \int \nabla\phi(x)^\top \left(\int \mathbf{H}_1 \tilde{k}(\varphi_t(x), \varphi_t(y)) d\mu(y) - \int \mathbf{H}_1 \tilde{k}(\varphi_t(x), y) d\pi(y) \right) \nabla\phi(x) d\mu(x), \end{aligned}$$

where \tilde{k} is the regularized kernel defined in (6) and $\tilde{\mathcal{H}}$ is the associated RKHS. Using Lemma B.3, the first term above can be upper-bounded by,

$$\begin{aligned} &\left| 2(1 + \lambda) \iint \nabla\phi(x)^\top \nabla_1 \nabla_2 \tilde{k}(\varphi_t(x), \varphi_t(y)) \nabla\phi(y) d\mu(x) d\mu(y) \right| \\ &\leq 2(1 + \lambda) \iint \|\nabla\phi(x)\| \left\| \nabla_1 \nabla_2 \tilde{k}(\varphi_t(x), \varphi_t(y)) \right\|_F \|\nabla\phi(y)\| d\mu(x) d\mu(y) \\ &\leq 2(1 + \lambda) \frac{K_{1d}}{\lambda} \left(\int \|\nabla\phi(y)\| d\mu(y) \right)^2 \leq 2(1 + \lambda) \frac{K_{1d}}{\lambda} \|\nabla\phi\|_{L^2(\mu)}^2. \end{aligned}$$

Using Lemma B.3 again, the second term can be upper-bounded by

$$\begin{aligned} &\left| \int \nabla\phi(x)^\top \left(\int \mathbf{H}_1 \tilde{k}(\varphi_t(x), \varphi_t(y)) d\mu(y) - \int \mathbf{H}_1 \tilde{k}(\varphi_t(x), y) d\pi(y) \right) \nabla\phi(x) d\mu(x) \right| \\ &\leq \left| \int \nabla\phi(x)^\top \left(\int \mathbf{H}_1 \tilde{k}(\varphi_t(x), \varphi_t(y)) d\mu(y) \right) \nabla\phi(x) d\mu(x) \right| \\ &\quad + \left| \int \nabla\phi(x)^\top \left(\int \mathbf{H}_1 \tilde{k}(\varphi_t(x), y) d\pi(y) \right) \nabla\phi(x) d\mu(x) \right| \\ &\leq \int \|\nabla\phi(x)\| \left(\int \left\| \mathbf{H}_1 \tilde{k}(\varphi_t(x), \varphi_t(y)) \right\|_F d\mu(y) \right) \|\nabla\phi(x)\| d\mu(x) \\ &\quad + \int \|\nabla\phi(x)\| \left(\int \left\| \mathbf{H}_1 \tilde{k}(\varphi_t(x), y) \right\|_F d\pi(y) \right) \|\nabla\phi(x)\| d\mu(x) \\ &\leq 2 \frac{\sqrt{K K_{2d}}}{\lambda} \|\nabla\phi\|_{L^2(\mu)}^2. \end{aligned}$$

Combining the above two inequalities to have

$$\left| \frac{d^2}{dt^2} \text{DrMMD}(\rho_t|\pi) \right| \leq 2(1 + \lambda) \frac{K_{1d}}{\lambda} \|\nabla\phi\|_{L^2(\mu)}^2 + 4(1 + \lambda) \frac{\sqrt{K K_{2d}}}{\lambda} \|\nabla\phi\|_{L^2(\mu)}^2. \quad (37)$$

Therefore, we have $\frac{d^2}{dt^2}\Big|_{t=\kappa} g(t) = \mathcal{O}(\|\nabla\phi\|_{L^2(\mu)}^2) = o(\|\nabla\phi\|_{L^2(\mu)})$ as $\|\nabla\phi\|_{L^2(\mu)} \rightarrow 0$. So (36) is proved, which means that $(1 + \lambda)\nabla h_{\mu, \pi}^*(x)$ is the Fréchet subdifferential of $\mathcal{F}_{\text{DrMMD}}$

evaluated at μ . According to Definition 11.1.1 of Ambrosio et al. (2005), there exists a solution $(\mu_t)_{t \geq 0}$ such that the following equation holds in the sense of distributions,

$$\partial_t \mu_t - (1 + \lambda) \nabla \cdot (\mu_t \nabla h_{\mu_t, \pi}^*) = 0,$$

and such $(\mu_t)_{t \geq 0}$ is indeed the DrMMD gradient flow, so existence is proved.

Next, we are going to prove uniqueness. (37) indicates that $\mathcal{F}_{\text{DrMMD}}$ is geodesically $-(1 + \lambda) \frac{4\sqrt{KK_{2d} + 2K_{1d}}}{\lambda}$ semiconvex. Therefore the uniqueness of $(\mu_t)_t$ follows from Theorem 11.2.1 of Ambrosio et al. (2005). (8) is proved in (37).

10.7 Proof of Proposition 4.2

Given $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ and $\phi \in C_c^\infty(\mathbb{R}^d)$, consider the path $(\rho_t)_{0 \leq t \leq 1}$ from μ to $(I + \nabla \phi)_\# \mu$ given by $\rho_t = (I + t \nabla \phi)_\# \mu$. $(\rho_t)_{0 \leq t \leq 1}$ is a constant-time geodesic in the Wasserstein-2 space by construction (Appendix A). Define $\varphi_t : \mathbb{R}^d \rightarrow \mathbb{R}^d, x \mapsto x + t \nabla \phi(x)$. We know from (Villani et al., 2009, Example 15.9) (by taking $m = 2$) that,

$$\begin{aligned} \frac{d^2}{dt^2} \Big|_{t=0} \chi^2(\rho_t \| \pi) &= \int \frac{d\mu}{d\pi}(x) \left(\nabla V(x)^\top \nabla \phi(x) - \nabla \cdot \nabla \phi(x) \right)^2 d\mu(x) \\ &\quad + \int \frac{d\mu}{d\pi}(x) \nabla \phi(x)^\top \mathbf{H} V(x) \nabla \phi(x) d\mu(x) \\ &\quad + \int \frac{d\mu}{d\pi}(x) \|\mathbf{H} \phi(x)\|_F^2 d\mu(x). \end{aligned} \quad (38)$$

V is twice differentiable, so ∇V as a function from \mathbb{R}^d to \mathbb{R}^d is continuous. $\phi \in C_c^\infty(\mathbb{R}^d)$ has compact support, so $x \mapsto \nabla V(x)^\top \nabla \phi(x) - \nabla \cdot \nabla \phi(x)$ is a continuous function over a compact domain, so its image is also compact and hence bounded. Using similar arguments, $x \mapsto \nabla \phi(x)^\top \mathbf{H} V(x) \nabla \phi(x)$ and $x \mapsto \|\mathbf{H} \phi(x)\|_F^2$ are also continuous functions over compact domains, so they are all bounded.

Since $\frac{d\mu}{d\pi} - 1 \in \mathcal{H} \subset L^2(\pi)$ and $\int \frac{d\mu}{d\pi}(x) d\mu(x) = \|\frac{d\mu}{d\pi} - 1\|_{L^2(\pi)}^2 + 1 < \infty$, we have

$$\frac{d^2}{dt^2} \Big|_{t=0} \chi^2(\rho_t \| \pi) < \infty. \quad (39)$$

Next, from Lemma B.9 we know that $\frac{d^2}{dt^2} \Big|_{t=0} \chi^2(\rho_t \| \pi)$ can be alternatively expressed as,

$$\begin{aligned} \frac{d^2}{dt^2} \Big|_{t=0} \chi^2(\rho_t \| \pi) &= 2 \int \left(\nabla \cdot (\nabla \phi(x) \mu(x)) \frac{1}{\pi(x)} \right)^2 \pi(x) dx \\ &\quad + 2 \int \nabla \phi(x)^\top \mathbf{H} \left(\frac{\mu(x)}{\pi(x)} \right) \nabla \phi(x) \mu(x) dx. \end{aligned} \quad (40)$$

Recall from Lemma B.6 that

$$\begin{aligned} &\frac{d^2}{dt^2} \Big|_{t=0} \text{DrMMD}(\rho_t \| \pi) \\ &= 2(1 + \lambda) \iint \nabla \phi(x)^\top \nabla_1 \nabla_2 \tilde{k}(x, y) \nabla \phi(y) d\mu(x) d\mu(y) \end{aligned}$$

$$+ 2(1 + \lambda) \int \nabla \phi(x)^\top \left(\int \mathbf{H}_1 \tilde{k}(x, y) d\mu(y) - \int \mathbf{H}_1 \tilde{k}(x, y) d\pi(y) \right) \nabla \phi(x) d\mu(x). \quad (41)$$

To prove (9), our aim then is to compare and bound the difference of $\frac{d^2}{dt^2}|_{t=0} \chi^2(\rho_t || \pi)$ in (40) and $\frac{1}{1+\lambda} \frac{d^2}{dt^2}|_{t=0} \text{DrMMD}(\rho_t || \pi)$ in (41), so we compare and bound their first and second term separately.

The first term of (40) can be rewritten as

$$\begin{aligned} & \int \left(\nabla \cdot (\nabla \phi(x) \mu(x)) \frac{1}{\pi(x)} \right)^2 \pi(x) dx = \sum_{i \geq 1} \left\langle \nabla \cdot (\nabla \phi \mu) \frac{1}{\pi}, e_i \right\rangle_{L^2(\pi)}^2 \\ & = \sum_{i \geq 1} \left(\int \nabla \cdot (\nabla \phi(x) \mu(x)) e_i(x) dx \right)^2 = \sum_{i \geq 1} \left(\int \nabla \phi(x)^\top \nabla e_i(x) \mu(x) dx \right)^2, \end{aligned} \quad (42)$$

where we use integration by parts in the last line since $\phi \in C_c^\infty(\mathbb{R}^d)$. And the first term of (41) after rescaling by $\frac{1}{1+\lambda}$ can be rewritten as,

$$\begin{aligned} & 2 \iint \nabla \phi(x)^\top \nabla_1 \nabla_2 \tilde{k}(x, y) \nabla \phi(y) d\mu(x) d\mu(y) \\ & = 2 \iint \nabla \phi(x)^\top \nabla_x \nabla_y \left(\sum_{i \geq 1} \frac{\varrho_i}{\varrho_i + \lambda} e_i(x) e_i(y) \right) \nabla \phi(y) d\mu(x) d\mu(y) \\ & = 2 \iint \left(\nabla \cdot (\nabla \phi(x) \mu(x)) \right) \left(\nabla \cdot (\nabla \phi(y) \mu(y)) \right) \sum_{i \geq 1} \frac{\varrho_i}{\varrho_i + \lambda} e_i(x) e_i(y) dx dy. \end{aligned} \quad (43)$$

Since $\pi \ll \mathcal{L}^d$ so (33) is true for all $x, y \in \mathbb{R}^d$, hence the second equality is true, and the last equality uses integration by parts since $\phi \in C_c^\infty(\mathbb{R}^d)$. Notice that

$$\begin{aligned} & \sum_{i \geq 1} \frac{\varrho_i}{\varrho_i + \lambda} \iint \left| e_i(x) e_i(y) \left(\nabla \cdot (\nabla \phi(x) \mu(x)) \right) \left(\nabla \cdot (\nabla \phi(y) \mu(y)) \right) \right| dx dy \\ & = \sum_{i \geq 1} \frac{\varrho_i}{\varrho_i + \lambda} \left(\int \left| \left(\nabla \cdot (\nabla \phi(x) \mu(x)) \right) e_i(x) \right| dx \right)^2 \\ & \leq \frac{1}{\lambda} \sum_{i \geq 1} \varrho_i \left(\int \left(\frac{\nabla \cdot (\nabla \phi(x) \mu(x))}{\mu(x)} \right)^2 \mu(x) dx \right) \left(\int e_i(x)^2 \mu(x) dx \right) \\ & = \frac{1}{\lambda} \left(\int \left(\nabla \log \mu(x)^\top \nabla \phi(x) + \nabla \cdot \nabla \phi(x) \right)^2 d\mu(x) \right) \left(\int \sum_{i \geq 1} \varrho_i e_i(x)^2 d\mu(x) \right) \\ & \leq \frac{K}{\lambda} \left(\int \left(\nabla \log \mu(x)^\top \nabla \phi(x) + \nabla \cdot \nabla \phi(x) \right)^2 d\mu(x) \right) < +\infty. \end{aligned}$$

The first inequality uses Cauchy-Schwartz, the second inequality uses $\sum_{i \geq 1} \varrho_i (e_i(x))^2 = k(x, x) \leq K$. The last quantity is finite because $x \rightarrow \nabla \log \mu(x)$ is a continuous function and $\phi \in C_c^\infty(\mathbb{R}^d)$ has compact support, hence the integral of a continuous function over a

compact domain is always finite. Then, by using Fubini's theorem (Rudin, 1976), we are allowed to interchange the infinite sum and integration of (43) to reach,

$$\begin{aligned}
& 2 \iint \left(\nabla \cdot (\nabla \phi(x) \mu(x)) \right) \left(\nabla \cdot (\nabla \phi(y) \mu(y)) \right) \sum_{i \geq 1} \frac{\varrho_i}{\varrho_i + \lambda} e_i(x) e_i(y) dx dy \\
&= 2 \sum_{i \geq 1} \frac{\varrho_i}{\varrho_i + \lambda} \iint \left(\nabla \cdot (\nabla \phi(x) \mu(x)) \right) \left(\nabla \cdot (\nabla \phi(y) \mu(y)) \right) e_i(x) e_i(y) dx dy \\
&= 2 \sum_{i \geq 1} \frac{\varrho_i}{\varrho_i + \lambda} \left(\int \left(\nabla \cdot (\nabla \phi(x) \mu(x)) \right) e_i(x) dx \right)^2 \\
&= 2 \sum_{i \geq 1} \frac{\varrho_i}{\varrho_i + \lambda} \left(\int \nabla \phi(x)^\top \nabla e_i(x) \mu(x) dx \right)^2,
\end{aligned}$$

where the last equality uses integration by parts.

So the difference between the first term of (40) and (41) rescaled by $\frac{1}{1+\lambda}$ is,

$$\begin{aligned}
& 2 \left| \sum_{i \geq 1} \left(\int \nabla \phi(x)^\top \nabla e_i(x) \mu(x) dx \right)^2 - \sum_{i \geq 1} \frac{\varrho_i}{\varrho_i + \lambda} \left(\int \nabla \phi(x)^\top \nabla e_i(x) \mu(x) dx \right)^2 \right| \\
&= 2 \sum_{i \geq 1} \frac{\lambda}{\varrho_i + \lambda} \left(\int \nabla \phi(x)^\top \nabla e_i(x) \mu(x) dx \right)^2.
\end{aligned} \tag{44}$$

Now we turn to the second term. The second term of (40) can be rewritten as

$$\begin{aligned}
& 2 \int \nabla \phi(x)^\top \mathbf{H} \left(\frac{d\mu}{d\pi}(x) \right) \nabla \phi(x) \mu(x) dx \\
&= 2 \int \nabla \phi(x)^\top \mathbf{H} \left(\frac{d\mu}{d\pi}(x) - 1 \right) \nabla \phi(x) \mu(x) dx \\
&= 2 \int \nabla \phi(x)^\top \mathbf{H} \left(\sum_{i \geq 1} \left\langle \frac{d\mu}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} e_i(x) \right) \nabla \phi(x) \mu(x) dx,
\end{aligned} \tag{45}$$

and the second term of (41) rescaled by $\frac{1}{1+\lambda}$ can be rewritten as,

$$\begin{aligned}
& 2 \int \nabla \phi(x)^\top \left(\int \mathbf{H}_1 \tilde{k}(x, y) d\mu(y) - \int \mathbf{H}_1 \tilde{k}(x, y) d\pi(y) \right) \nabla \phi(x) d\mu(x) \\
&= 2 \int \nabla \phi(x)^\top \mathbf{H} \left(\int \tilde{k}(x, y) d\mu(y) - \int \tilde{k}(x, y) d\pi(y) \right) \nabla \phi(x) d\mu(x) \\
&= 2 \int \nabla \phi(x)^\top \mathbf{H} \left(\sum_{i \geq 1} \frac{\varrho_i}{\varrho_i + \lambda} e_i(x) \int e_i(y) d(\mu - \pi)(y) \right) \nabla \phi(x) d\mu(x) \\
&= 2 \int \nabla \phi(x)^\top \mathbf{H} \left(\sum_{i \geq 1} \frac{\varrho_i}{\varrho_i + \lambda} \left\langle \frac{d\mu}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} e_i(x) \right) \nabla \phi(x) d\mu(x).
\end{aligned} \tag{46}$$

Since $\pi \ll \mathcal{L}^d$, so (33) is true for all $x, y \in \mathbb{R}^d$ hence the third equality is true. From Lemma B.3, we have $\int \tilde{k}(x, y) d\mu(y) \leq \frac{K}{\lambda}$, $x \mapsto \tilde{k}(x, y)$ is second-order differentiable, $\sup_x |\mathbf{H}_1 \tilde{k}(x, y)| \leq \frac{\sqrt{KK_{2d}}}{\lambda}$. So we are allowed to interchange integration and Hessian in the second equality using the differentiation lemma (Klenke, 2013, Theorem 6.28). Consider the difference of (46) and (45), we have

$$\begin{aligned} & 2 \left| \int \nabla \phi(x)^\top \mathbf{H} \left(\sum_{i \geq 1} \left\langle \frac{d\mu}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} e_i(x) \right) \nabla \phi(x) \mu(x) dx \right. \\ & \quad \left. - \int \nabla \phi(x)^\top \mathbf{H} \left(\sum_{i \geq 1} \frac{\varrho_i}{\varrho_i + \lambda} \left\langle \frac{d\mu}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} e_i(x) \right) \nabla \phi(x) d\mu(x) \right| \\ &= 2 \left| \int \nabla \phi(x)^\top \mathbf{H} \left(\sum_{i \geq 1} \frac{\lambda}{\varrho_i + \lambda} \left\langle \frac{d\mu}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} e_i(x) \right) \nabla \phi(x) \mu(x) dx \right|. \end{aligned} \quad (47)$$

Given $\frac{d\mu}{d\pi} - 1 \in \mathcal{H}$, there exists $q \in L^2(\pi)$ such that $\frac{d\mu}{d\pi} - 1 = \mathcal{T}_\pi^{1/2} q$ so that $\langle \frac{d\mu}{d\pi} - 1, e_i \rangle = \varrho_i^{1/2} \langle q, e_i \rangle$ for all i . For $j, r \in \{1, \dots, d\}$, we have

$$\begin{aligned} g_{M_0}(x) &:= \left| \sum_{i \geq M_0} \frac{\lambda}{\varrho_i + \lambda} \left\langle \frac{d\mu}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} \partial_j \partial_r e_i(x) \right| = \left| \sum_{i \geq M_0} \frac{\lambda}{\varrho_i + \lambda} \varrho_i^{1/2} \langle q, e_i \rangle_{L^2(\pi)} \partial_j \partial_r e_i(x) \right| \\ &\leq \left(\sum_{i \geq M_0} \left(\frac{\lambda}{\varrho_i + \lambda} \right)^2 \langle q, e_i \rangle_{L^2(\pi)}^2 \right)^{1/2} \left(\sum_{i \geq M_0} \varrho_i (\partial_j \partial_r e_i(x))^2 \right)^{1/2} \\ &\leq \left(\sum_{i \geq M_0} \langle q, e_i \rangle_{L^2(\pi)}^2 \right)^{1/2} \left(\sum_{i \geq M_0} \varrho_i (\partial_j \partial_r e_i(x))^2 \right)^{1/2} \leq \sqrt{K_{2d}} \left(\sum_{i \geq M_0} \langle q, e_i \rangle_{L^2(\pi)}^2 \right)^{1/2}. \end{aligned}$$

The final inequality holds because,

$$\begin{aligned} \sum_{j,r=1}^d \sum_{i \geq 1} \varrho_i (\partial_j \partial_r e_i(x))^2 &\leq \sum_{j,r=1}^d \sum_{i \geq 1} \varrho_i \langle \partial_j \partial_r k(x, \cdot), e_i \rangle_{\mathcal{H}}^2 = \sum_{j,r=1}^d \sum_{i \geq 1} \langle \partial_j \partial_r k(x, \cdot), \sqrt{\varrho_i} e_i \rangle_{\mathcal{H}}^2 \\ &= \sum_{j,r=1}^d \|\partial_j \partial_r k(x, \cdot)\|_{\mathcal{H}}^2 = \|\mathbf{H}_1 k(x, \cdot)\|_{\mathcal{H}^{d \times d}}^2 \leq K_{2d}. \end{aligned} \quad (48)$$

Since $\|sdq\|_{L^2(\pi)}$ is bounded, so $\sum_{i \geq M_0} \langle q, e_i \rangle_{L^2(\pi)}^2$ converges to 0 uniformly as $M_0 \rightarrow \infty$ and hence $g_{M_0}(x)$ converge to 0 uniformly. Therefore, we are allowed to interchange the Hessian and the infinite sum (Rudin, 1976) in (47) to achieve,

$$2 \left| \int \nabla \phi(x)^\top \mathbf{H} \left(\sum_{i \geq 1} \frac{\lambda}{\varrho_i + \lambda} \left\langle \frac{d\mu}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} e_i(x) \right) \nabla \phi(x) \mu(x) dx \right|$$

$$\begin{aligned}
&= 2 \left| \int \nabla \phi(x)^\top \left(\sum_{i \geq 1} \frac{\lambda}{\varrho_i + \lambda} \varrho_i^{1/2} \langle q, e_i \rangle_{L^2(\pi)} \mathbf{H} e_i(x) \right) \nabla \phi(x) \mu(x) dx \right| \\
&\leq 2 \left(\sum_{i \geq 1} \left(\frac{\lambda}{\varrho_i + \lambda} \right)^2 \langle q, e_i \rangle_{L^2(\pi)}^2 \right)^{1/2} \left(\sum_{i \geq 1} \varrho_i \left(\int \nabla \phi(x)^\top \mathbf{H} e_i(x) \nabla \phi(x) \mu(x) dx \right)^2 \right)^{1/2} \\
&\leq 2 \left(\sum_{i \geq 1} \left(\frac{\lambda}{\varrho_i + \lambda} \right)^2 \langle q, e_i \rangle_{L^2(\pi)}^2 \right)^{1/2} \left(\sum_{i \geq 1} \varrho_i \|\mathbf{H} e_i(x)\|_{op}^2 \left(\int \nabla \phi(x)^\top \nabla \phi(x) \mu(x) dx \right)^2 \right)^{1/2} \\
&\leq 2 \left(\sum_{i \geq 1} \left(\frac{\lambda}{\varrho_i + \lambda} \right)^2 \langle q, e_i \rangle_{L^2(\pi)}^2 \right)^{1/2} \left(\sum_{i \geq 1} \varrho_i \|\mathbf{H} e_i(x)\|_{op}^2 \right)^{1/2} \|\nabla \phi\|_{L^2(\mu)}^2 \\
&\leq 2 \left(\sum_{i \geq 1} \left(\frac{\lambda}{\varrho_i + \lambda} \right)^2 \langle q, e_i \rangle_{L^2(\pi)}^2 \right)^{1/2} \left(\sum_{i \geq 1} \varrho_i \|\mathbf{H} e_i(x)\|_F^2 \right)^{1/2} \|\nabla \phi\|_{L^2(\mu)}^2 \\
&\leq 2 \left(\sum_{i \geq 1} \left(\frac{\lambda}{\varrho_i + \lambda} \right)^2 \langle q, e_i \rangle_{L^2(\pi)}^2 \right)^{1/2} \sqrt{K_{2d}} \|\nabla \phi\|_{L^2(\mu)}^2. \tag{49}
\end{aligned}$$

The first inequality uses Cauchy Schwartz, the second last inequality uses that matrix operator norm is smaller than matrix Frobenius norm, and the last inequality uses (48). Combining together (44) and (49), we reach

$$\begin{aligned}
&\left| \frac{1}{1 + \lambda} \frac{d^2}{dt^2} \Big|_{t=0} \text{DrMMD}(\rho_t || \pi) - \frac{d^2}{dt^2} \Big|_{t=0} \chi^2(\rho_t || \pi) \right| \\
&\leq 2 \sum_{i \geq 1} \frac{\lambda}{\varrho_i + \lambda} \left(\int \nabla \phi(x)^\top \nabla e_i(x) \mu(x) dx \right)^2 \\
&\quad + 2 \left(\sum_{i \geq 1} \left(\frac{\lambda}{\varrho_i + \lambda} \right)^2 \langle q, e_i \rangle_{L^2(\pi)}^2 \right)^{1/2} \sqrt{K_{2d}} \|\nabla \phi\|_{L^2(\mu)}^2 \\
&=: \bar{R}(\lambda, \mu, \nabla \phi).
\end{aligned}$$

Therefore,

$$\begin{aligned}
\frac{d^2}{dt^2} \Big|_{t=0} \text{DrMMD}(\rho_t || \pi) &\geq (1 + \lambda) \frac{d^2}{dt^2} \Big|_{t=0} \chi^2(\rho_t || \pi) - (1 + \lambda) \bar{R}(\lambda, \mu, \nabla \phi) \\
&\geq (1 + \lambda) \int \frac{d\mu}{d\pi}(x) \nabla \phi(x)^\top \mathbf{H} V(x) \nabla \phi(x) d\mu(x) - (1 + \lambda) \bar{R}(\lambda, \mu, \nabla \phi) \\
&\geq (1 + \lambda) \alpha \int \frac{d\mu}{d\pi}(x) \|\nabla \phi(x)\|^2 d\mu(x) - (1 + \lambda) \bar{R}(\lambda, \mu, \nabla \phi),
\end{aligned}$$

where the second inequality is using (38) and the last inequality is using $\mathbf{H} V \succeq \alpha \mathbf{I}$. So (9) is proved.

Define $R(\lambda, \mu, \nabla\phi) := (1 + \lambda)\bar{R}(\lambda, \mu, \nabla\phi)$. The final thing to check is $\lim_{\lambda \rightarrow 0} R(\lambda, \mu, \nabla\phi) = 0$, which is equivalent to check that $\lim_{\lambda \rightarrow 0} \bar{R}(\lambda, \mu, \nabla\phi) = 0$. Since we know from (39) and (42) that

$$\sum_{i \geq 1} \left(\int \nabla\phi(x)^\top \nabla e_i(x) \mu(x) dx \right)^2 < \frac{d^2}{dt^2} \Big|_{t=0} \chi^2(\rho_t \| \pi) < \infty,$$

using the dominated convergence theorem (Rudin, 1976), we are allowed to interchange infinite sum and taking limits,

$$\begin{aligned} \lim_{\lambda \rightarrow 0} \sum_{i \geq 1} \frac{\lambda}{\varrho_i + \lambda} \left(\int \nabla\phi(x)^\top \nabla e_i(x) \mu(x) dx \right)^2 &= \sum_{i \geq 1} \lim_{\lambda \rightarrow 0} \frac{\lambda}{\varrho_i + \lambda} \left(\int \nabla\phi(x)^\top \nabla e_i(x) \mu(x) dx \right)^2 \\ &= 0. \end{aligned}$$

Similarly, because $\sum_{i \geq 1} \left(\frac{\lambda}{\varrho_i + \lambda} \right)^2 \langle q, e_i \rangle_{L^2(\pi)}^2 < \|q\|_{L^2(\pi)}^2 < \infty$, using dominated convergence theorem (Rudin, 1976) again, we have,

$$\lim_{\lambda \rightarrow 0} \sum_{i \geq 1} \left(\frac{\lambda}{\varrho_i + \lambda} \right)^2 \langle q, e_i \rangle_{L^2(\pi)}^2 = \sum_{i \geq 1} \lim_{\lambda \rightarrow 0} \left(\frac{\lambda}{\varrho_i + \lambda} \right)^2 \langle q, e_i \rangle_{L^2(\pi)}^2 = 0.$$

Therefore, we have that

$$\lim_{\lambda \rightarrow 0} R(\lambda, \mu, \nabla\phi) = \lim_{\lambda \rightarrow 0} \bar{R}(\lambda, \mu, \nabla\phi) = 0.$$

And the proof of the proposition is finished.

10.8 Proof of Theorem 4.1

Considering the time derivative of $\text{KL}(\mu_t \| \pi)$, we have

$$\begin{aligned} & \frac{d}{dt} \text{KL}(\mu_t \| \pi) \\ &= -(1 + \lambda) \int \nabla h_{\mu_t, \pi}^*(x)^\top \nabla \log \frac{d\mu_t}{d\pi}(x) \mu_t(x) dx \\ &= -(1 + \lambda) \int \nabla h_{\mu_t, \pi}^*(x)^\top \nabla \frac{d\mu_t}{d\pi}(x) \pi(x) dx \\ &= -(1 + \lambda) \int \left(\nabla h_{\mu_t, \pi}^*(x) - 2 \nabla \frac{d\mu_t}{d\pi}(x) \right)^\top \nabla \frac{d\mu_t}{d\pi}(x) \pi(x) dx - 2(1 + \lambda) \int \pi(x) \left\| \nabla \frac{d\mu_t}{d\pi}(x) \right\|^2 dx \\ &= -(1 + \lambda) \int \left(\nabla h_{\mu_t, \pi}^*(x) - 2 \nabla \left(\frac{d\mu_t}{d\pi}(x) - 1 \right) \right)^\top \nabla \frac{d\mu_t}{d\pi}(x) \pi(x) dx \\ &\quad - 2(1 + \lambda) \int \pi(x) \left\| \nabla \frac{d\mu_t}{d\pi}(x) \right\|^2 dx. \end{aligned} \tag{50}$$

Case one: $\int \pi(x) \left\| \nabla \frac{d\mu_t}{d\pi}(x) \right\|^2 dx < \infty$. We use integration by parts for the first term in (50) and we can safely ignore the boundary term due to condition 5 that for $i = 1, \dots, d$,

$\lim_{x \rightarrow \infty} \left(h_{\mu_t, \pi}^*(x) - 2 \frac{d\mu_t}{d\pi}(x) \right) \left(\partial_t \frac{d\mu_t}{d\pi}(x) \right) \pi(x) \rightarrow 0$. So, we obtain

$$\begin{aligned}
\frac{d}{dt} \text{KL}(\mu_t \| \pi) &= (1 + \lambda) \int \left(h_{\mu_t, \pi}^*(x) - 2 \left(\frac{d\mu_t}{d\pi}(x) - 1 \right) \right) \nabla \cdot \left(\pi(x) \nabla \frac{d\mu_t}{d\pi}(x) \right) dx \\
&\quad - 2(1 + \lambda) \int \pi(x) \left\| \nabla \frac{d\mu_t}{d\pi}(x) \right\|^2 dx \\
&= (1 + \lambda) \int \left(h_{\mu_t, \pi}^*(x) - 2 \left(\frac{d\mu_t}{d\pi}(x) - 1 \right) \right) \frac{\nabla \cdot \left(\pi(x) \nabla \frac{d\mu_t}{d\pi}(x) \right)}{\pi(x)} \pi(x) dx \\
&\quad - 2(1 + \lambda) \int \pi(x) \left\| \nabla \frac{d\mu_t}{d\pi}(x) \right\|^2 dx \\
&\leq (1 + \lambda) \left\| h_{\mu_t, \pi}^* - 2 \left(\frac{d\mu_t}{d\pi} - 1 \right) \right\|_{L^2(\pi)} \left\| \frac{\nabla \cdot \left(\pi \nabla \frac{d\mu_t}{d\pi} \right)}{\pi} \right\|_{L^2(\pi)} - \frac{2(1 + \lambda)}{C_P} \text{KL}(\mu_t \| \pi), \quad (51)
\end{aligned}$$

where the first part of the last inequality holds by using Cauchy Schwartz, and the second part holds by the fact that $\text{KL}(\mu_t \| \pi) \leq \chi^2(\mu_t \| \pi)$ (Van Erven and Harremos, 2014) and by applying the Poincaré inequality with $f = \frac{d\mu_t}{d\pi} - 1$ (notice that $\|\nabla f\|_{L^2(\pi)} < \infty$ from *Case one* and $\|f\|_{L^2(\pi)} < \infty$ from condition 3),

$$\int \pi(x) \left\| \nabla \frac{d\mu_t}{d\pi}(x) \right\|^2 dx \geq \frac{1}{C_P} \chi^2(\mu_t \| \pi) \geq \frac{1}{C_P} \text{KL}(\mu_t \| \pi).$$

Since $\frac{d\mu_t}{d\pi} - 1 \in \text{Ran}(\mathcal{T}_\pi^r)$ with $r > 0$, using Lemma B.5 we have

$$\left\| h_{\mu_t, \pi}^* - 2 \left(\frac{d\mu_t}{d\pi} - 1 \right) \right\|_{L^2(\pi)} \leq 2\lambda^r \|q_t\|_{L^2(\pi)}. \quad (52)$$

Then, notice that

$$\begin{aligned}
\left\| \frac{\nabla \cdot \left(\pi \nabla \frac{d\mu_t}{d\pi} \right)}{\pi} \right\|_{L^2(\pi)}^2 &= \int \frac{\left[\nabla \cdot \left(\pi(x) \nabla \frac{d\mu_t}{d\pi}(x) \right) \right]^2}{\pi(x)^2} d\pi(x) \\
&= \int \frac{\left(\nabla \pi(x)^\top \nabla \frac{d\mu_t}{d\pi}(x) + \pi(x) \nabla \cdot \nabla \frac{d\mu_t}{d\pi}(x) \right)^2}{\pi(x)^2} d\pi(x) \\
&= \int \left(\nabla \log \pi(x)^\top \nabla \frac{d\mu_t}{d\pi}(x) + \Delta \frac{d\mu_t}{d\pi}(x) \right)^2 d\pi(x) \\
&\leq \int 2 \left(\nabla \log \pi(x)^\top \nabla \frac{d\mu_t}{d\pi}(x) \right)^2 + 2 \left(\Delta \frac{d\mu_t}{d\pi}(x) \right)^2 d\pi(x) \\
&= 2 \left\| \nabla (\log \pi)^\top \nabla \left(\frac{d\mu_t}{d\pi} \right) \right\|_{L^2(\pi)}^2 + 2 \left\| \Delta \left(\frac{d\mu_t}{d\pi} \right) \right\|_{L^2(\pi)}^2 \leq 2\mathcal{J}_t^2 + 2\mathcal{I}_t^2. \quad (53)
\end{aligned}$$

Therefore, plugging (52) and (53) back to (51), we have

$$\frac{d}{dt} \text{KL}(\mu_t \| \pi) \leq 4(1 + \lambda) \lambda^r \|q_t\|_{L^2(\pi)} (\mathcal{J}_t + \mathcal{I}_t) - \frac{2(1 + \lambda)}{C_P} \text{KL}(\mu_t \| \pi). \quad (54)$$

Case two: $\int \pi(x) \left\| \nabla \frac{d\mu_t}{d\pi}(x) \right\|^2 dx = \infty$. The first term of (50) remains the same, and the second term of (50) now equals infinity which is larger than the finite $\frac{2(1+\lambda)}{C_P} \text{KL}(\mu_t \| \pi)$, so we also obtain (54) as in *Case one*.

Therefore, both *Case one* and *Case two* result in (54). Using the Gronwall lemma, we obtain that for any $T > 0$,

$$\begin{aligned} \text{KL}(\mu_T \| \pi) &\leq \exp\left(-\frac{2(1 + \lambda)}{C_P} T\right) \text{KL}(\mu_0 \| \pi) \\ &\quad + 4(1 + \lambda) \lambda^r \int_0^T \exp\left(-\frac{2(1 + \lambda)}{C_P} (T - t)\right) \|q_t\|_{L^2(\pi)} (\mathcal{J}_t + \mathcal{I}_t) dt, \end{aligned}$$

which concludes the proof of Theorem 4.1.

10.8.1 DERIVATION OF (14) UNDER STRONGER RANGE ASSUMPTION $r = \frac{1}{2}$.

Notice that for any $x \in \mathbb{R}^d$, since k is differentiable

$$\nabla \left(\frac{d\mu_t}{d\pi} - 1 \right) (x) = \left\langle \nabla k(x, \cdot), \frac{d\mu_t}{d\pi} - 1 \right\rangle_{\mathcal{H}^d} \leq \sqrt{K_{1d}} \left\| \frac{d\mu_t}{d\pi} - 1 \right\|_{\mathcal{H}}.$$

And since $\frac{d\mu_t}{d\pi} - 1 \in \mathcal{H}$, there exists $q_t \in L^2(\pi)$ such that $\left\langle \frac{d\mu_t}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} = \varrho_i^{1/2} \langle q_t, e_i \rangle_{L^2(\pi)}$ for all i , so

$$\left\| \frac{d\mu_t}{d\pi} - 1 \right\|_{\mathcal{H}}^2 = \sum_{i \geq 1} \frac{1}{\varrho_i} \left\langle \frac{d\mu_t}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)}^2 = \sum_{i \geq 1} \langle q_t, e_i \rangle_{L^2(\pi)}^2 = \|q_t\|_{L^2(\pi)}^2.$$

Combining the above two equations, we have

$$\left\| \nabla (\log \pi)^\top \nabla \left(\frac{d\mu_t}{d\pi} - 1 \right) \right\|_{L^2(\pi)} \leq \sqrt{K_{1d}} \|q_t\|_{L^2(\pi)} \|\nabla \log \pi\|_{L^2(\pi)}.$$

Also, for the other one, notice that

$$\begin{aligned} \Delta \left(\frac{d\mu_t}{d\pi} \right) (x) &\leq \left\| \mathbf{H} \left(\frac{d\mu_t}{d\pi} - 1 \right) (x) \right\|_F = \left\| \left\langle \mathbf{H}k(x, \cdot), \frac{d\mu_t}{d\pi} - 1 \right\rangle_F \right\|_F \\ &\leq \|\mathbf{H}k(x, \cdot)\|_{\mathcal{H}^{d \times d}} \left\| \frac{d\mu_t}{d\pi} - 1 \right\|_{\mathcal{H}} \leq \sqrt{K_{2d}} \|q_t\|_{L^2(\pi)}. \end{aligned}$$

Therefore,

$$\left\| \Delta \left(\frac{d\mu_t}{d\pi} \right) \right\|_{L^2(\pi)} \leq \sqrt{K_{2d}} \|q_t\|_{L^2(\pi)}.$$

10.9 Proof of Proposition 5.1

We know that $\mu_{n+1} = (\mathbf{I} - \gamma(1 + \lambda)\nabla h)_{\#}\mu_n$ where we drop the subscripts of the witness function $h_{\mu_n, \pi}$ when it causes no ambiguity. Denote $\rho_t = (\mathbf{I} - t(1 + \lambda)\nabla h)_{\#}\mu_n$, so $\rho_0 = \mu_n$, and $\rho_\gamma = \mu_{n+1}$. Consider the difference of KL divergence between the two iterates μ_{n+1} and μ_n along the time-discretized DrMMD flow:

$$\begin{aligned} \text{KL}(\mu_{n+1} \parallel \pi) - \text{KL}(\mu_n \parallel \pi) &= \text{KL}(\rho_\gamma \parallel \pi) - \text{KL}(\rho_0 \parallel \pi) \\ &= \left. \frac{d}{dt} \right|_{t=0} \text{KL}(\rho_t \parallel \pi) \gamma + \int_0^\gamma (\gamma - t) \frac{d^2}{dt^2} \text{KL}(\rho_t \parallel \pi) dt. \end{aligned} \quad (55)$$

For the first term of (55),

$$\begin{aligned} &\left. \frac{d}{dt} \right|_{t=0} \text{KL}(\rho_t \parallel \pi) \\ &= -(1 + \lambda) \mathbb{E}_{\mu_n} \left[\nabla \log \frac{d\mu_n}{d\pi}^\top \nabla h \right] = -(1 + \lambda) \mathbb{E}_\pi \left[\nabla \frac{d\mu_n}{d\pi}^\top \nabla h \right] \\ &= -2(1 + \lambda) \left\| \nabla \frac{d\mu_n}{d\pi} \right\|_{L^2(\pi)}^2 + (1 + \lambda) \mathbb{E}_\pi \left[\nabla \frac{d\mu_n}{d\pi}^\top \left(2\nabla \frac{d\mu_n}{d\pi} - \nabla h \right) \right] \\ &= -2(1 + \lambda) \left\| \nabla \frac{d\mu_n}{d\pi} \right\|_{L^2(\pi)}^2 + (1 + \lambda) \int \left(\nabla \frac{d\mu_n}{d\pi}(x) \right)^\top \left(2\nabla \left(\frac{d\mu_n}{d\pi}(x) - 1 \right) - \nabla h(x) \right) \pi(x) dx \\ &= -2(1 + \lambda) \left\| \nabla \frac{d\mu_n}{d\pi} \right\|_{L^2(\pi)}^2 - (1 + \lambda) \int \left(2 \left(\frac{d\mu_n}{d\pi}(x) - 1 \right) - h(x) \right) \nabla \cdot \left(\pi(x) \nabla \frac{d\mu_n}{d\pi}(x) \right) dx \\ &\leq -(1 + \lambda) \frac{2}{C_P} \chi^2(\mu_n \parallel \pi) + (1 + \lambda) \left\| h - 2 \left(\frac{d\mu_n}{d\pi} - 1 \right) \right\|_{L^2(\pi)} \left\| \frac{\nabla \cdot \left(\pi \nabla \frac{d\mu_n}{d\pi} \right)}{\pi} \right\|_{L^2(\pi)}, \end{aligned} \quad (56)$$

where the fourth equality uses an integration by parts, and the last inequality uses Poincaré inequality for the first term under similar arguments in Section 10.8 and uses Cauchy-Schwarz for the second term. Using Lemma B.5 and the derivations in (53), (56) can be further upper bounded by

$$\left. \frac{d}{dt} \right|_{t=0} \text{KL}(\rho_t \parallel \pi) \leq -(1 + \lambda) \frac{2}{C_P} \chi^2(\mu_n \parallel \pi) + 2(1 + \lambda) \lambda^r Q(\mathcal{J} + \mathcal{I}). \quad (57)$$

Then, for the second term of (55), we know from Example 15.9 of Villani et al. (2009) (taking $m = 1$) that,

$$\begin{aligned} \frac{d^2}{dt^2} \text{KL}(\rho_t \parallel \pi) &= (1 + \lambda)^2 \int \nabla h(x)^\top \mathbf{H} V(\varphi_t(x)) \nabla h(x) d\mu_n(x) \\ &\quad + (1 + \lambda)^2 \int \left\| \mathbf{H} h(x) (\mathbf{I} - t(1 + \lambda) \mathbf{H} h(x))^{-1} \right\|_F^2 d\mu_n(x). \end{aligned}$$

Because $2(1 + \lambda)\gamma \sqrt{\chi^2(\mu_0 \parallel \pi) \frac{K_{2d}}{\lambda}} \leq \frac{\zeta-1}{\zeta}$ for $1 < \zeta < 2$, applying Lemma B.7 we have,

$$\frac{d^2}{dt^2} \text{KL}(\rho_t \parallel \pi) \leq 4(1 + \lambda)^2 \beta \chi^2(\mu_n \parallel \pi) \frac{K_{1d}}{\lambda} + 4(1 + \lambda)^2 \zeta^2 \chi^2(\mu_n \parallel \pi) \frac{K_{2d}}{\lambda}. \quad (58)$$

Combining the above two inequalities (57) and (58) and plugging them back into (55), we obtain

$$\begin{aligned}
\text{KL}(\mu_{n+1} \|\pi) - \text{KL}(\mu_n \|\pi) &\leq -\frac{2}{C_P}(1+\lambda)\chi^2(\mu_n \|\pi)\gamma + 2(1+\lambda)\gamma\lambda^r Q(\mathcal{J} + \mathcal{I}) \\
&\quad + 2(1+\lambda)^2\gamma^2(\beta + \zeta^2)\chi^2(\mu_n \|\pi)\frac{K_{1d} + K_{2d}}{\lambda} \\
&\leq -\frac{2}{C_P}\chi^2(\mu_n \|\pi)\gamma \\
&\quad + 4\gamma\lambda^r Q(\mathcal{J} + \mathcal{I}) + 8\gamma^2(\beta + \zeta^2)\chi^2(\mu_n \|\pi)\frac{K_{1d} + K_{2d}}{\lambda},
\end{aligned}$$

where the last inequality holds by using $0 < \lambda \leq 1$, and the result follows.

10.10 Proof of Theorem 5.1

In order to use Proposition 5.1 in the proof, first we are going to show that Proposition 5.1 holds under the conditions of Theorem 5.1. Notice that conditions 1-4 of Proposition 5.1 are precisely the conditions 1-4 of Theorem 5.1. To use Proposition 5.1 in the proof of Theorem 5.1, the only thing left is to check that the condition of step size γ in (16) is satisfied.

In Theorem 5.1, λ_n is selected to be $\left(2\gamma\frac{\chi^2(\mu_n \|\pi)(\beta + \zeta^2)(K_{1d} + K_{2d})}{Q(\mathcal{J} + \mathcal{I})}\right)^{\frac{1}{r+1}} \wedge 1$. If λ_n is taken to be the former, then

$$\begin{aligned}
2\gamma(1 + \lambda_n)\sqrt{\chi^2(\mu_n \|\pi)\frac{K_{2d}}{\lambda_n}} &\leq 4\gamma\sqrt{\chi^2(\mu_n \|\pi)\frac{K_{2d}}{\lambda_n}} \\
&= 2(2\gamma)^{\frac{2r+1}{2r+2}}\chi^2(\mu_n \|\pi)^{\frac{r}{2r+2}}(Q(\mathcal{J} + \mathcal{I}))^{\frac{1}{2r+2}}\left(\frac{1}{K_{1d} + K_{2d}}\frac{1}{\beta + \zeta^2}\right)^{\frac{1}{2r+2}}K_{2d}^{\frac{1}{2}} \\
&\stackrel{(*)}{\leq} (8\gamma)^{\frac{2r+1}{2r+2}}Q^{\frac{2r+1}{2r+2}}(\mathcal{J} + \mathcal{I})^{\frac{1}{2r+2}}\left(\frac{1}{\beta + \zeta^2}\right)^{\frac{1}{2r+2}}K_{2d}^{\frac{r}{2r+2}} \\
&\leq \frac{\zeta - 1}{\zeta}\left(\frac{1}{\beta + \zeta^2}\right)^{\frac{r}{2r+2}} \leq \frac{\zeta - 1}{\zeta}.
\end{aligned} \tag{59}$$

(*) holds because $\frac{K_{2d}}{K_{1d} + K_{2d}} \leq 1$ and

$$\chi^2(\mu_n \|\pi) = \left\| \frac{d\mu_n}{d\pi} - 1 \right\|_{L^2(\pi)}^2 = \|\mathcal{T}_\pi^r q_n\|_{L^2(\pi)}^2 \leq K^{2r}Q^2 \leq Q^2. \tag{60}$$

The second last inequality of (59) holds due to the constraint on γ in (19), and the last inequality of (59) holds because $\beta + \zeta^2 \geq 1$.

On the other hand, if λ_n is chosen to be 1, similarly based on the constraint on γ in (19), we have

$$2\gamma(1 + \lambda_n)\chi^2(\mu_n \|\pi)\frac{K_{2d}}{\lambda_n} = 4\gamma\chi^2(\mu_n \|\pi)K_{2d} \leq \frac{\zeta - 1}{\zeta}.$$

Therefore, all the conditions of Proposition 5.1 have been verified. So, if we select $\lambda_n = \left(2\gamma \frac{\chi^2(\mu_n\|\pi)(\beta+\zeta^2)(K_{1d}+K_{2d})}{Q(\mathcal{J}+\mathcal{I})}\right)^{\frac{1}{r+1}} \wedge 1$, then

$$\begin{aligned} \text{KL}(\mu_{n+1}\|\pi) - \text{KL}(\mu_n\|\pi) &\leq -\gamma \frac{2}{C_P} \chi^2(\mu_n\|\pi) \\ &\quad + \underbrace{4\gamma \lambda_n^r Q(\mathcal{J} + \mathcal{I})}_{(\Delta_1)} + \underbrace{8\gamma^2(\beta + \zeta^2) \chi^2(\mu_n\|\pi) \frac{K_{1d} + K_{2d}}{\lambda_n}}_{(\Delta_2)}. \end{aligned} \quad (61)$$

By observing (61), the first term on the right-hand side $-\frac{2}{C_P} \chi^2(\mu_n\|\pi) \gamma$ is strictly negative and is decreasing KL divergence at each iteration n of the DrMMD gradient descent. In contrast, the second term $(\Delta_1) := 4\gamma \lambda_n^r Q(\mathcal{J} + \mathcal{I})$ and the third term $(\Delta_2) := 8\gamma^2(\beta + \zeta^2) \chi^2(\mu_n\|\pi) \frac{K_{1d} + K_{2d}}{\lambda_n}$ are positive and prevent the KL divergence from decreasing. Denote $G(\lambda_n) = (\Delta_1) + (\Delta_2)$ and the optimal λ_n is achieved by taking $\frac{d}{d\lambda_n} G(\lambda_n) = 0$, which leads to $\lambda_n = \left(2\gamma \frac{\chi^2(\mu_n\|\pi)(\beta+\zeta^2)(K_{1d}+K_{2d})}{Q(\mathcal{J}+\mathcal{I})}\right)^{\frac{1}{r+1}}$. Plugging the value of λ_n back to (61) to obtain,

$$\begin{aligned} &\text{KL}(\mu_{n+1}\|\pi) - \text{KL}(\mu_n\|\pi) \\ &\leq -\frac{2}{C_P} \chi^2(\mu_n\|\pi) \gamma + 4\gamma (2\gamma)^{\frac{r}{r+1}} (\chi^2(\mu_n\|\pi)(\beta + \zeta^2)(K_{1d} + K_{2d}))^{\frac{r}{r+1}} (Q(\mathcal{J} + \mathcal{I}))^{\frac{1}{r+1}} \\ &\leq -\frac{2}{C_P} \chi^2(\mu_n\|\pi) \gamma + 8\gamma^{1+\frac{r}{r+1}} (\chi^2(\mu_n\|\pi)(\beta + \zeta^2)(K_{1d} + K_{2d}))^{\frac{r}{r+1}} (Q(\mathcal{J} + \mathcal{I}))^{\frac{1}{r+1}} \\ &\leq -\frac{2}{C_P} \chi^2(\mu_n\|\pi) \gamma + 8\gamma^{1+\frac{r}{r+1}} Q^{\frac{2r+1}{r+1}} \left((K_{1d} + K_{2d})(\beta + \zeta^2)\right)^{\frac{r}{r+1}} (\mathcal{J} + \mathcal{I})^{\frac{1}{r+1}}, \end{aligned} \quad (62)$$

where the last inequality holds because of (60). Since $\chi^2(\mu_n\|\pi) \geq \text{KL}(\mu_n\|\pi)$ (Van Erven and Harremoës, 2014, Equation (7)), we have

$$\begin{aligned} \text{KL}(\mu_{n+1}\|\pi) &\leq \left(1 - \gamma \frac{2}{C_P}\right) \text{KL}(\mu_n\|\pi) \\ &\quad + 8\gamma^{1+\frac{r}{r+1}} Q^{\frac{2r+1}{r+1}} \left((K_{1d} + K_{2d})(\beta + \zeta^2)\right)^{\frac{r}{r+1}} (\mathcal{J} + \mathcal{I})^{\frac{1}{r+1}}. \end{aligned}$$

After iterating, we obtain

$$\begin{aligned} &\text{KL}(\mu_{n_{max}}\|\pi) \\ &\leq \left(1 - \gamma \frac{2}{C_P}\right)^{n_{max}} \text{KL}(\mu_0\|\pi) + 4\gamma^{\frac{r}{r+1}} C_P Q^{\frac{2r+1}{r+1}} \left((K_{1d} + K_{2d})(\beta + \zeta^2)\right)^{\frac{r}{r+1}} (\mathcal{J} + \mathcal{I})^{\frac{1}{r+1}} \\ &\leq \exp\left(-\frac{2n_{max}\gamma}{C_P}\right) \text{KL}(\mu_0\|\pi) + 4\gamma^{\frac{r}{r+1}} C_P Q^{\frac{2r+1}{r+1}} \left((K_{1d} + K_{2d})(\beta + \zeta^2)\right)^{\frac{r}{r+1}} (\mathcal{J} + \mathcal{I})^{\frac{1}{r+1}} \end{aligned}$$

and the result follows.

10.11 Proof of Theorem 5.2

The proof of Theorem 5.2 is also based on Proposition 5.1 proved in the last section. Recalling (62) yet with adaptive step size γ_n , we have

$$\text{KL}(\mu_{n+1}\|\pi) - \text{KL}(\mu_n\|\pi)$$

$$\leq -\frac{2}{C_P} \chi^2(\mu_n \|\pi) \gamma_n + 8\gamma_n^{1+\frac{r}{r+1}} (\chi^2(\mu_n \|\pi)(\beta + \zeta^2)(K_{1d} + K_{2d}))^{\frac{r}{r+1}} (Q(\mathcal{J} + \mathcal{I}))^{\frac{1}{r+1}}.$$

From (21), we have

$$8\gamma_n^{\frac{r}{r+1}} \left((K_{1d} + K_{2d})(\beta + \zeta^2) \right)^{\frac{r}{r+1}} (Q(\mathcal{J} + \mathcal{I}))^{\frac{1}{r+1}} \leq \frac{1}{C_P} \chi^2(\mu_n \|\pi)^{\frac{1}{r+1}},$$

so that we have

$$\text{KL}(\mu_{n+1} \|\pi) - \text{KL}(\mu_n \|\pi) \leq -\frac{1}{C_P} \chi^2(\mu_n \|\pi) \gamma_n \leq -\frac{1}{C_P} \text{KL}(\mu_n \|\pi) \gamma_n.$$

Hence

$$\text{KL}(\mu_{n+1} \|\pi) \leq \left(1 - \frac{1}{C_P} \gamma_n \right) \text{KL}(\mu_n \|\pi). \quad (63)$$

After iterating n from 1 to n_{\max} , the theorem is proved.

10.12 Proof of Theorem 6.1

In order to analyze the error of space discretization, we introduce another particle descent scheme using the population witness function $h_{\mu_n, \pi}^*$ defined in (15) starting from the same initialization as that of (24),

$$\bar{y}_{n+1}^{(i)} = \bar{y}_n^{(i)} - \gamma(1 + \lambda_n) \nabla h_{\mu_n, \pi}^*(\bar{y}_n^{(i)}), \quad \bar{y}_0^{(i)} = y_0^{(i)}. \quad (64)$$

The corresponding empirical distribution of the particles at time step n is defined as $\bar{\mu}_n = \frac{1}{N} \sum_{i=1}^N \bar{y}_n^{(i)}$. Note that (64) is an unbiased sampled version (since it is composed of N i.i.d. realizations) of (15). The following proposition shows that $W_2(\bar{\mu}_n, \hat{\mu}_n) \rightarrow 0$ as $N, M \rightarrow \infty$, i.e., with a sufficient number of samples from μ_0 and π , (24) can approximate (64) with arbitrary precision. The proof of Proposition 10.1 is provided in Section 10.13.

Proposition 10.1 *Suppose k satisfies Assumption 1 and 2. Given initial particles $\{y_0^{(i)}\}_{i=1}^N$ that are i.i.d sampled from μ_0 , a sequence $(\bar{\mu}_n)_{n \in \mathbb{N}}$ of empirical distributions arising from (64), and a sequence $(\hat{\mu}_n)_{n \in \mathbb{N}}$ arising from (24), then for all $n \geq 1$, we have*

$$\mathbb{E}[W_2(\hat{\mu}_n, \bar{\mu}_n)] \leq A \left(\frac{K}{\sqrt{M}\tilde{\lambda}} + \frac{1}{\sqrt{M}} + \frac{1}{\sqrt{N}} \right) \left(\exp \left(\gamma n \frac{2(1 + \tilde{\lambda})R}{\tilde{\lambda}} \right) - 1 \right),$$

where $A = \frac{2\sqrt{KK_{1d}}}{\sqrt{KK_{2d} + K_{1d}}}$ and $R = K_{1d} + \sqrt{KK_{2d}}$ are constants that only depend on the kernel, and $\tilde{\lambda} = \min_{i=1, \dots, n} \lambda_i$ denotes the smallest regularization coefficient.

Now we are ready to prove Theorem 6.1. By the triangle inequality, we have,

$$\mathbb{E}[W_2(\hat{\mu}_n, \pi)] \leq \mathbb{E}[W_2(\hat{\mu}_n, \bar{\mu}_n)] + \mathbb{E}[W_2(\bar{\mu}_n, \mu_n)] + W_2(\mu_n, \pi).$$

From Proposition 10.1, the first term is upper bounded by

$$\begin{aligned}\mathbb{E}[W_2(\hat{\mu}_n, \bar{\mu}_n)] &\leq A \left(\frac{K}{\sqrt{M\tilde{\lambda}}} + \frac{1}{\sqrt{M}} + \frac{1}{\sqrt{N}} \right) \left(\exp \left(n\gamma \frac{2(1+\tilde{\lambda})R}{\tilde{\lambda}} \right) - 1 \right) \\ &\leq A \left(\frac{1}{\sqrt{M\tilde{\lambda}}} + \frac{1}{\sqrt{N}} \right) \left(\exp \left(\frac{4n\gamma R}{\tilde{\lambda}} \right) - 1 \right),\end{aligned}$$

where the second inequality uses $\tilde{\lambda} \leq 1$ and $K \leq 1$. Since $(\mu_n)_n$ has finite fourth moment, then by taking $p = 2, q = 4$ in (Lei, 2020, Theorem 3.1) and (Fournier and Guillin, 2015, Theorem 1), the second term is upper bounded by,

$$\mathbb{E}[W_2(\bar{\mu}_n, \mu_n)] = \mathcal{O} \left(N^{-\frac{1}{d\sqrt{4}}} \right).$$

For the third term, since the Wasserstein-2 distance is upper bounded by the square root of the KL divergence, if the target π that satisfies Talagrand-2 inequality with constant C_T (Villani et al., 2009, Definition 22.1), we have

$$W_2(\mu_n, \pi) \leq \sqrt{2C_T} \sqrt{\text{KL}(\mu_n \parallel \pi)} \leq \sqrt{2C_T} \exp \left(-\frac{n\gamma}{C_P} \right) \sqrt{\text{KL}(\mu_0 \parallel \pi)} + \mathcal{O} \left(\gamma^{\frac{r}{2r+2}} \right),$$

where the last inequality follows from Theorem 5.1. Combining the above three terms, we obtain

$$\begin{aligned}\mathbb{E}[W_2(\hat{\mu}_n, \pi)] &\leq A \left(\frac{1}{\sqrt{M\tilde{\lambda}}} + \frac{1}{\sqrt{N}} \right) \exp \left(\frac{4n\gamma R}{\tilde{\lambda}} \right) + \mathcal{O} \left(N^{-\frac{1}{d\sqrt{4}}} \right) \\ &\quad + \sqrt{2C_T} \exp \left(-\frac{n\gamma}{C_P} \right) \sqrt{\text{KL}(\mu_0 \parallel \pi)} + \mathcal{O} \left(\gamma^{\frac{r}{2r+2}} \right),\end{aligned}\tag{65}$$

where $A = \frac{2\sqrt{KK_{1d}}}{\sqrt{KK_{2d}+K_{1d}}}$. Recall from Theorem 5.1 that $\lambda_i = (\gamma\chi^2(\mu_i \parallel \pi)Z)^{\frac{1}{r+1}} \wedge 1$ for $i = 1, \dots, n$ where Z is the constant that depends on $\beta, \zeta, K_{1d}, K_{2d}, Q, \mathcal{J}, \mathcal{I}$. From the condition on the number of samples M and N in (25), we obtain that if $\tilde{\lambda} = \lambda_j = (\gamma\chi^2(\mu_j \parallel \pi)Z)^{\frac{1}{r+1}}$ for some $j \in \{1, \dots, n\}$,

$$\begin{aligned}\sqrt{M} &\gtrsim \left(\frac{1}{\gamma} \right) \left(\frac{1}{\min_{i=1, \dots, n} \text{KL}(\mu_i \parallel \pi)Z \wedge 1} \right)^{\frac{1}{r+1}} \exp \left(\frac{4n\gamma^{\frac{r}{r+1}} R}{\left(\min_{i=1, \dots, n} \text{KL}(\mu_i \parallel \pi)Z \right)^{\frac{1}{r+1}} \wedge 1} \right) \\ &\geq \frac{1}{\gamma} \left(\frac{1}{\min_{i=1, \dots, n} \text{KL}(\mu_i \parallel \pi)Z} \right)^{\frac{1}{r+1}} \exp \left(\frac{4n\gamma^{\frac{r}{r+1}} R}{\left(\min_{i=1, \dots, n} \text{KL}(\mu_i \parallel \pi)Z \right)^{\frac{1}{r+1}}} \right) \\ &\geq \frac{1}{\gamma} \left(\frac{1}{\text{KL}(\mu_j \parallel \pi)Z} \right)^{\frac{1}{r+1}} \exp \left(\frac{4n\gamma^{\frac{r}{r+1}} R}{(\text{KL}(\mu_j \parallel \pi)Z)^{\frac{1}{r+1}}} \right)\end{aligned}$$

$$\begin{aligned}
&\geq \frac{1}{\gamma} \left(\frac{1}{\chi^2(\mu_j \|\pi) Z} \right)^{\frac{1}{r+1}} \exp \left(\frac{4n\gamma^{\frac{r}{r+1}} R}{(\chi^2(\mu_j \|\pi) Z)^{\frac{1}{r+1}}} \right) \\
&= A\gamma^{-\frac{r}{r+1}} \frac{1}{\tilde{\lambda}} \exp \left(\frac{4n\gamma R}{\tilde{\lambda}} \right) \geq \gamma^{-\frac{r}{2r+2}} \frac{1}{\tilde{\lambda}} \exp \left(\frac{4n\gamma R}{\tilde{\lambda}} \right).
\end{aligned}$$

On the other hand, if $\tilde{\lambda} = 1$, since $\gamma \leq 1$,

$$\begin{aligned}
\sqrt{M} &\gtrsim \left(\frac{1}{\gamma} \right) \left(\frac{1}{\min_{i=1,\dots,n} \text{KL}(\mu_i \|\pi) Z \wedge 1} \right)^{\frac{1}{r+1}} \exp \left(\frac{4n\gamma^{\frac{r}{r+1}} R}{\left(\min_{i=1,\dots,n} \text{KL}(\mu_i \|\pi) Z \right)^{\frac{1}{r+1}} \wedge 1} \right) \\
&\geq \frac{1}{\gamma} \exp \left(4n\gamma^{\frac{r}{r+1}} R \right) \geq \gamma^{-\frac{r}{2r+2}} \exp(4n\gamma R) = \gamma^{-\frac{r}{2r+2}} \frac{1}{\tilde{\lambda}} \exp \left(\frac{4n\gamma R}{\tilde{\lambda}} \right).
\end{aligned}$$

Similarly for N , we have

$$\sqrt{N} \gtrsim \gamma^{-\frac{r}{2r+2}} \exp \left(\frac{4n\gamma R}{\tilde{\lambda}} \right) \quad \text{and} \quad N^{-\frac{1}{d\sqrt{4}}} \lesssim \gamma^{\frac{r}{2r+2}}.$$

Plugging them back to (65), we obtain

$$\mathbb{E}[W_2(\hat{\mu}_n, \pi)] \leq \sqrt{2C_T} \exp \left(-\frac{n\gamma}{C_P} \right) \sqrt{\text{KL}(\mu_0 \|\pi)} + \mathcal{O} \left(\gamma^{\frac{r}{2r+2}} \right),$$

which completes the proof.

10.13 Proof of Proposition 10.1

Since the proof below works for any regularization coefficient λ , we use a fixed λ for the majority of the analysis and resort back to adaptive λ_n at the end of the proof. For empirical distributions $\hat{\mu}_n = \frac{1}{N} \sum_{i=1}^N y_n^{(i)}$ and $\bar{\mu}_n = \frac{1}{N} \sum_{i=1}^N \bar{y}_n^{(i)}$ defined in (24) and (64), note that

$$\mathbb{E} W_2^2(\hat{\mu}_n, \bar{\mu}_n) \leq \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\left\| y_n^{(i)} - \bar{y}_n^{(i)} \right\|^2 \right] := c_n^2.$$

Consider

$$\begin{aligned}
c_{n+1} &= \sqrt{\frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\left\| y_{n+1}^{(i)} - \bar{y}_{n+1}^{(i)} \right\|^2 \right]} \\
&= \sqrt{\frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\left\| y_n^{(i)} - \bar{y}_n^{(i)} - \gamma(1+\lambda) \left(\nabla h_{\hat{\mu}_n, \hat{\pi}}^*(y_n^{(i)}) - \nabla h_{\mu_n, \pi}^*(\bar{y}_n^{(i)}) \right) \right\|^2 \right]} \\
&\leq \sqrt{\frac{1}{N} \sum_{i=1}^N \mathbb{E} \left\| y_n^{(i)} - \bar{y}_n^{(i)} \right\|^2}
\end{aligned}$$

$$\begin{aligned}
& + \sqrt{\frac{1}{N} \mathbb{E} \sum_{i=1}^N \left\| \gamma(1+\lambda) \left(\nabla h_{\hat{\mu}_n, \hat{\pi}}^*(y_n^{(i)}) - \nabla h_{\mu_n, \pi}^*(\bar{y}_n^{(i)}) \right) \right\|^2} \\
& = c_n + \frac{\gamma(1+\lambda)}{\sqrt{N}} \sqrt{\sum_{i=1}^N \mathbb{E} \left\| \nabla h_{\hat{\mu}_n, \hat{\pi}}^*(y_n^{(i)}) - \nabla h_{\mu_n, \pi}^*(\bar{y}_n^{(i)}) \right\|^2},
\end{aligned}$$

where we used Minkowski's inequality,

$$\sqrt{\sum_{i=1}^N \|a_i + b_i\|^2} \leq \sqrt{\sum_{i=1}^N \|a_i\|^2} + \sqrt{\sum_{i=1}^N \|b_i\|^2}$$

in the above inequalities. Again, by Minkowski's inequality, we have

$$\begin{aligned}
c_{n+1} \leq c_n + \gamma(1+\lambda) & \underbrace{\left(\frac{1}{\sqrt{N}} \sqrt{\sum_{i=1}^N \mathbb{E} \left[\left\| \nabla h_{\hat{\mu}_n, \hat{\pi}}^*(y_n^{(i)}) - \nabla h_{\hat{\mu}_n, \hat{\pi}}^*(\bar{y}_n^{(i)}) \right\|^2 \right]} \right)}_{(i)} \\
& + \underbrace{\frac{1}{\sqrt{N}} \sqrt{\sum_{i=1}^N \mathbb{E} \left[\left\| \nabla h_{\hat{\mu}_n, \hat{\pi}}^*(\bar{y}_n^{(i)}) - \nabla h_{\mu_n, \pi}^*(\bar{y}_n^{(i)}) \right\|^2 \right]}}_{(ii)}.
\end{aligned}$$

10.13.1 CONTROLLING (i):

$$\begin{aligned}
\sum_{i=1}^N \mathbb{E} \left\| \nabla h_{\hat{\mu}_n, \hat{\pi}}^*(y_n^{(i)}) - \nabla h_{\hat{\mu}_n, \hat{\pi}}^*(\bar{y}_n^{(i)}) \right\|^2 & = \sum_{i=1}^N \mathbb{E} \left[\sum_{j=1}^d \left\langle \partial_j k(y_n^{(i)}, \cdot) - \partial_j k(\bar{y}_n^{(i)}, \cdot), h_{\hat{\mu}_n, \hat{\pi}}^* \right\rangle_{\mathcal{H}}^2 \right] \\
& \leq \sum_{i=1}^N \mathbb{E} \left[\sum_{j=1}^d \left\| \partial_j k(y_n^{(i)}, \cdot) - \partial_j k(\bar{y}_n^{(i)}, \cdot) \right\|_{\mathcal{H}}^2 \right] \|h_{\hat{\mu}_n, \hat{\pi}}^*\|_{\mathcal{H}}^2 \leq \frac{4KK_{2d}}{\lambda^2} \sum_{i=1}^N \mathbb{E} \|y_n^{(i)} - \bar{y}_n^{(i)}\|^2,
\end{aligned}$$

where the second inequality uses Cauchy-Schwarz inequality and the third follows from using Lemma B.3 and Lemma B.4. So we have

$$(i) = \frac{1}{\sqrt{N}} \sqrt{\sum_{i=1}^N \mathbb{E} \left\| \nabla h_{\hat{\mu}_n, \hat{\pi}}^*(y_n^{(i)}) - \nabla h_{\hat{\mu}_n, \hat{\pi}}^*(\bar{y}_n^{(i)}) \right\|^2} \leq \frac{2\sqrt{KK_{2d}}}{\lambda} c_n.$$

10.13.2 CONTROLLING (ii):

First, we introduce some auxiliary witness functions,

$$h_{\hat{\mu}_n, \hat{\pi}}^* = 2(\Sigma_{\hat{\pi}} + \lambda I)^{-1} (m_{\hat{\mu}_n} - m_{\hat{\pi}}), \quad \bar{\bar{h}}_n^* = 2(\Sigma_{\pi} + \lambda I)^{-1} (m_{\hat{\mu}_n} - m_{\hat{\pi}}),$$

and for completeness, we recall the witness function we are interested in:

$$h_{\mu_n, \pi}^* = 2(\Sigma_{\pi} + \lambda I)^{-1} (m_{\mu_n} - m_{\pi}), \quad h_{\hat{\mu}_n, \hat{\pi}}^* = 2(\Sigma_{\hat{\pi}} + \lambda I)^{-1} (m_{\hat{\mu}_n} - m_{\hat{\pi}}).$$

We know that

$$\begin{aligned}
(ii) &= \frac{1}{\sqrt{N}} \sqrt{\sum_{i=1}^N \mathbb{E} \left[\left\| \nabla h_{\hat{\mu}_n, \hat{\pi}}^* (\bar{y}_n^{(i)}) - \nabla h_{\bar{\mu}_n, \pi}^* (\bar{y}_n^{(i)}) \right\|^2 \right]} \\
&\leq \frac{1}{\sqrt{N}} \sqrt{\sum_{i=1}^N \mathbb{E} \left[\left\| \nabla h_{\hat{\mu}_n, \hat{\pi}}^* (\bar{y}_n^{(i)}) - \nabla h_{\bar{\mu}_n, \hat{\pi}}^* (\bar{y}_n^{(i)}) \right\|^2 \right]} \\
&\quad + \frac{1}{\sqrt{N}} \sqrt{\sum_{i=1}^N \mathbb{E} \left[\left\| \nabla h_{\bar{\mu}_n, \hat{\pi}}^* (\bar{y}_n^{(i)}) - \nabla \bar{h}_n^* (\bar{y}_n^{(i)}) \right\|^2 \right]} \\
&\quad + \frac{1}{\sqrt{N}} \sqrt{\sum_{i=1}^N \mathbb{E} \left[\left\| \nabla \bar{h}_n^* (\bar{y}_n^{(i)}) - \nabla h_{\mu_n, \pi}^* (\bar{y}_n^{(i)}) \right\|^2 \right]} \\
&\leq \frac{1}{\sqrt{N}} \left(\sqrt{\sum_{i=1}^N K_{1d} \mathbb{E} \left\| h_{\hat{\mu}_n, \hat{\pi}}^* - h_{\bar{\mu}_n, \hat{\pi}}^* \right\|_{\mathcal{H}}^2} + \sqrt{\sum_{i=1}^N K_{1d} \mathbb{E} \left\| h_{\bar{\mu}_n, \hat{\pi}}^* - \bar{h}_n^* \right\|_{\mathcal{H}}^2} \right. \\
&\quad \left. + \sqrt{\sum_{i=1}^N K_{1d} \mathbb{E} \left\| \bar{h}_n^* - h_{\mu_n, \pi}^* \right\|_{\mathcal{H}}^2} \right) \\
&= \sqrt{K_{1d}} \left(\sqrt{\mathbb{E} \left\| h_{\hat{\mu}_n, \hat{\pi}}^* - h_{\bar{\mu}_n, \hat{\pi}}^* \right\|_{\mathcal{H}}^2} + \sqrt{\mathbb{E} \left\| h_{\bar{\mu}_n, \hat{\pi}}^* - \bar{h}_n^* \right\|_{\mathcal{H}}^2} + \sqrt{\mathbb{E} \left\| \bar{h}_n^* - h_{\mu_n, \pi}^* \right\|_{\mathcal{H}}^2} \right),
\end{aligned}$$

where the first inequality follows from Minkowski's inequality, and the second inequality uses the fact that for $h_0, h_1 \in \mathcal{H}$,

$$\|\nabla h_1 - \nabla h_0\|_{\mathcal{H}^d}^2 \leq \|\nabla_1 k(x, \cdot)\|_{\mathcal{H}^d}^2 \|h_1 - h_0\|_{\mathcal{H}}^2 \leq K_{1d} \|h_1 - h_0\|_{\mathcal{H}}^2.$$

Next, we will bound $\sqrt{\mathbb{E} \left\| h_{\hat{\mu}_n, \hat{\pi}}^* - h_{\bar{\mu}_n, \hat{\pi}}^* \right\|_{\mathcal{H}}^2}$, $\sqrt{\mathbb{E} \left\| h_{\bar{\mu}_n, \hat{\pi}}^* - \bar{h}_n^* \right\|_{\mathcal{H}}^2}$, and $\sqrt{\mathbb{E} \left\| \bar{h}_n^* - h_{\mu_n, \pi}^* \right\|_{\mathcal{H}}^2}$ separately.

First, by noticing that $h_{\hat{\mu}_n, \hat{\pi}}^* = 2(\Sigma_{\hat{\pi}} + \lambda I)^{-1}(m_{\hat{\mu}_n} - m_{\hat{\pi}})$ is the witness function associated with $\text{DrMMD}(\hat{\mu}_n || \hat{\pi})$, and $h_{\bar{\mu}_n, \hat{\pi}}^* = 2(\Sigma_{\hat{\pi}} + \lambda I)^{-1}(m_{\bar{\mu}_n} - m_{\hat{\pi}})$ is the witness function associated with $\text{DrMMD}(\bar{\mu}_n || \hat{\pi})$, by using Lemma B.4, we have

$$\sqrt{\mathbb{E} \left\| h_{\hat{\mu}_n, \hat{\pi}}^* - h_{\bar{\mu}_n, \hat{\pi}}^* \right\|_{\mathcal{H}}^2} \leq \sqrt{\mathbb{E} \frac{4K_{1d}}{\lambda^2} W_2^2(\bar{\mu}_n, \hat{\mu}_n)} \leq \frac{2\sqrt{K_{1d}}}{\lambda} c_n. \quad (66)$$

Second,

$$\begin{aligned}
\sqrt{\mathbb{E} \left\| h_{\bar{\mu}_n, \hat{\pi}}^* - \bar{h}_n^* \right\|_{\mathcal{H}}^2} &= \sqrt{\mathbb{E} \left\| 2(\Sigma_{\hat{\pi}} + \lambda I)^{-1}(m_{\bar{\mu}_n} - m_{\hat{\pi}}) - 2(\Sigma_{\pi} + \lambda I)^{-1}(m_{\bar{\mu}_n} - m_{\hat{\pi}}) \right\|_{\mathcal{H}}^2} \\
&\leq 2\sqrt{\mathbb{E} \left\| (\Sigma_{\hat{\pi}} + \lambda I)^{-1} - (\Sigma_{\pi} + \lambda I)^{-1} \right\|_{\text{HS}}^2 \|m_{\bar{\mu}_n} - m_{\hat{\pi}}\|_{\mathcal{H}}^2} \\
&\leq 4\sqrt{K} \sqrt{\mathbb{E} \left\| (\Sigma_{\hat{\pi}} + \lambda I)^{-1} - (\Sigma_{\pi} + \lambda I)^{-1} \right\|_{\text{HS}}^2}
\end{aligned}$$

$$\begin{aligned}
&= 4\sqrt{K} \sqrt{\mathbb{E} \left\| (\Sigma_{\hat{\pi}} + \lambda I)^{-1} \left((\Sigma_{\hat{\pi}} + \lambda I) - (\Sigma_{\pi} + \lambda I) \right) (\Sigma_{\pi} + \lambda I)^{-1} \right\|_{\text{HS}}^2} \\
&\leq 4\sqrt{K} \frac{1}{\lambda^2} \sqrt{\mathbb{E} \|\Sigma_{\hat{\pi}} - \Sigma_{\pi}\|_{\text{HS}}^2} \leq 4\sqrt{K} \frac{1}{\lambda^2} \sqrt{\frac{K^2}{M}},
\end{aligned} \tag{67}$$

where the last inequality follows from using Lemma B.8 and the fact that $\|k(x, \cdot) \otimes k(x, \cdot)\|_{\text{HS}} \leq K$.

Third,

$$\begin{aligned}
&\sqrt{\mathbb{E} \left\| \bar{h}_n^* - h_{\mu_n, \pi}^* \right\|_{\mathcal{H}}^2} = \sqrt{\mathbb{E} \left\| 2(\Sigma_{\pi} + \lambda I)^{-1} (m_{\bar{\mu}_n} - m_{\hat{\pi}}) - 2(\Sigma_{\pi} + \lambda I)^{-1} (m_{\mu_n} - m_{\pi}) \right\|_{\mathcal{H}}^2} \\
&\leq \frac{2}{\lambda} \sqrt{\mathbb{E} \|(m_{\bar{\mu}_n} - m_{\hat{\pi}}) - (m_{\mu_n} - m_{\pi})\|_{\mathcal{H}}^2} \\
&\leq \frac{2}{\lambda} \left(\sqrt{\mathbb{E} \|m_{\bar{\mu}_n} - m_{\mu_n}\|_{\mathcal{H}}^2} + \sqrt{\mathbb{E} \|m_{\hat{\pi}} - m_{\pi}\|_{\mathcal{H}}^2} \right) \leq \frac{4}{\lambda} \left(\sqrt{\frac{K}{N}} + \sqrt{\frac{K}{M}} \right),
\end{aligned} \tag{68}$$

where the first inequality follows from Cauchy-Schwartz, and the last inequality from Lemma B.8 since $\|k(x, \cdot)\|_{\mathcal{H}} \leq \sqrt{K}$. Therefore, combining (66), (67) and (68), we have

$$(ii) \leq 2\sqrt{K_{1d}} \left(\frac{\sqrt{K_{1d}}}{\lambda} c_n + \frac{2K^{3/2}}{\sqrt{M}\lambda^2} + \frac{2\sqrt{K}}{\sqrt{N}\lambda} + \frac{2\sqrt{K}}{\sqrt{M}\lambda} \right).$$

Combining (i) and (ii), we have

$$c_{n+1} \leq c_n \left(1 + \gamma(1 + \lambda) \frac{2\sqrt{K}K_{2d} + 2K_{1d}}{\lambda} \right) + 2\gamma(1 + \lambda) \sqrt{K_{1d}} \left(\frac{2K^{3/2}}{\sqrt{M}\lambda^2} + \frac{2\sqrt{K}}{\sqrt{M}\lambda} + \frac{2\sqrt{K}}{\sqrt{N}\lambda} \right).$$

Denoting $A = \frac{2\sqrt{K}K_{1d}}{\sqrt{K}K_{2d} + K_{1d}}$ and $R = K_{1d} + \sqrt{K}K_{2d}$ as constants that only depend on the kernel, and using the discrete Gronwall lemma (Lemma 26 from Arbel et al. 2019) along with $c_0 = 0$, we obtain

$$c_{n_{\max}} \leq A \left(\frac{K}{\sqrt{M}\lambda} + \frac{1}{\sqrt{M}} + \frac{1}{\sqrt{N}} \right) \left(\exp \left(\gamma n_{\max} \frac{2(1 + \lambda)R}{\lambda} \right) - 1 \right).$$

Since $\mathbb{E} W_2(\hat{\mu}_n, \bar{\mu}_n) \leq \sqrt{\mathbb{E} W_2^2(\hat{\mu}_n, \bar{\mu}_n)} \leq c_n$, we reach

$$\mathbb{E} W_2(\hat{\mu}_{n_{\max}}, \bar{\mu}_{n_{\max}}) \leq A \left(\frac{K}{\sqrt{M}\lambda} + \frac{1}{\sqrt{M}} + \frac{1}{\sqrt{N}} \right) \left(\exp \left(\gamma n_{\max} \frac{2(1 + \lambda)R}{\lambda} \right) - 1 \right).$$

Finally, the proof is completed by noting that the r.h.s. is monotonically decreasing in λ and therefore the r.h.s. can be bounded by replacing λ with $\tilde{\lambda} = \min_{i=1, \dots, n_{\max}} \lambda_i$.

10.14 Proof of Proposition 6.1

By defining the following operators,

$$S_x : \mathcal{H} \rightarrow \mathbb{R}^M, \quad f \rightarrow \frac{1}{\sqrt{M}} \left[f(x^{(1)}), \dots, f(x^{(M)}) \right]^\top,$$

$$\begin{aligned}
S_x^* : \mathbb{R}^M &\rightarrow \mathcal{H}, \quad \alpha \rightarrow \frac{1}{\sqrt{M}} \sum_{i=1}^M \alpha_i k(x^{(i)}, \cdot), \\
S_y : \mathcal{H} &\rightarrow \mathbb{R}^N, \quad f \rightarrow \frac{1}{\sqrt{N}} \left[f(y^{(1)}), \dots, f(y^{(N)}) \right]^\top, \\
S_y^* : \mathbb{R}^N &\rightarrow \mathcal{H}, \quad \alpha \rightarrow \frac{1}{\sqrt{N}} \sum_{i=1}^N \alpha_i k(y^{(i)}, \cdot).
\end{aligned}$$

Then we have

$$\Sigma_{\hat{\pi}} = S_x^* S_x, \quad K_{xx} = M S_x S_x^*, \quad K_{xy} = \sqrt{MN} S_x S_y^*.$$

Using these, note that

$$\begin{aligned}
h_{\hat{\mu}, \hat{\pi}}^* &= 2 (\Sigma_{\hat{\pi}} + \lambda \mathbf{I})^{-1} (m_{\hat{\mu}} - m_{\hat{\pi}}) \\
&= 2 \left(\frac{1}{M} \sum_{i=1}^M k(x^{(i)}, \cdot) \otimes k(x^{(i)}, \cdot) + \lambda \mathbf{I} \right)^{-1} \left(\frac{1}{N} \sum_{i=1}^N k(y^{(i)}, \cdot) - \frac{1}{M} \sum_{i=1}^M k(x^{(i)}, \cdot) \right) \\
&= 2 (S_x^* S_x + \lambda \mathbf{I})^{-1} \left(\frac{1}{\sqrt{N}} S_y^* \mathbb{1}_N - \frac{1}{\sqrt{M}} S_x^* \mathbb{1}_M \right). \tag{69}
\end{aligned}$$

From the Woodbury inversion lemma, we have that

$$(S_x^* S_x + \lambda \mathbf{I})^{-1} = \frac{1}{\lambda} \mathbf{I} - \frac{1}{\lambda} S_x^* (S_x S_x^* + \lambda \mathbf{I})^{-1} S_x.$$

Plugging the above into (69), we obtain

$$\begin{aligned}
h_{\hat{\mu}, \hat{\pi}}^* &= 2 \left(\frac{1}{\lambda} \mathbf{I} - \frac{1}{\lambda} S_x^* (S_x S_x^* + \lambda \mathbf{I})^{-1} S_x \right) \left(\frac{1}{\sqrt{N}} S_y^* \mathbb{1}_N - \frac{1}{\sqrt{M}} S_x^* \mathbb{1}_M \right) \\
&= \frac{2}{\lambda} \left(\frac{1}{\sqrt{N}} S_y^* \mathbb{1}_N - \frac{1}{\sqrt{M}} S_x^* \mathbb{1}_M \right) - \frac{2}{\lambda} S_x^* \left(\frac{1}{M} K_{xx} + \lambda \mathbf{I} \right)^{-1} \frac{1}{N \sqrt{M}} K_{xy} \mathbb{1}_N \\
&\quad + \frac{2}{\lambda} S_x^* \left(\frac{1}{M} K_{xx} + \lambda \mathbf{I} \right)^{-1} \frac{1}{M \sqrt{M}} K_{xx} \mathbb{1}_M \\
&= \frac{2}{N \lambda} k(\cdot, y^{1:N}) \mathbb{1}_N - \frac{2}{M \lambda} k(\cdot, x^{1:M}) \mathbb{1}_M - \frac{2}{N \lambda} k(\cdot, x^{1:M}) (K_{xx} + M \lambda \mathbf{I})^{-1} K_{xy} \mathbb{1}_N \\
&\quad + \frac{2}{M \lambda} k(\cdot, x^{1:M}) (K_{xx} + M \lambda \mathbf{I})^{-1} K_{xx} \mathbb{1}_M. \tag{70}
\end{aligned}$$

Obtaining DrMMD($\hat{\mu} \parallel \hat{\pi}$) is then easy with $h_{\hat{\mu}, \hat{\pi}}^*$ shown in (70).

$$\begin{aligned}
\text{DrMMD}(\hat{\mu} \parallel \hat{\pi}) &= (1 + \lambda) \left\| (\Sigma_{\hat{\pi}} + \lambda \mathbf{I})^{-\frac{1}{2}} (m_{\hat{\mu}} - m_{\hat{\pi}}) \right\|_{\mathcal{H}}^2 \\
&= (1 + \lambda) \left\langle \frac{1}{2} h_{\hat{\mu}, \hat{\pi}}^*, m_{\hat{\mu}} - m_{\hat{\pi}} \right\rangle_{\mathcal{H}} \\
&= (1 + \lambda) \left\langle \frac{1}{2} h_{\hat{\mu}, \hat{\pi}}^*, \frac{1}{N} \sum_{i=1}^N k(y^{(i)}, \cdot) - \frac{1}{M} \sum_{i=1}^M k(x^{(i)}, \cdot) \right\rangle_{\mathcal{H}}
\end{aligned}$$

$$\begin{aligned}
&= \frac{1+\lambda}{\lambda} \left(\frac{1}{N^2} \mathbb{1}_N^\top K_{yy} \mathbb{1}_N + \frac{1}{M^2} \mathbb{1}_M^\top K_{xx} \mathbb{1}_M - \frac{2}{MN} \mathbb{1}_M^\top K_{xy} \mathbb{1}_N \right. \\
&\quad - \frac{1}{N^2} \mathbb{1}_N^\top K_{xy}^\top (K_{xx} + M\lambda I)^{-1} K_{xy} \mathbb{1}_N + \frac{2}{NM} \mathbb{1}_M^\top K_{xx} (K_{xx} + M\lambda I)^{-1} K_{xy} \mathbb{1}_N \\
&\quad \left. - \frac{1}{M^2} \mathbb{1}_M^\top K_{xx} (K_{xx} + M\lambda I)^{-1} K_{xx} \mathbb{1}_M \right). \tag{71}
\end{aligned}$$

Acknowledgments

We would like to thank Gabriele Steidl and Viktor Stein for fruitful discussions on DrMMD and related functionals. Zonghao Chen is supported by the Engineering and Physical Sciences Research Council (EPSRC) through grant [EP/S021566/1]. Aratrika Mustafi and Bharath K. Sriperumbudur are partially supported by the National Science Foundation (NSF) grant DMS-2413425 and NSF CAREER award DMS-1945396. Pierre Glaser and Arthur Gretton are supported by the Gatsby Charitable Foundation. Anna Korba thanks Google for their academic support in the form of a gift in support of her academic research.

References

- Rohit Agrawal and Thibaut Horel. Optimal bounds between f -divergences and integral probability metrics. *Journal of Machine Learning Research*, 22(1):5662–5720, 2021.
- Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient Flows: In Metric Spaces and in the Space of Probability Measures*. Springer Science & Business Media, 2005.
- Abdul Fatir Ansari, Ming Liang Ang, and Harold Soh. Refining deep generative models via discriminator gradient flow. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=Zbc-ue9p_rE.
- Michael Arbel, Danica J Sutherland, Mikolaj Binkowski, and Arthur Gretton. On gradient regularizers for MMD GANs. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- Michael Arbel, Anna Korba, Adil Salim, and Arthur Gretton. Maximum mean discrepancy gradient flow. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Nachman Aronszajn. Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68(3):337–404, 1950.
- Krishna Balasubramanian, Tong Li, and Ming Yuan. On the optimality of kernel-embedding based goodness-of-fit tests. *Journal of Machine Learning Research*, 22(1):1–45, 2021.

- Krishna Balasubramanian, Sinho Chewi, Murat A. Erdogdu, Adil Salim, and Shunshi Zhang. Towards a theory of non-log-concave sampling: First-order stationarity guarantees for Langevin Monte Carlo. In Po-Ling Loh and Maxim Raginsky, editors, *Conference on Learning Theory*, pages 2896–2923. PMLR, 2022.
- Sayan Banerjee, Krishna Balasubramanian, and Promit Ghosal. Improved finite-particle convergence rates for Stein variational gradient descent. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=sbG8qhMjkZ>.
- Frank Bauer, Sergei Pereverzev, and Lorenzo Rosasco. On regularization algorithms in learning theory. *Journal of Complexity*, 23(1):52–72, 2007.
- A Belafhal, Z Hricha, L Dalil-Essakali, and T Usman. A note on some integrals involving Hermite polynomials and their applications. *Advanced Mathematical Models and Applications*, 5(3):313–319, 2020.
- Jeremiah Birrell, Paul Dupuis, Markos A Katsoulakis, Yannis Pantazis, and Luc Rey-Bellet. (f, Γ) -divergences: Interpolating between f -divergences and integral probability metrics. *Journal of Machine Learning Research*, 23(39):1–70, 2022.
- Clément Bonet, Théo Uscidda, Adam David, Pierre-Cyril Aubin-Frankowski, and Anna Korba. Mirror and preconditioned gradient descent in Wasserstein space. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, editors, *Advances in Neural Information Processing Systems*, volume 37, pages 25311–25374. Curran Associates, Inc., 2024.
- Siwan Boufadène and François-Xavier Vialard. On the global convergence of Wasserstein gradient flow of the Coulomb discrepancy, 2024. URL <https://arxiv.org/abs/2312.00800>.
- Stephen P. Boyd and Lieven Vandenbergh. *Convex Optimization*. Cambridge University Press, 2004.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2019.
- Andrea Caponnetto and Ernesto De Vito. Optimal rates for the regularized least-squares algorithm. *Foundations of Computational Mathematics*, 7:331–368, 2007.
- Claudio Carmeli, Ernesto De Vito, Alessandro Toigo, and Veronica Umanitá. Vector-valued reproducing kernel Hilbert spaces and universality. *Analysis and Applications*, 8(01):19–61, 2010.

- Niladri Chatterji, Jelena Diakonikolas, Michael I. Jordan, and Peter Bartlett. Langevin Monte Carlo without smoothness. In S. Chiappa and R. Calandra, editors, *International Conference on Artificial Intelligence and Statistics*, pages 1716–1726. PMLR, 2020.
- Fan Chen, Yiqing Lin, Zhenjie Ren, and Songbo Wang. Uniform-in-time propagation of chaos for kinetic mean field Langevin dynamics. *Electronic Journal of Probability*, 29:1–43, 2024.
- Zonghao Chen, Toni Karvonen, Heishiro Kanagawa, François-Xavier Briol, and Chris Oates. Stationary MMD points for cubature. <https://arxiv.org/abs/2505.20754>, 2025.
- Sinho Chewi, Thibaut Le Gouic, Chen Lu, Tyler Maunu, and Philippe Rigollet. SVGD as a kernelized Wasserstein gradient flow of the chi-squared divergence. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 2098–2109. Curran Associates, Inc., 2020.
- Sinho Chewi, Murat A. Erdogdu, Mufan Li, Ruoqi Shen, and Matthew S. Zhang. Analysis of Langevin Monte Carlo from Poincare to log-Sobolev. *Foundations of Computational Mathematics*, pages 1–51, 2024.
- Lenaic Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- Katy Craig, Karthik Elamvazhuthi, Matt Haberland, and Olga Turanova. A blob method for inhomogeneous diffusion with applications to multi-agent control and sampling. *Mathematics of Computation*, 2023.
- Katy Craig, Matt Jacobs, and Olga Turanova. Nonlocal approximation of slow and fast diffusion. *Journal of Differential Equations*, 426:782–852, 2025.
- Felipe Cucker and Ding Xuan Zhou. *Learning Theory: An Approximation Theory Viewpoint*, volume 24. Cambridge University Press, 2007.
- Arnak Dalalyan. Further and stronger analogy between sampling and optimization: Langevin Monte Carlo and gradient descent. In S. Kale and O. Shamir, editors, *Conference on Learning Theory*, pages 678–689. PMLR, 2017.
- Arnak S Dalalyan and Avetik Karagulyan. User-friendly guarantees for the Langevin Monte Carlo with inaccurate gradient. *Stochastic Processes and their Applications*, 129(12): 5278–5311, 2019.
- Arnak S Dalalyan, Avetik Karagulyan, and Lionel Riou-Durand. Bounding the error of discretized Langevin algorithms for non-strongly log-concave targets. *Journal of Machine Learning Research*, 23(235):1–38, 2022.
- Andrew Duncan, Nikolas Nüsken, and Lukasz Szpruch. On the geometry of Stein variational gradient descent. *Journal of Machine Learning Research*, 24(56):1–39, 2023.

- Alain Durmus, Szymon Majewski, and Bl  zej Miasojedow. Analysis of Langevin Monte Carlo via convex optimization. *Journal of Machine Learning Research*, 20(1):2666–2711, 2019.
- Heinz Werner Engl, Martin Hanke, and Andreas Neubauer. *Regularization of Inverse Problems*, volume 375. Springer Science & Business Media, 1996.
- Vassiliy A Epanechnikov. Non-parametric estimation of a multivariate probability density. *Theory of Probability & Its Applications*, 14(1):153–158, 1969.
- Jean Feydy, Thibault S  journ  , Fran  ois-Xavier Vialard, Shun-ichi Amari, Alain Trouv  , and Gabriel Peyr  . Interpolating between optimal transport and MMD using Sinkhorn divergences. In K. Chaudhuri and M. Sugiyama, editors, *International Conference on Artificial Intelligence and Statistics*, pages 2681–2690. PMLR, 2019.
- Simon Fischer and Ingo Steinwart. Sobolev norm learning rates for regularized least-squares algorithms. *Journal of Machine Learning Research*, 21(205):1–38, 2020.
- Nicolas Fournier and Arnaud Guillin. On the rate of convergence in Wasserstein distance of the empirical measure. *Probability Theory and Related Fields*, 162(3-4):707–738, 2015.
- Jean-Yves Franceschi, Mike Gartrell, Ludovic Dos Santos, Thibaut Issenhuth, Emmanuel de B  zenac, Micka  l Chen, and Alain Rakotomamonjy. Unifying GANs and score-based diffusion as generative particle models. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, editors, *Advances in Neural Information Processing Systems*, volume 36. Curran Associates, Inc., 2024.
- Alexandre Galashov, Valentin De Bortoli, and Arthur Gretton. Deep MMD gradient flow without adversarial training. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=Pf85K2wtz8>.
- Yuan Gao, Yuling Jiao, Yang Wang, Yao Wang, Can Yang, and Shunkang Zhang. Deep generative learning via variational gradient flow. In K. Chaudhuri and R. Salakhutdinov, editors, *International Conference on Machine Learning*, pages 2093–2101. PMLR, 2019.
- Aude Genevay, Gabriel Peyr  , and Marco Cuturi. Learning generative models with Sinkhorn divergences. In A. Storkey and F. Perez-Cruz, editors, *International Conference on Artificial Intelligence and Statistics*, pages 1608–1617. PMLR, 2018.
- Pierre Glaser, Michael Arbel, and Arthur Gretton. KALE flow: A relaxed KL gradient flow for probabilities with disjoint support. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 8018–8031. Curran Associates, Inc., 2021.
- Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Sch  lkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(1):723–773, 2012.
- Thomas Hakon Gronwall. Note on the derivatives with respect to a parameter of the solutions of a system of differential equations. *Annals of Mathematics*, 20(4):292–296, 1919.

- Hyemin Gu, Panagiota Birmpa, Yannis Pantazis, Luc Rey-Bellet, and Markos A. Katsoulakis. Lipschitz-regularized gradient flows and generative particle algorithms for high-dimensional scarce data. *SIAM Journal on Mathematics of Data Science*, 6(4):1205–1235, 2024.
- Omar Hagrass, Bharath K. Sriperumbudur, and Bing Li. Spectral regularized kernel two-sample tests. *The Annals of Statistics*, 52(3):1076–1101, 2024a.
- Omar Hagrass, Bharath K. Sriperumbudur, and Bing Li. Spectral regularized kernel goodness-of-fit tests. *Journal of Machine Learning Research*, 25(309):1–52, 2024b.
- Zaïd Harchaoui, Francis Bach, and Eric Moulines. Testing for homogeneity with kernel Fisher discriminant analysis. In J. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2007.
- Ye He, Krishna Balasubramanian, Bharath K. Sriperumbudur, and Jianfeng Lu. Regularized Stein variational gradient flow. *Foundations of Computational Mathematics*, pages 1–59, 2024.
- Johannes Hertrich, Robert Beinert, Manuel Gräf, and Gabriele Steidl. Wasserstein gradient flows of the discrepancy with distance kernel on the line. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 431–443. Springer, 2023.
- Johannes Hertrich, Manuel Gräf, Robert Beinert, and Gabriele Steidl. Wasserstein steepest descent flows of discrepancies with Riesz kernels. *Journal of Mathematical Analysis and Applications*, 531(1):127829, 2024a.
- Johannes Hertrich, Christian Wald, Fabian Altekürer, and Paul Hagemann. Generative sliced MMD flows with riesz kernels. In *The Twelfth International Conference on Learning Representations*, 2024b. URL <https://openreview.net/forum?id=VdkGRV1vcf>.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc., 2020.
- Richard Jordan, David Kinderlehrer, and Felix Otto. The variational formulation of the Fokker–Planck equation. *SIAM Journal on Mathematical Analysis*, 29(1):1–17, 1998.
- Mark Kac. Foundations of kinetic theory. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, volume 3, pages 171–197, 1956.
- Achim Klenke. *Probability Theory: A Comprehensive Course*. Springer Science & Business Media, 2013.
- Benoit Kloeckner. Approximation by finitely supported measures. *ESAIM: Control, Optimisation and Calculus of Variations*, 18(2):343–359, 2012.
- Anna Korba, Adil Salim, Michael Arbel, Giulia Luise, and Arthur Gretton. A non-asymptotic analysis for Stein variational gradient descent. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 4672–4682. Curran Associates, Inc., 2020.

- Anna Korba, Pierre-Cyril Aubin-Frankowski, Szymon Majewski, and Pierre Ablin. Kernel Stein discrepancy descent. In M. Meila and T. Zhang, editors, *International Conference on Machine Learning*, pages 5719–5730. PMLR, 2021.
- A. Ya. Kruger. On Fréchet subdifferentials. *Journal of Mathematical Sciences*, 116(3): 3325–3358, 2003.
- Marc Lambert, Sinho Chewi, Francis Bach, Silvère Bonnabel, and Philippe Rigollet. Variational inference via Wasserstein gradient flows. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 14434–14447. Curran Associates, Inc., 2022.
- Jing Lei. Convergence and concentration of empirical measures under Wasserstein distance in unbounded functional spaces. *Bernoulli*, 26(1):767–798, 2020.
- Chun-Liang Li, Wei-Cheng Chang, Yu Cheng, Yiming Yang, and Barnabas Poczos. MMD GAN: Towards deeper understanding of moment matching network. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- Lingxiao Li, Qiang Liu, Anna Korba, Mikhail Yurochkin, and Justin Solomon. Sampling with mollified interaction energy descent. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=zWy7dq0ce1>.
- Tengyuan Liang and Hai Tran-Bach. Mehler’s formula, branching process, and compositional kernels of deep neural networks. *Journal of the American Statistical Association*, 117(539): 1324–1337, 2022.
- Qiang Liu. Stein variational gradient descent as gradient flow. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- Qiang Liu and Dilin Wang. Stein variational gradient descent: A general purpose Bayesian inference algorithm. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- Song Liu, Jiahao Yu, Jack Simons, Mingxuan Yi, and Mark Beaumont. Minimizing f -divergences by interpolating velocity fields. In R. Salakhutdinov, Z. Kolter, K. Heller, A. Weller, N. Oliver, J. Scarlett, and F. Berkenkamp, editors, *International Conference on Machine Learning*, pages 32308–32331. PMLR, 2024a.
- Tianle Liu, Promit Ghosal, Krishna Balasubramanian, and Natesh Pillai. Towards understanding the dynamics of Gaussian-Stein variational gradient descent. In *Advances in Neural Information Processing Systems*, volume 36, 2024b.
- Antoine Liutkus, Umut Simsekli, Szymon Majewski, Alain Durmus, and Fabian-Robert Stöter. Sliced-Wasserstein flows: Nonparametric generative modeling via optimal transport

- and diffusions. In K. Chaudhuri and R. Salakhutdinov, editors, *International Conference on Machine Learning*, pages 4104–4113. PMLR, 2019.
- Sebastian Mika, Gunnar Rätsch, Jason Weston, Bernhard Schölkopf, and Klaus-Robert Müller. Fisher discriminant analysis with kernels. In *Neural Networks for Signal Processing IX: Proceedings of the 1999 IEEE Signal Processing Society Workshop (Cat. No.98TH8468)*, pages 41–48, 1999.
- XuanLong Nguyen, Martin J. Wainwright, and Michael I. Jordan. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory*, 56(11):5847–5861, 2010.
- Atsushi Nitanda, Denny Wu, and Taiji Suzuki. Convex analysis of the mean field Langevin dynamics. In G. Camps-Valls, F. J. R. Ruiz, and I. Valera, editors, *International Conference on Artificial Intelligence and Statistics*, pages 9741–9757. PMLR, 2022.
- Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f -GAN: Training generative neural samplers using variational divergence minimization. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- Shin-ichi Ohta and Asuka Takatsu. Displacement convexity of generalized relative entropies. *Advances in Mathematics*, 228(3):1742–1787, 2011.
- Victor M Panaretos and Yoav Zemel. *An Invitation to Statistics in Wasserstein Space*. Springer, 2020.
- Loucas Pillaud-Vivien, Francis Bach, Tony Lelièvre, Alessandro Rudi, and Gabriel Stoltz. Statistical estimation of the Poincaré constant and application to sampling multimodal distributions. In S. Chiappa and R. Calandra, editors, *International Conference on Artificial Intelligence and Statistics*, pages 2753–2763. PMLR, 2020.
- Alfréd Rényi. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, volume 1: Contributions to the Theory of Statistics*, volume 4, pages 547–562. University of California Press, 1961.
- Walter Rudin. *Principles of Mathematical Analysis*, volume 3. McGraw-Hill New York, 1976.
- Filippo Santambrogio. {Euclidean, metric, and Wasserstein} gradient flows: An overview. *Bulletin of Mathematical Sciences*, 7:87–154, 2017.
- Bernhard Schölkopf and Alexander J Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT press, 2002.
- Dino Sejdinovic, Bharath Sriperumbudur, Arthur Gretton, and Kenji Fukumizu. Equivalence of distance-based and RKHS-based statistics in hypothesis testing. *The Annals of Statistics*, 41(5):2263–2291, 2013.

- Jiaxin Shi and Lester Mackey. A finite-particle convergence rate for Stein variational gradient descent. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, editors, *Advances in Neural Information Processing Systems*, volume 36. Curran Associates, Inc., 2024.
- Tao Shi, Mikhail Belkin, and Bin Yu. Data spectroscopy: Eigenspaces of convolution operators and clustering. *The Annals of Statistics*, pages 3960–3984, 2009.
- Carl-Johann Simon-Gabriel, Alessandro Barp, Bernhard Schölkopf, and Lester Mackey. Metrizing weak convergence with maximum mean discrepancies. *Journal of Machine Learning Research*, 24(184):1–20, 2023.
- Jack Simons, Song Liu, and Mark Beaumont. Variational likelihood-free gradient descent. In *Fourth Symposium on Advances in Approximate Bayesian Inference*, 2022.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.
- Bharath K. Sriperumbudur. On the optimal estimation of probability measures in weak and strong topologies. *Bernoulli*, 22(3):1839–1893, 2016.
- Bharath K. Sriperumbudur, Arthur Gretton, Kenji Fukumizu, Bernhard Schölkopf, and Gert R. G. Lanckriet. Hilbert space embeddings and metrics on probability measures. *Journal of Machine Learning Research*, 11:1517–1561, 2010.
- Bharath K. Sriperumbudur, Kenji Fukumizu, and Gert R. G. Lanckriet. Universality, characteristic kernels and RKHS embedding of measures. *Journal of Machine Learning Research*, 12(7):2389–2410, 2011.
- Viktor Stein, Sebastian Neumayer, Nicolaj Rux, and Gabriele Steidl. Wasserstein gradient flows for Moreau envelopes of f -divergences in reproducing kernel Hilbert spaces. *Analysis and Applications*, pages 1–45, 2025.
- Ingo Steinwart and Andreas Christmann. *Support Vector Machines*. Springer Science & Business Media, 2008.
- Ingo Steinwart and Clint Scovel. Mercer’s theorem on general domains: On the interaction between measures, kernels, and RKHSs. *Constructive Approximation*, 35:363–417, 2012.
- Taiji Suzuki, Atsushi Nitanda, and Denny Wu. Uniform-in-time propagation of chaos for the mean-field gradient Langevin dynamics. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=_JSUk9TBUn.
- Nicolas Garcia Trillos and Daniel Sanz-Alonso. The Bayesian update: Variational formulations and gradient flows. *Bayesian Analysis*, 15(1):29–56, 2020.
- Tim Van Erven and Peter Harremoës. Rényi divergence and Kullback-Leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014.

Santosh Vempala and Andre Wibisono. Rapid convergence of the unadjusted Langevin algorithm: Isoperimetry suffices. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

Cédric Villani et al. *Optimal Transport: Old and New*, volume 338. Springer, 2009.

Holger Wendland. *Scattered Data Approximation*, volume 17. Cambridge university press, 2004.

Appendix A: Further Background on $(\mathcal{P}_2(\mathbb{R}^d), W_2)$

Let μ and π be two probability measures in $\mathcal{P}_2(\mathbb{R}^d)$ and let $\Pi(\mu, \pi)$ denote the set of all admissible transport plans between μ and π , i.e., $\Pi(\mu, \pi) = \{\Gamma \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d); (\text{proj}_1)_\# \Gamma = \mu, (\text{proj}_2)_\# \Gamma = \pi\}$, where proj_1 and proj_2 respectively stand for projection maps $(x, y) \mapsto x$ and $(x, y) \mapsto y$, and $\#$ is the pushforward operator. The Wasserstein-2 distance between μ and π is then defined as

$$W_2(\mu, \pi) = \left(\inf_{\Gamma \in \Pi(\mu, \pi)} \int \|x - y\|^2 d\Gamma(x, y) \right)^{\frac{1}{2}},$$

and $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ is a metric space called the Wasserstein space (Panaretos and Zemel, 2020). Brenier’s theorem guarantees that if μ is an absolutely continuous measure, then the optimal transport map is unique and is of the form $\Gamma^* = (\text{I}, T)_\# \mu$, i.e., $T_\# \mu = \pi$ (Santambrogio, 2017). T can also be expressed as $T(x) = x + \nabla \phi(x)$, where ϕ is known as the Kantorovich potential function and is differentiable μ -a.e.

For an absolutely continuous μ and the optimal transport plan T such that $T_\# \mu = \pi$, the shortest path $(\rho_t)_{0 \leq t \leq 1}$ from μ to π is called the (Wasserstein) geodesic given by the following form:

$$\rho_t = ((1 - t)\text{I} + tT)_\# \mu = (\text{I} + t\nabla \phi)_\# \mu.$$

Therefore, in this paper, we always use $(\text{I} + t\nabla \phi)_\# \mu$ with $\phi \in C_c^\infty(\mathbb{R}^d)^8$ to define a geodesic curve that starts at μ . Define $\varphi_t : \mathbb{R}^d \rightarrow \mathbb{R}^d, x \mapsto x + t\nabla \phi(x)$, then $\omega_t : \mathbb{R}^d \rightarrow \mathbb{R}^d, x \mapsto [\nabla \phi \circ \varphi_t^{-1}](x)$ becomes the optimal transport map from ρ_t to π . Notice that $\|\omega_t\|_{L^2(\rho_t)} = \|\nabla \phi\|_{L^2(\mu)}$ for all $t \in [0, 1]$, so $(\rho_t)_{0 \leq t \leq 1}$ is also a constant-speed geodesic. The notion of a constant-speed geodesic is crucial in the introduction of geodesic convexity below.

A functional $\mathcal{F} : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ is geodesically convex if for any μ and π , the following inequality holds:

$$\mathcal{F}(\rho_t) \leq (1 - t)\mathcal{F}(\mu) + t\mathcal{F}(\pi), \quad \forall t \in [0, 1], \quad (\text{A.1})$$

where $(\rho_t)_{t \in [0, 1]}$ is the constant-speed geodesic between μ and π . The geodesic convexity of \mathcal{F} can be equivalently characterized through the Wasserstein Hessian (Villani et al., 2009). The geodesic convexity ought not to be confused with mixture convexity, which replaces displacement geodesic ρ_t in (A.1) with the mixture geodesic $\nu_t = (1 - t)\mu + t\pi$.

8. ϕ is compactly supported because the tangent space of $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ is $\overline{\{\nabla \psi, \psi \in C_c^\infty(\mathbb{R}^d)\}}^{L^2(\mu)}$ (Ambrosio et al., 2005, Definition 8.4.1).

Appendix B: Auxiliary Results

In this appendix, we collect all technical results required to prove the main results of the paper.

Lemma B.1 *For $\mu \ll \pi$, χ^2 -divergence admits the following variational form:*

$$\chi^2(\mu||\pi) = \sup_{h \in L^2(\pi)} \int h d\mu - \int \left(h + \frac{1}{4} h^2 \right) d\pi,$$

where it is sufficient to restrict the variational set to $L^2(\pi)$ in contrast to the set of all measurable functions for general f -divergences (Nowozin et al., 2016; Nguyen et al., 2010).

Proof For $\mu \ll \pi$, we have:

$$\begin{aligned} \chi^2(\mu||\pi) &= \sup_h \left\{ \int h d\mu - \int \left(\frac{h^2}{4} + h \right) d\pi \right\} \\ &= \sup_h \left\{ \int h \frac{d\mu}{d\pi} d\pi - \int \left(\frac{h^2}{4} + h \right) d\pi \right\} = \sup_h \left\{ \int \left(h \frac{d\mu}{d\pi} - \frac{h^2}{4} - h \right) d\pi \right\} \\ &= - \inf_h \left\{ \int \left[\frac{h^2}{4} - h \left(\frac{d\mu}{d\pi} - 1 \right) \right] d\pi \right\} \\ &= - \inf_h \left\{ \int \left[\frac{h}{2} - \left(\frac{d\mu}{d\pi} - 1 \right) \right]^2 d\pi \right\} + \int \left(\frac{d\mu}{d\pi} - 1 \right)^2 d\pi. \end{aligned}$$

Clearly, the above equation is minimized at $h^* = 2(\frac{d\mu}{d\pi} - 1)$ and $\chi^2(\mu||\pi) = \int (\frac{d\mu}{d\pi} - 1)^2 d\pi$ which is finite if and only if $\frac{d\mu}{d\pi} - 1 \in L^2(\pi)$. Therefore, it is sufficient to consider the above maximization over $L^2(\pi)$. \blacksquare

Lemma B.2 *Under Assumption 1 and 2, the mappings $x \mapsto k(x, \cdot)$ and $x \mapsto \nabla_1 k(x, \cdot)$ are differentiable and Lipschitz:*

$$\begin{aligned} \|k(x, \cdot) - k(y, \cdot)\|_{\mathcal{H}} &\leq \sqrt{K_{1d}} \|x - y\|, \\ \|\nabla_1 k(x, \cdot) - \nabla_1 k(y, \cdot)\|_{\mathcal{H}^d} &\leq \sqrt{K_{2d}} \|x - y\|. \end{aligned}$$

Proof This is Lemma 7 from Glaser et al. (2021). \blacksquare

Lemma B.3 *Under Assumption 1 and 2, the regularized kernel \tilde{k} defined in (6) satisfies the following properties:*

1. $|\tilde{k}(x, y)| \leq \frac{K}{\lambda}$;
2. $\partial_i \tilde{k}(x, y) = \langle (\Sigma_\pi + \lambda \mathbf{I})^{-1} \partial_i k(x, \cdot), k(y, \cdot) \rangle_{\mathcal{H}}$;
3. $\|\nabla_1 \tilde{k}(x, y)\|^2 = \sum_{i=1}^d \partial_i \tilde{k}(x, y)^2 \leq \frac{KK_{1d}}{\lambda^2}$;

4. $\partial_i \partial_{i+d} \tilde{k}(x, y) = \langle (\Sigma_\pi + \lambda I)^{-1} \partial_i k(x, \cdot), \partial_i k(y, \cdot) \rangle_{\mathcal{H}};$
5. $\|\nabla_1 \nabla_2 \tilde{k}(x, y)\|_F^2 = \sum_{i=1}^d \partial_i \partial_{i+d} \tilde{k}(x, y)^2 \leq \frac{K_{1d}}{\lambda^2};$
6. $\partial_i \partial_j \tilde{k}(x, y) = \langle (\Sigma_\pi + \lambda I)^{-1} \partial_i \partial_j k(x, \cdot), k(y, \cdot) \rangle_{\mathcal{H}};$
7. $\|\mathbf{H}_1 \tilde{k}(x, y)\|_F^2 = \sum_{i,j=1}^d \partial_i \partial_j \tilde{k}(x, y)^2 \leq \frac{KK_{2d}}{\lambda^2};$
8. $\|\nabla_1 \tilde{k}(x, x') - \nabla_1 \tilde{k}(y, y')\| \leq \frac{\sqrt{KK_{2d}}}{\lambda} (\|x - y\| + \|x' - y'\|).$

Proof Notice that

$$\tilde{k}(x, y) = \left\langle (\Sigma_\pi + \lambda I)^{-1} k(x, \cdot), k(y, \cdot) \right\rangle_{\mathcal{H}} \leq \frac{1}{\lambda} \|k(x, \cdot)\|_{\mathcal{H}} \|k(y, \cdot)\|_{\mathcal{H}} \leq \frac{K}{\lambda},$$

so the first bullet point is proved. Before we prove the second bullet point, we first prove the differentiability of $x \mapsto \tilde{k}(x, y)$. For $i \in \{1, \dots, d\}$, consider $h \in \mathbb{R}$ and denote $\Delta_i \in \mathbb{R}^d$ as a vector of all 0 except the value at i being equal to h . Then, for any $y \in \mathbb{R}^d$,

$$\begin{aligned} \lim_{h \rightarrow 0} \frac{\tilde{k}(x + \Delta_i, y) - \tilde{k}(x, y)}{h} &= \lim_{h \rightarrow 0} \frac{\langle (\Sigma_\pi + \lambda I)^{-1} (k(x + \Delta_i, \cdot) - k(x, \cdot)), k(y, \cdot) \rangle_{\mathcal{H}}}{h} \\ &\leq \lim_{h \rightarrow 0} \frac{\frac{\sqrt{K}}{\lambda} \|k(x, \cdot) - k(x + \Delta_i, \cdot)\|_{\mathcal{H}}}{h} \leq \lim_{h \rightarrow 0} \frac{\frac{\sqrt{K}}{\lambda} \sqrt{K_{1d}} \|\Delta_i\|}{h} = \frac{\sqrt{K}}{\lambda} \sqrt{K_{1d}}. \end{aligned}$$

So $x \mapsto \tilde{k}(x, y)$ is differentiable for any $y \in \mathbb{R}^d$. Since the kernel k is differentiable per Assumption 2, for any $f \in \mathcal{H}$, $\partial_{x_i} f(x) = \langle \partial_i k(x, \cdot), f \rangle_{\mathcal{H}}$. Hence,

$$\partial_i \tilde{k}(x, y) = \partial_{x_i} \left\langle k(x, \cdot), (\Sigma_\pi + \lambda I)^{-1} k(y, \cdot) \right\rangle_{\mathcal{H}} = \left\langle \partial_i k(x, \cdot), (\Sigma_\pi + \lambda I)^{-1} k(y, \cdot) \right\rangle_{\mathcal{H}}.$$

So the second bullet point is proved.

Next, notice that

$$\begin{aligned} \|\nabla_1 \tilde{k}(x, y)\|^2 &= \sum_{i=1}^d \left\langle \partial_i k(x, \cdot), (\Sigma_\pi + \lambda I)^{-1} k(y, \cdot) \right\rangle_{\mathcal{H}}^2 \leq \frac{1}{\lambda^2} \sum_{i=1}^d \|\partial_i k(x, \cdot)\|_{\mathcal{H}}^2 \|k(y, \cdot)\|_{\mathcal{H}}^2 \\ &\leq \frac{1}{\lambda^2} \|\nabla_1 k(x, \cdot)\|_{\mathcal{H}^d}^2 \|k(y, \cdot)\|_{\mathcal{H}}^2 \leq \frac{KK_{1d}}{\lambda^2}. \end{aligned}$$

So the third bullet point is proved. Similar arguments above lead to bullet points 4 to 7. Finally, to prove bullet point 8, notice that

$$\begin{aligned} \|\nabla_1 \tilde{k}(x, x') - \nabla_1 \tilde{k}(y, x')\|^2 &= \sum_{i=1}^d \left\langle (\partial_i k(x, \cdot) - \partial_i k(y, \cdot)), (\Sigma_\pi + \lambda I)^{-1} k(x', \cdot) \right\rangle_{\mathcal{H}}^2 \\ &\leq \frac{1}{\lambda^2} \sum_{j=1}^d \|\partial_i k(x, \cdot) - \partial_i k(y, \cdot)\|_{\mathcal{H}}^2 \|k(x', \cdot)\|_{\mathcal{H}}^2 \leq \frac{1}{\lambda^2} \|\nabla_1 k(x, \cdot) - \nabla_1 k(y, \cdot)\|_{\mathcal{H}^d}^2 \|k(x', \cdot)\|_{\mathcal{H}}^2 \\ &\leq \frac{KK_{2d}}{\lambda^2} \|x - y\|^2, \end{aligned}$$

where the last inequality uses Lemma B.2. Therefore,

$$\begin{aligned} \left\| \nabla_1 \tilde{k}(x, x') - \nabla_1 \tilde{k}(y, y') \right\| &\leq \left\| \nabla_1 \tilde{k}(x, x') - \nabla_1 \tilde{k}(y, x') \right\| + \left\| \nabla_1 \tilde{k}(y, x') - \nabla_1 \tilde{k}(y, y') \right\| \\ &\leq \frac{\sqrt{KK_{2d}}}{\lambda} (\|x - y\| + \|x' - y'\|). \end{aligned}$$

So the bullet point 8 is proved. ■

Lemma B.4 *For any two distributions $\mu_0, \mu_1 \in \mathcal{P}_2(\mathbb{R}^d)$, with associated DrMMD witness functions $h_{\mu_0, \pi}^*, h_{\mu_1, \pi}^*$ defined in Proposition 3.2, we have*

$$\|h_{\mu_1, \pi}^* - h_{\mu_0, \pi}^*\|_{\mathcal{H}} \leq \frac{2\sqrt{K_{1d}}}{\lambda} W_2(\mu_0, \mu_1), \quad \text{and} \quad \|h_{\mu_0, \pi}^*\|_{\mathcal{H}} \leq \frac{2\sqrt{K}}{\lambda}.$$

Both the witness function $h_{\mu, \pi}^*$ and its gradient $\nabla h_{\mu, \pi}^*$ are Lipschitz continuous, i.e.,

$$\begin{aligned} |h_{\mu_1, \pi}^*(x) - h_{\mu_0, \pi}^*(y)| &\leq L(W_2(\mu_0, \mu_1) + \|x - y\|); \\ \|\nabla h_{\mu_1, \pi}^*(x) - \nabla h_{\mu_0, \pi}^*(y)\| &\leq L(W_2(\mu_0, \mu_1) + \|x - y\|), \end{aligned}$$

where the constant $L = \frac{1}{\lambda} \max \{2\sqrt{KK_{1d}}, 2\sqrt{KK_{2d}}, 2K_{1d}\}$.

Proof Let $\gamma \in \Gamma(\mu_0, \mu_1)$ be the optimal coupling between μ_0 and μ_1 . Then

$$\begin{aligned} \|h_{\mu_1, \pi}^* - h_{\mu_0, \pi}^*\|_{\mathcal{H}}^2 &= 4 \left\| \left(\Sigma_{\pi} + \lambda I \right)^{-1} \left(\int k(x, \cdot) d(\mu_1 - \mu_0) \right) \right\|_{\mathcal{H}}^2 \\ &\leq \frac{4}{\lambda^2} \left\| \int k(x, \cdot) - k(y, \cdot) d\gamma(x, y) \right\|_{\mathcal{H}}^2 \leq \frac{4K_{1d}}{\lambda^2} \int \|x - y\|^2 d\gamma(x, y) = \frac{4K_{1d}}{\lambda^2} W_2^2(\mu_0, \mu_1), \end{aligned}$$

where the first inequality holds because Σ_{π} is a positive and self-adjoint operator, and the second inequality uses Lemma B.2. Also note that

$$\|h_{\mu_0, \pi}^*\|_{\mathcal{H}} = \left\| 2 \left(\Sigma_{\pi} + \lambda I \right)^{-1} \left(\int k(x, \cdot) d\mu_0 - \int k(x, \cdot) d\pi \right) \right\|_{\mathcal{H}} \leq \frac{2\sqrt{K}}{\lambda}.$$

So the first part has been proved. Furthermore, note that

$$\begin{aligned} |h_{\mu_1, \pi}^*(x) - h_{\mu_0, \pi}^*(y)| &\leq |h_{\mu_1, \pi}^*(x) - h_{\mu_0, \pi}^*(x)| + |h_{\mu_0, \pi}^*(x) - h_{\mu_0, \pi}^*(y)| \\ &\leq \|k(x, \cdot) - k(y, \cdot)\|_{\mathcal{H}} \|h_{\mu_1, \pi}^* - h_{\mu_0, \pi}^*\|_{\mathcal{H}} + \|k(x, \cdot) - k(y, \cdot)\|_{\mathcal{H}} \|h_{\mu_0, \pi}^*\|_{\mathcal{H}} \\ &\leq \frac{2\sqrt{KK_{1d}}}{\lambda} W_2(\mu_0, \mu_1) + \frac{2\sqrt{KK_{1d}}}{\lambda} \|x - y\| \\ &\leq \frac{2\sqrt{KK_{1d}}}{\lambda} (W_2(\mu_0, \mu_1) + \|x - y\|) \leq L(W_2(\mu_0, \mu_1) + \|x - y\|) \end{aligned}$$

and

$$\|\nabla h_{\mu_1, \pi}^*(x) - \nabla h_{\mu_0, \pi}^*(y)\| \leq \|\nabla h_{\mu_1, \pi}^*(x) - \nabla h_{\mu_0, \pi}^*(x)\| + \|\nabla h_{\mu_0, \pi}^*(x) - \nabla h_{\mu_0, \pi}^*(y)\|$$

$$\begin{aligned}
&\leq \|\nabla_1 k(x, \cdot)\|_{\mathcal{H}^d} \|h_{\mu_1, \pi}^* - h_{\mu_0, \pi}^*\|_{\mathcal{H}} + \|\nabla_1 k(x, \cdot) - \nabla_1 k(y, \cdot)\|_{\mathcal{H}^d} \|h_{\mu_0, \pi}^*\|_{\mathcal{H}} \\
&\leq \frac{2K_{1d}}{\lambda} W_2(\mu_0, \mu_1) + \frac{2\sqrt{K}K_{2d}}{\lambda} \|x - y\| \leq L(W_2(\mu_0, \mu_1) + \|x - y\|)
\end{aligned}$$

and the result follows. \blacksquare

Lemma B.5 *Given two probability measures $\mu \ll \pi$ that are both absolutely continuous with respect to Lebesgue measure. If $\frac{d\mu}{d\pi} - 1 \in \text{Ran}(\mathcal{T}_\pi^r)$ with $r > 0$, i.e., there exists $q \in L^2(\pi)$ such that $\frac{d\mu}{d\pi} - 1 = \mathcal{T}_\pi^r q$, then*

$$\left\| h - 2 \left(\frac{d\mu}{d\pi} - 1 \right) \right\|_{L^2(\pi)} \leq 2\lambda^r \|q\|_{L^2(\pi)},$$

where $h = 2(\Sigma_\pi + \lambda I)^{-1}(m_\pi - m_\mu)$ is defined in Proposition 3.2.

Proof Given the assumption that $\frac{d\mu}{d\pi} - 1 \in \text{Ran}(\mathcal{T}_\pi^r)$ with $r > 0$, there exists $q \in L^2(\pi)$ such that $\frac{d\mu}{d\pi} - 1 = \mathcal{T}_\pi^r q$ and $\langle \frac{d\mu}{d\pi} - 1, e_i \rangle_{L^2(\pi)} = \varrho_i^r \langle q, e_i \rangle_{L^2(\pi)}$. Since $\mu \ll \pi$, the Mercer decomposition of k in (1) holds for any x in the support of μ and in the support of π ,

$$m_\mu - m_\pi = \int k(x, \cdot) d(\mu - \pi)(x) = \sum_{i \geq 1} \varrho_i \left(\int e_i(x) d(\mu - \pi)(x) \right) e_i = \varrho_i \left\langle \frac{d\mu}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} e_i. \quad (\text{B.1})$$

So we have,

$$\begin{aligned}
&\left\| h - 2 \left(\frac{d\mu}{d\pi} - 1 \right) \right\|_{L^2(\pi)} = 2 \left\| (\Sigma_\pi + \lambda I)^{-1} (m_\mu - m_\pi) - \left(\frac{d\mu}{d\pi} - 1 \right) \right\|_{L^2(\pi)} \\
&= 2 \left\| \sum_{i \geq 1} \frac{\varrho_i}{\varrho_i + \lambda} \left\langle \frac{d\mu}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} e_i - \sum_{i \geq 1} \left\langle \frac{d\mu}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} e_i \right\|_{L^2(\pi)} \\
&= 2 \left\| \sum_{i \geq 1} \frac{\lambda}{\varrho_i + \lambda} \left\langle \frac{d\mu}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} e_i \right\|_{L^2(\pi)} = 2 \left\| \sum_{i \geq 1} \frac{\lambda \varrho_i^r}{\varrho_i + \lambda} \langle q, e_i \rangle_{L^2(\pi)} e_i \right\|_{L^2(\pi)} \\
&\leq 2\lambda^r \|q\|_{L^2(\pi)},
\end{aligned}$$

where the last inequality is obtained by using

$$\frac{\lambda \varrho_i^r}{\varrho_i + \lambda} = \left(\frac{\varrho_i}{\varrho_i + \lambda} \right)^r \left(\frac{\lambda}{\varrho_i + \lambda} \right)^{1-r} \lambda^r \leq \lambda^r.$$

Hence the proof. \blacksquare

Lemma B.6 *Let $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ and $\phi \in C_c^\infty(\mathbb{R}^d)$. Consider the path $(\rho_s)_{0 \leq s \leq 1}$ from ρ to $(I + \nabla \phi)_\# \rho$ given by $\rho_s = (I + s \nabla \phi)_\# \rho$. Define $\varphi_s : \mathbb{R}^d \rightarrow \mathbb{R}^d, x \mapsto x + s \nabla \phi(x)$. Let \tilde{k} be the regularized kernel defined in (6) along with its associated RKHS $\tilde{\mathcal{H}}$. The mapping $s \mapsto \text{DrMMD}(\rho_s \| \pi)$ is continuous and differentiable, and its first-order time derivative is given by*

$$\begin{aligned} & \frac{d}{ds} \text{DrMMD}(\rho_s \| \pi) \\ &= 2(1 + \lambda) \int \nabla \phi(x)^\top \left(\int \nabla_1 \tilde{k}(\varphi_s(x), \varphi_s(z)) d\rho(z) - \int \nabla_1 \tilde{k}(\varphi_s(x), z) d\pi(z) \right) d\rho(x). \end{aligned} \quad (\text{B.2})$$

Moreover,

$$\begin{aligned} & \left. \frac{d}{ds} \right|_{s=0} \text{DrMMD}(\rho_s \| \pi) \\ &= 2(1 + \lambda) \int \nabla \phi(x)^\top \left(\int \nabla_1 \tilde{k}(x, z) d\rho(z) - \int \nabla_1 \tilde{k}(x, z) d\pi(z) \right) d\rho(x). \end{aligned} \quad (\text{B.3})$$

Additionally, the mapping $s \rightarrow \frac{d}{ds} \text{DrMMD}(\rho_s \| \pi)$ is continuous and differentiable, and the second-order time derivative of $\text{DrMMD}(\rho_s \| \pi)$ is given by

$$\begin{aligned} & \frac{d^2}{ds^2} \text{DrMMD}(\rho_s \| \pi) = 2(1 + \lambda) \iint \nabla \phi(x)^\top \nabla_1 \nabla_2 \tilde{k}(\varphi_s(x), \varphi_s(z)) \nabla \phi(z) d\rho(x) d\rho(z) \\ & + 2(1 + \lambda) \int \nabla \phi(x)^\top \left(\int \mathbf{H}_1 \tilde{k}(\varphi_s(x), \varphi_s(z)) d\rho(z) - \int \mathbf{H}_1 \tilde{k}(\varphi_s(x), z) d\pi(z) \right) \nabla \phi(x) d\rho(x), \end{aligned} \quad (\text{B.4})$$

with

$$\begin{aligned} & \left. \frac{d^2}{ds^2} \right|_{s=0} \text{DrMMD}(\rho_s \| \pi) = 2(1 + \lambda) \iint \nabla \phi(x)^\top \nabla_1 \nabla_2 \tilde{k}(x, z) \nabla \phi(z) d\rho(x) d\rho(z) \\ & + 2(1 + \lambda) \int \nabla \phi(x)^\top \left(\int \mathbf{H}_1 \tilde{k}(x, z) d\rho(z) - \int \mathbf{H}_1 \tilde{k}(x, z) d\pi(z) \right) \nabla \phi(x) d\rho(x). \end{aligned} \quad (\text{B.5})$$

Proof Recall that DrMMD is, up to a multiplicative factor of $(1 + \lambda)$, MMD^2 of the regularized kernel \tilde{k} defined in (6) and the associated RKHS $\tilde{\mathcal{H}}$. From Lemma B.3, we know that assumptions (A) and (B) of Arbel et al. (2019) are satisfied for the regularized kernel \tilde{k} , so using Lemma 22 and Lemma 23 from Arbel et al. (2019), (B.2) and (B.4) are proved. Then (B.3) and (B.5) are subsequently proved by taking $s = 0$. \blacksquare

Lemma B.7 *For $\mu_0 \ll \pi$, define $h = 2(\Sigma_\pi + \lambda I)^{-1}(m_\pi - m_{\mu_0})$ and $\varphi_t : \mathbb{R}^d \rightarrow \mathbb{R}^d, x \mapsto x - t(1 + \lambda) \nabla h(x)$. Suppose $\pi \propto \exp(-V)$, $\mathbf{H}V \preceq \beta I$, and the step size γ satisfies*

$$2(1 + \lambda) \gamma \sqrt{\chi^2(\mu_0 \| \pi) \frac{K_{2d}}{\lambda}} \leq \frac{\zeta - 1}{\zeta}$$

for some constant $1 < \zeta < 2$. Then, for $t \in [0, \gamma]$, the following inequalities hold,

$$\int \nabla h(x)^\top \mathbf{H}V(\varphi_t(x)) \nabla h(x) d\mu_0(x) \leq 4\beta \chi^2(\mu_0 \| \pi) \frac{K_{1d}}{\lambda}; \quad (\text{B.6})$$

$$\int \left\| \mathbf{H}h(x) (I - t(1 + \lambda) \mathbf{H}h(x))^{-1} \right\|_F^2 d\mu_0(x) \leq 4\zeta^2 \chi^2(\mu_0 \| \pi) \frac{K_{2d}}{\lambda}. \quad (\text{B.7})$$

Proof First, recall (B.1),

$$\begin{aligned}\|\nabla h(x)\|^2 &= 4 \left\| \nabla (\Sigma_\pi + \lambda \mathbf{I})^{-1} (m_\pi - m_{\mu_0})(x) \right\|^2 \\ &= 4 \left\| \nabla \left(\sum_{i \geq 1} \frac{\varrho_i}{\varrho_i + \lambda} \left\langle \frac{d\mu_0}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} e_i(x) \right) \right\|^2.\end{aligned}\quad (\text{B.8})$$

For any $j \in \{1, \dots, d\}$, consider

$$\begin{aligned}g_{M_0}(x) &:= \sum_{i \geq M_0} \left| \varrho_i^{1/2} \left\langle \frac{d\mu_0}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} \partial_j e_i(x) \right| \\ &\leq \left(\sum_{i \geq M_0} \varrho_i \partial_j e_i(x)^2 \right)^{\frac{1}{2}} \left(\sum_{i \geq M_0} \left\langle \frac{d\mu_0}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)}^2 \right)^{\frac{1}{2}} \\ &\leq \left(\sum_{i \geq M_0} \varrho_i \partial_j e_i(x)^2 \right)^{\frac{1}{2}} \left\| \frac{d\mu_0}{d\pi} - 1 \right\|_{L^2(\pi)}.\end{aligned}$$

Since

$$\begin{aligned}\sum_{i \geq 1} \sum_{j=1}^d \varrho_i (\partial_j e_i(x))^2 &= \sum_{i \geq 1} \sum_{j=1}^d \varrho_i \langle \partial_j k(x, \cdot), e_i \rangle_{\mathcal{H}}^2 = \sum_{i \geq 1} \sum_{j=1}^d \langle \partial_j k(x, \cdot), \sqrt{\varrho_i} e_i \rangle_{\mathcal{H}}^2 \\ &= \sum_{j=1}^d \|\partial_j k(x, \cdot)\|_{\mathcal{H}}^2 = \|\nabla_1 k(x, \cdot)\|_{\mathcal{H}}^2 \leq K_{1d},\end{aligned}\quad (\text{B.9})$$

so $\sum_{i \geq M_0} \varrho_i (\partial_j e_i(x))^2$ converges uniformly to 0, and hence $g_{M_0}(x)$ also converges uniformly to 0. Therefore, we are allowed to interchange the derivative and the infinite sum (Rudin, 1976) in (B.8) to achieve,

$$\begin{aligned}\|\nabla h(x)\|^2 &= 4 \left\| \sum_{i \geq 1} \frac{\varrho_i}{\varrho_i + \lambda} \left\langle \frac{d\mu_0}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} \nabla e_i(x) \right\|^2 \\ &\leq 4 \left(\sum_{i \geq 1} \frac{\varrho_i}{(\varrho_i + \lambda)^2} \left\langle \frac{d\mu_0}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)}^2 \right) \left(\sum_{i \geq 1} \varrho_i \|\nabla e_i(x)\|^2 \right) \\ &\leq \frac{4}{\lambda} \left(\sum_{i \geq 1} \left\langle \frac{d\mu_0}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)}^2 \right) \left(\sum_{i \geq 1} \varrho_i \|\nabla e_i(x)\|^2 \right) \leq 4\chi^2(\mu_0\|\pi) \frac{K_{1d}}{\lambda}.\end{aligned}\quad (\text{B.10})$$

The first inequality follows from Cauchy Schwartz, the penultimate inequality follows by noting that $\frac{\varrho_i}{(\varrho_i + \lambda)^2} \leq \frac{1}{\lambda}$, and the last inequality follows from (B.9). Given $\mathbf{H}V \preceq \beta \mathbf{I}$, (B.6) is proved by the following,

$$\int \nabla h(x)^\top \mathbf{H}V(\varphi_t(x)) \nabla h(x) d\mu_0(x) \leq \beta \int \|\nabla h(x)\|^2 d\mu_0(x) \leq 4\beta\chi^2(\mu_0\|\pi) \frac{K_{1d}}{\lambda}.$$

We now turn to proving the second statement. Similarly to (B.10), we have

$$\|\mathbf{H}h(x)\|_F^2 \leq 4\chi^2(\mu_0\|\pi)\frac{K_{2d}}{\lambda}. \quad (\text{B.11})$$

Using $2(1+\lambda)\gamma\sqrt{\chi^2(\mu_0\|\pi)\frac{K_{2d}}{\lambda}} \leq \frac{\zeta-1}{\zeta}$ for some constant $1 < \zeta < 2$, the inverse of $\mathbf{I} - t(1+\lambda)\mathbf{H}h(x)$ can be represented by the Neumann series, and hence

$$\begin{aligned} \|(\mathbf{I} - t(1+\lambda)\mathbf{H}h(x))^{-1}\|_F &\leq \sum_{m \geq 0} \|t(1+\lambda)\mathbf{H}h(x)\|_F^m \leq \sum_{m \geq 0} \left(\gamma(1+\lambda)2\sqrt{\chi^2(\mu_0\|\pi)\frac{K_{2d}}{\lambda}} \right)^m \\ &\leq \sum_{m \geq 0} \left(\frac{\zeta-1}{\zeta} \right)^m = \zeta. \end{aligned} \quad (\text{B.12})$$

Therefore, (B.7) is proved by combining (B.11) and (B.12),

$$\begin{aligned} &\int \left\| \mathbf{H}h(x) (\mathbf{I} - t(1+\lambda)\mathbf{H}h(x))^{-1} \right\|_F^2 d\mu_0(x) \\ &\leq \int \|\mathbf{H}h(x)\|_F^2 \|(\mathbf{I} - t(1+\lambda)\mathbf{H}h(x))^{-1}\|_F^2 d\mu_0(x) \leq 4\zeta^2\chi^2(\mu_0\|\pi)\frac{K_{2d}}{\lambda} \end{aligned}$$

and the result follows. ■

Lemma B.8 *Let \mathcal{H} be a separable Hilbert space, $\xi_1, \dots, \xi_n : \Omega \rightarrow \mathcal{H}$ are n identical independent \mathcal{H} -valued random variables satisfying $\|\xi_i\|_{\mathcal{H}} \leq B$. Then*

$$\mathbb{E} \left\| \frac{1}{n} \sum_{i=1}^n \xi_i - \mathbb{E}[\xi_1] \right\|_{\mathcal{H}} \leq \frac{\sqrt{2\pi}B}{2\sqrt{n}}, \quad \text{and} \quad \mathbb{E} \left\| \frac{1}{n} \sum_{i=1}^n \xi_i - \mathbb{E}[\xi_1] \right\|_{\mathcal{H}}^2 \leq \frac{B^2}{n}.$$

Proof We know from Corollary 6.15 of Steinwart and Christmann (2008) that

$$\mathbb{P} \left(\left\| \frac{1}{n} \sum_{i=1}^n \xi_i - \mathbb{E}[\xi_1] \right\|_{\mathcal{H}} \geq t \right) \leq 2 \exp(-2nt^2/B^2).$$

Then, denote $R := \left\| \frac{1}{n} \sum_{i=1}^n \xi_i - \mathbb{E}[\xi_1] \right\|_{\mathcal{H}}$, we know that

$$\mathbb{E}[R] = \int_0^\infty \mathbb{P}(R \geq t) dt \leq 2 \int_0^\infty \exp(-2nt^2/B^2) dt = \frac{\sqrt{2\pi}B}{2\sqrt{n}}.$$

The other part is proved similarly. ■

Lemma B.9 (Wasserstein Hessian of \mathcal{F}_{χ^2}) *Let $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ and $\phi \in C_c^\infty(\mathbb{R}^d)$. Consider the path $(\rho_s)_{0 \leq s \leq 1}$ from ρ to $(\mathbf{I} + \nabla \phi)_\# \rho$ given by $\rho_s = (\mathbf{I} + s\nabla \phi)_\# \rho$. Define $\varphi_s : \mathbb{R}^d \rightarrow \mathbb{R}^d, x \mapsto x + s\nabla \phi(x)$ and $\omega_s : \mathbb{R}^d \rightarrow \mathbb{R}^d, x \mapsto [\nabla \phi \circ \varphi_s^{-1}](x)$. For $\pi(x) \propto \exp(-V(x))$, the second order derivative of $s \rightarrow \chi^2(\rho_s\|\pi)$ is given by*

$$\frac{d^2}{ds^2} \chi^2(\rho_s\|\pi) = \int \frac{\rho_s(x)}{\pi(x)} \left(\nabla V(x)^\top \omega_s(x) - \nabla \cdot \omega_s(x) \right)^2 d\rho_s(x)$$

$$+ \int \frac{\rho_s(x)}{\pi(x)} \omega_s(x)^\top \mathbf{H} V(x) \omega_s(x) d\rho_s(x) + \int \frac{\rho_s(x)}{\pi(x)} \|\nabla \omega_s(x)\|_F^2 d\rho_s(x). \quad (\text{B.13})$$

Equivalently, the second order derivative of $s \rightarrow \chi^2(\rho_s \|\pi)$ can also be written as

$$\begin{aligned} \frac{d^2}{ds^2} \chi^2(\rho_s \|\pi) &= 2 \int \left(\nabla \cdot (\omega_s(x) \rho_s(x)) \frac{1}{\pi(x)} \right)^2 \pi(x) dx \\ &\quad + 2 \int \omega_s(x)^\top \mathbf{H} \left(\frac{\rho_s(x)}{\pi(x)} \right) \omega_s(x) \rho_s(x) dx. \end{aligned} \quad (\text{B.14})$$

Proof (B.13) is provided in the Example 15.9 of Villani et al. (2009) by taking $m = 2$. Now, we are going to prove (B.14). For ease of notation in the following derivations, we are going to drop the function input x in ρ_s, π , and ω_s . We introduce colors to picture grouping of terms that will carry over during chains of calculation.

In order to prove (B.14), we need to expand the terms in (B.13) accordingly. We denote the three terms in the RHS of (B.13) as (A), (B) and (C). Consider

$$\begin{aligned} (A) &= \int \frac{\rho_s^2}{\pi} \left(\nabla V^\top \omega_s - \nabla \cdot \omega_s \right)^2 dx \\ &= \int \frac{\rho_s^2}{\pi} \left(-\frac{1}{\pi} \nabla \pi^\top \omega_s - \nabla \cdot \omega_s \right)^2 dx \\ &= \underbrace{\int \frac{\rho_s^2}{\pi^3} (\nabla \pi^\top \omega_s)^2 dx}_{(A_1)} + \underbrace{2 \int \frac{\rho_s^2}{\pi^2} (\nabla \pi^\top \omega_s) \nabla \cdot \omega_s dx}_{(A_2)} + \underbrace{\int \frac{\rho_s^2}{\pi} (\nabla \cdot \omega_s)^2 dx}_{(A_3)}. \end{aligned}$$

Then we are going to use integration by parts for (A₂) and (A₃).

$$\begin{aligned} (A_2) &= 2 \int \frac{\rho_s^2}{\pi^2} (\nabla \pi^\top \omega_s) \nabla \cdot \omega_s dx = -2 \int \omega_s^\top \nabla \left(\frac{\rho_s^2}{\pi^2} (\nabla \pi^\top \omega_s) \right) dx \\ &= -4 \int (\omega_s^\top \nabla \rho_s) (\nabla \pi^\top \omega_s) \frac{\rho_s}{\pi^2} dx + 4 \int (\omega_s^\top \nabla \pi)^2 \frac{\rho_s^2}{\pi^3} dx \\ &\quad - 2 \int (\omega_s^\top \mathbf{H} \pi \omega_s) \frac{\rho_s^2}{\pi^2} dx - 2 \int (\omega_s^\top \nabla \omega_s \nabla \pi) \frac{\rho_s^2}{\pi^2} dx, \\ (A_3) &= \int \frac{\rho_s^2}{\pi} (\nabla \cdot \omega_s)^2 dx = \int \nabla \cdot (\omega_s) \nabla \cdot \omega_s \frac{\rho_s^2}{\pi} dx = - \int \omega_s^\top \nabla \left(\nabla \cdot \omega_s \frac{\rho_s^2}{\pi} \right) dx \\ &= \underbrace{- \int \omega_s^\top \nabla (\nabla \cdot \omega_s) \frac{\rho_s^2}{\pi} dx}_{(A_{31})} - \underbrace{\int \nabla \cdot \omega_s (\omega_s^\top \nabla \rho_s) \frac{2\rho_s}{\pi} dx}_{(A_{32})} + \underbrace{\int \nabla \cdot \omega_s (\omega_s^\top \nabla \pi) \frac{\rho_s^2}{\pi^2} dx}_{(A_{33})}. \end{aligned}$$

Furthermore, we use integration by parts for (A₃₂) and (A₃₃), we have

$$\begin{aligned} (A_{32}) &= -2 \int \nabla \cdot \omega_s (\omega_s^\top \nabla \rho_s) \frac{\rho_s}{\pi} dx = 2 \int (\omega_s^\top \nabla \left((\omega_s^\top \nabla \rho_s) \frac{\rho_s}{\pi} \right)) dx \\ &= 2 \int (\omega_s^\top \nabla \omega_s \nabla \rho_s) \frac{\rho_s}{\pi} dx + 2 \int (\omega_s^\top \mathbf{H} \rho_s \omega_s) \frac{\rho_s}{\pi} dx + 2 \int (\omega_s^\top \nabla \rho_s)^2 \frac{1}{\pi} dx \end{aligned}$$

$$\begin{aligned}
& -2 \int (\omega_s^\top \nabla \rho_s)(\omega_s^\top \nabla \pi) \frac{\rho_s}{\pi^2} dx, \\
(A_{33}) &= \int \nabla \cdot \omega_s (\omega_s^\top \nabla \pi) \frac{\rho_s^2}{\pi^2} dx = - \int (\omega_s^\top \nabla \left((\omega_s^\top \nabla \pi) \frac{\rho_s^2}{\pi^2} \right)) dx \\
&= - \int (\omega_s^\top \nabla \omega_s \nabla \pi) \frac{\rho_s^2}{\pi^2} dx - \int (\omega_s^\top \mathbf{H} \pi \omega_s) \frac{\rho_s^2}{\pi^2} dx - 2 \int (\omega_s^\top \nabla \rho_s)(\omega_s^\top \nabla \pi) \frac{\rho_s}{\pi^2} dx \\
&\quad + 2 \int (\omega_s^\top \nabla \pi)^2 \frac{\rho_s^2}{\pi^3} dx.
\end{aligned}$$

Having completed (A), now we turn to (B).

$$(B) = \int \frac{\rho_s^2}{\pi} \omega_s^\top \mathbf{H} V \omega_s dx = - \int \omega_s^\top \mathbf{H} \pi \omega_s \frac{\rho_s^2}{\pi^2} dx + \int (\nabla \pi^\top \omega_s)^2 \frac{\rho_s^2}{\pi^3} dx.$$

So, combining (A), (B) and (C), we have

$$\begin{aligned}
& \frac{d^2}{ds^2} \chi^2(\rho_s \| \pi) = (A) + (B) + (C) = (A_1) + (A_2) + (A_{31}) + (A_{32}) + (A_{33}) + (B) + (C) \\
&= \int \frac{\rho_s^2}{\pi^3} (\nabla \pi^\top \omega_s)^2 dx \\
&\quad - 4 \int (\omega_s^\top \nabla \rho_s)(\nabla \pi^\top \omega_s) \frac{\rho_s}{\pi^2} dx + 4 \int (\omega_s^\top \nabla \pi)^2 \frac{\rho_s^2}{\pi^3} dx - 2 \int (\omega_s^\top \mathbf{H} \pi \omega_s) \frac{\rho_s^2}{\pi^2} dx \\
&\quad - 2 \int (\omega_s^\top \nabla \omega_s \nabla \pi) \frac{\rho_s^2}{\pi^2} dx + (A_{31}) + (A_{32}) + (A_{33}) \\
&\quad - \int (\omega_s^\top \mathbf{H} \pi \omega_s) \frac{\rho_s^2}{\pi^2} dx + \int (\nabla \pi^\top \omega_s)^2 \frac{\rho_s^2}{\pi^3} dx + \int \frac{\rho_s^2}{\pi} \|\nabla \omega_s\|_F^2 dx.
\end{aligned}$$

Since $\omega_s = \nabla \phi \circ \varphi_s^{-1}$ is a mapping from \mathbb{R}^d to \mathbb{R}^d , denote $\omega_{s,i} := [\nabla \phi \circ \varphi_s^{-1}]_i$ which is a mapping from \mathbb{R}^d to \mathbb{R} . Notice that $\omega_{s,i} \partial_j \omega_{s,i}$ vanishes at boundary because $\phi \in C_c^\infty$. Hence,

$$\begin{aligned}
0 &= \sum_{i,j} \int \partial_j \left(\omega_{s,i} \partial_j \omega_{s,i} \frac{\rho_s^2}{\pi} \right) dx \\
&= \sum_{i,j} \int \partial_j \omega_{s,i} \partial_j \omega_{s,i} \frac{\rho_s^2}{\pi} dx + \sum_{i,j} \int \omega_{s,i} \partial_{jj} \omega_{s,i} \frac{\rho_s^2}{\pi} dx \\
&\quad + \sum_{i,j} \int \omega_{s,i} \partial_j \omega_{s,i} \left(\frac{2\rho_s \partial_j \rho_s}{\pi} - \frac{\rho_s^2 \partial_j \pi}{\pi^2} \right) dx \\
&= \int \frac{\rho_s^2}{\pi} \|\nabla \omega_s\|_F^2 dx + \int \frac{\rho_s^2}{\pi} (\omega_s^\top \nabla (\nabla \cdot \omega_s)) dx + \int (\omega_s^\top \nabla \omega_s \nabla \rho_s) \frac{2\rho_s}{\pi} dx \\
&\quad - \int (\omega_s^\top \nabla \omega_s \nabla \pi) \frac{\rho_s^2}{\pi^2} dx.
\end{aligned}$$

Therefore, by replacing $\int \frac{\rho_s^2}{\pi} \|\nabla \omega_s\|_F^2 dx$, and noticing that $-\int \frac{\rho_s^2}{\pi} (\omega_s^\top \nabla (\nabla \cdot \omega_s)) dx$ is exactly (A₃₁), we have the following:

$$\frac{d^2}{ds^2} \chi^2(\rho_s \| \pi) = \int \frac{\rho_s^2}{\pi^3} (\nabla \pi^\top \omega_s)^2 dx - 4 \int (\omega_s^\top \nabla \rho_s)(\nabla \pi^\top \omega_s) \frac{\rho_s}{\pi^2} dx + 4 \int (\omega_s^\top \nabla \pi)^2 \frac{\rho_s^2}{\pi^3} dx$$

$$\begin{aligned}
& -2 \int (\omega_s^\top \mathbf{H} \pi \omega_s) \frac{\rho_s^2}{\pi^2} dx - 2 \int (\omega_s^\top \nabla \omega_s \nabla \pi) \frac{\rho_s^2}{\pi^2} dx + (A_{31}) + (A_{32}) + (A_{33}) \\
& - \int (\omega_s^\top \mathbf{H} \pi \omega_s) \frac{\rho_s^2}{\pi^2} dx + \int (\nabla \pi^\top \omega_s)^2 \frac{\rho_s^2}{\pi^3} dx \\
& - \int \frac{\rho_s^2}{\pi} (\omega_s^\top \nabla (\nabla \cdot \omega_s)) dx - \int (\omega_s^\top \nabla \omega_s \nabla \rho_s) \frac{2\rho_s}{\pi} dx + \int (\omega_s^\top \nabla \omega_s \nabla \pi) \frac{\rho_s^2}{\pi^2} dx.
\end{aligned}$$

Next, we combine the terms and obtain

$$\begin{aligned}
& = 2((A_{31}) + (A_{32}) + (A_{33})) + 6 \int \frac{\rho_s^2}{\pi^3} (\nabla \pi^\top \omega_s)^2 dx - 4 \int (\omega_s^\top \nabla \rho_s) (\nabla \pi^\top \omega_s) \frac{\rho_s}{\pi^2} dx \\
& - 3 \int (\omega_s^\top \mathbf{H} \pi \omega_s) \frac{\rho_s^2}{\pi^2} dx - \int (\omega_s^\top \nabla \omega_s \nabla \pi) \frac{\rho_s^2}{\pi^2} dx - 2 \int (\omega_s^\top \nabla \omega_s \nabla \rho_s) \frac{\rho_s}{\pi} dx - (A_{32}) - (A_{33}).
\end{aligned}$$

Recall that $(A_{31}) + (A_{32}) + (A_{33}) = (A_3)$, and replacing (A_{32}) and (A_{33}) with their expressions, we have

$$\begin{aligned}
& = 2 \int \frac{\rho_s^2}{\pi} (\nabla \cdot \omega_s)^2 dx + 6 \int \frac{\rho_s^2}{\pi^3} (\nabla \pi^\top \omega_s)^2 dx - 4 \int (\omega_s^\top \nabla \rho_s) (\nabla \pi^\top \omega_s) \frac{\rho_s}{\pi^2} dx \\
& - 3 \int (\omega_s^\top \mathbf{H} \pi \omega_s) \frac{\rho_s^2}{\pi^2} dx - \int (\omega_s^\top \nabla \omega_s \nabla \pi) \frac{\rho_s^2}{\pi^2} dx - 2 \int (\omega_s^\top \nabla \omega_s \nabla \rho_s) \frac{\rho_s}{\pi} dx \\
& - \left(2 \int (\omega_s^\top \nabla \omega_s \nabla \rho_s) \frac{\rho_s}{\pi} dx + 2 \int (\omega_s^\top \mathbf{H} \rho_s \omega_s) \frac{\rho_s}{\pi} dx + 2 \int (\omega_s^\top \nabla \rho_s)^2 \frac{1}{\pi} dx - 2 \int (\omega_s^\top \nabla \rho_s) (\omega_s^\top \nabla \pi) \frac{\rho_s}{\pi^2} dx \right) \\
& - \left(- \int (\omega_s^\top \nabla \omega_s \nabla \pi) \frac{\rho_s^2}{\pi^2} dx - \int (\omega_s^\top \mathbf{H} \pi \omega_s) \frac{\rho_s^2}{\pi^2} dx - 2 \int (\omega_s^\top \nabla \rho_s) (\omega_s^\top \nabla \pi) \frac{\rho_s}{\pi^2} dx + 2 \int (\omega_s^\top \nabla \pi)^2 \frac{\rho_s^2}{\pi^3} dx \right) \\
& = 2 \underbrace{\int (\omega_s^\top \mathbf{H} \rho_s \omega_s) \frac{\rho_s}{\pi} dx - 2 \int (\omega_s^\top \mathbf{H} \pi \omega_s) \frac{\rho_s^2}{\pi^2} dx - 4 \int (\omega_s^\top \nabla \rho_s) (\omega_s^\top \nabla \pi) \frac{\rho_s}{\pi^2} dx + 4 \int \frac{\rho_s^2}{\pi^3} (\nabla \pi^\top \omega_s)^2 dx}_{(M)} \\
& + 2 \underbrace{\int \frac{\rho_s^2}{\pi} (\nabla \cdot \omega_s)^2 dx - 2 \int (\omega_s^\top \nabla \rho_s)^2 \frac{1}{\pi} dx - 4 \int (\omega_s^\top \nabla \omega_s \nabla \rho_s) \frac{\rho_s}{\pi} dx}_{(N_1)} \\
& - 4 \underbrace{\int (\omega_s^\top \mathbf{H} \rho_s \omega_s) \frac{\rho_s}{\pi} dx + 4 \int (\omega_s^\top \nabla \rho_s) (\omega_s^\top \nabla \pi) \frac{\rho_s}{\pi^2} dx}_{(N_2)}.
\end{aligned}$$

Now we analyze (M) and $(N_1) + (N_2)$ separately. Notice that

$$(M) = 2 \int \left(\omega_s^\top \mathbf{H} \left(\frac{\rho_s}{\pi} \right) \omega_s \right) \rho_s dx,$$

and

$$\begin{aligned}
(N_1) + (N_2) & = 2 \int \frac{\rho_s^2}{\pi} (\nabla \cdot \omega_s)^2 dx + 2 \int (\omega_s^\top \nabla \rho_s)^2 \frac{1}{\pi} dx - 4 \int (\omega_s^\top \nabla \omega_s \nabla \rho_s) \frac{\rho_s}{\pi} dx \\
& - 4 \int (\omega_s^\top \mathbf{H} \rho_s \omega_s) \frac{\rho_s}{\pi} dx - 4 \int (\omega_s^\top \nabla \rho_s)^2 \frac{1}{\pi} dx + 4 \int (\omega_s^\top \nabla \rho_s) (\omega_s^\top \nabla \pi) \frac{\rho_s}{\pi^2} dx \\
& = 2 \int \frac{\rho_s^2}{\pi} (\nabla \cdot \omega_s)^2 dx + 2 \int (\omega_s^\top \nabla \rho_s)^2 \frac{1}{\pi} dx - 4 \int (\omega_s^\top \nabla (\omega_s^\top \nabla \rho_s)) \frac{\rho_s}{\pi} dx \\
& - 4 \int (\omega_s^\top \nabla \rho_s) (\omega_s^\top \nabla \frac{\rho_s}{\pi}) dx
\end{aligned}$$

$$\begin{aligned}
&= 2 \int \frac{\rho_s^2}{\pi} (\nabla \cdot \omega_s)^2 dx + 2 \int (\omega_s^\top \nabla \rho_s)^2 \frac{1}{\pi} dx + 4 \int \nabla \cdot \omega_s (\omega_s^\top \nabla \rho_s) \frac{\rho_s}{\pi} dx \\
&= 2 \int \left(\nabla \cdot (\omega_s \rho_s) \frac{1}{\pi} \right)^2 \pi dx.
\end{aligned}$$

Since $\frac{d^2}{ds^2} \chi^2(\rho_s || \pi) = (M) + (N_1) + (N_2)$, Lemma B.9 is proved. \blacksquare

Appendix C: An Illustrative Example for Explicit Forms of \mathcal{I}_t , \mathcal{J}_t , $\|q_t\|_{L^2(\pi)}$

Consider an illustrative example where we simulate the DrMMD gradient flow when the target is a one-dimensional Gaussian target distribution $\pi = \mathcal{N}(0, \bar{\sigma}^2)$ and the initialization is also a one-dimensional Gaussian $\mu_0 = \mathcal{N}(0, \frac{1}{2}\bar{\sigma}^2)$. We take a Gaussian kernel $k(x, y) = \exp(-\frac{1}{2}(x - y)^2)$ whose eigenvalues and eigenfunctions in its Mercer decomposition have the following closed form expressions (Shi et al., 2009, Proposition 1),

$$\varrho_i = \sqrt{\frac{\frac{1}{2\bar{\sigma}^2}}{\frac{1}{2\bar{\sigma}^2} + c + 0.5}} \left(\frac{\beta^2 - 1}{\beta^2 + 1} \right)^i, \quad e_i(x) = \sqrt{\frac{\beta}{i!2^i}} \exp(-cx^2) H_i \left(\sqrt{\frac{1}{2\bar{\sigma}^2}} \beta x \right), \quad (\text{C.1})$$

where $\beta = (1 + 4\bar{\sigma}^2)^{1/4}$, $c = \frac{\beta^2 - 1}{4\bar{\sigma}^2}$ and H_i is the i -th Hermite polynomial function. We pick $\bar{\sigma}^2 > 2$ so $\beta^2 > 3$. The Gaussian kernel is continuous, bounded and c_0 -universal as required in Assumption 1. It also possesses bounded first- and second-order derivatives, thereby satisfying Assumption 2.

Consider the DrMMD gradient flow $(\mu_t)_{t \geq 0}$ defined in (7) along with its particle update scheme $dx_t = -(1 + \lambda) \nabla h_{\mu_t, \pi}(x_t) dt$, where $h_{\mu_t, \pi}$ is the witness function defined in (32) and λ is a positive regularization parameter. Denote m_t, σ_t^2 as the mean and covariance of μ_t , respectively, then we have the following update scheme for m_t, σ_t^2 proved in Lemma C.1:

$$\begin{aligned}
dm_t &= -(1 + \lambda) \mathbb{E}_{x_t \sim \mu_t} [\nabla h_{\mu_t, \pi}(x_t)] dt, \\
d(\sigma_t^2) &= -2(1 + \lambda) \mathbb{E}_{x_t \sim \mu_t} [\nabla h_{\mu_t, \pi}(x_t) \cdot x_t] dt + 2(1 + \lambda) \mathbb{E}_{x_t \sim \mu_t} [\nabla h_{\mu_t, \pi}(x_t)] \cdot m_t dt,
\end{aligned}$$

where the expectations are taken over $x_t \sim \mu_t$. While the resulting distribution is not necessarily Gaussian, we may follow the existing analysis of Stein variational gradient descent (Liu et al., 2024b) and Langevin Monte Carlo dynamics (Lambert et al., 2022) and approximate $x_t \sim \mu_t$ with a Gaussian random variable $y_t \sim \nu_t = \mathcal{N}(m_t, \sigma_t^2)$; this yields the update scheme:

$$\begin{aligned}
dm_t &= (1 + \lambda) \mathbb{E}_{y_t \sim \nu_t} [\nabla h_{\nu_t, \pi}(y_t)] dt, \\
d(\sigma_t^2) &= -2(1 + \lambda) \mathbb{E}_{y_t \sim \nu_t} [\nabla h_{\nu_t, \pi}(y_t) \cdot y_t] dt + 2(1 + \lambda) \mathbb{E}_{y_t \sim \nu_t} [\nabla h_{\nu_t, \pi}(y_t)] \cdot m_t dt,
\end{aligned} \quad (\text{C.2})$$

which gives an evolution of Gaussian distributions ν_t . From (32), the witness function $h_{\nu_t, \pi}$ admits a decomposition with eigenvalues ϱ_i and eigenfunctions e_i : $h_{\nu_t, \pi}(y) = \sum_{i \geq 1} \frac{\varrho_i}{\varrho_i + \lambda} \langle \frac{d\nu_t}{d\pi} - 1, e_i \rangle_{L^2(\pi)} e_i(y)$. Therefore, the velocity field $\frac{d}{dy} h_{\nu_t, \pi}(y)$ can be written as,

$$\frac{d}{dy} h_{\nu_t, \pi}(y) = \sum_{i \geq 1} \frac{\varrho_i}{\varrho_i + \lambda} \left\langle \frac{d\nu_t}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} \frac{d}{dy} e_i(y).$$

Notice that if $m_t = 0$, then for odd i , $\langle \frac{d\nu_t}{d\pi} - 1, e_i \rangle_{L^2(\pi)} = \int e_i d\nu_t - \int e_i d\pi = 0$ because e_i is an odd function. When i is even, $\mathbb{E}_{\nu_t}[\frac{d}{dy} e_i(y)] = 0$ because $y \mapsto \frac{d}{dy} e_i(y)$ is an odd function. As a result, if $m_t = 0$, then $\mathbb{E}[\nabla h_{\nu_t, \pi}(y_t)] = 0$ and hence $\frac{dm_t}{dt} = 0$ from the update scheme in (C.2). Therefore, as long as we initialize the DrMMD gradient flow with $\nu_0 = \mathcal{N}(0, \frac{1}{2}\bar{\sigma}^2)$, a zero mean Gaussian distribution, the entire trajectory will remain a zero mean Gaussian distribution $\mathcal{N}(0, \sigma_t^2)$.

Next, observe that for a zero-mean Gaussian trajectory, if the initial variance satisfies $\sigma_0^2 < \bar{\sigma}^2$, it is natural to expect that the variance increases monotonically toward the target variance $\bar{\sigma}^2$ as the DrMMD flow evolves, i.e. $\sigma_0^2 < \sigma_t^2 \leq \bar{\sigma}^2$ for all t . In Lemma C.2, we provide a rigorous proof of this claim in the cases $\lambda = 0$ and $\lambda = \infty$, by showing that $(1 + \lambda) \mathbb{E}_{y_t \sim \nu_t}[\nabla h_{\nu_t, \pi}(y_t) \cdot y_t] < 0$, which implies that the variance update in (C.2) is monotone increasing. The cases $\lambda = 0$ and $\lambda = \infty$ correspond respectively to the χ^2 flow and the MMD flow regimes. For general $\lambda > 0$, however, we are unable to establish a rigorous proof. Our argument in Lemma C.2 relies heavily on Mehler's formula (Liang and Tran-Bach, 2022, Proposition 2.2), which requires exponential decay of the spectrum $(\varrho_i)_{i \geq 1}$. This condition is not satisfied for DrMMD, whose spectrum is modified to $(\frac{\varrho_i}{\varrho_i + \lambda})_{i \geq 1}$. Nevertheless, we conjecture that the monotonicity property continues to hold for all $\lambda > 0$.

Checking assumptions in Theorem 4.1 and Theorem 5.1: Now we check the assumptions from Theorem 4.1 and Theorem 5.1. The target π is a Gaussian distribution which automatically satisfies a Poincaré inequality, and its potential V is a quadratic function, hence satisfies $\mathbf{H}V \leq \beta \mathbf{I}$. ν_t, π are Gaussians, hence absolutely continuous with respect to Lebesgue on \mathbb{R} . And most importantly, we have $\frac{d\nu_t}{d\pi} - 1 \in \text{Ran}(\mathcal{T}_\pi^{0.25})$, i.e., there exists $q_t \in L^2(\pi)$ such that $\frac{d\nu_t}{d\pi} - 1 = \mathcal{T}_\pi^{0.25} q_t$. To see why, we first need to upper bound $\langle \frac{d\nu_t}{d\pi} - 1, e_i \rangle_{L^2(\pi)}^2$. From Belafhal et al. (2020, Corollary 2), we have the following closed-form expressions for $\langle \frac{d\nu_t}{d\pi} - 1, e_i \rangle_{L^2(\pi)}$:

$$\sqrt{\frac{\beta}{i!2^i}} \left(\sqrt{\frac{1}{1 + 2\sigma_t^2 c}} \left(1 - \frac{\frac{1}{2\bar{\sigma}^2} \beta^2}{c + \frac{1}{2\sigma_t^2}} \right)^{\frac{i}{2}} - \sqrt{\frac{1}{1 + 2\bar{\sigma}^2 c}} \left(1 - \frac{\frac{1}{2\bar{\sigma}^2} \beta^2}{c + \frac{1}{2\bar{\sigma}^2}} \right)^{\frac{i}{2}} \right) \mathbf{H}_i(0),$$

when i is even and 0 otherwise. Therefore, we have

$$\left\langle \frac{d\nu_t}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)}^2 \leq \frac{1}{1 + \bar{\sigma}^2 c} \cdot \frac{\beta}{i!2^i} |\mathbf{H}_i(0)|^2 \cdot 2 \left(\left| \frac{\frac{1}{2\bar{\sigma}^2} \beta^2}{c + \frac{1}{2\sigma_t^2}} - 1 \right|^i + \left| \frac{\frac{1}{2\bar{\sigma}^2} \beta^2}{c + \frac{1}{2\bar{\sigma}^2}} - 1 \right|^i \right).$$

By the monotonicity of the variance, $\sigma_0^2 = \frac{1}{2}\bar{\sigma}^2 \leq \sigma_t^2 < \bar{\sigma}^2$ and $\beta^2 > 3$, we have

$$0 < \frac{2\beta^2}{\beta^2 + 3} - 1 = \frac{\frac{1}{2\bar{\sigma}^2} \beta^2}{c + \frac{1}{\bar{\sigma}^2}} - 1 < \frac{\frac{1}{2\bar{\sigma}^2} \beta^2}{c + \frac{1}{2\sigma_t^2}} - 1 < \frac{\frac{1}{2\bar{\sigma}^2} \beta^2}{c + \frac{1}{2\bar{\sigma}^2}} - 1 = \frac{2\beta^2}{\beta^2 + 1} - 1.$$

So we have the following upper bound

$$\left\langle \frac{d\nu_t}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)}^2 \leq \frac{4}{1 + \bar{\sigma}^2 c} \cdot \frac{\beta}{i!2^i} |\mathbf{H}_i(0)|^2 \cdot \left(\frac{2\beta^2}{\beta^2 + 1} - 1 \right)^i.$$

Now we are ready to study the $L^2(\pi)$ -norm of q_t . Recall the formulas of ϱ_i in (C.1), we have

$$\begin{aligned} \|q_t\|_{L^2(\pi)} &= \sum_{i=1}^{\infty} \frac{\langle \frac{d\nu_t}{d\pi} - 1, e_i \rangle_{L^2(\pi)}^2}{\varrho_i^{0.5}} \leq \frac{4\beta}{1 + \bar{\sigma}^2 c} \sum_{i=1}^{\infty} \frac{1}{i!2^i} |\mathcal{H}_i(0)|^2 \left(\frac{\beta^2 - 1}{\beta^2 + 1} \right)^i \left(\left(\frac{\beta^2 + 1}{\beta^2 - 1} \right)^{0.5} \right)^i \\ &= \frac{4\beta}{1 + \bar{\sigma}^2 c} \sum_{i=1}^{\infty} \frac{1}{i!2^i} |\mathcal{H}_i(0)|^2 \left(\left(\frac{\beta^2 - 1}{\beta^2 + 1} \right)^{0.5} \right)^i = \frac{4\beta}{1 + \bar{\sigma}^2 c} \sqrt{\frac{\beta^2 + 1}{2}}. \end{aligned}$$

The last equality holds by Mehler's formula (Liang and Tran-Bach, 2022, Proposition 2.2). A quick sanity check for the above derivations is to take $r = 0.5$ and see that $\sum_{i=1}^{\infty} \varrho_i^{-1} \langle \frac{d\nu_t}{d\pi} - 1, e_i \rangle_{L^2(\pi)}^2$ is divergent. This indicates that $\frac{d\nu_t}{d\pi} - 1 \notin \mathcal{H}$, which verifies the fact that the Gaussian RKHS does not contain constant functions (Steinwart and Christmann, 2008, Corollary 4.44). Finally, \mathcal{I}_t and \mathcal{J}_t admit the following explicit formulas as well as uniform upper bounds

$$\begin{aligned} \mathcal{J}_t &= \left\| \nabla(\log \pi)^\top \nabla \left(\frac{d\nu_t}{d\pi} \right) \right\|_{L^2(\pi)}^2 = \frac{\frac{1}{\sigma_t^2} - \frac{1}{\bar{\sigma}^2}}{\sqrt{2\pi\bar{\sigma}^3\sigma_t^2}} \cdot \frac{\Gamma(7/2)}{\left(\frac{1}{\sigma_t^2} - \frac{1}{2\bar{\sigma}^2} \right)^{7/2}} \leq \frac{4\Gamma(7/2)}{\sqrt{2\pi} \left(\frac{3}{2} \right)^{7/2}} \\ \mathcal{I}_t &= \left\| \Delta \left(\frac{d\nu_t}{d\pi} \right) \right\|_{L^2(\pi)} = \left(\frac{\bar{\sigma}^2 \left(\frac{1}{\sigma_t^2} - \frac{1}{\bar{\sigma}^2} \right)^2}{\sigma_t^2 \sqrt{2\pi\bar{\sigma}^2}} \right) \cdot \left(\frac{\Gamma(1/2)}{\left(\frac{1}{\sigma_t^2} - \frac{1}{2\bar{\sigma}^2} \right)^{1/2}} + \frac{2\Gamma(3/2) \left(\frac{1}{\sigma_t^2} - \frac{1}{\bar{\sigma}^2} \right)}{\left(\frac{1}{\sigma_t^2} - \frac{1}{2\bar{\sigma}^2} \right)^{3/2}} \right. \\ &\quad \left. + \frac{\Gamma(5/2) \left(\frac{1}{\sigma_t^2} - \frac{1}{\bar{\sigma}^2} \right)^2}{\left(\frac{1}{\sigma_t^2} - \frac{1}{2\bar{\sigma}^2} \right)^{5/2}} \right) \leq \frac{2}{\sqrt{2\pi}\bar{\sigma}^4} \left(\frac{\Gamma(1/2)}{\left(\frac{3}{2} \right)^{5/2}} + \frac{\Gamma(3/2)}{\left(\frac{3}{2} \right)^{3/2}} + \frac{\Gamma(5/2)}{\left(\frac{3}{2} \right)^{5/2}} \right). \end{aligned}$$

Empirical verification: While the above derivations demonstrate that a (Gaussian projected) DrMMD gradient flow satisfies all the required assumptions, it is instructive to demonstrate that a particle implementation in the discrete time setting, without explicit Gaussian projection, shows the behavior consistent with the theory. We therefore simulate our finite particle DrMMD gradient descent $\hat{\mu}_n$ with a step size $\gamma = 0.01$ and particle number $N = 10,000$ and empirically inspect its convergence properties. We take $\bar{\sigma}^2 = 6$ so $\beta^2 = 5$. We estimate the density of the DrMMD descent μ_n from the particles with a kernel density estimator using a Gaussian kernel with lengthscale 0.1 (Epanechnikov, 1969). Based on the estimated densities μ_n , we compute the following two quantities: $\mathcal{I}_n = \|\nabla V^\top \nabla(\frac{d\mu_n}{d\pi})\|_{L^2(\pi)}$ and $\mathcal{J}_n = \|\Delta(\frac{d\mu_n}{d\pi})\|_{L^2(\pi)}$. Their evolution along the DrMMD gradient descent is shown in Figure 3. We observe that both quantities decrease over time as desired, which is a consequence of increasing smoothness of the density ratio $\frac{d\mu_n}{d\pi}$ as μ_n converges to π . We also report the evolution of the KL divergence $\widehat{\text{KL}}(\hat{\mu}_n \|\hat{\pi})$ along the flow, estimated from particles.

Lemma C.1 *Given the DrMMD gradient flow update scheme $dx_t = -(1 + \lambda)\nabla h_{\mu_t, \pi}(x_t)dt$, its mean $m_t = \mathbb{E}[x_t]$ and variance $\sigma_t^2 = \mathbb{E}[x_t^2] - \mathbb{E}[x_t]^2$ update scheme can be expressed as*

$$\begin{aligned} dm_t &= -(1 + \lambda) \mathbb{E}_{x_t \sim \mu_t} [\nabla h_{\mu_t, \pi}(x_t)] dt, \\ d(\sigma_t^2) &= -2(1 + \lambda) \mathbb{E}_{x_t \sim \mu_t} [\nabla h_{\mu_t, \pi}(x_t) \cdot x_t] dt + 2(1 + \lambda) \mathbb{E}_{x_t \sim \mu_t} [\nabla h_{\mu_t, \pi}(x_t)] \cdot m_t dt, \end{aligned}$$




Figure 3: Evolution of $\widehat{\text{KL}}(\hat{\mu}_n \parallel \hat{\pi})$, $\hat{\mathcal{L}}_n$ and $\hat{\mathcal{J}}_n$ along DrMMD particle descent, where all three terms are estimated with samples.

Proof For the mean update, $\frac{d}{dt}m_t = \mathbb{E}[\frac{d}{dt}x_t] = -(1+\lambda) \mathbb{E}_{x_t \sim \mu_t}[\nabla h_{\mu_t, \pi}(x_t)]$. For the variance update,

$$\begin{aligned} \frac{d}{dt}\sigma_t^2 &= \frac{d}{dt}(\mathbb{E}[x_t^2] - \mathbb{E}[x_t]^2) = 2\mathbb{E}[x_t \cdot \frac{d}{dt}x_t] - 2\mathbb{E}[x_t] \cdot \frac{d}{dt}\mathbb{E}[x_t] \\ &= -2(1+\lambda) \mathbb{E}_{x_t \sim \mu_t}[\nabla h_{\mu_t, \pi}(x_t) \cdot x_t] + 2(1+\lambda) \mathbb{E}_{x_t \sim \mu_t}[\nabla h_{\mu_t, \pi}(x_t)] \cdot m_t. \end{aligned}$$

Hence the result. ■

Lemma C.2 Let ν_t be a zero mean normal distribution $\mathcal{N}(0, \sigma_t^2)$ and $0 < \sigma_t^2 < \bar{\sigma}^2$. For the eigenvalues $(\varrho_i)_{i \geq 1}$ and eigenfunctions $(e_i)_{i \geq 1}$ defined in (C.1), we have

$$\begin{aligned} \sum_{i \geq 1} \left\langle \frac{d\nu_t}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} \mathbb{E}_{x \sim \nu_t} \left[\frac{d}{dx} e_i(x) x \right] &< 0, \\ \sum_{i \geq 1} \varrho_i \left\langle \frac{d\nu_t}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} \mathbb{E}_{x \sim \nu_t} \left[\frac{d}{dx} e_i(x) x \right] &< 0. \end{aligned}$$

which correspond to the variance update in (C.2) when $\lambda = 0$ and $\lambda = \infty$, respectively.

Proof For the inner product term, we know from Corollary 2 of Belafhal et al. (2020) that

$$\begin{aligned} &\left\langle \frac{d\nu_t}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} \\ &= H_i(0) \sqrt{\frac{\beta}{i!2^i}} \left(\sqrt{\frac{1}{1+2\sigma_t^2 c}} \left(1 - \frac{\frac{1}{2\bar{\sigma}^2} \beta^2}{c + \frac{1}{2\sigma_t^2}} \right)^{\frac{i}{2}} - \sqrt{\frac{1}{1+2\bar{\sigma}^2 c}} \left(1 - \frac{\frac{1}{2\bar{\sigma}^2} \beta^2}{c + \frac{1}{2\bar{\sigma}^2}} \right)^{\frac{i}{2}} \right). \end{aligned}$$

For the expectation term, notice that

$$\mathbb{E}_{x \sim \nu_t} \left[\frac{d}{dx} e_i(x) x \right] = \sqrt{\frac{\beta}{i!2^i}} \frac{1}{\sqrt{2\pi}\sigma_t} \int x \exp\left(-\frac{1}{2\sigma_t^2} x^2\right) \frac{d}{dx} e_i(x) dx$$

$$\begin{aligned}
&\stackrel{(*)}{=} -\sqrt{\frac{\beta}{i!2^i}} \frac{1}{\sqrt{2\pi}\sigma_t} \int \exp\left(-\frac{1}{2\sigma_t^2}x^2\right) \left(1 - \frac{x^2}{\sigma_t^2}\right) e_i(x) \, dx \\
&= -\sqrt{\frac{\beta}{i!2^i}} \frac{1}{\sqrt{2\pi}\sigma_t} \int \exp\left(-\frac{1}{2\sigma_t^2}x^2\right) \exp(-cx^2) \text{H}_i\left(\sqrt{\frac{1}{2\sigma^2}}\beta x\right) \, dx \\
&\quad + \sqrt{\frac{\beta}{i!2^i}} \frac{1}{\sqrt{2\pi}\sigma_t^3} \int \exp\left(-\frac{1}{2\sigma_t^2}x^2\right) \exp(-cx^2) \text{H}_i\left(\sqrt{\frac{1}{2\sigma^2}}\beta x\right) x^2 \, dx \\
&=: E_{i,1} + E_{i,2}.
\end{aligned}$$

Here in (*), we use integration by parts in which the boundary term vanishes, because $\lim_{x \rightarrow \infty} \exp(-x^2) \text{H}_i(x) = 0$. Next, notice that the second term $E_{i,2}$ equals precisely the derivative of $E_{i,1}$ with respect to c , rescaled by $\frac{1}{\sigma_t^2}$. To distinguish it from the other c that will show up later in $\langle \frac{d\nu_t}{d\pi} - 1, e_i \rangle_{L^2(\pi)}$, we denote it as \mathfrak{c} . The original $\sum_{i \geq 1} \langle \frac{d\nu_t}{d\pi} - 1, e_i \rangle_{L^2(\pi)} \mathbb{E}_{x \sim \nu_t}[\frac{d}{dx} e_i(x)x]$ can be written as the sum of two components $F_1 + F_2$:

$$\begin{aligned}
F_1 &= \sum_{i \geq 1} \text{H}_i(0) \sqrt{\frac{\beta}{i!2^i}} \left(\sqrt{\frac{1}{1+2\sigma_t^2 c}} \left(1 - \frac{\frac{1}{2\sigma^2}\beta^2}{c + \frac{1}{2\sigma_t^2}}\right)^{\frac{i}{2}} - \sqrt{\frac{1}{1+2\sigma^2 c}} \left(1 - \frac{\frac{1}{2\sigma^2}\beta^2}{c + \frac{1}{2\sigma^2}}\right)^{\frac{i}{2}} \right) E_{i,1}, \\
F_2 &= \sum_{i \geq 1} \text{H}_i(0) \sqrt{\frac{\beta}{i!2^i}} \left(\sqrt{\frac{1}{1+2\sigma_t^2 c}} \left(1 - \frac{\frac{1}{2\sigma^2}\beta^2}{c + \frac{1}{2\sigma_t^2}}\right)^{\frac{i}{2}} - \sqrt{\frac{1}{1+2\sigma^2 c}} \left(1 - \frac{\frac{1}{2\sigma^2}\beta^2}{c + \frac{1}{2\sigma^2}}\right)^{\frac{i}{2}} \right) E_{i,2}.
\end{aligned}$$

From Corollary 2 of Belafhal et al. (2020), we know

$$E_{i,1} = -\sqrt{\frac{\beta}{i!2^i}} \sqrt{\frac{1}{1+2\sigma_t^2 \mathfrak{c}}} \left(1 - \frac{\frac{1}{2\sigma^2}\beta^2}{\mathfrak{c} + \frac{1}{2\sigma_t^2}}\right)^{\frac{i}{2}} \text{H}_i(0),$$

and $E_{i,2} = \frac{1}{\sigma_t^2} \frac{d}{dc} E_{i,1}$. The relation between $E_{i,1}$ and $E_{i,2}$ would carry over to F_1 and F_2 as well, i.e., $F_2 = \frac{1}{\sigma_t^2} \frac{d}{dc} F_1$. Next, we compute the first half of F_1 .

$$\begin{aligned}
&\sum_{i \geq 1} \text{H}_i(0) \sqrt{\frac{\beta}{i!2^i}} \sqrt{\frac{1}{1+2\sigma_t^2 c}} \left(1 - \frac{\frac{1}{2\sigma^2}\beta^2}{c + \frac{1}{2\sigma_t^2}}\right)^{\frac{i}{2}} E_{i,1} \\
&= -\sqrt{\frac{1}{1+2\sigma_t^2 c}} \sqrt{\frac{1}{1+2\sigma_t^2 \mathfrak{c}}} \sum_{i \geq 1} \frac{\beta}{i!2^i} \left(1 - \frac{\frac{1}{2\sigma^2}\beta^2}{c + \frac{1}{2\sigma_t^2}}\right)^{\frac{i}{2}} \left(1 - \frac{\frac{1}{2\sigma^2}\beta^2}{\mathfrak{c} + \frac{1}{2\sigma_t^2}}\right)^{\frac{i}{2}} |\text{H}_i(0)|^2 \\
&= -\beta \sqrt{\frac{1}{1+2\sigma_t^2 c}} \sqrt{\frac{1}{1+2\sigma_t^2 \mathfrak{c}}} \left(1 - \left(1 - \frac{\frac{1}{2\sigma^2}\beta^2}{c + \frac{1}{2\sigma_t^2}}\right) \cdot \left(1 - \frac{\frac{1}{2\sigma^2}\beta^2}{\mathfrak{c} + \frac{1}{2\sigma_t^2}}\right)\right)^{-\frac{1}{2}} \\
&= -\beta \left((1+2\sigma_t^2 c)(1+2\sigma_t^2 \mathfrak{c}) - (1+2\sigma_t^2 c - \frac{\sigma_t^2}{\sigma^2}\beta^2)(1+2\sigma_t^2 \mathfrak{c} - \frac{\sigma_t^2}{\sigma^2}\beta^2) \right)^{-\frac{1}{2}} \\
&= -\left(\frac{\sigma_t^2}{\sigma^2} (2 + 2\sigma_t^2 c + 2\sigma_t^2 \mathfrak{c}) - \left(\frac{\sigma_t^2}{\sigma^2} \beta\right)^2 \right)^{-\frac{1}{2}}.
\end{aligned}$$

The second last equality holds by the Mehler's formula (Liang and Tran-Bach, 2022, Proposition 2.2). Similarly, we can also compute the second half of F_1 .

$$\begin{aligned} & \sum_{i \geq 1} H_i(0) \sqrt{\frac{\beta}{i!2^i}} \sqrt{\frac{1}{1+2\bar{\sigma}^2 c}} \left(1 - \frac{\frac{1}{2\bar{\sigma}^2} \beta^2}{c + \frac{1}{2\bar{\sigma}^2}}\right)^{\frac{i}{2}} E_{i,1} \\ &= -\beta \left((1+2\bar{\sigma}^2 c)(1+2\sigma_t^2 \mathfrak{c}) - (1+2\bar{\sigma}^2 c - \beta^2)(1+2\sigma_t^2 \mathfrak{c} - \frac{\sigma_t^2}{\bar{\sigma}^2} \beta^2) \right)^{-\frac{1}{2}} \\ &= - \left(1+2\sigma_t^2 \mathfrak{c} + (1+2\bar{\sigma}^2 c) \frac{\sigma_t^2}{\bar{\sigma}^2} - \frac{\sigma_t^2}{\bar{\sigma}^2} \beta^2 \right)^{-\frac{1}{2}}. \end{aligned}$$

Combining the above two equations, we obtain the following formula of F_1 .

$$F_1 = - \left(\frac{\sigma_t^2}{\bar{\sigma}^2} (2+2\sigma_t^2 c + 2\sigma_t^2 \mathfrak{c}) - \left(\frac{\sigma_t^2}{\bar{\sigma}^2} \beta \right)^2 \right)^{-\frac{1}{2}} + \left(1+2\sigma_t^2 \mathfrak{c} + (1+2\bar{\sigma}^2 c) \frac{\sigma_t^2}{\bar{\sigma}^2} - \frac{\sigma_t^2}{\bar{\sigma}^2} \beta^2 \right)^{-\frac{1}{2}}.$$

Since $F_2 = \frac{1}{\sigma_t^2} \frac{d}{d\mathfrak{c}} F_1$, we have

$$F_2 = \left(\frac{\sigma_t^2}{\bar{\sigma}^2} (2+2\sigma_t^2 c + 2\sigma_t^2 \mathfrak{c}) - \left(\frac{\sigma_t^2}{\bar{\sigma}^2} \beta \right)^2 \right)^{-\frac{3}{2}} \frac{\sigma_t^2}{\bar{\sigma}^2} - \left(1+2\sigma_t^2 \mathfrak{c} + (1+2\bar{\sigma}^2 c) \frac{\sigma_t^2}{\bar{\sigma}^2} - \frac{\sigma_t^2}{\bar{\sigma}^2} \beta^2 \right)^{-\frac{3}{2}}.$$

Recall that the original $\sum_{i \geq 1} \langle \frac{d\nu_t}{d\pi} - 1, e_i \rangle_{L^2(\pi)} \mathbb{E}_{x \sim \nu_t} [\frac{d}{dx} e_i(x)x]$ can be written as the sum of two components $F_1 + F_2$, and recall that $\mathfrak{c} = c$ and $\beta^2 = 1 + 4\bar{\sigma}^2 c$ by definition. We obtain

$$\sum_{i \geq 1} \left\langle \frac{d\nu_t}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} \mathbb{E}_{x \sim \nu_t} \left[\frac{d}{dx} e_i(x)x \right] = F_1 + F_2 = \left(\frac{\sigma_t^2}{\bar{\sigma}^2} \left(2 - \frac{\sigma_t^2}{\bar{\sigma}^2} \right) \right)^{-\frac{1}{2}} \cdot \left(-1 + \left(2 - \frac{\sigma_t^2}{\bar{\sigma}^2} \right)^{-1} \right).$$

which is negative when $0 < \sigma_t < \bar{\sigma}$. So we have concluded the proof of the first claim. Now, we are about to prove the second claim. Following the same derivations as above, we can write $\sum_{i \geq 1} \langle \varrho_i \frac{d\nu_t}{d\pi} - 1, e_i \rangle_{L^2(\pi)} \mathbb{E}_{x \sim \nu_t} [\frac{d}{dx} e_i(x)x]$ as the sum of two terms G_1 and G_2 .

$$\begin{aligned} G_1 &= \sum_{i \geq 1} \varrho_i H_i(0) \sqrt{\frac{\beta}{i!2^i}} \left(\sqrt{\frac{1}{1+2\sigma_t^2 c}} \left(1 - \frac{\frac{1}{2\sigma_t^2} \beta^2}{c + \frac{1}{2\sigma_t^2}} \right)^{\frac{i}{2}} - \sqrt{\frac{1}{1+2\bar{\sigma}^2 c}} \left(1 - \frac{\frac{1}{2\bar{\sigma}^2} \beta^2}{c + \frac{1}{2\bar{\sigma}^2}} \right)^{\frac{i}{2}} \right) E_{i,1}, \\ G_2 &= \sum_{i \geq 1} \varrho_i H_i(0) \sqrt{\frac{\beta}{i!2^i}} \left(\sqrt{\frac{1}{1+2\sigma_t^2 c}} \left(1 - \frac{\frac{1}{2\sigma_t^2} \beta^2}{c + \frac{1}{2\sigma_t^2}} \right)^{\frac{i}{2}} - \sqrt{\frac{1}{1+2\bar{\sigma}^2 c}} \left(1 - \frac{\frac{1}{2\bar{\sigma}^2} \beta^2}{c + \frac{1}{2\bar{\sigma}^2}} \right)^{\frac{i}{2}} \right) E_{i,2}. \end{aligned}$$

Since $\varrho_i \propto (\frac{\beta^2-1}{\beta^2+1})^i$ defined in (C.1) has exponential decay, hence Mehler's formula still hold. Up to some positive multiplier coefficient that do not change the sign, G_1 can be written as the following formula

$$G_1 = - \left((2\bar{\sigma}^2 c)^{-2} (1+2\sigma_t^2 c)^2 + 2 \frac{\sigma_t^2}{\bar{\sigma}^2} - \frac{\sigma_t^4}{\bar{\sigma}^4} \right)^{-\frac{1}{2}} + ((2\bar{\sigma}^2 c)^{-2} (1+2\bar{\sigma}^2 c)(1+2\sigma_t^2 \mathfrak{c}) + 1)^{-\frac{1}{2}}.$$

And similar to F_1, F_2 , we have $G_2 = \frac{1}{\sigma_t^2} \frac{d}{d\mathfrak{c}} G_1$. As a result,

$$G_2 = \left((2\bar{\sigma}^2 c)^{-2} (1 + 2\sigma_t^2 c)^2 + 2\frac{\sigma_t^2}{\bar{\sigma}^2} - \frac{\sigma_t^4}{\bar{\sigma}^4} \right)^{-\frac{3}{2}} \cdot \left(\frac{1 + 2\sigma_t^2 c}{(2\bar{\sigma}^2 c)^2} + \frac{\sigma_t^2}{\bar{\sigma}^2} \right) \\ - \left((2\bar{\sigma}^2 c)^{-2} (1 + 2\bar{\sigma}^2 c)(1 + 2\sigma_t^2 \mathfrak{c}) + 1 \right)^{-\frac{3}{2}} \cdot \left(\frac{1 + 2\bar{\sigma}^2 c}{(2\bar{\sigma}^2 c)^2} + 1 \right).$$

Combined, we obtain

$$G_1 + G_2 = - \left((2\bar{\sigma}^2 c)^{-2} (1 + 2\sigma_t^2 c)^2 + 2\frac{\sigma_t^2}{\bar{\sigma}^2} - \frac{\sigma_t^4}{\bar{\sigma}^4} \right)^{-\frac{3}{2}} \cdot \left((1 + 2\sigma_t^2 c) \frac{2\sigma_t^2 c}{(2\bar{\sigma}^2 c)^2} + \frac{\sigma_t^2}{\bar{\sigma}^2} - \frac{\sigma_t^4}{\bar{\sigma}^4} \right) \\ + \left((2\bar{\sigma}^2 c)^{-2} (1 + 2\bar{\sigma}^2 c)(1 + 2\sigma_t^2 c) + 1 \right)^{-\frac{3}{2}} \cdot (1 + 2\bar{\sigma}^2 c) \frac{2\sigma_t^2 c}{(2\bar{\sigma}^2 c)^2}.$$

When $0 < \sigma_t^2 < \bar{\sigma}^2$, we have the following relations

$$(2\bar{\sigma}^2 c)^{-2} (1 + 2\sigma_t^2 c)^2 + 2\frac{\sigma_t^2}{\bar{\sigma}^2} - \frac{\sigma_t^4}{\bar{\sigma}^4} < (2\bar{\sigma}^2 c)^{-2} (1 + 2\bar{\sigma}^2 c)(1 + 2\sigma_t^2 c) + 1, \\ (1 + 2\sigma_t^2 c) \frac{2\sigma_t^2 c}{(2\bar{\sigma}^2 c)^2} + \frac{\sigma_t^2}{\bar{\sigma}^2} - \frac{\sigma_t^4}{\bar{\sigma}^4} > (1 + 2\bar{\sigma}^2 c) \frac{2\sigma_t^2 c}{(2\bar{\sigma}^2 c)^2}.$$

Therefore,

$$\sum_{i \geq 1} \varrho_i \left\langle \frac{d\nu_t}{d\pi} - 1, e_i \right\rangle_{L^2(\pi)} \mathbb{E}_{x \sim \nu_t} \left[\frac{d}{dx} e_i(x) x \right] = G_1 + G_2 < 0,$$

proving the second claim. ■