
On the Convergence of Gradient Descent in GANs: MMD GAN As a Gradient Flow

Youssef Mroueh*
IBM Research

Truyen Nguyen*
University of Akron

Abstract

We consider the maximum mean discrepancy (MMD) GAN problem and propose a parametric kernelized gradient flow that mimics the min-max game in gradient regularized MMD GAN. We show that this flow provides a descent direction minimizing the MMD on a statistical manifold of probability distributions. We then derive an explicit condition which ensures that gradient descent on the parameter space of the generator in gradient regularized MMD GAN is globally convergent to the target distribution. Under this condition, we give non asymptotic convergence results of gradient descent in MMD GAN. Another contribution of this paper is the introduction of a dynamic formulation of a regularization of MMD and demonstrating that the parametric kernelized descent for MMD is the gradient flow of this functional with respect to the new Riemannian structure. Our obtained theoretical result allows ones to treat gradient flows for quite general functionals and thus has potential applications to other types of variational inferences on a statistical manifold beyond GANs. Finally, numerical experiments suggest that our parametric kernelized gradient flow stabilizes GAN training and guarantees convergence.

1 Introduction

Generative Adversarial Networks (GANs) were introduced in [1] and have attracted growing attention in the machine learning community. Implicit Generative models such as GANs can be seen as learning a

distribution via optimizing a functional defined on a statistical manifold. The statistical manifold refers to the parametrization of the *Generator*. There are a plethora of works on functionals that are optimized in GANs for example the Jensen-Shanon divergence in the original work [1]; general ϕ divergences in [2]; the neural net distance in [3]; integral probability metrics such as the Wasserstein order 1 distance considered in [4]; the maximum mean discrepancy [5] considered in [6–10]. Despite their striking empirical success, the rigorous understanding of the convergence in *distributional sense* of gradient descent in GANs remains less understood. Much of the theoretical analysis has been dedicated to the stability of the min-max game via the introduction of gradient regularizers [11–17]. Min-max convergence rates for a large class of GANs was studied in [18, 19], however these bounds are not specific to gradient descent in GANs. In this work we aim at understanding the distributional convergence properties of gradient descent in the context of MMD GANs. The work closest to ours is [20] that establishes global convergence of the generator using the Wasserstein 1 distance. Nevertheless, gradient descent is not explicitly considered in [20].

We summarize our main contributions in this work as follows:

- We introduce in Section 3 a new gradient regularizer for the MMD. This regularizer has the form of a parametric energy, where the gradient is taken with respect to the generator parameters, instead of the input space as usually considered in previous works. We call the new proposed discrepancy $\text{MMD}_{\alpha,\beta}$.
- We consider the MMD GAN problem in Section 4 and propose a new descent direction in terms of the witness function of the parametric regularized $\text{MMD}_{\alpha,\beta}$. We give in this section detailed descriptions and properties of the corresponding continuous flow.
- We analyze in Section 5 the non-asymptotic distributional convergence properties of gradient descent in MMD GAN when using the $\text{MMD}_{\alpha,\beta}$ witness functions to drive the generator updates.

Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume 130. Copyright 2021 by the author(s).

• We derive in Section 6 a dynamic formulation of the MMD on the statistical manifold of probability distributions and use it to propose a novel regularization of the MMD, which we call $d_{\alpha,\beta}$. We show that $d_{\alpha,\beta}$ admits a Riemannian metric tensor and investigate gradient flows for general functionals w.r.t. this structure. Intriguingly, we show that gradient descent in MMD GANs driven by the witness function of $\text{MMD}_{\alpha,\beta}$ coincides with the gradient flow of the MMD w.r.t. this new geometric structure.

• Finally, we discuss related works in Section 7, and validate experimentally our theoretical findings in Section 8.

2 Preliminaries

Let Ω be an open region in \mathbb{R}^d and let \mathcal{H} be a reproducing kernel Hilbert space (RKHS) generated by a kernel $k(x, y)$ on $\Omega \times \Omega$. Let $\mathcal{Z} \subset \mathbb{R}^m$ be an open region on a lower dimensional space endowed with a probability distribution ν on \mathcal{Z} . Let Θ be a parameter space in \mathbb{R}^p , and $(\theta, z) \rightarrow G_\theta(z) = (G_\theta^1(z), \dots, G_\theta^d(z)) \in \Omega$ be a generator function defined on Θ . We assume that k is bounded, G_θ is differentiable in θ , and

$$\begin{aligned} \|k(x, \cdot) - k(y, \cdot)\|_{\mathcal{H}} &\leq L \|x - y\|, \\ \|G_\theta(z) - G_{\theta'}(z)\| &\leq M(z) \|\theta - \theta'\| \end{aligned} \quad (1)$$

for some constant $L > 0$ and function $M : \mathcal{Z} \rightarrow [0, \infty)$ with $\mathbb{E}_\nu[M(z)^2] < \infty$. In (1) and throughout the paper, $\langle \cdot, \cdot \rangle$ and $\|\cdot\|$ denote the standard Euclidean inner product and norm, while $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ and $\|\cdot\|_{\mathcal{H}}$ denote the inner product and norm on \mathcal{H} . We use ∇ and ∇_w to respectively denote the standard gradient and the standard gradient with respect to the variable w . For two vectors \mathbf{a} and \mathbf{b} , by $\mathbf{a} \otimes \mathbf{b}$ we mean the matrix whose (i, j) entry is $\mathbf{a}_i \mathbf{b}_j$. For a probability distribution ρ on Ω , let $\mu_\rho(x) := \int k(x, y) \rho(dy)$ denote its kernel mean embedding. As k is bounded, we have $\mu_\rho \in \mathcal{H}$ and $\int f(x) \rho(dx) = \langle f, \mu_\rho \rangle_{\mathcal{H}}$ for every $f \in \mathcal{H}$ (see [21]). Note that $\mu_{\rho_1 - \rho_2} = \mu_{\rho_1} - \mu_{\rho_2}$ by linearity.

We consider the following statistical manifold of probability distributions:

$$\mathcal{P}_\Theta := \{q_\theta = (G_\theta)_\# \nu, \theta \in \Theta\}.$$

The main objective functions considered in this paper are parametric energy regularizations of $\text{MMD}(p, q_\theta) := \|\mu_p - \mu_{q_\theta}\|_{\mathcal{H}}$. The introduction of the following operators plays an important role in understanding these objective functions and associated gradient flows.

Definition 2.1 (Matrix Mass). *Let $J_\theta G_\theta(z) = \left(\frac{\partial G_\theta^i(z)}{\partial \theta_j}\right)_{ij} \in \mathbb{R}^{p \times d}$ denote the Jacobian of G_θ with respect to θ . Then for $\theta \in \Theta$, we define a matrix valued*

kernel Γ_θ on $\mathcal{Z} \times \mathcal{Z}$ as follows, for $(z, z') \in \mathcal{Z} \times \mathcal{Z}$:

$$\Gamma_\theta(z, z') := J_\theta G_\theta(z)^\top J_\theta G_\theta(z') \in \mathbb{R}^{d \times d}.$$

Definition 2.2 (Parametric Grammian - Mass Corrected Grammian of Derivatives). *For $\theta \in \Theta$, let $L_\theta : \mathcal{H} \rightarrow \mathbb{R}^p$ be the operator given by*

$$L_\theta(f) := \int J_\theta G_\theta(z) \nabla f(G_\theta(z)) \nu(dz)$$

and $L_\theta^\top : \mathbb{R}^p \rightarrow \mathcal{H}$ be the operator given by

$$\begin{aligned} L_\theta^\top(v) &:= \int \langle \nabla_\theta [k(G_\theta(z), \cdot), v] \rangle \nu(dz) \\ &= \int \langle \nabla_x k(G_\theta(z), \cdot), J_\theta G_\theta(z)^\top v \rangle \nu(dz). \end{aligned}$$

Then the parametric Grammian $D_\theta : \mathcal{H} \rightarrow \mathcal{H}$ is defined by $D_\theta := L_\theta^\top L_\theta$.

The main properties of these operators are summarized as follows. Proofs are given in the Appendix.

Proposition 2.3. *For each $\theta \in \Theta$, we have*

- i) L_θ^\top is the adjoint operator of L_θ , i.e., $\langle L_\theta f, v \rangle = \langle f, L_\theta^\top v \rangle_{\mathcal{H}}$ for $f \in \mathcal{H}$ and $v \in \mathbb{R}^p$.
- ii) D_θ is symmetric, i.e., $\langle D_\theta f, g \rangle_{\mathcal{H}} = \langle f, D_\theta g \rangle_{\mathcal{H}}$ for $f, g \in \mathcal{H}$.
- iii) $\langle f, D_\theta f \rangle_{\mathcal{H}} = \|\nabla_\theta \int f(x) q_\theta(dx)\|^2 = \|\nabla_\theta [\langle f, \mu_{q_\theta} \rangle_{\mathcal{H}}]\|^2 \geq 0$. In particular, D_θ is a positive operator and hence its spectrum is contained in $[0, \infty)$.
- iv) For $f \in \mathcal{H}$, we have $(D_\theta f)(x) = \langle D(x, \cdot), f \rangle_{\mathcal{H}}$ with $D(y, y') := \iint \langle \partial_\theta k(G_\theta(z), y), \partial_\theta k(G_\theta(z'), y') \rangle \nu(dz) \nu(dz')$, where: $\partial_\theta k(G_\theta(z), y) = J_\theta G_\theta(z) \nabla_x k(G_\theta(z), y)$. Equivalently $D(y, y') = \iint \text{Trace}(\nabla_x k(G_\theta(z), y) \otimes \Gamma_\theta(z, z') \nabla_x k(G_\theta(z'), y')) \nu(dz) \nu(dz')$.

3 A Novel Parametric Energy Regularization of MMD

Let us introduce a parametric energy regularization of MMD and this notion of discrepancy will play a central role in this paper. For parameters $\alpha, \beta \geq 0$, the regularized discrepancy between a given probability distribution p on Ω and a parametric distribution $q_\theta \in \mathcal{P}_\Theta$ is defined by

$$\begin{aligned} \text{MMD}_{\alpha,\beta}(p, q_\theta) &:= \\ \sup_{f \in E_{\alpha,\beta}} &\left\{ \int_\Omega f(x) p(dx) - \int_\Omega f(x) q_\theta(dx) \right\} \end{aligned} \quad (2)$$

with $E_{\alpha,\beta} := \left\{ f \in \mathcal{H} : \alpha \|\nabla_\theta \int f(x) q_\theta(dx)\|_{\mathbb{R}^p}^2 + \beta \|f\|_{\mathcal{H}}^2 \leq \frac{1}{2} \right\}$. The case $\alpha = 0$ and $\beta = 1/2$ corresponds to the MMD [5], while the case $\alpha = 1/2$ and $\beta = 0$ shares some similarity with the usual kernelized Sobolev discrepancy [22]. The main difference in this definition with the Sobolev discrepancy is that the parametric energy $\|\nabla_\theta \langle f, \mu_{q_\theta} \rangle_{\mathcal{H}}\|_{\mathbb{R}^p}^2 = \langle f, D_\theta f \rangle_{\mathcal{H}}$ is used in place of the standard energy $\|\nabla f\|_{L^2_{q_\theta}}^2$.

Remark 3.1. *Note that our regularization for MMD while it shares similarities with WGAN-GP [11], it is different since our gradient penalty is with respect to the generator parameter whereas it is with respect to the input in WGAN-GP.*

Hereafter, $I : \mathcal{H} \rightarrow \mathcal{H}$ denotes the identity operator. Then it follows from property iii) in Proposition 2.3 that $\alpha D_\theta + \beta I$ is invertible whenever $\alpha \geq 0$ and $\beta > 0$. Proposition 2.3 also allows us to express the constraint $E_{\alpha,\beta}$ in (2) as

$$E_{\alpha,\beta} = \left\{ f \in \mathcal{H} : 2 \langle f, (\alpha D_\theta + \beta I) f \rangle_{\mathcal{H}} \leq 1 \right\}. \quad (3)$$

This constraint can be interpreted as a regularization through the following unconstrained formulation.

Proposition 3.2. *Let $\Delta_\theta(f) := \int f(x) p(dx) - \int f(x) q_\theta(dx)$. For $\alpha \geq 0$ and $\beta > 0$, we have*

$$\begin{aligned} \text{MMD}_{\alpha,\beta}(p, q_\theta)^2 &= \sup_{f \in \mathcal{H}} \left\{ \Delta_\theta(f) - \frac{\alpha}{2} \|\nabla_\theta \langle f, \mu_{q_\theta} \rangle_{\mathcal{H}}\|_{\mathbb{R}^p}^2 - \frac{\beta}{2} \|f\|_{\mathcal{H}}^2 \right\} \\ &= \frac{1}{2} \langle \mu_{p-q_\theta}, (\alpha D_\theta + \beta I)^{-1} \mu_{p-q_\theta} \rangle_{\mathcal{H}}. \end{aligned}$$

Moreover, the witness function f^* realizing the above supremum is given by:

$$(\alpha D_\theta + \beta I) f^* = \mu_{p-q_\theta}. \quad (4)$$

The next result shows that the regularized $\text{MMD}_{\alpha,\beta}$ is upper bounded by the MMD, and gives a characterization on distributions for which the two discrepancies are the same.

Corollary 3.3. *For $\alpha \geq 0$ and $\beta > 0$, we have*

$$\sqrt{2\beta} \text{MMD}_{\alpha,\beta}(p, q_\theta) \leq \text{MMD}(p, q_\theta).$$

In case $\alpha > 0$, the equality happens if and only if $D_\theta \mu_{p-q_\theta} = 0$.

4 Generative Adversarial Networks via Parametric Regularized Flows

Let p be the target distribution which is a probability measure on Ω . Consider the functional $\mathcal{F}(q_\theta) =$

$\frac{1}{2} \text{MMD}^2(p, q_\theta) = \frac{1}{2} \|\mu_{p-q_\theta}\|_{\mathcal{H}}^2$. We now focus on the MMD GAN problem:

$$\begin{aligned} \min_{q \in \mathcal{P}_\Theta} \mathcal{F}(q) &= \min_{\theta \in \Theta} \mathcal{F}(q_\theta) \\ &= \min_{\theta \in \Theta} \frac{1}{2} \|\mathbb{E}_p k(x, \cdot) - \mathbb{E}_{z \sim \nu} k(G_\theta(z), \cdot)\|_{\mathcal{H}}^2. \end{aligned}$$

This problem has been investigated in several works [6–8, 22, 23]. In this section we propose a new descent direction in the parameter space of the generator.

Continuous Descent. We consider the following dynamic for any sequence of functions $f_t \in \mathcal{H}$, $t \geq 0$:

$$\frac{d\theta_t}{dt} = L_{\theta_t}(f_t) = \int J_{\theta_t} G_{\theta_t}(z') \nabla_x f_t(G_{\theta_t}(z')) \nu(dz'). \quad (5)$$

For a given $z \in \mathcal{Z}$, the dynamic of the generator is as follows

$$\begin{aligned} \frac{dG_{\theta_t}(z)}{dt} &= J_{\theta_t} G_{\theta_t}(z)^\top \frac{d\theta_t}{dt} \\ &= \int J_{\theta_t} G_{\theta_t}(z)^\top J G_{\theta_t}(z') \nabla_x f_t(G_{\theta_t}(z')) \nu(dz') \\ &= \int \Gamma_{\theta_t}(z, z') \nabla_x f_t(G_{\theta_t}(z')) \nu(dz'). \end{aligned} \quad (7)$$

While the dynamic of particles is usually given by velocities defined at each particle, the generator's dynamic at a given z is the average of the mass corrected velocities of all other samples from the generator. The mass correction is driven by the matrix valued kernel Γ_{θ_t} that defines a similarity in the hidden space \mathcal{Z} . For example in particles descent such as Sobolev Descent [22] the dynamic of particles X_t is given by:

$$\frac{dX_t}{dt} = \nabla_x f_t(X_t),$$

and this simple advection dynamic is to be contrasted with the generator dynamic (7).

Using the dynamic of the generator in (7) and item iv) in Proposition 2.3, we also have the following dynamic of the mean embedding:

$$\begin{aligned} \frac{d}{dt} \mu_{q_{\theta_t}} &= \frac{d}{dt} \int k(G_{\theta_t}(z), \cdot) \nu(dz) \\ &= \int \left\langle \nabla_x k(G_{\theta_t}(z), \cdot), \frac{dG_{\theta_t}(z)}{dt} \right\rangle \nu(dz) \\ &= D_{\theta_t} f_t. \end{aligned} \quad (8)$$

Thanks to (8), it is easy to derive the dynamic of the MMD distance:

$$\begin{aligned} \frac{d\mathcal{F}(q_{\theta_t})}{dt} &= \left\langle \mu_{p-q_{\theta_t}}, -\frac{d}{dt} \mu_{q_{\theta_t}} \right\rangle_{\mathcal{H}} \\ &= -\langle \mu_{p-q_{\theta_t}}, D_{\theta_t} f_t \rangle_{\mathcal{H}}. \end{aligned} \quad (9)$$

Let us consider the following choices for the sequence f_t :

- **Witness functions of MMD.** In that case, we set $f_t = \mu_{p-q_{\theta_t}}$. Using the generator updates given in (5), we have therefore:

$$\frac{d\mathcal{F}(q_{\theta_t})}{dt} = -\langle \mu_{p-q_{\theta_t}}, D_{\theta_t} \mu_{p-q_{\theta_t}} \rangle_{\mathcal{H}} \leq 0.$$

This is a valid descent direction, and is similar in spirit to the MMD flows of [24]. Nevertheless as shown for the particles case in [24], it does not lead to convergence. In the discrete case, [24] introduced a noising scheme that has convergence guarantees.

- **Witness functions of MMD_{1,0}.** In that case, let us assume that solutions f_t of $D_{\theta_t} f_t = \mu_{p-q_{\theta_t}}$ exist. Then by using $d\theta_t = L_{\theta_t}(f_t)dt$ we obtain:

$$\frac{d\mathcal{F}(q_{\theta_t})}{dt} = -\langle \mu_{p-q_{\theta_t}}, \mu_{p-q_{\theta_t}} \rangle = -2\mathcal{F}(q_{\theta_t}).$$

While this seems to be the ideal choice as it gives us an exponential convergence, it comes with the caveat that D_{θ_t} may be singular and hence we have either no solution or infinitely many solutions for f_t . These derivations and the singularity issue of operator D_{θ} motivated the introduction of MMD _{α,β} , and we consider hereafter its flows.

- **Witness functions of MMD _{α,β} .** Let f_t be the unique witness function of MMD _{α,β} between p and $q_{\theta_t} = (G_{\theta_t})_{\#}\nu$ given by

$$(\alpha D_{\theta_t} + \beta I)f_t = \mu_{p-q_{\theta_t}}. \quad (10)$$

Theorem 4.1 below gives the dynamic of the MMD when the generator parameters are updated according to Equation (5) with f_t being the MMD _{α,β} witness functions given in Equation (10).

Theorem 4.1 (Parametric Regularized Flows Decrease the MMD Distance). *Assume that $\alpha, \beta > 0$. Then the dynamic (5)–(10) defined by the witness function of the parametric regularized MMD decreases the functional $\mathcal{F}(q_{\theta})$:*

$$\frac{d\mathcal{F}(q_{\theta_t})}{dt} = -\frac{2}{\alpha} \left[\mathcal{F}(q_{\theta_t}) - \beta \text{MMD}_{\alpha,\beta}(p, q_{\theta_t})^2 \right] \leq 0. \quad (11)$$

Moreover, we have $\frac{d\mathcal{F}(q_{\theta_t})}{dt} < 0$ if and only if $D_{\theta_t} \mu_{p-q_{\theta_t}} \neq 0$.

We see from Theorem 4.1 that MMD _{α,β} witness functions alleviate the singularity issue of D_{θ_t} , but slows down the convergence by introducing a damping term that is proportional to $\text{MMD}_{\alpha,\beta}^2$.

5 Non-Asymptotic Convergence Of Gradient Descent In MMD GAN

In Section 4, we showed that MMD _{α,β} witness functions provide descent directions for continuous MMD GAN. In this section we turn to (discrete) gradient descent in the parameter space of the generator, and give non asymptotic convergence results for gradient descent in regularized MMD GANs.

Discrete Descent Directions. We would like to identify directions of θ along which the functional $\mathcal{F}(q_{\theta})$ decreases its value. For this, let us compute the rate $\frac{d}{d\varepsilon} \Big|_{\varepsilon=0} \mathcal{F}(q_{\theta+\varepsilon v})$ for each vector $v \in \mathbb{R}^p$.

Lemma 5.1. *Let $\theta \in \Theta$. Then for any vector $v \in \mathbb{R}^p$, we have*

$$\frac{d}{d\varepsilon} \mathcal{F}(q_{\theta+\varepsilon v}) = -\langle \mu_{p-q_{\theta+\varepsilon v}}, L_{\theta+\varepsilon v}^{\top} v \rangle_{\mathcal{H}} \text{ for every } \varepsilon \geq 0.$$

Lemma 5.1 implies that v is a descent direction of $\mathcal{F}(q_{\theta})$ if and only if $\langle \mu_{p-q_{\theta}}, L_{\theta}^{\top} v \rangle_{\mathcal{H}} > 0$. The next result gives one such direction.

Proposition 5.2. *Let $\theta \in \Theta$, and assume that $D_{\theta} \mu_{p-q_{\theta}} \neq 0$. Let $v^* = L_{\theta} f$ with $f \in \mathcal{H}$ being the solution of $(\alpha D_{\theta} + \beta I)f = \mu_{p-q_{\theta}}$. Then v^* is a descent direction of $\mathcal{F}(q_{\theta})$. Precisely, we have*

$$\frac{d}{d\varepsilon} \Big|_{\varepsilon=0} \mathcal{F}(q_{\theta+\varepsilon v^*}) = -\left[\alpha \|D_{\theta} f\|_{\mathcal{H}}^2 + \beta \langle f, D_{\theta} f \rangle_{\mathcal{H}} \right] < 0.$$

In particular, $\mathcal{F}(q_{\theta+\varepsilon v^*}) < \mathcal{F}(q_{\theta})$ if $\varepsilon > 0$ is small enough.

Discrete Time Descent for MMD GAN. Hereafter, we denote $\|h\|_{\mathcal{H}} := \left(\sum_{i=1}^d \|h_i\|_{\mathcal{H}}^2 \right)^{\frac{1}{2}}$ for $h = (h_1, \dots, h_d) \in \mathcal{H}^d$ and $\|A\| := \left(\sum_{i,j} a_{ij}^2 \right)^{\frac{1}{2}}$ for a matrix $A = (a_{ij})$. The next result holds under the following extra conditions for k and G :

$$\|\nabla_x k(z_1, \cdot) - \nabla_x k(z_2, \cdot)\|_{\mathcal{H}} \leq \tilde{L} \|z_1 - z_2\|, \quad (12)$$

$$\|J_{\theta} G_{\theta}(z) - J_{\theta'} G_{\theta'}(z)\| \leq \tilde{M}(z) \|\theta - \theta'\| \quad (13)$$

with $\tilde{L} > 0$ being a constant and $\mathbb{E}_{\nu}[\tilde{M}(z)] < \infty$.

In Theorem 5.3 below we find conditions under which we can achieve global convergence in MMD of gradient descent in MMD GAN. When the kernel is characteristic, this is equivalent to the global weak convergence of gradient descent in MMD GAN to the target distribution. Gradient descent for MMD GAN is given by the following updates: for all $\ell \geq 1$, the witness functions between p and $q_{\theta_{\ell}} = (G_{\theta_{\ell}})_{\#}\nu$ update:

$$f_{\ell} := (\alpha_{\ell} D_{\theta_{\ell}} + \beta_{\ell} I)^{-1} \mu_{p-q_{\theta_{\ell}}}, \quad (14)$$

and the generator update:

$$\theta_{\ell+1} := \theta_{\ell} + \varepsilon_{\ell} L_{\theta_{\ell}}(f_{\ell}), \quad (15)$$

where α_ℓ, β_ℓ are sequences of regularization parameters, and ε_ℓ is a sequence of learning rates.

Theorem 5.3. *Assume that k and G satisfy (1) and (12)–(13). Let $\lambda_i(\theta) > 0$ be the smallest non-zero eigenvalue of D_θ , and define $a(\theta, f) := 1 - \frac{\|P_{\text{Null}(D_\theta)} f\|_{\mathcal{H}}^2}{\|f\|_{\mathcal{H}}^2}$ with $P_{\text{Null}(D_\theta)}$ being the projection to the null space of D_θ . Let $\mathcal{F}(q_\theta) = \frac{1}{2} \text{MMD}(p, q_\theta)^2$. Consider the gradient descent updates given in (14) and (15). Let $0 < \tau < 1$ and $\theta_1 \in \Theta$ be the starting point chosen such that:*

$$a(\theta_1, f_1) > \tau. \quad (16)$$

The sequence ε_ℓ is chosen so that the following two conditions are satisfied for each ℓ :

$$a(\theta_{\ell+1}, f_{\ell+1}) > \tau \quad (17)$$

and

$$2C(2\beta_\ell)^{-1}(1 + \sqrt{\mathcal{F}(q_{\theta_1})}) \leq \varepsilon_\ell^{-1} \quad (18)$$

with $C > 0$ depending only on the constants C_1, C_2, C_3, C_4 given in Lemma B.1.

Under the condition (17) on ε_ℓ we have $\tau < a(\theta_\ell, f_\ell) \leq 1$. Let $\chi_j := \frac{\lambda_i(\theta_j)a(\theta_j, f_j)}{\alpha_j \lambda_i(\theta_j)a(\theta_j, f_j) + \beta_j} > 0$. Then we have:

$$\mathcal{F}(q_{\theta_{\ell+1}}) \leq \mathcal{F}(q_{\theta_1}) \exp\left(-\sum_{j=1}^{\ell} \varepsilon_j \chi_j\right) \quad \forall \ell \geq 1.$$

In particular for $\alpha_\ell \leq \frac{\tau}{2}$ and $\beta_\ell = \alpha_\ell \lambda_i(\theta_\ell)$, we obtain $\chi_\ell \geq 1$ and it follows that:

$$\mathcal{F}(q_{\theta_{\ell+1}}) \leq \mathcal{F}(q_{\theta_1}) \exp\left(-\sum_{j=1}^{\ell} \varepsilon_j\right).$$

Consequently if $\sum_{j=1}^{\infty} \varepsilon_j = +\infty$ (meaning that ε_j decays as $1/\sqrt{j}$ for e.g) and conditions (17) and (18) hold, then we obtain $\mathcal{F}(\theta_\ell) \rightarrow 0$ as $\ell \rightarrow \infty$. As we see from Theorem 5.3, not all learning rates are admissible. At a given iteration ℓ , we select a learning rate ε_ℓ so that $f_{\ell+1} \notin \text{Null}(D_{\theta_{\ell+1}})$ (ensured by condition (17)) and so that (18) holds as well. To understand condition (17), recall from Theorem 4.1 that having f_ℓ not in the null space of D_{θ_ℓ} means that we have a strict descent. Notice that $f_\ell \notin \text{Null}(D_{\theta_\ell})$ if and only if $\mu_{p-q_{\theta_\ell}} \notin \text{Null}(D_{\theta_\ell})$.

6 Parametric Kernelized Flows for a General Functional

The flow of the MMD functional (i.e. $\frac{1}{2} \text{MMD}(p, q_\theta)^2$) analyzed in the previous sections is driven by the gradient of the witness function between p and q_θ of the discrepancy $\text{MMD}_{\alpha, \beta}$. In this section we discover a

Riemannian structure on the statistical manifold of probability distributions and show that the continuous gradient descent in MMD GANs described in Section 4 coincides with the gradient flow of the functional with respect to this new geometric structure. We also develop a rigorous theory for treating gradient flows of general functionals and thus open a way for other types of variational inferences beyond GANs.

6.1 Dynamic Formulation For MMD on a Statistical Manifold

The following result gives a dynamic formulation of MMD and allows us to discover a Riemannian structure associated to MMD. This is analogous to Benamou–Brenier dynamic formulation of the Wasserstein of order 2 [25]:

$$W_2^2(p, q) = \inf_{(q_t, f_t)} \int_0^1 \int \|\nabla_x f_t(x)\|^2 q_t(dx) dt$$

$$\text{s.t. } \frac{\partial q_t(x)}{\partial t} = -\text{div}(q_t \nabla_x f_t(x)), \quad q_0 = q, \quad q_1 = p. \quad (19)$$

The main difference is that their flows are with respect to the standard $L_{q_\theta}^2$ energy, while ours as explained in Section 6.2, are driven by the parametric energy $\|\nabla_\theta[\langle f, \mu_{q_\theta} \rangle_{\mathcal{H}}]\|^2$.

Theorem 6.1 (Dynamic MMD on a Statistical Manifold). *Assume that for any $\theta_0, \theta_1 \in \Theta$, there exists a path $(\theta_t, f_t)_{t \in [0,1]}$ such that $\theta_{t=0} = \theta_0$, $\theta_{t=1} = \theta_1$, $f_t \in \mathcal{H}$, and*

$$\partial_t \theta_t = L_{\theta_t} f_t \text{ and } D_{\theta_t} f_t = 2\mu_{q_{\theta_1} - q_{\theta_t}} \quad \forall t \in [0, 1]. \quad (20)$$

Let q_{θ_0} and q_{θ_1} be two probability measures in \mathcal{P}_Θ . Then we have the following dynamic form of MMD between distributions defined on the statistical manifold \mathcal{P}_Θ :

$$\text{MMD}^2(q_{\theta_0}, q_{\theta_1}) = \min_{(\theta_t, f_t)} \left\{ \int_0^1 \|D_{\theta_t} f_t\|_{\mathcal{H}}^2 dt \right\},$$

$$\partial_t \theta_t = L_{\theta_t} f_t, \quad f_t \in \mathcal{H}, \quad \theta_{t=0} = \theta_0, \quad \theta_{t=1} = \theta_1.$$

Theorem 6.1 is proven under Assumption 20 that guarantees the existence of a solution. This assumption is not realistic since D_{θ_t} can be singular. Nevertheless, we state this theorem to motivate the introduction in the next section of the dynamic form of a regularized MMD that alleviates this singularity issue.

6.2 Regularized MMD and Gradient Flows on a Statistical Manifold

Motivated by the result in Theorem 6.1 and to alleviate the singularity issue in Assumption (20), we define the following regularized version of MMD:

Definition 6.2 (Regularized MMD on a Statistical Manifold). Let $\alpha, \beta > 0$. Define

$$d_{\alpha,\beta}(q_{\theta_0}, q_{\theta_1})^2 = \min_{(\theta_t, f_t)} \int_0^1 \left(\alpha \|D_{\theta_t} f_t\|_{\mathcal{H}}^2 + \beta \langle f_t, D_{\theta_t} f_t \rangle_{\mathcal{H}} \right) dt,$$

$$\partial_t \theta_t = L_{\theta_t} f_t, f_t \in \mathcal{H}, \theta_{t=0} = \theta_0, \theta_{t=1} = \theta_1. \quad \text{where}$$

Note that the regularization we introduced here is the parametric energy $\|\nabla_{\theta}[\langle f_t, \mu_{q_{\theta}} \rangle_{\mathcal{H}}]\|_{\theta=\theta_t}^2 = \langle f_t, D_{\theta_t} f_t \rangle_{\mathcal{H}}$, which plays a similar role as the kinetic energy in Benamou-Brenier's formula. The evolution of θ_t in our form is analogous to the continuity equation in the \mathcal{W}_2 dynamic form [25]. Conditions on the kernel and the generator family under which we can guarantee existence of the solution for this problem are out of the scope of this work, since our interest in this heuristic Riemannian structure is solely in order to define an appropriate tangent space and Riemannian metric tensor. We leave the analysis of $d_{\alpha,\beta}$ to a future work.

This dynamic formulation gives rise to the following Riemannian metric tensor on the tangent space of Θ : for $\theta \in \Theta$, let

$$g_{\theta}(\xi_1, \xi_2) := \alpha \langle D_{\theta} \varphi_1, D_{\theta} \varphi_2 \rangle_{\mathcal{H}} + \beta \langle \varphi_1, D_{\theta} \varphi_2 \rangle_{\mathcal{H}}$$

$$= \langle (\alpha D_{\theta} + \beta I) \varphi_1, D_{\theta} \varphi_2 \rangle_{\mathcal{H}}$$

where $\xi_i = L_{\theta}(\varphi_i) = \nabla_{\theta}[\langle \varphi_i, \mu_{q_{\theta}} \rangle_{\mathcal{H}}] \in \mathbb{R}^p$ with $\varphi_i \in \mathcal{H}$ ($i = 1, 2$). We note that $g_{\theta}(\xi_1, \xi_2) = g_{\theta}(\xi_2, \xi_1)$ due to the symmetry of D_{θ} (see property ii) of Proposition 2.3). Then it follows that

$$d_{\alpha,\beta}(q_{\theta_0}, q_{\theta_1})^2 = \min_{(\theta_t, f_t)} \left\{ \int_0^1 g_{\theta}(\partial_t \theta_t, \partial_t \theta_t) dt \right\},$$

$$\partial_t \theta_t = L_{\theta_t} f_t, f_t \in \mathcal{H}, \theta_{t=0} = \theta_0, \theta_{t=1} = \theta_1.$$

Let us assume that $\alpha, \beta > 0$ from now on. For a functional $\mathcal{F} : \mathcal{P}_{\Theta} \rightarrow \mathbb{R}$, let $\text{grad}_{d_{\alpha,\beta}} \mathcal{F}(q_{\theta})$ denote the gradient of \mathcal{F} with respect to the metric $d_{\alpha,\beta}$. That is, $\text{grad}_{d_{\alpha,\beta}} \mathcal{F}(q_{\theta})$ is a vector in \mathbb{R}^p satisfying

$$\left. \frac{d\mathcal{F}(q_{\theta_t})}{dt} \right|_{t=0} = g_{\theta}(\text{grad}_{d_{\alpha,\beta}} \mathcal{F}(q_{\theta}), \xi) \quad (21)$$

for every differentiable curve $t \mapsto \theta_t \in \Theta$ with $\theta_{t=0} = \theta$ and $\partial_t \theta_t|_{t=0} = \xi = L_{\theta} \varphi$ for some $\varphi \in \mathcal{H}$. The next theorem shows us how to compute such gradient.

Theorem 6.3. Let $\mathcal{F}(q_{\theta})$ be a functional depending only on the kernel mean embedding of q_{θ} . Precisely, assume that $\mathcal{F}(q_{\theta}) = H(\mu_{q_{\theta}})$ for some functional H with the chain rule property

$$\partial_{\theta_i} [H(\mu_{q_{\theta}})] = \langle h_{\theta}, \partial_{\theta_i} [\mu_{q_{\theta}}] \rangle_{\mathcal{H}} \quad (22)$$

for some function $h_{\theta} \in \mathcal{H}$ and for all $\theta \in \Theta$. Then the gradient of \mathcal{F} w.r.t. the discrepancy $d_{\alpha,\beta}$ is given by

$$\text{grad}_{d_{\alpha,\beta}} \mathcal{F}(q_{\theta}) = L_{\theta} u,$$

$$(\alpha D_{\theta} + \beta I) u = h_{\theta}. \quad (23)$$

Let $\mathcal{F}(q_{\theta})$ be the functional as in Theorem 6.3, and consider the gradient flow of $\mathcal{F}(q_{\theta})$ with respect to $d_{\alpha,\beta}$. We note that this is a gradient regularized flow. According to Theorem 6.3, the equation of this flow is given by

$$\partial_t \theta_t = -\text{grad}_{d_{\alpha,\beta}} \mathcal{F}(q_{\theta_t}) = -L_{\theta_t} u_t, \quad (24)$$

where

$$(\alpha D_{\theta_t} + \beta I) u_t = h_{\theta_t}. \quad (25)$$

The following proposition shows that these gradient flows are indeed descent directions of the functional defined on the statistical manifold \mathcal{P}_{Θ} :

Proposition 6.4. Along the gradient flow (24)–(25) of $\mathcal{F}(q_{\theta})$, we have

$$\begin{aligned} \frac{d}{dt} \mathcal{F}(q_{\theta_t}) &= -[\alpha \|D_{\theta_t} u_t\|_{\mathcal{H}}^2 + \beta \|L_{\theta_t} u_t\|^2] \\ &= -\frac{1}{\alpha} [\|h_{\theta_t}\|_{\mathcal{H}}^2 - \beta \langle h_{\theta_t}, (\alpha D_{\theta_t} + \beta I)^{-1} h_{\theta_t} \rangle_{\mathcal{H}}] \\ &\leq 0, \end{aligned}$$

where h_{θ} is defined by (22). Moreover, we have $\frac{d\mathcal{F}(q_{\theta_t})}{dt} < 0$ if and only if $D_{\theta_t} h_{\theta_t} \neq 0$.

Intuition on the role of D_{θ} in the gradient flow. The operator D_{θ} plays a central role in our framework and we give here an intuitive interpretation of its role from the gradient flow lens. Let $\{(\lambda_j(\theta), v_j(\theta))\}_{j=1}^{\infty}$ be the eigenvalues and eigenfunctions of the operator D_{θ} . Then as we have $u_t = \sum_{j=0}^{\infty} \frac{1}{\alpha \lambda_j(\theta_t) + \beta} \langle h_{\theta_t}, v_j(\theta_t) \rangle_{\mathcal{H}} v_j(\theta_t)$, the flow equation can be written as follows: $\partial_t \theta_t = -\text{grad}_{d_{\alpha,\beta}} \mathcal{F}(q_{\theta_t}) = -\sum_{j=0}^{\infty} \frac{1}{\alpha \lambda_j(\theta_t) + \beta} \langle h_{\theta_t}, v_j(\theta_t) \rangle_{\mathcal{H}} L_{\theta_t}(v_j(\theta_t))$. The eigenfunctions of D_{θ_t} provide the descent directions $L_{\theta_t}(v_j(\theta_t))$, that are linearly combined according to the similarity of $v_j(\theta_t)$ and h_{θ_t} and weighted by a factor $1/(\alpha \lambda_j(\theta_t) + \beta)$. Small eigenvalues are noisy directions and spectral filtering them via α and β favors descent directions with larger eigenvalues.

6.3 Gradient Flows of Particular Functionals: MMD GAN as Gradient Flow

MMD GAN as a Gradient Flow w.r.t. $d_{\alpha,\beta}$. The next result shows that the flow (5)–(10) of the parametric gradient regularized MMD GAN coincides with the

gradient flow of $\mathcal{F}(q_\theta) = \frac{1}{2}\text{MMD}(p, q_\theta)^2$ with respect to $d_{\alpha, \beta}$.

Corollary 6.5. *For $\mathcal{F}(q_\theta) = \frac{1}{2}\text{MMD}(p, q_\theta)^2$, we have $\text{grad}_{d_{\alpha, \beta}} \mathcal{F}(q_\theta) = L_\theta u$ with $(\alpha D_\theta + \beta I)u = -\mu_{p-q_\theta}$. Consequently, the parametric regularized flow (5)–(10) on the statistical manifold is the gradient flow of \mathcal{F} with respect to $d_{\alpha, \beta}$.*

Proof. This follows immediately from Theorem 6.3 noting that $\partial_{\theta_i} [\mathcal{F}(q_\theta)] = -\langle \mu_{p-q_\theta}, \partial_{\theta_i} [\mu_{q_\theta}] \rangle_{\mathcal{H}}$. \square

Gradient flows of Generic Functionals on the Statistical Manifold. Our framework is not limited to the MMD functional. In the corollary below we exhibit additional examples of functional $\mathcal{F}(q_\theta)$ defined on the statistical manifold along with their gradient flows w.r.t. $d_{\alpha, \beta}$.

Corollary 6.6. • *For the (potential) energy functional $\mathcal{F}_1(q_\theta) = \int V(x)q_\theta(dx)$ with $V \in \mathcal{H}$, we have $\text{grad}_{d_{\alpha, \beta}} \mathcal{F}_1(q_\theta) = L_\theta u$, where $(\alpha D_\theta + \beta I)u = V$.*

- *For the (entropy) functional $\mathcal{F}_1(q_\theta) = \int f(\mu_{q_\theta}(x))dx$ with $f : \mathbb{R} \rightarrow \mathbb{R}$ being continuously differentiable, we have $\text{grad}_{d_{\alpha, \beta}} \mathcal{F}_1(q_\theta) = L_\theta u$, where $(\alpha D_\theta + \beta I)u = \int f'(\mu_{q_\theta}(x))k(x, \cdot)dx$.*
- *For the (interaction) functional $\mathcal{F}_3(q_\theta) = \int \int f(x)g(y)q_\theta(dx)q_\theta(dy)$ with $f, g \in \mathcal{H}$, we have $\text{grad}_{d_{\alpha, \beta}} \mathcal{F}_3(q_\theta) = L_\theta u$, where $(\alpha D_\theta + \beta I)u = \langle f, \mu_{q_\theta} \rangle_{\mathcal{H}} g + \langle g, \mu_{q_\theta} \rangle_{\mathcal{H}} f$.*

Continuous Time Descent In this section we analyze the convergence properties of the kernelized parametric gradient flows of functionals defined on the statistical manifold w.r.t. to $d_{\alpha, \beta}$. The following proposition studies the convergence behavior of the parametric flows.

Proposition 6.7 (Convergence up to a barrier). *Let $\mathcal{F}(q_\theta)$ be the functional as in Theorem 6.3. Assume that there exists a continuous function $\gamma : \Theta \rightarrow [0, \infty)$ such that*

$$\|h_\theta\|_{\mathcal{H}}^2 \geq \gamma(\theta) \mathcal{F}(q_\theta) \quad \forall \theta \in \Theta. \quad (26)$$

Consider the dynamic $t \in [0, +\infty) \rightarrow \theta_t \in \Theta$ of the gradient flow (24)–(25). Then we have for every $t \geq 0$:

$$\mathcal{F}(q_{\theta_t}) \leq \mathcal{F}(q_{\theta_0}) e^{-\int_0^t \frac{\lambda_i(\theta_s) a(\theta_s, u_s) \gamma(\theta_s)}{\alpha \lambda_i(\theta_s) a(\theta_s, u_s) + \beta} ds}, \quad (27)$$

where $\lambda_i(\theta)$ and $a(\theta, u)$ are defined in Theorem 5.3.

Remark 6.8. *Note that condition (26) is satisfied by the MMD for $\gamma(\theta) = 2$, since $h_\theta = -\mu_{p-q_\theta}$, and hence $\|h_\theta\|_{\mathcal{H}}^2 = 2\mathcal{F}(q_\theta)$.*

7 Related Work

MMD GAN. Since their introduction in [1], many cost functions have been proposed for training GANs. Related to our work is MMD GAN introduced in [6, 7] and later improved in [8].

Gradient Regularized GANs. GAN’s training is notoriously known to be unstable. Wasserstein GAN [4] tackled that issue by considering \mathcal{W}_1 as a loss function for training. Imposing lipchitzity in practice is challenging and a gradient penalty in the input space was introduced in WGAN-GP [11] as means to impose lipchitzity. Sobolev GAN [12] connected this gradient penalty to advection and semi-Sobolev norms. The work [16] considered this input gradient regularizer in MMD GANs.

Closely related to our study are the gradient penalties on the parameter space of the generator and the discriminator considered in [13]. It was demonstrated in [14, 15] that these gradient regularizers ensure stability and convergence of the min/max game. [17] showed that GANs with a variant of a gradient penalty on the parameter space of the discriminator is locally stable. Even though these works studied the stability of the min/max game and its convergence to a saddle point, they did not show the global convergence of the discrete gradient flow in the distributional sense which is what we prove in this paper thanks to the newly proposed Riemannian structure.

Gradient Flows. Wasserstein \mathcal{W}_2 flows for minimizing functionals over probabilities (see the excellent introduction [26]) are related to our work. These flows are non parametric and are given by the continuity equation. Recently Wasserstein flows were extended to statistical manifold [27], and a kernelized Wasserstein natural gradient flow was introduced in [24]. Kernelized particle flows such as Stein Descent [28–30], Sobolev Descent [22] and MMD flows [23] are not defined on a statistical manifold and only act on the particle level. [29, 30] introduce similar Riemannian structures to ours for the Stein geometry, nevertheless they are not defined on a statistical manifold. It would be interesting to derive a dynamical form of the Stein metric on a manifold of parametric explicit densities. [31] studied minimum stein estimators and introduced a natural Stein gradient, and it would be interesting to connect their study to a definition of an appropriate dynamic Stein structure on a parametric manifold.

8 Numerical Experiments

We consider a finite dimensional RKHS defined by a random feature map [32, 33] $\Phi(x) : \Omega \rightarrow \mathbb{R}^m$. Φ

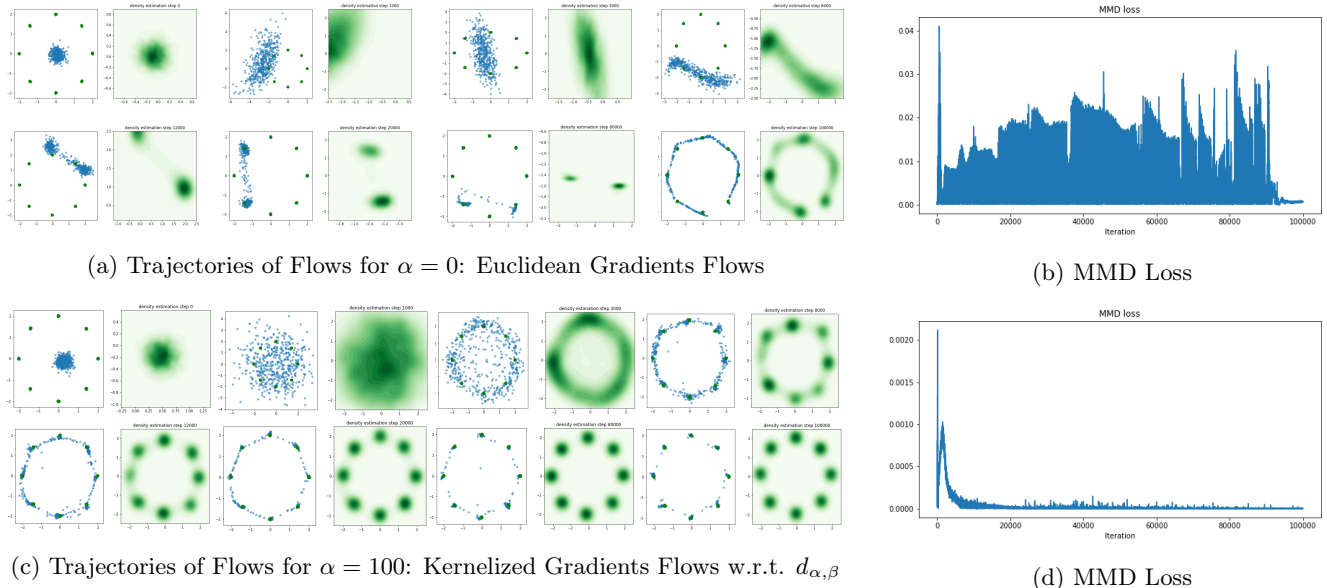


Figure 1: Trajectories of Kernelized Gradient flows of the MMD functional for $\alpha = 0$ (no gradient regularization) and $\alpha = 100$. It is clear that the Riemannian structure induced by $d_{\alpha, \beta}$, $\alpha > 0$ guarantees the convergence, while GAN suffer from cycles and mode collapse for $\alpha = 0$.

in our case is a 4 layers Relu Network with hidden dimension $m = 512$, with weights sampled from standard Gaussian and then fixed [34]. The space $\mathcal{H} = \{f(x) = \langle w, \Phi(x) \rangle, w \in \mathbb{R}^m\}$.

MMD GAN as a Kernelized Gradient Flow.

In parametric gradient regularized MMD GANs, we find the $\text{MMD}_{\alpha, \beta}(p, q_{\theta_\ell})$ witness function f_ℓ between p and q_{θ_ℓ} , and then update the generator according to: $\theta_{\ell+1} = \theta_\ell + \varepsilon_\ell L_{\theta_\ell} f_\ell$. We showed that this coincides exactly with the gradient flow of the MMD functional w.r.t. the geometric structure $d_{\alpha, \beta}$ introduced in this paper, and we proved its global convergence under mild assumptions. We show here an example, where the target distribution is a two dimensional mixture of Gaussians with 8 modes. In this example, \mathcal{Z} is 256 dimensional space and ν is standard Gaussian on \mathcal{Z} . The generator G_θ is a 2 layer Relu Network (hidden size 128) and two dimensional output. We fix the mini-batch size in the training to 512, the learning rate of the witness function and generator is set to $1e^{-4}$. We train the witness function f_ℓ of $\text{MMD}_{\alpha, \beta}$ with stochastic gradient (RmsProp [35]) for 5 iterations. Note that in this case only the last linear layer is updated in the witness function, and Φ is kept fixed. We then update the parameter θ with gradient descent on $-\mathbb{E}_{q_{\theta_\ell}} f_\ell$. We use also RmsProp to ensure that the learning rates are adaptive. Note that given the witness function f_ℓ , the discrete flow update given in Eq. (15) is exactly the gradient descent update on $-\mathbb{E}_{q_{\theta_\ell}} f_\ell$.

We give in Figure 1 the trajectories of the flows for $\alpha = 100$ and $\alpha = 0$ for $L = 100K$ iterations. Since in

our implementation (Appendix E) we use stochastic gradient for learning the witness function f_ℓ , the desired regularization effect of β is ensured by early stopping and SGD [36, 37]. The case $\alpha = 100$ corresponds to the gradient flow w.r.t. Riemannian structure $d_{\alpha, \beta}$, the case $\alpha = 0$ corresponds to unregularized MMD GAN that uses Euclidean gradients on the objective $\text{MMD}(p, q_\theta)$ [6–8]. We see from Figure 1 that the kernelized flows induced by $d_{\alpha, \beta}$ for $\alpha > 0$ ensure graceful convergence of the generator to the target distribution as predicted by Theorem 5.3, while GAN suffers from cycles and mode collapse for $\alpha = 0$.

In Appendix F, we give trajectories for neural witness functions (i.e. Φ is learned as well), and we see similar behavior to the fixed kernel. We leave for future work the extension of $d_{\alpha, \beta}$ and its gradient flows to neural function spaces as in [3, 38, 39]. Analysing the dynamic of the min-max optimization of parametric gradient regularized MMD GAN descent in the neural case will be interesting to conduct in the spirit of [40].

9 Conclusion

We propose an energy regularized gradient descent to MMD GAN training and derive a condition guaranteeing its global convergence. We demonstrate that the resulting flow coincides with the gradient flow of MMD with respect to a newly proposed Riemannian structure on the statistical manifold of probability distributions. Our investigation deepens the role of gradient regularization in GANs. Future directions include more understanding about the relationship between the static

discrepancy and the dynamic discrepancy introduced in this paper, and investigation of other variational problems on a statistical manifold using the proposed Riemannian structure.

References

- [1] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014.
- [2] Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. In *NIPS*, 2016.
- [3] Tengyu Ma. Generalization and equilibrium in generative adversarial nets (gans) (invited talk). *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing - STOC 2018*, 2018.
- [4] Martin Arjovsky, Soumith Chintala, and Leon Bottou. Wasserstein generative adversarial networks. *ICML*, 2017.
- [5] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *JMLR*, 2012.
- [6] Yujia Li, Kevin Swersky, and Richard Zemel. Generative moment matching networks. In *ICML*, 2015.
- [7] Gintare Karolina Dziugaite, Daniel M. Roy, and Zoubin Ghahramani. Training generative neural networks via maximum mean discrepancy optimization. In *UAI*, 2015.
- [8] Chun-Liang Li, Wei-Cheng Chang, Yu Cheng, Yiming Yang, and Barnabás Póczos. Mmd gan: Towards deeper understanding of moment matching network. *Advances in Neural Information Processing Systems*, 2017.
- [9] Youssef Mroueh, Tom Sercu, and Vaibhava Goel. Mcgan: Mean and covariance feature matching gan. *ICML*, 2017.
- [10] Cicero Nogueira dos Santos, Youssef Mroueh, Inkit Padhi, and Pierre Dognin. Learning implicit generative models by matching perceptual features. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [11] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of wasserstein gans. *Advances in Neural Information Processing Systems*, 2017.
- [12] Youssef Mroueh, Chun-Liang Li, Tom Sercu, Anant Raj, and Yu Cheng. Sobolev gan. *ICLR*, 2018.
- [13] Lars Mescheder, Sebastian Nowozin, and Andreas Geiger. The numerics of gans. *Advances in neural information processing systems*, 2017.
- [14] Kevin Roth, Aurelien Lucchi, Sebastian Nowozin, and Thomas Hofmann. Stabilizing training of generative adversarial networks through regularization. In *Advances in neural information processing systems*, pages 2018–2028, 2017.
- [15] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? *Proceedings of the 35th International Conference on Machine Learning*, 2018.
- [16] Michael Arbel, Dougal J. Sutherland, Mikolaj Binkowski, and Arthur Gretton. On gradient regularizers for mmd gans. *Advances in Neural Information Processing Systems*, 2018.
- [17] Vaishnavh Nagarajan and J Zico Kolter. Gradient descent gan optimization is locally stable. In *Advances in neural information processing systems*, pages 5585–5595, 2017.
- [18] Ananya Uppal, Shashank Singh, and Barnabas Poczos. Nonparametric density estimation & convergence rates for gans under besov ipm losses. In *Advances in Neural Information Processing Systems*, pages 9086–9097, 2019.
- [19] Tengyuan Liang. On how well generative adversarial networks learn distributions. *Preprint arXiv:1811.03179*, 2020.
- [20] Leon Bottou, Martin Arjovsky, David Lopez-Paz, and Maxime Oquab. Geometrical insights for implicit generative modeling. In *Braverman Readings in Machine Learning*, pages 229–268, 2017.
- [21] Krikamol Muandet, Kenji Fukumizu, Bharath Sriperumbudur, and Bernhard Schölkopf. Kernel mean embedding of distributions: A review and beyonds. *Foundations and Trends in Machine Learning*, pages 1–141, 2017.
- [22] Youssef Mroueh, Tom Sercu, and Anant Raj. Sobolev descent. In *Aistats, Proceedings of Machine Learning Research*, 2019.
- [23] Michael Arbel, Anna Korba, Adil Salim, and Arthur Gretton. Maximum mean discrepancy gradient flow. *Advances in Neural Information Processing Systems*, 2019.

-
- [24] M Arbel, A Gretton, W Li, and G Montufar. Kernelized wasserstein natural gradient. In *International Conference on Learning Representations*, 2020.
- [25] Jean-David Benamou and Yann Brenier. A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numerische Mathematik*, 2000.
- [26] Filippo Santambrogio. {Euclidean, metric, and Wasserstein} gradient flows: an overview. *Bulletin of Mathematical Sciences*, 7(1):87–154, 2017.
- [27] Yifan Chen and Wuchen Li. Optimal transport natural gradient for statistical manifolds with continuous sample space. *Information Geometry*, 2018.
- [28] Qiang Liu and Dilin Wang. Stein variational gradient descent: A general purpose bayesian inference algorithm. In *Advances in Neural Information Processing Systems 29*, 2016.
- [29] Qiang Liu. Stein variational descent as a gradient flow. *Advances in neural information processing systems*, 2017.
- [30] A. Duncan, N. Nuesken, and L. Szpruch. On the geometry of stein variational gradient descent. *Preprint arXiv:1912.00894*, 2019.
- [31] Alessandro Barp, Francois-Xavier Briol, Andrew Duncan, Mark Girolami, and Lester Mackey. Minimum stein discrepancy estimators. In *Advances in Neural Information Processing Systems*, pages 12964–12976, 2019.
- [32] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. *Advances in neural information processing systems*, 2007.
- [33] Francis Bach. Breaking the curse of dimensionality with convex neural networks. *The Journal of Machine Learning Research*, 18(1):629–681, 2017.
- [34] Youngmin Cho and Lawrence K. Saul. Kernel methods for deep learning. *Advances in Neural Information Processing Systems 22*, pages 342–350, 2009.
- [35] T. Tieleman and G. Hinton. Lecture 6.5—RmsProp: Divide the gradient by a running average of its recent magnitude. COURSERA: Neural Networks for Machine Learning, 2012.
- [36] Yuan Yao, Lorenzo Rosasco, and Andrea Caponnetto. On early stopping in gradient descent learning. *Constructive Approximation*, 2007.
- [37] Léon Bottou, Frank E Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *SIAM Rev.*, 60(2), 2018.
- [38] Lenaïc Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. In *Advances in neural information processing systems*, pages 3036–3046, 2018.
- [39] Arthur Jacot, Franck Gabriel, and Clement Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in Neural Information Processing Systems 31*, pages 8571–8580, 2018.
- [40] Carles Domingo-Enrich, Samy Jelassi, Arthur Mensch, Grant Rotskoff, and Joan Bruna. A mean-field analysis of two-player zero-sum games. *Preprint arXiv:2002.06277*, 2020.
- [41] David G. Luenberger. *Optimization by Vector Space Methods*. John Wiley & Sons, Inc., USA, 1st edition, 1997.