
DDEQs: Distributional Deep Equilibrium Models through Wasserstein Gradient Flows

Jonathan Geuter
Harvard University

Clément Bonet
CREST, ENSAE, IP Paris

Anna Korba
CREST, ENSAE, IP Paris

David Alvarez-Melis
Harvard University

Abstract

Deep Equilibrium Models (DEQs) are a class of implicit neural networks that solve for a fixed point of a neural network in their forward pass. Traditionally, DEQs take sequences as inputs, but have since been applied to a variety of data. In this work, we present Distributional Deep Equilibrium Models (DDEQs), extending DEQs to discrete measure inputs, such as sets or point clouds. We provide a theoretically grounded framework for DDEQs. Leveraging Wasserstein gradient flows, we show how the forward pass of the DEQ can be adapted to find fixed points of discrete measures under permutation-invariance, and derive adequate network architectures for DDEQs. In experiments, we show that they can compete with state-of-the-art models in tasks such as point cloud classification and point cloud completion, while being significantly more parameter-efficient.

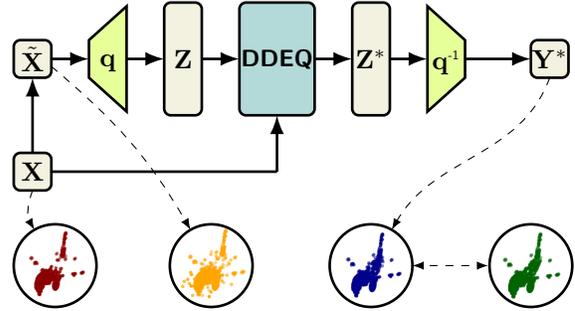


Figure 1: Point cloud completion with DDEQs for an airplane. We add random particles to the partial input point cloud \mathbf{X} (red) to create the DDEQ input $\tilde{\mathbf{X}}$ (orange), which we upscale by an invertible layer q . The DDEQ outputs a prediction \mathbf{Y}^* (blue), which is compared against the target (green) with a MMD loss.

Jacobian-based linear fixed-point equations [Bai et al., 2019] or *Phantom Gradients* [Geng et al., 2021] are used to approximate the gradient, without the need to backpropagate through a large number of stacked layers, which would be prohibitively expensive. Originally, DEQs were introduced for sequence-to-sequence tasks such as language modelling [Bai et al., 2019], since the implicit nature of the network mandates the outputs to be of the same shape as the inputs. They have since been applied to tasks such as computer vision [Bai et al., 2020] leveraging multi-scale network outputs, and exhibit competitive performance with state-of-the-art explicit models.

In this paper, we apply DEQs to data that comes in the shape of discrete probability measures seen as a set of particles. Various types of data, such as sets, graphs, or point clouds, can be viewed through this distributional lens. It arises in areas such as autonomous driving, robotics, or augmented reality [Mao et al., 2023, Zhu et al., 2024a], and often stems from sensors scanning their environments with discrete measurements. Various architectures for distributional data have been proposed, ranging from deep networks with simple feature transformations [Zaheer et al., 2017, Qi et al., 2017a,b] to attention based models [Lee et al., 2019,

1 INTRODUCTION

Implicit neural networks such as Deep Equilibrium Models (DEQs) [Bai et al., 2019] or Neural ODEs [Chen et al., 2018] have recently emerged as a promising class of networks which, instead of explicitly computing outputs of a neural network, find equilibria of certain dynamics implicitly. They simulate “infinitely deep” networks with adaptive depth, while maintaining a constant memory footprint thanks to their implicit nature.

DEQs solve for the equilibrium point of a model in their forward pass, and in the backward pass,

Proceedings of the 28th International Conference on Artificial Intelligence and Statistics (AISTATS) 2025, Mai Khao, Thailand. PMLR: Volume 258. Copyright 2025 by the author(s).

[Zhao et al., 2021]. In this work, we present DDEQs, short for Distributional Deep Equilibrium Models, which process distributional data and feature very flexible architectures that can be applied to a variety of tasks. The main challenge in using DEQs on measure inputs lies in the fixed point solver: traditional solvers such as fixed point iteration, Newton’s method, or Anderson acceleration [Bai et al., 2019] cannot produce fixed points under permutation invariance of the particles. We propose to minimize a discrepancy between probability distributions to find fixed points in the forward pass. To this end, we rely on Wasserstein gradient flows (WGFs) [Ambrosio et al., 2008, Santambrogio, 2017], which are gradient flows of a functional in the space of probability measures.

While many permutation-equivariant measure-to-measure architectures exist [Zaheer et al., 2017, Qi et al., 2017a, Lee et al., 2019], DDEQs are, to our knowledge, the first class of measure-to-measure neural networks whose forward pass is agnostic to permutations, which is an important step towards properly treating inputs as measures.

We provide a theoretically grounded framework for DDEQs, proving that they enjoy the same backward pass as traditional DEQs, deriving the Wasserstein gradient of the forward pass, and proving a convergence result for the average gradient of the optimization loss. In experiments on point cloud classification and point cloud completion, we show that DDEQs can compete with state-of-the-art architectures, and believe they can open the door to a new class of neural networks for point cloud processing.

2 RELATED WORK

Deep Equilibrium Models. Since their introduction [Bai et al., 2019], DEQs have been extended in several key directions. To improve representational capabilities, multiscale DEQs [Bai et al., 2020] introduce hidden states at different resolutions. Jacobian regularization [Bai et al., 2021] and Lyapunov-stable DEQs were shown to stabilize training. One of the main drawbacks of DEQs, their slow training, is addressed by *Phantom Gradients* [Geng et al., 2021], an approximation of the true gradient direction, which we also utilize in training DDEQs. Existence guarantees for fixed points in DEQs, as well as convergence guarantees, are difficult to obtain in theory however, and typically involve imposing restrictions on the model weights or activations [Winston and Kolter, 2020, El Ghaoui et al., 2021, Gabor et al., 2024, Ling et al., 2024]. These works also study the uniqueness of the DEQ fixed point under the assumption that the model takes the form $f(z, x) = \phi(Wz + Ux + b)$ for weights

W , U , and b . This setting only covers single-layer affine networks with fixed-size inputs; in particular, it does not include attention-based networks like transformers as we are using in this work. This goes to show that studying questions of uniqueness and existence of fixed points for DEQs is challenging, but establishing results for attention-based networks is beyond the scope of this work.

Bilevel Optimization. DEQ training can be seen as a particular instance of a *bilevel optimization* problem [Ji et al., 2021], where the goal is to minimize an outer objective

$$\mathcal{L}(\theta) = f(\theta, y_\theta^*)$$

over parameters $\theta \in \mathbb{R}^w$, with the inner problem

$$y_\theta^* = \arg \min_{y \in \mathbb{R}^p} g(\theta, y).$$

Typically, θ are the parameters of a neural network, and the inner problem corresponds to an empirical loss. The gradient of the outer problem w.r.t. the parameters θ can be computed implicitly with the inverse function theorem [Bai et al., 2019, Ye et al., 2024], whereas the inner loop is often solved iteratively with root solving algorithms [Bai et al., 2019]. However, errors in the inner loop can propagate to the outer loop, leading to inexact gradients [Ye et al., 2024]. Various approaches to reduce the inner loop error have been proposed in the literature, such as sparse fixed point corrections [Bai et al., 2022], warm starting [Bai et al., 2022, Geuter et al., 2025, Thornton and Curi, 2023, Amos et al., 2023], and preconditioning or reparametrizing the inner problem [Ye et al., 2024]. Other studies have investigated the role of the number of iterations for the inner loop at train versus inference time [Ramzi et al., 2023]. Recently, it was also proposed to jointly solve the inner and outer problem in a single loop for some bilevel optimization problems [Dagr eou et al., 2022, Marion et al., 2024]. We leave studying the effect of these strategies on DDEQs for future research.

Point Cloud Networks. Neural architectures acting on point cloud inputs have seen significant advancements in recent years, and been applied to a multitude of tasks. For classification, PointNet [Qi et al., 2017a] was groundbreaking in directly processing unordered point clouds using shared MLPs and a max-pooling operation, while PointNet++ [Qi et al., 2017b] improved upon this by capturing local structures using hierarchical feature learning. Set Transformers [Lee et al., 2019] introduced attention mechanisms to improve permutation invariance, and Point Transformers [Zhao et al., 2021] further leveraged attention for capturing both local and global features, achieving state-of-the-art performance. Notably, the attention layer in Point Trans-

formers uses vector attention instead of scalar attention (i.e. where attention scores are vector-valued), and features positional encodings which learn distances between input particles. While our architecture also builds on attention networks (see Section 3.6), we use scalar attention without positional encodings, as we did not find the attention network from Point Transformers to work particularly well for DDEQs.

For point cloud completion [Fei et al., 2022], several approaches have emerged. AtlasNet [Groueix et al., 2018] introduced a method that learns to map 2D parametrizations to 3D surfaces to reconstruct complete shapes, while PCN [Yuan et al., 2018] features a coarse-to-fine reconstruction strategy. PMP-Net [Wen et al., 2021] learns paths between particles in the input and the target, and moves input particles along these paths. However, none of these methods exactly preserve the input point cloud as part of the output. Instead, they predict a new point set that approximates the complete shape. Preserving inputs can be important in applications where exact input fidelity is required. We show that DDEQs are architecturally well-suited to predict point cloud completions that exactly preserve inputs.

Wasserstein Gradient Flows. WGFs offer an elegant way to minimize functionals with respect to probability distributions, making them an increasingly popular choice for various applications ranging from generative modeling [Fan et al., 2022, Choi et al., 2024] and sampling [Wibisono, 2018, Lambert et al., 2022, Huix et al., 2024] to reinforcement learning [Ziesche and Rozo, 2023] and optimizing machine learning algorithms such as gradient boosting [Matsubara, 2024]. With suitable discretization of the flows, various objectives for minimization have been studied, including the Kullback-Leibler divergence [Wibisono, 2018], the Maximum Mean Discrepancy (MMD) [Arbel et al., 2019, Altekruiger et al., 2023, Hertrich et al., 2024] and their variants [Glaser et al., 2021, Chen et al., 2024b, Neumayer et al., 2024, Chazal et al., 2024], Sliced-Wasserstein distances [Liutkus et al., 2019, Du et al., 2023, Bonet et al., 2025] or Sinkhorn divergences [Carlier et al., 2024, Zhu et al., 2024b].

3 DDEQs

Notation. Denote by $\mathcal{P}_2(\mathbb{R}^d)$ the space of probability measures on \mathbb{R}^d with finite second moments; for $\mu \in \mathcal{P}_2(\mathbb{R}^p)$ and a function $F : \mathbb{R}^p \rightarrow \mathbb{R}^p$, denote by $F_{\#}\mu \in \mathcal{P}_2(\mathbb{R}^p)$ the push-forward measure of μ by F , and for $\mathcal{G} : \mathcal{P}_2(\mathbb{R}^p) \rightarrow \mathbb{R}$ denote by $\nabla_W \mathcal{G}(\mu) : \mathbb{R}^d \rightarrow \mathbb{R}^p$ the Wasserstein gradient of \mathcal{G} at μ , if it exists. $\frac{\delta \mathcal{G}}{\delta \mu} : \mathbb{R}^d \rightarrow \mathbb{R}$ denotes the first variation of \mathcal{G} . For any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, we denote by $L^2(\mu)$ the space of functions

$f : \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that $\|f\|_{L^2(\mu)}^2 := \int \|f\|^2 d\mu < \infty$ and by $\text{Id} \in L^2(\mu)$ the identity map. Background on measure theory, optimal transport, and Wasserstein gradient flows, as well as all missing proofs, can be found in Appendix B.

3.1 Deep Equilibrium Models

Given an input data $x \in \mathbb{R}^d$, a DEQ initializes a latent variable z and aims to find a fixed point $z^* \in \mathbb{R}^p$ of the function $F_\theta(\cdot, x)$, where F_θ is a neural network parameterized by weights $\theta \in \mathbb{R}^w$. Adding the latent variable z instead of directly operating on x has several advantages, such as parametrizing a unique fixed point function $F_\theta(\cdot, x)$ for each input x instead of using the same function $F_\theta(\cdot)$ for all inputs, which improves the richness of the fixed points, and the dimension of the latent z can be chosen arbitrarily, increasing the flexibility and complexity of the model. One way to find fixed points would be via function iterations:

$$z^{l+1} = F_\theta(z^l, x), \quad l = 0, 1, \dots, L-1,$$

for a large L . This sequence could be backpropagated through to find gradient directions; however, this would be prohibitively expensive. Hence, root solvers such as Newton’s or Anderson’s method are used to find roots of the function $g_\theta(z) = F_\theta(z, x) - z$. Then, the gradient for the backward pass can be computed through the inverse function theorem [Blondel and Roulet, 2024]. Suppose the fixed point z^* is passed through a non-parametric function h , and a loss l between $h(z^*)$ and the target y is computed as $l(h(z^*), y)$. Then the following formula holds [Bai et al., 2019]:

$$\frac{dl}{d\theta} = \frac{\partial l(h(z_\theta^*))}{\partial h} \frac{\partial h(z_\theta^*)}{\partial z_\theta^*} \left(I - \frac{\partial F_\theta(z_\theta^*, x)}{\partial z_\theta^*} \right)^{-1} \frac{\partial F_\theta(z_\theta^*, x)}{\partial \theta}. \quad (1)$$

Importantly, this is *independent* of how the fixed point was attained, hence any black-box root solver can be used. However, due to the inverse Jacobian, (1) is costly to compute precisely, but efficient approximation methods have been developed [Bai et al., 2019, Geng et al., 2021].

3.2 DDEQs as Bilevel Optimization

We will now turn to distributional inputs and latent variables. Let $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ be the input, $\mu \in \mathcal{P}_2(\mathbb{R}^p)$ be the latent variable, and $F_\theta : \mathcal{P}_2(\mathbb{R}^p) \times \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathcal{P}_2(\mathbb{R}^p)$ be a function parameterized by some $\theta \in \mathbb{R}^w$. Let \mathcal{Y} be the target space (which can be Euclidean or itself be a probability measure space), i.e., our dataset consists of samples $(\rho, y) \in \mathcal{P}_2(\mathbb{R}^d) \times \mathcal{Y}$. Depending on the task, we might further have to process a fixed

point $\mu_{\theta,\rho}^*$ of $F_\theta(\cdot, \rho)$, hence we let $h_\theta : \mathcal{P}_2(\mathbb{R}^p) \rightarrow \mathcal{Y}$, which maps latent fixed points $\mu_{\theta,\rho}^*$ to predictions such as class labels, also be parametrized by θ . Furthermore, let $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ be a loss function and $D : \mathcal{P}_2(\mathbb{R}^p) \times \mathcal{P}_2(\mathbb{R}^p) \rightarrow \mathbb{R}$ be a distance or divergence between probability distributions. We can write, for any input data $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ and $\theta \in \mathbb{R}^w$,

$$\mu_{\theta,\rho}^* = \arg \min_{\mu \in \mathcal{P}_2(\mathbb{R}^p)} D(\mu, F_\theta(\mu, \rho)) := \mathcal{G}_{\theta,\rho}(\mu), \quad (2)$$

since the optimizers of the functional $\mathcal{G}_{\theta,\rho}$ are precisely the fixed points of $F_\theta(\cdot, \rho)$. Note that existence of fixed points in DEQs is difficult to prove, and typically involves imposing restrictions on the model weights [Winston and Kolter, 2020, Gabor et al., 2024]. We leave deriving such guarantees for DDEQs for future work.

A natural optimization scheme for minimizing the loss

$$\mathcal{L}(\theta) := \mathbb{E}_{\rho, y \sim P_{\text{data}}} [\ell(h_\theta(\mu_{\theta,\rho}^*), y)] \quad (3)$$

with respect to $\theta \in \mathbb{R}^w$ is gradient descent, i.e. for all $k \geq 0$, $\theta^{k+1} = \theta^k - \gamma \nabla \mathcal{L}(\theta^k)$. However, since

$$\nabla \mathcal{L}(\theta) = \mathbb{E}_{\rho, y \sim P_{\text{data}}} [\nabla_\theta \ell(h_\theta(\mu_{\theta,\rho}^*), y)],$$

we need to solve the following two problems: *i*) find a fixed point $\mu_{\theta,\rho}^*$, given θ and ρ (*inner loop*) and *ii*) compute $\nabla_\theta \ell(h_\theta(\mu_{\theta,\rho}^*), y)$, given $\mu_{\theta,\rho}^*$ and y . This will enable us to solve (3) (*outer loop*). Hence, minimizing \mathcal{L} is a typical instance of a bilevel optimization problem. In section 3.3, we will see how to solve the inner problem (2), and in section 3.4, we will investigate computing the gradient $\nabla_\theta \ell(h_\theta(\mu_{\theta,\rho}^*), y)$ and solving the outer problem.

3.3 Inner Optimization

In this section, we focus on the inner problem (2). Let $\mathcal{G}_{\theta,\rho}(\mu) := D(\mu, F_\theta(\mu, \rho))$ for an input $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ and parameters $\theta \in \mathbb{R}^w$ (we will sometimes drop the dependency on θ and ρ for ease of reading). In the following, we assume for simplicity that we can write $F_\theta(\mu, \rho) = F_{\theta\#\mu}$, i.e., as some push-forward (dropping the dependency on a fixed ρ for ease of notation)¹. Note that this is not possible in general, as the push-forward itself might depend on μ (i.e., $F_\theta(\mu, \rho) = F_\theta(\mu)\#\mu$, as used in [Furuya et al., 2024, Castin et al., 2024]). However, it is a reasonable assumption in our setting (for more details, see Appendix D.3). Since (2) is an objective over probability measures, the natural dynamic to minimize it is following its Wasserstein gradient flow $t \mapsto \mu_t$, which solves

the following continuity equation:

$$\partial_t \mu_t = \text{div}(\mu_t \nabla_W \mathcal{G}_{\theta,\rho}(\mu_t)), \quad (4)$$

where $\nabla_W \mathcal{G}_{\theta,\rho}(\mu)$ denotes the Wasserstein gradient of $\mathcal{G}_{\theta,\rho}$ at μ . Its existence, of course, depends on the nature of \mathcal{G} , i.e., the choice of D . In practice, we use $D = \frac{1}{2} \text{MMD}^2$, the squared maximum mean discrepancy [Gretton et al., 2012], which is defined as

$$\text{MMD}^2(\mu, \nu) = \iint k(x, y) d(\mu - \nu)(x) d(\mu - \nu)(y),$$

for a symmetric positive definite kernel $k : \mathbb{R}^p \times \mathbb{R}^p \rightarrow \mathbb{R}$. Hence, our inner optimization loss takes the form

$$\mathcal{G}_{\theta,\rho}(\mu) = \frac{1}{2} \text{MMD}^2(\mu, F_\theta(\mu, \rho)).$$

The MMD is comparably fast to compute, with a time complexity of $\mathcal{O}(n^2)$ where n is the number of particles (when μ, ν are discrete and both supported on n particles), whereas e.g. the time complexity of the Sinkhorn algorithm [Cuturi, 2013] is $\mathcal{O}(n^2 \log(n)/\epsilon^2)$ [Dvurechensky et al., 2018], where ϵ is the regularization parameter and typically fairly small. Furthermore, the MMD is well defined between empirical distributions with different support (unlike, for example, the KL Divergence), and a Wasserstein gradient of the MMD to a fixed target measure exists and can be evaluated in closed form in quadratic time [Arbel et al., 2019] for smooth kernels. We will now show that this result also extends to the setting where the target is a push-forward of the source measure. To this end, we first derive the Wasserstein gradient of a functional of the form $\mu \mapsto \mathcal{F}(T\#\mu)$ for a μ -almost everywhere (a.e.) differentiable push-forward operator $T : \mathbb{R}^p \rightarrow \mathbb{R}^p$.

Proposition 1. *Let $\mu \in \mathcal{P}_2(\mathbb{R}^p)$, $\mathcal{F} : \mathcal{P}_2(\mathbb{R}^p) \rightarrow \mathbb{R}$, $T : \mathbb{R}^p \rightarrow \mathbb{R}^p \in L^2(\mu)$ a μ -a.e. differentiable map and define $\tilde{\mathcal{F}}(\mu) := \mathcal{F}(T\#\mu)$. Assume $\sup_x \|\nabla T(x)\|_{\text{op}} < +\infty$. If the Wasserstein gradient of \mathcal{F} at $T\#\mu$ exists, then the Wasserstein gradient of $\tilde{\mathcal{F}}$ at μ also exists, and it holds:*

$$\nabla_W \tilde{\mathcal{F}}(\mu) = \nabla T(\nabla_W \mathcal{F}(T\#\mu) \circ T) = \nabla \left(\frac{\delta \mathcal{F}}{\delta \mu}(T\#\mu) \circ T \right).$$

This allows us to derive the Wasserstein gradient for our MMD objective from the chain rule.

Corollary 2. *Let $\mathcal{G}(\mu) := \frac{1}{2} \text{MMD}^2(\mu, T\#\mu)$. Define the witness function f_μ as*

$$f_\mu := \int k(\cdot, y) d\mu(y) - \int k(\cdot, y) d(T\#\mu)(y),$$

where $k : \mathbb{R}^p \times \mathbb{R}^p \rightarrow \mathbb{R}$ is the kernel of the MMD, which we assume smooth. Then the Wasserstein gradient of \mathcal{G} at μ exists, and is equal to

$$\nabla_W \mathcal{G}(\mu) = \nabla f_\mu - \nabla(f_\mu \circ T). \quad (5)$$

¹This is also a slight abuse of notation, as F_θ in $F_{\theta\#\mu}$ is now a map from \mathbb{R}^p to \mathbb{R}^p .

Note that without convexity assumptions, we cannot hope to find global optimizers of (2), as with most optimization problems. Yet, following (minus) the Wasserstein gradient (5) lets us find a local minimizer. In practice, we approximate (4) by discretizing it using the forward Euler scheme, which is commonly referred to as the Wasserstein gradient descent [Bonet et al., 2024], for a step-size $\eta > 0$:

$$\mu_{k+1} = (\text{Id} - \eta \nabla_W \mathcal{G}_{\theta, \rho})(\mu_k)_{\#} \mu_k,$$

for all $k \geq 0$, where Id is the identity map on $L^2(\mu_k)$.

3.4 Outer Optimization

In this section, we derive a formula to compute the derivative of the loss $\ell := \ell(h_{\theta}(\mu_{\theta, \rho}^*, y))$ w.r.t. θ , hence of the outer loss \mathcal{L} given in (3).

Theorem 3. (Informal.) *Let $\mathcal{Y} = \mathbb{R}^n$ or $\mathcal{Y} = \mathcal{P}_2(\mathbb{R}^y)$, $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ and $y \in \mathcal{Y}$ be fixed, and $\mu_{\theta}^* \in \mathcal{P}_2(\mathbb{R}^p)$ be a fixed point of $F(\cdot, \theta) := F_{\theta}(\cdot, \rho)$. Assume that $(\text{Id} - D_{\mu_{\theta, \rho}^*} F(\mu_{\theta, \rho}^*, \theta))^{-1} \circ D_{\theta} F(\mu_{\theta, \rho}^*, \theta) : \mathbb{R}^p \rightarrow \mathbb{R}^w$ exists, where Id is the identity map in \mathbb{R}^p . Then under suitable differentiability assumptions it holds:*

$$\frac{d\ell}{d\theta} = D_h \ell(h(\mu_{\theta, \rho}^*, \theta), y) \circ [D_{\theta} h(\mu_{\theta, \rho}^*, \theta) + D_{\mu_{\theta, \rho}^*} h(\mu_{\theta, \rho}^*, \theta) \circ (\text{Id} - D_{\mu_{\theta, \rho}^*} F(\mu_{\theta, \rho}^*, \theta))^{-1} \circ D_{\theta} F(\mu_{\theta, \rho}^*, \theta)].$$

The formal version and proof of Theorem 3 is deferred to Appendix B.2 along with the definitions of the differential of F relative to μ , which we define following [Lessel and Schick, 2020]. Similar to Theorem 1 in [Bai et al., 2019], Theorem 3 shows that the gradient of ℓ w.r.t. θ can be computed implicitly, without backpropagating through the inner loop. But as is commonly done, we do not use this formula directly in the backward pass, but instead use the phantom gradient and autodifferentiation.

By analyzing the following continuous bilevel problem, similarly to [Marion et al., 2024],

$$\begin{cases} \forall \rho \in \mathcal{P}_2(\mathbb{R}^d), \partial \mu_{t, \rho} = \text{div}(\mu_{t, \rho} \nabla_W \mathcal{G}_{\theta_t, \rho}(\mu_{t, \rho})) \\ d\theta_t = -\epsilon_t \nabla \mathcal{L}(\theta_t) dt, \end{cases}$$

where $\epsilon_t > 0$ corresponds to the ratio of learning rates between the inner and the outer problems, we prove a convergence result for the average of the outer optimization gradients in $\mathcal{O}(\log(T)^2/\sqrt{T})$ in Appendix B.3, which holds under suitable regularity assumptions on D and F_{θ} and a Polyak-Łojasiewicz inequality. However, we note that these might not always hold for the specific choice of D (i.e., an MMD) and F_{θ} in our case; we refer the reader to the appendix for a more detailed discussion.

3.5 Training Algorithm

Our training procedure is described in Algorithm 1. It takes as input the number of iterations $K \in \mathbb{N}$ for the outer loop, the number of iterations $L \in \mathbb{N}$ of the inner loop, a (data) batch size $B \in \mathbb{N}$, inner loop learning rates η_l for $l \in [L]$, and outer loop learning rates γ_k for $k \in [K]$. In practice, our latents μ will be discrete measures and have a finite number of particles, where the number of particles J_{μ} is itself a hyperparameter². We denote these particles by $\mu_{\theta^k, i}(j)$ for $j = 1, \dots, J_{\mu_{\theta^k, i}}$, where $\mu_{\theta^k, i}$ is the latent corresponding to an input sample ρ_i . The initialization of $\mu_{\theta^k, i}$ as well as the dimension p of particles in the latents $\mu_{\theta^k, i}$ are chosen depending on the task, see Section 4. To speed up computations, we pad all samples in the batch with zeros (to align the number of particles) and compute the inner loop over l for all B inputs in parallel by leveraging adequate masking. We use the Riesz kernel $k(x, y) = -\|x - y\|$ for the MMD, as it has been shown to have better convergence than Gaussian kernels [Hertrich et al., 2024, Hagemann et al., 2024]. Note this is not a positive definite kernel (neither a smooth one), but the resulting distance (the energy distance) is equivalent to an MMD [Sejdicinovic et al., 2013]. Further details on the implementation can be found in Appendix C.

Algorithm 1 DDEQ Training Procedure

- 1: Initialize $\theta^0 \in \mathbb{R}^w$
 - 2: Input $K, L, B \in \mathbb{N}$, $\eta^l > 0$, $l = 0, \dots, L$; $\gamma_k > 0$, $k = 0, \dots, K - 1$; $p \in \mathbb{N}$
 - 3: **for** $k \in \{0, \dots, K - 1\}$ **do**
 - 4: sample $(\rho_1, y_1), \dots, (\rho_B, y_B) \sim P_{\text{data}}$
 - 5: Init $\mu_{\theta^k, i}$, $i = 1, \dots, B$ (e.g. $\mu_{\theta^k, i}(j) \sim \mathcal{N}(0, I_p)$ for $j = 1, \dots, J_{\mu_{\theta^k, i}}$)
 - 6: **for** $i = 1, \dots, B$ (in parallel) **do**
 - 7: **for** $l = 0, \dots, L - 1$ **do**
 - 8: $\mu_{\theta^k, i}(j) \leftarrow \mu_{\theta^k, i}(j) - \eta_l \nabla_W \mathcal{G}_{\theta^k, \rho_i}(\mu_{\theta^k, i})(\mu_{\theta^k, i}(j))$ for all particles $\mu_{\theta^k, i}(j)$,
 - 9: **end for**
 - 10: **end for**
 - 11: $\theta^{k+1} \leftarrow \theta^k - \gamma_k \frac{1}{B} \sum_{i=1}^B [\nabla_{\theta} \ell(h_{\theta}(\mu_{\theta^k, i}), y_i)]$
 - 12: **end for**
-

3.6 Architecture

In this section, we will describe our architecture and show that it fulfills certain criteria desirable for point cloud processing. We want to emphasize, however,

²For classification, this could be a fixed number, while for point cloud completion, it could be an estimate of the number of particles in the complete point cloud. More details can be found in section 4.

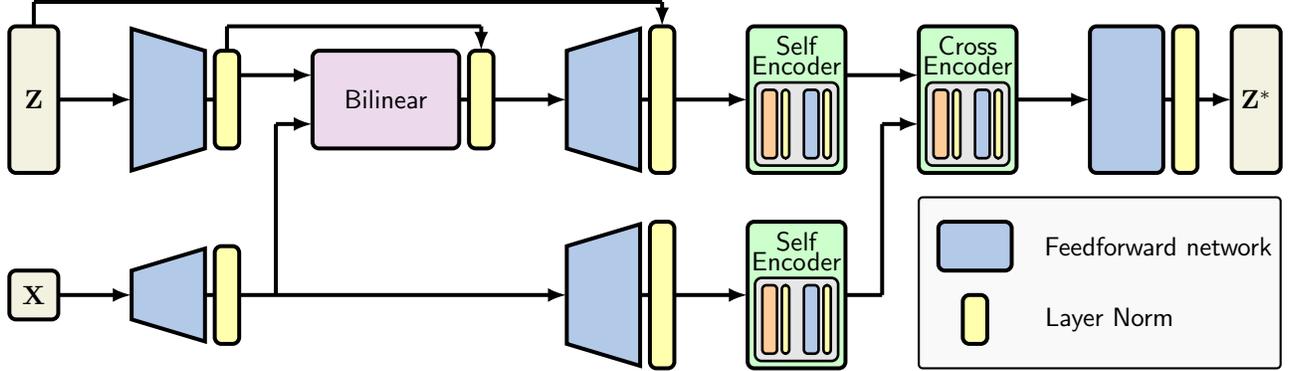


Figure 2: DDEQ Network architecture. There are two residual connection to the second and third layer norm in the top row. All encoders are standard multi-head attention network (see appendix for more details).

that DDEQs are largely architecture-independent, as long as the architecture is suitable for processing point clouds, and more carefully designed architectures could improve their performance. We leave the design of such architectures for future research.

Recall that $F = F_\theta : \mathcal{P}_2(\mathbb{R}^p) \times \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathcal{P}_2(\mathbb{R}^p)$, $(\mu, \rho) \mapsto F(\mu, \rho)$ (in this section, we remove dependency on θ for simplicity). Computationally, we will represent empirical measures as matrices by stacking the particles along the first dimension. That is, we represent $\rho = \sum_{i=1}^M \frac{1}{M} \delta_{\mathbf{x}_i} \in \mathcal{P}_2(\mathbb{R}^d)$ as a matrix $\mathbf{X} \in \mathbb{R}^{M \times d}$, and $\mu = \sum_{i=1}^N \frac{1}{N} \delta_{\mathbf{z}_i} \in \mathcal{P}_2(\mathbb{R}^p)$ as a matrix $\mathbf{Z} \in \mathbb{R}^{N \times p}$. Note that M depends on the data sample ρ and N depends on the latent μ here; in particular, the number of rows in the matrices varies across samples. However, all architectures we discuss in this section allow for variable numbers of particles in their inputs.

We will now derive a property which we will call the *EI Property* that induces a useful inductive bias to the network in our setting. A common approach to design neural networks that act on matrices encoding data without inherent ordering (e.g., sets or empirical distributions), and that make predictions such as classifications, is to make them *invariant* under row permutations. This corresponds to the fact that the prediction should be independent under the *ordering* of the points in the point cloud. Similarly, in the setting where the output is another empirical measure of the same size, making the network *equivariant* under permutations of the rows of the input (i.e., if the input is permuted in a certain way, the output is permuted in the same way) is a common approach. Unlike in most well-known architectures acting on set inputs, such as Deep Sets [Zaheer et al., 2017], PointNet [Qi et al., 2017a], or Set Transformer [Lee et al., 2019], we are now dealing with *two* variables \mathbf{X} and \mathbf{Z} instead of one, which makes the analysis a bit more delicate. Since our network outputs measures that have

the same number of particles as the latent \mathbf{Z} , it is natural to consider networks $F : \mathbb{R}^{N \times p} \times \mathbb{R}^{M \times d} \rightarrow \mathbb{R}^{N \times p}$, $(\mathbf{Z}, \mathbf{X}) \mapsto F(\mathbf{Z}, \mathbf{X})$ (abusing notation a bit, as we now let F act on matrices) which are *row-equivariant* under \mathbf{Z} and *row-invariant* under \mathbf{X} . While row-invariance in \mathbf{Z} would be closer to treating \mathbf{Z} as a set, this is difficult to accomplish with standard network architectures. However, we get a type of row-invariance in \mathbf{Z} “for free” by using a Wasserstein gradient flow forward solver which finds fixed points up to permutations, as opposed to classical forward solvers.

Definition 4 (EI Property). Let $F : \mathbb{R}^{N \times p} \times \mathbb{R}^{M \times d} \rightarrow \mathbb{R}^{N \times p}$, $(\mathbf{Z}, \mathbf{X}) \mapsto F(\mathbf{Z}, \mathbf{X})$. The function F is said to be *permutation-Equivariant* in \mathbf{Z} and *permutation-Invariant* in \mathbf{X} , if for any permutations $\sigma_N \in \mathfrak{S}_N$ and $\sigma_M \in \mathfrak{S}_M$ of the rows of \mathbf{Z} and \mathbf{X} , resp., it holds that

$$F(\sigma_N(\mathbf{Z}), \sigma_M(\mathbf{X})) = \sigma_N(F(\mathbf{Z}, \mathbf{X})).$$

In that case, we say that F fulfills the *EI Property*.

By ensuring that our architecture fulfills the EI property, we induce an inductive bias particularly well-suited for point cloud data. In the following, we show that a wide range of network building blocks fulfills this property.

Proposition 5. Let $F^{\text{bil}} : \mathbb{R}^{N \times p} \times \mathbb{R}^{M \times d} \rightarrow \mathbb{R}^{N \times p}$, $(\mathbf{Z}, \mathbf{X}) \mapsto F^{\text{bil}}(\mathbf{Z}, \mathbf{X})$. Then F^{bil} is a bilinear map that fulfills the EI property if and only if F^{bil} is of the form

$$F^{\text{bil}}(\mathbf{Z}, \mathbf{X}) = \mathbf{Z} \alpha \mathbf{X}^\top \mathbf{1}_M + \mathbf{1}_N \mathbf{1}_N^\top \mathbf{Z} \beta \mathbf{X}^\top \mathbf{1}_M$$

for some tensors $\alpha, \beta \in \mathbb{R}^{p \times p \times d}$, where $\mathbf{1}_N, \mathbf{1}_M$ denote the vectors with ones in \mathbb{R}^N and \mathbb{R}^M .

EI also extends to self- and cross-attention, as well as layers that process each particle separately, such as Layer Normalization [Ba et al., 2016].

Proposition 6. Denote by $\text{MultiHead}(tgt, src)$ multi-head attention [Vaswani et al., 2017] between a target and source sequence. Then $\text{MultiHead}(\mathbf{Z}, \mathbf{Z})$ and $\text{MultiHead}(\mathbf{Z}, \mathbf{X})$ both fulfill the EI property.

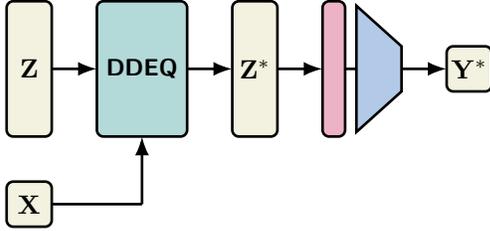


Figure 3: The point cloud classification pipeline, where \mathbf{Z} is initialized independently of \mathbf{X} , and the DDEQ is followed by a max pool (purple) and linear (blue) layer.

We can now define our network architecture for F_θ , which also fulfills the EI property, being a composition of functions that do. The core building blocks are two self-encoder networks which compute the self-attention on \mathbf{X} and \mathbf{Z} , resp., which are followed by a cross-encoder network computing the cross-attention between the target \mathbf{Z} and the source \mathbf{X} . All encoders are standard multi-head attention encoders (see Appendix C.1). An overview of the complete network architecture can be seen in Figure 2. Here, feedforward networks are two-layer linear networks with one ReLU activation, which are applied on a per-particle basis. Note how modules in the architecture have different input and output heights, which correspond to the particle dimension. The network features the dimension of \mathbf{X} (typically 2 or 3), the encoder dimension which is the same as the dimension of \mathbf{Z} , and the dimension of the bilinear layer which is typically much smaller than the encoder dimension. Details on the implementation can be found in section 4.2. Depending on the task, we pre- and post-process the in- and outputs to the DDEQ in different ways.

Classification Pipeline. For classification, we initialize $\mathbf{Z} \sim \mathcal{N}(0, I)$. We set $h_\theta(\mathbf{Z}^*) = \text{Lin}_{h_\theta}(\text{MaxPool}(\mathbf{Z}^*))$, where Lin_{h_θ} denotes a single linear layer, i.e. we first apply max-pooling along the particle dimension and then a single linear layer mapping from the latent dimension to the number of classes. The loss is the cross entropy between the predicted and target labels. An overview is found in Figure 3.

Completion Pipeline. In point cloud completion [Fei et al., 2022], the task is to predict a target point cloud \mathbf{Y} given a partial input point cloud \mathbf{X} . Hence, we could initialize \mathbf{Z} to be equal to \mathbf{X} , add a certain number of “free” particles, and then let the network learn to move the free particles to the right locations, while keeping the particles that were initialized at \mathbf{X} fixed. However, this would mean that particles in \mathbf{Z} are of the same dimension as \mathbf{X} . Instead we allow for a higher hidden dimension of the particles to increase expressiveness of the network. We first add a number

of free particles³ to \mathbf{X} and call this $\tilde{\mathbf{X}}$. Then we pad $\tilde{\mathbf{X}}$ with zeros to match the latent dimension; call this $\tilde{\mathbf{Z}}$. We then feed $\tilde{\mathbf{Z}}$ through an invertible neural network q which takes the following form:

$$\begin{aligned} \mathbf{Z} &= q(\tilde{\mathbf{Z}}) = [\tilde{\mathbf{Z}}_1, \tilde{\mathbf{Z}}_2 e^{\phi(\tilde{\mathbf{Z}}_1)} + \psi(\tilde{\mathbf{Z}}_1)] \\ q^{-1}(\mathbf{Z}) &= [\mathbf{Z}_1, (\mathbf{Z}_2 - \psi(\mathbf{Z}_1)) e^{-\phi(\mathbf{Z}_1)}], \end{aligned}$$

where ϕ and ψ are two-layer neural networks, and $\tilde{\mathbf{Z}} = [\tilde{\mathbf{Z}}_1, \tilde{\mathbf{Z}}_2]$ is a partition of $\tilde{\mathbf{X}}$ into two equal-sized chunks along the particle dimension. In this sense, q^{-1} corresponds to h_θ . The invertible network is a parametrized coupling layer [Dinh et al., 2014] which is typically used in normalizing flows [Kobyzev et al., 2021] and invertible by definition; further details can be found in the appendix. In the DDEQ forward pass, we keep the particles in \mathbf{Z} corresponding to inputs \mathbf{X} fixed. Once a fixed point \mathbf{Z}^* is found, we apply q^{-1} on \mathbf{Z}^* . This ensures the output still contains the input particles \mathbf{X} . The final loss is the squared MMD (with the same Riesz kernel as before) between the prediction and the target point cloud. Crucially, at no point in training is the model given *any information* about the class label. A schematic overview can be seen in Figure 1. We highlight two key differences between our approach and common point cloud completion algorithms from the literature, such as AtlasNet [Groueix et al., 2018], PCN [Yuan et al., 2018], or PMP-Net [Wen et al., 2021]: firstly, we recover the input point cloud *exactly* as part of the output, and secondly, we choose the number of target particles *adaptively* based on the input. Both of these properties can be crucial in faithfully recovering completed point clouds.

4 EXPERIMENTS

We evaluate DDEQs on point cloud classification and point cloud completion. Note that tasks such as point cloud segmentation [Qi et al., 2017a] that require identifying output particles with input particles are not suitable for our architecture, because the WGF solver inherently treats inputs as measures and removes any “ordering” from the particles.

More extensive experimental results, including ablation studies, a computational complexity analysis, and results on the nature of the fixed points can be found in Appendix D. The code is available at: <https://github.com/j-geuter/DDEQs>.

³as the target is not known, the number of free particles is empirically chosen such that the total number of particles in \mathbf{Z} matches the number of particles in the target on average; for more details, refer to section 4.

4.1 Datasets

We train our model on two different datasets: **MNIST point clouds**, a 2D dataset which features point clouds of digits with up to 350 particles per sample, a train split with 60,000 samples, and a test split with 10,000 samples, which we will call *MNIST-pc*; and on **ModelNet40**, a 3D dataset which features point clouds of objects of 40 different categories with up to nearly 200,000 particles per sample, a train split with 9,843 samples, and a test split with 2,468 samples. Since samples with large numbers of particles are highly redundant and too memory intense to handle, we employ voxel based downscaling, such that each sample contains no more than 800 particles, and dub the resulting dataset *ModelNet40-s*. We normalize all samples from both datasets to have zero mean and unit variance. For point cloud completion, we select two particles per (normalized) sample at random, and remove all particles in a radius r around them to create the partial point clouds, where $r = 0.6$ for MNIST and $r = 0.8$ for ModelNet. We do not impose any restrictions on the distance between the two selected particles. For ModelNet40, while we use all classes for point cloud classification, we select a subset of eight classes (airplane, bathtub, bowl, car, chair, cone, toilet, vase) for point cloud completion. We call the resulting datasets *MNIST-pc-partial* and *ModelNet40-s-partial*.

4.2 Implementation Details

We set the dimension of \mathbf{Z} and the hidden dimension of the encoders to 128. The latent dimension of the bilinear model is 16. The cross-encoder has three layers, while the self-encoders have a single layer⁴. This results in a total number of 1.17M parameters for the DDEQ. The linear classifier for MNIST-pc has 1.3k parameters, the one for ModelNet40-s 5k; the invertible coupling layer for point cloud completion has 16.6k. For classification tasks, we set the number of particles in \mathbf{Z} to 10. Note that in theory, for classification, even the edge case of just a single particle in \mathbf{Z} works, which would effectively reduce the inner loop to standard gradient descent. We found that increasing this number to 10 slightly improves performance, but further increasing it was detrimental. This shows that a DDEQ forward pass with a discretized MMD flow can be helpful even in the classification setting. An ablation study of this hyperparameter can be found in Appendix D.2. For point cloud completion, we set the number of free particles to be 27.5% of the number of particles in the input, which equals the average amount of particles

⁴For classification on MNIST-pc, we reduce the number of layers in the cross-encoder from three to one, which reduces the number of parameters to 776k, since the performance gains of more layers are marginal.

Table 1: Accuracies on Point Cloud Classification

Models (size)	MNIST-pc	ModelNet40-s
PointNet (1.6M)	97.5	77.3
PT (3.5M)	98.6	79.2
DDEQ (776k/1.2M)	98.1	78.2

removed to create the partial input point clouds.

We train the model for 5 epochs on both MNIST datasets, and for 20 epochs on ModelNet40-s for classification and 100 epochs for completion (since we select just eight classes for point cloud completion, the dataset contains only 2868 samples), both with a batch size of 64. To batch together samples with varying numbers of particles, we adequately pad samples with zeros and mask out gradients. For the inner loop, we use SGD with a learning rate of 5 and 200 iterations. The outer loop uses an Adam optimizer with initial learning rate 0.001, and a scheduler which reduces the learning rate by 90% after 40% resp. 80% of the total epochs. On MNIST-pc, we also reduce the number of layers in the cross-encoder from 3 to 1, as additional layers only provide marginal performance gains. We implement DDEQs with the `torchdeq` library [Geng and Kolter, 2023], which utilizes phantom gradients [Geng et al., 2021] for the backward pass.

Point Cloud Classification. We compare against PointNet [Qi et al., 2017a] and Point Transformer (PT) [Zhao et al., 2021], both of which we train from scratch on our datasets with the same number of epochs and with the same learning rate scheduler as DDEQs. Other hyperparameters, such as the optimizers, are taken from the original papers. The top-1 accuracies on both datasets are reported in Table 1. We see that DDEQs achieve results almost identical to PointNet, and only marginally worse than PT, while being significantly more parameter efficient.

Table 2: Comparison of W2 distances for DDEQ and PCN under noisy training data for MNIST-pc-partial.

Noisy Particles	0%	5%	20%	80%
W2 (DDEQ)	0.33	0.17	0.19	0.33
W2 (PCN)	0.49	0.51	0.49	0.53

Point Cloud Completion. We compare DDEQs against PCN [Yuan et al., 2018] with a grid size of one and 256 points per sample on MNIST-pc-partial and 768 on ModelNet40-s-partial, which yielded the best results amongst a range of hyperparameters. Again, we trained this baseline on our dataset from scratch

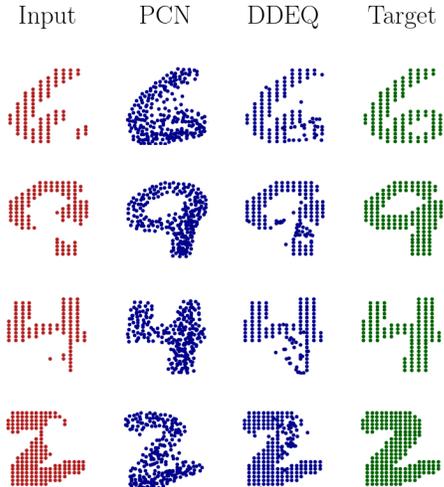


Figure 4: MNIST-pc-partial samples.

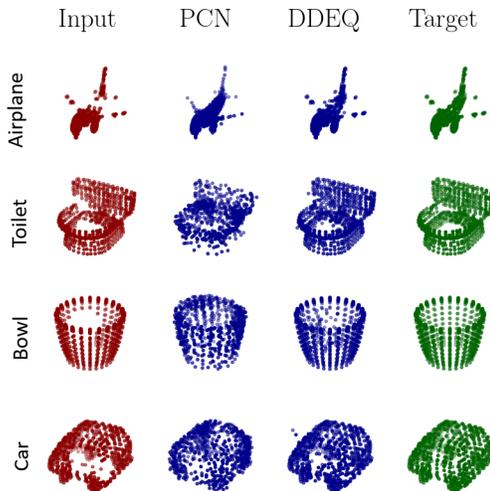


Figure 5: ModelNet40-s-partial samples.

for the same number of epochs as DDEQs⁵. While DDEQs allow for input batches to be padded with zeros, most existing frameworks require all inputs to have the same number of points, hence we filled input batches for PCN with randomly sampled points of the input point clouds which improves performance. Sample results can be seen in Figures 4 and 5; additional samples and results can be found in Appendix D.8. As we can see from the Figures, the DDEQ tends to be good at moving free particles to the right locations, but is having trouble spreading them out evenly in those locations. PCN, on the other hand, spreads out particles very evenly, but tends to produce outputs that are too diffuse, and clearly do not faithfully recover the partial input. We observed that DDEQs benefit from adding Gaussian noise from $\mathcal{N}(0, 1)$ to a small portion of the input points, as can be seen in Table 2, where we quantitatively compare DDEQs and PCN in terms of the Wasserstein-2 loss over varying noising schedules. To make the comparison more fair, we compute the Wasserstein-2 distance only over the free particles for DDEQs, whereas we compute it over all particles for PCN⁶. We hypothesize this makes training more robust, and consequently we trained DDEQs on 5% noisy samples for point cloud completion.

⁵We also tried comparing against PMP-Net [Wen et al., 2021] and PMP-Net++ [Wen et al., 2022] using the official [GitHub repo](#), but could not get it to produce predictions that differ from the input point cloud on our dataset, thus we are not including it.

⁶Arguably, this comparison slightly favours PCN, as the loss for PCN is also averaged over particles that “correspond” to input particles in some way, but it is more fair than averaging over all particles for DDEQs, as DDEQs retain the input point cloud by design.

5 CONCLUSION

In this paper, we presented DDEQs, a type of Deep Equilibrium Model which features discrete probability measure inputs and outputs, and provide a rigorous theoretical framework of the model. We derived a suitable forward pass, which we implement using a discretized Wasserstein gradient flow on the squared MMD between a latent sample and its push-forward under the neural network. We carefully analyzed the forward pass, and proposed a suitable network architecture based on multi-head attention encoders. Through experiments on point cloud classification and point cloud completion, we showed the versatility of the proposed approach, and how to apply it end-to-end on these tasks. Our model achieves competitive results on challenging tasks, such as point cloud classification on ModelNet40, while being significantly more parameter efficient. However, a more extensive empirical analysis would be instructive, and we believe that through more carefully designed architectures and hyperparameters, DDEQs can successfully be applied to a wide range of applications. One of the main drawbacks is the slow convergence of the forward pass, as is typical for MMD flows. Future research includes applying DDEQs to other tasks that require measure-to-measure learning, and to derive methods to speed up the forward pass.

ACKNOWLEDGEMENTS

CB acknowledges the support of ANR PEPR PDE-AI. JG and DAM acknowledge support from the Kempner Institute, the Aramont Fellowship Fund, and the FAS Dean’s Competitive Fund for Promising Scholarship.

References

- Fabian Altekriiger, Johannes Hertrich, and Gabriele Steidl. Neural Wasserstein Gradient Flows for Discrepancies with Riesz Kernels. In *International Conference on Machine Learning*, 2023.
- Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2008.
- Brandon Amos, Giulia Luise, Samuel Cohen, and Ievgen Redko. Meta Optimal Transport. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 791–813. PMLR, 23–29 Jul 2023.
- Michael Arbel, Anna Korba, Adil Salim, and Arthur Gretton. Maximum Mean Discrepancy Gradient Flow. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer Normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Shaojie Bai, J Zico Kolter, and Vladlen Koltun. Deep equilibrium models. *Advances in neural information processing systems*, 32, 2019.
- Shaojie Bai, Vladlen Koltun, and J Zico Kolter. Multiscale deep equilibrium models. *Advances in neural information processing systems*, 33:5238–5250, 2020.
- Shaojie Bai, Vladlen Koltun, and J. Zico Kolter. Stabilizing Equilibrium Models by Jacobian Regularization. In *International Conference on Machine Learning*, 2021.
- Shaojie Bai, Zhengyang Geng, Yash Savani, and J Zico Kolter. Deep Equilibrium Optical Flow Estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 620–630, 2022.
- Adrien Blanchet and Jérôme Bolte. A family of functional inequalities: Łojasiewicz inequalities and displacement convex functions. *Journal of Functional Analysis*, 275(7):1650–1673, 2018.
- Mathieu Blondel and Vincent Roulet. The elements of differentiable programming. *arXiv preprint arXiv:2403.14606*, 2024.
- Clément Bonet, Théo Uscidda, Adam David, Pierre-Cyril Aubin-Frankowski, and Anna Korba. Mirror and Preconditioned Gradient Descent in Wasserstein Space. In *Advances in Neural Information Processing Systems*, 2024.
- Clément Bonet, Lucas Drumetz, and Nicolas Courty. Sliced-Wasserstein Distances and Flows on Cartan-Hadamard Manifolds. *Journal of Machine Learning Research*, 26(32):1–76, 2025.
- Benoît Bonnet. A Pontryagin Maximum Principle in Wasserstein spaces for constrained optimal control problems. *ESAIM: Control, Optimisation and Calculus of Variations*, 25:52, 2019.
- Guillaume Carlier, Lénaïc Chizat, and Maxime Laborde. Displacement smoothness of entropic optimal transport. *ESAIM: Control, Optimisation and Calculus of Variations*, 30:25, 2024.
- Valérie Castin, Pierre Ablin, and Gabriel Peyré. How Smooth Is Attention? In *International Conference on Machine Learning*, 2024.
- Clémentine Chazal, Anna Korba, and Francis Bach. Statistical and Geometrical properties of regularized Kernel Kullback-Leibler divergence. In *Advances in Neural Information Processing Systems*, 2024.
- Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- Yihong Chen, Xiangxiang Xu, Yao Lu, Pontus Stenertorp, and Luca Franceschi. Jet Expansions of Residual Computation. *arXiv preprint arXiv:2410.06024*, 2024a.
- Zonghao Chen, Aratrika Mustafi, Pierre Glaser, Anna Korba, Arthur Gretton, and Bharath K Sriperumbudur. (De)-regularized Maximum Mean Discrepancy Gradient Flow. *arXiv preprint arXiv:2409.14980*, 2024b.
- Sinho Chewi, Jonathan Niles-Weed, and Philippe Rigollet. Statistical optimal transport. *arXiv preprint arXiv:2407.18163*, 2024.
- Lenaïc Chizat and Francis Bach. On the Global Convergence of Gradient Descent for Overparameterized Models using Optimal Transport. *Advances in Neural Information Processing Systems*, 31, 2018.
- Jaemoo Choi, Jaewoong Choi, and Myungjoo Kang. Scalable Wasserstein Gradient Flow for Generative Modeling through Unbalanced Optimal Transport. In *Forty-first International Conference on Machine Learning*, 2024.
- Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in Neural Information Processing Systems*, 26, 2013.
- Mathieu Dagréou, Pierre Ablin, Samuel Vaïter, and Thomas Moreau. A framework for bilevel optimization that enables stochastic and global variance re-

- duction algorithms. *Advances in Neural Information Processing Systems*, 35:26698–26710, 2022.
- Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- Sever Silvestru Dragomir. Some Gronwall type inequalities and applications. *Science Direct Working Paper*, (S1574-0358):04, 2003.
- Chao Du, Tianbo Li, Tianyu Pang, Shuicheng Yan, and Min Lin. Nonparametric Generative Modeling with Conditional Sliced-Wasserstein Flows. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 8565–8584. PMLR, 23–29 Jul 2023.
- Pavel Dvurechensky, Alexander Gasnikov, and Alexey Kroshnin. Computational Optimal Transport: Complexity by Accelerated Gradient Descent Is Better Than by Sinkhorn’s Algorithm. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1367–1376. PMLR, 10–15 Jul 2018.
- Laurent El Ghaoui, Fangda Gu, Bertrand Travacca, Armin Askari, and Alicia Tsai. Implicit Deep Learning. *SIAM Journal on Mathematics of Data Science*, 3(3):930–958, 2021.
- Jiaojiao Fan, Qinsheng Zhang, Amirhossein Taghvaei, and Yongxin Chen. Variational Wasserstein gradient flow. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 6185–6215. PMLR, 17–23 Jul 2022.
- Ben Fei, Weidong Yang, Wen-Ming Chen, Zhijun Li, Yikang Li, Tao Ma, Xing Hu, and Lipeng Ma. Comprehensive Review of Deep Learning-Based 3D Point Cloud Completion Processing and Analysis. *IEEE Transactions on Intelligent Transportation Systems*, 23(12):22862–22883, December 2022. ISSN 1558-0016. doi: 10.1109/tits.2022.3195555.
- Takashi Furuya, Maarten V de Hoop, and Gabriel Peyré. Transformers are Universal In-context Learners. *arXiv preprint arXiv:2408.01367*, 2024.
- Mateusz Gabor, Tomasz Piotrowski, and Renato LG Cavalcante. Positive concave deep equilibrium models. *arXiv preprint arXiv:2402.04029*, 2024.
- Zhengyang Geng and J Zico Kolter. Torchdeq: A library for deep equilibrium models. *arXiv preprint arXiv:2310.18605*, 2023.
- Zhengyang Geng, Xin-Yu Zhang, Shaojie Bai, Yisen Wang, and Zhouchen Lin. On Training Implicit Models. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021.
- Jonathan Geuter, Gregor Kornhardt, Ingimar Tomasson, and Vaios Laschos. Universal Neural Optimal Transport. *arXiv preprint arXiv:2212.00133*, 2025.
- Pierre Glaser, Michael Arbel, and Arthur Gretton. KALE flow: A relaxed KL gradient flow for probabilities with disjoint support. *Advances in Neural Information Processing Systems*, 34:8018–8031, 2021.
- Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012.
- Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. A papier-mâché approach to learning 3d surface generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 216–224, 2018.
- Paul Hagemann, Johannes Hertrich, Fabian Altekrüger, Robert Beinert, Jannis Chemseddine, and Gabriele Steidl. Posterior Sampling Based on Gradient Flows of the MMD with Negative Distance Kernel. In *The Twelfth International Conference on Learning Representations*, 2024.
- Johannes Hertrich, Christian Wald, Fabian Altekrüger, and Paul Hagemann. Generative Sliced MMD Flows with Riesz Kernels. In *The Twelfth International Conference on Learning Representations*, 2024.
- Tom Huix, Anna Korba, Alain Durmus, and Eric Moulines. Theoretical Guarantees for Variational Inference with Fixed-Variance Mixture of Gaussians. *International Conference on Machine Learning*, 2024.
- Kaiyi Ji, Junjie Yang, and Yingbin Liang. Bilevel Optimization: Convergence Analysis and Enhanced Design. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 4882–4892. PMLR, 18–24 Jul 2021.
- Ivan Kobyzev, Simon J.D. Prince, and Marcus A. Brubaker. Normalizing Flows: An Introduction and Review of Current Methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11):3964–3979, November 2021. ISSN 1939-3539. doi: 10.1109/tpami.2020.2992934.
- Marc Lambert, Sinho Chewi, Francis Bach, Silvère Bonnabel, and Philippe Rigollet. Variational infer-

- ence via Wasserstein gradient flows. *Advances in Neural Information Processing Systems*, 35:14434–14447, 2022.
- Nicolas Lanzetti, Saverio Bolognani, and Florian Dörfler. First-Order Conditions for Optimization in the Wasserstein Space. *SIAM Journal on Mathematics of Data Science*, 7(1):274–300, 2025.
- John M. Lee. *Introduction to Smooth Manifolds*. Graduate Texts in Mathematics. Springer New York, NY, 1 edition, March 2003. ISBN 978-0-387-21752-9. doi: 10.1007/978-0-387-21752-9. Published: 09 March 2013 (eBook).
- John M. Lee. *Introduction to Riemannian Manifolds*. Graduate Texts in Mathematics. Springer Cham, 2 edition, January 2019. ISBN 978-3-319-91754-2. doi: 10.1007/978-3-319-91755-9. Published: 14 January 2019 (Hardcover), Published: 05 August 2021 (Softcover).
- Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosior, Seungjin Choi, and Yee Whye Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In *International conference on machine learning*, pages 3744–3753. PMLR, 2019.
- Bernadette Lessel and Thomas Schick. Differentiable maps between Wasserstein spaces. *arXiv preprint arXiv:2010.02131*, 2020.
- Zenan Ling, Longbo Li, Zhanbo Feng, Yixuan Zhang, Feng Zhou, Robert C Qiu, and Zhenyu Liao. Deep Equilibrium Models are Almost Equivalent to Not-so-deep Explicit Models for High-dimensional Gaussian Mixtures. *arXiv preprint arXiv:2402.02697*, 2024.
- Antoine Liutkus, Umut Simsekli, Szymon Majewski, Alain Durmus, and Fabian-Robert Stöter. Sliced-Wasserstein flows: Nonparametric generative modeling via optimal transport and diffusions. In *International Conference on Machine Learning*, pages 4104–4113. PMLR, 2019.
- Jiageng Mao, Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. 3D object detection for autonomous driving: A comprehensive survey. *International Journal of Computer Vision*, 131(8):1909–1963, 2023.
- Pierre Marion, Anna Korba, Peter Bartlett, Mathieu Blondel, Valentin De Bortoli, Arnaud Doucet, Felipe Llinares-López, Courtney Paquette, and Quentin Berthet. Implicit Diffusion: Efficient Optimization through Stochastic Sampling. *arXiv preprint arXiv:2402.05468*, 2024.
- Takuo Matsubara. Wasserstein gradient boosting: A framework for distribution-valued supervised learning. In *Advances in Neural Information Processing Systems*, 2024.
- Leland McInnes, John Healy, and James Melville. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- Sebastian Neumayer, Viktor Stein, and Gabriele Steidl. Wasserstein Gradient Flows for Moreau Envelopes of f-Divergences in Reproducing Kernel Hilbert Spaces. *arXiv preprint arXiv:2402.04613*, 2024.
- Felix Otto. The geometry of dissipative evolution equations: The porous medium equation. *Communications in Partial Differential Equations*, 26(1-2): 101–174, 2001.
- Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017a.
- Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017b.
- Zaccharie Ramzi, Pierre Ablin, Gabriel Peyré, and Thomas Moreau. Test like you Train in Implicit Deep Learning. *arXiv preprint arXiv:2305.15042*, 2023.
- Filippo Santambrogio. Optimal transport for applied mathematicians. *Birkhäuser, NY*, 55(58-63):94, 2015.
- Filippo Santambrogio. {Euclidean, metric, and Wasserstein} gradient flows: an overview. *Bulletin of Mathematical Sciences*, 7:87–154, 2017.
- Dino Sejdinovic, Bharath Sriperumbudur, Arthur Gretton, and Kenji Fukumizu. Equivalence of distance-based and RKHS-based statistics in hypothesis testing. *The annals of statistics*, pages 2263–2291, 2013.
- James Thornton and Marco Cuturi. Rethinking initialization of the Sinkhorn algorithm. In *International Conference on Artificial Intelligence and Statistics*, pages 8682–8698. PMLR, 2023.
- Laurens van der Maaten and Geoffrey Hinton. Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz

- Kaiser, and Illia Polosukhin. Attention is All you Need. 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.
- Santosh Vempala and Andre Wibisono. Rapid convergence of the unadjusted Langevin algorithm: Isoperimetry suffices. *Advances in neural information processing systems*, 32, 2019.
- Cédric Villani. *Topics in optimal transportation*, volume 58. American Mathematical Soc., 2008.
- Xin Wen, Peng Xiang, Zhizhong Han, Yan-Pei Cao, Pengfei Wan, Wen Zheng, and Yu-Shen Liu. Pmpnet: Point cloud completion by learning multi-step point moving paths. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7443–7452, 2021.
- Xin Wen, Peng Xiang, Zhizhong Han, Yan-Pei Cao, Pengfei Wan, Wen Zheng, and Yu-Shen Liu. PMP-Net++: Point Cloud Completion by Transformer-Enhanced Multi-step Point Moving Paths. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1):852–867, 2022.
- Andre Wibisono. Sampling as optimization in the space of measures: The Langevin dynamics as a composite optimization problem. In *Conference on Learning Theory*, pages 2093–3027. PMLR, 2018.
- Ezra Winston and J Zico Kolter. Monotone operator equilibrium networks. *Advances in neural information processing systems*, 33:10718–10728, 2020.
- Zhenzhang Ye, Gabriel Peyré, Daniel Cremers, and Pierre Ablin. Enhancing Hypergradients Estimation: A Study of Preconditioning and Reparameterization. In *International Conference on Artificial Intelligence and Statistics*, pages 955–963. PMLR, 2024.
- Wentao Yuan, Tejas Khot, David Held, Christoph Mertz, and Martial Hebert. PCN: Point completion network. In *2018 international conference on 3D vision (3DV)*, pages 728–737. IEEE, 2018.
- Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola. Deep sets. *Advances in neural information processing systems*, 30, 2017.
- Valentin F Zaitsev and Andrei D Polyinin. *Handbook of exact solutions for ordinary differential equations*. Chapman and Hall/CRC, 2002.
- Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 16259–16268, 2021.
- Haoyi Zhu, Yating Wang, Di Huang, Weicai Ye, Wanli Ouyang, and Tong He. Point Cloud Matters: Rethinking the Impact of Different Observation Spaces on Robot Learning. *arXiv preprint arXiv:2402.02500*, 2024a.
- Huminhao Zhu, Fangyikang Wang, Chao Zhang, Hanbin Zhao, and Hui Qian. Neural Sinkhorn Gradient Flow. *arXiv preprint arXiv:2401.14069*, 2024b.
- Hanna Ziesche and Leonel Rozo. Wasserstein gradient flows for optimizing Gaussian mixture policies. *Advances in Neural Information Processing Systems*, 36, 2023.

Supplementary Materials

A BACKGROUND

In this section, we provide background on some of the mathematical concepts used in the paper. In Section A.1, we recall two definitions from measure theory used throughout the paper; in Section A.2, we recall some of the basic definitions from OT (see [Santambrogio, 2015, Villani, 2008] for more details). Section A.3 defines the Wasserstein distance and MMD between probability measures, and in Section A.4 we define Wasserstein gradient flows (see [Villani, 2008, Chewi et al., 2024, Lee, 2019, 2003] for more details).

A.1 Measure Theory

All functions and sets used throughout the paper are considered to be Borel measurable, and all measures are considered to be Borel measures, without explicitly mentioning it; e.g., when we say “*measurable*”, we mean Borel measurable.

Definition A.1 (Space of Measures with bounded second Moments). Denote by $\mathcal{P}(\mathbb{R}^d)$ the space of probability measures over \mathbb{R}^d with respect to the Borel- σ -algebra over \mathbb{R}^d . Then the space of probability measures of bounded second moments, denoted by $\mathcal{P}_2(\mathbb{R}^d) \subset \mathcal{P}(\mathbb{R}^d)$, contains all measures $\mu \in \mathcal{P}(\mathbb{R}^d)$ such that

$$\int \|x\|^2 d\mu(x) < \infty.$$

Definition A.2 (Pushforward). Let $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, and $F : \mathbb{R}^d \rightarrow \mathbb{R}^d$ a (measurable) function. Then the pushforward of μ under F , denoted by $F_{\#}\mu$, is a probability measure defined by

$$F_{\#}\mu(B) = \mu(F^{-1}(B))$$

for measurable $B \subset \mathbb{R}^d$.

A.2 Optimal Transport

In this section, we define the optimal transport problem on \mathbb{R}^d . Let $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$.

Definition A.3 (Coupling). The set of couplings between μ and ν is

$$\Pi(\mu, \nu) = \{\gamma \in \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d), \pi_{\#}^1 \gamma = \mu, \pi_{\#}^2 \gamma = \nu\}$$

where π^i denotes the projection on the i^{th} coordinate.

Definition A.4 (Optimal Transport Problem). Let $c : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be a cost function. The optimal transport problem is defined as:

$$\inf_{\gamma \in \Pi(\mu, \nu)} \int c(x, y) d\gamma(x, y) \tag{6}$$

The infimum in (6) is called the *transport cost*, and the minimizer γ , if it exists, the *optimal transport plan*.

Under minimal assumptions on c , such as lower semicontinuity and boundedness from below, one can prove existence of an optimal transport plan; this result extends to the more general case of Polish spaces as well [Villani, 2008].

A.3 Discrepancies between Probability Measures

Different discrepancies can be used to compare probability measures, and each comes with its own advantages and disadvantages, e.g. from a computational point of view or in terms of their theoretical properties. In this work, we focus mainly on two discrepancies: the Wasserstein distance [Villani, 2008] and the Maximum Mean Discrepancy (MMD) [Gretton et al., 2012].

Definition A.5 (Wasserstein Distance). For $1 \leq p < \infty$, the Wasserstein- p distance W_p is the p^{th} root of the transport cost of the optimal transport problem (6) with cost $c(x, y) = \|x - y\|_p^p$, i.e.

$$W_p^p(\mu, \nu) = \inf_{\gamma \in \Pi(\mu, \nu)} \int \|x - y\|_p^p d\gamma(x, y).$$

It can be shown that Wasserstein distances are indeed distances on $\mathcal{P}_p(\mathbb{R}^d)$. In particular, $\mathcal{P}_2(\mathbb{R}^d)$ endowed with W_2 , denoted by $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ and called the *Wasserstein space*, has a Riemannian structure [Otto, 2001], which allows defining notions such as tangent spaces or gradients. However, it is costly to compute in practice. In this work, we leverage the Riemannian structure of the Wasserstein space to minimize the MMD between probability distributions, which we define now.

Definition A.6 (Reproducing Kernel Hilbert Space). A Reproducing Kernel Hilbert Space (RKHS) is a Hilbert space \mathcal{H} of functions from a space \mathcal{X} to \mathbb{R} in which point evaluations are continuous linear functions. By the Riesz representation theorem, this is equivalent to the existence of a symmetric, positive definite kernel $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ such that

$$f(x) = \langle f, k(x, \cdot) \rangle_{\mathcal{H}} \quad \forall f \in \mathcal{H}.$$

Definition A.7 (Maximum Mean Discrepancy). Let \mathcal{H} be a RKHS with kernel $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$. The MMD is defined as [Gretton et al., 2012]

$$\begin{aligned} \text{MMD}(\mu, \nu) &= \iint k(x, y) d(\mu - \nu)(x) d(\mu - \nu)(y) \\ &= \sup_{f, \|f\|_{\mathcal{H}} \leq 1} \left| \int f d\mu - \int f d\nu \right|. \end{aligned}$$

The MMD can be shown to be a distance if and only if the *kernel mean embedding* $\mu \mapsto \int k(x, \cdot) d\mu(x)$ is injective, which is e.g. the case for the Gaussian kernel $k(x, y) = e^{-\frac{\|x-y\|^2}{2\sigma^2}}$ [Chewi et al., 2024]. The Riesz kernel $k(x, y) = -\|x - y\|$ used in this work is not positive definite, hence does not define an RKHS and is thus technically not an MMD. However, it can be shown to be a distance [Sejdinovic et al., 2013].

The MMD between discrete probability measures can be evaluated in $\mathcal{O}(n^2)$, where n is the number of particles. In contrast, computing the Wasserstein distance requires $\mathcal{O}(n^3 \log n)$ computations [Peyré et al., 2019].

A.4 Wasserstein Gradient Flows

In this section, we recall some definitions and properties of Wasserstein gradient flows. As $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ is a metric space that admits a quasi-Riemannian structure [Otto, 2001] (i.e., there exists a natural scalar product in each tangent space, even though it is not a Riemannian manifold per se), we can define a tangent space at $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ as

$$T_{\mu} \mathcal{P}_2(\mathbb{R}^d) = \overline{\{\nabla \psi \mid \psi : \mathbb{R}^d \rightarrow \mathbb{R} \text{ compactly supported, smooth}\}}^{L^2(\mu)} \subset L^2(\mu), \quad (7)$$

where the closure is taken in $L^2(\mu)$, see [Ambrosio et al., 2008, Definition 8.4.1]. Moreover, the tangent space is endowed with the scalar product

$$\langle v, w \rangle_{\mu} := \int \langle v, w \rangle d\mu$$

for $v, w \in T_{\mu} \mathcal{P}_2(\mathbb{R}^d)$. Intuitively, the vector fields in the tangent space can be thought of as pointing into the direction that mass moves to at any point $x \in \mathbb{R}^d$.

One can show that the Wasserstein distance is indeed the distance induced by this structure, i.e. that the squared Wasserstein distance between two distributions $\mu_0, \mu_1 \in \mathcal{P}_2(\mathbb{R}^d)$ is equal to the minimum length of a constant-speed geodesic connecting μ_0 and μ_1 (*Benamou-Brenier dynamic formulation*):

$$W_2^2(\mu_0, \mu_1) = \inf \left\{ \int_0^1 \|v_t\|_{\mu_t}^2 dt \mid (\mu_t, v_t) \text{ solves } \partial_t \mu_t + \operatorname{div}(\mu_t v_t) = 0 \right\}. \quad (8)$$

Here,

$$\partial_t \mu_t + \operatorname{div}(\mu_t v_t) = 0 \quad (9)$$

is the so-called *continuity equation*, which is satisfied whenever a path $(\mu_t)_t \subset \mathcal{P}_2(\mathbb{R}^d)$ evolves according to a time vector field $(v_t)_t$, i.e. whenever for $X_t \sim \mu_t$, it holds that

$$\dot{X}_t = v_t(X_t).$$

Hence, the integral in equation (8) equals the length of a geodesic $(\mu_t)_t$ which connects μ_0 to μ_1 along a vector field $(v_t)_t$. In general, the vector field $(v_t)_t$ satisfying the continuity equation for $(\mu_t)_t$ is not unique. However, we will see below that there exists a unique such vector field which lies in the tangent space of μ_t at all times.

Now, let $\mathcal{G} : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ be a functional.

Let us introduce the notion of sub- and super-differentiability on the Wasserstein space (see *e.g.* [Bonnet, 2019, Lanzetti et al., 2025]).

Definition A.8. Let $\mu \in \mathcal{P}_2(\mathbb{R}^d)$. A map $\xi \in L^2(\mu)$ belongs to the sub-differential $\partial^- \mathcal{G}(\mu)$ of \mathcal{F} at μ if for all $\nu \in \mathcal{P}_2(\mathbb{R}^d)$,

$$\mathcal{G}(\nu) \geq \mathcal{F}(\mu) + \sup_{\gamma \in \Pi_o(\mu, \nu)} \int \langle \xi(x), y - x \rangle d\gamma(x, y) + o(W_2(\mu, \nu)),$$

with $\Pi_o(\mu, \nu)$ the set of optimal couplings between μ and ν . Similarly, $\xi \in L^2(\mu)$ belongs to the super-differential $\partial^+ \mathcal{G}(\mu)$ of \mathcal{G} at μ if $-\xi \in \partial^-(-\mathcal{G})(\mu)$.

We also say that a functional is Wasserstein differentiable if it admits sub- and super-differentials which coincide.

Definition A.9. \mathcal{G} is called *Wasserstein differentiable* at $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ if $\partial^- \mathcal{G}(\mu) \cap \partial^+ \mathcal{G}(\mu) \neq \emptyset$. In this case, we say that $\nabla_W \mathcal{G}(\mu) \in \partial^- \mathcal{G}(\mu) \cap \partial^+ \mathcal{G}(\mu)$ is a *Wasserstein gradient* of \mathcal{G} at μ , and it satisfies for any $\nu \in \mathcal{P}_2(\mathbb{R}^d)$, $\gamma \in \Pi_o(\mu, \nu)$,

$$\mathcal{G}(\nu) = \mathcal{G}(\mu) + \int \langle \nabla_W \mathcal{G}(\mu)(x), y - x \rangle d\gamma(x, y) + o(W_2(\mu, \nu)). \quad (10)$$

If a Wasserstein gradient exists, then there is always a unique Wasserstein gradient in the tangent space $T_\mu \mathcal{P}_2(\mathbb{R}^d)$ [Lanzetti et al., 2025, Proposition 2.5], and we restrict to it in practice. Examples of Wasserstein differentiable functionals include potential energies $\mathcal{G}(\mu) = \int V d\mu$ or interaction energies $\mathcal{G}(\mu) = \frac{1}{2} \iint W(x, y) d\mu(x) d\mu(y)$ for $V : \mathbb{R}^d \rightarrow \mathbb{R}$ and $W : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ differentiable and smooth, see [Lanzetti et al., 2025, Section 2.4].

We also recall the following result of Lanzetti et al. [2025], which states that the Wasserstein gradients can be used for the Taylor expansion of functionals for any coupling.

Proposition A.10 (Proposition 2.6 in [Lanzetti et al., 2025]). *Let $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$, $\gamma \in \Pi(\mu, \nu)$ any coupling and let $\mathcal{F} : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ be Wasserstein differentiable at μ with Wasserstein gradient $\nabla_W \mathcal{F}(\mu) \in T_\mu \mathcal{P}_2(\mathbb{R}^d)$. Then,*

$$\mathcal{F}(\nu) = \mathcal{F}(\mu) + \int \langle \nabla_W \mathcal{F}(\mu)(x), y - x \rangle d\gamma(x, y) + o\left(\sqrt{\int \|x - y\|_2^2 d\gamma(x, y)}\right).$$

It turns out that under suitable regularity assumptions, the Wasserstein gradient coincides with the gradient of the first variation of the functional [Chewi et al., 2024].

Definition A.11 (First Variation). If it exists, the first variation of \mathcal{G} at μ , denoted by $\frac{\delta \mathcal{G}}{\delta \mu}(\mu) : \mathbb{R}^d \rightarrow \mathbb{R}$, is defined as the unique map such that

$$\lim_{\epsilon \rightarrow 0} \frac{\mathcal{G}(\mu + \epsilon \chi) - \mathcal{G}(\mu)}{\epsilon} = \int \frac{\delta \mathcal{G}}{\delta \mu}(\mu) d\chi$$

for all perturbations χ such that $\int d\chi = 0$.

Proposition A.12 (Lemma 10.4.1 in [Ambrosio et al., 2008]). *Assume that μ is absolutely continuous with respect to the Lebesgue measure, and that its density lies in $C^1(\mathbb{R}^d)$. Also assume that $\mathcal{G}(\mu) < \infty$. Then the Wasserstein gradient of \mathcal{G} at μ is equal to the gradient of the first variation at μ , i.e.:*

$$\nabla_W \mathcal{G}(\mu) = \nabla \frac{\delta \mathcal{G}}{\delta \mu}(\mu).$$

Moreover, for any curve $(\mu_t)_{t \geq 0}$ in $\mathcal{P}_2(\mathbb{R}^d)$ with associated tangent vectors $(v_t)_{t \geq 0}$, it holds:

$$\partial_t \mathcal{G}(\mu_t) = \langle \nabla_W \mathcal{G}(\mu_t), v_t \rangle_{\mu_t}. \quad (11)$$

Indeed, applying Proposition A.10 between μ_t and $(\text{Id} + sv_t) \# \mu_t$ for any $s > 0$, with $\gamma = (\text{Id}, \text{Id} + sv_t) \# \mu_t \in \Pi(\mu_t, (\text{Id} + sv_t) \# \mu_t)$, we get

$$\frac{\mathcal{G}((\text{Id} + sv_t) \# \mu_t) - \mathcal{G}(\mu_t)}{s} = \int \langle \nabla_W \mathcal{G}(\mu_t)(x), v_t(x) \rangle d\mu_t(x) + o(1) \xrightarrow{s \rightarrow 0} \langle \nabla_W \mathcal{G}(\mu_t), v_t \rangle_{\mu_t},$$

and thus $\partial_t \mathcal{G}(\mu_t) = \langle \nabla_W \mathcal{G}(\mu_t), v_t \rangle_{\mu_t}$.

Definition A.13 (Wasserstein Gradient Flow). A Wasserstein gradient flow (WGF) of \mathcal{G} is a path $(\mu_t)_{t \geq 0} \subset \mathcal{P}_2(\mathbb{R}^d)$ starting at some μ_0 and moving along the negative Wasserstein gradient of \mathcal{G} . In terms of the continuity equation, this reads (in the distributional sense)

$$\partial_t \mu_t + \text{div}(\mu_t v_t) = 0,$$

where $-v_t$ is a subgradient of \mathcal{G} at μ_t , i.e. for all $\nu \in \mathcal{P}_2(\mathbb{R}^d)$,

$$\mathcal{G}(\nu) \geq \mathcal{G}(\mu_t) + \sup_{\gamma \in \Pi_o(\mu_t, \nu)} \int \langle -v_t(x), y - x \rangle d\gamma(x, y) + o(W_2(\mu_t, \nu)),$$

with $\Pi_o(\mu, \nu)$ the set of optimal couplings.

Note that the velocity field generating the path $t \mapsto \mu_t$ in Definition A.13 is not unique. However, the Wasserstein gradient $\nabla_W \mathcal{G}(\mu_t)$ is the unique element amongst all such velocity fields $(v_t)_{t \geq 0}$ which lies in the tangent space of μ_t , i.e. for which $v_t \in T_{\mu_t} \mathcal{P}_2(\mathbb{R}^d)$ holds for all t .

We now recall the Wasserstein gradient of the MMD to a fixed target measure, which can be written as a sum of potential and interaction energies, see [Arbel et al., 2019] for more details.

Proposition 7. *Let $\nu \in \mathcal{P}_2(\mathbb{R}^d)$, and define $\mathcal{F}(\mu) = \frac{1}{2} \text{MMD}(\mu, \nu)^2$ for $\mu \in \mathcal{P}_2(\mathbb{R}^d)$. Then, \mathcal{F} is Wasserstein differentiable at any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, with Wasserstein gradient*

$$\nabla_W \mathcal{F}(\mu) = \int \nabla_1 k(\cdot, y) d\mu(y) - \int \nabla_1 k(\cdot, z) d\nu(z),$$

where ∇_1 denotes the gradient w.r.t. the first argument.

B PROOFS

B.1 Wasserstein Gradient

Proposition 1. *Let $\mu \in \mathcal{P}_2(\mathbb{R}^p)$, $\mathcal{F} : \mathcal{P}_2(\mathbb{R}^p) \rightarrow \mathbb{R}$, $T : \mathbb{R}^p \rightarrow \mathbb{R}^p \in L^2(\mu)$ a μ -almost everywhere differentiable map with ∇T bounded in operator norm, i.e. $\sup_x \|\nabla T(x)\|_{\text{op}} < +\infty$, and define $\tilde{\mathcal{F}}(\mu) := \mathcal{F}(T \# \mu)$. Up to a set of measure zero, if the Wasserstein gradient of \mathcal{F} at $T \# \mu$ exists, then the Wasserstein gradient of $\tilde{\mathcal{F}}$ at μ also exists, and it holds:*

$$\nabla_W \tilde{\mathcal{F}}(\mu) = (\nabla T) \cdot \nabla_W \mathcal{F}(T \# \mu) \circ T = \nabla \left(\frac{\delta \mathcal{F}}{\delta \mu}(T \# \mu) \circ T \right).$$

Proof. Let $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^p)$ and $\gamma \in \Pi_o(\mu, \nu)$ an optimal coupling between μ and ν . Let $\tilde{\gamma} = (T, T)_\# \gamma \in \Pi(T_\# \mu, T_\# \nu)$ a coupling between $T_\# \mu$ and $T_\# \nu$. Since \mathcal{F} is Wasserstein differentiable at $T_\# \mu$, we have by Proposition A.10 [Lanzetti et al., 2025, Proposition 2.6]

$$\begin{aligned} \tilde{\mathcal{F}}(\nu) - \tilde{\mathcal{F}}(\mu) &= \mathcal{F}(T_\# \nu) - \mathcal{F}(T_\# \mu) \\ &= \int \langle \nabla_W \mathcal{F}(T_\# \mu)(x), y - x \rangle d\tilde{\gamma}(x, y) + o\left(\sqrt{\int \|x - y\|^2 d\tilde{\gamma}(x, y)}\right) \\ &= \int \langle \nabla_W \mathcal{F}(T_\# \mu)(T(x)), T(y) - T(x) \rangle d\gamma(x, y) + o\left(\sqrt{\int \|T(x) - T(y)\|^2 d\gamma(x, y)}\right). \end{aligned}$$

Then, by the Jet expansion which generalizes the Taylor expansion for maps $T : \mathbb{R}^p \rightarrow \mathbb{R}^p$ (see e.g. [Chen et al., 2024a]), we have for $x, y \in \mathbb{R}^p$,

$$T(y) = T(x) + \nabla T(x)^T(y - x) + O(\|x - y\|^2).$$

Assuming that the norm operator of ∇T is bounded, we can write $\|\nabla T(x)^T(y - x)\| \leq \|\nabla T(x)\|_{\text{op}} \|y - x\| \leq M \|y - x\|$ for $M > 0$, and thus $\int \|\nabla T(x)^T(y - x)\|^2 d\gamma(x, y) \leq MW_2^2(\mu, \nu)$. Therefore, we obtain

$$\begin{aligned} \tilde{\mathcal{F}}(\nu) - \tilde{\mathcal{F}}(\mu) &= \int \langle \nabla_W \mathcal{F}(T_\# \mu)(T(x)), \nabla T(x)^T(y - x) \rangle d\gamma(x, y) + o(W_2(\mu, \nu)) \\ &= \int \langle \nabla T(x) \nabla_W \mathcal{F}(T_\# \mu)(T(x)), y - x \rangle d\gamma(x, y) + o(W_2(\mu, \nu)). \end{aligned}$$

Thus, we conclude that $\nabla_W \tilde{\mathcal{F}}(\mu) = (\nabla T) \cdot \nabla_W \mathcal{F}(T_\# \mu) \circ T$ by Definition A.9. □

Corollary 2. Let $\mathcal{G}(\mu) := \frac{1}{2} \text{MMD}^2(\mu, T_\# \mu)$. Define the witness function f_μ as

$$f_\mu := \int k(\cdot, y) d\mu(y) - \int k(\cdot, y) d(T_\# \mu)(y),$$

where $k : \mathbb{R}^p \times \mathbb{R}^p \rightarrow \mathbb{R}$ is the kernel of the MMD, which ought to be differentiable almost everywhere. Then the Wasserstein gradient of \mathcal{G} exists, and is equal to

$$\nabla_W \mathcal{G}(\mu) = \nabla f_\mu - \nabla(f_\mu \circ T). \quad (12)$$

Proof. First, recall that for $\mathcal{F}(\mu) = \frac{1}{2} \text{MMD}^2(\mu, \nu)$, $\nabla_W \mathcal{F}(\mu) = \nabla f_{\mu, \nu}$ with $f_{\mu, \nu} = \int k(\cdot, y) d\mu(y) - \int k(\cdot, z) d\nu(z)$ [Arbel et al., 2019]; note that in this notation, f_μ from above corresponds to $f_{\mu, T_\# \mu}$.

For $\mathcal{G}(\mu) = \frac{1}{2} \text{MMD}^2(\mu, T_\# \mu)$, we have $\nabla_W \mathcal{G}(\mu) = \nabla_W \mathcal{G}_1(\mu) + \nabla_W \mathcal{G}_2(\mu)$ with $\mathcal{G}_1(\nu) = \frac{1}{2} \text{MMD}^2(\nu, T_\# \mu)$ and $\mathcal{G}_2(\nu) = \frac{1}{2} \text{MMD}^2(T_\# \nu, \mu)$. Moreover,

$$\begin{aligned} \nabla_W \mathcal{G}_1(\nu) &= \nabla f_{\nu, T_\# \mu} \\ &= \int \nabla_1 k(\cdot, x') d\nu(x') - \int \nabla_1 k(\cdot, y) d(T_\# \mu)(y) \\ &= \int \nabla_1 k(\cdot, x') d\nu(x') - \int \nabla_1 k(\cdot, T(y)) d\mu(y). \end{aligned}$$

Furthermore, using Proposition 1,

$$\begin{aligned} \nabla_W \mathcal{G}_2(\nu) &= \nabla T(\nabla f_{T_\# \nu, \mu} \circ T) \\ &= \int (\nabla T) \nabla_1 k(T(\cdot), T(x')) d\nu(x') - \int (\nabla T) \nabla_1 k(T(\cdot), y) d\mu(y) \\ &= -\nabla(f_{\mu, T_\# \nu} \circ T) \end{aligned}$$

Thus,

$$\begin{aligned} \nabla_W \mathcal{G}(\mu)(x) &= \nabla_W \mathcal{G}_1(\mu)(x) + \nabla_W \mathcal{G}_2(\mu)(x) \\ &= \nabla f_{\mu, T_\# \mu}(x) - \nabla(f_{\mu, T_\# \mu} \circ T)(x). \end{aligned}$$

□

B.2 Implicit Gradient

In this section, we prove Theorem 3. A rigorous proof of this statement is more involved than one might think, since it requires the derivative of a map $F : \mathcal{P}_2(\mathbb{R}^p) \rightarrow \mathcal{P}_2(\mathbb{R}^p)$, as well as of maps $f : \mathbb{R}^w \rightarrow \mathcal{P}_2(\mathbb{R}^p)$, which we will carefully define in the following, following the construction in [Lessel and Schick, 2020] for maps $F : \mathcal{P}_2(\mathbb{R}^p) \rightarrow \mathcal{P}_2(\mathbb{R}^p)$, and extending it to maps $f : \mathbb{R}^w \rightarrow \mathcal{P}_2(\mathbb{R}^p)$.

Definition B.1 (Absolutely Continuous Curve). Let (X, d) be a metric space, and $I \subset \mathbb{R}$ an interval. A curve $\gamma : I \rightarrow X$ is called *absolutely continuous* (a.c.), if there exists a function $f \in L^1(I)$ such that

$$d(\gamma(t), \gamma(s)) \leq \int_t^s f(r) dr, \quad \forall s, t \in I, t \leq s.$$

We have seen in section A.4 that given a functional and its Wasserstein gradient, there exists a unique vector field which induces this gradient (in the sense of the continuity equation) which lies in the tangent space. This statement can be extended to any absolutely continuous curve μ_t of measures, cmp. [Lessel and Schick, 2020].

Definition B.2 (Tangent Couple). Let $\mu_t : [0, 1] \rightarrow \mathcal{P}_2(\mathbb{R}^p)$ be an absolutely continuous curve. Then there exists a unique vector field $v_t : [0, 1] \times \mathbb{R}^p \rightarrow \mathbb{R}^p$, such that (μ_t, v_t) fulfills the continuity equation (9), and such that v_t lies in the tangent space $T_{\mu_t} \mathcal{P}_2(\mathbb{R}^p)$ for (almost all) times t . We will call such a couple (μ_t, v_t) a *tangent couple*.

Similarly, a tangent couple in Euclidean space is a pair (θ_t, h_t) , where $\theta_t : [0, 1] \rightarrow \mathbb{R}^q$, $h_t : [0, 1] \rightarrow \mathbb{R}^q$ such that $h_t = \partial_t \theta_t$ (which one could extend to a vector field h_t that coincides with $\partial_t \theta_t$ on the support of θ_t , which would be a closer equivalent to the measure-theoretic definition above).

Definition B.3 (Absolutely Continuous Map). Let $X = \mathbb{R}^q$ or $X = \mathcal{P}_2(\mathbb{R}^p)$, and $F : X \rightarrow \mathcal{P}_2(\mathbb{R}^d)$. F is called *absolutely continuous* if for any a.c. curve μ_t in X , the curve $F(\mu_t) \subset \mathcal{P}_2(\mathbb{R}^d)$ is a.c. (up to redefining it on a set of measure 0).

We are now ready to define the differential of a map $F : \mathcal{P}_2(\mathbb{R}^p) \rightarrow \mathcal{P}_2(\mathbb{R}^d)$ (which we will use for $d = p$), following Definition 27 from [Lessel and Schick, 2020], which we can also naturally extend to functions $F : \mathbb{R}^q \rightarrow \mathcal{P}_2(\mathbb{R}^d)$.

Definition B.4 (Differentiable Map). Let $X = \mathbb{R}^q$ or $X = \mathcal{P}_2(\mathbb{R}^p)$, and $F : X \rightarrow \mathcal{P}_2(\mathbb{R}^d)$. We say that F is *differentiable* if there exist bounded linear maps $DF_\mu : T_\mu X \rightarrow T_{F(\mu)} \mathcal{P}_2(\mathbb{R}^d)$ such that for every tangent couple (μ_t, v_t) in X , the pair $(F(\mu_t), DF_{\mu_t}(v_t))$ fulfills the continuity equation and is a tangent couple in $\mathcal{P}_2(\mathbb{R}^d)$.

As noted in [Lessel and Schick, 2020], it is difficult to define pointwise differentiability for maps between Wasserstein spaces, as tangent vector fields v_t are not defined pointwise. Hence, to be precise, all pointwise statements in the following will only hold almost everywhere.

We can now state the chain rule from [Lessel and Schick, 2020], and extend it to compositions between Euclidean and Wasserstein space.

Proposition B.5 (Chain Rule). Let $X = \mathbb{R}^q$ or $X = \mathcal{P}_2(\mathbb{R}^q)$ and $Z = \mathbb{R}^d$ or $Z = \mathcal{P}_2(\mathbb{R}^d)$. Let $G : X \rightarrow \mathcal{P}_2(\mathbb{R}^p)$, and $F : \mathcal{P}_2(\mathbb{R}^p) \rightarrow Z$ be differentiable. Then $F \circ G : X \rightarrow Z$ is also differentiable, and

$$D(F \circ G)_\mu(v) = (DF_{G(\mu)} \circ DG_\mu)(v).$$

Proof. The case where both X and Z are Wasserstein spaces is proven in [Lessel and Schick, 2020, Corollary 33. 3)]. Note that the proof works exactly the same if X or Z is Euclidean using the right notion of tangent couple. \square

Remark B.6. The result from Proposition B.5 recovers equation (11) in the case $X = Z = \mathbb{R}$. To see this, let $F : \mathcal{P}_2(\mathbb{R}^p) \rightarrow \mathbb{R}$ be a Wasserstein differentiable map, $(\mu_t)_{t \geq 0}$ a curve in $\mathcal{P}_2(\mathbb{R}^p)$ with associated tangent vectors $(v_t)_{t \geq 0}$. By (11) we get the chain rule

$$\partial_t F(\mu_t) = \langle \nabla_W F(\mu_t), v_t \rangle_{\mu_t}.$$

Thus $(F(\mu_t), \langle \nabla_W F(\mu_t), v_t \rangle_{\mu_t})$ is a tangent couple on \mathbb{R} , and by Definition B.4 this implies

$$DF_{\mu_t}(v_t) = \langle \nabla_W F(\mu_t), v_t \rangle_{\mu_t}.$$

Moreover, with $G : t \mapsto \mu_t$ we get that $DG_t(v_t) = v_t$ is the differential of G at t by Definition B.4 since (μ_t, v_t) is a tangent couple. By Proposition B.5, we have

$$\partial_t F(\mu_t) = D(F \circ G)_t(v_t) = DF_{G(t)}(DG_t(v_t)) = DF_{\mu_t}(v_t) = \langle \nabla_W F(\mu_t), v_t \rangle_{\mu_t}, \quad (13)$$

which is the formula from equation (11).

We are now ready to prove Theorem 3. In the following, the notation $D_x f(x, y)$ will denote the derivative of f w.r.t. its first argument at x .

Theorem 3. *Let $\mathcal{Y} = \mathbb{R}^n$ or $\mathcal{Y} = \mathcal{P}_2(\mathbb{R}^n)$, $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ and $y \in \mathcal{Y}$ be fixed, and $\mu_{\theta, \rho}^* \in \mathcal{P}_2(\mathbb{R}^p)$ be a fixed point of $F(\cdot, \theta) := F_\theta(\cdot, \rho)$ found by a given fixed point solver, where we assume that the fixed point solver $\mathbb{R}^w \rightarrow \mathcal{P}_2(\mathbb{R}^p)$ is deterministic, i.e. given $\theta \in \mathbb{R}^w$, it finds a unique fixed point $\mu_{\theta, \rho}^*$ differentiable w.r.t. θ . Let $\ell = \ell(h(\mu_{\theta, \rho}^*), y)$ be the differentiable loss, where $h : \mathcal{P}_2(\mathbb{R}^d) \times \mathbb{R}^w \rightarrow \mathcal{Y}$ is differentiable. Assume further that F is differentiable, and that $\text{Id} - D_\mu F$ is injective, where Id denotes the identity on $T_{\mu_{\theta, \rho}^*} \mathcal{P}_2(\mathbb{R}^p)$, and where $D_\mu F$ is short-hand notation for $D_\mu F(\mu, \theta)$, i.e. the derivative of $F(\mu, \theta)$ w.r.t. μ . Then ℓ is differentiable w.r.t. θ , and it holds:*

$$\frac{d\ell}{d\theta} = D_h \ell(h(\mu_{\theta, \rho}^*), y) \circ \left[D_\theta h(\mu_{\theta, \rho}^*), \theta \right] + D_{\mu_{\theta, \rho}^*} h(\mu_{\theta, \rho}^*), \theta \circ \left(\text{Id} - D_{\mu_{\theta, \rho}^*} F(\mu_{\theta, \rho}^*), \theta \right)^{-1} \circ D_\theta F(\mu_{\theta, \rho}^*), \theta \Big].$$

Proof. We have by the chain rule, Proposition B.5,

$$\frac{d\ell}{d\theta} = D_h \ell(h(\mu_{\theta, \rho}^*), y) \circ \left[D_\theta h(\mu_{\theta, \rho}^*), \theta \right] + D_{\mu_{\theta, \rho}^*} h(\mu_{\theta, \rho}^*), \theta \circ D_\theta \mu_{\theta, \rho}^*. \quad (14)$$

Furthermore, by the chain rule once more,

$$D_\theta \mu_{\theta, \rho}^* = D_{\mu_{\theta, \rho}^*} F(\mu_{\theta, \rho}^*), \theta \circ D_\theta \mu_{\theta, \rho}^* + D_\theta F(\mu_{\theta, \rho}^*), \theta.$$

Using the fact that $\text{Id} - D_{\mu_{\theta, \rho}^*} F(\mu_{\theta, \rho}^*), \theta$ is injective, we can rewrite this as

$$D_\theta \mu_{\theta, \rho}^* = \left(\text{Id} - D_{\mu_{\theta, \rho}^*} F(\mu_{\theta, \rho}^*), \theta \right)^{-1} \circ D_\theta F(\mu_{\theta, \rho}^*), \theta.$$

Note that by definition, $D_{\mu_{\theta, \rho}^*} F(\mu_{\theta, \rho}^*), \theta$ is a map $T_{\mu_{\theta, \rho}^*} \mathcal{P}_2(\mathbb{R}^p) \mapsto T_{F(\mu_{\theta, \rho}^*), \theta} \mathcal{P}_2(\mathbb{R}^p)$. However, since $\mu_{\theta, \rho}^* = F(\mu_{\theta, \rho}^*), \theta$, it can be viewed as a map $T_{\mu_{\theta, \rho}^*} \mathcal{P}_2(\mathbb{R}^p) \mapsto T_{\mu_{\theta, \rho}^*} \mathcal{P}_2(\mathbb{R}^p)$, hence the above expression makes sense. Plugging this back into equation (14) yields the result. \square

An important question regarding Theorem 3 is under what conditions we get differentiability. In [Lessel and Schick, 2020], a sufficient condition is given for pushforward maps. The study of the differentiability of non pushforward maps is left for future works.

Proposition B.7. *Let $F : \mathcal{P}_2(\mathbb{R}^p) \rightarrow \mathcal{P}_2(\mathbb{R}^p)$ be such that $F(\mu) = f_{\#} \mu$ for a function $f : \mathbb{R}^p \rightarrow \mathbb{R}^p$ that is proper (i.e., preimages of compact sets are compact), smooth, and such that $\sup_x \|Df_x\| < \infty$. Then F is differentiable.*

B.3 Loss Convergence

Let $F_\theta : \mathcal{P}_2(\mathbb{R}^d) \times \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathcal{P}_2(\mathbb{R}^d)$ and $\mathcal{G}_{\theta, \rho}(\mu) = D(\mu, F_\theta(\mu, \rho))$ with D a differentiable divergence, assumed symmetric for simplicity. Recall that we want to minimize $\mathcal{L}(\theta) = \mathbb{E}_{\rho, y} [\ell(h(\mu_{\theta, \rho}^*), y)]$ with $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ a differentiable loss function and $\mu_{\theta, \rho}^* = \arg \min_{\mu} \mathcal{G}_{\theta, \rho}(\mu)$.

We propose to solve this problem by alternating between a gradient descent step on θ and a Wasserstein gradient descent step on $\mathcal{G}_{\theta, \rho}$. Inspired from [Dagr eou et al., 2022, Marion et al., 2024], we discuss the convergence of this bilevel optimization problem by analyzing the continuous process

$$\begin{cases} \forall \rho \in \mathcal{P}_2(\mathbb{R}^d), \partial \mu_{t, \rho} = \text{div}(\mu_{t, \rho} \nabla_W \mathcal{G}_{\theta_t, \rho}(\mu_{t, \rho})) \\ d\theta_t = -\varepsilon_t \nabla \mathcal{L}(\theta_t) dt, \end{cases}$$

with $t \mapsto \varepsilon_t$ a decreasing and positive curve. We analyze the convergence under several assumptions which we describe now.

Assumption B.8. For any $\rho \in \mathcal{P}_2(\mathbb{R}^d)$, $\theta \in \mathbb{R}^w$, the point fixes $\mu_{\theta,\rho}^*$ are absolutely continuous with respect to the Lebesgue measure and $\mu_{\theta,\rho}^* \propto e^{-V_\rho(\cdot,\theta)}$, with $V_\rho : \mathbb{R}^p \times \mathbb{R}^w \rightarrow \mathbb{R}^p$ continuously differentiable and satisfying $\|\nabla_2 V_\rho(x, \theta)\| \leq C_V$.

The assumption of absolute continuity allows us to write $\nabla \mathcal{L}(\theta) = \mathbb{E}_{\rho,y}[\Gamma(\mu_{\theta,\rho}^*, \theta)]$ with $\Gamma(\mu, \theta) = \nabla_1 \ell(h(\mu), y) \cdot \text{Cov}_\mu\left(\frac{\delta h}{\delta \mu}(\mu)(Z), \nabla_2 V_\rho(Z, \theta)\right)$. We also assume boundedness of Γ and Lipschitzness with respect to D .

Assumption B.9. For any $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$, $\theta \in \mathbb{R}^w$, $\|\Gamma(\mu, \theta) - \Gamma(\nu, \theta)\| \leq K_\Gamma D(\mu, \nu)$ (K_Γ -Lipschitzness) and $\|\Gamma(\mu, \nu)\| \leq C_\Gamma$.

Now, let us define $\mathcal{F}_\nu(\mu) := D(\mu, \nu)$. In this notation, ν is assumed to be fixed, meaning that in expressions like $\nabla_W \mathcal{F}_\nu(\mu)$, we differentiate only w.r.t. the argument μ . We assume that $\mathcal{F}_{\mu_{\theta,\rho}^*}$ satisfies the Polyak-Lojasiewicz (PL) inequality, which provides a sufficient condition for the convergence of the gradient flows of $\mathcal{F}_{\mu_{\theta,\rho}^*}$ [Blanchet and Bolte, 2018].

Assumption B.10. For all $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, $\mathcal{F}_{\mu_{\theta,\rho}^*}(\mu) \leq C_{PL} \|\nabla_W \mathcal{F}_{\mu_{\theta,\rho}^*}(\mu)\|_{L^2(\mu)}^2$ for $C_{PL} \geq 0$ (PL inequality). Additionally, we suppose $\|\frac{\delta \mathcal{F}_{\mu_t}}{\delta \mu_t}(\mu_{\theta_t,\rho}^*)\|_{L^2(\mu_{\theta_t,\rho}^*)} \leq C_{\mathcal{F}}$.

Finally, we also need to control the discrepancy between the directions given by the gradient of $\mathcal{G}_{\theta,\rho}$ and the gradients of $\mathcal{F}_{\mu_{\theta,\rho}^*}$ towards the objective. We note that this assumption also depends on properties of the map F_θ .

Assumption B.11. There exists some constant $M > 0$ such that for all $t \geq 0$,

$$\|\nabla_W \mathcal{G}_{\theta_t,\rho}(\mu_t) - \nabla_W \mathcal{F}_{\mu_{\theta_t,\rho}^*}(\mu_t)\|_{L^2(\mu_t)}^2 \leq M \varepsilon_t. \quad (15)$$

We now show the convergence of the average of the objective gradients, which is the best we can hope for since we do not make any convexity assumptions.

Theorem B.12. Take $\varepsilon_t = \min(1, t^{-\frac{1}{2}})$, and assume Assumptions B.8, B.9, B.10 and B.11. Then for $T > 0$ and some constant $c > 0$,

$$\frac{1}{T} \int_0^T \|\nabla \mathcal{L}(\theta_t)\|^2 dt \leq c \frac{(\log T)^2}{\sqrt{T}}.$$

Proof. First, by Assumption B.8, $d\mu_{\theta,\rho}^* = \frac{e^{-V_\rho(\cdot,\theta)}}{Z_{\theta,\rho}} d\text{Leb}$ with $Z_{\theta,\rho} = \int e^{-V_\rho(z,\theta)} dz$. Moreover, observe that

$$\frac{\partial \mu_{\theta,\rho}^*}{\partial \theta} = (-\nabla_2 V_\rho(\cdot, \theta) + \mathbb{E}_{\mu_{\theta,\rho}^*}[\nabla_2 V_\rho(Z, \theta)]) \mu_{\theta,\rho}^*,$$

where we write $\mu_{\theta,\rho}^*$ for the density with an abuse of notation.

Then, taking the gradient of $\mathcal{L}(\theta) = \mathbb{E}_{\rho,y}[\ell(h(\mu_\theta^*), y)]$, we obtain

$$\begin{aligned} \nabla \mathcal{L}(\theta) &= \mathbb{E}_{\rho,y} \left[\nabla_1 \ell(h(\mu_{\theta,\rho}^*), y) \int \frac{\delta h}{\delta \mu}(\mu_{\theta,\rho}^*)(z) \frac{\partial \mu_{\theta,\rho}^*}{\partial \theta}(z) dz \right] \\ &= \mathbb{E}_{\rho,y} \left[\nabla_1 \ell(h(\mu_{\theta,\rho}^*), y) \cdot \int \frac{\delta h}{\delta \mu}(\mu_{\theta,\rho}^*)(z) \left(-\frac{\nabla_2 V_\rho(z, \theta) e^{-V_\rho(z, \theta)}}{Z_{\theta,\rho}} + \frac{e^{-V_\rho(z, \theta)} \int \nabla_2 V_\rho(y, \theta) e^{-V_\rho(y, \theta)} dy}{Z_{\theta,\rho}^2} \right) dz \right] \\ &= \mathbb{E}_{\rho,y} \left[\nabla_1 \ell(h(\mu_{\theta,\rho}^*), y) \cdot \left(-\mathbb{E}_{\mu_{\theta,\rho}^*} \left[\frac{\delta h}{\delta \mu}(\mu_{\theta,\rho}^*)(Z) \nabla_2 V_\rho(Z, \theta) \right] + \mathbb{E}_{\mu_{\theta,\rho}^*}[\nabla_2 V_\rho(Z, \theta)] \mathbb{E}_{\mu_{\theta,\rho}^*} \left[\frac{\delta h}{\delta \mu}(\mu_{\theta,\rho}^*)(Z) \right] \right) \right] \\ &= \mathbb{E}_{\rho,y} \left[\nabla_1 \ell(h(\mu_{\theta,\rho}^*), y) \cdot \text{Cov}_{\mu_{\theta,\rho}^*} \left(\frac{\delta h}{\delta \mu}(\mu_{\theta,\rho}^*)(Z), \nabla_2 V_\rho(Z, \theta) \right) \right]. \end{aligned}$$

Let $\Gamma(\mu, \theta) = \nabla_1 \ell(h(\mu), y) \cdot \text{Cov}_\mu\left(\frac{\delta h}{\delta \mu}(\mu)(Z), \nabla_2 V_\rho(Z, \theta)\right)$, then we have $\nabla \mathcal{L}(\theta) = \mathbb{E}_{\rho,y}[\Gamma(\mu_{\theta,\rho}^*, \theta)]$.

Now, we compute $\frac{d}{dt}\mathcal{L}(\theta_t)$ by first applying the chain rule, using that $\frac{d\theta_t}{dt} = -\varepsilon_t \mathbb{E}_{\rho,y}[\Gamma(\mu_t, \theta_t)]$ and remembering that $\nabla\mathcal{L}(\theta_t) = \mathbb{E}_{\rho,y}[\Gamma(\mu_{\theta_t}^*, \theta_t)]$.

$$\begin{aligned} \frac{d}{dt}\mathcal{L}(\theta_t) &= \left\langle \nabla\mathcal{L}(\theta_t), \frac{d\theta_t}{dt} \right\rangle \\ &= -\varepsilon_t \langle \nabla\mathcal{L}(\theta_t), \mathbb{E}_{\rho,y}[\Gamma(\mu_{t,\rho}, \theta_t)] \rangle \\ &= -\varepsilon_t \langle \nabla\mathcal{L}(\theta_t), \mathbb{E}_{\rho,y}[\Gamma(\mu_{\theta_t,\rho}^*, \theta_t)] \rangle + \varepsilon_t \langle \nabla\mathcal{L}(\theta_t), \mathbb{E}_{\rho,y}[\Gamma(\mu_{\theta_t,\rho}^*, \theta_t) - \Gamma(\mu_{t,\rho}, \theta_t)] \rangle \\ &= -\varepsilon_t \|\nabla\mathcal{L}(\theta_t)\|^2 + \varepsilon_t \mathbb{E}_{\rho,y}[\langle \nabla\mathcal{L}(\theta_t), \Gamma(\mu_{\theta_t,\rho}^*, \theta_t) - \Gamma(\mu_{t,\rho}, \theta_t) \rangle]. \end{aligned} \quad (16)$$

Then, applying the Cauchy-Schwartz inequality, and the inequality $ab \leq \frac{1}{2}(a^2 + b^2) \iff ab - a^2 \leq \frac{1}{2}b^2 - \frac{1}{2}a^2$ with $a = \|\nabla\mathcal{L}(\theta_t)\|$ and $b = \mathbb{E}_{\rho,y}[\|\Gamma(\mu_{\theta_t,\rho}^*, \theta_t) - \Gamma(\mu_{t,\rho}, \theta_t)\|]$, we get

$$\begin{aligned} \frac{d}{dt}\mathcal{L}(\theta_t) &\leq -\varepsilon_t \|\nabla\mathcal{L}(\theta_t)\|^2 + \varepsilon_t \|\nabla\mathcal{L}(\theta_t)\| \mathbb{E}_{\rho,y}[\|\Gamma(\mu_{\theta_t,\rho}^*, \theta_t) - \Gamma(\mu_{t,\rho}, \theta_t)\|] \\ &\leq -\frac{\varepsilon_t}{2} \|\nabla\mathcal{L}(\theta_t)\|^2 + \frac{\varepsilon_t}{2} \mathbb{E}_{\rho,y}[\|\Gamma(\mu_{\theta_t,\rho}^*, \theta_t) - \Gamma(\mu_{t,\rho}, \theta_t)\|]^2 \\ &\leq -\frac{\varepsilon_t}{2} \|\nabla\mathcal{L}(\theta_t)\|^2 + \frac{\varepsilon_t}{2} \mathbb{E}_{\rho,y}[\|\Gamma(\mu_{\theta_t,\rho}^*, \theta_t) - \Gamma(\mu_{t,\rho}, \theta_t)\|^2], \end{aligned}$$

where we applied the Jensen inequality in the last line.

Then, since $t \mapsto \varepsilon_t$ is non-increasing (and thus $\varepsilon_T \leq \varepsilon_t$), integrating from $t = 0$ to $t = T$, we get

$$\begin{aligned} \frac{\varepsilon_T}{2} \int_0^T \|\nabla\mathcal{L}(\theta_t)\|^2 dt &\leq \int_0^T \frac{\varepsilon_t}{2} \|\nabla\mathcal{L}(\theta_t)\|^2 dt \\ &\leq \mathcal{L}(\theta_0) - \mathcal{L}(\theta_T) + \int_0^T \frac{\varepsilon_t}{2} \mathbb{E}_{\rho,y}[\|\Gamma(\mu_{\theta_t,\rho}^*, \theta_t) - \Gamma(\mu_{t,\rho}, \theta_t)\|^2] dt. \end{aligned}$$

And we finally obtain:

$$\frac{1}{T} \int_0^T \|\nabla\mathcal{L}(\theta_t)\|^2 dt \leq \frac{2}{\varepsilon_T T} (\mathcal{L}(\theta_0) - \inf_{\theta} \mathcal{L}(\theta)) + \frac{1}{\varepsilon_T T} \int_0^T \varepsilon_t \mathbb{E}_{\rho,y}[\|\Gamma(\mu_{\theta_t,\rho}^*, \theta_t) - \Gamma(\mu_{t,\rho}, \theta_t)\|^2] dt. \quad (17)$$

Using Assumption B.9, $\|\Gamma(\mu_{\theta_t,\rho}^*, \theta_t) - \Gamma(\mu_{t,\rho}, \theta_t)\|^2 \leq K_{\Gamma}^2 D(\mu_{t,\rho}, \mu_{\theta_t,\rho}^*)^2$. Let us note $\mathcal{F}_1^{\nu}(\mu) = D(\mu, \nu)^2$ and $\mathcal{F}_2^{\nu}(\mu) = D(\nu, \mu)^2$, and therefore let us bound $d(t) = \mathcal{F}_1^{\mu_{\theta_t,\rho}^*}(\mu_{t,\rho})$.

By differentiating, we obtain

$$d'(t) = \frac{d}{dt} D(\mu_{t,\rho}, \mu_{\theta_t,\rho}^*) = \underbrace{\int \frac{\delta \mathcal{F}_1^{\mu_{\theta_t,\rho}^*}}{\delta \mu}(\mu_t) \frac{\partial \mu_t}{\partial t}}_{(1)} + \underbrace{\int \frac{\delta \mathcal{F}_2^{\mu_{t,\rho}}}{\delta \mu}(\mu_{\theta_t,\rho}^*) \frac{\partial \mu_{\theta_t,\rho}^*}{\partial t}}_{(2)}.$$

On one hand, using that $\frac{\partial \mu_{t,\rho}}{\partial t} = \text{div}(\mu_{t,\rho} \nabla_W \mathcal{G}_{\theta_t,\rho}(\mu_{t,\rho}))$, and doing an integration by parts, we get

$$\begin{aligned} (1) &= \int \frac{\delta \mathcal{F}_1^{\mu_{\theta_t,\rho}^*}}{\delta \mu}(\mu_{t,\rho}) \text{div}(\mu_{t,\rho} \nabla_W \mathcal{G}_{\theta_t,\rho}(\mu_{t,\rho})) \\ &= - \int \left\langle \nabla_W \mathcal{F}_1^{\mu_{\theta_t,\rho}^*}(\mu_{t,\rho}), \nabla_W \mathcal{G}_{\theta_t,\rho}(\mu_t) \right\rangle d\mu_{t,\rho} \\ &= - \|\nabla_W \mathcal{F}_1^{\mu_{\theta_t,\rho}^*}(\mu_{t,\rho})\|_{L^2(\mu_{t,\rho})}^2 + \langle \nabla_W \mathcal{F}_1^{\mu_{\theta_t,\rho}^*}(\mu_{t,\rho}), \nabla_W \mathcal{F}_1^{\mu_{\theta_t,\rho}^*}(\mu_{t,\rho}) - \nabla_W \mathcal{G}_{\theta_t,\rho}(\mu_{t,\rho}) \rangle_{L^2(\mu_{t,\rho})}. \end{aligned}$$

Applying the inequality $\langle a, b \rangle \leq \|a\|^2 + \frac{1}{4}\|b\|^2$ (see *e.g.* [Vempala and Wibisono, 2019, Equation (33)]) and Assumptions B.10 and B.11, we get

$$\begin{aligned} (1) &\leq -\frac{3}{4} \|\nabla_W \mathcal{F}_1^{\mu_{\theta_t,\rho}^*}(\mu_{t,\rho})\|_{L^2(\mu_{t,\rho})}^2 + \|\nabla_W \mathcal{F}_1^{\mu_{\theta_t,\rho}^*}(\mu_{t,\rho}) - \nabla_W \mathcal{G}_{\theta_t,\rho}(\mu_{t,\rho})\|_{L^2(\mu_{t,\rho})}^2 \\ &\leq -\frac{3}{4C_{PL}} \mathcal{F}_1^{\mu_{\theta_t,\rho}^*}(\mu_{t,\rho}) + M\varepsilon_t \\ &= -\frac{3}{4C_{PL}} d(t) + M\varepsilon_t. \end{aligned}$$

On the other hand, for (2), we have

$$\begin{aligned} \frac{\partial \mu_{\theta_t, \rho}^*}{\partial t} &= \left\langle \frac{\partial \mu_{\theta_t, \rho}^*}{\partial \theta}, \frac{d\theta_t}{dt} \right\rangle \\ &= \left\langle \nabla_2 V_\rho(\cdot, \theta_t) - \mathbb{E}_{\mu_{\theta_t, \rho}^*} [\nabla_2 V_\rho(Z, \theta_t)], -\varepsilon_t \mathbb{E}_{\rho, y} [\Gamma(\mu_{t, \rho}, \theta_t)] \right\rangle \mu_{\theta_t, \rho}^*. \end{aligned}$$

Thus, using that $\|\nabla_2 V_\rho(x, \theta)\| \leq C_V$ by Assumption B.8, $\|\Gamma(\mu, \nu)\| \leq C_\Gamma$ by Assumption B.9 and $\|\frac{\delta \mathcal{F}_2^{\mu_{t, \rho}}}{\delta \mu}(\mu_{\theta_t, \rho}^*)\|_{L^2(\mu_{\theta_t, \rho}^*)} \leq C_{\mathcal{F}}$ by Assumption B.10, we get

$$\begin{aligned} (2) &= -\varepsilon_t \left\langle \int \frac{\delta \mathcal{F}_2}{\delta \mu}(\mu_{\theta_t, \rho}^*)(x) \cdot (\nabla_2 V_\rho(x, \theta_t) - \mathbb{E}_{\mu_{\theta_t, \rho}^*} [\nabla_2 V_\rho(Z, \theta_t)]) \, d\mu_{\theta_t, \rho}^*(x), \mathbb{E}_{\rho, y} [\Gamma(\mu_{t, \rho}, \theta_t)] \right\rangle \\ &\leq \varepsilon_t \|\mathbb{E}_{\rho, y} [\Gamma(\mu_{t, \rho}, \theta_t)]\| \int \left| \frac{\delta \mathcal{F}_2}{\delta \mu}(\mu_{\theta_t, \rho}^*)(x) \right| \cdot \|\nabla_2 V_\rho(x, \theta_t) - \mathbb{E}_{\mu_{\theta_t, \rho}^*} [\nabla_2 V_\rho(Z, \theta_t)]\| \, d\mu_{\theta_t, \rho}^*(x) \\ &\leq 2\varepsilon_t C_\Gamma C_V \int \left| \frac{\delta \mathcal{F}_2}{\delta \mu}(\mu_{\theta_t, \rho}^*) \right| \, d\mu_{\theta_t, \rho}^* \\ &\leq 2\varepsilon_t C_\Gamma C_V C_{\mathcal{F}}. \end{aligned} \tag{18}$$

Putting everything together, we obtain

$$d'(t) \leq -\frac{3}{4C_{PL}} d(t) + (M + 2C_\Gamma C_V C_{\mathcal{F}}) \varepsilon_t.$$

Now, note $\tilde{C} = M + 2C_\Gamma C_V C_{\mathcal{F}}$, and let's apply the Grönwall lemma [Dragomir, 2003]:

$$d(t) \leq d(0) e^{-\frac{3}{4C_{PL}} t} + \tilde{C} \int_0^t \varepsilon_s e^{\frac{3}{4C_{PL}}(s-t)} \, ds.$$

Plugging it into (17), we get

$$\frac{1}{T} \int_0^T \|\nabla \mathcal{L}(\theta_t)\|^2 \, dt \leq \frac{2}{\varepsilon_T T} (\mathcal{L}(\theta_0) - \inf_\theta \mathcal{L}(\theta)) + \frac{K_\Gamma^2 \mathbb{E}_{\rho, y} [d(0)]}{\varepsilon_T T} \int_0^T \varepsilon_t e^{-\frac{3}{4C_{PL}} t} \, dt + \frac{K_\Gamma^2 \tilde{C}}{\varepsilon_T T} \int_0^T \varepsilon_t \int_0^t \varepsilon_s e^{\frac{3}{4C_{PL}}(s-t)} \, ds \, dt.$$

Recall that $\varepsilon_t = \min(1, t^{-\frac{1}{2}})$, thus $T\varepsilon_T = \sqrt{T}$. The two first terms converge, hence they lie in $O(\sqrt{T})$. For the third one, we use the same trick as in [Marion et al., 2024, Proof of Proposition 4.5]. Let $T_0 \geq 2$ such that $\frac{1}{2}T_0 \geq \frac{4C_{PL}}{3} \log(T_0)$. For $t \geq T_0$, let $\alpha(t) = t - \frac{4C_{PL}}{3} \log(t)$. Notice that $\alpha(t) = \frac{t}{2} + \frac{t}{2} - \frac{4C_{PL}}{3} \log(t) \geq \frac{t}{2}$ for all $t \geq T_0$. Then, observe that for $t \geq T_0$,

$$\begin{aligned} \int_0^t \varepsilon_s e^{\frac{3}{4C_{PL}}(s-t)} \, ds &= \int_0^{\alpha(t)} \varepsilon_s e^{\frac{3}{4C_{PL}}(s-t)} \, ds + \int_{\alpha(t)}^t \varepsilon_s e^{\frac{3}{4C_{PL}}(s-t)} \, ds \\ &\leq \varepsilon_0 e^{-\frac{3}{4C_{PL}}} \int_0^{\alpha(t)} e^{\frac{3}{4C_{PL}} s} \, ds + (t - \alpha(t)) \varepsilon_{\alpha(t)} \\ &\leq \frac{4C_{PL} \varepsilon_0}{3} e^{\frac{3}{4C_{PL}}(\alpha(t)-t)} + \frac{4C_{PL}}{3} \log(t) \varepsilon_{\alpha(t)} \\ &\leq \frac{4C_{PL} \varepsilon_0}{3t} + \frac{4C_{PL}}{3} \log(t) \varepsilon_{t/2}. \end{aligned}$$

Thus,

$$\begin{aligned} \int_0^T \varepsilon_t \int_0^t \varepsilon_s e^{\frac{3}{4C_{PL}}(s-t)} \, ds \, dt &\leq \int_0^{T_0} \varepsilon_t \varepsilon_0 T_0 \, dt + \int_{T_0}^T \frac{4C_{PL} \varepsilon_t \varepsilon_0}{3t} \, dt + \int_{T_0}^T \frac{4C_{PL} \varepsilon_t \varepsilon_{t/2} \log(t)}{3} \, dt \\ &\leq \varepsilon_0 T_0^2 + \frac{4C_{PL}}{3} (\log T - \log T_0) + \frac{4C_{PL}}{3} \log T \cdot \int_{T_0}^T \frac{2}{t} \, dt \\ &= \mathcal{O}((\log T)^2). \end{aligned}$$

Putting everything together, we obtain that there exists a constant $c > 0$ such that

$$\frac{1}{T} \int_0^T \|\nabla \mathcal{L}(\theta_t)\|^2 dt \leq c \frac{(\log T)^2}{\sqrt{T}}.$$

□

We note that the MMD that we use has only been proven to fulfill a PL inequality with a different power [Arbel et al., 2019]. We provide this result as it gives a first step towards deriving guarantees of the method, e.g. for other divergence which would satisfy these assumptions. We leave for future work deriving the PL inequality for our MMD, or deriving alternative divergences which satisfy these assumptions. In Section D.1, we show empirically that Assumptions B.10 and B.11, as well as Theorem B.12, hold true in practice. Note that a weaker version of Theorem B.12, which does not give a specific rate of convergence for the average gradient, holds true for any gradient flow:

Theorem B.13. *For a gradient flow $t \mapsto \theta_t$ on a non-negative loss $\mathcal{L}(\theta)$ with bounded gradients, it holds that*

$$\frac{1}{T} \int_0^T \|\nabla \mathcal{L}(\theta_t)\|^2 dt \xrightarrow{T \rightarrow \infty} 0.$$

Proof. Note that

$$\partial_t \mathcal{L}(\theta_t) = \langle \nabla_{\theta} \mathcal{L}(\theta_t), \partial_t \theta_t \rangle = \langle \nabla_{\theta} \mathcal{L}(\theta_t), -\nabla_{\theta} \mathcal{L}(\theta_t) \rangle = -\|\nabla_{\theta} \mathcal{L}(\theta_t)\|^2 \leq 0.$$

Since $\mathcal{L}(\theta_t) \geq 0$, this implies that $\|\nabla_{\theta} \mathcal{L}(\theta_t)\|^2 \rightarrow 0$ as $t \rightarrow \infty$. Now for any $\epsilon > 0$, we can find a $t_{\epsilon} > 0$ such that for any $t > t_{\epsilon}$, $\|\nabla_{\theta} \mathcal{L}(\theta_t)\|^2 \leq \epsilon$. This means that for all $T \geq t_{\epsilon}$,

$$\frac{1}{T} \int_0^T \|\nabla \mathcal{L}(\theta_t)\|^2 dt \leq \frac{1}{T} (M t_{\epsilon} + (T - t_{\epsilon}) \epsilon) \leq \frac{M t_{\epsilon}}{T} + \epsilon \xrightarrow{T \rightarrow \infty} \epsilon$$

for some constant $M > 0$, and as this holds for any $\epsilon > 0$, the claim follows. □

We now discuss in more details the MMD case.

MMD case. Let $D = \frac{1}{2} \text{MMD}^2$ with k a kernel satisfying $k(x, x) \leq 1$. We first notice that it is currently unknown whether the MMD satisfies the PL inequality from Assumption B.10. However, it was shown in [Arbel et al., 2019, Proposition 7], that under some conditions, $\mathcal{F}_{\mu_{\theta, \rho}^*}(\mu)^2 \leq \frac{C}{4} \|\nabla_W \mathcal{F}_{\mu_{\theta, \rho}^*}(\mu)\|_{L^2(\mu)}^2$.

For this, we first need to recall the definition of the weighted negative Sobolev distance, e.g. [Arbel et al., 2019, Definition 5].

Definition B.14. Let $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ and denote by $\|\cdot\|_{\dot{H}(\mu)}$ its corresponding weighted Sobolev semi-norm, defined for all differentiable $f \in L^2(\mu)$ as $\|f\|_{\dot{H}(\mu)}^2 = \int \|\nabla f(x)\|^2 d\mu(x)$. The weighted negative Sobolev distance between $p, q \in \mathcal{P}_2(\mathbb{R}^d)$ is defined as

$$\|p - q\|_{\dot{H}^{-1}(\mu)} = \sup_{f \in L^2(\mu), \|f\|_{\dot{H}(\mu)} \leq 1} \left| \int f dp - \int f dq \right|.$$

This allows us to derive the following lemma.

Lemma B.15 (PL Inequality). *Assume k is continuously differentiable with L -Lipschitz gradient, i.e. for all $x, y, x', y' \in \mathbb{R}^d$, $\|\nabla k(x, x') - \nabla k(y, y')\| \leq L(\|x - x'\| + \|y - y'\|)$. Then, if for all $t \geq 0$, $\|\mu_{\theta_t, \rho}^* - \mu_{t, \rho}\|_{\dot{H}^{-1}(\mu_t)} \leq C$,*

$$\mathcal{F}_{\mu_{\theta_t, \rho}^*}(\mu)^2 \leq \frac{C}{4} \|\nabla_W \mathcal{F}_{\mu_{\theta_t, \rho}^*}(\mu)\|_{L^2(\mu)}^2.$$

Proof. See [Arbel et al., 2019, Proof of Proposition 7]. □

We now discuss what changes in the MMD case compared to the setting developed above. First, we will also assume Assumptions B.8, B.9 and B.11. Assumption B.10 is replaced by the assumptions of Lemma B.15 to get the PL inequality for the MMD.

Under Assumption B.8, we recover (17), i.e.

$$\frac{1}{T} \int_0^T \|\nabla \mathcal{L}(\theta_t)\|^2 dt \leq \frac{2}{\varepsilon_T T} (\mathcal{L}(\theta_0) - \inf_{\theta} \mathcal{L}(\theta)) + \frac{1}{\varepsilon_T T} \int_0^T \varepsilon_t \mathbb{E}_{\rho, y} [\|\Gamma(\mu_{\theta_t, \rho}^*, \theta_t) - \Gamma(\mu_t, \rho, \theta_t)\|^2] dt.$$

Remember also that Γ can be written as

$$\Gamma(\mu, \theta) = \nabla_1 \ell(h(\mu), y) \cdot \text{Cov}_{\mu} \left(\frac{\delta h}{\delta \mu}(\mu)(Z), \nabla_2 V(Z, \theta) \right).$$

First, we need to assume $\|\Gamma(\mu, \theta) - \Gamma(\nu, \theta)\| \leq \frac{K_{\Gamma}}{2} \text{MMD}(\mu, \nu)$. Thus, we need to upper bound $d : t \mapsto \frac{1}{2} \text{MMD}^2(\mu_t, \mu_{\theta_t, \rho}^*)$.

Let $d(t) = \frac{1}{2} \text{MMD}^2(\mu_t, \mu_{\theta_t, \rho}^*)$. By differentiating, we also get

$$d'(t) = \frac{d}{dt} D(\mu_t, \mu_{\theta_t, \rho}^*) = \underbrace{\int \frac{\delta \mathcal{F}_1^{\mu_{\theta_t, \rho}^*}}{\delta \mu}(\mu_t) \frac{\partial \mu_t}{\partial t}}_{(1)} + \underbrace{\int \frac{\delta \mathcal{F}_2^{\mu_t, \rho}}{\delta \mu}(\mu_{\theta_t, \rho}^*) \frac{\partial \mu_{\theta_t, \rho}^*}{\partial t}}_{(2)}.$$

By following the same computations as in Theorem B.12, assuming that Assumption B.11 holds and using Lemma B.15, we have

$$(1) \leq -\frac{3}{C_{PL}} d(t)^2 + M \varepsilon_t. \quad (19)$$

For (2), we have by using (18),

$$\begin{aligned} (2) &= -\varepsilon_t \left\langle \int \frac{\delta \mathcal{F}_2}{\delta \mu}(\mu_{\theta_t, \rho}^*)(x) \cdot (\nabla_2 V_{\rho}(x, \theta_t) - \mathbb{E}_{\mu_{\theta_t, \rho}^*}[\nabla_2 V_{\rho}(Z, \theta)]) d\mu_{\theta_t, \rho}^*(x), \mathbb{E}_{\rho, y}[\Gamma(\mu_t, \rho, \theta_t)] \right\rangle \\ &\leq \varepsilon_t \|\mathbb{E}_{\rho, y}[\Gamma(\mu_t, \rho, \theta_t)]\| \int \left| \frac{\delta \mathcal{F}_2}{\delta \mu}(\mu_{\theta_t, \rho}^*)(x) \right| \cdot \|\nabla_2 V_{\rho}(x, \theta_t) - \mathbb{E}_{\mu_{\theta_t, \rho}^*}[\nabla_2 V_{\rho}(Z, \theta_t)]\| d\mu_{\theta_t, \rho}^*(x) \\ &\leq 2\varepsilon_t C_{\Gamma} C_V \int \left| \frac{\delta \mathcal{F}_2}{\delta \mu}(\mu_{\theta_t, \rho}^*) \right| d\mu_{\theta_t, \rho}^* \\ &= 2\varepsilon_t C_{\Gamma} C_V \int \left| \int k(x, x') d\mu_{\theta_t, \rho}^*(x') - \int k(x, y) d\mu_t(y) \right| d\mu_{\theta_t, \rho}^* \\ &\leq 2\varepsilon_t C_{\Gamma} C_V \cdot \frac{1}{2} \text{MMD}(\mu_{\theta_t, \rho}^*, \mu_t) \\ &= 2\varepsilon_t C_{\Gamma} C_V d(t) \\ &\leq 2\varepsilon_t C_{\Gamma} C_V (1 + d(t)^2), \end{aligned} \quad (20)$$

where we used that for $\mu_{\theta_t, \rho}$ -almost every x , $k(x, \cdot) \in \mathcal{H}$ and $\|k(x, \cdot)\|_{\mathcal{H}} = k(x, x)^{\frac{1}{2}} \leq 1$ and taking the supremum.

Putting everything together, we get,

$$d'(t) \leq \left(2C_{\Gamma} C_V - \frac{3}{C_{PL}} \right) d(t)^2 + (M + 2C_{\Gamma} C_V) \varepsilon_t. \quad (21)$$

Here, we observe that we need to assume $2C_{\Gamma} C_V - \frac{3}{C_{PL}} \leq 0$ for the rate to be converging.

To apply the Grönwall lemma to (21), we need to solve the following Riccati equation $y'(t) = -ay^2(t) + b\varepsilon(t)$ ⁷ for $a = \frac{3}{C_{PL}} - 2C_{\Gamma} C_V \geq 0$, $b = M + 2C_{\Gamma} C_V$, which however does not have a simple solution for $b \neq 0$ [Zaitsev and Polyagin, 2002].

⁷math.stackexchange.com/questions/4773818/generalized-gronwall-inequality-covering-many-different-applications

For $b = 0$, we would have

$$d(t) \leq \frac{d(0)}{1 + ad(0)t},$$

and

$$\begin{aligned} \int_0^T \varepsilon_t d(t) dt &\leq d(0) \int_0^T \frac{1}{1 + ad(0)t} dt \\ &= d(0) \cdot \frac{\log(1 + ad(0)T)}{ad(0)} \\ &= \mathcal{O}(\log T). \end{aligned}$$

Thus, the rate would be in $\mathcal{O}(\log T/\sqrt{T})$.

If we bound ε_t by 1, then we can show that

$$d(t) \leq \frac{\sqrt{ba}}{a} \cdot \frac{1}{\tanh(\sqrt{bat} + K)},$$

for $K = \operatorname{arctanh}\left(\frac{\sqrt{ba}}{ad(0)}\right)$. Observing that

$$\begin{aligned} \int_0^T \varepsilon_t \frac{1}{\tanh(\sqrt{bat} + K)} dt &\leq \int_0^T \frac{1}{\tanh(\sqrt{bat} + K)} dt \\ &= \frac{1}{\sqrt{ba}} (\log(\sinh(\sqrt{ba}T + K)) - \log(\sinh(K))) \\ &= \mathcal{O}(T). \end{aligned}$$

However, this bound is too big, and will diverge.

B.4 EI Property

Proposition 5. *Let $F^{\text{bil}} : \mathbb{R}^{N \times p} \times \mathbb{R}^{M \times d} \rightarrow \mathbb{R}^{N \times p}$, $(\mathbf{Z}, \mathbf{X}) \mapsto F^{\text{bil}}(\mathbf{Z}, \mathbf{X})$. Then F^{bil} is a bilinear map that fulfills the EI property if and only if F^{bil} is of the form*

$$F^{\text{bil}}(\mathbf{Z}, \mathbf{X}) = \mathbf{Z} \alpha \mathbf{X}^\top \mathbf{1}_M + \mathbf{1}_N \mathbf{1}_N^\top \mathbf{Z} \beta \mathbf{X}^\top \mathbf{1}_M \quad (22)$$

for some tensors $\alpha, \beta \in \mathbb{R}^{p \times p \times d}$, where $\mathbf{1}_N, \mathbf{1}_M$ denote the vectors with ones in \mathbb{R}^N and \mathbb{R}^M .

Proof. F^{bil} is bilinear if it can be written as

$$F^{\text{bil}}(Z, X)_{ij} = \sum_{klmn} F_{ijklmn}^{\text{bil}} Z_{kl} X_{mn}$$

for some $F_{ijklmn}^{\text{bil}} \in \mathbb{R}$. By definition, F^{bil} is equivariant in Z and invariant in X if and only if for any σ_N, σ_M , it holds that

$$F^{\text{bil}}(Z, X)_{\sigma_N(i)j} = \sum_{klmn} F_{ijklmn}^{\text{bil}} Z_{\sigma_N(k)l} X_{\sigma_M(m)n},$$

which is equivalent to

$$F^{\text{bil}}(Z, X)_{ij} = \sum_{klmn} F_{\sigma_N^{-1}(i)j\sigma_N^{-1}(k)l\sigma_M^{-1}(m)n}^{\text{bil}} Z_{kl} X_{mn}.$$

Hence, F^{bil} has the desired properties if and only if for any σ_N and σ_M , it holds that for all indices i, j, k, l, m, n ,

$$F_{ijklmn}^{\text{bil}} = F_{\sigma_N(i)j\sigma_N(k)l\sigma_M(m)n}^{\text{bil}}.$$

(Note we can replace the inverse permutations by permutations since the set of all inverse permutations is the same as the set of all permutations.) This is the case if and only if the following two conditions hold:

- $F_{ijklmn}^{\text{bil}} = F_{\sigma_N(i)j\sigma_N(k)l\sigma_M(m)n}^{\text{bil}}$. For fixed i, k , consider two cases:

- if $i \neq k$, then for all $i' \neq k'$, we can find a σ_N s.t. $\sigma_N(i) = i'$ and $\sigma_N(k) = k'$, which implies $F_{ijklmn}^{\text{bil}} = F_{i'jk'lmn}^{\text{bil}}$
- if $i = k$, then $\sigma_N(i) = \sigma_N(k)$ for any σ_N , hence for all i' , $F_{ijilmn}^{\text{bil}} = F_{i'jilmn}^{\text{bil}}$ (since we can always find σ_N such that $\sigma_N(i) = i'$).

This means we can write $F_{ijklmn}^{\text{bil}} = \tilde{\alpha}_{l_jmn} \delta_{ik} + \tilde{\beta}_{l_jmn} (1 - \delta_{ik})$ for some $\tilde{\alpha}$ and $\tilde{\beta}$ in $\mathbb{R}^{p \times p \times d}$, where δ_{ik} is the Kronecker delta;

- $F_{ijklmn}^{\text{bil}} = F_{ijkl\sigma_M(m)n}^{\text{bil}}$, which is equivalent to $F_{ijklmn}^{\text{bil}} = F_{ijklm'n}^{\text{bil}}$ for all m, m' . This means we can write $F_{ijklmn}^{\text{bil}} = F_{ijkln}^{\text{bil}}$.

Putting these two things together, this means F^{bil} has the desired properties if and only if it can be written as

$$F_{ijklmn}^{\text{bil}} = \tilde{\alpha}_{l_jn} \delta_{ik} + \tilde{\beta}_{l_jn} (1 - \delta_{ik})$$

for some $\tilde{\alpha}$ and $\tilde{\beta}$ in $\mathbb{R}^{p \times p \times d}$. If we let $\alpha := \tilde{\alpha} - \tilde{\beta}$ and $\beta := \tilde{\beta}$, this is equivalent to

$$F_{ijklmn}^{\text{bil}} = \alpha_{l_jn} \delta_{ik} + \beta_{l_jn}.$$

By writing out equation (22) with indices, one can verify that this is equivalent. \square

In order to proof that MultiHead attention layers fulfill the EI property, we prove two easy, but helpful lemmas.

Lemma B.16. *Let $F : \mathbb{R}^{N \times p} \times \mathbb{R}^{M \times d} \rightarrow \mathbb{R}^{N \times p}$, such that it can be written as*

$$F(\mathbf{Z}, \mathbf{X}) = F(\mathbf{Z}) = \begin{bmatrix} f(z_1) \\ f(z_2) \\ \vdots \\ f(z_N) \end{bmatrix}$$

for some function $f : \mathbb{R}^p \rightarrow \mathbb{R}^p$. Then F fulfills the EI property.

Proof. F is trivially invariant in \mathbf{X} , as it does not depend on \mathbf{X} . Furthermore, it is clear that F is equivariant in \mathbf{Z} , as the i^{th} entry of F only depends on z_i . \square

Lemma B.17. *The EI property is preserved under compositions of functions that each fulfill the EI property.*

Proof. This follows immediately from the fact that both row-equivariance as well as -invariance are preserved under compositions. \square

Proposition 6. *Denote by $\text{MultiHead}(tgt, src)$ multi-head attention [Vaswani et al., 2017] between a target and source sequence. Then $\text{MultiHead}(\mathbf{Z}, \mathbf{Z})$ and $\text{MultiHead}(\mathbf{Z}, \mathbf{X})$ both fulfill the EI property.*

Proof. Recall that single-head cross-attention between \mathbf{X} and \mathbf{Z} is defined as

$$\text{Att}(\mathbf{Z}, \mathbf{X})_i = \text{LayerNorm} \left(\frac{\sum_j e^{\tau \langle Qz_i, Kx_j \rangle} Vx_j}{\sum_j e^{\tau \langle Qz_i, Kx_j \rangle}} \right)$$

for some matrices $Q \in \mathbb{R}^{h \times p}$, $K \in \mathbb{R}^{h \times d}$, $V \in \mathbb{R}^{p \times d}$, where $h \in \mathbb{R}$ (typically, $h = p = d$), and some $\tau \in \mathbb{R}$; here, the z_i denote the rows of \mathbf{Z} , and the x_i denote the rows of \mathbf{X} . The term $\frac{\sum_j e^{\tau \langle Qz_i, Kx_j \rangle} Vx_j}{\sum_j e^{\tau \langle Qz_i, Kx_j \rangle}}$ only depends on z_i but none of the other z 's, and is clearly invariant under permutations of the x_i , which makes it fulfill the EI property. Layer normalization fulfills the EI property by lemma B.16, and thus by lemma B.17 single-head cross-attention does as well. Again by lemma B.17, the result holds for multi-head cross-attention as well. Similarly, we can show that multi-head self-attention on \mathbf{Z} fulfills the EI property, since it is trivially invariant under permutations of \mathbf{X} (as it does not depend on \mathbf{X}). \square

C Training Details

C.1 Architecture

In this section, we provide additional details about the architecture. The encoder blocks are stacked multi-head attention layers [Vaswani et al., 2017], as can be seen in Figure 6.

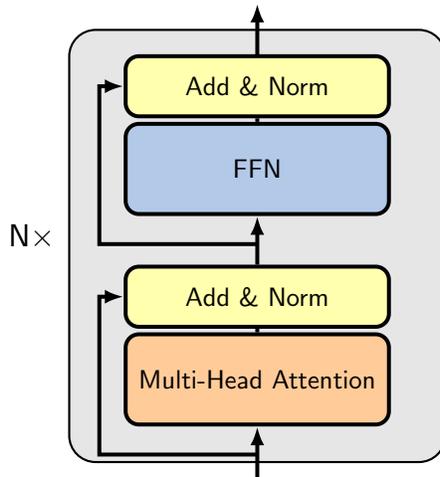


Figure 6: Our encoders are stacked multi-head attention layers [Vaswani et al., 2017].

Layer Sizes. The hidden dimensions of the first two feedforward networks (the ones before the bilinear layer) are the same as the dimension of the bilinear layer (16). The hidden dimensions of the feedforward networks right after the bilinear layer is the same as the encoder dimension (128). The feedforward network at the end is a typical transformer feedforward layer, where the hidden dimension is four times that of the encoder dimension (i.e. 512). We would like to highlight here that our head size (of four), which we tuned extensively, is extremely small, compared to what is common in the literature. Further investigation into what such a small head size means for the model and its training is outstanding.

Invertible Network. The invertible network used in the point cloud completion pipeline takes the form described in the paper, i.e. for the padded $\tilde{\mathbf{Z}}$ (which is \mathbf{X} with the free particles, padded with zeros to match the encoder dimension):

$$\begin{aligned}\mathbf{Z} = q(\tilde{\mathbf{Z}}) &= [\tilde{\mathbf{Z}}_1, \tilde{\mathbf{Z}}_2 e^{\phi(\tilde{\mathbf{Z}}_1)} + \psi(\tilde{\mathbf{Z}}_1)] \\ q^{-1}(\mathbf{Z}) &= [\mathbf{Z}_1, (\mathbf{Z}_2 - \psi(\mathbf{Z}_1))e^{-\phi(\mathbf{Z}_1)}],\end{aligned}$$

where ϕ and ψ are two-layer neural networks, and $\tilde{\mathbf{Z}} = [\tilde{\mathbf{Z}}_1, \tilde{\mathbf{Z}}_2]$ is a partition of \mathbf{X} into two equal-sized chunks along the particle dimension. Here, both the input and output dimensions, as well as the hidden dimensions of ϕ and ψ are equal to 64, i.e. half the encoder dimension. It can be readily seen that these layers are indeed invertible, as we have

$$q^{-1}(q(\mathbf{Z})) = [\mathbf{Z}_1, (\mathbf{Z}_2 e^{\phi(\mathbf{Z}_1)} + \psi(\mathbf{Z}_1) - \psi(\mathbf{Z}_1))e^{-\phi(\mathbf{Z}_1)}] = [\mathbf{Z}_1, \mathbf{Z}_2] = \mathbf{Z}.$$

We have tried various different methods for the invertible layer. This includes other invertible neural networks, but also non-parametric methods, such as an orthogonal matrix, where for the inverse we use the pseudo-inverse; using the real Fourier transform the map from a lower- to a higher-dimensional space; and also just using a regular, non-invertible, trainable linear layer, and simply setting the output particles equal to the input particles manually once a fixed point has been found. Except for the latter, all of these worked comparably well, but the one we opted for seemed to have a slight edge in practice. Since all methods that are provably invertible are equivalent on the input points, as they will faithfully recover the input points, the only difference in performance can stem from how they process points that are not contained in the partial input point cloud. Possibly, the architecture we opted for is naturally well suited for this type of task; this is supported by the fact

that it originated in the invertible normalizing flow literature [Dinh et al., 2014], which also deals with mapping distributions to distributions.

Masking. Since different samples have different numbers of particles, our batches are padded with zeros. To account for this, all intermediate states of the network are masked accordingly. For point cloud completion, we fix the particles that correspond to the input. We do this by setting the outputs of all encoders at the respective positions equal to the inputs to the network at those positions, and also by masking out the gradient in the MMD flow with zeros for the fixed particles. This ensures that the particles remain unchanged, and are converted to the original low-dimensional particles by the inverse invertible layer.

C.2 Training

In this section, we provide some more details regarding the training procedure.

Approximating the Wasserstein Gradient. In practice, we are not directly using the formula for the Wasserstein gradient we derived, but instead rely on autodifferentiation in PyTorch. This is both for simplicity, as well as to account for the fact that, as discussed in Section D.3, our network architecture cannot strictly be written as a pushforward operator on the particles, hence the Wasserstein gradient formula is not exact⁸. To account for the fact that in the discrete measure corresponding to \mathbf{Z} , each particle is assigned a mass of $1/N_{\mathbf{Z}}$, where $N_{\mathbf{Z}}$ denotes the number of particles in \mathbf{Z} , we rescale the gradient obtained from autodifferentiation by the number of particles in \mathbf{Z} , similar to [Chizat and Bach, 2018, Definition 2.2].

Phantom Gradient. As described in the paper, we rely on the phantom gradient to compute the gradients in the backward pass of the DDEQ. For a definition of the phantom gradient, we refer the reader to the original paper [Geng et al., 2021]. We use the DEQSliced implementation from the torchdeq library [Geng and Kolter, 2023], and use just a single gradient step for the phantom gradient, as well as a single state (where increasing the number of states corresponds to sampling multiple fixed points from the forward solver, and using the best one).

D Additional Experiments

D.1 Empirical Verification of Theorem B.12

In this section, we empirically verify that Theorem B.12, as well as Assumptions B.10 and B.11 hold true in practice across all four datasets.

Assumption B.10. We plot

$$\frac{\mathcal{F}_{\mu_{\hat{\theta}_t, \rho}^*}(\mu)}{\left\| \nabla_W \mathcal{F}_{\mu_{\hat{\theta}_t, \rho}^*}(\mu) \right\|_{L^2(\mu)}^2}$$

averaged across 8 randomly created $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ over the course of training. The plots can be seen in Figure 7. We see that across all four datasets, the above quantity is bounded, which verifies Assumption B.10.

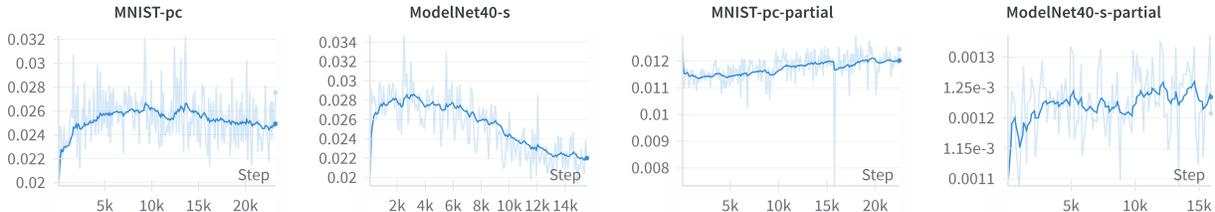


Figure 7: Empirical Verification of Assumption B.10 across all four datasets, with a smoothed EMA.

⁸We found autodifferentiation to have slightly better performance than using the formula for the Wasserstein gradient, which can likely be accounted for by the fact that our network is only “almost” a pushforward operator.

Assumption B.11. We plot

$$\frac{\left\| \nabla_W \mathcal{G}_{\theta_t, \rho}(\mu_t) - \nabla_W \mathcal{F}_{\mu_{\theta_t, \rho}^*}(\mu_t) \right\|_{L^2(\mu_t)}^2}{\epsilon_t},$$

where we set $\epsilon_t = \min(1, t^{-1/2})$, as in Theorem B.12, and choose $\mu_t = \mu_{\theta_t, \rho}^*$ for simplicity; the plots can be seen in Figure 8. Again, we have boundedness across all four datasets.

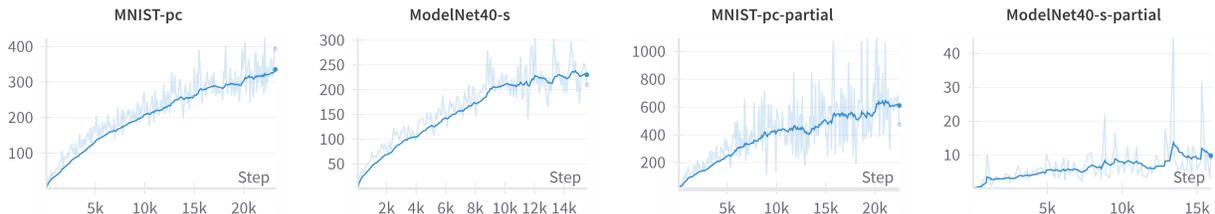


Figure 8: Empirical Verification of Assumption B.11 across all four datasets, with a smoothed EMA.

Theorem B.12.

We plot

$$\frac{\int_0^T \|\nabla \mathcal{L}(\theta_t)\|^2 dt}{\sqrt{T}(\log T)^2}$$

over the course of training across all three datasets in Figure 9; here, integration is discrete over the training steps. We see that across datasets, the above quantity is bounded, which verifies Theorem B.12 empirically.

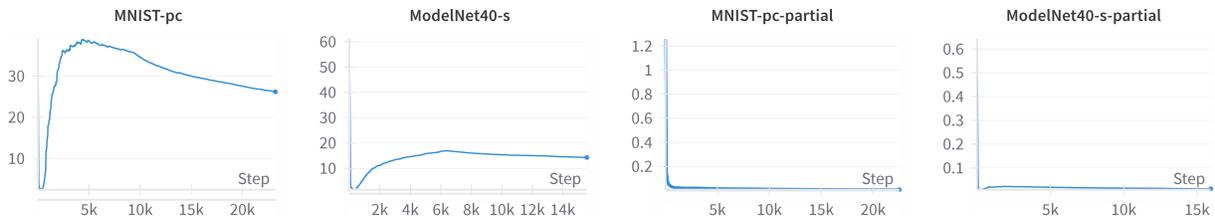


Figure 9: Empirical Verification of Theorem B.12 across all four datasets.

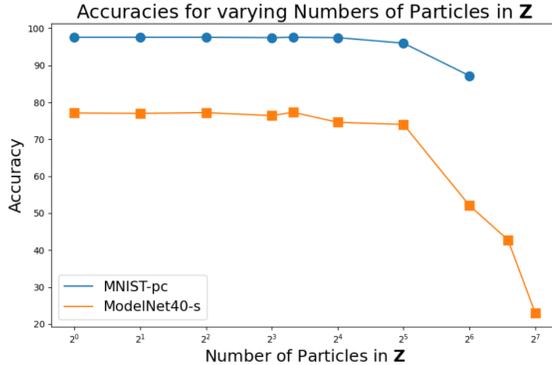
D.2 Ablation Study on Number of Particles in \mathbf{Z}

We provide an ablation study on the number of particles in \mathbf{Z} . While for point cloud completion, this number is dictated by the number of particles in the input/target, for point cloud classification, this is a free hyperparameter which can be tuned. In Figure 10, we see the average accuracy of a DDEQ for the two classification datasets ModelNet40-s and MNIST-pc over varying numbers of particles in \mathbf{Z} .

We can see that for both datasets, the accuracy tends to be very similar for between one and 16 particles in \mathbf{Z} (with a very slight increase in accuracy in the range of four to ten particles), and then quickly drops when increasing the number of particles even further.

D.3 Ablation Study on Pushforward Architecture

In the derivation of the Wasserstein gradient of $\mathcal{G}(\mu) = \frac{1}{2} \text{MMD}^2(\mu, F_{\theta\#}\mu)$, as well as in the proof of Theorem B.12, we assumed that we can write $F_{\theta}(\mu) = F_{\theta\#}(\mu)$ as a pushforward operator. This means there cannot be any interactions between particles in how F_{θ} processes them. However, in the architecture we use in practice, two parts do not comply with this assumption, namely the bilinear layer and the self-encoder on \mathbf{Z} (which corresponds to μ). In this section, we study how making our architecture a pushforward operator on \mathbf{Z} , namely by changing the bilinear layer from

Figure 10: Average accuracy on ModelNet40-s and MNIST-pc for varying numbers of particles in \mathbf{Z} .

$$F^{\text{bil}}(\mathbf{Z}, \mathbf{X}) = \mathbf{Z}\alpha\mathbf{X}^\top \mathbf{1}_M + \mathbf{1}_N \mathbf{1}_N^\top \mathbf{Z}\beta\mathbf{X}^\top \mathbf{1}_M$$

to

$$F^{\text{bil}}(\mathbf{Z}, \mathbf{X}) = \mathbf{Z}\alpha\mathbf{X}^\top \mathbf{1}_M,$$

i.e. removing the part that introduces interactions between particles in \mathbf{Z} , and by removing the self-encoder on \mathbf{Z} (while leaving the self-encoder on \mathbf{X} and the cross-encoder unchanged) affects the performance.

Table 3: Accuracies and test losses of DDEQs, for the pushforward architecture compared against the complete architecture we use in practice.

	MNIST-pc			ModelNet40-s			MNIST-pc-partial
	Acc.	Cross-Entr.	Loss	Acc.	Cross-Entr.	Loss	MMD Loss
pushforward	97.7	0.101		75.8	0.938		0.0034
full architecture	98.1	0.086		77.3	0.888		0.0033

In Table 3, we can see the differences in accuracy and test loss between the complete architecture, and the one that can be written as a pushforward. We see that using the architecture that’s strictly a pushforward is overall comparable to the full architecture, but slightly worse in all metrics.

D.4 Computational Complexity

In this section, we address the computational complexity of DDEQs, and how it changes w.r.t. the number of forward iterations in the inner loop, as well as the number of points in the input measures. The results are for point cloud completion on *MNIST-pc-partial* trained on a single H100 GPU. When reducing the number of forward pass iterations of the inner loop, we increase the inner loop learning rate accordingly, such that $\text{learning rate} \times \text{forward iterations}$ remains constant.

Number of Iterations	10	50	100	200	500	1000
Runtime (1 Epoch)	13m	47m	1h32m	3h1m	8h12m	14h53m

Table 4: Runtime per epoch for different numbers of iterations.

In contrast, the runtime for PCN is 4min. This shows that the number of forward iterations is crucial in controlling the runtime of DDEQs. It can be reduced to around 50 iterations without significant loss in performance, but deriving methods to improve accuracy for fewer iterations is an important line of future research, as well as deriving other methods to speed up training, such as improving the optimizer, possibly with momentum.

Input Particles (%)	20%	40%	60%	80%	100%
Runtime (1 Epoch)	1h16m	1h29m	1h51m	2h20m	3h1m

Table 5: Runtime per epoch for different input particle percentages.

Fixed point acceleration methods such as Anderson acceleration that exist in a Euclidean setting could be derived for WGFs; another approach would be to use the sliced MMD [Hertrich et al., 2024], which reduces the computational complexity of computing (the gradient of) the MMD from $O(n^2)$ to $O(n \log n)$ (where n is the number of particles).

D.5 Comparing DDEQs and DEQs

To demonstrate that DDEQs outperform standard DEQs on point cloud completion, we trained classical DEQs with the same architecture, but with (1) Broyden, (2) Anderson, and (3) fixed point iteration forward solvers, three widely used forward solvers for DEQs. All three perform similarly and achieve a test loss, measured as the squared MMD between prediction and target, of $13\times$ that of DDEQs, and an average Wasserstein-2 loss of $4\times$ that of DDEQs, with very poor visual image quality. However, we want to mention that DEQs achieve (almost) comparable performance on point cloud classification, highlighting the fact that while DDEQs are capable of point cloud classification, this is not a task that necessitates them, whereas point cloud completion clearly is.

D.6 Fixed Points

In this section, we show what the fixed points for the classification tasks look like, and how close they are to being real fixed points. In Figures 11 and 12, we see the fixed points \mathbf{Z}^* obtained for both datasets, alongside $F_\theta(\mathbf{Z}^*)$ – for perfect fixed points, these two should be identical. Since \mathbf{Z} is high dimensional, we project it to two dimensions with UMAP [McInnes et al., 2018], both with two, as well as 30 neighbors. Clearly, with both projection methods \mathbf{Z}^* and $F_\theta(\mathbf{Z}^*)$ tend to be visibly fairly close, but noticeable differences remain.

D.7 Class Embeddings

In our classification models, the classification head is a single linear layer which gets as input only the MaxPooling of the fixed points. Thus, the DDEQ needs to embed the class information in the MaxPooling of the fixed points. In Figures 13 and 14, we see the output of the MaxPooling for the samples of three batches of the test split, respectively.

As can be seen, the embeddings for MNIST-pc classes are very distinct, and the embeddings for ModelNet40-s are distinct for most classes, but for some, there is significant overlap. Note that these are visualizations of 128-dimensional points in two dimensions, and suffer from the limitations of the representation techniques used.

D.8 Additional ModelNet40-s-partial Samples

Figure 15 contains additional samples for point cloud completion on ModelNet40-s-partial with DDEQ and PCN. We see from the samples that DDEQ excels at certain classes, such as airplane, but other classes, such as bowls or cones, can be failure modes. PCN tends to do better on airplanes as well, but generally produces point clouds that are too diffuse. In Table 6, we report the average OT distances (with Euclidean distance as cost function) between the predictions and ground truth for DDEQ and PCN. For DDEQ, we compute both the distance between the complete prediction and ground truth, as well as only for the "free particles" and the corresponding particles in the target.

	DDEQ		PCN
	free particles	full point cloud	
OT Distance	1.27 ± 0.04	0.154 ± 0.002	0.61 ± 0.02

Table 6: OT distances between predictions and targets for DDEQ and PCN on ModelNet40-s-partial.

D.9 Visualization of MMD Gradient Flow With Rotational Target

In the literature, MMD gradient flows are typically studied with a fixed target measure, i.e. gradient flows of $\mu \mapsto \text{MMD}^2(\mu, \nu)$ for a fixed measure ν . To our knowledge, gradient flows of functionals of the form $\mu \mapsto \text{MMD}^2(\mu, F(\mu))$, where F is a functional, have not been studied yet. To illustrate what interesting emergent properties these gradient flows can have even for very simple target functions, we show what the gradient flow looks like when F is a rotation of $2\pi/8$ resp. $2\pi/5$ around the origin, in Figures 16 and 17. We initialize $x \sim \mathcal{N}\left(\begin{bmatrix} -10 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 12 \end{bmatrix}\right)$, such that the initial particles are not centered around the origin, and again use the Riesz kernel, with $\eta = 10$ and $\gamma_\eta = 0.999$.

We see that instead of converging to a point cloud uniformly distributed around the origin, the MMD flow converges to diverse geometric patterns which are perfectly symmetric.

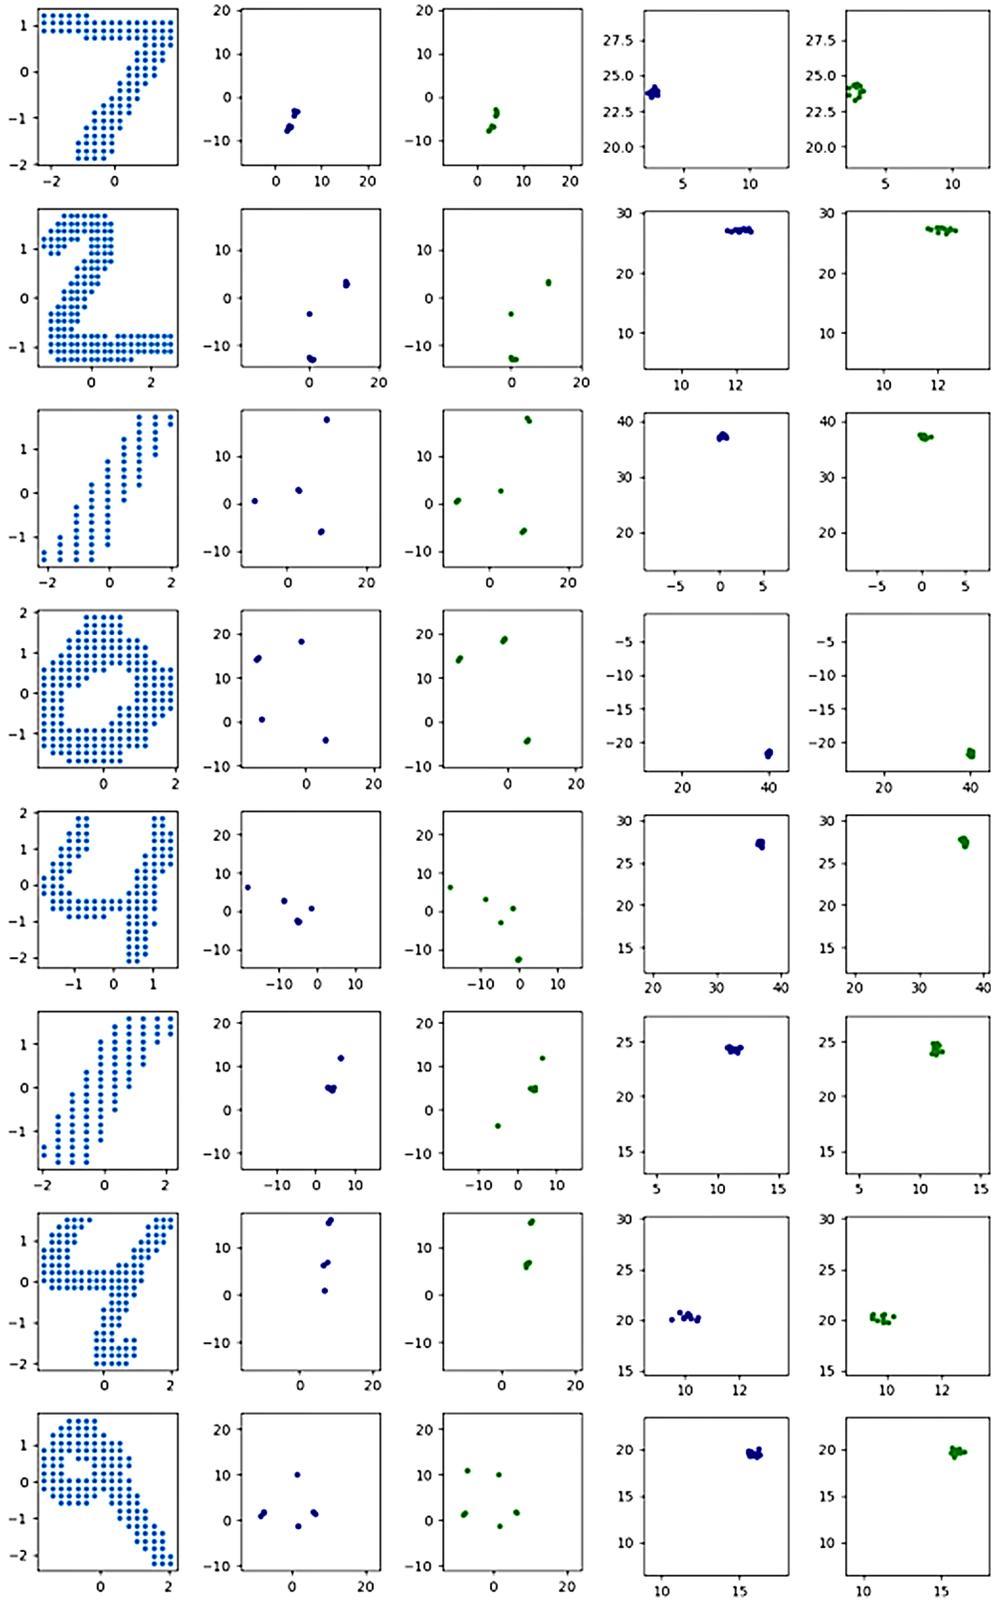


Figure 11: Fixed points for MNIST-10 classification. First column is inputs, second and third column are \mathbf{Z}^* resp. $F_\theta(\mathbf{Z}^*)$, mapped with UMAP with $n_neighbors=2$; and the fourth and fifth column are \mathbf{Z}^* resp. $F_\theta(\mathbf{Z}^*)$, mapped with UMAP with $n_neighbors=30$.

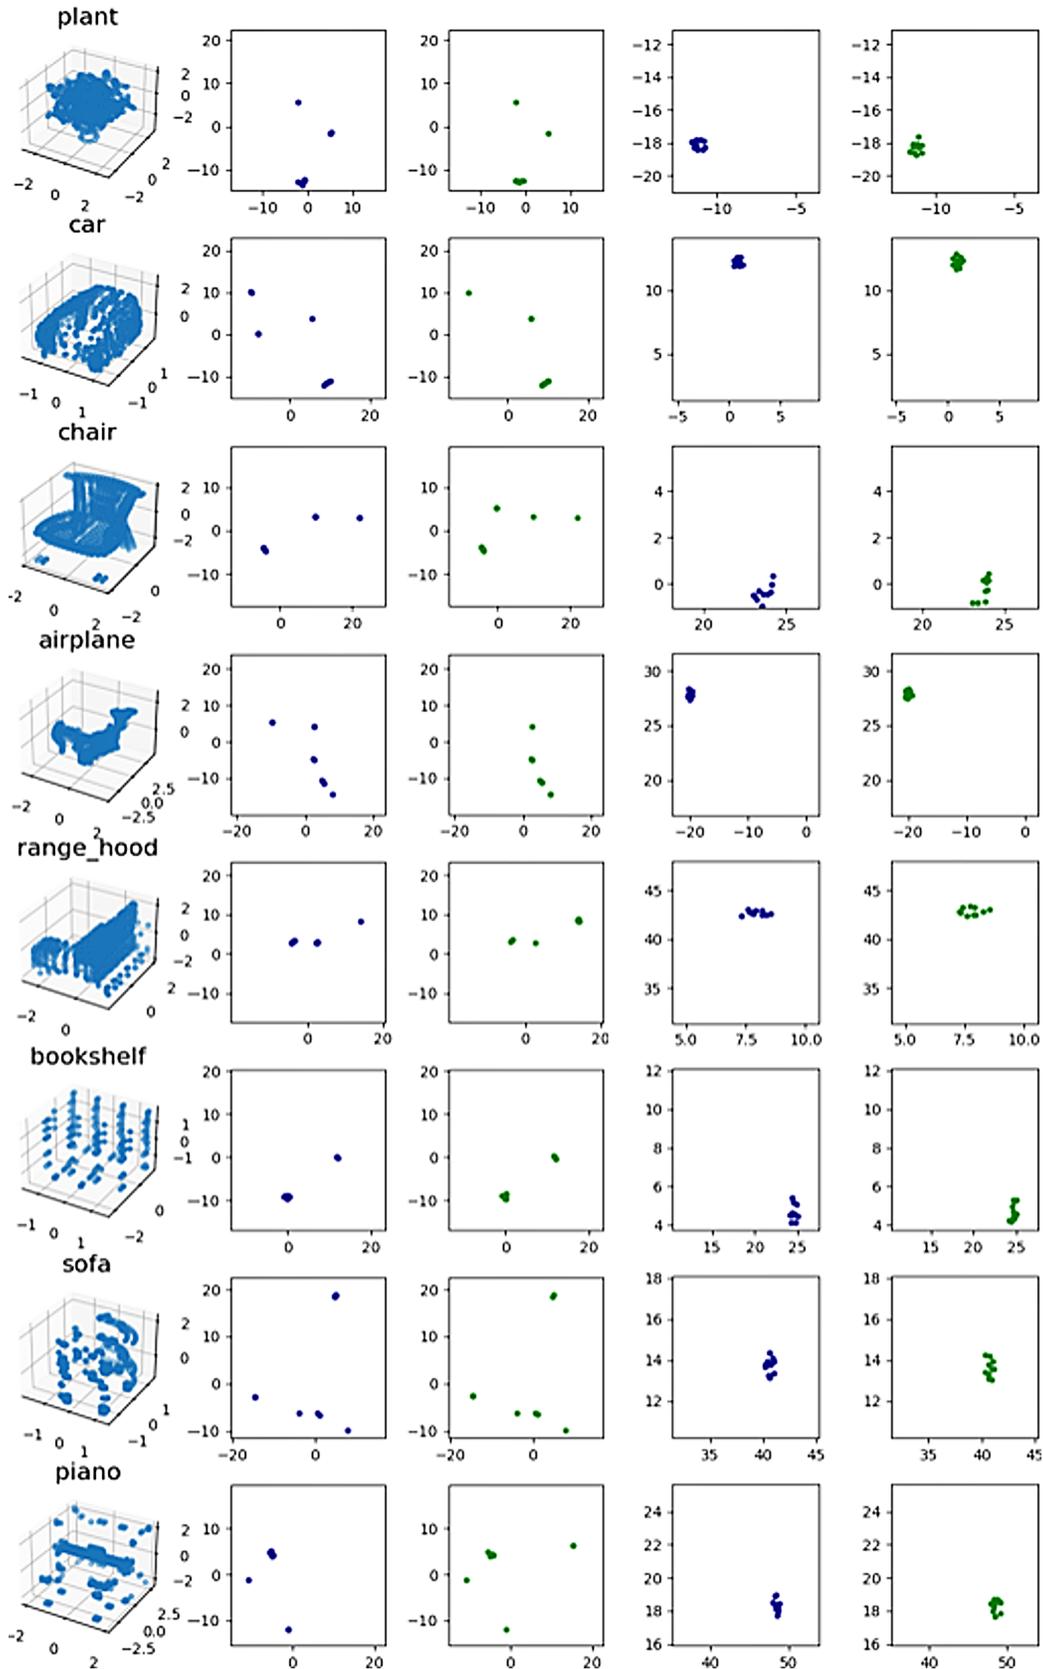


Figure 12: Fixed points for ModelNet40-s classification. First column is inputs, second and third column are \mathbf{Z}^* resp. $F_\theta(\mathbf{Z}^*)$, mapped with UMAP with `n_neighbors=2`; and the fourth and fifth column are \mathbf{Z}^* resp. $F_\theta(\mathbf{Z}^*)$, mapped with UMAP with `n_neighbors=30`.

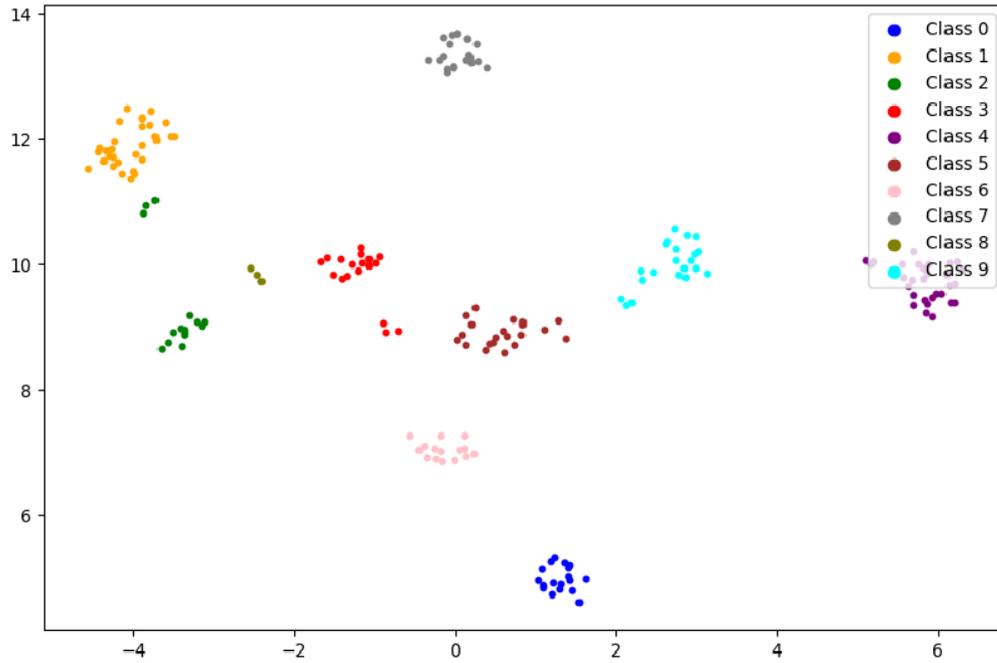


Figure 13: Distribution of the MaxPooling of the fixed points on MNIST-pc, visualized with UMAP with `n_neighbors=80`.

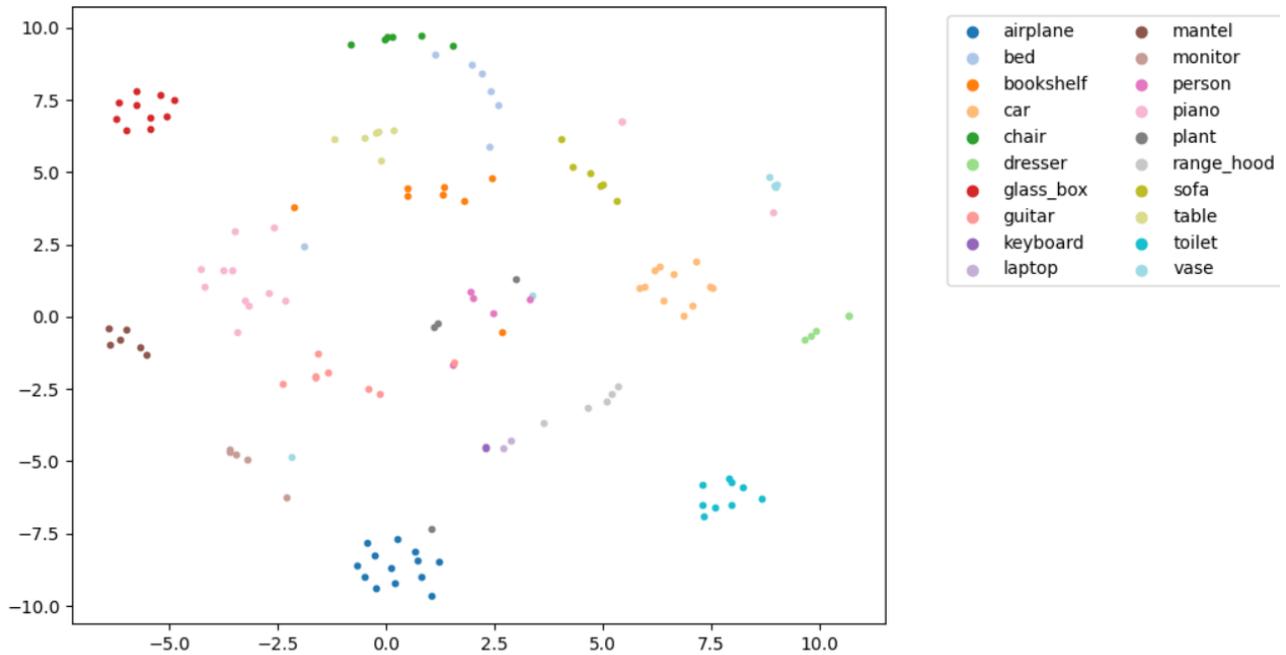


Figure 14: Distribution of the MaxPooling of the fixed points on ModelNet40-s, for a subset of 20 of the 40 classes. Visualized with TSNE [van der Maaten and Hinton, 2008] with `perplexity=30`.

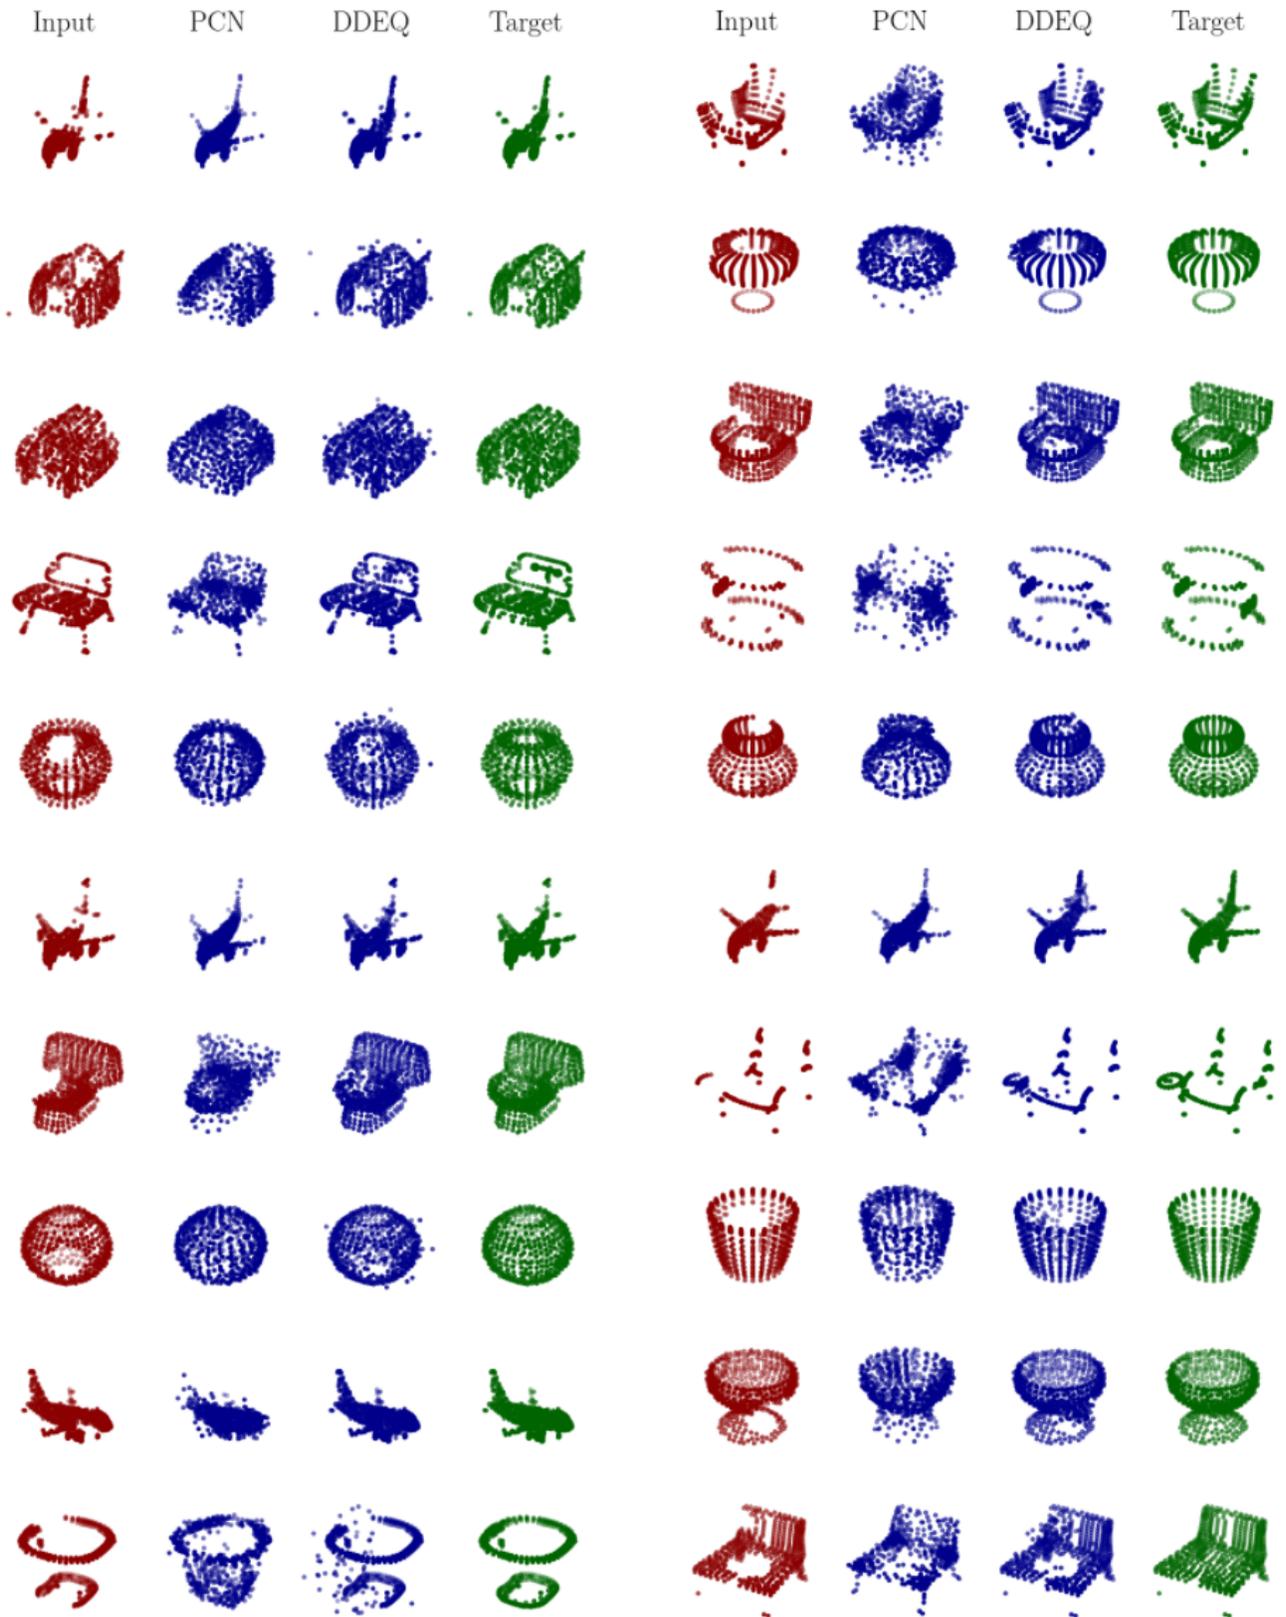


Figure 15: ModelNet40-s-partial samples.

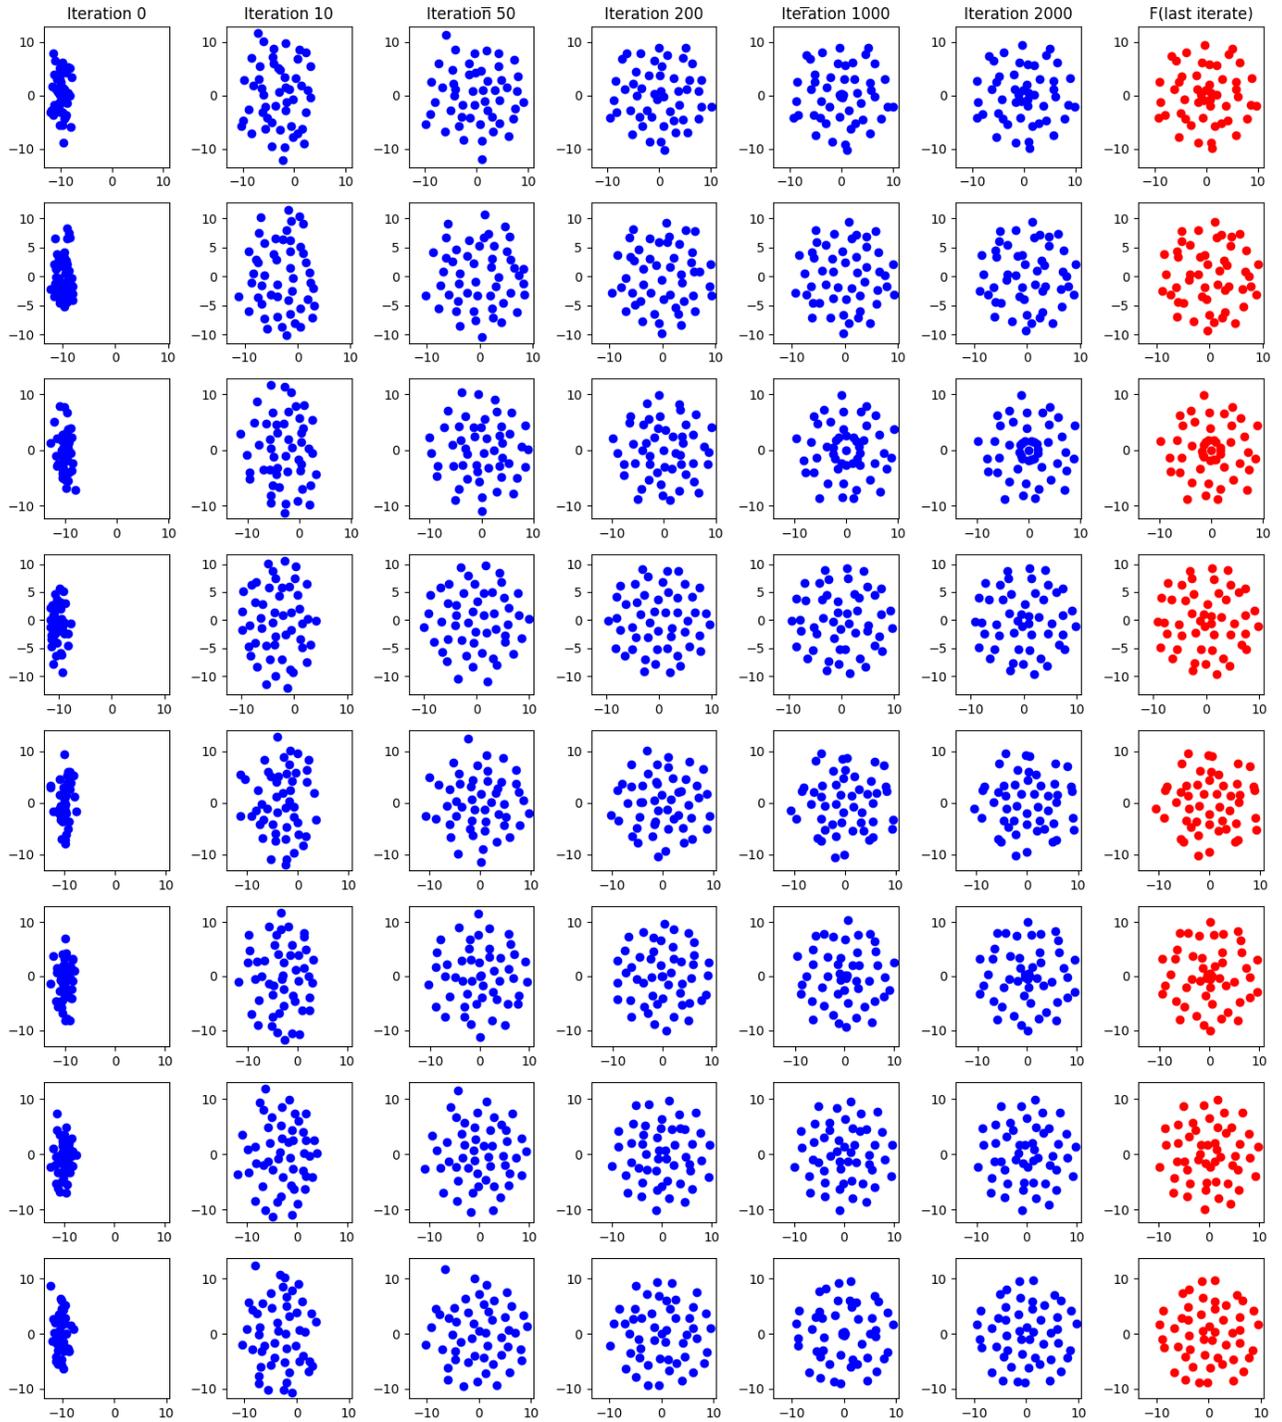


Figure 16: MMD gradient flow where the target functional is a rotation by $2\pi/5$. The columns show the flow at initialization, after 10, 50, 200, 1000, and 2000 iterations, and the rotation applied to the last iterate (in red).

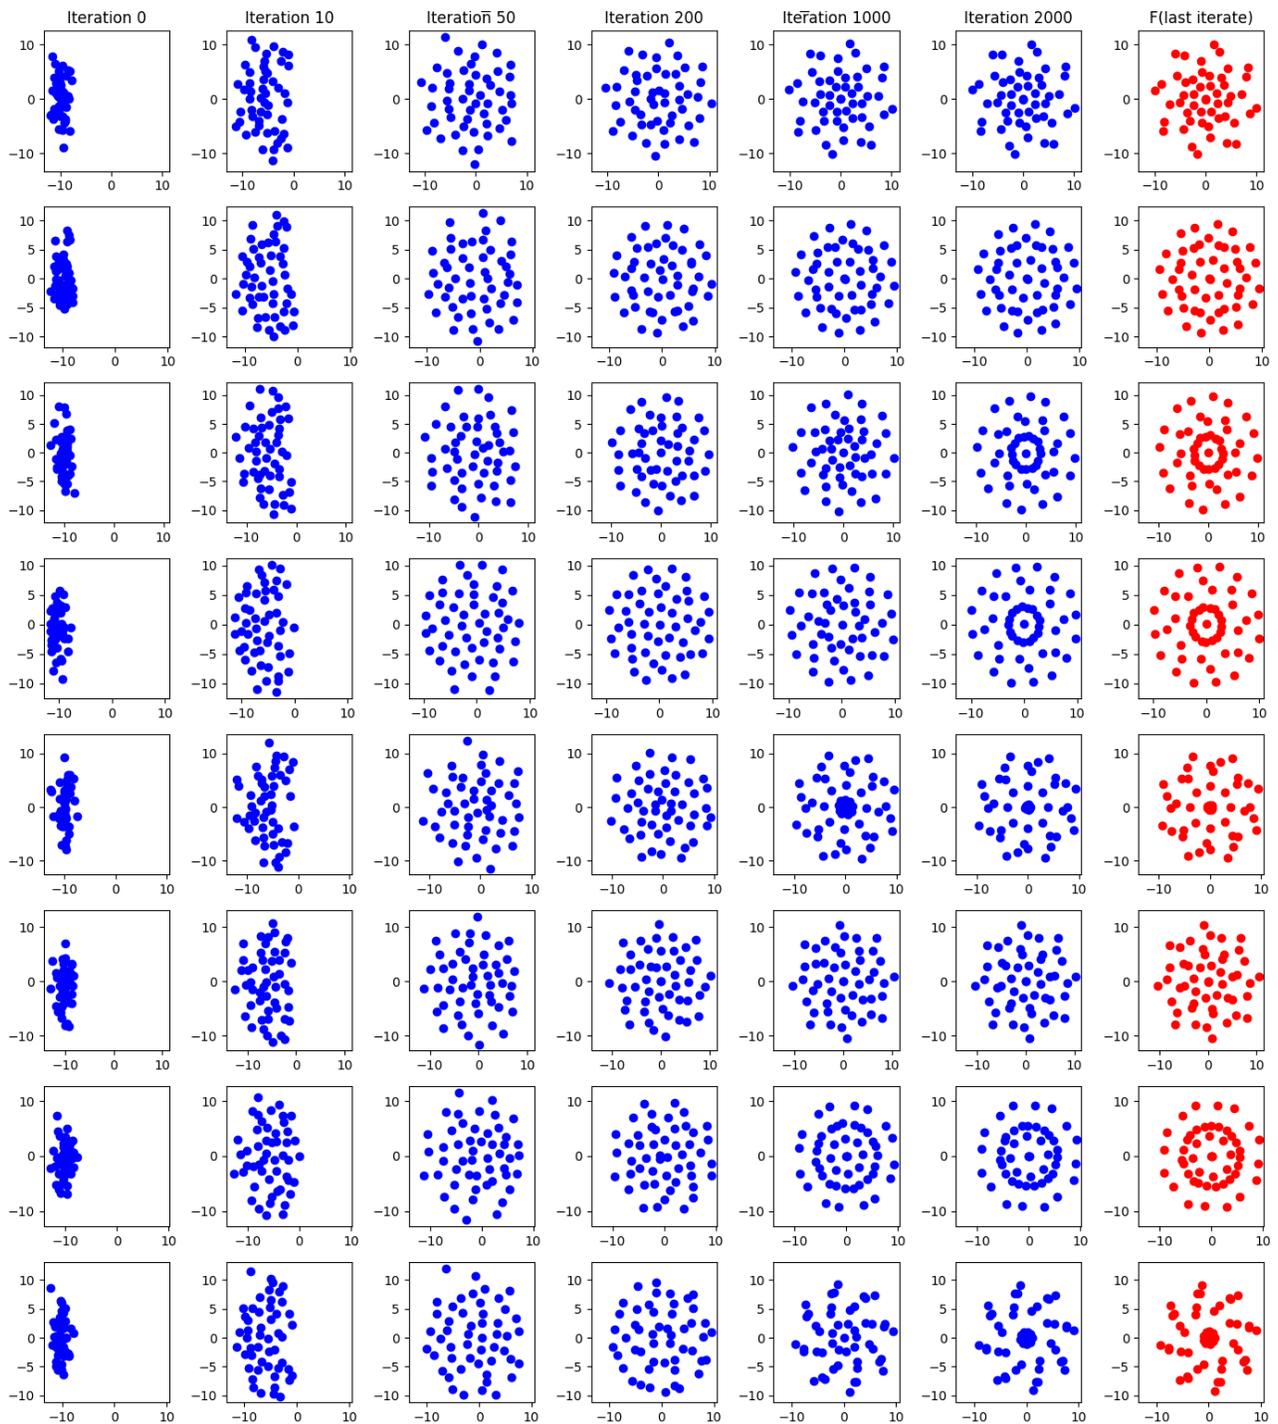


Figure 17: MMD gradient flow where the target functional is a rotation by $2\pi/8$. The columns show the flow at initialization, after 10, 50, 200, 1000, and 2000 iterations, and the rotation applied to the last iterate (in red).