

A Gradual, Semi-Discrete Approach to Generative Network Training via Explicit Wasserstein Minimization

Yucheng Chen¹, Matus Telgarsky¹, Chao Zhang¹, Bolton Bailey¹,
Daniel Hsu², and Jian Peng¹

¹University of Illinois at Urbana-Champaign, Urbana, IL

²Columbia University, New York, NY

Abstract

This paper provides a simple procedure to fit generative networks to target distributions, with the goal of a small Wasserstein distance (or other optimal transport cost). The approach is based on two principles: (a) if the source randomness of the network is a continuous distribution (the “semi-discrete” setting), then the Wasserstein distance is realized by a deterministic optimal transport mapping; (b) given an optimal transport mapping between a generator network and a target distribution, the Wasserstein distance may be decreased via a regression between the generated data and the mapped target points. The procedure here therefore alternates these two steps, forming an optimal transport and regressing against it, gradually adjusting the generator network towards the target distribution. Mathematically, this approach is shown to minimize the Wasserstein distance to both the empirical target distribution, and also its underlying population counterpart. Empirically, good performance is demonstrated on the training and testing sets of the MNIST and Thin-8 data. The paper closes with a discussion of the unsuitability of the Wasserstein distance for certain tasks, as has been identified in prior work [Arora et al., 2017, Huang et al., 2017].

1 Introduction

A generative network g models a distribution by first sampling $x \sim \mu$ from some simple distribution μ (e.g., a multivariate Gaussian), and thereafter outputting $g(x)$; this sampling procedure defines a *pushforward distribution* $g\#\mu$. A common training procedure to fit $g\#\mu$ to a target distribution $\hat{\nu}$ is to minimize a divergence $\mathcal{D}(g\#\mu, \hat{\nu})$ over a collection of parameters defining g .

The original algorithms for this framework, named *generative adversarial networks*, alternatively optimized both the generator network g , as well as a second *discriminator* of *adversarial* network: first the discriminator was fixed and the generator was optimized to fool it, and second the generator was fixed and the discriminator was optimized to distinguish it from $\hat{\nu}$. This procedure was originally constructed to minimize a Jensen-Shannon Divergence via a game-theoretic derivation [Goodfellow et al., 2014]. Subsequent work derived the adversarial relationship in other ways, for instance the Wasserstein GAN used duality properties of the Wasserstein distance [Arjovsky et al., 2017].

This paper proposes a simple non-adversarial but still alternating procedure to fit generative networks to target distributions. The procedure explicitly optimizes the Wasserstein- p distance between the generator $g\#\mu$ and the target distribution $\hat{\nu}$. As depicted in Figure 1, it alternates two steps: given a current generator g_i , an *Optimal Transport Solver (OTS)* associates (or “labels”) g_i ’s

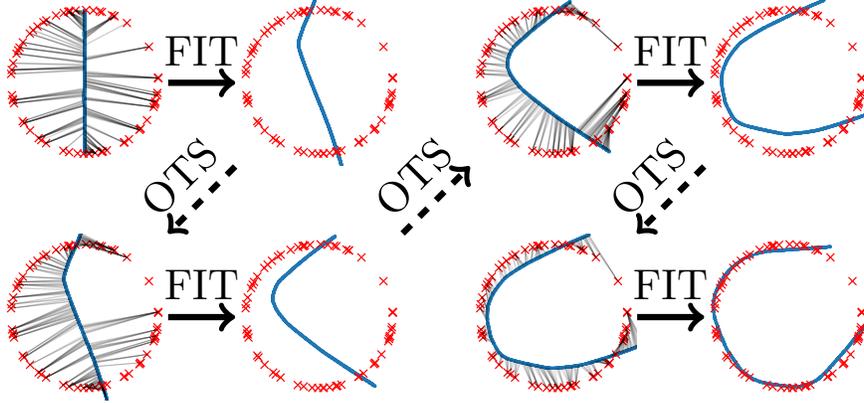


Figure 1: The goal in this example is to fit the initial distribution (the blue central line) to the target distribution (the red outer ring). The algorithm alternates OTS and FIT steps, first (OTS) associating input distribution samples with target distribution samples, and secondly (FIT) shifting input samples towards their targets, thereafter repeating the process. Thanks to being gradual, and not merely sticking to the first or second OTS, the process has a hope of constructing a simple generator which generalizes well.

probability mass with that of the target distribution $\hat{\nu}$, and then *FIT* uses this labeling to find a new generator g_{i+1} via a standard regression.

The effectiveness of this procedure hinges upon two key properties: it is *semi-discrete*, meaning the generators always give rise to continuous distributions, and it is *gradual*, meaning the generator is slowly shifted towards the target distribution. The key consequence of being semi-discrete is that the underlying optimal transport can be realized with a deterministic mapping. Solvers exist for this problem and construct a transport between the continuous distribution $g\#\mu$ and the target $\hat{\nu}$; by contrast, methods forming only batch-to-batch transports using samples from $g\#\mu$ are biased and do not exactly minimize the Wasserstein distance [Bellemare et al., 2017, Genevay et al., 2018, Salimans et al., 2018, Liu et al., 2018].

The procedure also aims to be gradual, as in Figure 1, slowly deforming the source distribution into the target distribution. While it is not explicitly shown that this gradual property guarantees a simple generator, promising empirical results measuring Wasserstein distance *to a test set* suggest that the learned generators generalize well.

Section 2 and Section 3 detail the method along with a variety of theoretical guarantees. Foremost amongst these are showing that the Wasserstein distance is indeed minimized, and secondly that it is minimized with respect to not just the dataset distribution $\hat{\nu}$, but moreover the underlying ν from which $\hat{\nu}$ was sampled. This latter property can be proved via the triangle inequality for Wasserstein distances, however such an approach introduces the Wasserstein distance between ν and $\hat{\nu}$, namely $W(\nu, \hat{\nu})$, which is exponential in dimension even in simple cases [Sriperumbudur et al., 2012, Arora et al., 2017]. Instead, we show that when a parametric model captures the distributions well, then bounds which are polynomial in dimension are possible.

Empirical results are presented in Section 4. We find that our method generates both quantitatively and qualitatively better digits than the compared baselines on MNIST, and the performance is consistent on both training and test datasets. We also experiment with the Thin-8 dataset [Huang et al., 2017], which is considered challenging for methods without a parametric loss. We discuss limitations of our method and conclude with some future directions.

2 Algorithm

We present our alternating procedure in this section. We first describe the OTS and FIT steps in detail, and then give the overall algorithm.

2.1 Optimal Transport Solver (OTS)

The Wasserstein- p distance W_p between two probability measures μ', ν' in a metric space (X, d) is defined as

$$W_p(\mu', \nu') := \left(\inf_{\gamma \in \Gamma(\mu', \nu')} \int_X d(x, y)^p d\gamma(x, y) \right)^{1/p},$$

where $\Gamma(\mu', \nu')$ is the collection of probability measures on $X \times X$ with marginal distributions μ' and ν' . $W_p(\mu', \nu')$ is equal to the $1/p$ -th power of the *optimal transport cost*

$$\mathcal{T}_c(\mu', \nu') := \inf_{\gamma \in \Gamma(\mu', \nu')} \int_X c(x, y) d\gamma(x, y)$$

with cost function $c(x, y) := d(x, y)^p$.

By *Kantorovich duality* [Villani, 2003, Chapter 1],

$$\mathcal{T}_c(\mu', \nu') = \sup_{\varphi, \psi \in \Phi_c} \int_X \varphi(x) d\mu'(x) + \int_X \psi(y) d\nu'(y), \quad (2.1)$$

where Φ_c is the collection of (φ, ψ) where $\varphi \in L_1(d\mu')$ and $\psi \in L_1(d\nu')$ (which means φ and ψ are absolutely Lebesgue integrable functions with respect to μ', ν'), and $\varphi(x) + \psi(y) \leq c(x, y)$ for almost all x, y).

In our generative modeling case, $\mu' = g\#\mu$ is a pushforward of the simple distribution μ by g , and $\nu' = \hat{\nu}$ is an *empirical measure*, meaning the uniform distribution on a training set $\{y_1, \dots, y_N\}$. When $g\#\mu$ is continuous, the optimal transport problem here becomes *semi-discrete*, and the maximizing choice of φ can be solved analytically, transforming the problem to optimization over a vector in \mathbb{R}^N [Genevay et al., 2016, Peyré and Cuturi, 2018]:

$$\begin{aligned} \mathcal{T}_c(g\#\mu, \hat{\nu}) &= \sup_{\varphi, \psi \in \Phi_c} \int_X \varphi(x) dg\#\mu(x) + \frac{1}{N} \sum_{i=1}^N \psi(y_i) \\ &= \sup_{\varphi \in \Phi'_{c, \hat{\psi}}, \hat{\psi} \in \mathbb{R}^N} \int_X \varphi(x) dg\#\mu(x) + \frac{1}{N} \sum_{i=1}^N \hat{\psi}_i \\ &= \sup_{\hat{\psi} \in \mathbb{R}^N} \int_X \min_i (c(x, y_i) - \hat{\psi}_i) dg\#\mu(x) + \frac{1}{N} \sum_{i=1}^N \hat{\psi}_i, \end{aligned} \quad (2.2)$$

where $\hat{\psi}_i := \psi(y_i)$, and $\Phi'_{c, \hat{\psi}}$ is the collection of functions $\varphi \in L_1(d(g\#\mu))$ such that $\varphi(x) + \hat{\psi}_i \leq c(x, y_i)$ for almost all x and $i = 1, \dots, N$. The third equality comes from the maximizing choice of φ : $\varphi(x) = \min_i (c(x, y_i) - \hat{\psi}_i) := \psi^c(x)$, the *c-transform* of ψ .

Our OTS solver, presented in Algorithm 1, uses SGD to maximize eq. (2.2), or rather to minimize its negation

$$- \int_X \min_i (c(x, y_i) - \hat{\psi}_i) dg\#\mu(x) - \frac{1}{N} \sum_{i=1}^N \hat{\psi}_i. \quad (2.3)$$

OTS is similar to Algorithm 2 of Genevay et al. [2016], but without averaging. Note as follows that Algorithm 1 is convex in $\hat{\psi}$, and thus a convergence theory of OTS could be developed, although this direction is not pursued here.

Algorithm 1 Optimal Transport Solver (OTS)

Input: continuous generated distribution $g\#\mu$, training dataset (y_1, \dots, y_n) corresponding to $\hat{\nu}$, cost function c , batch size B , learning rate η .

Output: $\hat{\psi} = (\hat{\psi}_1, \dots, \hat{\psi}_N)$

Initialize $t := 0$ and $\hat{\psi}^{(0)} \in \mathbb{R}^N$.

repeat

 Generate samples $\mathbf{x} = (x_1, \dots, x_B)$ from $g\#\mu$.

 Define loss $l(\mathbf{x}) := \frac{1}{B} \sum_{j=1}^B \min_i (c(x_j, y_i) - \hat{\psi}_i) + \frac{1}{N} \sum_{i=1}^N \hat{\psi}_i$.

 Update $\hat{\psi}^{(t+1)} := \hat{\psi}^{(t)} + \eta \cdot \nabla_{\hat{\psi}} l(\mathbf{x})$.

 Update $t := t + 1$.

until Stopping criterion is satisfied.

return $\hat{\psi}^{(t)}$.

Proposition 2.4. Equation (2.3) is a convex function of $\hat{\psi}$.

Proof. It suffices to note that $\min_i (c(g(x), y_i) - \hat{\psi}_i)$ is a minimum of concave functions and thus concave; eq. (2.3) is therefore concave since it is a convex combination of concave functions with an additional linear term. \square

In the *semi-discrete* setting where $g\#\mu$ is continuous and $\hat{\nu}$ is discrete, it can be proved that the *Kantorovich optimal transference plan* computed via eq. (2.1) is indeed a *Monge optimal transference plan* characterized by $\arg \min_i (x, y_i) - \hat{\psi}_i$, which is a deterministic mapping providing the regression target for our FIT step.

Proposition 2.5. Assume $X = \mathbb{R}^d$, $g\#\mu$ is continuous, and the cost function $c(x, y)$ takes the form of $c(x - y)$ and is a strictly convex, superlinear function on \mathbb{R}^d . Given the optimal $\hat{\psi}$ for eq. (2.3), then $T(x) := y_{\arg \min_i c(x, y_i) - \hat{\psi}_i}$, which is a Monge transference plan, is the unique Kantorovich optimal transference plan from $g\#\mu$ to ν .

(The proof is technical, and appears in the Appendix.)

We give some remarks to Proposition 2.5. First, Wasserstein- p distances on ℓ_p metric with $p > 1$ satisfy strict convexity and superlinearity [Gangbo and McCann, 1996], while $p = 1$ does not. On the other hand, in practice we have found that for $p = 1$, Algorithm 1 still converges to near-optimal transference plans, and this particular choice of metric generates crisper images than others.

Second, the continuity of $g\#\mu$, which is required for the uniqueness of OT plans may be violated in practice, but can be theoretically circumvented by adding an arbitrarily small perturbation to g 's output. On the other hand, since the optimal plan lies in the set $\varphi(x) + \psi(y) = c(x, y)$ almost surely [Villani, 2009, Theorem 5.10], it has to have the ‘‘argmin form’’ in Proposition 2.5. Thus, the only condition in which $T(x) := y_{\arg \min_i c(x, y_i) - \hat{\psi}_i}$ does not characterize a unique Monge OT plan, is the existence of ties when computing arg min, which does not happen in practice. We give some additional discussion about this issue in the Appendix.

A drawback of Algorithm 1 is that computing minimum on the whole dataset has $O(N)$ complexity, which is costly for extremely large datasets. We will revisit this issue in Section 4.

2.2 Fitting the Optimal Transference Plan (FIT)

Given an initial generator g , and an optimal transference plan T between $g\#\mu$ and $\hat{\nu}$ thanks to OTS, we update g to obtain a new generator g' by simply sampling $z \sim \mu$ and regressing the *new* generated sample $g'(z)$ towards the *old* OT plan $T(g(z))$, as detailed in Algorithm 2.

Under a few assumptions detailed in Section 3.1, Algorithm 2 is guaranteed to return a generator g' with strictly lesser optimal transport cost $\mathcal{T}_c(g'\#\mu, \hat{\nu}) \leq \mathbb{E}_{x \sim g'\#\mu} c(x, T(x)) < \mathbb{E}_{x \sim g\#\mu} c(x, T(x)) = \mathcal{T}_c(g\#\mu, \hat{\nu})$, where T denotes an exact optimal plan between $g\#\mu$ and $\hat{\nu}$; Section 3.1 moreover considers the case of a merely approximately optimal T , as returned by OTS.

Algorithm 2 Fitting Optimal Transport Plan (FIT)

Input: sampling distribution μ , old generator g with parameter θ , transference plan T , cost function c , batch size B , learning rate η .

Output: new generator g' with parameter θ' .

Initialize $t := 0$ and g' with parameter $\theta^{(0)} = \theta$.

repeat

 Generate random noise $\mathbf{z} = (z_1, \dots, z_B)$ from μ .

 Define loss $l(\mathbf{z}) := \frac{1}{B} \sum_{j=1}^B c(g'(z_j), T(g(z_j)))$.

 Update $\theta^{(t+1)} := \theta^{(t)} - \eta \cdot \nabla_{\theta'} l(\mathbf{z})$.

 Update $t := t + 1$.

until Stopping criterion is satisfied.

return g' with parameter $\theta^{(t)}$.

2.3 The Overall Algorithm

The overall algorithm, presented in Algorithm 3, alternates between OTS and FIT: during iteration i , OTS solves for the optimal transport map T between old generated distribution $g_i\#\mu$ and $\hat{\nu}$, then FIT regresses $g_{i+1}\#\mu$ towards $T\#g_i\#\mu$ to obtain lower Wasserstein distance.

Algorithm 3 Overall Algorithm

Input: sampling distribution μ , training dataset (y_1, \dots, y_n) corresponding to $\hat{\nu}$, initialized generator g_0 with parameter θ_0 , cost function c , batch size B , learning rate η .

Output: final generator g with parameter θ .

Initialize $i := 0$ and g_0 with parameter $\theta^{(0)} = \theta_0$.

repeat

 Compute $\hat{\psi}_i := \text{OTS}(g_i\#\mu, \hat{\nu}, c, B, \eta)$.

 Get T_i as $T_i(x) := \arg \min_{y_i} c(x, y_i) - \hat{\psi}_i$.

 Compute $g_{i+1} := \text{FIT}(\mu, g_i, T_i, c, B, \eta)$ with parameter $\theta^{(i+1)}$.

 Update $i := i + 1$.

until Stopping criterion is satisfied.

return g with parameter $\theta^{(i)}$.

3 Theoretical Analysis

We now analyze the optimization and generalization properties of our Algorithm 3: we will show that the method indeed minimizes the empirical transport cost, meaning $\mathcal{T}_c(g_i\#\mu, \hat{\nu}) \rightarrow 0$, and also

generalizes to the transport cost over the underlying distribution, meaning $\mathcal{T}_c(g_i\#\mu, \nu) \rightarrow 0$.

3.1 Optimization Guarantee: $\mathcal{T}_c(g_i\#\mu, \hat{\nu}) \rightarrow 0$.

Our analysis works for costs \mathcal{T}_c whose β -th powers satisfy the triangle inequality, such as $\mathcal{T}_c := W_p^p$ over any metric space and $\beta = 1/p$, if $p \geq 1$ [Villani, 2003, Theorem 7.3].

Our method is parameterized by a scalar $\alpha \in (0, 1/2)$ whose role is to determine the relative precisions of OTS and FIT, controlling the gradual property of our method. Defining $C_\mu(f, g) := \int c(f(x), g(x)) d\mu(x)$, we assume that for each round i , there exist error terms $\epsilon_{\text{ot1}}, \epsilon_{\text{ot2}}, \epsilon_{\text{fit}}$ such that:

- (1) Round i of OTS finds transport T_i satisfying

$$\mathcal{T}_c^\beta(g_i\#\mu, \hat{\nu}) \leq C_\mu^\beta(T_i \circ g_i, g_i) \leq \mathcal{T}_c^\beta(g_i\#\mu, \hat{\nu})(1 + \epsilon_{\text{ot1}})$$

(approximate optimality), and $\mathcal{T}_c^\beta(T_i\#g_i\#\mu, \hat{\nu}) \leq \epsilon_{\text{ot2}} \leq \alpha\mathcal{T}_c^\beta(g_i\#\mu, \hat{\nu})$ (approximate pushforward);

- (2) Round i of FIT finds g_{i+1} satisfying

$$\begin{aligned} C_\mu^\beta(T_i \circ g_i, g_{i+1}) &\leq \epsilon_{\text{fit}} \leq \frac{1 - 2\alpha}{1 + \epsilon_{\text{ot1}}} C_\mu^\beta(T_i \circ g_i, g_i) \\ &\leq (1 - 2\alpha)\mathcal{T}_c^\beta(g_i\#\mu, \hat{\nu}) \text{ (progress of FIT)}. \end{aligned}$$

In addition, we assume each $g_i\#\mu$ is continuous to guarantee the existence of Monge transport plan.

(1) is satisfied by Algorithm 1 since it represents a convex problem; moreover, it is necessary in practice to assume only approximate solutions. (2) holds when there is still room for the generative network to improve Wasserstein distance: otherwise, the training process can be stopped.

α is a tunable parameter of our overall algorithm: a large α relaxes the optimality requirement of OTS (which allows early stopping of Algorithm 1) but requires large progress of FIT (which prevents early stopping of Algorithm 2), and vice versa. This gives us a principled way to determine the stopping criteria of OTS and FIT.

Given the assumptions, we now show $\mathcal{T}_c(g_i\#\mu, \hat{\nu}) \rightarrow 0$. By triangle inequality,

$$\begin{aligned} &\mathcal{T}_c^\beta(g_{i+1}\#\mu, \hat{\nu}) \\ &\leq \mathcal{T}_c^\beta(g_{i+1}\#\mu, T_i\#g_i\#\mu) + \mathcal{T}_c^\beta(T_i\#g_i\#\mu, \hat{\nu}) \\ &\leq \mathcal{T}_c^\beta(g_{i+1}\#\mu, T_i\#g_i\#\mu) + \epsilon_{\text{ot2}}. \end{aligned}$$

Since $g_{i+1}\#\mu$ is continuous, $\mathcal{T}_c(g_{i+1}\#\mu, T_i\#g_i\#\mu)$ is realized by some deterministic transport T'_i satisfying $T'_i\#g_{i+1}\#\mu = T_i\#g_i\#\mu$, whereby

$$\begin{aligned} \mathcal{T}_c^\beta(g_{i+1}\#\mu, T_i\#g_i\#\mu) &= C_\mu^\beta(T'_i \circ g_{i+1}, g_{i+1}) \\ &= C_\mu^\beta(T_i \circ g_i, g_{i+1}) \leq \epsilon_{\text{fit}}. \end{aligned}$$

Combining these steps with the upper bounds on ϵ_{ot2} and ϵ_{fit} ,

$$\begin{aligned} \mathcal{T}_c^\beta(g_{i+1}\#\mu, \hat{\nu}) &\leq \epsilon_{\text{ot2}} + \epsilon_{\text{fit}} \leq (1 - \alpha)\mathcal{T}_c^\beta(g_i\#\mu, \hat{\nu}) \\ &\leq e^{-\alpha}\mathcal{T}_c^\beta(g_i\#\mu, \hat{\nu}). \end{aligned}$$

Summarizing these steps and iterating this inequality gives the following bound on $\mathcal{T}_c(g_t\#\mu, \hat{\nu})$, which goes to 0 as $t \rightarrow 0$.

Theorem 3.1. *Suppose (as discussed above) that \mathcal{T}_c^β satisfies the triangle inequality, each $g_i\#\mu$ is continuous, and the OTS and FIT iterations satisfy (1) (2), then $\mathcal{T}_c(g_t\#\mu, \hat{\nu}) \leq e^{-t\alpha/\beta}\mathcal{T}_c(g_0\#\mu, \hat{\nu})$.*

3.2 Generalization Guarantee: $\mathcal{T}_c(g_i\#\mu, \nu) \rightarrow 0$.

In the context of generative modeling, *generalization* means that the model fitted via training dataset $\hat{\nu}$ not only has low divergence $\mathcal{D}(g_i\#\mu, \hat{\nu})$ to $\hat{\nu}$, but also low divergence to ν , the underlying distribution from which $\hat{\nu}$ is drawn *i.i.d.*. If \mathcal{T}_c satisfies triangle inequality, then

$$\mathcal{T}_c(g_i\#\mu, \nu) \leq \mathcal{T}_c(g_i\#\mu, \hat{\nu}) + \mathcal{T}_c(\hat{\nu}, \nu),$$

and the second term goes to 0 with sample size $n \rightarrow \infty$, but the sample complexity depends exponentially on the dimensionality [Sriperumbudur et al., 2012, Arora et al., 2017, Huang et al., 2017]. To remove this exponential dependence, we make parametric assumptions about the underlying distribution ν ; a related idea was investigated in detail in parallel work [Bai et al., 2019].

Our approach is to assume the *Kantorovich potential* $\hat{\psi}$, defined on $\hat{\nu}$, is induced from a function $\psi \in \Psi$ defined on ν , where Ψ is a function class with certain approximation and generalization guarantees. Since neural networks are one such function classes (as will be discussed later), this is an empirically verifiable assumption (by fitting a neural network to approximate $\hat{\psi}$), and is indeed verified in the Appendix.

For this part we use slightly different notation: for a fixed sample size n , let $(g_n, T_n, \hat{\nu}_n)$ denote the earlier $(g, T, \hat{\nu})$. We first suppose the following *approximation condition*: Suppose that for any $\epsilon > 0$, there exists a class of functions Ψ so that

$$\sup_{\psi \in L_1(\nu)} \int \psi^c d\mu + \int \psi d\nu \leq \epsilon + \sup_{\psi \in \Psi} \int \psi^c d\mu + \int \psi d\nu; \quad (3.2)$$

thanks to the extensive literature on function approximation with neural networks [Hornik et al., 1989, Cybenko, 1989, Yarotsky, 2016], there are various ways to guarantee this, for example increasing the depth of the network. A second assumption is a *generalization condition*: given any sample size n and function class Ψ , suppose there exists $D_{n, \Psi} \geq 0$ so that with probability at least $1 - \delta$ over a draw of n examples from ν (giving rise to empirical measure $\hat{\nu}_n$), every $\psi \in \Psi$ satisfies

$$\int \psi d\nu \leq D_{n, \Psi} + \int \psi d\hat{\nu}_n; \quad (3.3)$$

thanks to the extensive theory of neural network generalization, there are in turn various ways to provide such a guarantee [Anthony and Bartlett, 1999], for example through VC-dimension of neural networks.

Combining these two assumptions,

$$\begin{aligned} \mathcal{T}_c(g_n\#\mu, \nu) &= \sup_{\psi \in L_1(\nu)} \left\{ \int \psi^c d(g_n\#\mu) + \int \psi d\nu \right\} \\ &\leq \epsilon + \sup_{\psi \in \Psi} \left\{ \int \psi^c d(g_n\#\mu) + \int \psi d\nu \right\} \\ &\leq D_{n, \Psi} + \epsilon + \sup_{\psi \in \Psi} \left\{ \int \psi^c d(g_n\#\mu) + \int \psi d\hat{\nu}_n \right\} \\ &\leq D_{n, \Psi} + \epsilon + \sup_{\psi \in L_1(\hat{\nu}_n)} \left\{ \int \psi^c d(g_n\#\mu) + \int \psi d\hat{\nu}_n \right\} \\ &\leq D_{n, \Psi} + \epsilon + \mathcal{T}_c(g_n\#\mu, \hat{\nu}_n). \end{aligned}$$

This can be summarized as follows.

Theorem 3.4. *Let $\epsilon > 0$ be given, and suppose assumptions eqs. (3.2) and (3.3) hold. Then, with probability at least $1 - \delta$ over the draw of n examples from ν ,*

$$\mathcal{T}_c(g_n \# \mu, \nu) \leq D_{n, \Psi} + \epsilon + \mathcal{T}_c(g_n \# \mu, \hat{\nu}_n).$$

By the earlier discussion, $D_{n, \Psi} \rightarrow 0$ and $\epsilon \rightarrow 0$ as $n \rightarrow \infty$, whereas the third term goes to 0 as discussed in Section 3.1.

4 Experimental Results

4.1 Experimental Setup

We briefly describe the datasets, baselines, and metrics in our experiments. Detailed descriptions of network architectures and the computations of evaluation metrics can be found in the Appendix.

Datasets. We evaluate our generative model on the MNIST and Thin-8 128×128 datasets [Lecun et al., 1998, Huang et al., 2017]. On MNIST, we use the original test/train split [Lecun et al., 1998], and each model is trained on the training set and evaluated on both training and testing sets. For Thin-8 we use the full dataset for training since the number of samples is limited.

Baselines. We compare our model against several neural net generative models: (1) WGAN [Arjovsky et al., 2017]; (2) WGANGP [Gulrajani et al., 2017]; (3) variational autoencoder (VAE) [Kingma and Welling, 2014]; (4) Wasserstein autoencoder (WAE) [Tolstikhin et al., 2017]. We experiment with both MLP and DCGAN as the generator architecture [Radford et al., 2015], and use DCGAN as the default discriminator/encoder architecture as it achieves better results for these baselines. Our method and WAE allow optimizing general optimal transport costs, and we choose to optimize the Wasserstein-1 distance on the ℓ_1 metric both for fair comparison with WGAN, and also since we observed clearer images on both MNIST and Thin-8.

Metrics. We use the following metrics to quantify the performance of different generative models: (1) Neural net distance (NND-WC, NND-GP) [Arora et al., 2017] based on DCGAN with weight clipping and gradient penalty respectively; (2) Wasserstein-1 distance (WD) on ℓ_1 metric; (3) Inception score (IS) [Salimans et al., 2016]; and (4) Fréchet Inception distance (FID) [Heusel et al., 2017].

We chose the above metrics because they capture different aspects of a generative model, and none of them is a one-size-fit-all evaluation measure. Among them, NND-WC and NND-GP are based on the adversarial game and thus biased in favor of WGAN and WGANGP. WD measures the Wasserstein distance between the generated distribution and the real dataset, and favors WAE and our method. IS and FID can be considered as neutral evaluation metrics, but they require labeled data or pretrained models to measure the performance of different models.

4.2 Qualitative Study

We first qualitatively investigate our generative model and compare the samples generated by different models.

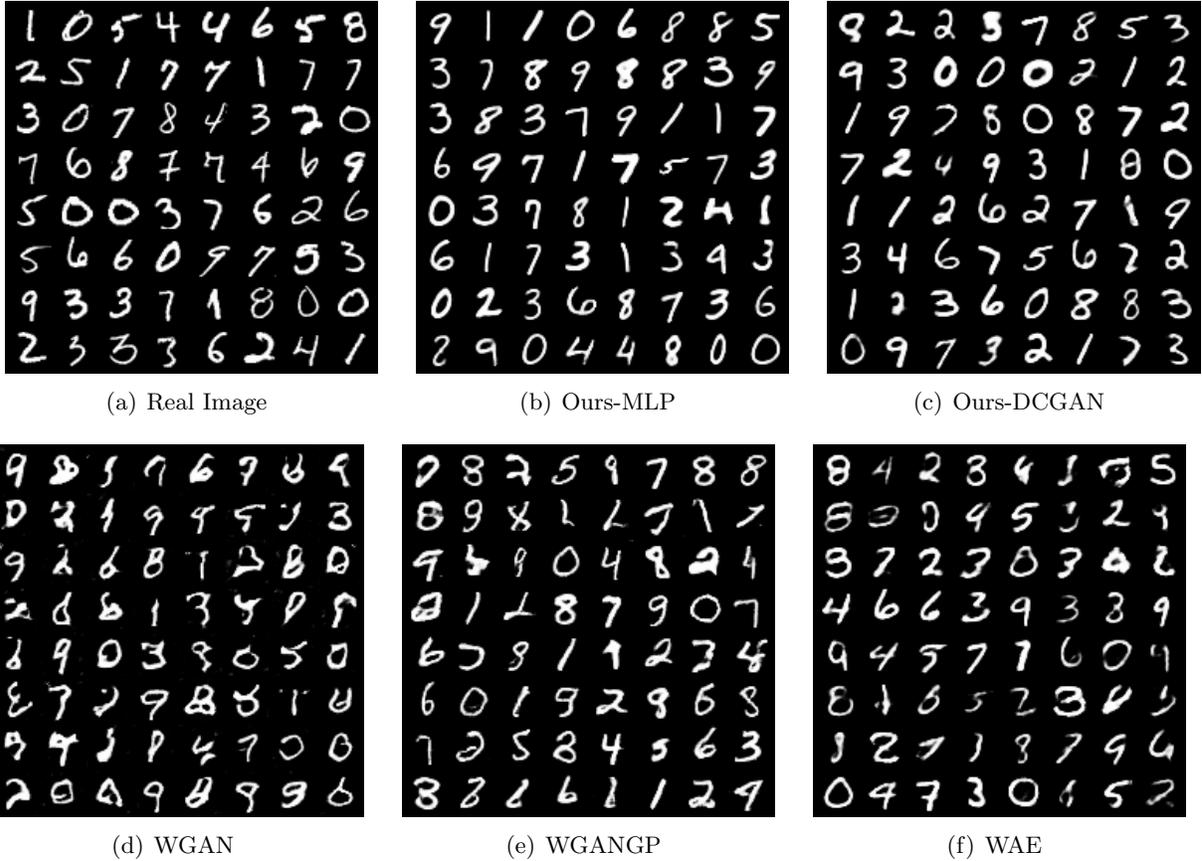


Figure 2: Real and generated samples on the MNIST dataset: (a) real samples; (b) samples generated by our model with MLP as the generator network; (c) samples generated by our model with DCGAN as the generator network; (d) samples generated by WGAN; (e) samples generated by WGANGP; (f) samples generated by WAE. DCGAN is used as the generator architecture in (d)(e)(f).

Samples of generated images. Figure 2 shows samples generated by different models on the MNIST dataset. The results show that our method with MLP (Figure 2(b)) and DCGAN (Figure 2(c)) both generate digits with better visual quality than the baselines with the DCGAN architecture. Figure 3 shows the generated samples on Thin-8 by our method, WGANGP, and WAE. The results of WGAN and VAE are omitted as they are similar to both WGANGP and WAE consistently on Thin-8. When MLP is used as the generator architecture, our method again outperforms WGANGP and WAE in terms of the visual quality of the generated samples. When DCGAN is used, the digits generated by our method have slightly worse quality than WGANGP, but better than WAE.

Importance of alternating procedure. We use this set of experiments to verify the importance of the alternating procedure that gradually improves the generative network. Figure 4 shows: (a) the samples generated by our model; and (b) the samples generated by a weakened version of our model that does not employ the alternating procedure. The non-alternating counterpart derives an optimal transport plan in the first run, and then fits towards the derived plan. It can be seen clearly the samples generated with such a non-alternating procedure have considerably lower visual quality.

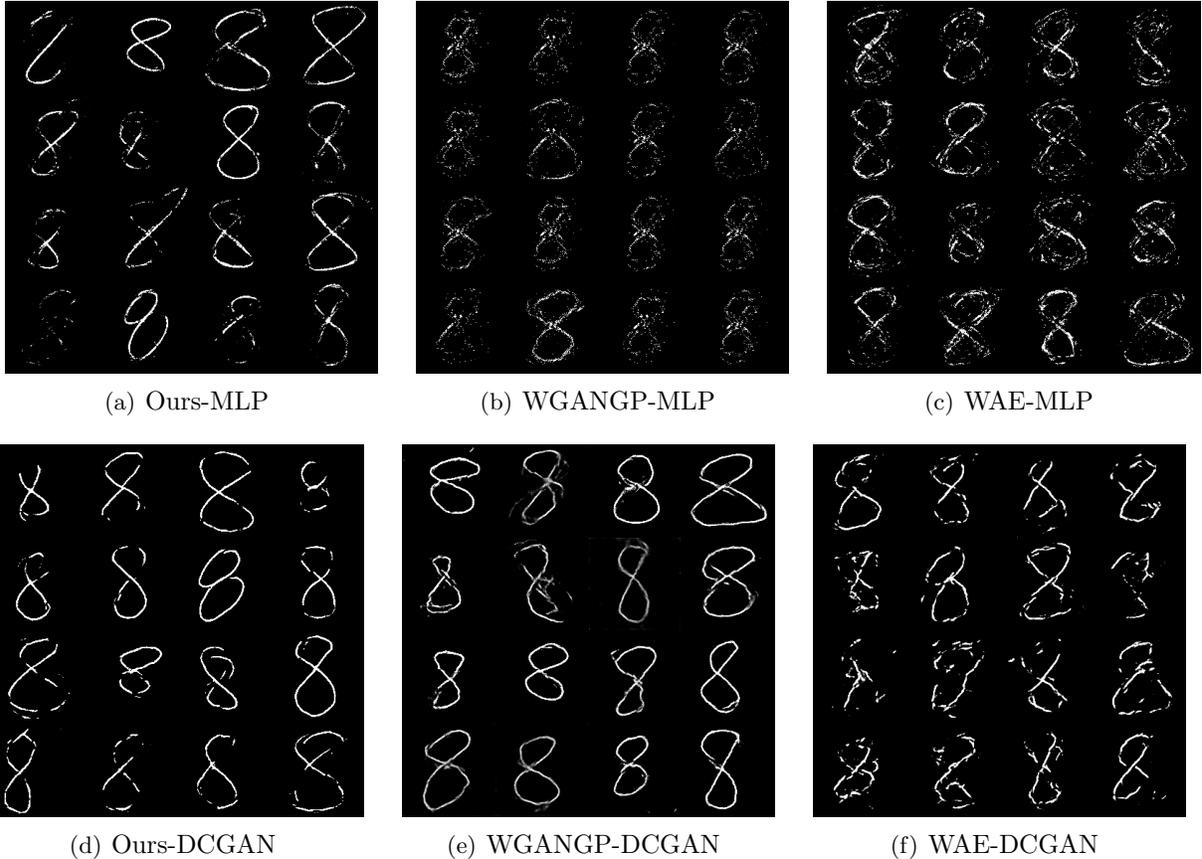


Figure 3: Generated samples on the Thin-8 dataset. (a)(b)(c) are samples generated by different methods using MLP as the generator; and (d)(e)(f) are samples when using DCGAN as the generator.

This verifies the importance of the alternating training procedure: fitting the generator towards the initial OT plan does not provide good enough gradient direction to produce a high-quality generator.

4.3 Quantitative Results

We proceed to measure the quantitative performance of the compared models.

MNIST results. Table 1 shows the performance of different models on the MNIST training and testing sets. In the first part when MLP is used to instantiate generators, our model achieves the best performance in terms of all the five metrics. The results on neural network distances (NND-WC and NND-GP) are particularly interesting: even though neural network distances are biased in favor of GAN-based models because the adversarial game explicitly optimizes such distances, our model still outperforms GAN-based models without adversarial training. The second part shows the results when DCGAN is the generator architecture. Under this setting, our method achieves the best results among all the metrics except for neural network distances. Comparing the performance of our method on the training and testing sets, one can observe its consistent performance and similar comparisons against baselines. This phenomenon empirically verifies that our method does not overfit.

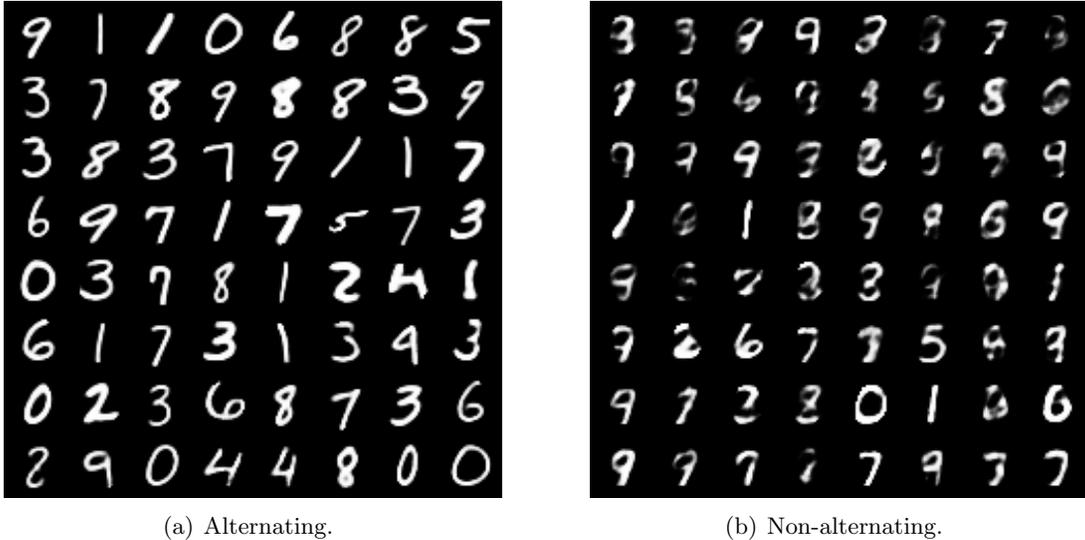


Figure 4: Generated samples on MNIST with and without the alternating procedure.

Table 1: Quantitative results on the MNIST training and testing sets. Note that the Wasserstein distance on training and testing sets of different sizes are not directly comparable.

Method	Arch	MNIST Training					MNIST Test			
		NND-WC	NND-GP	WD	IS	FID	NND-WC	NND-GP	WD	FID
WGAN	MLP	0.29	5.82	140.710	7.51	31.28	0.29	6.05	142.48	31.91
WGANGP		0.13	2.61	107.61	8.89	8.46	0.12	3.02	112.22	8.99
VAE		0.53	4.26	101.06	7.10	52.42	0.52	4.42	110.49	51.88
WAE		0.18	3.64	90.91	8.42	11.12	0.15	3.80	101.46	11.49
Ours		0.11	2.56	66.68	9.77	3.21	0.10	2.79	82.87	3.56
WGAN	DCGAN	0.11	4.69	125.63	7.02	27.64	0.10	4.86	132.97	28.44
WGANGP		0.08	0.83	93.61	8.65	4.65	0.07	1.66	104.15	5.45
VAE		0.48	3.68	106.63	6.96	42.10	0.46	3.89	115.59	41.95
WAE		0.18	3.29	90.96	8.35	12.28	0.15	3.53	101.02	12.66
Ours		0.10	2.28	70.13	9.54	3.76	0.09	2.55	82.79	4.18

Thin-8 results. There are no meaningful classifiers to compute IS and FID on the Thin-8 dataset. We thus only use NND-WC, NND-GP and WD as the quantitative metrics, and Table 2 shows the results. Our method obtains the best results among all the metrics with both the MLP and DCGAN architectures. For NND-WC, all methods except ours have similar results of around 3.1: we suspect this is due to the weight clipping effect, which is verified by tuning the clipping factor in our exploration. NND-GP and WD have consistent correlations for all the methods. This phenomenon is expected on a small-sized but high-dimensional dataset like Thin-8, because the discriminator neural network has enough capacity to approximate the Lipschitz-1 function class on the samples. The result comparison between NND-WC and NND-GP directly supports the claim Gulrajani et al. [2017] that gradient-penalized neural networks (NND-GP) has much higher approximation power than weight-clipped neural networks (NND-WC).

It is interesting to see that WGAN and WGANGP lead to the largest neural net distance and Wasserstein distance, yet their generated samples still have the best visual qualities on Thin-8.

This suggests that the success of GAN-based models cannot be solely explained by the restricted approximation power of discriminator [Arora et al., 2017, Huang et al., 2017].

Table 2: Quantitative results on the Thin-8 dataset.

Method	Arch	NND-WC	NND-GP	WD
WGAN	MLP	3.12	258.05	3934
WGANGP		3.12	144.38	2235
VAE		3.11	105.22	1950
WAE		3.07	111.79	1945
Ours		2.87	80.48	1016
WGAN	DCGAN	3.10	157.84	2481
WGANGP		3.04	79.47	1909
VAE		3.02	81.38	1820
WAE		3.11	88.04	1950
Ours		2.92	47.59	923

Time cost. Table 3 reports the training time of different models on MNIST. For moderate sized datasets such as MNIST, our method is faster than WGAN and WGANGP but slower than VAE and WAE. Compared with GAN-based models, our method does not have a discriminator which saves time. On the other hand, the loss function of eq. (2.3) requires computing $c(x, y_i) - \hat{\psi}_i$ for the whole dataset, which can be costly. That being said, a useful trick to accelerate our model is to take gradient of $\hat{\psi}$ over a moderately-sized subsample (for example 1%) of dataset in the first iterations, then gradually increase the subsample data size to cover the whole dataset.

5 Discussion of Limitations

We have also run our method on the CelebA and CIFAR10 datasets Liu et al. [2015], Krizhevsky [2009]. On CelebA, our method generates clear faces with good visual quality and with meaningful latent space interpolation, as shown in Figure 5(a) and Figure 5(b). However, we observe that the good visual quality partly comes from the average face effect: the expressions and backgrounds of generated images lack diversity compared with GAN-based methods.

Figure 6 shows the results of our method on CIFAR10. As shown, our method generates identifiable objects, but they are more blurry than GAN-based models. VAE generates objects that are also blurry but less identifiable. We compute the Wasserstein-1 distance of the compared methods on CIFAR10: our method (655), WGAN-GP (849) and VAE (745). Our method achieves the lowest Wasserstein distance but does not have better visual quality than GAN-based models on CIFAR10.

Table 3: Training time per iteration on MNIST.

Method	WGAN	WGANGP	VAE	WAE	Ours
Time (ms)	26.17	47.03	7.38	7.22	11.08

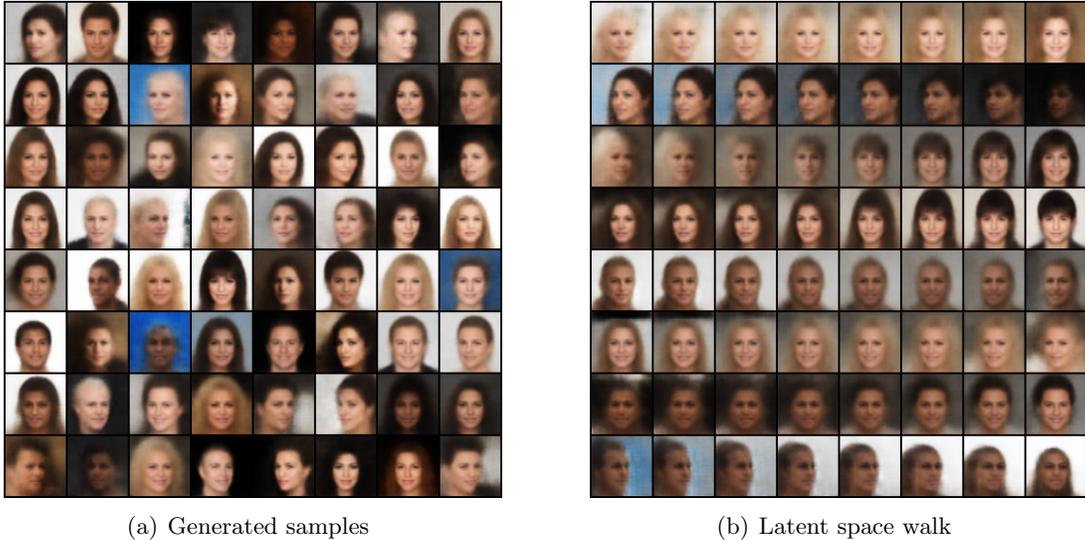


Figure 5: Samples generated by our method on CelebA, and a latent space interpolation.

Analyzing these results, we conjecture that minimizing Wasserstein distances on pixel-wise metrics such as ℓ_1 and ℓ_2 leads to a mode-collapse-free regularization effect. For models that minimize the Wasserstein distance, the primary task inherently tends to cover all the modes disregarding the sharpness of the generated samples. This is because not covering all the modes will result in huge transport cost. In GANs, the primary task is to generate sharp images which can fool the discriminator, and some modes can be dropped towards this goal. Consequently, the objective of our model naturally prevents it from mode collapse, but at the cost of generating more blurry samples. We propose two potential remedies to the blurriness issue: one is to use a perceptual loss [Bojanowski et al., 2018]; and the other is to incorporate adversarial metric into the framework. We leave them as future work.

6 Related Work

Optimal Transport. Optimal transport is an old yet vibrant topic [Villani, 2003, 2009, Peyré and Cuturi, 2018]. Genevay et al. [2016] give the stochastic formulation of semi-discrete optimal transport used in our solver. Lei et al. [2019] provide a geometric view of semi-discrete optimal transport. Seguy et al. [2018] propose to parameterize the Kantorovich potential via neural networks. They also propose to train generative networks by fitting towards the optimal transport plan between latent code and data, which can be considered as a special case of the non-alternating procedure we discussed earlier.

Generative models and OT. Combining generative modeling and optimal transport has been studied extensively. One line of research comes from the dual representation of Wasserstein distance via the Kantorovich-Rubinstein equality [Arjovsky et al., 2017, Gulrajani et al., 2017]. Another line optimizes the primal form of Wasserstein distance by relaxing it to a penalized form [Bousquet et al., 2017, Tolstikhin et al., 2017]. Genevay et al. [2017] give a comparison of WGAN and WAE from the view of optimal transport. By contrast, our work evaluate Wasserstein distance in dual space, and then optimizes it using its primal form.

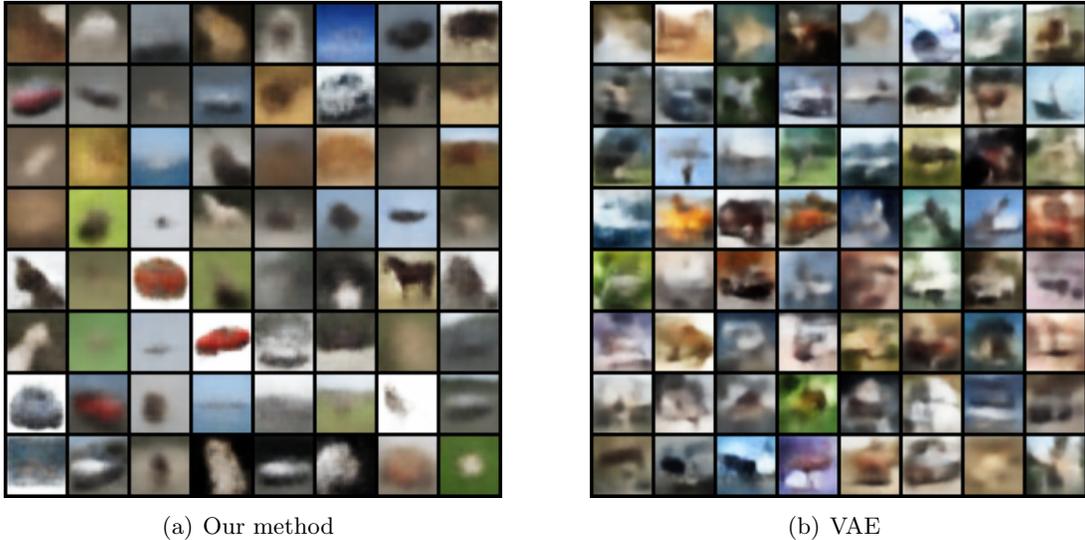


Figure 6: Samples generated by our method and VAE on CIFAR10.

The idea of computing optimal transport between batches of generated and real samples has been used in both non-adversarial generative modeling [Genevay et al., 2018, Xie et al., 2018], as well as adversarial generative modeling [Salimans et al., 2018, Liu et al., 2018]. However, minimizing batch-to-batch transport distance does not lead to the minimization of the Wasserstein distance between the generated and target distributions [Bellemare et al., 2017]. Instead, our method computes the whole-to-whole optimal transport via the semi-discrete formulation.

Generalization properties of Wasserstein distance. Sriperumbudur et al. [2012] analyze the sample complexity of evaluating integral probability metrics. Arora et al. [2017] show that KL-divergence and Wasserstein distances do not generalize well in high dimensions, but their neural net distance counterparts do generalize. Huang et al. [2017] give reasons for the advantage of GAN over VAE, and collect the Thin-8 dataset to demonstrate the advantage of GANs, which is used in our experiments. In this work we have compared both our theoretical and empirical findings with those in [Arora et al., 2017, Huang et al., 2017].

7 Conclusion and Future Work

We have proposed a simple alternating procedure to generative modeling by explicitly optimizing the Wasserstein distance between the generated distribution and real data. We show theoretically and empirically that our method does optimize Wasserstein distance to the training dataset, and generalizes to the underlying distribution. Interesting future work includes combining our method with adversarial and perceptual losses, and theoretical analysis on how gradual fitting contributes to smoother manifolds and better generalization.

Acknowledgements

BB and MJT are grateful for support from the NSF under grant IIS-1750051. DH acknowledges support from the Sloan Foundation. JP acknowledges support from Tencent AI Lab Rhino-Bird Gift Fund.

References

- Martin Anthony and Peter L. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, 1999.
- Martín Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *ICML*, 2017.
- Sanjeev Arora, Rong Ge, Yingyu Liang, Tengyu Ma, and Yi Zhang. Generalization and equilibrium in generative adversarial nets (GANs). In *ICML*, 2017.
- Yu Bai, Tengyu Ma, and Andrej Risteski. Approximability of discriminators implies diversity in GANs. In *ICLR*, 2019.
- Marc G. Bellemare, Ivo Danihelka, Will Dabney, Shakir Mohamed, Balaji Lakshminarayanan, Stephan Hoyer, and Rémi Munos. The cramer distance as a solution to biased wasserstein gradients. 2017. [arXiv:1705.10743](https://arxiv.org/abs/1705.10743) [cs.LG].
- Piotr Bojanowski, Armand Joulin, David Lopez-Pas, and Arthur Szlam. Optimizing the latent space of generative networks. In *ICML*, 2018.
- Olivier Bousquet, Sylvain Gelly, Ilya Tolstikhin, Carl-Johann Simon-Gabriel, and Bernhard Schölkopf. From optimal transport to generative modeling: the VEGAN cookbook, 2017. [arXiv:1705.07642](https://arxiv.org/abs/1705.07642) [stat.ML].
- George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(4):303–314, 1989.
- Wilfrid Gangbo and Robert J. McCann. The geometry of optimal transportation. *Acta Math.*, 177(2):113–161, 1996. doi: 10.1007/BF02392620. URL <https://doi.org/10.1007/BF02392620>.
- Aude Genevay, Marco Cuturi, Gabriel Peyré, and Francis R. Bach. Stochastic optimization for large-scale optimal transport. In *NIPS*, 2016.
- Aude Genevay, Gabriel Peyré, and Marco Cuturi. GAN and VAE from an optimal transport point of view. 2017. [arXiv:1706.01807](https://arxiv.org/abs/1706.01807) [stat.ML].
- Aude Genevay, Gabriel Peyré, and Marco Cuturi. Learning generative models with sinkhorn divergences. In *AISTATS*, 2018.
- Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014.
- Ishaan Gulrajani, Faruk Ahmed, Martín Arjovsky, Vincent Dumoulin, and Aaron C. Courville. Improved training of wasserstein GANs. In *NIPS*, 2017.
- Hartin Heusel, Hubert Ramsauer, Thomas Unterthiner, and Bernhard Nessler. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *NIPS*, 2017.
- K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, july 1989.
- Gabriel Huang, Gauthier Gidel, Hugo Berard, Ahmed Touati, and Simon Lacoste-Julien. Adversarial divergences are good task losses for generative modeling. 2017. [arXiv:1708.02511](https://arxiv.org/abs/1708.02511) [cs.LG].

- D. P. Kingma and M. Welling. Auto-encoding variational bayes. In *ICLR*, 2014.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- Yann Lecun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, pages 2278–2324, 1998.
- Na Lei, Kehua Su, Li Cui, Shing-Tung Yau, and Xianfeng David Gu. A geometric view of optimal transportation and generative model. *Computer Aided Geometric Design*, 68:1–21, 2019.
- Chunyuan Li, Hao Liu, Changyou Chen, Yunchen Pu, Liqun Chen, Ricardo Henao, and Lawrence Carin. ALICE: Towards understanding adversarial learning for joint distribution matching. In *NIPS*, 2017.
- Huidong Liu, GU Xianfeng, and Dimitris Samaras. A two-step computation of the exact gan wasserstein distance. In *ICML*, 2018.
- Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *ICCV*, 2015.
- Gabriel Peyré and Marco Cuturi. Computational optimal transport. 2018. [arXiv:1803.00567 \[stat.ML\]](#).
- Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. 2015. [arXiv:1511.06434 \[cs.LG\]](#).
- Tim Salimans, Ian J. Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training GANs. 2016. [arXiv:1606.03498 \[cs.LG\]](#).
- Tim Salimans, Han Zhang, Alec Radford, and Dimitris Metaxas. Improving GANs using optimal transport. In *ICLR*, 2018.
- Vivien Seguy, Bharath Bhushan Damodaran, Rémi Flamary, Nicolas Courty, Antoine Rolet, and Mathieu Blondel. Large-scale optimal transport and mapping estimation. In *ICLR*, 2018.
- B. Sriperumbudur, K. Fukumizu, A. Gretton, B. Schölkopf, and G. Lanckriet. On the empirical estimation of integral probability metrics. *Electronic Journal of Statistics*, 6:1550–1599, 2012.
- Ilya Tolstikhin, Olivier Bousquet, Sylvain Gelly, and Bernhard Schölkopf. Wasserstein auto-encoders. 2017. [arXiv:1711.01558](#).
- Cédric Villani. *Topics in Optimal Transportation*. American Mathematical Society, 2003.
- Cédric Villani. *Optimal Transport: Old and New*. Springer, 2009.
- Yujia Xie, Xiangfeng Wang, Ruijia Wang, and Hongyuan Zha. A fast proximal point method for wasserstein distance. 2018. [arXiv:1802.04307 \[stat.ML\]](#).
- Dmitry Yarotsky. Error bounds for approximations with deep relu networks. 2016. [arXiv:1610.01145 \[cs.LG\]](#).

A Proof of Proposition 2.5

Proof. By [Villani, 2003, Theorem 2.44], if $T(x)$ can be uniquely determined by $T(x) = x - \nabla c^*(\nabla\varphi(x))$, where φ is defined in eq. (2.1), then $T(x)$ is the unique Monge-Kantorovich optimal transference plan. Defining $m := \arg \min_i c(x - y_i) - \hat{\psi}_i$ for convenience, we have

$$\begin{aligned}
 & T(x) = x - \nabla c^*(\nabla\varphi(x)) \\
 \implies & \quad x - T(x) = \nabla c^*(\nabla\varphi(x)) \\
 & \quad = \nabla c^*(\nabla_x \min_i c(x - y_i) - \hat{\psi}_i) \\
 & \quad = \nabla c^* \nabla_x (c(x - y_m) - \hat{\psi}_m) \\
 & \quad = \nabla c^* \nabla c(x - y_m) \\
 & \quad = x - y_m \\
 \implies & \quad T(x) = y_m \\
 & \quad = y_{\arg \min_i c(x - y_i) - \hat{\psi}_i}.
 \end{aligned}$$

□

B An Example of Non-Gradual Training

We present in Figure 7 a similar synthetic example to Figure 1, except that FIT step will iterate itself until stuck at a local optimum. The local optimum after the first OTS-FIT run is a zig-zag shaped curve which does not generalize. While alternating the OTS-FIT is able to further push the generated samples to targets, the learned manifold is not as smooth as in Figure 1.

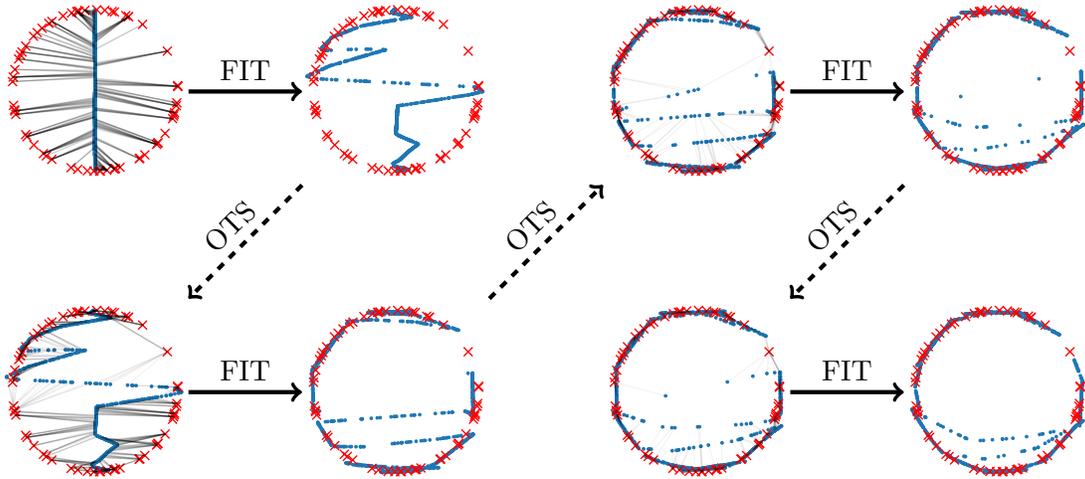


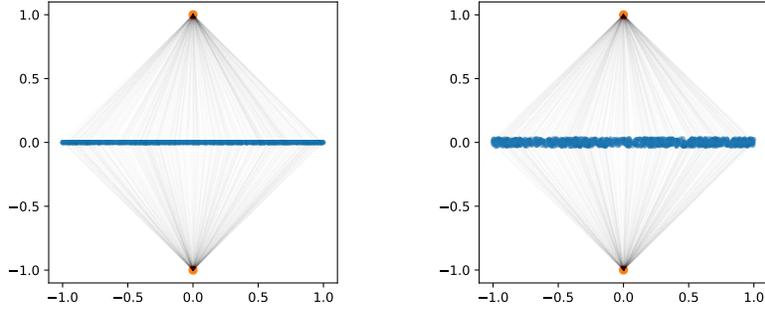
Figure 7: Example of non-gradual training.

C Further Discussion on the Uniqueness of OT Plan

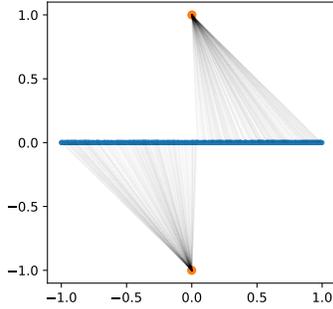
Continuity of $g\#\mu$ is required in [Villani, 2003, Theorem 2.44], which is violated if $g\#\mu$ lies in a low-dimensional manifold in \mathbb{R}^d , or the activation function is not strictly increasing (such as ReLU).

When $g\#\mu$ is not continuous, the optimal transference plan is not unique in general, and moreover it is possible that $T(x) := y_{\arg \min_i c(x, y_i) - \hat{\psi}_i}$ is not a transference plan at all (fig. 8(a)).

However, as discussed in Proposition 2.5, the only condition in which the “argmin form” does not characterize a unique Monge OT plan is the existence of ties. To break the ties, one can add arbitrarily small d -dimensional Gaussian noise to the outputs of g (fig. 8(b)). Furthermore, this makes the distribution continuous and satisfies the conditions of [Villani, 2003, Theorem 2.44]. On the other hand, supposing that the dataset $\hat{\nu}$ is drawn from a continuous distribution in \mathbb{R}^d , the event that two samples achieves exactly the same minimum $\min_i c(x, y_i) - \hat{\psi}_i$ happens with probability zero (fig. 8(c)).



(a) Example of non-uniqueness. (b) Any perturbation breaks the ties.



(c) When the samples are shifted.

Figure 8: (a) Suppose $g\#\mu$ is a uniform distribution on $[-1, 1] \times \{0\}$, and the samples are $y_1(0, 1), y_2(0, -1)$, then any equal-sized partition of $[-1, 1]$ would result to an optimal Monge plan. (b) Adding a small perturbation would make the optimal Monge plan unique. (c) If the samples become $y_1(0.001, 1), y_2(-0.001, -1)$, the ties are also broken.

D Verifying the Assumption in Section 3.2

We train $\hat{\psi}$ between our fitted generating distribution $g\#\mu$ and MNIST full dataset $\hat{\nu}$, then fit an MLP ψ with 4 hidden layers of 512 neurons to $\hat{\psi}$ on training dataset $\hat{\nu}_{\text{train}}$, and evaluated on both $\hat{\nu}_{\text{train}}$ and test dataset $\hat{\nu}_{\text{test}}$ (ψ takes 784-dimensional images as input and outputs a scalar). Figure 9 shows that the training error goes to zero and ψ has almost the same value as $\hat{\psi}$ when evaluated on $\hat{\nu}$.

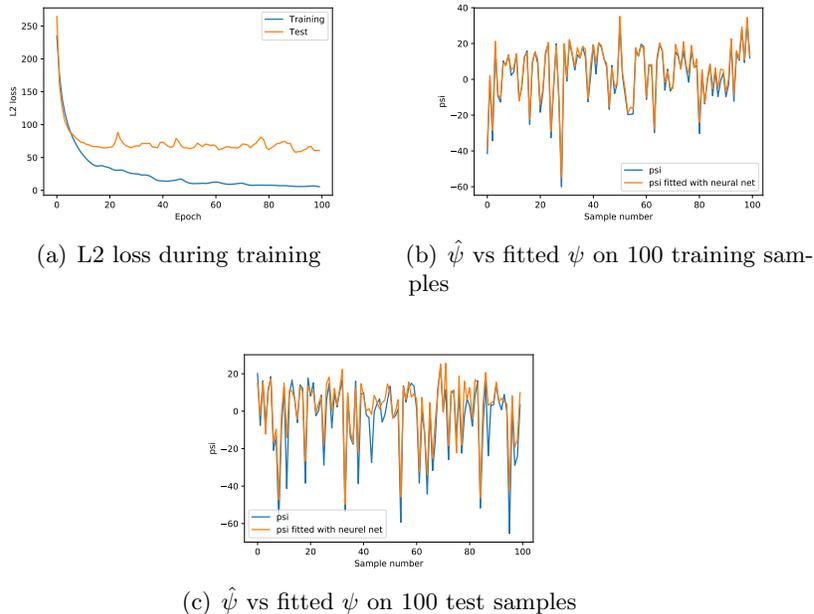


Figure 9: $\hat{\psi}$ on MNIST fitted by neural network.

We point out some concrete cases where the exponential lower bound is defeated by introducing a dependence on the measures μ and ν , without such empirical verification. If μ and ν are both piecewise constant densities with few pieces, then the supremum in Equation (3.2) can be well-approximated with piecewise affine functions with few pieces. In the worst case, the number of pieces can be exponential in dimension, but we will depend explicitly on the measure and its properties.

Another way to achieve the generalization result without the verification process, is to parametrize ψ as a neural network in the optimal transport algorithm, an idea related to [Seguy et al., 2018]. In Algorithm 1, we choose to optimize the vectorized $\hat{\psi}$ since it is a convex programming formulation guaranteed to converge to global minimum.

E Experimental Details

E.1 Evaluation Metrics

Neural net distance (NND-WC, NND-GP). Arora et al. [2017] define the *neural net distance* between the generated distribution $g\#\mu$ and dataset ν as:

$$\mathcal{D}_{\mathcal{F}}(g\#\mu, \nu) := \sup_{f \in \mathcal{F}} \mathbb{E}_{x \sim g\#\mu} f(x) - \mathbb{E}_{y \sim \nu} f(y),$$

where \mathcal{F} is a neural network function class. We use DCGAN with weight clipping at ± 0.01 , and DCGAN with gradient penalty with $\lambda = 10$ as two choices of \mathcal{F} . We call the corresponding neural net distances NND-WC and NND-GP respectively.

Wasserstein-1 distance (WD). This refers to the exact Wasserstein distance on ℓ_1 metric between the generated distribution μ and dataset ν , computed with our Algorithm 1.

Inception score (IS). Salimans et al. [2016] assume there exists a pretrained external classifier outputting label distribution $p(y|x)$ given sample x . The score is defined as $\text{IS}_p(\mu) := \exp\{\mathbb{E}_{x \sim \mu} \text{KL}(p(y|x) \parallel p(y))\}$.

Fréchet Inception distance (FID). Heusel et al. [2017] give this improvement over IS, which compares generated and real samples by the activations of a certain layer in a pretrained classifier. Assuming the activations follow Multivariate Gaussian distribution of mean μ_g, μ_r and covariance Σ_g, Σ_r , FID is defined as:

$$\text{FID}(\mu_g, \mu_r, \Sigma_g, \Sigma_r) := \|\mu_g - \mu_r\|^2 + \text{Tr}(\Sigma_r + \Sigma_g - 2(\Sigma_r \Sigma_g)^{1/2}).$$

Note that the naming of IS and FID is used since the classifier is an Inception-v3 network pretrained on Imagenet. Since the pretrained Inception network is not suitable for classifying handwritten digits, we follow an idea due to Li et al. [2017] and pretrain a CNN classifier on MNIST dataset with 99.1% test accuracy.

E.2 Neural Network Architectures and Hyperparameters

E.2.1 Generator

MLP:

$$\text{Input}(100) \xrightarrow{FC} \text{Hidden}(512) \xrightarrow{FC} \text{Hidden}(512) \xrightarrow{FC} \text{Hidden}(512) \xrightarrow{FC} \text{Output}(D).$$

DCGAN MNIST 28×28 :

$$\begin{aligned} \text{Input}(100, 1, 1) &\xrightarrow{\text{ConvT}(7,1,0)+\text{BN}} \text{Hidden}(128, 7, 7) \xrightarrow{\text{ConvT}(4,2,1)+\text{BN}} \text{Hidden}(64, 14, 14) \\ &\xrightarrow{\text{ConvT}(4,2,1)} \text{Output}(1, 28, 28). \end{aligned}$$

DCGAN Thin-8 128×128 :

$$\begin{aligned} \text{Input}(100, 1, 1) &\xrightarrow{\text{ConvT}(4,1,0)+\text{BN}} \text{Hidden}(256, 4, 4) \xrightarrow{\text{ConvT}(4,2,1)+\text{BN}} \text{Hidden}(128, 8, 8) \\ &\xrightarrow{\text{ConvT}(4,2,1)+\text{BN}} \text{Hidden}(64, 16, 16) \xrightarrow{\text{ConvT}(4,2,1)+\text{BN}} \text{Hidden}(32, 32, 32) \xrightarrow{\text{ConvT}(4,2,1)} \text{Hidden}(16, 64, 64) \\ &\xrightarrow{\text{ConvT}} \text{Output}(1, 128, 128). \end{aligned}$$

E.2.2 Discriminator/Encoder

DCGAN Discriminator 28×28 :

$$\text{Input}(1, 28, 28) \xrightarrow{\text{Conv}(4,2,1)} \text{Hidden}(64, 14, 14) \xrightarrow{\text{ConvT}(4,2,1)+\text{BN}} \text{Hidden}(128, 7, 7) \xrightarrow{FC} \text{Output}(1).$$

Thin-8 Discriminator 128×128 :

$$\begin{aligned} \text{Input}(1, 128, 128) &\xrightarrow{\text{Conv}(4,2,1)} \text{Hidden}(16, 64, 64) \xrightarrow{\text{ConvT}(4,2,1)+\text{BN}} \text{Hidden}(32, 32, 32) \\ &\xrightarrow{\text{ConvT}(4,2,1)+\text{BN}} \text{Hidden}(64, 16, 16) \xrightarrow{\text{ConvT}(4,2,1)+\text{BN}} \text{Hidden}(128, 8, 8) \xrightarrow{\text{ConvT}(4,2,1)+\text{BN}} \text{Hidden}(256, 4, 4) \\ &\xrightarrow{FC} \text{Output}(1). \end{aligned}$$

Encoder architectures are the same as discriminator architectures except for the output layers.

For WGANGP, batch normalization is not used since it affects computation of gradients. All activations used are ReLU. Learning rate is 10^{-4} for WGAN, WGANGP and our method, and 10^{-3} for VAE and WAE.

E.3 Training Details of Our Method

To get the transport mapping T in OTS, we memorize the sampled batches and their transportation targets, and reuse these batches in FIT. By this trick we avoid recomputing the maximum over the whole dataset.

Our empirical stopping criterion relies upon keeping a histogram of transportation targets in memory: if the histogram of targets is close to a uniform distribution (which is the distribution of training dataset), we stop OTS. This stopping criterion is grounded by our analysis in Section 3.1.