

Sliced-Wasserstein Autoencoder: An Embarrassingly Simple Generative Model

Abstract

In this paper we study generative modeling via autoencoders while using the elegant geometric properties of the optimal transport (OT) problem and the Wasserstein distances. We introduce Sliced-Wasserstein Autoencoders (SWAE), which are generative models that enable one to shape the distribution of the latent space into any samplable probability distribution without the need for training an adversarial network or defining a closed-form for the distribution. In short, we regularize the autoencoder loss with the sliced-Wasserstein distance between the distribution of the encoded training samples and a predefined samplable distribution. We show that the proposed formulation has an efficient numerical solution that provides similar capabilities to Wasserstein Autoencoders (WAE) and Variational Autoencoders (VAE), while benefiting from an embarrassingly simple implementation.

I. INTRODUCTION

Scalable generative models that capture the rich and often nonlinear distribution of high-dimensional data, (i.e., image, video, and audio), play a central role in various applications of machine learning, including transfer learning [14, 25], super-resolution [16, 21], image inpainting and completion [35], and image retrieval [7], among many others. The recent generative models, including Generative Adversarial Networks (GANs) [1, 2, 11, 30] and Variational Autoencoders (VAE) [5, 15, 24] enable an unsupervised and end-to-end modeling of the high-dimensional distribution of the training data.

Learning such generative models boils down to minimizing a dissimilarity measure between the data distribution and the output distribution of the generative model. To this end, and following the work of Arjovsky et al. [1] and Bousquet et al. [5] we approach the problem of generative modeling from the optimal transport point of view. The optimal transport problem [18, 34] provides a way to measure the distances between probability distributions by transporting (i.e., morphing) one distribution into another. Moreover, and as opposed to the common information theoretic dissimilarity measures (e.g., f -divergences), the p -Wasserstein dissimilarity measures that arise from the optimal transport problem: 1) are true distances, and 2) metrize a weak convergence of probability measures (at least on compact spaces). Wasserstein distances have recently attracted a lot of interest in the learning community [1, 5, 9, 12, 18] due to their exquisite geometric characteristics [31]. See the supplementary material for an intuitive example showing the benefit of the Wasserstein distance over commonly used f -divergences.

In this paper, we introduce a new type of autoencoders for generative modeling (Algorithm 1), which we call Sliced-Wasserstein Autoencoders (SWAE), that minimize the sliced-Wasserstein distance between the distribution of the encoded samples and a predefined samplable distribution. Our work is most closely related to the recent work by Bousquet et al. [5] and the follow-up work by Tolstikhin et al. [33]. However, our approach avoids the need to perform costly adversarial training in the encoding space and is not restricted to closed-form distributions, while still benefiting from a Wasserstein-like distance measure in the encoding space that permits a simple numerical solution to the problem.

In what follows we first provide an extensive review of the preliminary concepts that are needed for our formulation. In Section 3 we formulate our proposed method. The proposed

numerical scheme to solve the problem is presented in Section 4. Our experiments are summarized in Section 5. Finally, our work is Concluded in Section 6.

II. NOTATION AND PRELIMINARIES

Let X denote the compact domain of a manifold in Euclidean space and let $x_n \in X$ denote an individual input data point. Furthermore, let ρ_X be a Borel probability measure defined on X . We define the probability density function $p_X(x)$ for input data x to be:

$$d\rho_X(x) = p_X(x)dx$$

Let $\phi : X \rightarrow Z$ denote a deterministic parametric mapping from the input space to a latent space Z (e.g., a neural network encoder). Utilizing a technique often used in the theoretical physics community (See [10]), known as Random Variable Transformation (RVT), the probability density function of the encoded samples z can be expressed in terms of ϕ and p_X by:

$$p_Z(z) = \int_X p_X(x)\delta(z - \phi(x))dx, \quad (1)$$

where δ denotes the Dirac distribution function. The main objective of Variational Auto-Encoders (VAEs) is to encode the input data points $x \in X$ into latent codes $z \in Z$ such that: 1) x can be recovered/approximated from z , and 2) the probability density function of the encoded samples, p_Z , follows a prior distribution q_Z . Similar to classic auto-encoders, a decoder $\psi : Z \rightarrow X$ is required to map the latent codes back to the original space such that

$$p_Y(y) = \int_X p_X(x)\delta(y - \psi(\phi(x)))dx, \quad (2)$$

where y denotes the decoded samples. It is straightforward to see that when $\psi = \phi^{-1}$ (i.e. $\psi(\phi(\cdot)) = id(\cdot)$), the distribution of the decoder p_Y and the input distribution p_X are identical. Hence, the objective of a variational auto-encoder simplifies to learning ϕ and ψ such that they minimize a dissimilarity measure between p_Y and p_X , and between p_Z and q_Z . Defining and implementing the dissimilarity measure is a key design decision, and is one of the main contributions of this work, and thus we dedicate the next section to describing existing methods for measuring these dissimilarities.

A. Dissimilarity between p_X and p_Y

We first emphasize that the VAE work in the literature often assumes stochastic encoders and decoders [15], while we consider the case of only deterministic mappings. Different dissimilarity measures have been used between p_X and p_Y in various work in the literature. Most notably, Nowozin et al. [26] showed that for the general family of f -divergences, $D_f(p_X, p_Y)$, (including the KL-divergence, Jensen-Shannon, etc.), using the Fenchel conjugate of the convex function f and minimizing $D_f(p_X, p_Y)$ leads to a min-max problem that is equivalent to the *adversarial training* widely used in the generative modeling literature [11, 23, 24].

Others have utilized the rich mathematical foundation of the OT problem and Wasserstein distances [1, 5, 12, 33]. In Wasserstein-GAN, [1] utilized the Kantorovich-Rubinstein duality for the 1-Wasserstein distance, $W_1(p_X, p_Y)$, and reformulated the problem as a min-max optimization that is solved through an adversarial training scheme. In a different approach, [5] utilized the autoencoding nature of the problem and showed that $W_c(p_X, p_Y)$ could be simplified as:

$$W_c(p_X, p_Y) = \inf_{\phi, \psi} \int_X p_X(x)c(x, \psi(\phi(x)))dx \quad (3)$$

Note that Eq. (3) is equivalent to Theorem 1 in [5] for deterministic encoder-decoder pair, and also note that ϕ and ψ are parametric differentiable models (e.g. neural networks). Furthermore, Eq. (3) supports a simple implementation where for i.i.d samples of the input distribution $\{x_n\}_{n=1}^N$ the minimization can be written as:

$$W_c(p_X, p_Y) = \inf_{\phi, \psi} \frac{1}{N} \sum_{n=1}^N c(x_n, \psi(\phi(x_n))) \quad (4)$$

We emphasize that Eq. (3) (and consequently Eq. (4)) takes advantage of the fact that the pairs x_n and $y_n = \psi(\phi(x_n))$ are available, hence calculating the transport distance coincides with summing the transportation costs between all pairs (x_n, y_n) . For example, the total transport distance may be defined as the sum of Euclidean distances between all pairs of points. In this paper, we also use $W_c(p_X, p_Y)$ following Eq. (4) to measure the discrepancy between p_X and p_Y . Next, we review the methods used for measuring the discrepancy between p_Z and q_Z .

B. Dissimilarity between p_Z and q_Z

If q_Z is a known distribution with an explicit formulation (e.g. Normal distribution) the most straightforward approach for measuring the (dis)similarity between p_Z and q_Z is the log-likelihood of $z = \phi(x)$ with respect to q_Z , formally:

$$\sup_{\phi} \int_X p_X(x) \log(q_Z(\phi(x))) dx \quad (5)$$

maximizing the log-likelihood is equivalent to minimizing the KL-divergence between p_Z and q_Z , $D_{KL}(p_Z, q_Z)$ (see supplementary material for more details and derivation of Equation (5)). This approach has two major limitations: 1) The KL-Divergence and in general f -divergences do not provide meaningful dissimilarity measures for distributions supported on non-overlapping low-dimensional manifolds [1, 19] (see supplementary material), which is common in hidden layers of neural networks, and therefore they do not provide informative gradients for training ϕ , and 2) we are limited to distributions q_Z that have known explicit formulations, which is very restrictive because it eliminates the ability to use the much broader class of distributions were we know how to sample from them, but do not know their explicit form.

Various alternatives exist in the literature to address the above-mentioned limitations. These methods often sample $\tilde{Z} = \{\tilde{z}_j\}_{j=1}^N$ from q_Z and $Z = \{z_n = \phi(x_n)\}_{n=1}^N$ from p_X and measure the discrepancy between these sets (i.e. point clouds). Note that there are no one-to-one correspondences between \tilde{z}_j s and z_n s. Tolstikhin et al. [33] for instance, proposed two different approaches for measuring the discrepancy between \tilde{Z} and Z , namely the GAN-based and the *maximum mean discrepancy* (MMD)-based approaches. The GAN-based approach proposed in [33] defines a discriminator network, $D_Z(p_Z, q_Z)$, to classify \tilde{z}_j s and z_n s as coming from ‘true’ and ‘fake’ distributions correspondingly and proposes a min-max adversarial optimization for learning ϕ and D_Z . This approach could be thought as a Fenchel conjugate of some f -divergence between p_Z and q_Z . The MMD-based approach, on the other hand, utilizes a positive-definite reproducing kernel $k : Z \times Z \rightarrow \mathbb{R}$ to measure the discrepancy between \tilde{Z} and Z , however, the choice of the kernel remain a data-dependent design parameter.

An interesting alternative approach is to use the Wasserstein distance between p_Z and q_Z . The reason being that Wasserstein metrics have been shown to be particularly beneficial for measuring the distance between distributions supported on non-overlapping low-dimensional manifolds. Following the work of Arjovsky et al. [1], this can be accomplished utilizing the Kantorovich-Rubinstein duality and through introducing a min-max problem, which leads to yet another adversarial training scheme similar the GAN-based method in [33]. Note that, since

elements of $\tilde{\mathcal{Z}}$ and \mathcal{Z} are not paired an approach similar to Eq. (4) could not be used to calculate the Wasserstein distance. In this paper, we propose to use the sliced-Wasserstein metric, [3, 6, 17, 19, 28, 29], to measure the discrepancy between p_Z and q_Z . We show that using the sliced-Wasserstein distance ameliorates the need for training an adversary network, and provides an efficient but yet simple numerical implementation.

Before explaining our proposed approach, it is worthwhile to point out the benefits of learning autoencoders as generative models over GANs. In GANs, one needs to minimize a distance between $\{\psi(\tilde{z}_j) | \tilde{z}_j \sim q_Z\}_{j=1}^M$ and $\{x_n\}_{n=1}^M$ which are high-dimensional point clouds for which there are no correspondences between $\psi(\tilde{z}_j)$ s and x_n s. For the autoencoders, on the other hand, there exists correspondences between the high-dimensional point clouds $\{x_n\}_{n=1}^M$ and $\{y_n = \psi(\phi(x_n))\}_{n=1}^M$, and the problem simplifies to matching the lower-dimensional point clouds $\{\phi(x_n)\}_{n=1}^M$ and $\{\tilde{z}_j \sim q_Z\}_{j=1}^M$. In other words, the encoder performs a nonlinear dimensionality reduction, that enables us to solve a much simpler problem compared to GANs. Next we introduce the details of our approach.

III. PROPOSED METHOD

In what follows we first provide a brief review of the necessary equations to understand the Wasserstein and sliced-Wasserstein distances and then present our Sliced Wasserstein Autoencoders (SWAE).

A. Wasserstein distances

The Wasserstein distance between probability measures ρ_X and ρ_Y , with corresponding densities $d\rho_X = p_X(x)dx$ and $d\rho_Y = p_Y(y)dy$ is defined as:

$$W_c(p_X, p_Y) = \inf_{\gamma \in \Gamma(\rho_X, \rho_Y)} \int_{X \times Y} c(x, y) d\gamma(x, y) \quad (6)$$

where $\Gamma(\rho_X, \rho_Y)$ is the set of all transportation plans (i.e. joint measures) with marginal densities p_X and p_Y , and $c : X \times Y \rightarrow \mathbb{R}^+$ is the transportation cost. Eq. (6) is known as the Kantorovich formulation of the optimal mass transportation problem, which seeks the optimal transportation plan between p_X and p_Y . If there exist diffeomorphic mappings, $f : X \rightarrow Y$ (i.e. transport maps) such that $y = f(x)$ and consequently,

$$p_Y(y) = \int_X p_X(x) \delta(y - f(x)) dx \xrightarrow[\text{a diffeomorphism}]{\text{When } f \text{ is}} p_Y(y) = \det(Df^{-1}(y)) p_X(f^{-1}(y)) \quad (7)$$

where $\det(D \cdot)$ is the determinant of the Jacobian, then the Wasserstein distance could be defined based on the Monge formulation of the problem (see [34] and [18]) as:

$$W_c(p_X, p_Y) = \min_{f \in MP} \int_X c(x, f(x)) d\rho_X(x) \quad (8)$$

where MP is the set of all diffeomorphisms that satisfy Eq. (7). As can be seen from Eqs. (6) and (8), obtaining the Wasserstein distance requires solving an optimization problem. Various efficient optimization techniques have been proposed in the past (e.g. [8, 27, 32]).

The case of one dimensional probability densities, p_X and p_Y , is specifically interesting as the Wasserstein distance has a closed-form solution. Let P_X and P_Y be the cumulative distributions of one-dimensional probability distributions p_X and p_Y , correspondingly. The Wasserstein distance can then be calculated as:

$$W_c(p_X, p_Y) = \int_0^1 c(P_X^{-1}(\tau), P_Y^{-1}(\tau)) d\tau \quad (9)$$

The closed-form solution of Wasserstein distance for one-dimensional probability densities motivates the definition of sliced-Wasserstein distances.

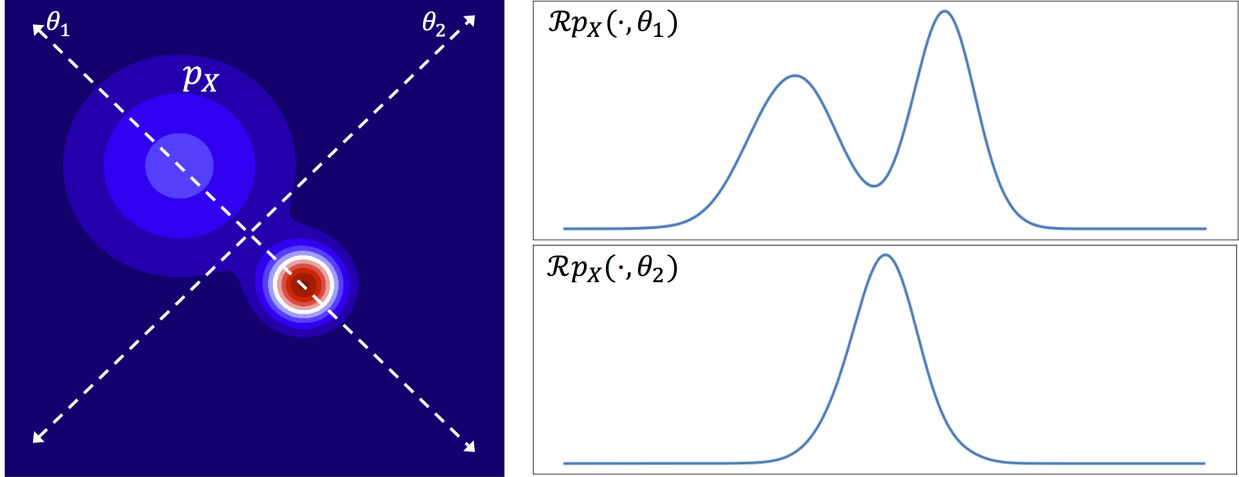


Fig. 1: Visualization of the slicing process defined in Eq. (10)

B. Sliced-Wasserstein distances

The interest in the sliced-Wasserstein distance is due to the fact that it has very similar qualitative properties as the Wasserstein distance, but it is much easier to compute, since it only depends on one-dimensional computations. The sliced-Wasserstein distance was used in [28, 29] to calculate barycenter of distributions and point clouds. Bonneel et al. [3] provided a nice theoretical overview of barycentric calculations using the sliced-Wasserstein distance. Kolouri et al. [17] used this distance to define positive definite kernels for distributions and Carriere et al. [6] used it as a distance for persistence diagrams. Sliced-Wasserstein was also recently used for learning Gaussian mixture models [19].

The main idea behind the sliced-Wasserstein distance is to slice (i.e. project) higher-dimensional probability densities into sets of one-dimensional distributions and compare their one-dimensional representations via Wasserstein distance. The slicing/projection process is related to the field of Integral Geometry and specifically the Radon transform [13]. The relevant result to our discussion is that a d -dimensional probability density p_X could be uniquely represented as the set of its one-dimensional marginal distributions following the Radon transform and the Fourier slice theorem [13]. These one dimensional marginal distributions of p_X are defined as:

$$\mathcal{R}p_X(t; \theta) = \int_{\mathcal{X}} p_X(x) \delta(t - \theta \cdot x) dx, \quad \forall \theta \in \mathbb{S}^{d-1}, \forall t \in \mathbb{R} \quad (10)$$

where \mathbb{S}^{d-1} is the d -dimensional unit sphere. Note that for any fixed $\theta \in \mathbb{S}^{d-1}$, $\mathcal{R}p_X(\cdot; \theta)$ is a one-dimensional slice of distribution p_X . In other words, $\mathcal{R}p_X(\cdot; \theta)$ is a marginal distribution of p_X that is obtained from integrating p_X over the hyperplane orthogonal to θ (See Figure 1). Utilizing the one-dimensional marginal distributions in Eq. (10), the sliced Wasserstein distance could be defined as:

$$SW_c(p_X, p_Y) = \int_{\mathbb{S}^{d-1}} W_c(\mathcal{R}p_X(\cdot; \theta), \mathcal{R}p_Y(\cdot; \theta)) d\theta \quad (11)$$

Given that $\mathcal{R}p_X(\cdot; \theta)$ and $\mathcal{R}p_Y(\cdot; \theta)$ are one-dimensional the Wasserstein distance in the integrand has a closed-form solution as demonstrated in (9). The fact that SW_c is a distance comes from W_c being a distance. Moreover, the two distances also induce the same topology, at least on compact sets [31].

A natural transportation cost that has extensively studied in the past is the ℓ_2^2 , $c(x, y) = \|x - y\|_2^2$, for which there are theoretical guarantees on existence and uniqueness of transportation plans and maps (see [31] and [34]). When $c(x, y) = \|x - y\|_2^2$ the following inequality bounds hold for the SW distance:

$$SW_2(p_X, p_Y) \leq W_2(p_X, p_Y) \leq \alpha SW_2^\beta(p_X, p_Y) \quad (12)$$

where α is a constant. Chapter 5 in [4] proves this inequality with $\beta = (2(d + 1))^{-1}$ (See [31] for more details). The inequalities in (12) is the main reason we can use the sliced Wasserstein distance, SW_2 , as an approximation for W_2 .

C. Sliced-Wasserstein auto-encoder

Our proposed formulation for the SWAE is as follows:

$$\operatorname{argmin}_{\phi, \psi} W_c(p_X, p_Y) + \lambda SW_c(p_Z, q_Z) \quad (13)$$

where ϕ is the encoder, ψ is the decoder, p_X is the data distribution, p_Y is the data distribution after encoding and decoding (Eq. (2)), p_Z is the distribution of the encoded data (Eq. (1)), q_Z is the predefined distribution (or a distribution we know how to sample from), and λ is a hyperparameter that identifies the relative importance of the loss functions.

To further clarify why we use the Wasserstein distance to measure the difference between p_X and p_Y , but the *sliced-Wasserstein* distance to measure the difference between p_Z and q_Z , we reiterate that the Wasserstein distance for the first term can be solved via Eq. (4) due to the existence of correspondences between y_n and x_n (i.e., we desire $x_n = y_n$), however, for p_Z and q_Z , analogous correspondences between the \tilde{z}_i s and z_j s do not exist and therefore calculation of the Wasserstein distance requires an additional optimization step (e.g., in the form of an adversarial network). To avoid this additional optimization, while maintaining the favorable characteristics of the Wasserstein distance, we use the sliced-Wasserstein distance to measure the discrepancy between p_Z and q_Z .

IV. NUMERICAL OPTIMIZATION

A. Numerical implementation of the Wasserstein distance in 1D

The Wasserstein distance between two one-dimensional distributions p_X and p_Y is obtained from Eq. (9). The integral in Eq. (9) could be numerically calculated using $\frac{1}{M} \sum_{m=1}^M a_m$, where $a_m = c(P_X^{-1}(\tau_m), P_Y^{-1}(\tau_m))$ and $\tau_m = \frac{2m-1}{2M}$ (see Fig. 2). In scenarios where only samples from the distributions are available, $x_m \sim p_X$ and $y_m \sim p_Y$, the empirical distributions can be estimated as $p_X = \frac{1}{M} \sum_{m=1}^M \delta_{x_m}$ and $p_Y = \frac{1}{M} \sum_{m=1}^M \delta_{y_m}$, where δ_{x_m} is the Dirac delta function centered at x_m . Therefore the corresponding empirical cumulative distribution of p_X is $P_X(t) = \frac{1}{M} \sum_{m=1}^M u(t - x_m)$ where $u(\cdot)$ is the step function (P_Y is defined similarly). Sorting x_m s in an ascending order, such that $x_{i[m]} \leq x_{i[m+1]}$ and where $i[m]$ is the index of the sorted x_m s, it is straightforward to confirm that $P_X^{-1}(\tau_m) = x_{i[m]}$ (see Fig. 2 for a visual confirmation). Therefore, the Wasserstein distance can be approximated by first sorting x_m s and y_m s and then calculating:

$$W_c(p_X, p_Y) = \frac{1}{M} \sum_{m=1}^M c(x_{i[m]}, y_{j[m]}) \quad (14)$$

Eq. (14) turns the problem of calculating the Wasserstein distance for two one-dimensional probability densities from their samples into a sorting problem that can be solved efficiently ($\mathcal{O}(M)$ best case and $\mathcal{O}(M \log(M))$ worst case).

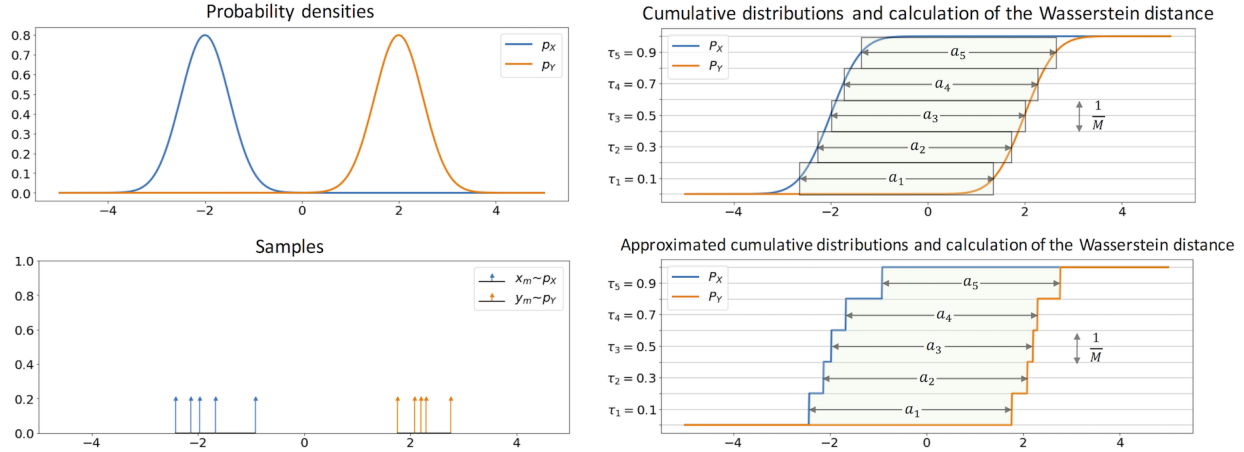


Fig. 2: The Wasserstein distance for one-dimensional probability distributions p_X and p_Y (top left) is calculated based on Eq. (9). For a numerical implementation, the integral in Eq. (9) is substituted with $\frac{1}{M} \sum_{m=1}^M a_m$ where, $a_m = c(P_X^{-1}(\tau_m), P_Y^{-1}(\tau_m))$ (top right). When only samples from the distributions are available $x_n \sim p_X$ and $y_n \sim Y$ (bottom left), the Wasserstein distance is approximated by sorting x_m s and y_m s and letting $a_m = c(x_{i[m]}, y_{j[m]})$, where $i[m]$ and $j[m]$ are the sorted indices (bottom right).

B. Slicing empirical distributions

In scenarios where only samples from the d -dimensional distribution, p_X , are available, $x_m \sim p_X$, the empirical distribution can be estimated as $p_X = \frac{1}{M} \sum_{m=1}^M \delta_{x_m}$. Following Eq. (10) it is straightforward to show that the marginal distributions (i.e. slices) of the empirical distribution, p_X , are obtained from:

$$\mathcal{R}p_X(t, \theta) = \frac{1}{M} \sum_{m=1}^M \delta(t - x_m \cdot \theta), \quad \forall \theta \in \mathbb{S}^{d-1}, \text{ and } \forall t \in \mathbb{R} \quad (15)$$

see the supplementary material for a proof.

C. Minimizing sliced-Wasserstein via random slicing

Minimizing the sliced-Wasserstein distance (i.e. as in the second term of Eq. 13) requires an integration over the unit sphere in \mathbb{R}^d , i.e., \mathbb{S}^{d-1} . In practice, this integration is substituted by a summation over a finite set $\Theta \subset \mathbb{S}^{d-1}$,

$$\min_{\phi} SW_c(p_Z, q_Z) \approx \min_{\phi} \frac{1}{|\Theta|} \sum_{\theta_l \in \Theta} W_c(\mathcal{R}p_Z(\cdot; \theta_l), \mathcal{R}q_Z(\cdot; \theta_l))$$

Note that $SW_c(p_Z, q_Z) = \mathbb{E}_{\mathbb{S}^{d-1}}(W_c(\mathcal{R}p_Z(\cdot; \theta), \mathcal{R}q_Z(\cdot; \theta)))$. Moreover, the global minimum for $SW_c(p_Z, q_Z)$ is also a global minimum for each $W_c(\mathcal{R}p_Z(\cdot; \theta_l), \mathcal{R}q_Z(\cdot; \theta_l))$. A fine sampling of \mathbb{S}^{d-1} is required for a good approximation of $SW_c(p_Z, q_Z)$. Such sampling, however, becomes prohibitively expensive as the dimension of the embedding space grows. Alternatively, following the approach presented by Rabin and Peyré [28], and later by Bonneel et al. [3] and subsequently by Kolouri et al. [19], we utilize random samples of \mathbb{S}^{d-1} at each minimization step to approximate the sliced-Wasserstein distance. Intuitively, if p_Z and q_Z are similar, then their projections with respect to any finite subset of \mathbb{S}^{d-1} would also be similar. This leads to a stochastic gradient descent scheme where in addition to the random sampling of the input data, we also random sample the projection angles from \mathbb{S}^{d-1} .

Algorithm 1 Sliced-Wasserstein Auto-Encoder (SWAE)

Require: Regularization coefficient λ , and number of random projections, L .

Initialize the parameters of the encoder, ϕ , and decoder, ψ

while ϕ and ψ have not converged **do**

 Sample $\{x_1, \dots, x_M\}$ from training set (i.e. p_X)

 Sample $\{\tilde{z}_1, \dots, \tilde{z}_M\}$ from q_Z

 Sample $\{\theta_1, \dots, \theta_L\}$ from \mathbb{S}^{K-1}

 Sort $\theta_l \cdot \tilde{z}_M$ such that $\theta_l \cdot \tilde{z}_{i[m]} \leq \theta_l \cdot \tilde{z}_{i[m+1]}$

 Sort $\theta_l \cdot \phi(x_m)$ such that $\theta_l \cdot \phi(x_{j[m]}) \leq \theta_l \cdot \phi(x_{j[m+1]})$

 Update ϕ and ψ by descending

$$\sum_{m=1}^M c(x_m, \psi(\phi(x_m))) + \lambda \sum_{l=1}^L \sum_{m=1}^M c(\theta_l \cdot \tilde{z}_{i[m]}, \theta_l \cdot \phi(x_{j[m]}))$$

end while

D. Putting it all together

To optimize the proposed SWAE objective function in Eq. (13) we use a stochastic gradient descent scheme as described here. In each iteration, let $\{x_m \sim p_X\}_{m=1}^M$ and $\{\tilde{z}_m \sim q_Z\}_{m=1}^M$ be i.i.d random samples from the input data and the predefined distribution, q_Z , correspondingly. Let $\{\theta_l\}_{l=1}^L$ be randomly sampled from a uniform distribution on \mathbb{S}^{d-1} . Then using the numerical approximations described in this section, the loss function in Eq. (13) can be rewritten as:

$$\mathcal{L}(\phi, \psi) = \frac{1}{M} \sum_{m=1}^M c(x_m, \psi(\phi(x_m))) + \frac{\lambda}{LM} \sum_{l=1}^L \sum_{m=1}^M c(\theta_l \cdot \tilde{z}_{i[m]}, \theta_l \cdot \phi(x_{j[m]})) \quad (16)$$

where $i[m]$ and $j[m]$ are the indices of sorted $\theta_l \cdot \tilde{z}_m$ s and $\theta_l \cdot \phi(x_m)$ with respect to m , correspondingly. The steps of our proposed method are presented in Algorithm 1. It is worth pointing out that sorting is by itself an optimization problem (which can be solved very efficiently), and therefore the sorting followed by the gradient descent update on ϕ and ψ is in essence a min-max problem, which is being solved in an alternating fashion.

V. EXPERIMENTS

Here we show the results of SWAE for two mid-size image datasets, namely the MNIST dataset [20], and the CelebFaces Attributes Dataset (CelebA) [22]. For the encoder and the decoder we used mirrored classic deep convolutional neural networks with 2D average poolings and leaky rectified linear units (Leaky-ReLU) as the activation functions. The implementation details are included in the Supplementary material.

For the MNIST dataset, we designed a deep convolutional encoder that embeds the hand-written digits into a two-dimensional embedding space (for visualization). To demonstrate the capability of SWAE on matching distributions p_Z and q_Z in the embedding/encoder space we chose four different q_Z s, namely the ring distribution, the uniform distribution, a circle distribution, and a bowl distribution. Figure 3 shows the results of our experiment on the MNIST dataset. The left column shows samples from q_Z , the middle column shows $\phi(x_n)$ s for the trained ϕ and the color represent the labels (note that the labels were only used for visualization). Finally, the right column depicts a 25×25 grid in $[-1, 1]^2$ through the trained decoder ψ . As can be seen, the embedding/encoder space closely follows the predefined q_Z , while the space remains decodable. The implementation details are included in the supplementary material.

The CelebA face dataset contains a higher degree of variations compared to the MNIST dataset and therefore a two-dimensional embedding space does not suffice to capture the variations in this dataset. Therefore, while the SWAE loss function still goes down and the network achieves a good match between p_Z and q_Z the decoder is unable to match p_X and p_Y . Therefore, a higher-dimensional embedding/encoder space is needed. In our experiments for this dataset we chose a ($K = 128$)-dimensional embedding space. Figure 4 demonstrates the outputs of trained SWAEs with $K = 2$ and $K = 128$ for sample input images. The input images were resized to 64×64 and then fed to our autoencoder structure.

For CelebA dataset we set q_Z to be a ($K = 128$)-dimensional uniform distribution and trained our SWAE on the CelebA dataset. Given the convex nature of q_Z , any linear combination of the encoded faces should also result in a new face. Having that in mind, we ran two experiments in the embedding space to check that in fact the embedding space satisfies this convexity assumption. First we calculated linear interpolations of sampled pairs of faces in the embedding space and fed the interpolations to the decoder network to visualize the corresponding faces. Figure 5, left column, shows the interpolation results for random pairs of encoded faces. It is clear that the interpolations remain faithful as expected from a uniform q_Z . Finally, we performed Principle Component Analysis (PCA) of the encoded faces and visualized the faces corresponding to these principle components via ψ . The PCA components are shown on the left column of Figure 5. Various interesting modes including, hair color, skin color, gender, pose, etc. can be observed in the PC components.

VI. CONCLUSIONS

We introduced Sliced Wasserstein Autoencoders (SWAE), which enable one to shape the distribution of the encoded samples to any samplable distribution. We theoretically showed that utilizing the sliced Wasserstein distance as a dissimilarity measure between the distribution of the encoded samples and a predefined distribution ameliorates the need for training an adversarial network in the embedding space. In addition, we provided a simple and efficient numerical scheme for this problem, which only relies on few inner products and sorting operations in each SGD iteration. We further demonstrated the capability of our method on two mid-size image datasets, namely the MNIST dataset and the CelebA face dataset and showed results comparable to the techniques that rely on additional adversarial trainings. Our implementation is publicly available ¹.

ACKNOWLEDGMENTS

This work was partially supported by NSF (CCF 1421502). The authors would like to thank Drs. Dejan Slepčev, and Heiko Hoffmann for their invaluable inputs and many hours of constructive conversations.

REFERENCES

- [1] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein GAN. *arXiv preprint arXiv:1701.07875*, 2017.
- [2] David Berthelot, Tom Schumm, and Luke Metz. Began: Boundary equilibrium generative adversarial networks. *arXiv preprint arXiv:1703.10717*, 2017.
- [3] Nicolas Bonneel, Julien Rabin, Gabriel Peyré, and Hanspeter Pfister. Sliced and Radon Wasserstein barycenters of measures. *Journal of Mathematical Imaging and Vision*, 51(1):22–45, 2015.

¹<https://github.com/skolouri/swae>

- [4] Nicolas Bonnotte. *Unidimensional and evolution methods for optimal transportation*. PhD thesis, Paris 11, 2013.
- [5] Olivier Bousquet, Sylvain Gelly, Ilya Tolstikhin, Carl-Johann Simon-Gabriel, and Bernhard Schoelkopf. From optimal transport to generative modeling: the VEGAN cookbook. *arXiv preprint arXiv:1705.07642*, 2017.
- [6] Mathieu Carriere, Marco Cuturi, and Steve Oudot. Sliced wasserstein kernel for persistence diagrams. *arXiv preprint arXiv:1706.03358*, 2017.
- [7] Antonia Creswell and Anil Anthony Bharath. Adversarial training for sketch retrieval. In *European Conference on Computer Vision*, pages 798–809. Springer, 2016.
- [8] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in neural information processing systems*, pages 2292–2300, 2013.
- [9] Charlie Frogner, Chiyuan Zhang, Hossein Mobahi, Mauricio Araya, and Tomaso A Poggio. Learning with a wasserstein loss. In *Advances in Neural Information Processing Systems*, pages 2053–2061, 2015.
- [10] Daniel T Gillespie. A theorem for physicists in the theory of random variables. *American Journal of Physics*, 51(6):520–533, 1983.
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [12] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems*, pages 5769–5779, 2017.
- [13] Sigurdur Helgason. The radon transform on \mathbb{R}^n . In *Integral Geometry and Radon Transforms*, pages 1–62. Springer, 2011.
- [14] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. *arXiv preprint*, 2017.
- [15] Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [16] Soheil Kolouri and Gustavo K Rohde. Transport-based single frame super resolution of very low resolution face images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4876–4884, 2015.
- [17] Soheil Kolouri, Yang Zou, and Gustavo K Rohde. Sliced Wasserstein kernels for probability distributions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5258–5267, 2016.
- [18] Soheil Kolouri, Se Rim Park, Matthew Thorpe, Dejan Slepcev, and Gustavo K Rohde. Optimal mass transport: Signal processing and machine-learning applications. *IEEE Signal Processing Magazine*, 34(4):43–59, 2017.
- [19] Soheil Kolouri, Gustavo K Rohde, and Heiko Hoffman. Sliced Wasserstein distance for learning Gaussian mixture models. *arXiv preprint arXiv:1711.05376*, 2017.
- [20] Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- [21] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. *arXiv preprint*, 2016.
- [22] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2015.
- [23] Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, Ian Goodfellow, and Brendan Frey.

- Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*, 2015.
- [24] Lars Mescheder, Sebastian Nowozin, and Andreas Geiger. Adversarial variational Bayes: Unifying variational autoencoders and generative adversarial networks. *arXiv preprint arXiv:1701.04722*, 2017.
 - [25] Zak Murez, Soheil Kolouri, David Kriegman, Ravi Ramamoorthi, and Kyungnam Kim. Image to image translation for domain adaptation. *arXiv preprint arXiv:1712.00479*, 2017.
 - [26] Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f-GAN: Training generative neural samplers using variational divergence minimization. In *Advances in Neural Information Processing Systems*, pages 271–279, 2016.
 - [27] Adam M Oberman and Yuanlong Ruan. An efficient linear programming method for optimal transportation. *arXiv preprint arXiv:1509.03668*, 2015.
 - [28] Julien Rabin and Gabriel Peyré. Wasserstein regularization of imaging problem. In *Image Processing (ICIP), 2011 18th IEEE International Conference on*, pages 1541–1544. IEEE, 2011.
 - [29] Julien Rabin, Gabriel Peyré, Julie Delon, and Marc Bernot. Wasserstein barycenter and its application to texture mixing. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 435–446. Springer, 2011.
 - [30] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
 - [31] Filippo Santambrogio. Optimal transport for applied mathematicians. *Birkäuser, NY*, pages 99–102, 2015.
 - [32] Justin Solomon, Fernando De Goes, Gabriel Peyré, Marco Cuturi, Adrian Butscher, Andy Nguyen, Tao Du, and Leonidas Guibas. Convolutional wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Transactions on Graphics (TOG)*, 34(4): 66, 2015.
 - [33] Ilya Tolstikhin, Olivier Bousquet, Sylvain Gelly, and Bernhard Schoelkopf. Wasserstein auto-encoders. *arXiv preprint arXiv:1711.01558*, 2017.
 - [34] Cédric Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.
 - [35] Raymond A Yeh, Chen Chen, Teck Yian Lim, Alexander G Schwing, Mark Hasegawa-Johnson, and Minh N Do. Semantic image inpainting with deep generative models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5485–5493, 2017.

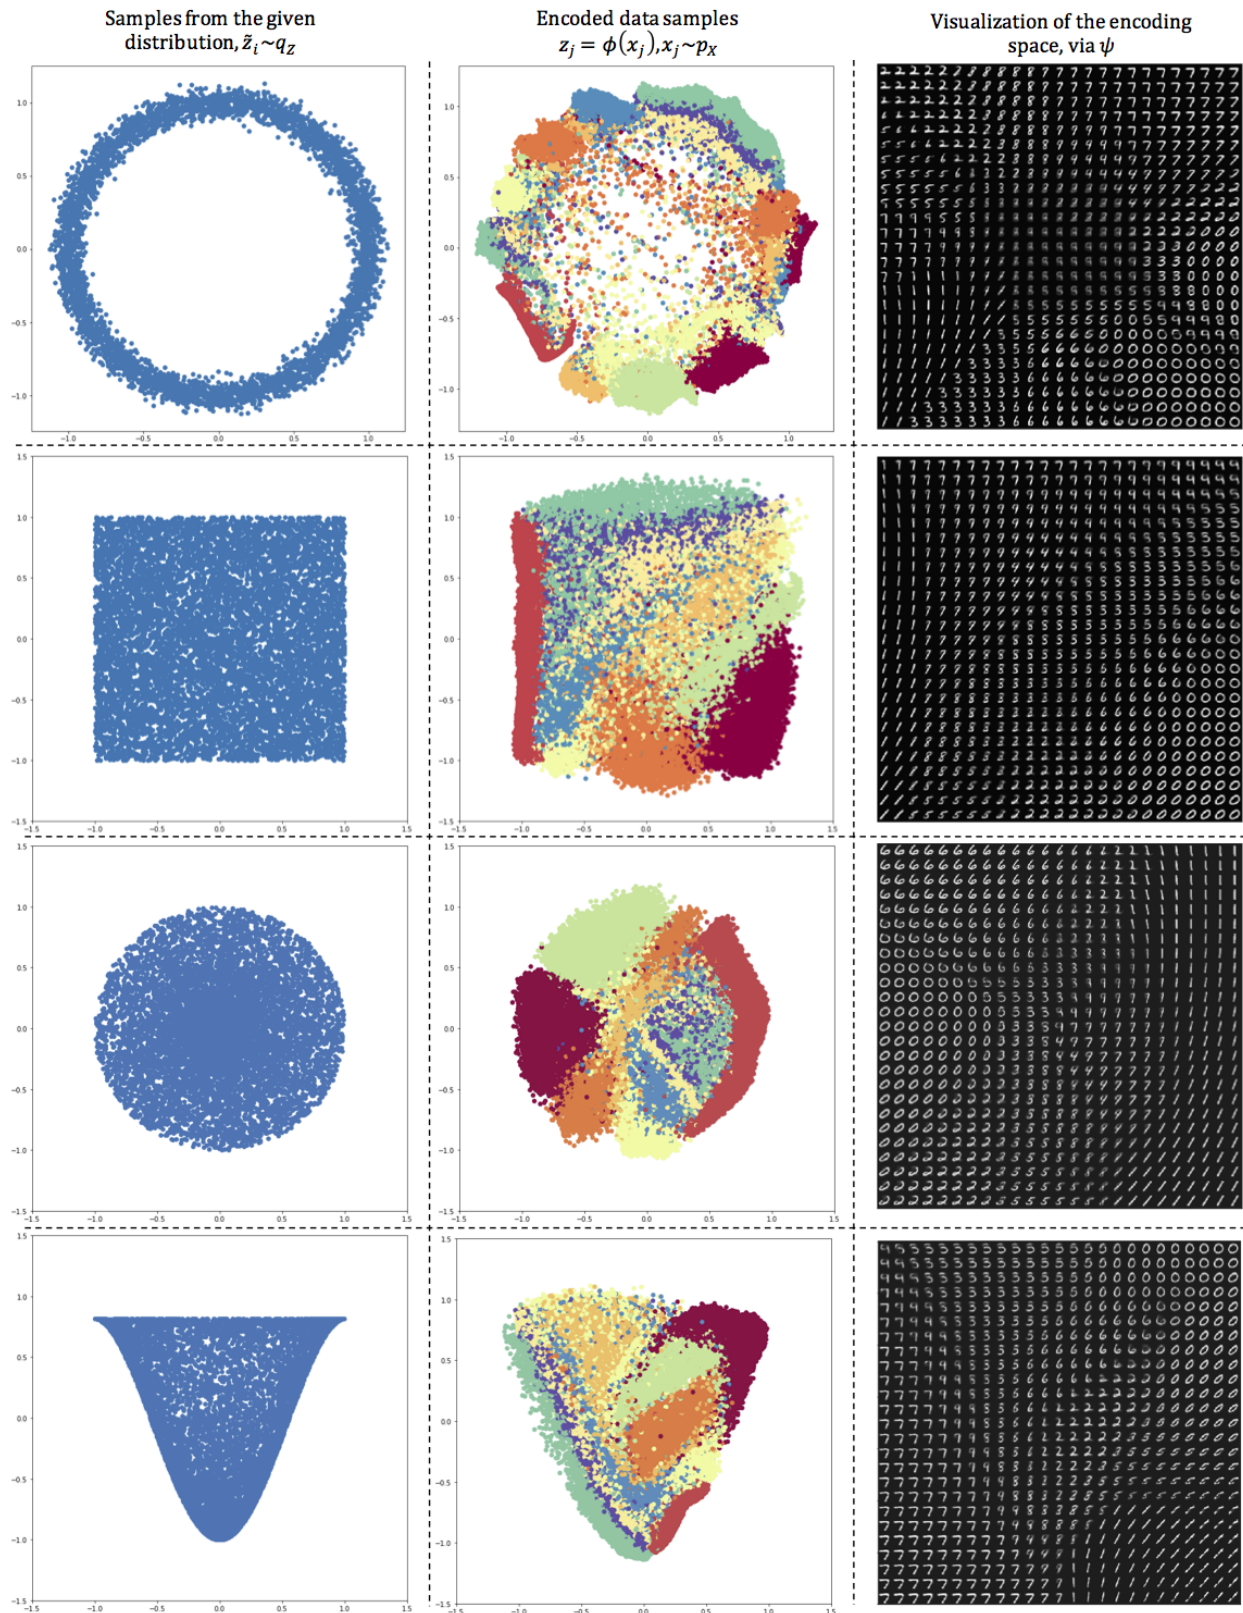


Fig. 3: The results of SWAE on the MNIST dataset for three different distributions as q_Z , namely the ring distribution, the uniform distribution, and the circle distribution. Note that the far right visualization is showing the decoding of a 25×25 grid in $[-1, 1]^2$ (in the encoding space).




Fig. 4: Trained SWAE outputs for sample input images with different embedding spaces of size $K = 2$ and $K = 128$.

Sample linear interpolations in the encoding space



$$\psi(\alpha\phi(x_1) + (1 - \alpha)\phi(x_2)), \quad \alpha \in [0,1]$$

PCA modes calculated in the encoding space



$$-3\sigma \quad -2\sigma \quad -\sigma \quad 0 \quad \sigma \quad 2\sigma \quad 3\sigma$$

Fig. 5: The results of SWAE on the CelebA face dataset with a 128-dimensional uniform distribution as q_Z . Linear interpolation in the encoding space for random samples (on the right) and the first 10 PCA components calculated in the encoding space.

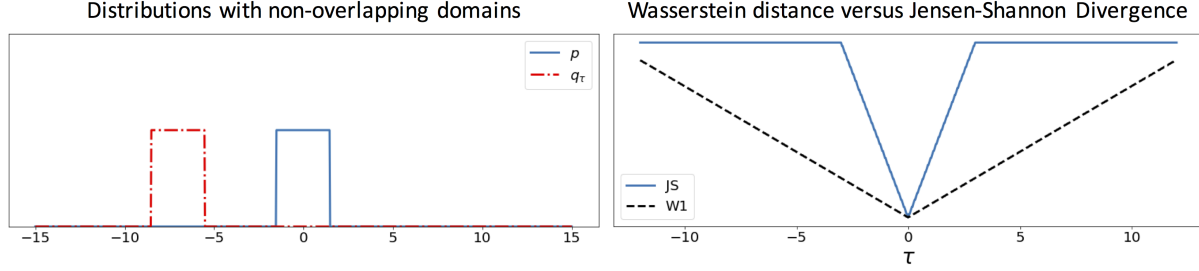


Fig. 6: These plots show $W_1(p, q_\tau)$ and $JS(p, q_\tau)$ where p is a uniform distribution around zero and $q_\tau(x) = p(x - \tau)$. It is clear that JS divergence does not provide a usable gradient when distributions are supported on non-overlapping domains.

SUPPLEMENTARY MATERIAL

Comparison of different distances

Following the example by Arjovsky et al. [1] and later Kolouri et al. [19] here we show a simple example comparing the Jensen-Shannon divergence with the Wasserstein distance. First note that the Jensen-Shannon divergence is defined as,

$$JS(p, q) = KL(p, \frac{p+q}{2}) + KL(q, \frac{p+q}{2})$$

where $KL(p, q) = \int_X p(x) \log(\frac{p(x)}{q(x)}) dx$ is the Kullback-Leibler divergence. Now consider the following densities, $p(x)$ be a uniform distribution around zero and let $q_\tau(x) = p(x - \tau)$ be a shifted version of the p . Figure 6 show $W_1(p, q_\tau)$ and $JS(p, q_\tau)$ as a function of τ . As can be seen the JS divergence fails to provide a useful gradient when the distributions are supported on non-overlapping domains.

Log-likelihood

To maximize (minimize) the similarity (dissimilarity) between p_Z and q_Z , we can write :

$$\begin{aligned} \operatorname{argmax}_\phi \int_Z p_Z(z) \log(q_Z(z)) dz &= \int_Z \int_X p_X(x) \delta(z - \phi(x)) \log(q_Z(z)) dx dz \\ &= \int_X p_X(x) \log(q_Z(\phi(x))) dx \end{aligned}$$

where we replaced p_Z with Eq. (1). Furthermore, it is straightforward to show:

$$\begin{aligned} \operatorname{argmax}_\phi \int_Z p_Z(z) \log(q_Z(z)) dz &= \operatorname{argmax}_\phi \int_Z p_Z(z) \log\left(\frac{q_Z(z)}{p_Z(z)}\right) dz \\ &= \operatorname{argmin}_\phi D_{KL}(p_Z, q_Z) \end{aligned}$$

Slicing empirical distributions

Here we calculate a Radon slice of the empirical distribution $p_X(x) = \frac{1}{M} \sum_{m=1}^M \delta(x - x_m)$ with respect to $\theta \in \mathbb{S}^{d-1}$. Using the definition of the Radon transform in Eq. (10) and RVT in Eq. (1)




Fig. 7: Different runs of SWAE to embed a 3D nonlinear manifold into a 2D uniform distribution.

we have:

$$\begin{aligned}
 \mathcal{R}p_X(t, \theta) &= \int_X p_X(x) \delta(t - \theta \cdot x) dx \\
 &= \frac{1}{M} \sum_{m=1}^M \int_X \delta(x - x_m) \delta(t - \theta \cdot x) dx \\
 &= \frac{1}{M} \sum_{m=1}^M \delta(t - \theta \cdot x_m)
 \end{aligned}$$

Simple manifold learning experiment

Figure 7 demonstrates the results of SWAE with random initializations to embed a 2D manifold in \mathbb{R}^3 to a 2D uniform distribution.

The implementation details of our algorithm

The following text walks you through the implementation of our Sliced Wasserstein Autoencoders (SWAE).

To run this notebook you'll require the following packages:

- Numpy
- Matplotlib
- tensorflow
- Keras

```

In [1]: import numpy as np
import keras.utils
from keras.layers import Input, Dense, Flatten, Activation
from keras.models import load_model, Model
from keras.layers import Conv2D, UpSampling2D, AveragePooling2D
from keras.layers import LeakyReLU, Reshape
from keras.preprocessing.image import ImageDataGenerator
from keras.optimizers import RMSprop
from keras.datasets import mnist
from keras.models import save_model
from keras import backend as K

```

```

import tensorflow as tf
import matplotlib.pyplot as plt
from IPython import display
import time

```

Using TensorFlow backend.

A. Define three helper function

- generateTheta(L,dim) -> Generates L random sampels from S^{dim-1}
- generateZ(batchsize,endum) -> Generates 'batchsize' samples 'endum' dimensional samples from q_Z
- stitchImages(I,axis=0) -> Helps us with visualization

```

In [2]: def generateTheta(L,endum):
        theta_=np.random.normal(size=(L,endum))
        for l in range(L):
            theta_[l,:]=theta_[l,:]/np.sqrt(np.sum(theta_[l,:]**2))
        return theta_
def generateZ(batchsize,endum):
    z_=2*(np.random.uniform(size=(batchsize,endum))-0.5)
    return z_
def stitchImages(I,axis=0):
    n,N,M,K=I.shape
    if axis==0:
        img=np.zeros((N*n,M,K))
        for i in range(n):
            img[i*N:(i+1)*N,:]=I[i,:,:,:]
    else:
        img=np.zeros((N,M*n,K))
        for i in range(n):
            img[:,i*M:(i+1)*M,:]=I[i,:,:,:]
    return img

```

B. Defining the Encoder/Decoder as Keras graphs

```

In [3]: img=Input((28,28,1)) #Input image
        interdim=128 # This is the dimension of intermediate latent
                    # (variable after convolution and before embedding)
        endim=2 # Dimension of the embedding space
        embedd=Input((endum,)) #Keras input to Decoder
        depth=16 # This is a design parameter.
        L=50 # Number of random projections
        batchsize=500

```

1) Define Encoder:

```

In [4]: x=Conv2D(depth*1, (3, 3), padding='same')(img)
        x=LeakyReLU(alpha=0.2)(x)
        # x=BatchNormalization(momentum=0.8)(x)
        x=Conv2D(depth*1, (3, 3), padding='same')(x)
        x=LeakyReLU(alpha=0.2)(x)
        # x=BatchNormalization(momentum=0.8)(x)
        x=AveragePooling2D((2, 2), padding='same')(x)
        x=Conv2D(depth*2, (3, 3), padding='same')(x)
        x=LeakyReLU(alpha=0.2)(x)
        # x=BatchNormalization(momentum=0.8)(x)
        x=Conv2D(depth*2, (3, 3), padding='same')(x)

```



```

x=LeakyReLU(alpha=0.2)(x)
# x=BatchNormalization(momentum=0.8)(x)
x=AveragePooling2D((2, 2), padding='same')(x)
x=Conv2D(depth*4, (3, 3), padding='same')(x)
x=LeakyReLU(alpha=0.2)(x)
# x=BatchNormalization(momentum=0.8)(x)
x=Conv2D(depth*4, (3, 3), padding='same')(x)
x=LeakyReLU(alpha=0.2)(x)
# x=BatchNormalization(momentum=0.8)(x)
x=AveragePooling2D((2, 2), padding='same')(x)
x=Flatten()(x)
x=Dense(interdim,activation='relu')(x)
encoded=Dense(endim)(x)

encoder=Model(inputs=[img],outputs=[encoded])
encoder.summary()

```

Layer (type)	Output Shape	Param #
input_1 (InputLayer)	(None, 28, 28, 1)	0
conv2d_1 (Conv2D)	(None, 28, 28, 16)	160
leaky_re_lu_1 (LeakyReLU)	(None, 28, 28, 16)	0
conv2d_2 (Conv2D)	(None, 28, 28, 16)	2320
leaky_re_lu_2 (LeakyReLU)	(None, 28, 28, 16)	0
average_pooling2d_1 (Average	(None, 14, 14, 16)	0
conv2d_3 (Conv2D)	(None, 14, 14, 32)	4640
leaky_re_lu_3 (LeakyReLU)	(None, 14, 14, 32)	0
conv2d_4 (Conv2D)	(None, 14, 14, 32)	9248
leaky_re_lu_4 (LeakyReLU)	(None, 14, 14, 32)	0
average_pooling2d_2 (Average	(None, 7, 7, 32)	0
conv2d_5 (Conv2D)	(None, 7, 7, 64)	18496
leaky_re_lu_5 (LeakyReLU)	(None, 7, 7, 64)	0
conv2d_6 (Conv2D)	(None, 7, 7, 64)	36928
leaky_re_lu_6 (LeakyReLU)	(None, 7, 7, 64)	0
average_pooling2d_3 (Average	(None, 4, 4, 64)	0
flatten_1 (Flatten)	(None, 1024)	0
dense_1 (Dense)	(None, 128)	131200

dense_2 (Dense)	(None, 2)	258
=====		
Total params: 203,250		
Trainable params: 203,250		
Non-trainable params: 0		

2) Define Decoder:

```
In [5]: x=Dense(interdim)(embedd)
x=Dense(depth*64,activation='relu')(x)
# x=BatchNormalization(momentum=0.8)(x)
x=Reshape((4,4,4*depth))(x)
x=UpSampling2D((2,2))(x)
x=Conv2D(depth*4,(3,3),padding='same')(x)
x=LeakyReLU(alpha=0.2)(x)
# x=BatchNormalization(momentum=0.8)(x)
x=Conv2D(depth*4,(3,3),padding='same')(x)
x=LeakyReLU(alpha=0.2)(x)
x=UpSampling2D((2,2))(x)
x=Conv2D(depth*4,(3,3),padding='valid')(x)
x=LeakyReLU(alpha=0.2)(x)
# x=BatchNormalization(momentum=0.8)(x)
x=Conv2D(depth*4,(3,3),padding='same')(x)
x=LeakyReLU(alpha=0.2)(x)
x=UpSampling2D((2,2))(x)
x=Conv2D(depth*2,(3,3),padding='same')(x)
x=LeakyReLU(alpha=0.2)(x)
# x=BatchNormalization(momentum=0.8)(x)
x=Conv2D(depth*2,(3,3),padding='same')(x)
x=LeakyReLU(alpha=0.2)(x)
# x=BatchNormalization(momentum=0.8)(x)
# x=BatchNormalization(momentum=0.8)(x)
decoded=Conv2D(1,(3,3),padding='same',activation='sigmoid')(x)

decoder=Model(inputs=[embedd],outputs=[decoded])
decoder.summary()
```

Layer (type)	Output Shape	Param #
input_2 (InputLayer)	(None, 2)	0
dense_3 (Dense)	(None, 128)	384
dense_4 (Dense)	(None, 1024)	132096
reshape_1 (Reshape)	(None, 4, 4, 64)	0
up_sampling2d_1 (UpSampling2D)	(None, 8, 8, 64)	0
conv2d_7 (Conv2D)	(None, 8, 8, 64)	36928
leaky_re_lu_7 (LeakyReLU)	(None, 8, 8, 64)	0

conv2d_8 (Conv2D)	(None, 8, 8, 64)	36928
leaky_re_lu_8 (LeakyReLU)	(None, 8, 8, 64)	0
up_sampling2d_2 (UpSampling2D)	(None, 16, 16, 64)	0
conv2d_9 (Conv2D)	(None, 14, 14, 64)	36928
leaky_re_lu_9 (LeakyReLU)	(None, 14, 14, 64)	0
conv2d_10 (Conv2D)	(None, 14, 14, 64)	36928
leaky_re_lu_10 (LeakyReLU)	(None, 14, 14, 64)	0
up_sampling2d_3 (UpSampling2D)	(None, 28, 28, 64)	0
conv2d_11 (Conv2D)	(None, 28, 28, 32)	18464
leaky_re_lu_11 (LeakyReLU)	(None, 28, 28, 32)	0
conv2d_12 (Conv2D)	(None, 28, 28, 32)	9248
leaky_re_lu_12 (LeakyReLU)	(None, 28, 28, 32)	0
conv2d_13 (Conv2D)	(None, 28, 28, 1)	289
=====		
Total params: 308,193		
Trainable params: 308,193		
Non-trainable params: 0		

Here we define Keras variables for θ and sample z s.

```
In [6]: #Define a Keras Variable for \theta_ls
theta=K.variable(generateTheta(L,endum))
#Define a Keras Variable for samples of z
z=K.variable(generateZ(batchsize,endum))
```

Put encoder and decoder together to get the autoencoder

```
In [7]: # Generate the autoencoder by combining encoder and decoder
aencoded=encoder(img)
ae=decoder(aencoded)
autoencoder=Model(inputs=[img],outputs=[ae])
autoencoder.summary()
```

Layer (type)	Output Shape	Param #
input_1 (InputLayer)	(None, 28, 28, 1)	0
model_1 (Model)	(None, 2)	203250
model_2 (Model)	(None, 28, 28, 1)	308193

```
=====
Total params: 511,443
Trainable params: 511,443
Non-trainable params: 0
=====
```

```
In [8]: # Let projae be the projection of the encoded samples
projae=K.dot(aencoded,K.transpose(theta))
# Let projz be the projection of the $q_Z$ samples
projz=K.dot(z,K.transpose(theta))
# Calculate the Sliced Wasserstein distance by sorting
# the projections and calculating the L2 distance between
W2=(tf.nn.top_k(tf.transpose(projae),k=batchsize).values-
    tf.nn.top_k(tf.transpose(projz),k=batchsize).values)**2


In [9]: crossEntropyLoss= (1.0)*K.mean(K.binary_crossentropy(K.flatten(img),
                                                                K.flatten(ae)))
L1Loss= (1.0)*K.mean(K.abs(K.flatten(img)-K.flatten(ae)))
W2Loss= (10.0)*K.mean(W2)
# I have a combination of L1 and Cross-Entropy loss
# for the first term and then and W2 for the second term
vae_Loss=crossEntropyLoss+L1Loss+W2Loss
autoencoder.add_loss(vae_Loss) # Add the custom loss to the model

In [10]: #Compile the model
autoencoder.compile(optimizer='rmsprop',loss='')

3) Load the MNIST dataset:


In [11]: (x_train,y_train),(x_test,_)=mnist.load_data()
x_train=np.expand_dims(x_train.astype('float32')/255.,3)

In [12]: plt.imshow(np.squeeze(x_train[0,...]))
plt.show()
```



C. Optimize the Loss

```
In [13]: loss=[]
        for epoch in range(20):
            ind=np.random.permutation(x_train.shape[0])
            for i in range(int(x_train.shape[0]/batchsize)):
                Xtr=x_train[ind[i*batchsize:(i+1)*batchsize],...]
                theta_=generateTheta(L, endim)
                z_=generateZ(batchsize, endim)
                K.set_value(z, z_)
                K.set_value(theta, theta_)
                loss.append(autoencoder.train_on_batch(x=Xtr, y=None))
            plt.plot(np.asarray(loss))
            display.clear_output(wait=True)
            display.display(plt.gcf())
            time.sleep(1e-3)
```




D. Encode and decode x_{train}

```
In [15]: # Test autoencoder
en=encoder.predict(x_train) # Encode the images
dec=decoder.predict(en) # Decode the encodings


In [16]: # Sanity check for the autoencoder
# Note that we can use a more complex autoencoder that results
# in better reconstructions. Also the autoencoders used in the
# literature often use a much larger latent space (we are using only 2!)
fig, [ax1, ax2]=plt.subplots(2, 1, figsize=(100, 10))
I_temp=(stitchImages(x_train[:10, ...], 1)*255.0).astype('uint8')
Idec_temp=(stitchImages(dec[:10, ...], 1)*255.0).astype('uint8')
ax1.imshow(np.squeeze(I_temp))
```

```
ax1.set_xticks([])
ax1.set_yticks([])
ax2.imshow(np.squeeze(Idec_temp))
ax2.set_xticks([])
ax2.set_yticks([])
plt.show()
```



E. Visualize the encoding space

```
In [17]: # Distribution of the encoded samples
plt.figure(figsize=(10,10))
plt.scatter(en[:,0],-en[:,1],c=10*y_train, cmap=plt.cm.Spectral)
plt.xlim([-1.5,1.5])
plt.ylim([-1.5,1.5])
plt.show()
```




1) *Sample a grid in the encoding space and decode it to visualize this space:*

```
In [18]: #Sample the latent variable on a Nsample x Nsample grid
Nsample=25
hiddenv=np.meshgrid(np.linspace(-1,1,Nsample),np.linspace(-1,1,Nsample))
v=np.concatenate((np.expand_dims(hiddenv[0].flatten(),1),
                  np.expand_dims(hiddenv[1].flatten(),1)),1)
# Decode the grid
decodeimg=np.squeeze(decoder.predict(v))

In [19]: #Visualize the grid
count=0
img=np.zeros((Nsample*28,Nsample*28))
for i in range(Nsample):
    for j in range(Nsample):
        img[i*28:(i+1)*28,j*28:(j+1)*28]=decodeimg[count,...]
        count+=1

In [20]: fig=plt.figure(figsize=(10,10))
plt.imshow(img)
plt.show()
```



```
In [21]: #Visualize the z samples
plt.figure(figsize=(10,10))
Z=generateZ(10000,2)
plt.scatter(Z[:,0],Z[:,1])
plt.xlim([-1.5,1.5])
plt.ylim([-1.5,1.5])
plt.show()
```