

Manifold-Valued Image Generation with Wasserstein Generative Adversarial Nets

Zhiwu Huang,[†] Jiqing Wu,[†] Luc Van Gool^{†‡}

[†]Computer Vision Lab, ETH Zurich, Switzerland [‡]VISICS, KU Leuven, Belgium
{zhiwu.huang, jwu, vangool}@vision.ee.ethz.ch

Abstract

Generative modeling over natural images is one of the most fundamental machine learning problems. However, few modern generative models, including Wasserstein Generative Adversarial Nets (WGANs), are studied on manifold-valued images that are frequently encountered in real-world applications. To fill the gap, this paper first formulates the problem of generating manifold-valued images and exploits three typical instances: hue-saturation-value (HSV) color image generation, chromaticity-brightness (CB) color image generation, and diffusion-tensor (DT) image generation. For the proposed generative modeling problem, we then introduce a theorem of optimal transport to derive a new Wasserstein distance of data distributions on complete manifolds, enabling us to achieve a tractable objective under the WGAN framework. In addition, we recommend three benchmark datasets that are CIFAR-10 HSV/CB color images, ImageNet HSV/CB color images, UCL DT image datasets. On the three datasets, we experimentally demonstrate the proposed manifold-aware WGAN model can generate more plausible manifold-valued images than its competitors.

Introduction

Building generative models of natural images has been a fundamental problem in machine learning. Such generative models are expected to estimate the probability distributions over the natural images. One of the most striking techniques is the family of generative adversarial networks (GANs) such as (Goodfellow et al. 2014; Radford, Metz, and Chintala 2015; Zhao, Mathieu, and LeCun 2016; Mao et al. 2017). The state-of-the-art GANs like (Arjovsky, Chintala, and Bottou 2017; Gulrajani et al. 2017; Wei et al. 2018; Miyato et al. 2018) are good at approximating the distributions of Euclidean-valued (or real-valued) images implicitly by optimizing Wasserstein distance between the distributions of generated images and real images under an adversarial training framework.

While the GAN techniques have made great success for real-valued image generation, they are rarely applied to manifold-valued images that are of much interest in a variety of applications. For example, one of the most popular applications is processing phase-valued images, whose data live on either the Cycle \mathbb{S}^1 or the Sphere \mathbb{S}^2 , in hue-saturation-value (HSV) and chromaticity-brightness (CB)

color spaces. Since such HSV/CB spaces are more adapted to human color perception than the RGB space, many works like (Bergmann et al. 2014; Bergmann and Weinmann 2015; Bansal and Tatu 2015; Bergmann and Weinmann 2016; Bacák et al. 2016; Laus et al. 2017) have studied that the image processing models based on HSV/CB components can surpass the competitors developed in the RGB space. Analogously, for a better understanding of color semantics, producing HSV/CB images in an unsupervised manner would be a good alternative for the regular image generation over the RGB space. Another good application is for diffusion-tensor magnetic resonance imaging (DT-MRI). In many DT-MRI works like (Pennec, Fillard, and Ayache 2006; Arsigny et al. 2007; Hasan et al. 2011; Jayasumana et al. 2013; Priya and Nair 2015), DT images are generally processed on the Riemannian manifold of symmetric positive definite 3×3 matrices SPD(3). In practice, researchers often have a severe lack of DT images for better analysis. Hence, generating photo-realistic DT images would have a high potential to benefit this field.

With the motivation in mind, we focus on generative modeling over manifold-valued images. In general, it proposes new challenges. One of the most important issues is the generalization of the classic distribution distance to manifolds, while the other major challenge is the exploration on the manifold setting of the regular GAN objective. To address the first issue, we leverage the well-studied optimal transport (OT) theory (Fathi and Figalli 2010; Loeper 2011; De Philippis and Figalli 2014; Fitschen, Laus, and Schmitzer 2017) on manifolds to introduce a new Wasserstein distance on complete manifolds. By adopting the proposed Wasserstein distance, we exploit a new model of Wasserstein GANs (Arjovsky, Chintala, and Bottou 2017; Gulrajani et al. 2017; Wei et al. 2018) for manifold-valued image generation. Without loss of generality, we take three typical types of manifold-valued images for study, and suggest the proposed manifold-aware Wasserstein GAN approach to produce plausible samples residing on such manifolds. In summary, the paper gives rise to two original contributions:

- To the best of our knowledge, this paper proposes the novel problem of unsupervised manifold-valued image generation for the first time in the existing literature.
- We generalize the Wasserstein GAN methodology to the

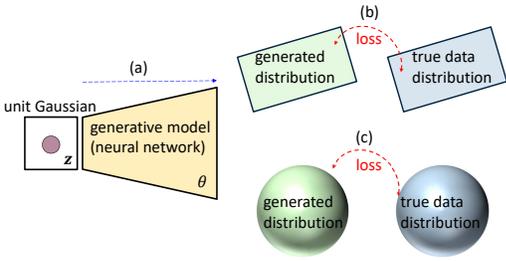


Figure 1: Conceptual illustration of real-valued image generation (a)→(b) where the data lie in Euclidean spaces, and manifold-valued image generation (a)→(c) where the data reside on Riemannian manifolds.

Riemannian manifold setting so as to address the proposed problem of manifold-valued image generation.

Background

Real-valued Image Generation

Natural images are examples of what our visual world looks like and we refer to these as “samples from the true data distribution” \mathbb{P}_r . Typically, this kind of image values are Euclidean-valued (or real-valued). In other words, their geometrical structure can be respected well with using classical Euclidean metric. Accordingly, the image generation problem is concerned with learning the true data probability distribution. The classical solution is to learn a probability density, which may not exist in practice. In contrast, there also exists another typical approach that directly generates samples following a certain distribution \mathbb{P}_g , which approximates the true data distribution. As studied in (Arjovsky, Chintala, and Bottou 2017) it is useful in two ways. First of all, this approach is able to represent distributions confined to a low dimensional manifold, which is easier to estimate. Second, the capability to generate samples easily is often more useful than estimating the numerical value of the density.

Formally, one can consider a dataset of examples $\mathbf{x}_1, \dots, \mathbf{x}_n$ sampled from a real data distribution $\mathbb{P}_r(\mathbf{x})$. As shown in Fig.1 (a)→(b), the blue region shows the part of the image space that, with a high probability consists of real images, and the elements in the space indicate the data points (each is one image). Now, the model also represents a distribution $\mathbb{P}_g(\mathbf{x})$ (green) that is defined implicitly by taking points from a unit Gaussian distribution (red) and mapping them through a neural network named the generative model (yellow). The network is a function with parameters θ , and learning these parameters will learn the generated distribution of images. The goal is then to seek parameters θ that produce a distribution that matches the real data distribution closely. Therefore, the core of the generation process is how to make the distribution of generated data close to the true data distribution, that is how to define the distribution distance. One of the most important measurements is Wasserstein-1 or Earth-Mover distance

$$W(\mathbb{P}_r, \mathbb{P}_g) = \inf_{\gamma \in \Pi(\mathbb{P}_r, \mathbb{P}_g)} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \gamma} [\|\mathbf{x} - \mathbf{y}\|], \quad (1)$$

where $\Pi(\mathbb{P}_r, \mathbb{P}_g)$ denotes the set of all joint distributions $\gamma(\mathbf{x}, \mathbf{y})$ whose marginals are $\mathbb{P}_r, \mathbb{P}_g$ respectively. Intuitively, $\gamma(\mathbf{x}, \mathbf{y})$ indicates how much mass should be transported from \mathbf{x} to \mathbf{y} in order to transform the distributions \mathbb{P}_r into the distribution \mathbb{P}_g . The Wasserstein distance then is the cost of the optimal transport plan.

Wasserstein GANs

For real-valued image generation, the state-of-the-art techniques are Generative Adversarial Networks (GANs) (Goodfellow et al. 2014; Radford, Metz, and Chintala 2015; Zhao, Mathieu, and LeCun 2016; Mao et al. 2017). The standard GAN framework establishes a min-max adversarial game between two competing networks. The generator network G maps a source of noise to the input space. The discriminator network D receives either a generated sample or a true data sample and must distinguish between the two. The generator is trained to fool the discriminator. Theoretically, the original GAN framework actually minimizes Jensen-Shannon divergence between the true data distribution and generated sample distribution.

By contrast, Wasserstein GANs (Arjovsky, Chintala, and Bottou 2017; Gulrajani et al. 2017; Wei et al. 2018; Miyato et al. 2018) studied that minimizing a reasonable approximation of the Wasserstein-1 distance is able to reach the state-of-the-art GANs. To approximate the Wasserstein-1 distance, the original Wasserstein GAN imposed weight clipping constraints on the critic (referred to as the discriminator pre-Wasserstein) such that the optimal map of the discriminator is Lipschitz continuous. However, as proved in (Gulrajani et al. 2017), the set of functions satisfying this constraint is merely a subset of the k -Lipschitz functions for some k which depends on the clipping threshold and the critic architecture and thus inevitably causes some training failures. To address this issue, improved training of Wasserstein GAN (Gulrajani et al. 2017) enables a more stable GAN training by penalizing the norm of the gradient of the critic with respect to its inputs instead of clipping weights. In particular, this gradient penalty is simply added to the basic Wasserstein GAN loss for the following full objective:

$$\min_G \max_D \mathbb{E}_{\mathbf{x} \sim \mathbb{P}_r} [D(\mathbf{x})] - \mathbb{E}_{G(\mathbf{z}) \sim \mathbb{P}_g} [D(G(\mathbf{z}))] + \lambda \mathbb{E}_{\hat{\mathbf{x}} \sim \mathbb{P}_{\hat{\mathbf{x}}}} [(\|\nabla_{\hat{\mathbf{x}}} D(\hat{\mathbf{x}})\|_2 - 1)^2], \quad (2)$$

where \mathbf{z} is random noise, $\hat{\mathbf{x}}$ is random samples following the distribution $\mathbb{P}_{\hat{\mathbf{x}}}$ that is sampled uniformly along straight lines between pairs of points sampled from \mathbb{P}_r and \mathbb{P}_g , $\nabla_{\hat{\mathbf{x}}} D(\hat{\mathbf{x}})$ is the gradient with respect to $\hat{\mathbf{x}}$, $G(\cdot)$, $D(\cdot)$ denotes the functions of generator and discriminator respectively.

Manifold-valued Image Generation

Manifold-valued Images

In this paper, we concentrate on three typical instances of manifold-valued images, which are commonly encountered in computer vision and medical imaging.

HSV Images: Each pixel in the HSV color model can be represented by a triple which specifies hue, saturation and value respectively. As hue value is phase-based, the HSV data

actually live on the product manifold of a Cyclic manifold and vector spaces $\mathcal{H} = \mathbb{S}^1 \times [0, 1]^2$. Since it is known that both the Cyclic manifold \mathbb{S}^1 and the vector space $[0, 1]^2$ are compact Riemannian manifolds, their product $\mathcal{H} = \mathbb{S}^1 \times [0, 1]^2$ is also a compact manifold.

CB Images: Each pixel in the CB color model contains chromaticity and brightness components. Since our focus is on manifold-valued data, the whole paper studies the chromaticity (spherical) component lying on the compact manifold \mathbb{S}^2 for CB images.

DT Images: In diffusion tensor (DT) images, each voxel is represented with a 3×3 tensor, that is symmetric positive definite (SPD) matrix. Hence, the data of DT images reside on the manifold $\text{SPD}(3)$ of SPD matrices, which is known as a convex cone instead of a compact manifold.

Problem Formulation

In analogy to real-valued image generation, the task of manifold-valued image generation is to synthesize samples respecting a certain distribution for learning the distribution of real manifold-valued data. As shown in Fig.1 (a)→(c), the manifold-valued data lie on Riemannian manifolds rather than a Euclidean space. As the definition of manifold-valued data distribution is different of that of real-valued data distribution, it is infeasible to apply the traditional Wasserstein distance Eqn.1 directly to measure the distance of such non-Euclidean data distribution, and the traditional Wasserstein GANs (Arjovsky, Chintala, and Bottou 2017; Gulrajani et al. 2017; Wei et al. 2018; Miyato et al. 2018) are very likely to fail. Accordingly, for the new image generation task, we should consider two critical problems: 1) the definition on the distribution of manifold-valued data, and 2) the generalization of the distribution (Wasserstein) distance to manifolds.

Fortunately, several Riemannian geometry studies (Pennec, Fillard, and Ayache 2006; Arsigny et al. 2006; 2007; Huang et al. 2015) have addressed the first issue well. In particular, they introduce various Riemannian metrics to define probability density functions on the underlying Riemannian manifolds. Formally, let \mathbf{x}, \mathbf{y} be two points of the manifold that we consider as a local reference and $\mathbf{v} = \mathbf{y}\hat{\mathbf{x}}$ a vector of the tangent space $T_{\mathbf{y}}\mathcal{M}$ at the point \mathbf{y} . The smoothly varying family of inner products in each tangent space is known as the Riemannian metric. According to the theory of second-order differential equations, there has one and only one geodesic starting from that point with the tangent vector. This allows us to span the curved manifold in the flat tangent space along the geodesics (think of rolling a sphere along its tangent plane at a given point). The geodesics going through the reference point \mathbf{y} are transformed into straight lines and the distance along these geodesics is preserved (at least in a neighborhood of \mathbf{y}). The function that maps to each tangent vector $\mathbf{v} \in T_{\mathbf{y}}\mathcal{M}$ the point \mathbf{x} of the manifold that is achieved by the geodesic starting at \mathbf{y} with this tangent vector is named the exponential map. This map is defined in the whole tangent space $T_{\mathbf{y}}\mathcal{M}$ but it is generally one-to-one only locally around 0 in the tangent space (i.e., around the reference point \mathbf{y} in the manifold). For mapping the manifold data to the tangent space that respects Euclidean geometry,

we denote $\mathbf{v} = \log_{\mathbf{y}}(\mathbf{x})$ as the inverse of the exponential map: this is the smallest vector such that $\mathbf{x} = \exp_{\mathbf{y}} \mathbf{v}$.

For the second issue, we suggest to generalize the Wasserstein-1 distance to Riemannian manifolds. In theory, we need to first study whether there exists an optimal map for the manifold setting. For optimal mass transport on real-valued data, (Evans 1997; Villani 2008; Lei et al. 2017) proved the existence of the optimal map, which has a convex potential (i.e. $D(\mathbf{x}) = \nabla\phi(\mathbf{x})$ with $\phi(\mathbf{x})$ being convex) and is shown to be the only map with the convex potential. Analogously, (McCann 2001; De Philippis and Figalli 2014; Fathi and Figalli 2010) also studied that there exists an optimal map for the manifold case. Furthermore, (Fathi and Figalli 2010) proved that the optimal transport problem can be solved for the square Riemannian metric on any complete manifolds without any assumption on the compactness or curvature, with the usual restriction on the measures. The following theorem is given for the existence of the optimal map on complete manifolds. For its proof, we refer readers to (Fathi and Figalli 2010) (mainly in Page 18-20).

Theorem 1. *Let \mathcal{M} be a Riemannian manifold, and consider the cost $c = d^2/2$, with d being the Riemannian distance metric that is required to be complete. Given two probability distributions \mathbb{P}_r and \mathbb{P}_g supported on the manifold \mathcal{M} , there exists a convex potential function $\phi : \mathcal{M} \rightarrow \mathbb{R} \cup \{+\infty\}$ such that $D(\mathbf{x}) = \exp_{\mathbf{y}}(\nabla_g\phi(\mathbf{x}))$ is the unique optimal transport map sending \mathbb{P}_r to \mathbb{P}_g , where ∇_g indicates the gradient with respect to the Riemannian metric $g = \langle \log_{\mathbf{y}}(\mathbf{x}_1), \log_{\mathbf{y}}(\mathbf{x}_2) \rangle_{\mathbf{y}}$ on the manifold \mathcal{M} .*

Following **Theorem 1**, we turn to study the complete Riemannian metrics for the existence of optimal transport map on our studied manifolds where the HSV, CB and DT image data reside on. It is known that any Riemannian metrics defined on compact Riemannian manifolds are complete. Therefore, for HSV and CB images, Riemannian metrics defined on them are all complete. As to the case of DT images, whose data reside on a non-compact manifold, we restore to employing the well-studied Log-Euclidean metric (Arsigny et al. 2006; 2007; Huang et al. 2015), which is actually an inner product distance and thus the resulting metric space is complete (Minh and Murino 2016). For a more general guidance to satisfy the completeness requirement, readers are suggested to follow the Hopf–Rinow theorem (Hopf and Rinow 1931).

Consequently, inspired by the original Wasserstein distance Eqn.1 in a Euclidean setting, we present a new Wasserstein distance on the underlying Riemannian manifold. By leveraging the complete Riemannian metric that generally applies classic logarithm map to yield Euclidean representations, the new Wasserstein distance can be derived by:

$$W(\mathbb{P}_r, \mathbb{P}_g) = \inf_{\gamma \in \Pi(\mathbb{P}_r, \mathbb{P}_g)} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \gamma} [\|\log_{\mathbf{y}}(\mathbf{x}) - \log_{\mathbf{y}}(\mathbf{y})\|], \quad (3)$$

where \mathbf{x}, \mathbf{y} are points on the Riemannian manifold, $\Pi(\mathbb{P}_r, \mathbb{P}_g)$ indicates the set of all joint distributions $\gamma(\mathbf{x}, \mathbf{y})$, and $\gamma(\mathbf{x}, \mathbf{y})$ denotes how much mass will be transported from \mathbf{x} to \mathbf{y} .

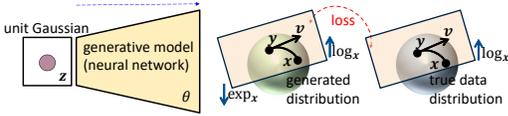


Figure 2: Overview of the proposed manifold-aware Wasserstein GAN (manifoldWGAN) approach for manifold-valued image generation. The manifoldWGAN introduces logarithm map and exponential map between Riemannian manifold and tangent space to the setting of Wasserstein GANs.

Manifold-aware Wasserstein GAN

According to **Theorem 1** and the derived Wasserstein distance Eqn.3, we generalize the objective function Eqn.2 of Wasserstein GAN¹ to Riemannian manifolds

$$\begin{aligned} \min_G \max_D \mathbb{E}_{\mathbf{x} \sim \mathbb{P}_r} [D(\log_{\mathbf{y}}(\mathbf{x}))] \\ - \mathbb{E}_{G(\mathbf{z}) \sim \mathbb{P}_g} [D(\log_{\mathbf{y}}(\exp_{\mathbf{y}}(G(\mathbf{z}))))] \quad (4) \\ + \lambda \mathbb{E}_{\hat{\mathbf{x}} \sim \mathbb{P}_{\hat{\mathbf{x}}}} [(\|\nabla_{\hat{\mathbf{x}}} D(\hat{\mathbf{x}})\|_2 - 1)^2], \end{aligned}$$

where $\log_{\mathbf{y}}(\cdot)$ and $\exp_{\mathbf{y}}(\cdot)$ are the logarithm and exponential maps respectively for the underlying manifold. As the logarithm map $\log_{\mathbf{y}}(\cdot)$ projects the manifold-valued data to Euclidean space, any regular networks can be applied directly to the resulting data. To generate valid manifold-valued data using the generative network G , we employ exponential map $\exp_{\mathbf{y}}(\cdot)$ to transform the data back to the manifold. An overview of the proposed manifold-aware Wasserstein GAN (manifoldWGAN) approach is shown in Fig.2. As done in (Gulrajani et al. 2017), $\mathbb{P}_{\hat{\mathbf{x}}}$ is also defined sampling uniformly along straight lines between pairs of points sampled from the data distribution \mathbb{P}_r and the generator distribution \mathbb{P}_g . In the context of manifold-valued data, we apply logarithm and exponential maps to the linear sampling:

$$\hat{\mathbf{x}} = (1 - t) \log_{\mathbf{y}}(\mathbf{x}) + t \log_{\mathbf{y}}(\exp_{\mathbf{y}}(G(\mathbf{z}))), \quad (5)$$

where $0 \leq t \leq 1$.

In this paper, our focus is on the generation of HSV, CB and DT images. Therefore, we consider the data on the product manifold $\mathcal{H} = \mathbb{S}^1 \times [0, 1]^2$, the spherical data on \mathbb{S}^2 and the SPD data on $\text{SPD}(3)$. Specially, we study the typical forms for their resulting exponential map and logarithm map to achieve the objective Eqn.4.

HSV case: Following (Bergmann et al. 2014; Bergmann and Weinmann 2016), we adopt the representation system $\mathbb{S}^1 \cong [-\pi, \pi)$ to interpret any point $\mathbf{x} \in \mathcal{H} = \mathbb{S}^1 \times [0, 1]^2$ as consisting of two components $(\mathbf{x}_{\mathbb{S}}, \mathbf{x}_{\mathbb{R}})$. As a result, the corresponding logarithm and exponential maps can be derived as

$$\begin{aligned} \log_{\mathbf{y}}(\mathbf{x}) &= [\text{mod}(\mathbf{x}_{\mathbb{S}} - \mathbf{y}_{\mathbb{S}}, 2\pi), \mathbf{x}_{\mathbb{R}} - \mathbf{y}_{\mathbb{R}}], \\ \exp_{\mathbf{y}}(\mathbf{v}) &= [\text{mod}(\mathbf{v}_{\mathbb{S}} + \mathbf{y}_{\mathbb{S}}, 2\pi), \mathbf{v}_{\mathbb{R}} + \mathbf{y}_{\mathbb{R}}], \end{aligned} \quad (6)$$

¹As studying the generalization of all the Wasserstein GANs (Arjovsky, Chintala, and Bottou 2017; Gulrajani et al. 2017; Wei et al. 2018) is beyond the scope of this paper, we choose the most representative Wasserstein GAN (Gulrajani et al. 2017) for study.

Algorithm 1 Manifold-aware Wasserstein GAN (manifoldWGAN), our proposed algorithm. All the experiments in the paper used the default values $\lambda = 10$, $n_{\text{critic}} = 5$.

Require: α , learning rate. m , the batch size. n_{critic} , the critic iterations per generation iteration. λ , the balance parameter of gradient norm penalty, \mathbf{w}_0 , initial critic parameters. θ_0 , initial generator's parameter.

```

1: while  $\theta$  has not converged do
2:   for  $t=0, \dots, n_{\text{critic}}$  do
3:     Sample  $\{\mathbf{x}^{(i)}\}_{i=1}^m \sim \mathbb{P}_r$  a batch from the real data.
4:     Sample  $\{\mathbf{z}^{(i)}\}_{i=1}^m \sim \mathbb{P}_g$  a batch of prior samples.
5:      $D_w \leftarrow \nabla_w [\mathcal{L}]$  where  $\mathcal{L}$  is computed by Eqn.4.
6:      $\mathbf{w} \leftarrow \mathbf{w} + \alpha \cdot \text{AdamOptimizer}(\mathbf{w}, D_w)$ 
7:   end for
8:   Sample  $\{\mathbf{z}^{(i)}\}_{i=1}^m \sim \mathbb{P}_g$  a batch of prior samples.
9:    $G_\theta \leftarrow \nabla_\theta [-\frac{1}{m} \sum_{i=1}^m D_w(\log_{\mathbf{y}}(\exp_{\mathbf{y}}(G_\theta(\mathbf{z}^{(i)}))))]$ 
10:   $\theta \leftarrow \theta + \alpha \cdot \text{AdamOptimizer}(\theta, G_\theta)$ 
11: end while

```

where mod denotes the modulo operation.

CB case: As studied in (Bhattacharya and Patrangenaru 2002; Lee, Abbott, and Araman 2007; Bacák et al. 2016; Laus et al. 2017), on the sphere manifold \mathbb{S}^n , the logarithm and exponential maps can be expressed by

$$\begin{aligned} \log_{\mathbf{y}}(\mathbf{x}) &= \frac{d(\mathbf{x}, \mathbf{y})}{\|p_{\mathbf{y}}(\mathbf{x} - \mathbf{y})\|_F} p_{\mathbf{y}}(\mathbf{x} - \mathbf{y}), \\ \exp_{\mathbf{y}}(\mathbf{v}) &= \cos(\|\mathbf{v}\|_F) \mathbf{y} + \frac{\sin(\|\mathbf{v}\|_F)}{\|\mathbf{v}\|_F} \mathbf{v}, \end{aligned} \quad (7)$$

where $\|\cdot\|_F$ indicates the Frobenius norm operation, $p_{\mathbf{y}}(\mathbf{H}) = \mathbf{H} - \text{trace}(\mathbf{y}^T \mathbf{H}) \mathbf{y}$.

DT case: (Arsigny et al. 2006; 2007; Huang et al. 2015; Huang and Van Gool 2017) studied that the Log-Euclidean metric for the SPD manifold $\text{SPD}(n)$ is derived by employing the Lie group structure. The resulting logarithmic and exponential maps is expressed with matrix logarithms and exponential operations:

$$\begin{aligned} \log_{\mathbf{y}}(\mathbf{x}) &= D_{\log(\mathbf{y})} \exp \cdot (\log(\mathbf{y}) - \log(\mathbf{x})), \\ \exp_{\mathbf{y}}(\mathbf{v}) &= \exp(\log(\mathbf{y}) + D_{\mathbf{y}} \log \cdot \mathbf{v}), \end{aligned} \quad (8)$$

where $D_{\log(\mathbf{y})} \exp \cdot = (D_{\mathbf{y}} \log \cdot)^{-1}$ is achieved by the differentiation of the equality $\log \circ \exp = \mathbf{I}$, and here \mathbf{I} is the identity matrix.

Following such basic logarithm and exponential maps² on the corresponding manifolds of the processed manifold-valued samples, the algorithm of the proposed manifoldWGAN is presented in Algorithm.1.

²For the HSV and CB case, the computational cost of log and exp operations are cheap as they only include simple operations like modulo, cosine and sine. For the DT case, the dominant computational complexity of the used log and exp operations lies on computing the matrix logarithm and exponential on SPD matrix valued DT image pixels. As the size of each SPD matrix is only of 3×3 , the log and exp computation is also cheap for network training.

Experiment

For the studied manifold-valued image generation problem, we suggest three benchmark evaluations that use the HSV and CB images of the well-known CIFAR-10 (Krizhevsky and Hinton 2009), ImageNet (Oord, Kalchbrenner, and Kavukcuoglu 2016), and the popular UCL DT image dataset (Cook et al. 2006). As our focus is on evaluating the generative models for image generation, we mainly compare the proposed manifoldWGAN against the state-of-the-art GAN techniques including deep convolutional GAN (DCGAN) (Radford, Metz, and Chintala 2015), least square GAN (LSGAN) (Mao et al. 2017) and Wasserstein GAN (WGAN)³ (Gulrajani et al. 2017) that has proved to improve the original WGAN (Arjovsky, Chintala, and Bottou 2017).

In all the evaluations, we follow (Gulrajani et al. 2017) to use residual network for all compared models. The setting of the employed residual networks is the same as the ones used in (Gulrajani et al. 2017). Specially, we utilize pre-activation residual blocks with two 3×3 convolutional layers and ReLU nonlinearity. Some residual blocks perform downsampling (in the critic) using mean pooling after the second convolutional layer, or nearest-neighbor upsampling (in the generator) before the second convolution. As done by (Gulrajani et al. 2017), we employ batch normalization in the generator but not the critic for WGAN and our proposed manifoldWGAN. We finally optimize the network using Adam with learning rate 0.0002, decayed linearly to 0 over 100K generator iterations, and batch size 64. For further architectural details, please refer to the official implementation of (Gulrajani et al. 2017).

HSV Image Generation

For HSV color image generation, we choose the CIFAR-10 and ImageNet datasets, both of which are standard benchmarks in real-valued image generation. The CIFAR-10 dataset consists of 60000 32×32 colour images in 10 classes, with 6000 images per class. We use the 64×64 version of ImageNet, which contains 1,281,149 training images and 49,999 images for testing. To gather HSV color images whose values are on $\mathcal{H} = \mathbb{S}^1 \times [0, 1]^2$, we transfer the images from the RGB space to the HSV space.

For generating valid elements on $\mathcal{H} = \mathbb{S}^1 \times [0, 1]^2$, the proposed manifoldWGAN suggests to use the corresponding logarithm and exponential maps Eqn.6 (where each pixel value of the reference point \mathbf{y} is set to $[\pi, 0, 0]^4$) during the optimization of the objective Eqn.4, while the state-of-the-art generation technique WGAN treats the data as real-valued data with its original Wasserstein-based GAN loss.

Fig.3 qualitatively shows the generation results of WGAN and the proposed manifoldWGAN on CIFAR-10 and ImageNet. The results justify that our manifoldWGAN can generate more visually pleasing HSV images, whose resulting RGB images appear to be better in both terms of image quality and semantic.

³The official code is available at https://github.com/igul222/improved_wgan_training

⁴We empirically find varying the value of anchor point impacts the performance of our model slightly.

Method	CIFAR-10	ImageNet
DCGAN-RGB	37.7	95.5
LSGAN-RGB	31.9	82.4
WGAN-RGB	29.3	71.2
WGAN-HSV	38.4	76.7
manifoldWGAN-HSV	27.2	68.4
WGAN-CB	80.7	101.3
manifoldWGAN-CB	59.4	89.2

Table 1: FIDs of compared methods performing generation on RGB images or HSV images of CIFAR-10 and ImageNet. The FIDs of all the HSV/CB-based models are computed by translating their produced HSV/CB images to RGB images.

In addition, we follow (Heusel et al. 2017; Lucic et al. 2017) to adopt Fréchet inception distance (FID)⁵ to compare the WGAN and our manifoldWGAN quantitatively on the CIFAR-10 dataset. As the original FID metric is computed on RGB images⁶, we transfer the produced HSV images of generative models to the corresponding RGB images. Besides, we also present the FIDs of the state-of-the-art GANs (Radford, Metz, and Chintala 2015; Mao et al. 2017; Gulrajani et al. 2017)⁷ that are all originally designed for RGB images. The results reported in Tab.1 demonstrate our manifoldWGAN performs better than WGAN for HSV image generation on both of the CIFAR-10 and ImageNet datasets. More interestingly, the performances of our manifoldWGAN-HSV on such two datasets are also consistently better than those of the state-of-the-art GANs (e.g., WGAN) that work for RGB images directly. This is because our manifoldWGAN for HSV-based image generation targets a better understanding on color semantics. From this perspective, studying HSV-based image generation would be a good direction to enhance the regular image generation.

We finally show the training curves of negative Wasserstein GAN loss (Wasserstein distance) and generated samples at different stages of training in Fig.4. As we can see, the loss decreases consistently as training progresses and sample quality increases. This verifies a strong correlation between lower loss and better sample quality.

CB Image Generation

To evaluate our manifoldWGAN for image generation on the sphere manifold \mathbb{S}^2 , we collect CB images from CIFAR-10 and ImageNet. The basic settings on them are the same with

⁵The FID uses the statistics (i.e., mean and covariance) of real world samples and compare it to the statistics of synthetic samples.

⁶While it might be possible to adapt the FID computation on manifold-valued images, we leave it as one of future works due to the absence of effective manifold-valued inception model.

⁷We also evaluated spectral normalized WGAN (SN-WGAN) (Miyato et al. 2018) and our manifold-enhanced SN-WGAN. Due to the space limit, we only report that their FIDs are 31.3 and 24.2 resp. for the CIFAR-10 HSV case, showing our method can enable other WGANs to work well for manifold-valued image generation.

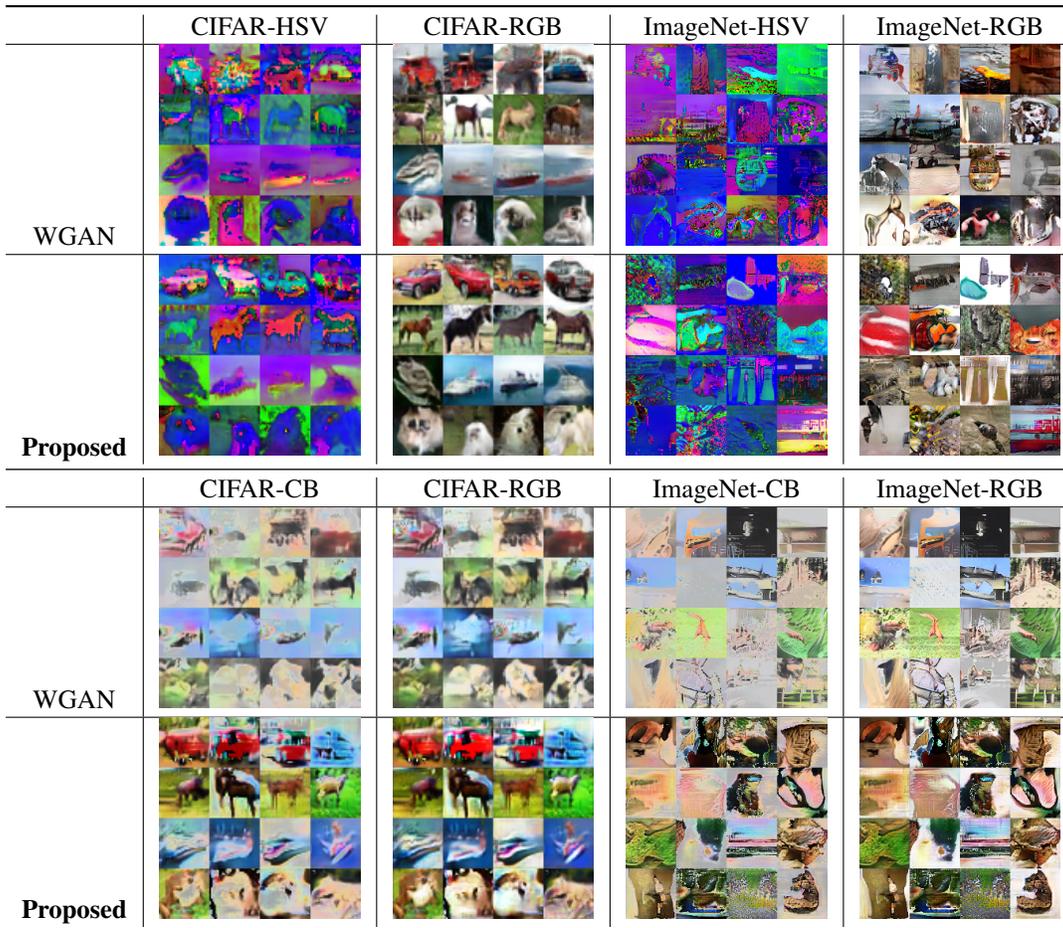


Figure 3: Generated samples of WGAN (Gulrajani et al. 2017) and the proposed manifoldWGAN for HSV/CB image generation on CIFAR-10/ImageNet with translated RGB images. Please refer to supplementary material for more visual results.

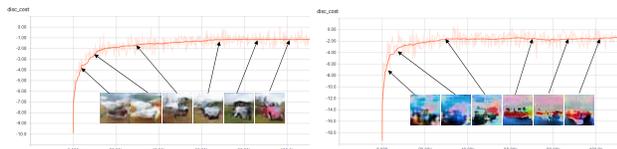


Figure 4: Training curves and generated samples at different training stages of our manifoldWGAN on CIFAR-10 HSV image set (left) and CB image set (right). The value on y-axis denotes the negative critic cost (Wasserstein distance), and the value on the x-axis indicates the iteration number.

those of the last evaluation. For the task, we extract the chromaticity component of the CB images so that the generation problem on pure spherical data can be better studied.

As done in the last evaluation, we compare our manifoldWGAN against the state-of-the-art GAN technique WGAN. Our manifoldWGAN proposes to make use of the manifold geometry-aware logarithm and exponential maps Eqn.7 (where each pixel value of the anchor point \mathbf{y} is set to

$[\frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}]^8$) on the sphere manifold during the sample generation. As shown in Fig.3, we discover that the WGAN approach generates samples with low-chromaticity, while the proposed manifoldWGAN can generate higher-quality samples. This shows the clear advantage of our method.

Similar to the last evaluation, in Tab.1 we also employ the FID metric to compare our manifoldWGAN against the state-of-the-art methods quantitatively for CIFAR-10 and ImageNet. For the FID computation, we translate CB images back to RGB images. By comparing with the models that work for RGB images directly, we can see both the WGAN-CB and our manifoldWGAN-CB performs relatively worse. This is reasonable because the CB-to-RGB translation is non-trivial and lossy. Nevertheless, our proposed manifoldWGAN-CB can still surpass the real competitor WGAN-CB, showing the clear superiority of our proposed model when working for CB images.

We also show the correlation between the proposed manifold-ware Wasserstein GAN loss and the quality of gen-

⁸The performance of our model changes slightly when varying the reference point.

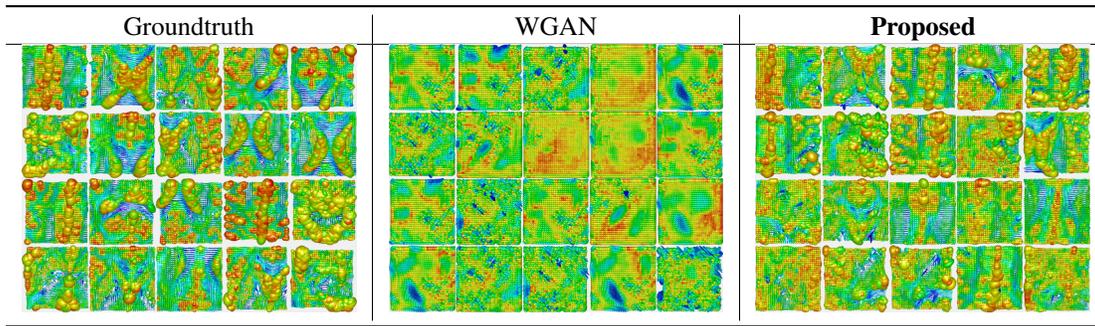


Figure 5: Groundtruth DT images (left), Generated samples of WGAN (Gulrajani et al. 2017) (middle) and the proposed manifoldWGAN (right) for DT image generation on the UCL DT image dataset. The figure is better viewed via zooming in. For more visual results, please refer to supplementary material.

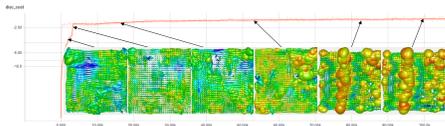


Figure 6: Training curves and generated samples at different stages of training of the proposed manifoldWGAN on UCL DT image set. The values on y-axis and x-axis denote the negative critic cost (Wasserstein distance) and the iteration number respectively.

erated CB images in Fig.4. From the result, we can observe that the quality of generated CB samples tend to be better as the iteration number increases.

DT image Generation

We employ Camino brain DT image set (Cook et al. 2006) to evaluate our proposed manifoldWGAN for the SPD-valued image generation on SPD(3). The Camino project (Cook et al. 2006) provides a dataset of a Diffusion Tensor Magnetic Resonance Image (DT-MRI) of human heads, which is freely available⁹. The UCL DT image database contains 50 brain DT images in total. For data augmentation, we extract 26,080 DT slices of size 32×32 from the original DT images. Some ground truth DT slice images are listed in Fig.5 (a). Among the sampled DT slice images, there are around 0.6% non-SPD voxels in each DT slice on average, and only 2,399 samples contain no non-SPD voxels.

For generating visually favorable DT images, our manifoldWGAN suggests to introduce logarithm map and exponential map Eqn.8 (where each pixel value of the reference point \mathbf{y} is set to the identity matrix of size 3×3) on SPD manifolds to the Wasserstein GAN setting. In contrast, the original WGAN treats the SPD data as real-valued data. In order to make the generated samples of WGAN to become valid DT (SPD-valued) images, we additionally employ our presented exponential map to transform the outputs of WGAN to the

data on SPD manifolds. In Fig.5, we quantitatively compare these two different approaches on the UCL DT image set. Note that, to plot the generated DT images, we leverage the Manifold-valued Image Restoration Toolbox¹⁰. From the results, we can discover that the generated samples of WGAN is not plausible, while our manifoldWGAN can often synthesize highly realistic DT samples, which would be very useful to the field of DT-MRI that generally lacks of DT images for analysis.

As the last two evaluations, we also study the relationship between the proposed manifold-aware Wasserstein loss and the generated sample quality in Fig.6. As the network training iterations increases, the negative Wasserstein loss goes up quickly and the visual quality of the generated DT images is also improved, again showing there are clear correlations among our manifoldWGAN loss, the convergence property and the quality of its generated samples.

Discussion and Outlook

We introduced a new generative modeling problem over manifold-valued images, which are often encountered in a wide variety of applications including HSV/CB color and DT-MRI image processing. To address the new problem, we generalize the Wasserstein GAN methodology to manifold-valued generation by introducing the manifold geometry-aware optimal transport theory. We finally suggested three benchmarks to evaluate the generative power of our proposed method, and clearly show the superiority of our approach over the state-of-the-art Wasserstein GAN.

For future works, inspired by (Shrivastava et al. 2017) that leverage generative modeling to improve other visual tasks, we will further exploit the practical benefits of our proposed generative models to a broad number of classical manifold-valued data processing tasks like denoising, inpainting and segmentation. In addition, we hope our work could open the path for more manifold-valued image generation tasks like shape and Grassmann manifold-valued data generation. Last but not least, like real-valued image generation, we also encourage more works to extend our framework to conditional

⁹To obtain the data, one can follow the tutorial at <http://cmic.cs.ucl.ac.uk/camino//index.php?n=Tutorials.DTI>

¹⁰The toolbox is available at <http://www.mathematik.uni-kl.de/imagepro/members/bergmann/mvirt/>.

settings and even semi-supervised learning.

Acknowledgement: We would like to thank Nvidia for donating the GPUs used in this work.

References

- Arjovsky, M.; Chintala, S.; and Bottou, L. 2017. Wasserstein generative adversarial networks. In *ICML*.
- Arsigny, V.; Fillard, P.; Pennec, X.; and Ayache, N. 2006. Log-Euclidean metrics for fast and simple calculus on diffusion tensors. *Magnetic resonance in medicine* 56(2):411–421.
- Arsigny, V.; Fillard, P.; Pennec, X.; and Ayache, N. 2007. Geometric means in a novel vector space structure on symmetric positive-definite matrices. *SIAM journal on matrix analysis and applications* 29(1):328–347.
- Bacák, M.; Bergmann, R.; Steidl, G.; and Weinmann, A. 2016. A second order nonsmooth variational model for restoring manifold-valued images. *SIAM Journal on Scientific Computing* 38(1):A567–A597.
- Bansal, S., and Tatu, A. 2015. Active contour models for manifold valued image segmentation. *Journal of Mathematical Imaging and Vision* 52(2):303–314.
- Bergmann, R., and Weinmann, A. 2015. Inpainting of cyclic data using first and second order differences. *EMMCVPR, Lecture Notes in Computer Science* 155–168.
- Bergmann, R., and Weinmann, A. 2016. A second-order TV-type approach for inpainting and denoising higher dimensional combined cyclic and vector space data. *Journal of Mathematical Imaging and Vision* 55(3):401–427.
- Bergmann, R.; Laus, F.; Steidl, G.; and Weinmann, A. 2014. Second order differences of cyclic data and applications in variational denoising. *SIAM Journal on Imaging Sciences* 7(4):2916–2953.
- Bhattacharya, R., and Patrangenaru, V. 2002. Nonparametric estimation of location and dispersion on Riemannian manifolds. *Journal of Statistical Planning and Inference* 108(1):23–35.
- Cook, P.; Bai, Y.; Nedjati-Gilani, S.; Seunarine, K.; Hall, M.; Parker, G.; and Alexander, D. 2006. Camino: open-source diffusion-mri reconstruction and processing. In *14th scientific meeting of the international society for magnetic resonance in medicine*, volume 2759. Seattle WA, USA.
- De Philippis, G., and Figalli, A. 2014. The monge–ampère equation and its link to optimal transportation. *Bulletin of the American Mathematical Society* 51(4):527–580.
- Evans, L. C. 1997. Partial differential equations and Monge-Kantorovich mass transfer. *Current developments in mathematics* 1997(1):65–126.
- Fathi, A., and Figalli, A. 2010. Optimal transportation on non-compact manifolds. *Israel Journal of Mathematics* 175(1):1–59.
- Fitschen, J. H.; Laus, F.; and Schmitzer, B. 2017. Optimal transport for manifold-valued images. In *International Conference on Scale Space and Variational Methods in Computer Vision*, 460–472. Springer.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *NIPS*.
- Gulrajani, I.; Ahmed, F.; Arjovsky, M.; Dumoulin, V.; and Courville, A. 2017. Improved training of Wasserstein GANs. *NIPS*.
- Hasan, K. M.; Walimuni, I. S.; Abid, H.; and Hahn, K. R. 2011. A review of diffusion tensor magnetic resonance imaging computational methods and software tools. *Computers in biology and medicine* 41(12):1062–1072.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *NIPS*.
- Hopf, H., and Rinow, W. 1931. Über den begriff der vollständigen differentialgeometrischen fläche. *Commentarii Mathematici Helvetici* 3(1):209–225.
- Huang, Z., and Van Gool, L. 2017. A Riemannian network for SPD matrix learning. In *AAAI*.
- Huang, Z.; Wang, R.; Shan, S.; Li, X.; and Chen, X. 2015. Log-Euclidean metric learning on symmetric positive definite manifold with application to image set classification. In *ICML*.
- Jayasumana, S.; Hartley, R.; Salzmann, M.; Li, H.; and Harandi, M. 2013. Kernel methods on the riemannian manifold of symmetric positive definite matrices. In *CVPR*.
- Krizhevsky, A., and Hinton, G. 2009. Learning multiple layers of features from tiny images. *Technical report, University of Toronto*.
- Laus, F.; Nikolova, M.; Persch, J.; and Steidl, G. 2017. A nonlocal denoising algorithm for manifold-valued images using second order statistics. *SIAM Journal on Imaging Sciences* 10(1):416–448.
- Lee, S.-M.; Abbott, A. L.; and Araman, P. A. 2007. Dimensionality reduction and clustering on statistical manifolds. In *CVPR*.
- Lei, N.; Su, K.; Cui, L.; Yau, S.-T.; and Gu, D. X. 2017. A geometric view of optimal transportation and generative model. *arXiv preprint arXiv:1710.05488*.
- Loeper, G. 2011. Regularity of optimal maps on the sphere: The quadratic cost and the reflector antenna. *Archive for rational mechanics and analysis* 199(1):269–289.
- Lucic, M.; Kurach, K.; Michalski, M.; Gelly, S.; and Bousquet, O. 2017. Are GANs created equal? a large-scale study. *arXiv preprint arXiv:1711.10337*.
- Mao, X.; Li, Q.; Xie, H.; Lau, R. Y.; Wang, Z.; and Smolley, S. P. 2017. Least squares generative adversarial networks. In *ICCV*.
- McCann, R. J. 2001. Polar factorization of maps on riemannian manifolds. *Geometric & Functional Analysis GAFA* 11(3):589–608.
- Minh, H. Q., and Murino, V. 2016. From covariance matrices to covariance operators: Data representation from finite to infinite-dimensional settings. In *Algorithmic Advances in Riemannian Geometry and Applications*. Springer. 115–143.
- Miyato, T.; Kataoka, T.; Koyama, M.; and Yoshida, Y. 2018. Spectral normalization for generative adversarial networks. *ICLR*.
- Oord, A. v. d.; Kalchbrenner, N.; and Kavukcuoglu, K. 2016. Pixel recurrent neural networks. *arXiv preprint arXiv:1601.06759*.
- Pennec, X.; Fillard, P.; and Ayache, N. 2006. A Riemannian framework for tensor computing. *IJCV* 66(1):41–66.
- Priya, U. S., and Nair, J. J. 2015. Denoising of dt-mr images with an iterative pca. *Procedia Computer Science* 58:603–613.
- Radford, A.; Metz, L.; and Chintala, S. 2015. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*.
- Shrivastava, A.; Pfister, T.; Tuzel, O.; Susskind, J.; Wang, W.; and Webb, R. 2017. Learning from simulated and unsupervised images through adversarial training. In *CVPR*.
- Villani, C. 2008. *Optimal transport: old and new*, volume 338. Springer Science & Business Media.
- Wei, X.; Gong, B.; Liu, Z.; Lu, W.; and Wang, L. 2018. Improving the improved training of Wasserstein GANs: A consistency term and its dual effect. In *ICLR*.
- Zhao, J.; Mathieu, M.; and LeCun, Y. 2016. Energy-based generative adversarial network. *arXiv preprint arXiv:1609.03126*.