

Manifold-Aware CycleGAN for High-Resolution Structural-to-DTI Synthesis^{*}

Benoit Anctil-Robitaille¹, Christian Desrosiers¹, Herve Lombaert¹

Department of IT and Software Engineering, ETS Montreal, Canada
benoit.anctil-robitaille.1@ens.etsmtl.ca

Abstract. Unpaired image-to-image translation has been applied successfully to natural images but has received very little attention for manifold-valued data such as in diffusion tensor imaging (DTI). The non-Euclidean nature of DTI prevents current generative adversarial networks (GANs) from generating plausible images and has mainly limited their application to diffusion MRI scalar maps, such as fractional anisotropy (FA) or mean diffusivity (MD). Even if these scalar maps are clinically useful, they mostly ignore fiber orientations and therefore have limited applications for analyzing brain fibers. Here, we propose a manifold-aware CycleGAN that learns the generation of high-resolution DTI from unpaired T1w images. We formulate the objective as a Wasserstein distance minimization problem of data distributions on a Riemannian manifold of symmetric positive definite 3×3 matrices SPD(3), using adversarial and cycle-consistency losses. To ensure that the generated diffusion tensors lie on the SPD(3) manifold, we exploit the theoretical properties of the exponential and logarithm maps of the Log-Euclidean metric. We demonstrate that, unlike standard GANs, our method is able to generate realistic high-resolution DTI that can be used to compute diffusion-based metrics and potentially run fiber tractography algorithms. To evaluate our model’s performance, we compute the cosine similarity between the generated tensors principal orientation and their ground-truth orientation, the mean squared error (MSE) of their derived FA values and the Log-Euclidean distance between the tensors. We demonstrate that our method produces 2.5 times better FA MSE than a standard CycleGAN and up to 30% better cosine similarity than a manifold-aware Wasserstein GAN while synthesizing sharp high-resolution DTI.

Keywords: Manifold-Valued Data Learning · High-Resolution DTI · Log-Euclidean Metric.

1 Introduction

Unpaired image-to-image translation and image synthesis have been widely used in medical imaging [21]. Whether they are employed to generate missing modal-

^{*} This work was supported financially by the Réseau de Bio-Imagerie du Québec (RBIQ), the Research Council of Canada (NSERC), the Fonds de Recherche du Québec (FQRNT), ETS Montreal, and NVIDIA with the donation of a GPU.

ities, normalize images or enhance images quality and resolution, generative adversarial networks (GANs) [4] have been proven effective in multiple challenging medical image analysis tasks. However, they have been mainly studied on real-valued images, thus impeding the development of applications for manifold-valued data such as diffusion tensor images (DTI). Despite the growing interest in the brain’s structural connectivity, applications of GANs to DTI have been mostly limited to generating derived scalar maps like fractional anisotropy (FA) and mean diffusivity (MD), which ignore the fibers’ orientation and provide limited insights on their structural organization.

Among the literature, [5] investigates the generation of diffusion MRI scalar maps from T1w images using a CycleGAN [23]. The authors show that the structural and diffusion spaces share a sufficient amount of information to be able to synthesize realistic FA and MD maps from downsampled T1w images. In [22], dual GANs and Markovian discriminators are used to harmonize multi-site FA and MD maps of neonatal brains. They demonstrate that using a GAN-like architecture can better capture the complex non-linear relations between multiple domains than standard normalization methods.

While the previously mentioned works present applications of GANs on DTI-derived metrics, they do not tackle the challenge of generating DTI. Being able to synthesize such images would unlock a vast amount of useful methods that are already well studied on real-valued modalities, while preserving all the geometrical information encoded in the diffusion tensors. However, DTI data is manifold-valued: the data of each voxel lies on a Riemannian manifold of symmetric positive definite 3×3 matrices, i.e., the SPD(3) manifold. The non-Euclidean nature of DTI prevents standard GANs from generating plausible images as there is no guarantee that the generated diffusion tensors lie on the SPD(3) manifold. A solution presented in [2] is to employ the Log-Euclidean metric to accurately process data on the SPD(n) manifold. By using the log and exp projections proposed in [2], one can apply Euclidean operations on tensors and guarantee that the resulting tensors will lie on such manifold. Those computationally efficient mapping operations form an interesting framework for manifold-valued data learning, and have been used in [7] to develop a deep neural network called SPDNet which learns discriminative SPD matrices. With the help of the matrix backpropagation of spectral layers defined in [9], they designed a network that learns data on SPD(n). Nonetheless, SPDNet [7] is limited to single SPD matrix learning and cannot help in learning multiple spatially-organized SPD matrices as it is the case with DTI.

Related to our work, [8] proposes a manifold-aware Wasserstein GAN for manifold-valued data generation, which leverages the aforementioned log and exp mappings. In their work, they generate plausible slices of DTI from noise vectors. By comparing the produced images of their network with those produced by a regular GAN, one can clearly see that the manifold mappings are necessary to produce credible diffusion tensors. However, the proposed manifold-aware GAN could not provide any additional clinical insights, nor help in understanding the brain’s connectivity as the generated images are not conditioned by any real

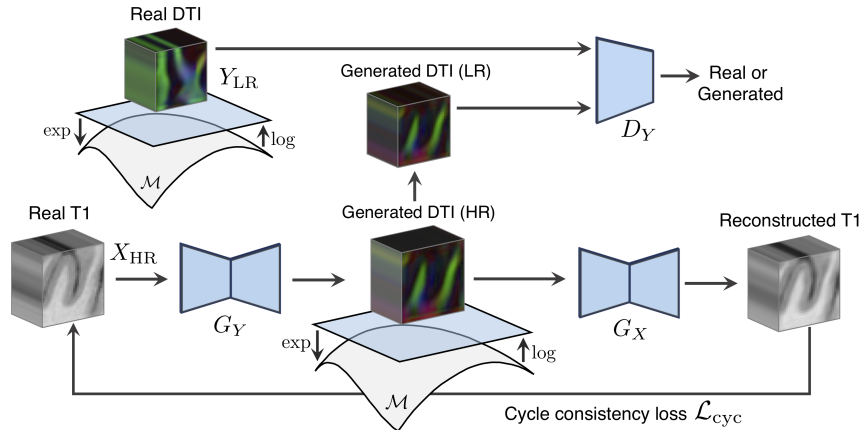


Fig. 1. The forward cycle of our manifold-aware CycleGAN: G_Y generates high-resolution DTIs on the SPD(3) manifold by projecting its prediction using the \exp_{Id} and \log_{Id} mappings consecutively. D_Y assesses the downsampled generated images and provides adversarial feedback to G_Y . G_X tries to reconstruct the original T1w images from $G_Y(\mathbf{x})$ and supplies high-resolution gradient information to G_Y .

contextual information such as T1w images. Furthermore, [8] only focuses on the generation of 2D DTIs, which is of limited application for the assessment of the structural organization of the brain’s fibers.

This paper presents a novel manifold-aware Wasserstein CycleGAN that generates high-resolution (HR) DTI from unpaired T1w images. Our method leverages the detailed structural information provided by T1w images while constraining the synthesized diffusion tensors to lie on the SPD(3) manifold using the Log-Euclidean metric. Specifically, the contributions of this work are as follows:

- We present the first CycleGAN model for the unpaired mapping between images and SPD(3) manifold-valued data.
- This is also the first deep learning model to generate DTI data from structural MRI. As mentioned before, previous approaches have focused on generating diffusion scalar maps like FA or MD, and not diffusion tensors as in this work.

Our proposed manifold-aware CycleGAN method is presented in the next section.

2 Method

Let X_{HR} be the domain of high-resolution structural images and Y_{HR} be the domain of high-resolution diffusion tensor images. Our goal is to learn mapping functions $G_Y : X_{\text{HR}} \mapsto Y_{\text{HR}}$ and $G_X : Y_{\text{HR}} \mapsto X_{\text{HR}}$ that translate the real-valued




Fig. 2. (**top row**) Real high-resolution T1w patches and (**bottom row**) recovered T1w patches.

domain X_{HR} into the manifold-valued domain Y_{HR} and the other way around. However, as it is often the case, we do not have access to high-resolution DTI. Thus, we train our model with unpaired training samples $\{\mathbf{x}_i\}_{i=1}^N$ where $\mathbf{x}_i \in X_{\text{HR}}$ is a 3D structural image (e.g., T1w), and $\{\mathbf{y}_j\}_{j=1}^M$ where $\mathbf{y}_j \in Y_{\text{LR}}$ is a DT image with lower resolution. We employ the \log_{Id} and \exp_{Id} mapping to project the generated and the real DTI on the tangent plane at the 3×3 identity matrix, to ensure that $G_Y(\mathbf{x})$ lies on the $\text{SPD}(3)$ manifold and to compare the manifold-valued data distributions as in [8]. Two discriminators, D_X and D_Y , assess the quality of the generated images $G_X(\mathbf{y})$ and downsampled generated DTI $\downarrow G_Y(\mathbf{x})$ with respect to their real data distributions $G_X(\mathbf{y}) \sim \mathbb{P}_{X_{\text{HR}}}$ and $\downarrow G_Y(\mathbf{x}) \sim \mathbb{P}_{\log_{\text{Id}}(Y_{\text{LR}})}$. We formulate the objective as a Wasserstein distance minimization problem on the $\text{SPD}(3)$ manifold with adversarial and cycle-consistent losses. The adversarial portion of the objective helps G_X and G_Y generate images that match the target distribution. On the other hand, the cycle-consistent losses provide high-resolution gradients that are necessary to generate high-resolution DTI with a proper structure.

2.1 Log-Euclidean Metric

Diffusion tensor matrices are well defined in the Log-Euclidean metric, where a matrix logarithm and exponential can be conveniently processed in one metric and always be mapped back to valid symmetric diffusion tensors [2]. Let $\mathbf{M} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{U}^\top$ be the eigendecomposition of a symmetric matrix \mathbf{M} . The computation of the logarithm and the exponential of a tensor noted as \log_{Id} and \exp_{Id} are defined as follows:

$$\forall \mathbf{P} \in S_{++}^*, \log_{\text{Id}}(\mathbf{P}) = \mathbf{U} \log(\boldsymbol{\Sigma}) \mathbf{U}^\top \in \mathcal{T}_{\text{Id}} \quad (1)$$

$$\forall \mathbf{S} \in \mathcal{T}_{\text{Id}}, \exp_{\text{Id}}(\mathbf{S}) = \mathbf{U} \exp(\boldsymbol{\Sigma}) \mathbf{U}^\top \in S_{++}^* \quad (2)$$

We use these maps throughout our work to project the generated and real DTI on the $\text{SPD}(3)$ manifold and on the tangent plane at the 3×3 identity matrix \mathcal{T}_{Id} . Moreover, the Log-Euclidean distance between two tensors \mathbf{P}_1 and \mathbf{P}_2 is defined as:

$$\text{dist}(\mathbf{P}_1, \mathbf{P}_2) = \|\log_{\text{Id}}(\mathbf{P}_1) - \log_{\text{Id}}(\mathbf{P}_2)\|_2. \quad (3)$$

In our framework, we use this distance to measure the similarity between predicted and real DTI in the tangent plane.

2.2 Adversarial Loss

In a traditional GAN setup [4], a generator G and a discriminator D compete in a minimax game where G tries to generate data close to a true data distribution so that D cannot identify if the generated data is real or not. In [1], D is replaced by a discriminator that leverages the Wasserstein distance to estimate the similarity between the real and generated data distributions. The Wasserstein GAN (WGAN) architecture tends to stabilize the training as the Wasserstein distance never saturates, and thus always provides relevant gradients to G . The adversarial part of our objective follows the WGAN framework and is divided in two separate loss terms, $\mathcal{L}_{\text{GAN}_X}$ and $\mathcal{L}_{\text{GAN}_Y}$, respectively for structural images and DTI:

$$\begin{aligned} \mathcal{L}_{\text{GAN}_X}(G_X, D_X, Y_{\text{LR}}, X_{\text{HR}}) \\ = \mathbb{E}_{\mathbf{x} \sim \mathbb{P}_{X_{\text{HR}}}} [D_X(\mathbf{x})] - \mathbb{E}_{\mathbf{y} \sim \mathbb{P}_{Y_{\text{LR}}}} [D_X(G_X(\uparrow \log_{\text{Id}}(\mathbf{y})))] \end{aligned} \quad (4)$$

$$\begin{aligned} \mathcal{L}_{\text{GAN}_Y}(G_Y, D_Y, X_{\text{HR}}, Y_{\text{LR}}) \\ = \mathbb{E}_{\mathbf{y} \sim \mathbb{P}_{Y_{\text{LR}}}} [D_Y(\log_{\text{Id}}(\mathbf{y}))] - \mathbb{E}_{\mathbf{x} \sim \mathbb{P}_{X_{\text{HR}}}} [D_Y(\downarrow G_Y(\mathbf{x}))] \end{aligned} \quad (5)$$

where \uparrow and \downarrow indicates trilinear up and downsampling.

In $\mathcal{L}_{\text{GAN}_X}$, G_X generates 3D structural images from high-resolution DTI projected on the identity-based tangent plane using the \log_{Id} mapping. Since we only have samples from the low-resolution data distribution $\mathbf{y} \sim \mathbb{P}_{Y_{\text{LR}}}$, the real DTI data is upsampled via trilinear interpolation before being fed to G_X . In the same loss, D_X measures the Wasserstein distance between the data distribution of generated and real structural images. Likewise, in $\mathcal{L}_{\text{GAN}_Y}$, D_Y computes the Wasserstein distance between the distribution of downsampled generated DTI and real low-resolution DTI in the tangent plane, using the Log-Euclidean metric [2].

2.3 Cycle Consistency Loss

The adversarial loss alone is not sufficient to drive the generation of high-resolution DTIs. Indeed, the discriminator D_Y only assesses downsampled DT images so its feedback cannot help G_Y improve beyond a certain level of detail. To mitigate this problem, we introduce a second loss that enforces the forward




Fig. 3. On the left, our forward cycle: (1) A T1w image is translated into a high-resolution DTI where each voxel belongs to the SPD(3) manifold, (2) the tensors are projected to \mathcal{T}_{Id} using the \log_{Id} map and (3) the image is translated back to the T1w domain where \mathcal{L}_{cyc} is computed. On the right, our backward cycle: (1) An upsampled DTI on \mathcal{T}_{Id} is translated to the T1w domain, (2) the generated T1w image is translated back to DTI and (3) the recovered DTI is projected to \mathcal{T}_{Id} where \mathcal{L}_{cyc} is computed.

and backward cycle consistency of the network [23], and provides high-resolution gradient information. In our case, the forward cycle consistency ensures that, from the translated DTI $G_Y(\mathbf{x})$, we are able to reconstruct the corresponding original structural images $\mathbf{x} \sim \mathbb{P}_{X_{\text{HR}}}$ which are originally in high resolution. The backward cycle consistency ensures that we are able to reconstruct the original upsampled DTI $\mathbf{y} \sim \mathbb{P}_{\log_{\text{Id}}(Y_{\text{LR}})}$ from the translated structural images $G_X(\uparrow \log_{\text{Id}}(\mathbf{y}))$. The total cycle consistency loss is defined as

$$\begin{aligned} \mathcal{L}_{\text{cyc}}(G_Y, G_X) &= \lambda_{\text{cyc}_X} \mathbb{E}_{\mathbf{x} \sim \mathbb{P}_{X_{\text{HR}}}} [\|G_X(G_Y(\mathbf{x})) - \mathbf{x}\|_1] \\ &+ \lambda_{\text{cyc}_Y} \mathbb{E}_{\mathbf{y} \sim \mathbb{P}_{Y_{\text{LR}}}} [\|G_Y(G_X(\uparrow \log_{\text{Id}}(\mathbf{y}))) - \uparrow \log_{\text{Id}}(\mathbf{y})\|_1] \end{aligned} \quad (6)$$

Here, λ_{cyc_X} and λ_{cyc_Y} balance the contribution of the forward and backward cycles respectively and have been empirically set to a value of 3 and 1. Note that ℓ_1 norm is employed instead of ℓ_2 to make the loss less sensitive to large reconstruction errors.

2.4 Manifold-Aware Wasserstein CycleGAN

Our full objective is

$$\begin{aligned} \mathcal{L}(G_X, G_Y, D_X, D_Y) &= \lambda_{\text{GAN}_X} \mathcal{L}_{\text{GAN}_X}(G_X, D_X, Y_{\text{LR}}, X_{\text{HR}}) \\ &+ \lambda_{\text{GAN}_Y} \mathcal{L}_{\text{GAN}_Y}(G_Y, D_Y, X_{\text{HR}}, Y_{\text{LR}}) + \mathcal{L}_{\text{cyc}}(G_Y, G_X) \end{aligned} \quad (7)$$

The adversarial parts $\mathcal{L}_{\text{GAN}}(G_X, D_X, Y_{\text{LR}}, X_{\text{HR}})$ and $\mathcal{L}_{\text{GAN}}(G_Y, D_Y, X_{\text{HR}}, Y_{\text{LR}})$ of our full objective guide the generators G_X and G_Y towards the synthesis of images close to their real data distributions using a Wasserstein distance on the SPD(3) manifold. The cycle consistency denoted as $\mathcal{L}_{\text{cyc}}(G_Y, G_X)$ gives fine-grained retro-action that helps generate HR DTI while preventing mode collapse. Once training is done, the high-resolution DTI \mathbf{y}_{HR} of a structural image \mathbf{x} can be obtained by applying the exponential map to the DTI generator output: $\mathbf{y}_{\text{HR}} = \exp_{\text{Id}}(G_Y(\mathbf{x}))$.

3 Experiments

Data We employ the pre-processed T1-weighted (T1w) and diffusion MRI (dMRI) data of 1,065 patients from the HCP1200 release of the Human Connectome Project [20] to evaluate our manifold-aware CycleGAN. The T1w (0.7 mm isotropic, FOV = 224mm, matrix = 320, 256 sagittal slices in a single slab) and diffusion (sequence = Spin-echo EPI, repetition time (TR) = 5520 ms, echo time (TE) = 89.5ms, resolution = $1.25 \times 1.25 \times 1.25$ mm³ voxels) data acquisition was done using a Siemens Skyra 3T scanner [18] and pre-processed following [3]. The diffusion tensors were fitted using DSI Studio toolbox [13]. Both T1w and DTI were decomposed in overlapping patches of 32^3 voxels centered on the foreground.

Experiments Setup We used 50,000 unpaired T1w and DTI patches randomly selected among 1,055 subjects as our training set. For the validation set and the test set, we took paired T1w and DTI patches covering the full brain of respectively 3 and 2 randomly chosen subjects. We compared our manifold-aware CycleGAN (MA-CycleGAN) method with two baselines: a manifold-aware Wasserstein GAN without cycle (MA-GAN) and a Wasserstein CycleGAN without the \log_{Id} and \exp_{Id} mappings. These baseline methods allow us to assess the impact of both the cycle consistency and the manifold mapping. We measure the quality of the generated HR DTI by computing three metrics: 1) the mean cosine similarity between the principal eigenvectors of the generated images and their ground-truth, 2) the mean squared error between the FA of the generated images and their ground-truth, and 3) the mean Log-Euclidean distance between the generated tensors and their ground-truth following Equation (3). Because the principal eigenvector’s direction is more relevant at voxels with higher anisotropy, cosine similarity is measured at three different FA thresholds taken on the ground-truth images: $FA \geq 0$ (all voxels), only voxels with $FA \geq 0.2$, and only voxels with $FA \geq 0.5$. An FA threshold near 0.2 is commonly used for tract-based analysis of white matter [19].

While the cosine similarity allows us to evaluate the precision of the predicted orientation of the generated tensors, the mean squared error on the derived FA highlights the network’s ability to estimate local diffusion anisotropy. As for the Log-Euclidean distance, it takes into account both the orientation and the anisotropy of the tensors. Furthermore, a qualitative inspection of the generated tensors and FA is performed in Figure 4.

Implementation Details Both generators are based on the Unet implementation from [16] where the convolutions have been changed to 3D convolutions. In addition, we changed the last activation layers to fit the scale of the generated data and the number of channels with respect to our inputs and outputs. For G_X , we use a sigmoid as the final activation function to generate values in the range [0,1]. In the case of G_Y , we use a hard hyperbolic tangent activation function. Furthermore, to guarantee that the synthesized tensors can be




Fig. 4. (top row) Real low-resolution FA, (second row) generated high-resolution FA, (third row) real low-resolution color orientation tensors, and (bottom row) generated high-resolution color orientation tensors. Best viewed in color.

decomposed following Equations (1) and (2), we convert the 9-channels output of G_Y into a 3×3 matrix \mathbf{Y} and make this matrix symmetric as follows: $\mathbf{Y}' = \frac{1}{2}(\mathbf{Y} + \mathbf{Y}^\top)$. For our discriminators, we employ a Resnet-18 architecture [6] where all convolutions have been changed to 3D convolutions. Both the generators and discriminators were trained for 30 epochs with the Adam optimizer [14] and a batch size of 8. A starting learning rate of 1×10^{-4} was used jointly with a reduce-on-plateau strategy. To stabilize the training of our network, we pre-trained the generators independently with 25,000 paired patches randomly sampled from 5 subjects during 10 epochs. The paired patches have been computed from aligned high-resolution structural T1w images and upsampled DTIs. The alignment was performed using FLIRT [10, 12] from FSL [11].

Fibers Orientation Analysis To evaluate the predicted tensor orientation, we compute the cosine similarity between the principal orientation of each generated tensor and its ground-truth, for FA threshold 0.0, 0.2 and 0.5. Results in Table 1 show that our method performs better than the two baselines, yielding average improvements in cosine similarity (for FA \geq 0.0, 0.2, 0.5) of 0.133, 0.085, 0.134

Table 1. Fractional anisotropy mean squared error (FA MSE), Log-Euclidean distance, and cosine similarity between principal eigenvectors of compared methods: manifold-aware Wasserstein GAN (MA-GAN), Wasserstein CycleGAN without the \log_{Id} and \exp_{Id} mappings (CycleGAN), and our manifold-aware Wasserstein CycleGAN (MA-CycleGAN). A smaller FA MSE and Log-Euclidean distance corresponds to a superior performance while a higher cosine similarity is better. Since the principal eigenvector’s direction is more relevant at voxels with higher diffusivity, we report cosine similarity at increasing FA thresholds of 0, 0.2 and 0.5.

Models	FA MSE	Log Distance	Cosine Similarity		
			FA ≥ 0	FA ≥ 0.2	FA ≥ 0.5
MA-GAN	0.0277	0.6067	0.4512	0.6369	0.6700
CycleGAN	0.0431	0.5699	0.5371	0.6717	0.7064
MA-CycleGAN (ours)	0.0172	0.5515	0.5846	0.7217	0.8041




Fig. 5. Pairs of derived high-resolution FA and their low-resolution ground-truth. The HR FA shows sharper edges and lower partial volume effect.

compared to the manifold-aware GAN without cycle and 0.048, 0.050, 0.100 compared to the CycleGAN without manifold mapping. As mentioned before, a good estimation of main diffusion orientation is generally more important at voxels with greater diffusivity. Hence, the cosine similarity for $\text{FA} \geq 0.2$ and $\text{FA} \geq 0.5$ is a better indicator of performance than for $\text{FA} \geq 0.0$. We see in Table 1 that our model’s estimation of fiber orientation improves with a higher FA threshold, reaching a similarity of 0.804 for voxels with $\text{FA} \geq 0.5$. This can be observed in Figure 6, which gives the FA MSE, Log-Euclidean distance and cosine similarity at each voxel of sagittal, coronal and axial slices from the same subject. As can be seen, orientations in regions with typically high FA, like the corpus callosum, are better predicted by our model than those in regions with lower FA.




Fig. 6. Metrics on the sagittal, coronal and axial slices between the generated HR DTI of a random evaluation subject and its interpolated ground-truth. (**top row**) FA MSE, (**middle row**) Log-Euclidean distance and (**bottom row**) cosine similarity.

FA Analysis Next, we evaluate the fractional anisotropy (FA) of generated HR DTI. FA is one of the most commonly used DT-derived metrics, thus an accurate prediction of this metric is critical. Table 1 shows the mean squared error (FA MSE) obtained by our method and the two baselines. Once more, our manifold-aware CycleGAN outperforms the manifold-aware GAN and the standard CycleGAN with an MSE of 0.172 compared to 0.0277 and 0.0431. Furthermore, we observe that the two methods using manifold mapping perform better than the standard CycleGAN. This performance gap is due to the fact that, without projecting the generated data on the $\text{SPD}(3)$ manifold using the Log-Euclidean metric, there is no guarantee that the generated tensors lie on such manifold. Consequently, the tensors eigenvalues used in the computation of FA are not strictly positive, which increases the differences between the generated FA and the ground-truth. In addition, as seen in Figures 4 and 5, generating high-

resolution DTI helps reducing partial volume effect which is known to impact subsequent analysis [15]. However, as shown in Figure 6, our model tends to underestimate the FA in white matter regions where the FA is further away from the mean value, despite the good estimation of fiber orientation.

4 Discussion and Conclusion

In this paper we proposed a novel manifold-aware CycleGAN that successfully leverages the Log-Euclidean metric and the structural information of T1w images to generate realistic high-resolution DTI. Our method outperformed the manifold-aware GAN and the standard CycleGAN architecture in terms of tensor principal orientation estimation, Log-Euclidean distance and MSE of derived FA. These results not only confirm that projecting the generated DTI on the SPD(3) manifold helps producing plausible diffusion tensors but also that the extra structural information provided by the T1w data is necessary to synthesize high-resolution DTI. Although physiological evidence remains limited, it is shown in [17] that fiber orientations have an impact on T1w image intensities. Our results further suggest that T1w images may contain information on the high-level geometry of fiber tracts, which can be learned by the network to estimate the diffusion properties and orientation. However, a deeper investigation is required to validate this hypothesis.

We believe that our method is an important contribution to medical image computing as it unlocks a vast number of applications on manifold-valued data. As future work, we plan on extending our method to other Riemannian manifolds such as the statistical manifold for orientation distribution function estimation. Furthermore we will investigate the integrity of our generated data with common downstream tasks such as tractography and fiber bundles segmentation.

References

1. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein GAN (2017)
2. Arsigny, V., Fillard, P., Pennec, X., Ayache, N.: Log-Euclidean metrics for fast and simple calculus on diffusion tensors. *Magnetic Resonance in Medicine* (2006)
3. Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., Van Essen, D.C., Jenkinson, M.: The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage* (2013)
4. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *Advances in Neural Information Processing Systems*. Neural information processing systems foundation (2014)
5. Gu, X., Knutsson, H., Nilsson, M., Eklund, A.: Generating Diffusion MRI Scalar Maps from T1 Weighted Images Using Generative Adversarial Networks. *Tech. rep.* (2019)
6. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. IEEE Computer Society (2016)

7. Huang, Z., Van Gool, L.: A riemannian network for SPD matrix learning. 31st AAAI Conference on Artificial Intelligence, AAAI 2017 (2017)
8. Huang, Z., Wu, J., Van Gool, L.: Manifold-Valued Image Generation with Wasserstein Generative Adversarial Nets. Proceedings of the AAAI Conference on Artificial Intelligence pp. 3886–3893 (2019)
9. Ionescu, C., Vantzos, O., Sminchisescu, C.: Matrix backpropagation for deep networks with structured layers. Tech. rep. (2015)
10. Jenkinson, M., Bannister, P., Brady, M., Smith, S.: Improved Optimization for the Robust and Accurate Linear Registration and Motion Correction of Brain Images. *NeuroImage* (2002)
11. Jenkinson, M., Beckmann, C.F., Behrens, T.E.J., Woolrich, M.W., Smith, S.M.: Review FSL. *NeuroImage* (2012)
12. Jenkinson, M., Smith, S.: A global optimisation method for robust affine registration of brain images. *Medical Image Analysis* (2001)
13. Jiang, H., Van Zijl, P.C.M., Kim, J., Pearlson, G.D., Mori, S.: DtiStudio: Resource program for diffusion tensor computation and fiber bundle tracking (2005)
14. Kingma, D.P., Ba, J.L.: Adam: A method for stochastic optimization. In: 3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings. International Conference on Learning Representations, ICLR (2015)
15. Pfefferbaum, A., Sullivan, E.V.: Increased brain white matter diffusivity in normal adult aging: relationship to anisotropy and partial voluming. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* (2003)
16. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) (2015)
17. Schyboll, F., Jaekel, U., Weber, B., Neeb, H.: The impact of fibre orientation on t1-relaxation and apparent tissue water content in white matter. *Magnetic Resonance Materials in Physics, Biology and Medicine* **31**(4), 501–510 (2018)
18. Sotiropoulos, S.N., Jbabdi, S., Xu, J., Andersson, J.L., Moeller, S., Auerbach, E.J., Glasser, M.F., Hernandez, M., Sapiro, G., Jenkinson, M., Feinberg, D.A., Yacoub, E., Lenglet, C., Van Essen, D.C., Ugurbil, K., Behrens, T.E.: Advances in diffusion MRI acquisition and processing in the Human Connectome Project. *NeuroImage* (2013)
19. Taoka, T., Morikawa, M., Akashi, T., Miyasaka, T., Nakagawa, H., Kiuchi, K., Kishimoto, T., Kichikawa, K.: Fractional anisotropy–threshold dependence in tract-based diffusion tensor analysis: evaluation of the uncinate fasciculus in alzheimer disease. *American Journal of Neuroradiology* **30**(9), 1700–1703 (2009)
20. Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K.: The WU-Minn Human Connectome Project: An overview. *NeuroImage* (2013)
21. Yi, X., Walia, E., Babyn, P.: Generative adversarial network in medical imaging: A review. *Medical Image Analysis* **58** (2019)
22. Zhong, J., Wang, Y., Li, J., Xue, X., Liu, S., Wang, M., Gao, X., Wang, Q., Yang, J., Li, X.: Inter-site harmonization based on dual generative adversarial networks for diffusion tensor imaging: Application to neonatal white matter development. *BioMedical Engineering Online* **19** (2020)
23. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. Proceedings of the IEEE International Conference on Computer Vision (2017)