

When Diffusion MRI Meets Diffusion Model: A Novel Deep Generative Model for Diffusion MRI Generation

Xi Zhu¹, Wei Zhang¹, Yijie Li¹, Lauren J. O'Donnell², and Fan Zhang¹(✉)

¹ University of Electronic Science and Technology of China, Chengdu, China

² Harvard Medical School and Brigham and Women's Hospital, Boston, USA
fan.zhang@uestc.edu.cn

Abstract. Diffusion MRI (dMRI) is an advanced imaging technique characterizing tissue microstructure and white matter structural connectivity of the human brain. The demand for high-quality dMRI data is growing, driven by the need for better resolution and improved tissue contrast. However, acquiring high-quality dMRI data is expensive and time-consuming. In this context, deep generative modeling emerges as a promising solution to enhance image quality while minimizing acquisition costs and scanning time. In this study, we propose a novel generative approach to perform dMRI generation using deep diffusion models. It can generate high dimension (4D) and high resolution data preserving the gradients information and brain structure. We demonstrated our method through an image mapping task aimed at enhancing the quality of dMRI images from 3T to 7T. Our approach demonstrates highly enhanced performance in generating dMRI images when compared to the current state-of-the-art (SOTA) methods. This achievement underscores a substantial progression in enhancing dMRI quality, highlighting the potential of our novel generative approach to revolutionize dMRI imaging standards.

Keywords: Diffusion model · Diffusion MRI · RISH feature.

1 Introduction

Diffusion MRI (dMRI) is an advanced neuroimaging tool to characterize the underlying brain tissue microstructure [1] and is widely used for studying the brains [2, 3]. Currently, there is an increasing interest in high-quality dMRI data for better resolutions and enhanced tissue contrast such as dMRI data from a 7T scanner [4–7]. However, acquiring such high-quality dMRI data necessitates advanced MRI scanners and/or acquisition protocols, which are not always accessible and thus remain impractical in real-world applications.

Generation of dMRI using machine learning offers high promise to improve image quality while reducing acquisition costs and scanning time. This task generally involves image-to-image translation to learn a mapping from low-quality

(source) to high-quality (target) data, which can subsequently predict (or generate) high-quality data when only low-quality data is available. Traditional methods have used techniques such as random forest [8, 9] to map voxel patches from low-quality to high-quality data. With the advances in deep learning, many studies have used deep networks for dMRI generation [10–13]. For instance, Karayumak et al. introduced a convolutional neural network (CNN) approach [11] and Ranjan Jha et al. employed a more sophisticated generative-adversarial network (GAN) approach [13] to generate high-quality dMRI data from 3T to 7T.

Recently, diffusion models have demonstrated remarkable results for generative modeling in medical imaging [14], which may provide a powerful tool for dMRI data generation. In brief, a diffusion model comprises a forward diffusion stage, where input data is progressively perturbed by Gaussian noise, followed by a reverse diffusion stage aimed at gradually reverting the process to recover the original input. The Denoising Diffusion Probabilistic Model (DDPM) [15] is one representative diffusion model for image generation. Many variations of DDPM have been proposed. For example, the same research group of DDPM introduced the concept of classifier-free guidance [16] to remove complex classifiers and make the model simpler. The Denoising Diffusion Implicit Model (DDIM) [17] skips some steps to accelerate the sampling speed in DDPM. The Latent Diffusion Model (LDM) [18] made DDPM work on latent space, which can handle high-resolution images to increase computational efficiency. Currently, diffusion models show great advantages in medical imaging, such as anomaly detection [19], signal reconstruction [20], and image generation [21]. In the dMRI field, one recent study has successfully used the diffusion model for data denoising [22]. Yet, there is no work for dMRI generation using diffusion models.

The application of diffusion models for dMRI generation is a challenging task due to the uniqueness of dMRI data. First, dMRI is a unique, multi-dimensional image dataset that describes not only the strength but also the orientation of water diffusion. Applying diffusion models to high-dimensional data presents significant challenges, leading recent research to focus primarily on slices or single volumes. This approach overlooks the crucial 3D orientation information offered by dMRI, which is essential for understanding the complex spatial relationships and orientations in the data. Secondly, training diffusion models for data quality enhancement necessitates the availability of both standard and high-quality data, such as datasets from 3T and 7T scanners. However, acquiring high-quality data is challenging, often resulting in limited data to restrict the potential for using generative models for a large scale data analysis.

In this paper, we present a novel generative approach for dMRI generation using deep diffusion models. To the best of our knowledge, this is the first work to apply diffusion models specifically for enhancing the quality of dMRI data. Our method has the following contributions: 1) proposing using LDM for the high-quality 7T rotation invariant spherical harmonic (RISH) features generation and reconstructing the 4D dMRI data, 2) designing a transfer learning strategy for autoencoder training to address the scarcity of high-quality 7T data, and 3) a super-resolution module to remove resolution differences. We demonstrated




Fig. 1. Overview of the proposed method.

our method through an image mapping task aimed at enhancing the quality of dMRI images from 3T to 7T. Our approach demonstrates highly enhanced performance in generating dMRI images when compared to several compared methods, indicating a notable advancement in dMRI quality improvement.

2 Method

Fig. 1 gives an overview of our method. First, RISH features [23, 24] are computed for an efficient and compact representation of the input 3T and 7T dMRI data. Next, we train a deep generative model to learn the RISH features of 7T. This step includes: 1) two autoencoders that, respectively, learn latent features of the 3T and 7T RISH features, where we design a fine-tuning strategy to address the scarcity of high-quality 7T training data; and 2) a classifier-free guidance DDPM to generate 7T-like latent features from 3T, where we introduce a super-resolution module to enable simultaneous dMRI signal generation and spatial resolution enhancement. Finally, during inference, the RISH features of a testing 3T dataset are encoded into the latent space using the 3T encoder, followed by a DDIM process to generate 7T-like RISH features for reconstructing a high-quality 7T dMRI dataset.

2.1 dMRI Datasets

We used the dMRI data provided in the Human Connectome Project (HCP) [25]. In total, data from 1065 subjects were used, of which 171 had both 3T and 7T dMRI data, and 894 had only 3T data. The acquisition parameters of 3T dMRI data were: $TE = 89.5ms$, $TR = 520ms$, and voxel size = $1.25 \times 1.25 \times 1.25mm^3$, 18 baseline images, and 270 diffusion-weighted images distributed evenly at $b = 1000/2000/3000 s/mm^3$; and those of the 7T data were: $TE = 71.2ms$,

$TR = 7000ms$, and voxel size= $1.05 \times 1.05 \times 1.05mm^3$, 15 baseline images, and 128 diffusion-weighted images distributed evenly at $b = 1000/2000 s/mm^3$. The provided dMRI data has been preprocessed as in [26]. In our study, for simplicity, we used only the single shell $b = 1000$ data in both 3T and 7T data.

2.2 dMRI signal representation and reconstruction using RISH

In dMRI, the signal S of each voxel can be represented in a basis of spherical harmonics (SH) [27]: $S \approx \sum_i \sum_j C_{ij} Y_{ij}$, where C_{ij} is the coefficient of SH basis function Y_{ij} at order i and degree j . Then, from the SH coefficients, the RISH features at each order i can be computed as follows:

$$\|C_i\|^2 = \sum_{j=1}^{2i+1} (C_{ij})^2 \quad (1)$$

One of the benefits of the RISH features is that they can be appropriately scaled to modify the dMRI signals without changing the principal directions of the fibers [23]. In addition, the RISH features give a compact and uniform representation of the dMRI data regardless of the number of gradient directions. In our study, we computed the RISH features for each subject’s 3T and 7T images with SH orders of $i = \{0, 2, 4\}$ as suggested in [28].

For dMRI data generation from 3T to 7T, during training stage (see Section 2.3), we can learn 7T-like RISH features by computing so-called scale maps between the two datasets, as:

$$\lambda_i = \sqrt{\frac{\|C_i\|_{7T}^2}{\|C_i\|_{3T}^2 + \tau}} \quad (2)$$

where τ is a constant with a very small value. Then, during inference when only 3T data is available, the scale maps can be predicted via the learned models, which can be subsequently applied to the SH coefficients from 3T images to generate 7T-like RISH features, as follows:

$$\hat{C}_{ij} = \lambda_i C_{ij} \quad (3)$$

where $\hat{C}_{ij}$ is the predicted SH coefficients of 7T. Finally, a high-quality 7T dMRI dataset can be generated by reconstructed dMRI signals through Eq. (1) with the predicted $\hat{C}_{ij}$ and the SH basis function Y_{ij} .

2.3 Latent diffusion model

To fully leverage the 3D properties of the dMRI RISH features and tackle the issue of limited availability of high-quality dMRI data, we propose a new architecture based on Latent Diffusion Models (LDM) complemented by a fine-tuning strategy. In detail, we use the Vector Quantised-Variational AutoEncoder (VQ-VAE) [29] to compress the whole brain image to the latent space. VQ-VAE

quantifies the latent representation of images to get a better latent presentation. To address the differences between 3T and 7T MRI datasets effectively, we train two separate VQ-VAE models, one for each dataset type. However, the VQ-VAE model struggles to produce satisfactory outcomes for 7T data due to the limited volume of available high-quality data. To overcome this limitation, we introduce the application of transfer learning. We first train a model extensively on the abundant 3T dataset and subsequently fine-tune this model using 7T data. Both of the models use MAE loss and quantization loss during the training process.

Then, the data generation via the diffusion process is performed in the latent space and takes the output x of the VQ-VAE's encoder as input. The process can be generally divided into two parts: the forward noising process and the backward denoising process. The forward noising process q is defined as follows:

$$q(x_t|x_{t-1}) = N(x_t; \sqrt{1 - \beta_t}x_{t-1}; \beta_t I) \quad (4)$$

Where x_t is the noisy latent features that are obtained by an iterative process of noise addition, $\{\beta_1, \beta_2, \dots, \beta_t, \dots, \beta_T\}$ is a series constant, and $t \in \{0, \dots, T\}$ is a moment during the noise addition process. Noisy features at the moment t can be written as:

$$x_t = \sqrt{\alpha_t}x_0 + \sqrt{1 - \alpha_t}\epsilon, \text{ with } \epsilon \in N(0, I) \quad (5)$$

with $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$. The denoising process p_θ relies on a U-Net to predict x_{t-1} from x_t by optimizing the U-Net's parameters θ . It can be given as

$$p_\theta(x_{t-1}|x_t) = N(x_{t-1}; \mu_\theta(x_t, t); \sum_{\theta} (x_t, t)) \quad (6)$$

The U-Net can be denoted as ϵ_θ , and MSE loss is used to train this model, as:

$$L = \|\epsilon - \epsilon_\theta(\sqrt{\alpha_t}x_0 + \sqrt{1 - \alpha_t}\epsilon)\|_2^2, \text{ with } \epsilon \in N(0, I) \quad (7)$$

In the sampling process, we can encode the latent features by adding noise based on the model's output at step t :

$$x_{t+1} = x_t + \sqrt{\alpha_{t+1}}[(\sqrt{\frac{1}{\alpha_t}} - \sqrt{\frac{1}{\alpha_{t+1}}})x_t + (\sqrt{\frac{1}{\alpha_{t+1}}} - 1 - \sqrt{\frac{1}{\alpha_t}} - 1)\epsilon_\theta(x_t, t)] \quad (8)$$

In the above diffusion model, the class labels (3T and 7T) are used to control the direction of diffusion model's generation, they can be encoded to class-embeddings and introduced to the U-Net backbone through the cross-attention mechanism implementing $Attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d}}\right) \cdot V$, with

$$Q = W_Q^{(i)} \cdot \varphi_i(z_t), K = W_K^{(i)} \cdot \tau_\theta(c), V = W_V^{(i)} \cdot \tau_\theta(c). \quad (9)$$

with c is the class label, τ_θ represents a specific encoder that encodes c to embedding and φ_i denotes a intermediate representation of the U-Net.

The generation of controlled diffusion models is divided into the sum of an unconditional generation process $\epsilon_\theta(x_t)$ and a conditional generation process

$\epsilon_\theta(x_t, c)$. So, we turn certain labels into uncertain ones with a probability which is set to a hyperparameter in the training process. In the inference process, we turn the 3T latent features into the 7T features with class-embedding guidance and the U-Net’s prediction can be given as:

$$\bar{\epsilon}_\theta(x_t, c) = (1 + \omega)\epsilon_\theta(x_t, c) - \omega\epsilon_\theta(x_t) \quad (10)$$

the ω represents the guidance scale of class embedding.

Finally, we use the dataset generated by LDM to train the super-resolution module located at the end of 7T VQ-VAE. The architecture of it includes two convolution layers and middle residual layers. We apply a similar training strategy with SR-CNN and use the MSE loss to optimize module.

3 Experimental Comparisons

We compare the performance of the proposed method with CNN-based [11] and GAN-based network architecture [18]. The CNN-based method designed a deep 3D convolutional network, and the GAN-based method contained an autoencoder with an attention mechanism trained by an adversarial framework. These methods also used 4 orders of RISH features as input and reconstruction of each method was the same. The synthesis quality was evaluated using normalized mean squared error (NMSE), and structural similarity index (SSIM) across multiple scales. 17 subjects with both 3T and 7T data were randomly selected and left for testing, while the remaining were used for training and validation of the autoencoders and the DDPM. There were 1065 subjects’ data for 3T VQ-VAE training and 171 7T data for fine-tuning 7T VQ-VAE. Finally, 342 subjects both having 7T and 3T data were used to train LDM. All of these datasets were divided into training sets and test sets at a ratio of 9 : 1, and every test set included 17 test subjects. All metrics are calculated over 3D volumes to ensure comprehensive analysis.


The implementations of VQ-VAE and LDM were done using Pytorch [30] and MONAI [31] framework. All the RISH features were downsampled to $96 \times 96 \times 96$ before inputting to the LDM, and so were other test methods. For the hyperparameters, we set the number of embedding dimensions to 32 and the number of embeddings to 256 for training VQ-VAE. As for LDM, we choose different guidance scales and levels of noise addition for RISH features of each order during sampling steps. We used the AdamW optimizer with a learning rate of 1×10^{-4} and set training epochs to 200 and 1000 for VQ-VAE and LDM respectively. For the detailed architecture, we used attention heads at the third and fourth layers of U-Net used in LDM to predict noise. All computation was conducted on RTX 3090 GPUs.

4 Results

SOTA comparison. Table 1 gives the mean NMSE and SSIM for the RISH features and the FA images, where we can see that our method outperforms the

Table 1. Comparison of NMSE and SSIM in RISH and FA across different methods.

NMSE↓:	RISH_L0	RISH_L2	RISH_L4	FA
CNN	0.126 ± 0.014	0.143 ± 0.011	0.495 ± 0.107	0.053 ± 0.007
GAN	0.129 ± 0.029	0.427 ± 0.051	1.652 ± 0.360	0.118 ± 0.009
Diffusion	0.105 ± 0.026	0.102 ± 0.017	0.158 ± 0.031	0.044 ± 0.008
SSIM↑:				
CNN	0.889 ± 0.008	0.959 ± 0.006	0.956 ± 0.016	0.958 ± 0.006
GAN	0.915 ± 0.012	0.893 ± 0.010	0.943 ± 0.004	0.902 ± 0.010
Diffusion	0.922 ± 0.009	0.961 ± 0.007	0.967 ± 0.002	0.966 ± 0.007

**Fig. 2.** Results for the RISH features and FA generated by different methods.

other two methods in quantitative metrics. Fig. 2 provides a visual comparison of various RISH features produced by the different methods and the FA images against the target data. Our approach is distinguished by producing images that more closely resemble the ground truth than the other compared methods. The CNN-based method tends to uniformly increase the intensity of all voxels across the input images, leading to a loss of contrast information between different regions. Meanwhile, the GAN-based method fails to preserve some of the structural details in higher-order L2 and L4 RISH features. Figure 3 illustrates the difference maps of FA between the predicted and target data, demonstrating




Fig. 3. Difference maps of FA images.

Table 2. Ablation study results

Fine-tuning	Super-resolution	NMSE↓	SSIM↑
-	-	0.046 ± 0.008	0.962 ± 0.008
✓	-	0.044 ± 0.008	0.966 ± 0.007
✓	✓	0.042 ± 0.004	0.967 ± 0.007

that our method achieves the closest resemblance to the target data, further highlighting its accuracy in generating high-quality dMRI images.

Ablation experiment. To explore the effects of the fine-tuning and the super-resolution module proposed in our method, we conducted ablation studies from these two aspects. We adopted a network the same as the 3T VQ-VAE and trained it on the collection of the 7T datasets. The DWI data were reconstructed with the same process which excluded the super-resolution module, and then compared the FA images with the NMSE. For the experiments examining the super-resolution module, we first applied B-spline interpolation to upsample the 3T data to the same resolution as 7T and register it to the 7T space. Following this, we enhance 3T data with our method and data acquired by 7T scanner. Table. 2 presents the ablation study results. We can observe a notable improvement in performance post-fine-tuning. Moreover, the introduction led to a further reduction in the NMSE of the FA images.

To investigate the impact of the fine-tuning and super-resolution components within our method, we carried out the following ablation studies. First, we utilized a network identical to the 3T VQ-VAE and trained it exclusively with the 7T dataset without fine-tuning. Second, we performed a method without using the super-resolution module, instead using a B-spline interpolation to upscale the 3T data to match the resolution of 7T data. The same quantitative measures NMSE and SSIM were used for experimental comparison. Table 2 shows the comparison results, showing a significant improvement in using the fine-tuning process and the super-resolution module.

5 Conclusion

We present a novel framework that leverages the latent diffusion model and rotation invariant spherical harmonic to generate high-quality dMRI data. We applied the proposed method for image generation on the HCP dataset and successfully generated the 7T-like dMRI image from 3T. Our method largely outperforms current SOTA methods in generating dMRI images, marking a major advancement in dMRI quality enhancement. This underscores the potential of our innovative generative method to transform dMRI imaging standards.

Acknowledgments. This work is in part supported by the National Key R&D Program of China (No. 2023YFE0118600), the National Natural Science Foundation of China (No. 62371107), and the National Institutes of Health (R01MH125860, R01MH119222, R01MH132610, R01NS125781).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Bassar, P.J., Mattiello, J., LeBihan, D.: MR diffusion tensor spectroscopy and imaging. *Biophysical journal* **66**(1), 259–267 (1994)
2. Pannek, K., Scheck, S.M., Colditz, P.B., Boyd, R.N., Rose, S.E.: Magnetic resonance diffusion tractography of the preterm infant brain: a systematic review. *Developmental Medicine & Child Neurology* **56**(2), 113–124 (2014)
3. Zhang, F., Daducci, A., He, Y., Schiavi, S., Seguin, C., Smith, R.E., Yeh, C.H., Zhao, T., O'Donnell, L.J.: Quantitative mapping of the brain's structural connectivity using diffusion MRI tractography: A review. *Neuroimage* **249**, 118870 (2022)
4. Chilla, G.S., Tan, C.H., Xu, C., Poh, C.L.: Diffusion weighted magnetic resonance imaging and its recent trend—a survey. *Quantitative imaging in medicine and surgery* **5**(3), 407 (2015)
5. Sotiropoulos, S.N., Hernández-Fernández, M., Vu, A.T., Andersson, J.L., Moeller, S., Yacoub, E., Lenglet, C., Ugurbil, K., Behrens, T.E., Jbabdi, S.: Fusion in diffusion MRI for improved fibre orientation estimation: An application to the 3T and 7T data of the Human Connectome Project. *Neuroimage* **134**, 396–409 (2016)
6. Ramos-Llordén, G., Ning, L., Liao, C., Mukhometzianov, R., Michailovich, O., Set-sompop, K., Rathi, Y.: High-fidelity, accelerated whole-brain submillimeter in vivo diffusion MRI using gSlider-spherical ridgelets (gSlider-SR). *Magnetic resonance in medicine* **84**(4), 1781–1795 (2020)
7. Vu, A.T., Auerbach, E., Lenglet, C., Moeller, S., Sotiropoulos, S.N., Jbabdi, S., Andersson, J., Yacoub, E., Ugurbil, K.: High resolution whole brain diffusion imaging at 7T for the Human Connectome Project. *Neuroimage* **122**, 318–331 (2015)
8. Alexander, D.C., Zikic, D., Zhang, J., Zhang, H., Criminisi, A.: Image quality transfer via random forest regression: applications in diffusion mri. In: *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2014*. pp. 225–232. Springer (2014)

9. Nedjati-Gilani, G.L., Schneider, T., Hall, M.G., Wheeler-Kingshott, C.A., Alexander, D.C.: Machine learning based compartment models with permeability for white matter microstructure imaging. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2014. pp. 257–264. Springer (2014)
10. Tanno, R., Worrall, D.E., Ghosh, A., Kaden, E., Sotiropoulos, S.N., Criminisi, A., Alexander, D.C.: Bayesian image quality transfer with CNNs: exploring uncertainty in dMRI super-resolution. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2017. pp. 611–619. Springer (2017)
11. Cetin Karayumak, S., Kubicki, M., Rath, Y.: Harmonizing Diffusion MRI Data Across Magnetic Field Strengths. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2018. pp. 116–124. Springer (2018)
12. Hirte, A.U., Platscher, M., Joyce, T., Heit, J.J., Tranvinh, E., Federau, C.: Realistic generation of diffusion-weighted magnetic resonance brain images with deep generative models. *Magnetic Resonance Imaging* **81**, 60–66 (2021)
13. Jha, R.R., Kumar, B.R., Pathak, S.K., Bhavsar, A., Nigam, A.: TrGANet: Transforming 3T to 7T dMRI using Trapezoidal Rule and Graph based Attention Modules. *Medical Image Analysis* **87**, 102806 (2023)
14. Kazerouni, A., Aghdam, E.K., Heidari, M., Azad, R., Fayyaz, M., Hacıhaliloglu, I., Merhof, D.: Diffusion models in medical imaging: A comprehensive survey. *Medical Image Analysis* p. 102846 (2023)
15. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
16. Ho, J., Salimans, T.: Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598* (2022)
17. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502* (2020)
18. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
19. Wolleb, J., Bieder, F., Sandkühler, R., Cattin, P.C.: Diffusion models for medical anomaly detection. In: International Conference on Medical image computing and computer-assisted intervention. pp. 35–45. Springer (2022)
20. Ozturkler, B., Liu, C., Eckart, B., Mardani, M., Song, J., Kautz, J.: Smrd: Sure-based robust mri reconstruction with diffusion models. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 199–209. Springer (2023)
21. Pinaya, W.H., Tudosi, P.D., Dafflon, J., Da Costa, P.F., Fernandez, V., Nachev, P., Ourselin, S., Cardoso, M.J.: Brain imaging generation with latent diffusion models. In: MICCAI Workshop on Deep Generative Models. pp. 117–126. Springer (2022)
22. Xiang, T., Yurt, M., Syed, A.B., Setsompop, K., Chaudhari, A.: DDM2: Self-Supervised Diffusion MRI Denoising with Generative Diffusion Models. *arXiv preprint arXiv:2302.03018* (2023)
23. Mirzaalian, H., Ning, L., Savadjiev, P., Pasternak, O., Bouix, S., Michailovich, O., Grant, G., Marx, C.E., Morey, R.A., Flashman, L.A., et al.: Inter-site and inter-scanner diffusion MRI data harmonization. *NeuroImage* **135**, 311–323 (2016)
24. Karayumak, S.C., Bouix, S., Ning, L., James, A., Crow, T., Shenton, M., Kubicki, M., Rath, Y.: Retrospective harmonization of multi-site diffusion MRI data acquired with different acquisition parameters. *Neuroimage* **184**, 180–200 (2019)

25. Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K., Consortium, W.M.H., et al.: The WU-Minn human connectome project: an overview. *Neuroimage* **80**, 62–79 (2013)
26. 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., et al.: The minimal preprocessing pipelines for the Human Connectome Project. *Neuroimage* **80**, 105–124 (2013)
27. Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R.: Regularized, fast, and robust analytical Q-ball imaging. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* **58**(3), 497–510 (2007)
28. De Luca, A., Karayumak, S.C., Leemans, A., Rath, Y., Swinnen, S., Gooijers, J., Clauwaert, A., Bahr, R., Sandmo, S.B., Sochen, N., et al.: Cross-site harmonization of multi-shell diffusion MRI measures based on rotational invariant spherical harmonics (RISH). *NeuroImage* **259**, 119439 (2022)
29. Van Den Oord, A., Vinyals, O., et al.: Neural discrete representation learning. *Advances in neural information processing systems* **30** (2017)
30. Imambi, S., Prakash, K.B., Kanagachidambaresan, G.: PyTorch. Programming with TensorFlow: Solution for Edge Computing Applications pp. 87–104 (2021)
31. Pinaya, W.H., Graham, M.S., Kerfoot, E., Tudosi, P.D., Dafflon, J., Fernandez, V., Sanchez, P., Wolleb, J., da Costa, P.F., Patel, A., et al.: Generative AI for medical imaging: extending the monai framework. *arXiv preprint arXiv:2307.15208* (2023)