



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Phy-Diff: Physics-guided Hourglass Diffusion Model for Diffusion MRI Synthesis^{*}

Juanhua Zhang¹, Ruodan Yan², Alessandro Perelli¹, Xi Chen³, and Chao Li^{1,2,4,5}✉

¹ School of Science and Engineering, University of Dundee

² Department of Applied Mathematics and Theoretical Physics, University of Cambridge

³ Department of Computer Science, University of Bath

⁴ Department of Clinical Neurosciences, University of Cambridge

⁵ School of Medicine, University of Dundee

c1647@cam.ac.uk

Abstract. Diffusion MRI (dMRI) is an important neuroimaging technique with high acquisition costs. Deep learning approaches have been used to enhance dMRI and predict diffusion biomarkers through under-sampled dMRI. To generate more comprehensive raw dMRI, generative adversarial network based methods are proposed to include b-values and b-vectors as conditions, but they are limited by unstable training and less desirable diversity. The emerging diffusion model (DM) promises to improve generative performance. However, it remains challenging to include essential information in conditioning DM for more relevant generation, i.e., the physical principles of dMRI and white matter tract structures. In this study, we propose a physics-guided diffusion model to generate high-quality dMRI. Our model introduces the physical principles of dMRI in the noise evolution in the diffusion process and introduces a query-based conditional mapping within the diffusion model. In addition, to enhance the anatomical fine details of the generation, we introduce the XTRACT atlas as a prior of white matter tracts by adopting an adapter technique. Our experiment results show that our method outperforms other state-of-the-art methods and has the potential to advance dMRI enhancement.

Keywords: Diffusion MRI · Image synthesis · Hourglass diffusion model · Physics informed deep learning.

1 Introduction

Diffusion MRI (dMRI) is an essential neuroimaging modality due to its sensitivity in probing white matter fibers and characterizing neurological diseases [13,26]. However, acquiring high-quality dMRI for sufficient fidelity remains a challenge due to expensive acquisition costs. For example, the neurite orientation dispersion and density imaging (NODDI) sequence requires 90 diffusion

^{*} J. Zhang and R. Yan contribute equally in this work.

gradients, which is time-consuming and infeasible for patients with severe conditions [27]. There is an increasing need to develop effective algorithms for dMRI enhancement to reduce acquisition costs.

Deep learning has achieved encouraging performance in medical image synthesis. Based on conventional MRI, various frameworks have been proposed with success to synthesize multimodal MRI [11], super-resolve images [16], and generate brain connectivities [14,25], supporting the feasibility of deep learning dMRI enhancement. Several studies have proposed frameworks to estimate dMRI parametric maps from undersampled high-resolution dMRI. For instance, a CNN-based model was proposed using a few dMRI directions to generate diffusion kurtosis imaging (DKI) and NODDI maps [9]. Further, other methods were proposed based on Graph Transformer and atlas-aided U-Net to synthesize dMRI biomarkers [3,12]. However, these methods use sparsely sampled dMRI to generate specific parametric maps without leveraging b-values and b-vectors from raw dMRI for more comprehensive image generation. Therefore, these methods require retraining when applied to generating other parametric maps.

To mitigate this challenge, generative adversarial network (GAN) based methods are proposed using b-values and b-vectors as conditions, promising to generate comprehensive raw dMRI. The generated images are used to calculate downstream parametric maps [20]. Despite the success, GAN-based architectures are often limited by their inability to balance fine details and visual quality [1], which could be suboptimal for clinical diagnosis. In addition, GANs capture less diversity in large-volume generation [5], which cannot meet real-world clinical needs to address patient individuality. In addition, GANs often demonstrate unstable training, challenged by gradient vanishing and mode collapse [2].

As a type of likelihood-based model, diffusion models (DM) promise to achieve better generation performance while addressing the limitations of GANs [5]. Mounting research has shown that DMs are successful in generating high-quality, complex, and diverse images [10,17,19], with properties of fixed training targets and easy scalability [15]. Despite the advantages, there remains an unmet need to enhance dMRI effectively due to several challenges: 1) dMRI is distinct from conventional MRI in underlying physics principles. Therefore, generating dMRI may require additional physics-informed modeling beyond image geometry. However, typical DMs only add or remove random Gaussian noise in the diffusion process; 2) It remains a significant challenge for DMs to preserve the anatomical structures that are essential to the efficacy of dMRI in assessing tissue microstructures and characterizing diseases. This is particularly important for the white matter tract regions that dMRI mainly reflects; 3) It is desirable to improve DMs to generate dMRI in a controllable manner with multiple conditions required, which is essential for restoring anatomical details in dMRI [5].

We propose a DM-based generative model tailored for dMRI, promoting image translation from the baseline images without gradient applied, i.e., b_0 images. Specifically, we introduce a physics-guided noise evolution and a query-based condition mapping. To better preserve the structural integrity of the white matter tracts in the synthesized dMRI, we introduce local conditions of the

XTRACT atlas, with 42 anatomical tracts constructed from more than 1,000 subjects [24]. Our main contributions include:

- We develop a physics-guided diffusion model to synthesize high-quality dMRI. As far as we know, this is the first attempt to generate raw dMRI using the conditional diffusion model.
- We propose a physics-guided noise evolution process in the forward diffusion to reflect the noise characteristics inherent in dMRI, which could improve on traditional DMs adding Gaussian noise.
- We design a query-based conditional mapping in the reverse diffusion process, incorporating q-space sampling to enable continuous diffusion gradient mapping, allowing for better directional dMRI synthesis at a given gradient.
- We introduce XTRACT atlas [24] as anatomical prior in conditioning DM for preserving the fine details of tract. Inspired by the T2I-Adapter [18], we improve the model adaptability for more controllable image generation while allowing to incorporate multiple conditions.

According to experimental results, our model achieves state-of-the-art performance and could serve as a promising tool for dMRI enhancement.

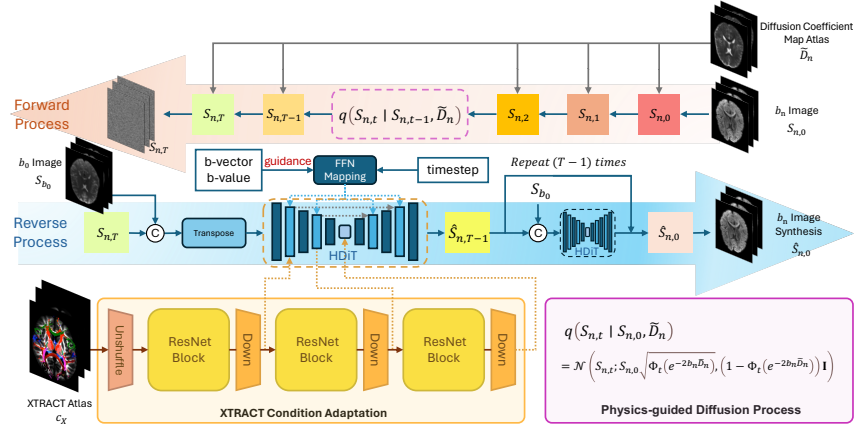


Fig. 1. Model architecture. The physics information ADC atlas \tilde{D}_n interacts with the original image b_n during the forward process by the function $q(S_{n,t} | S_{n,0}, \tilde{D}_n)$. After timesteps T , the image turns into the $S_{n,T}$. During the reverse process, HDiT predicts the added noise in each timestep in the pixel space, while using the b-vector, b-value as conditions via query-based conditional mapping. After training the noise prediction network, the XTRACT adapter is used to incorporate tractography structure c_X . By sampling from the learned pattern, the synthesized image $\hat{S}_{n,0}$ is obtained.

2 Methodology

Overview. Fig. 1 illustrates the model architecture. We propose a physics-guided diffusion process based on the framework of Denoising Diffusion Probabilistic Models (DDPM) [10], including a physics-guided noise evolution process and a query-based conditional mapping mechanism in the inference process. In the noise evolution process, Phy-Diff incorporates Apparent Diffusion Coefficient (ADC) maps to learn the underlying noise representation in the synthesized dMRI data (b_n) (see 2.1). During the reverse process, we employ a query-based conditional mapping technique that incorporates q-space sampling to guide the generation of b_n under the Hourglass Diffusion Transformer (HDiT) backbone, facilitating the more precise directional dMRI synthesis (see 2.2). Further, a XTRACT condition adaptation is designed to provide anatomical conditions that guide the structural integrity of tracts in synthesized dMRI (see 2.3).

2.1 Physics-guided Noise Evolution

The proposed physics-guided noise evolution process learn the underlying noise representation by leveraging an ADC atlas, where atlas referring to an average map. We first consider x_t as the t -th image representation, and the approximate posterior $q(x_t | x_0)$, representing the *diffusion process*, can be expressed as:

$$q(x_t | x_0) = \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t}x_0, (1 - \bar{\alpha}_t)\mathbf{I}) \quad (1)$$

where, x_t is the noisy image at the timestep t , x_0 is the image that is expected to be generated, and $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$, and α_i are hyper-parameters relevant to the variance. Thus, x_t can be computed by:

$$x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, \quad \text{with } \epsilon \sim \mathcal{N}(0, \mathbf{I}) \quad (2)$$

where ϵ represents the Gaussian noise schedule. By gradually increasing the proportion of noise and decreasing the proportion of content on the image, the final image will become completely noisy. The backbone network of DM gradually removes noise by predicting the noise changes at each timestep, ultimately restoring the original image.

However, the denoising mechanism of the standard DDPM may be more suitable for natural image generation, by multiplying the Gaussian noise with the uniform distributed variance, less considering the specific structure of dMRI. To address this, we incorporate the diffusion physics inherent in dMRI into the process. Specifically, the diffusion physics information is integrated into the forward phase of DDPM:

$$S_n = S_0 e^{-b_n \mathcal{D}_n} + R(\mathcal{D}_n) \quad (3)$$

where \mathcal{D}_n represents diffusion coefficient, b_n is an arbitrary b-value, and S_n and S_0 denote the signal intensities corresponding to the diffusion-weighted images under b-values b_n and b_0 , respectively. In order to approach the formula of DDPM

as closely as possible, we developed a residual term $R(\mathcal{D}_n)$ to effectively correct the modeling errors caused by signal errors.

We then integrate this physical process (Eq. 3) into the diffusion model architecture by assuming $\bar{\alpha}_t = \Phi_t(e^{-2b_n\mathcal{D}_n})$ and $R(\mathcal{D}_n) = \sqrt{1 - \Phi_t(e^{-2b_n\mathcal{D}_n})}\epsilon$, where ϵ is the residual schedule. Specifically, $\Phi_t(\cdot)$ refers to a transformation used to map the range of the function argument, making the diffusion process more stable. The goal is to solve for S_n . Therefore, we relax the restriction on \mathcal{D}_n and estimate that in the original equation as $\tilde{\mathcal{D}}_n$ by introducing ADC atlas, which is expressed as:

$$\tilde{\mathcal{D}}_n = \sum_{i=1}^N \mathcal{D}_n^i = \frac{N \ln S_0 - \ln \prod_{i=1}^N S_n^i}{b_n} \quad (4)$$

where N denotes the number of all dMRI images $\{S_n^0, S_n^1, \dots, S_n^N\}$ under the same b-value b_n , and S_0 refers to the dMRI image under $b = 0$.

Consequently, our physics-guided noising mechanism is formulated as:

$$S_{n,t} = S_{n,0} \sqrt{\Phi_t(e^{-2b_n\tilde{\mathcal{D}}_n})} + \sqrt{1 - \Phi_t(e^{-2b_n\tilde{\mathcal{D}}_n})}\epsilon \quad (5)$$

The reverse process is also similar. Ultimately, the simplified form of the goal we need to optimize should be:

$$\mathcal{L}_{\text{simple}}(\theta) = \mathbb{E}_{\epsilon \sim \mathcal{N}(0,1)} \left[\left\| \epsilon - D_\theta \left(S_{n,t}, t, \vec{b}_n, p_s, b_n \right) \right\|_2^2 \right] \quad (6)$$

where, t is the diffusion timestep. In addition, function $D_\theta(\cdot)$ represents the backbone neural network model for noise prediction, and \vec{b}_n, p_s as well as b_n are the b-vector, slice index and the b-value guidance, which will all be explained in detail in the following text.

2.2 Query-based Conditional Mapping

We introduce a query-based conditional mapping technique within the HDiT framework, to predict the noise at each timestep t . We employ HDiT due to its efficiency in generating high-quality images in the pixel space because of the use of Neighborhood Attention (NA) [4].

The q-space condition is defined as the b-vector $\vec{b}_n = (b_x, b_y, b_z)$ and its magnitude b-value $|\vec{b}_n|$, thus the q-space sampling, unlike integer scalar index categories, e.g., real number or vector categories, requires tailored embedding to use as condition. Here, we develop a category embedding method for query-based conditional mapping, to further capture the directional information of the physics guidance \vec{D}_n . In detail, b-vectors are normalized and fed into a multi-layer perceptron (MLP) based positional embedding for stability, but overlooking the magnitude information specifying dMRI intensity, so the b-values are also introduced as guidance. Since traditional class embedding only applies to the input

of integer scalar indexes as categories, ignoring the signal strength information represented by real numbers, an embedding method that indicates magnitude categories based on real numbers is proposed. The real number embedding function $\varepsilon_{real}(\cdot)$ can be briefly represented as:

$$\varepsilon_{real}(x) = f_{MLP} \left[\frac{x}{|x|} \ln(|x| + 1) \right] \quad (7)$$

where $f_{MLP}(\cdot)$ represents the MLP function. This approach could benefit numerical stability.

Subsequently, guidance embeddings, including embedded slice index, are integrated into the Transformer with N feed-forward network (FFN) blocks, using GEGLU [23]. The process includes enhancing the image with an MLP, three NA HDiT blocks for downsampling, a self-attention Transformer, and three more NA HDiT blocks for upsampling, all interacting with the guidances [4].

2.3 XTRACT Condition Adaptation

It is essential to preserve tract information for dMRI generation, thus we introduce the XTRACT, which includes 42 anatomically dissected tracts derived from over 1,000 human subjects. To better incorporate these tracts into the DM conditioning, we employ an adapter technique inspired by [24], which could improve the adaptability of the model for controllable image generation while incorporating multiple conditions.

Particularly, in order to fully adapt the XTRACT atlas to DM, it is necessary to exclude meaningless contents. We note that not all the image slices have tracts on the XTRACT atlas. Therefore, we develop a module to enrich the empty positions. For those empty XTRACT, we combine S_{b_0} with the sum of the non-empty tracts, which can be simply expressed as:

$$c'_{X(i,j)} = \left(\sum_{k=1}^{42} c_{X(i,k)} + \xi \right) \cdot f_N(S_{b_0}) \quad (8)$$

where $c_{X(i,j)}$ means the XTRACT condition in the i -th slice and the j -th tract, $c'_{X(i,j)}$ is its replacement, ξ is a correction factor, and $f_N(\cdot)$ represents the min-max normalization.

Subsequently, as illustrated in Fig. 1, we use a pixel unshuffle with a downsampling factor of 4 to downsample the input XTRACT atlas as conditions. Then, three classic ResNet blocks with each followed downsampling blocks are adopted. All features from each downsampling procedure are permuted to be consistent with each corresponding HDiT block for further addition.

3 Experiments and Results

3.1 Datasets and Preprocessing

We collected dMRI data from 9 randomly selected subjects from the Human Connectome Project (HCP) S1200 [6]. These subjects were divided into training,

Table 1. Performance on the HCP S1200 dataset. **Boldface** marks the top models. The b-values between arbitrary b_n , $b_n = 1000 \text{ s/mm}^2$, $b_n = 2000 \text{ s/mm}^2$ and $b_n = 3000 \text{ s/mm}^2$ as conditions are tested.

Model	Arbitrary b_n		$b_n = 1000 \text{ s/mm}^2$		$b_n = 2000 \text{ s/mm}^2$		$b_n = 3000 \text{ s/mm}^2$	
	SSIM%	PSNR	SSIM%	PSNR	SSIM%	PSNR	SSIM%	PSNR
U-Net	61.92±11.3	23.90±3.08	66.07±10.5	25.48±2.19	65.64±11.2	25.00±2.16	65.91±10.6	25.20±2.84
U-Net++	61.89±9.34	23.21±2.83	65.74±10.8	25.20±2.43	66.31±10.3	24.49±2.19	65.18±12.3	24.21±3.03
q-DL	78.59±7.16	13.24±2.65	80.45±8.54	15.54±2.20	79.97±5.09	14.92±1.80	78.57±11.1	15.11±2.34
q-DL 2D	63.41±21.25	21.93±3.38	67.88±18.6	25.41±2.85	70.63±16.66	25.52±2.94	66.10±21.7	24.78±2.70
SMRI2DWI (CGAN)	32.12±8.89	20.71±3.15	33.32±7.48	24.53±2.38	34.66±6.10	24.02±2.56	34.67±7.80	24.06±2.78
Latent Diffusion	70.69±8.66	21.31±2.11	69.96±9.08	21.00±1.47	68.69±7.59	20.89±1.51	69.26±8.50	21.31±1.66
DDPM (HDiT)	81.92±5.41	22.70±2.26	86.78±3.94	25.72±1.95	82.71±4.92	23.36±2.02	81.19±6.51	22.26±3.01
Phy-Diff (Ours)	84.27±6.17	24.46±2.12	87.64±4.14	25.83±2.20	83.10±5.83	23.68±2.55	81.44±5.82	23.59±1.54

validation and testing sets in a ratio of 7:1:1. Each subject contains images across b-value shells at 1000, 2000 and 3000 s/mm^2 , with each shell comprising 90 diffusion directions along with b_0 images. Furthermore, 110 2D slices are extracted from each 3D images.

To ensure spatial consistency with XTRACT atlas, we first registered dMRI images to their corresponding T1-weighted images, followed by non-linearly transformed to the standard space by registering them to MNI-305 [7] using the Advanced Normalization Tools (ANTs). Additionally, we normalized all dMRI images, ADC atlas and XTRACT atlas, as well as b-vectors into the range of $[-1, 1]$, and padded all input images to ensure their resolution at 256×256 .

3.2 Implementation Details

We implement the following hyperparameters: linear noise schedule; 1000 diffusion steps; noise variances at the range of $\beta_1 = 10^{-4}$ and $\beta_T = 0.02$; patch size of 4×4 ; batch size of 32; AdamW optimizer with learning rate at 5.0×10^{-4} , weight decay at 10^{-4} , betas at 0.9 and 0.95, eps at 10^{-8} . The HDiT are implemented with three layers with dimensions of 128, 256, 512; 2 Transformer blocks in each layer. The Adapter convolution channels are the same as the HDiT dimensions, with a kernel size of 1. The model is trained in a NVIDIA A100-PCIE-40GB, and implemented through PyTorch. Maximum epochs of 80 is set and early stopping is used.

3.3 Comparisons with SOTA and Baselines

We compared our model with 5 state-of-the-art DWI synthesis approaches (Table 1): U-Net [22], U-Net++ [28], q-DL [9], q-DL (2D) [8], SMRI2DWI [20], LDM [21] and the original HDiT [4]. As conditioning is essential for our task, all models underwent slight modifications to incorporate basic conditioning, enabling them to utilize b-vectors and slice indexes as conditions. Specifically, they were fed into the vanilla MLP and further embedded into the model backbones.

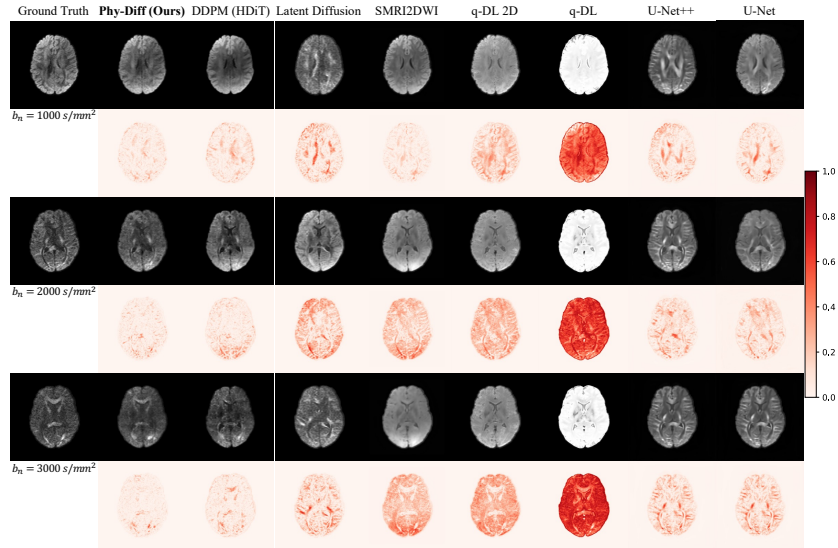


Fig. 2. Examples of the generated dMRI. In each b-value block, the first row contains the ground truth and synthesized images and the second row is the error maps.

Table 2. Ablation studies on arbitrary b_n . **Boldface** marks the top performance.

	SSIM%	PSNR
w/o Physics-guided noise evolution	84.24±5.86	24.12±2.53
w/o XTRACT Adaptation	82.63±9.94	24.34±2.54
w/o Query-based conditional mapping	75.31±15.4	22.84±2.46
Phy-Diff (Ours)	84.27±6.17	24.46±2.12

Quantitative Results. We generated dMRI based on three different b-values, and their combined set (arbitrary b_n). The results showed that for any b-value, Phy-Diff outperforms other models, with SSIM up to 52.15% and PSNR exceeding 11.22 dB, which is the same as in the situation when $b_n = 1000 s/mm^2$; In the cases of $b_n = 2000 s/mm^2$ and $b_n = 3000 s/mm^2$, SSIM also surpasses other models. However, PSNR is slightly inferior. Those high PSNR images often appear overly smooth with lower SSIM, losing important structural details. In clinical diagnosis, preserving these details is crucial.

Qualitative Results. Each various b_n block in Fig. 2 contains ground truth as well as synthesis results illustrated in the first row, and the error maps depicted in the second row. From the generated images, Phy-Diff shows the most comparable results to ground truth, with minimal errors shown in the heat maps, especially when b_n increases. Notably, Phy-Diff excels in reproducing white tracts, showcasing superior preservation of structural integrity.

3.4 Ablation Experiments

We evaluate the effectiveness of each component on the Phy-Diff by individually removing it. Our results (Table 2) show that each component contributes to the performance improvement, with query-based conditional mapping bringing a PSNR increase of 1.62 *dB* and SSIM of 8.96%.

4 Conclusion

In this paper, we present Phy-Diff, a DM-based dMRI generation model, with physics-guided noise evolution, query-based category condition mapping, and XTRACT condition adapter. Our experiments demonstrate that Phy-Diff achieves optimal performance under multiple b-values. Moreover, compared with other state-of-the-art models, Phy-Diff demonstrates more robust performance when accepting inputs with diverse b-vectors, b-values, and slices. Phy-Diff could serve as a tool for dMRI generation, offering the potential to reduce acquisition time while maintaining high-quality results.

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

References

1. Bau, D., Zhu, J.Y., Wulff, J., Peebles, W., Strobelt, H., Zhou, B., Torralba, A.: Seeing what a gan cannot generate (2019), <https://arxiv.org/abs/1910.11626>
2. Berard, H., Gidel, G., Almahairi, A., Vincent, P., Lacoste-Julien, S.: A closer look at the optimization landscapes of generative adversarial networks. arXiv preprint arXiv:1906.04848 (2019)
3. Chen, G., Jiang, H., Liu, J., Ma, J., Cui, H., Xia, Y., Yap, P.T.: Hybrid Graph Transformer for Tissue Microstructure Estimation with Undersampled Diffusion MRI Data. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2022. pp. 113–122. Lecture Notes in Computer Science, Springer Nature Switzerland, Cham (2022). https://doi.org/10.1007/978-3-031-16431-6_11
4. Crowson, K., Baumann, S.A., Birch, A., Abraham, T.M., Kaplan, D.Z., Shippole, E.: Scalable high-resolution pixel-space image synthesis with hourglass diffusion transformers (2024), <https://arxiv.org/abs/2401.11605>
5. Dhariwal, P., Nichol, A.: Diffusion Models Beat GANs on Image Synthesis (Jun 2021). <https://doi.org/10.48550/arXiv.2105.05233>, <http://arxiv.org/abs/2105.05233>, arXiv:2105.05233 [cs, stat]
6. Essen, D.C.V., Smith, S.M., Barch, D.M., Behrens, T.E.J., Yacoub, E., Ugurbil, K.: The WU-Minn Human Connectome Project: An overview. *NeuroImage* **80**, 62–79 (2013). <https://doi.org/https://doi.org/10.1016/j.neuroimage.2013.05.041>, <https://www.sciencedirect.com/science/article/pii/S1053811913005351>

7. Fonov, V., Evans, A.C., Botteron, K., Almli, C.R., McKinstry, R.C., Collins, D.L.: Unbiased average age-appropriate atlases for pediatric studies. *NeuroImage* **54**(1), 313–327 (2011). <https://doi.org/10.1016/j.neuroimage.2010.07.033>, <https://www.sciencedirect.com/science/article/pii/S1053811910010062>
8. Gibbons, E.K., Hodgson, K.K., Chaudhari, A.S., Richards, L.G., Majersik, J.J., Adluru, G., DiBella, E.V.: Simultaneous noddi and gfa parameter map generation from subsampled q-space imaging using deep learning. *Magnetic Resonance in Medicine* **81**(4), 2399–2411 (2019). <https://doi.org/https://doi.org/10.1002/mrm.27568>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.27568>
9. Golkov, V., Dosovitskiy, A., Sperl, J.I., Menzel, M.I., Czisch, M., Sämann, P., Brox, T., Cremers, D.: q-Space Deep Learning: Twelve-Fold Shorter and Model-Free Diffusion MRI Scans. *IEEE Transactions on Medical Imaging* **35**(5), 1344–1351 (May 2016). <https://doi.org/10.1109/TMI.2016.2551324>, <https://ieeexplore.ieee.org/document/7448418>, conference Name: IEEE Transactions on Medical Imaging
10. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *arXiv:2006.11239* (2020)
11. Jiang, L., Mao, Y., Chen, X., Wang, X., Li, C.: Cola-diff: Conditional latent diffusion model for multi-modal mri synthesis (2023), <https://arxiv.org/abs/2303.14081>
12. Karimi, D., Gholipour, A.: Atlas-powered deep learning (adl) – application to diffusion weighted mri (2022), <https://arxiv.org/abs/2205.03210>
13. Le Bihan, D., Johansen-Berg, H.: Diffusion MRI at 25: Exploring brain tissue structure and function. *NeuroImage* **61**(2), 324–341 (Jun 2012). <https://doi.org/10.1016/j.neuroimage.2011.11.006>, <https://www.sciencedirect.com/science/article/pii/S1053811911012857>
14. Li, C., Wei, Y., Chen, X., Schonlieb, C.B.: Brainnetgan: Data augmentation of brain connectivity using generative adversarial network for dementia classification (2021), <https://arxiv.org/abs/2103.08494>
15. Li, H., Yang, Y., Chang, M., Chen, S., Feng, H., Xu, Z., Li, Q., Chen, Y.: Srdiff: Single image super-resolution with diffusion probabilistic models. *Neurocomputing* **479**, 47–59 (2022)
16. Liu, P., Li, C., Schönlieb, C.B.: Ganredl: Medical image enhancement using a generative adversarial network with real-order derivative induced loss functions. In: Shen, D., Liu, T., Peters, T.M., Staib, L.H., Essert, C., Zhou, S., Yap, P.T., Khan, A. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*. pp. 110–117. Springer International Publishing, Cham (2019)
17. Mao, Y., Jiang, L., Chen, X., Li, C.: Disc-diff: Disentangled conditional diffusion model for multi-contrast mri super-resolution (2023), <https://arxiv.org/abs/2303.13933>
18. Mou, C., Wang, X., Xie, L., Wu, Y., Zhang, J., Qi, Z., Shan, Y., Qie, X.: T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models (2023), <https://arxiv.org/abs/2302.08453>
19. Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: *International Conference on Machine Learning*. pp. 8162–8171. PMLR (2021)
20. Ren, M., Kim, H., Dey, N., Gerig, G.: Q-space conditioned translation networks for directional synthesis of diffusion weighted images from multi-modal structural mri (2021), <https://arxiv.org/abs/2106.13188>
21. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models (2022), <https://arxiv.org/abs/2112.10752>

22. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation (2015), <https://arxiv.org/abs/1505.04597>
23. Shazeer, N.: Glu variants improve transformer (2020), <https://arxiv.org/abs/2002.05202>
24. Warrington, S., Bryant, K.L., Khrapitchev, A.A., Sallet, J., Charquero-Ballester, M., Douaud, G., Jbabdi, S., Mars, R.B., Sotiropoulos, S.N.: XTRACT - Standardised protocols for automated tractography in the human and macaque brain. *NeuroImage* **217**, 116923 (Aug 2020). <https://doi.org/10.1016/j.neuroimage.2020.116923>, <https://www.sciencedirect.com/science/article/pii/S1053811920304092>
25. Wei, Y., Chen, X., Zhu, L., Zhang, L., Schönlieb, C.B., Price, S., Li, C.: Multi-modal learning for predicting the genotype of glioma. *IEEE Transactions on Medical Imaging* **42**(11), 3167–3178 (2023)
26. Wei, Y., Li, C., Price, S.J.: Quantifying structural connectivity in brain tumor patients. In: de Bruijne, M., Cattin, P.C., Cotin, S., Padoy, N., Speidel, S., Zheng, Y., Essert, C. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*. pp. 519–529. Springer International Publishing, Cham (2021)
27. Zhang, H., Schneider, T., Wheeler-Kingshott, C.A., Alexander, D.C.: NODDI: Practical *in vivo* neurite orientation dispersion and density imaging of the human brain. *NeuroImage* **61**(4), 1000–1016 (Jul 2012). <https://doi.org/10.1016/j.neuroimage.2012.03.072>, <https://www.sciencedirect.com/science/article/pii/S1053811912003539>
28. Zhou, Z., Siddiquee, M.M.R., Tajbakhsh, N., Liang, J.: Unet++: A nested u-net architecture for medical image segmentation (2018), <https://arxiv.org/abs/1807.10165>