UFO-3: Unsupervised Three-Compartment Learning for Fiber Orientation Distribution Function Estimation

Xueqing Gao[1,2,◕], Rizhong Lin[3,◕], Jianhui Feng[1], Yonggang Shi[4,5], and Yuchuan Qiao[1(✉)]

[1] Institute of Science and Technology for Brain-Inspired Intelligence, Fudan University, Shanghai, China
yuchuanqiao@fudan.edu.cn
[2] School of Physics Science and Engineering, Tongji University, Shanghai, China
[3] School of Computer and Communication Sciences, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland
[4] Stevens Neuroimaging and Informatics Institute, Keck School of Medicine, University of Southern California (USC), Los Angeles, CA, USA
[5] Ming Hsieh Department of Electrical and Computer Engineering, Viterbi School of Engineering, University of Southern California (USC), Los Angeles, CA, USA

Abstract. Fiber orientation distribution function (fODF) estimation from diffusion MRI is crucial for mapping brain connectivity but often requires extensive multi-shell acquisitions and complex computational methods. While supervised deep learning approaches have shown promise in accelerating this process, they typically require large training datasets and face challenges with domain shifts and interpretability. We present UFO-3, an unsupervised framework that combines a three-compartment biophysical model with deep learning for fODF estimation from single-shell data. The method leverages a U-Net architecture to simultaneously estimate fiber orientations and tissue microstructure parameters while maintaining physical constraints through an optimization-based reconstruction. Evaluated on synthetic data across 2500 test cases, UFO-3 achieves superior angular accuracy (MAE $< 10^\circ$ at infinite SNR) and correlation (ACC $> 91\,\%$) compared to existing methods, particularly in resolving challenging fiber crossings. On *in vivo* human brain data, UFO-3 produces fODF reconstructions comparable to multi-shell reference methods while providing interpretable tissue parameter estimates. The framework requires a one-time, subject-specific training of about 30 min on a single consumer GPU and enables fast inference (< 10 s per subject), improving throughput compared to other unsupervised approaches that require hours or days of training. Our results demonstrate that UFO-3 effectively balances reconstruction accuracy, biological interpretability, and computational performance without requiring extensive training data or multi-shell acquisitions.

Keywords: Diffusion MRI · Fiber orientation distribution function · Unsupervised learning · Brain microstructure · Three-compartment model

◕ X. Gao and R. Lin contributed equally to this work.

1 Introduction

Diffusion-weighted magnetic resonance imaging (dMRI) has emerged as a powerful non-invasive imaging modality for mapping white matter microstructure and connectivity [18]. Fiber orientation distribution function (fODF) estimation, which characterizes the directional distribution of nerve fibers, is crucial for understanding complex white matter architecture. Unlike diffusion tensor imaging (DTI), fODF can resolve crossing fibers, enabling more accurate mapping of brain connectivity [29]. Classical fODF estimation methods include constrained spherical deconvolution (CSD) and its advanced variants [1, 2, 4, 5, 12, 28, 29] and compartment models [30]. CSD models the diffusion signal as a spherical convolution of a single-fiber response function and the fODF, whereas compartment models integrate intra-/extra-axonal and trapped water components to jointly estimate microstructure and fODFs specifically within white matter. Despite their robust performance in resolving complex fiber configurations, clinical adoption is limited by high computational demands and the usual need for extensive multi-shell acquisitions [12]—these prolonged scan times increase susceptibility to motion artifacts that can degrade image quality, particularly problematic for restless or uncooperative subjects.

Machine learning approaches have emerged to address these limitations, broadly categorized into supervised and unsupervised methods. Supervised approaches leveraging deep neural networks [14, 21, 27, 31, 32] have demonstrated the potential for fODF estimation from reduced diffusion gradient measurements with notable computational efficiency, including applications in challenging clinical scenarios [15, 16]. However, these data-driven methods rely on classical methods to generate training references, which in turn require densely sampled acquisitions. This interdependence presents methodological challenges, including domain shifts between training and test data [19, 20] and limited interpretability of learned representations [13], raising considerations for clinical translation.

Recent unsupervised methods have explored alternative approaches to eliminate ground truth requirements by incorporating physical constraints and symmetries directly into neural architectures. Particularly, equivariant networks that respect both spatial and rotational symmetries of diffusion signals [6–8] learn fODF estimation through signal reconstruction. Alternative approaches using neural fields [3] provide a more computationally efficient solution yet rely on parameterized models limited by assumptions, reducing flexibility and potentially introducing biases. Despite their improved accuracy in resolving crossing fibers and elimination of supervision requirements, these methods face practical challenges, including heavy computational overhead, lengthy training phases, and complex multi-step pipelines that limit clinical utility.

To address these limitations, we propose **UFO-3**, an **U**nsupervised framework for **F**iber **O**rientation distribution function estimation using a **3**-compartment model that integrates parametric tissue modeling with deep learning. Our method enables accurate fODF estimation from single-shell data without requiring precomputed ground truth for training. The framework maintains interpretability through biophysics-based constraints while offering a favorable balance between

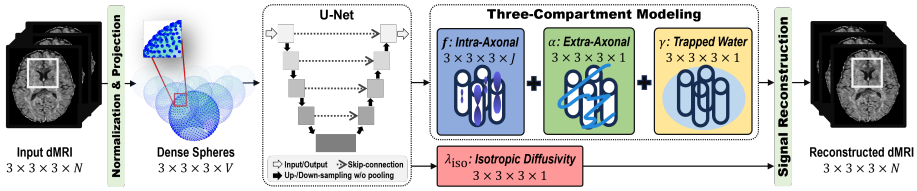

Fig. 1. UFO-3 framework overview. The input dMRI volume is first normalized and projected onto V dense nodes. A U-Net subsequently estimates four key components: intra-axonal fODF ($\boldsymbol{f}$), extra-axonal fraction (α), trapped water fraction (γ), and isotropic diffusivity (λ_{iso}). The final reconstruction integrates these estimates via physics-constrained optimization that enforces sparsity and non-negativity.

computational demands and reconstruction accuracy, making it suitable for both research and clinical applications.

2 Methods

Our UFO-3 framework combines biophysical modeling and unsupervised learning for fiber orientation estimation, balancing data-driven learning with physical principles, as illustrated in Fig. 1.

2.1 Biophysical three-compartment model

The input dMRI signal measurements are first normalized by $b = 0$ images and projected onto a dense spherical graph with V uniformly distributed nodes using HEALPix [11], transforming each voxel's representation from N to V dimensions. This projection operation, denoted as $\mathcal{P}(\boldsymbol{s})$, maps diffusion-weighted signals $\boldsymbol{s}$ in q-space to points on a uniform HEALPix sphere, providing a consistent representation for the neural network input.

Following Tran and Shi [30], we model the diffusion signal $\boldsymbol{s} \in \mathbb{R}^N$ as:

$$\boldsymbol{s} = \mathbf{A}_+ \left[\boldsymbol{f}\ \alpha\ \gamma\right]^\top + n, \tag{1}$$

where $\boldsymbol{f} \in \mathbb{R}^J$ represents intra-axonal fODF coefficients, α extra-axonal fraction, γ trapped water fraction, and n noise.

The augmented system matrix $\mathbf{A}_+ = [\mathbf{A}\ \boldsymbol{\beta}\ \boldsymbol{e}]$ combines the diffusion matrix $\mathbf{A} \in \mathbb{R}^{N\times J}$, isotropic terms $\boldsymbol{\beta} = [\mathrm{e}^{-b_1\lambda_{\text{iso}}} \cdots \mathrm{e}^{-b_N\lambda_{\text{iso}}}]^\top$, and unity vector $\boldsymbol{e} = \mathbf{1}_N$. The core matrix $\mathbf{A} = \mathbf{Y} \odot \mathbf{G}$ integrates spherical harmonics $\mathbf{Y} = [Y_j(u_i)]_{i=1,\dots,N}^{j=1,\dots,J} \in \mathbb{R}^{N\times J}$ and diffusion properties $\mathbf{G} = [G_l(b_i)]_{i=1,\dots,N}^{l=0,\dots,L} \in \mathbb{R}^{N\times J}$ matrices, where b is the b-value and u the gradient direction. Y_j represents the real spherical harmonics (SH) basis functions of even order $l = 0, 2, \dots, L$ with index $j = m + (^{(l^2+l+2)}/_2)$. Each $G_l(b_i) = 2\pi \int_{-1}^{1} P_l(t)\mathrm{e}^{-b_i\lambda_\parallel t^2}\,\mathrm{d}t$ is repeated $2l+1$ times in row i of $\mathbf{G}$ for $l = 0, \dots, L$, where $P_l(t)$ is the Legendre polynomial and $\lambda_\parallel = 0.0017\,\mathrm{mm}^2/\mathrm{s}$ the axonal diffusivity [33].

2.2 Unsupervised learning framework

We implement a U-Net [26] with pointwise convolutions ($1 \times 1 \times 1$) that operates on each voxel independently, enabling parallel processing. The network processes $3 \times 3 \times 3$ voxel patches solely for computational batching efficiency; as the convolutions are strictly pointwise, this patch-based approach does not introduce any spatial dependencies between voxels. The network $\mathcal{N}_\theta$ takes the projected signal $\mathcal{P}(\boldsymbol{s})$ as input and estimates all tissue parameters:

$$\left[\boldsymbol{f}\ \alpha\ \gamma\ \lambda_{\text{iso}}\right] = \mathcal{N}_\theta(\mathcal{P}(\boldsymbol{s}))^\top, \tag{2}$$

with *Softplus* activation ensuring non-negativity of α, γ, and λ_{iso}.

Physical constraints are enforced via compartment normalization $\boldsymbol{I}\boldsymbol{f}+\alpha+\gamma = 1$ with $\boldsymbol{I} = [\sqrt{4\pi}, 0, \cdots, 0] \in \mathbb{R}^J$ and non-negativity $\mathbf{C}_M\boldsymbol{f} \geq \mathbf{0}$. For sparsity concerns, we employ hierarchical remeshing to construct projection matrix $\mathbf{C}_M \in \mathbb{R}^{M\times J}$ from hemisphere point sets $\mathcal{V} = \{V_1, V_2, \dots\}$, where each set contains M uniformly distributed points on the unit hemisphere [24]. Here, M controls the sampling density for non-negativity enforcement.

The complete optimization objective combines signal reconstruction with physical constraints:

$$\min \quad \frac{1}{2}\left\|\boldsymbol{s} - \mathbf{A}_+\left(\lambda_{\text{iso}}\right)\left[\boldsymbol{f}\ \alpha\ \gamma\right]^\top\right\|_2^2 + \xi\,\text{Reg}\left(\mathbf{C}_M\boldsymbol{f}\right), \tag{3}$$

where $\mathbf{A}_+(\lambda_{\text{iso}})$ depends on the network-estimated extra-axonal diffusivity λ_{iso}.

3 Experimental results

Here, we evaluate UFO-3 against conventional biophysics-based (CSD [28]) and unsupervised learning methods (NODF [3], RT-ESD [6], and SHD-TV [7]). Implemented in PyTorch [23], UFO-3 was trained on a single NVIDIA RTX 3090 GPU optimized by Adam [17] (batch size: 128, patch size: 3^3, learning rate: 1.5×10^{-3}). With $V = 3072$ and $M = 258$, training converged within 30 min per subject for all experimental setups, with inference taking under 10 s per subject. The code is publicly available at `https://github.com/tensor2023/ufo-3`.

3.1 Synthetic fODF reconstruction analysis

Data. To simulate realistic conditions, we generate ground-truth (GT) fODFs for five fiber configurations: (1) a single fiber; (2–4) two crossing fibers with angles between them of 90°, 60°, and 45°, respectively; and (5) three crossing fibers with an angle of 60° between each pair of adjacent fibers. For each configuration, we generated 100 unique instances with randomly sampled orientations per SNR level, resulting in a total of 2500 test cases (5 configurations × 5 SNRs × 100 orientations). These signals are synthesized w.r.t. the GT fODFs using Dipy [9], following the biophysical model in Section 2.1. The synthesis incorporates diffusion gradient directions from the $b = 1000\,\text{s/mm}^2$ shell of five subjects from

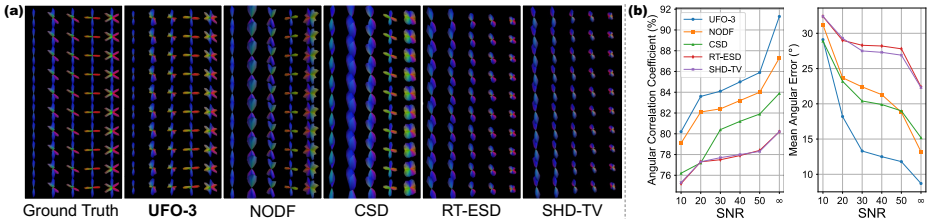

Fig. 2. Synthetic data evaluation results. (a) Qualitative comparison of fiber reconstruction performance. Each method's reconstruction is shown for five distinct fiber configurations: single fiber, two fibers crossing at 90°, 60°, and 45°, and three fibers crossing at 60°. (b) Quantitative performance analysis showing Angular Correlation Coefficient (left) and Mean Angular Error (right) across SNR levels from 10 to infinite (∞). An infinite SNR means no noise is added to the signal.

the Chinese Human Connectome Project (CHCP) [10]. The resulting data are validated against real-world brain data as described by Tran and Shi [30] and further augmented with Rician noise at signal-to-noise ratios (SNRs) of 10, 20, 30, 40, and 50 dB.

Evaluation setup. We adopt two metrics from Zeng et al. [32]: (1) mean angular error (MAE) to quantify the deviation between estimated and true fODFs, and (2) angular correlation coefficient (ACC) to assess fODF similarity.

Results. Visual comparison (Fig. 2a) demonstrates UFO-3's superior reconstruction performance across tested configurations, particularly in preserving sharp features and angular precision of crossing fibers in challenging 60° and 45° cases, while other methods show varying degrees of fiber orientation blurring or artifacts.

Quantitative analysis (Fig. 2b) shows consistently lower mean angular error across SNR levels, with UFO-3 achieving approximately 18° MAE at SNR = 10 dB compared to 28° to 32° for other methods. This advantage persists through higher SNR levels, reaching below 10° at infinite SNR.

The angular correlation coefficient results reinforce these findings, with UFO-3 reaching above 91 % ACC at infinite SNR, compared to 76 % to 87 % for competing approaches. Notably, even under severe noise (SNR = 10 dB), UFO-3 maintains 80 % ACC, with potential implications for clinical scenarios where high SNR cannot be guaranteed.

3.2 *In vivo* human brain fODF estimation analysis

Data. We analyze five CHCP subjects [10] acquired on a 3T Siemens Prisma MRI scanner. The data were collected using a multi-shell protocol with b-values of 0, 1000, and 2000 s/mm^2, with an isotropic spatial resolution of 1.5 mm, covering a field of view of $121 \times 145 \times 121$ voxels. Standard preprocessing was applied,

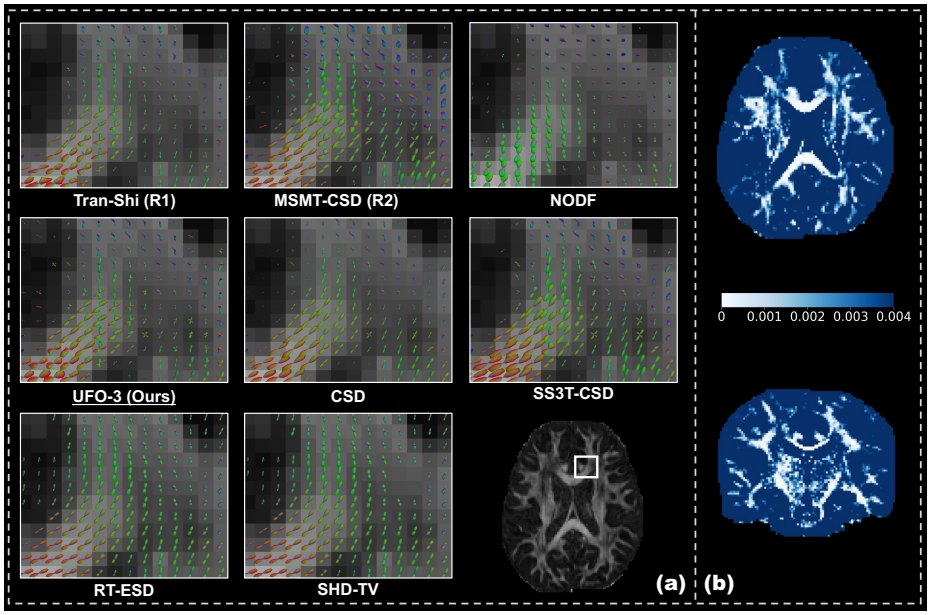

Fig. 3. Qualitative evaluation on *in vivo* human brain data. (a) Comparison of fODF reconstructions overlaid on fractional anisotropy (FA) maps in a region containing crossing fibers between the corpus callosum and superior longitudinal fasciculus (location indicated in the reference image). Multi-shell references (R1: Tran-Shi, R2: MSMT-CSD) are compared against single-shell methods. (b) Estimated diffusivity of the extra-axonal compartment (λ_{iso}) shown in coronal (top) and axial (bottom) views.

including denoising, motion, and distortion correction. Our network uses 46 gradient directions at $b = 1000\,\text{s/mm}^2$ as input for all experiments.

For qualitative analysis of the *in vivo* data where GT fODFs are unavailable, we compute two references on the full multi-shell sequence of dMRI using the compartment model [30] and multi-shell multi-tissue CSD (MSMT-CSD) [12], denoted as R1 and R2, respectively. We additionally include single-shell three-tissue CSD (SS3T-CSD) [4] for comparison.

Results. Fig. 3a shows fODF reconstructions in a representative region containing crossing fibers between the corpus callosum and superior longitudinal fasciculus. While the figure displays one subject for clarity, the observed patterns and performance are consistent across all five analyzed subjects. The multi-shell references (R1, R2) clearly delineate two distinct fiber populations with sharp angular profiles and minimal spurious orientations. UFO-3 produces reconstructions that closely match these reference patterns, particularly in preserving the sharp angular separation between crossing bundles and maintaining consistent orientation estimates across the region.

Among single-shell methods, notable differences emerge in their ability to resolve this complex architecture. NODF shows orientation blurring at crossing regions, particularly evident in the transition zone where the two fiber populations meet. Traditional CSD, while capturing the primary orientations, introduces spurious peaks and shows increased angular uncertainty. SS3T-CSD demonstrates improved crossing fiber resolution compared to CSD but exhibits some orientation dispersion in regions where the reference methods show more coherent fiber populations. RT-ESD and SHD-TV achieve better suppression of spurious peaks than CSD, though still show some degree of orientation uncertainty in complex crossing areas.

An advantage of UFO-3 is its ability to estimate tissue microstructure properties alongside fiber orientations. Fig. 3b visualizes the extra-axonal compartment diffusivity (λ_{iso}) in coronal and axial views. The diffusivity maps reveal anatomically plausible tissue properties: lower values in coherent white matter regions (corpus callosum, corticospinal tract) and higher values in regions with complex fiber crossings (centrum semiovale, thalamus). The maps show expected contrasts between tissue types, with the highest diffusivity in CSF regions and lower diffusivity in gray matter due to cellular density. White matter exhibits intermediate diffusivity values that vary with fiber configuration complexity.

These results demonstrate UFO-3's capability to simultaneously recover accurate fiber orientations and meaningful tissue microstructure parameters from single-shell data, approaching the quality of multi-shell reference methods.

3.3 Tractography reconstruction analysis

To evaluate the practical utility of fODF estimation in structural connectivity mapping, we assess the tractography reconstruction performance of UFO-3.

Data. The Tractometer dataset [22] provides a single-volume simulation of a real human brain with a low-angular resolution acquisition protocol (1 shell, 32 gradient directions), including ground truth tractography bundles.

Results. We perform a qualitative evaluation of tractography results using UFO-3-derived fODFs compared to traditional single-shell approaches. For visualization, we focus on the brainstem projection system (BPS). Fig. 4a shows tractography reconstructions from different methods compared to the ground truth configuration. SS3T-CSD and CSD appear less sensitive to fiber curvature, producing fewer streamlines in the right region and failing to capture directional changes. UFO-3 provides more voluminous streamline reconstruction in most regions, which may result from (1) capturing high-curvature regions that deterministic GT tracking could miss, and (2) mild over-smoothing due to the convolutional architecture. It is important to note that even ground truth tractography can be imperfect, as discussed in [22]. For quantitative comparison, we compute bundle overlap rates following the standard Tractometer scoring protocol. As shown in Fig. 4b, UFO-3 achieves comparable overlap rates to other

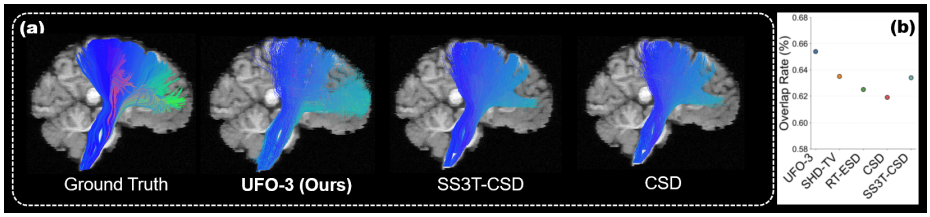

Fig. 4. Tractography evaluation on the Tractometer dataset. (a) Visualization of the brainstem projection system (BPS) reconstructed using different methods compared to the ground truth. (b) Bundle overlap rates for the methods. Values for UFO-3, CSD, and SS3T-CSD are from our implementation, while SHD-TV and RT-ESD values are from [7].

single-shell methods, with SHD-TV and RT-ESD values cited from their original reported results in [7].

4 Discussion

UFO-3 demonstrates the potential of combining biophysical modeling with unsupervised learning for accurate fODF estimation from single-shell dMRI. Our experiments show that UFO-3 achieves comparable or superior performance to existing methods while maintaining interpretability through its physics-based design. The framework's ability to simultaneously estimate microstructure parameters provides additional biological insights without requiring multi-shell acquisitions.

While our U-Net architecture might appear conventional, it represents a deliberate choice balancing accuracy and efficiency. Since its introduction, U-Net has proven remarkably effective across diverse medical imaging applications, with its skip connections particularly valuable for preserving fine structural details in fODF estimation. The 3^3 patch size similarly balances local context capture with computational efficiency, enabling multi-batch inference while maintaining spatial coherence.

A key advantage of UFO-3 is its subject-specific training approach, which eliminates the need for large-scale training datasets and mitigates potential domain shift issues common in supervised learning methods. However, this design choice necessitates that model parameters be optimized for each subject independently. The computational time required for this step (approximately 30 min per subject on a single NVIDIA RTX 3090 GPU) is a trade-off for achieving a personalized model without the need for large, pre-existing training datasets.

4.1 Limitations and future directions

Our approach has several limitations. First, while our $V = 3072$ spherical grid effectively balances resolution and computational demands, lower resolutions

caused unstable optimization due to poor spherical operator conditioning. Future work could explore more efficient sampling strategies, maintaining stability with fewer nodes. Second, UFO-3 currently handles single-shell acquisitions at $b = 1000\,\mathrm{s/mm^2}$ to match clinical settings. Extension to multi-shell or higher b-value protocols would leverage richer diffusion characteristics.

While our validation is comprehensive within scope, testing across more diverse datasets would strengthen results. Specifically, validation on the DiSCo dataset [25], with both microscopic (axon diameter, myelin) and macroscopic (fiber trajectory) ground truth, would enable more thorough evaluation of orientation accuracy and tissue parameter estimates. This would address concerns regarding volumetric differences in tractography results compared to established methods. Integration with advanced tractography algorithms could better demonstrate the method's utility for connectivity analysis.

4.2 Conclusion

We have presented UFO-3, an unsupervised framework that effectively combines biophysical modeling with deep learning for fODF estimation. The method achieves robust performance on single-shell data while maintaining interpretability and computational efficiency. While there remain opportunities for improvement, particularly in multi-tissue modeling and parameter stability, UFO-3 represents a promising step toward more practical and reliable fiber orientation estimation for clinical applications.

Acknowledgments. Xueqing Gao, Jianhui Feng, and Yuchuan Qiao were supported by the National Natural Science Foundation of China (Grant No. 82102002).

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

References

1. Canales-Rodríguez, E.J. et al.: Spherical deconvolution of multichannel diffusion MRI data with non-gaussian noise models and spatial regularization. PLOS ONE **10**(10), e0138910 (2015)
2. Canales-Rodríguez, E.J. et al.: Sparse wars: A survey and comparative study of spherical deconvolution algorithms for diffusion MRI. NeuroImage **184**, 140–160 (2019)
3. Consagra, W. et al.: Neural orientation distribution fields for estimation and uncertainty quantification in diffusion MRI. Medical Image Analysis **93**, 103105 (2024)
4. Dhollander, T., Connelly, A.: A novel iterative approach to reap the benefits of multi-tissue CSD from just single-shell (+b=0) diffusion MRI data. In: ISMRM, p. 3010 (2016)

5. Dhollander, T. et al.: Unsupervised 3-tissue response function estimation from single-shell or multi-shell diffusion MR data without a co-registered T1 image. In: ISMRM Workshop on Breaking the Barriers of Diffusion MRI, p. 5 (2016)
6. Elaldi, A. et al.: $E(3) \times SO(3)$-Equivariant networks for spherical deconvolution in diffusion MRI. In: MIDL, pp. 301–319 (2023)
7. Elaldi, A. et al.: Equivariant spatio-hemispherical networks for diffusion MRI deconvolution. In: NeurIPS, pp. 52095–52126 (2024)
8. Elaldi, A. et al.: Equivariant spherical deconvolution: Learning sparse orientation distribution functions from spherical data. In: IPMI, pp. 267–278 (2021)
9. Garyfallidis, E. et al.: Dipy, a library for the analysis of diffusion MRI data. Frontiers in Neuroinformatics **8** (2014)
10. Ge, J. et al.: Increasing diversity in connectomics with the Chinese Human Connectome Project. Nature Neuroscience **26**(1), 163–172 (2023)
11. Górski, K.M. et al.: HEALPix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. The Astrophysical Journal **622**(2), 759 (2005)
12. Jeurissen, B. et al.: Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion MRI data. NeuroImage **103**, 411–426 (2014)
13. Karimi, D., Warfield, S.K.: Diffusion MRI with machine learning. Imaging Neuroscience **2**, 1–55 (2024)
14. Karimi, D. et al.: Learning to estimate the fiber orientation distribution function from diffusion-weighted MRI. NeuroImage **239**, 118316 (2021)
15. Kebiri, H. et al.: Deep learning microstructure estimation of developing brains from diffusion MRI: A newborn and fetal study. Medical Image Analysis **95**, 103186 (2024)
16. Kebiri, H. et al.: Robust estimation of the microstructure of the early developing brain using deep learning. In: MICCAI, pp. 293–303 (2023)
17. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015)
18. Le Bihan, D. et al.: Diffusion tensor imaging: Concepts and applications. Journal of Magnetic Resonance Imaging **13**(4), 534–546 (2001)
19. Lin, R. et al.: Cross-age and cross-site domain shift impacts on deep learning-based white matter fiber estimation in newborn and baby brains. In: ISBI (2024)
20. Lin, R. et al.: Ground-truth effects in learning-based fiber orientation distribution estimation in neonatal brains. In: CDMRI, pp. 24–34 (2024)
21. Lin, Z. et al.: Fast learning of fiber orientation distribution function for MR tractography using convolutional neural network. Medical Physics **46**(7), 3101–3116 (2019)
22. Maier-Hein, K.H. et al.: The challenge of mapping the human connectome based on diffusion tractography. Nature Communications **8**(1), 1349 (2017)

23. Paszke, A. et al.: PyTorch: An imperative style, high-performance deep learning library. In: NeurIPS (2019)
24. Peyré, G., Cohen, L.D.: Geodesic remeshing using front propagation. International Journal of Computer Vision **69**(1), 145–156 (2006)
25. Rafael-Patino, J. et al.: The diffusion-simulated connectivity (DiSCo) dataset. Data in Brief **38**, 107429 (2021)
26. Ronneberger, O. et al.: U-Net: Convolutional networks for biomedical image segmentation. In: MICCAI, pp. 234–241 (2015)
27. da Silva, M.O. et al.: FOD-Swin-Net: angular super resolution of fiber orientation distribution using a transformer-based deep model. In: ISBI (2024)
28. Tournier, J.-D. et al.: Robust determination of the fibre orientation distribution in diffusion MRI: Non-negativity constrained super-resolved spherical deconvolution. NeuroImage **35**(4), 1459–1472 (2007)
29. Tournier, J.-D. et al.: Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. NeuroImage **23**(3), 1176–1185 (2004)
30. Tran, G., Shi, Y.: Fiber orientation and compartment parameter estimation from multi-shell diffusion imaging. IEEE Transactions on Medical Imaging **34**(11), 2320–2332 (2015)
31. Yao, T. et al.: Robust fiber orientation distribution function estimation using deep constrained spherical deconvolution for diffusion-weighted magnetic resonance imaging. Journal of Medical Imaging **11**(1), 014005 (2024)
32. Zeng, R. et al.: FOD-Net: A deep learning method for fiber orientation distribution angular super resolution. Medical Image Analysis **79**, 102431 (2022)
33. Zhang, H. et al.: NODDI: Practical *in vivo* neurite orientation dispersion and density imaging of the human brain. NeuroImage **61**(4), 1000–1016 (2012)