NEURAL FIBER OPERATORS: LEARNING ON ANTIPODALLY SYMMETRIC SPHERICAL FIBER FUNCTIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

Diffusion MRI (dMRI) is a critical non-invasive technique for characterizing white matter integrity. Predicting high angular resolution fiber orientation distribution (FOD) is crucial in clinical practice. Existing learning-based methods enable such FOD angular super resolution task but ignore the spherical-native geometry of FOD. To address this, we first reformulate FOD angular super resolution as a spherical operator learning problem, rather than conventional FOD regression. Based on this, we propose Neural Fiber Operator (NFO), a novel antipodal-symmetry-based neural operator, which efficiently learns the solution operator from FOD while preserving the antipodal symmetry of FOD. To exploit NFO, we design NFONet, a patch-based neural operator model for whole-brain or partial-brain FOD modeling. Extensive experiments are conducted on three complementary diffusion MRI datasets covering in-domain reconstruction, zero-shot generalization to independent cohorts and acquisition protocols, and ex-vivo biological validation. NFO consistently achieves state-of-the-art performance while using substantially fewer parameters. More importantly, NFO demonstrates robust generalization across progressively increasing distribution shifts, suggesting that the learned operator captures intrinsic properties of FOD angular super resolution. These results establish NFO as a principled framework for dMRI reconstruction and FOD modeling, and suggest that learning physically consistent solution operators provides a natural formulation for brain structure.

1 INTRODUCTION

Diffusion magnetic resonance imaging (dMRI) is an important non-invasive neuroimaging technique characterizing brain microstructure by measuring the direction-dependent diffusion of water molecules (Shamir & Assaf, 2025). Fiber bundle reconstruction based on dMRI is key to measuring fiber integrity for understanding brain physiology and pathology (Shamir & Assaf, 2025; Dell'Acqua & Tournier, 2019; Jeurissen et al., 2019). Reliable estimation of local fiber orientations is therefore a critical prerequisite for downstream tractography and connectome reconstruction (Jeurissen et al., 2019). Traditional methods, such as diffusion tensor imaging, estimate overall fiber directions efficiently (Basser et al., 1994) but fails to model complex structures like crossings and bendings due to the single-fiber-per-voxel assumption (Tuch, 2004).

Fiber orientation distribution (FOD) modeling has emerged as one of the most effective representations for resolving complex intra-voxel fiber configurations (Dell'Acqua & Tournier, 2019). FOD represents the angular distribution of fiber populations as a spherical function defined on $\mathbb{S}^2$. Advanced methods based on high angular resolution diffusion imaging (HARDI), offer better characterization of complex structures (Dell'Acqua & Tournier, 2019). Based on spherical harmonics (SH), constrained spherical deconvolution methods are developed to estimate FOD (Tournier et al., 2004; 2007; Khan et al., 2020; Jeurissen et al., 2014). Especially, single-shell multi-tissue CSD (SSMT-CSD) uses single-shell dMRI to estimate FOD (Khan et al., 2020). In contrast, multi-shell multi-tissue CSD (MSMT-CSD) (Jeurissen et al., 2014) utilizes multi-shell HARDI data for more accurate FOD estimation by incorporating more b-values and b-vectors. However, it brings the drawback of longer scan time due to the acquisition of more b-values and gradient directions, limiting its utility in clinical practice (Zeng et al., 2022).

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

In clinical practice, single-shell dMRI with lower b-values and fewer gradient directions, known as low angular resolution diffusion imaging (LARDI), is often used due to scanner limitations and acquisition-time constraints (Zeng et al., 2022), which yet is limited in accurate FOD estimation. As a result, there is an unmet need to improve FOD estimation from clinically feasible LARDI data for wider clinical impact.

Deep learning approaches promise to meet this demand. Existing methods, such as FOD-Net (Zeng et al., 2022), FOD-Swin-Net (Da Silva et al., 2024) and SDNet (Bartlett et al., 2024), formulate FOD super-resolution as a patch-based regression problem, where the local 3D LAR-FOD patch is used to predict the SH coefficients of HAR-FOD. This strategy benefits from local anatomical continuity and avoids directly modeling raw dMRI signals whose dimensionality varies across acquisition protocols (Zeng et al., 2022). However, such methods mostly treat FODs as finite-dimensional SH coefficient vectors or multi-channel volumes, resulting in the neglect of spherical geometry. More fundamentally, the FOD estimation problem is not merely a regression task. According to the spherical deconvolution formulation (Tournier et al., 2007; Khan et al., 2020; Jeurissen et al., 2014), FOD is a well-defined spherical function whose recovery is governed by a spherical convolutional inverse problem. Therefore, the FOD modeling should be considered as a function-to-function solution on $\mathbb{S}^2$. Neural operators provide a natural paradigm for this purpose (Kovachki et al., 2024; Li et al., 2020; Bonev et al., 2023), yet existing approaches overlook the antipodal structure of FODs and are computationally inefficient for large-scale whole-brain fiber modeling.

To address these challenges and solution the target FOD, we propose *Neural Fiber Operators* (NFO), a antipodal-symmetric neural operator framework for FOD angular super-resolution. Instead of treating FOD modeling as regression task, we reformulate the task as learning a solution operator from FOD functions on the sphere. Inspired by CSD and Green's function theory, we define the LAR-FOD as the input function f of a spherical system and derive the HAR-FOD g with strict antipodal symmetry as the operator solution. This formulation preserves the intrinsic geometry of FODs, enables global angular modeling in the spherical harmonic domain, and provides a principled bridge between classical spherical deconvolution and modern neural operator learning. To fully exploit NFO for whole-brain or partial-brain FOD modeling, we further design a full ptach-based operator network architecture, termed *NFONet*. The main contributions are summarized as follows:

1. We propose a spherical fiber operator system, which reformulates FOD angular super resolution as a spherical operator learning problem beyond conventional FOD regression.
2. We propose Neural Fiber Operator (NFO), a novel antipodal-symmetry-informed neural operator, which efficiently learns the operator solution from the fiber operator system while preserving the spherical geometry.
3. We design NFONet, a patch-based neural operator network architecture based on NFO for efficient whole-brain or partial-brain FOD modeling.
4. Extensive experiments demonstrate that the learned operator generalizes across different tissue compositions, independent cohorts, acquisition protocols and even microscopy-correlated ex-vivo specimens, indicating its transferable spherical fiber modeling ability.

2 Related Work

Learning-based FOD modeling. Learning-based FOD reconstruction has been broadly explored (Nath et al., 2020; Jha et al., 2020; Karimi et al., 2021). Some methods directly estimate FODs from diffusion MRI signals using spherical neural network to account for the directional structure of diffusion measurements (Sedlar et al., 2021; Yao et al., 2024; Jha et al., 2022; Snoussi & Karimi, 2025). However, such direct signal-to-FOD learning can be sensitive to variations in acquisition protocols, b-values, and scanner characteristics. Recent methods instead learn the mapping from low-angular-resolution FODs (LAR-FODs) to high-angular-resolution FODs (HAR-FODs), which exploit local spatial information through patch-based learning and directly predict the target FOD (Zeng et al., 2022; Bartlett et al., 2024; Da Silva et al., 2024).

Neural operators. Neural operators learn mappings between function spaces and have emerged as a general framework for operator learning (Kovachki et al., 2024; Lu et al., 2019; 2021; Bhattacharya et al., 2021; Li et al., 2020). Recent spherical neural operators extend this paradigm to functions defined on $\mathbb{S}^2$ through spherical harmonic representations and have demonstrated effectiveness in

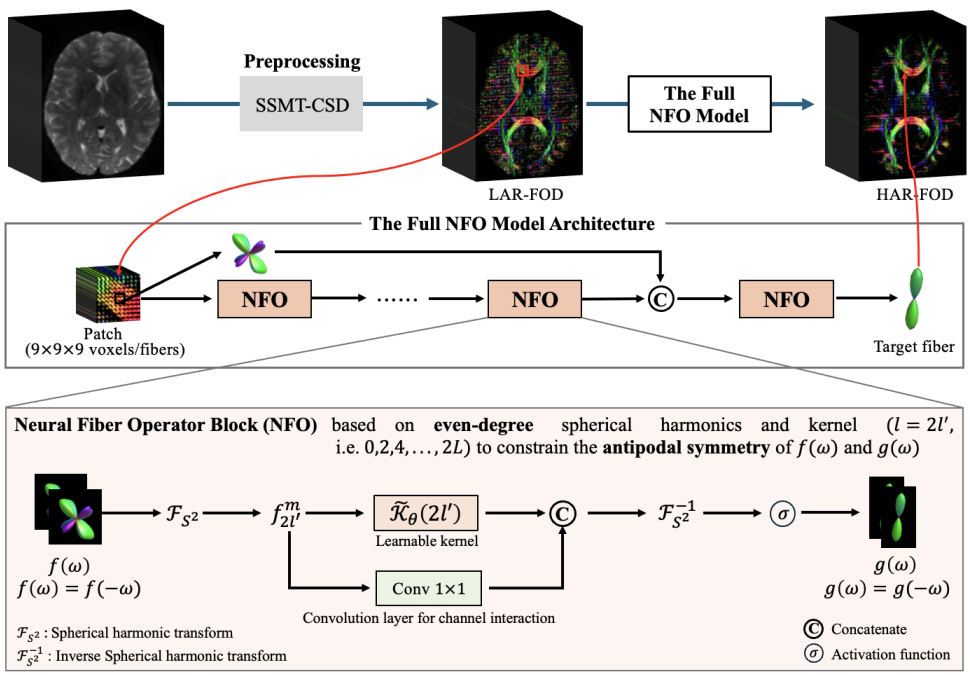


Figure 1: The full architecture of neural fiber operator model (NFO). Full patch-based framework (Top): Each target HAR-FOD voxel is predicted by the full NFO model architecture, which consists of multiple NFO blocks. NFO enables operator learning on antipodally symmetric spherical fiber functions (Bottom).

spherical PDEs and geoscience applications (Bonev et al., 2023; Lin et al., 2023; Tang et al., 2026; Mahesh et al., 2024a;b; Hu et al., 2025).

3 Preliminaries

3.1 Spherical Harmonics

For a square-integrable function $f \in L^2(\mathbb{S}^2)$, the spherical harmonic transform ($\mathcal{F}_{\mathbb{S}^2}$) projects f onto the spherical harmonic basis (Müller, 2006):

$$\hat{f}_l^m = \mathcal{F}_{\mathbb{S}^2}[f](l,m) = \int_{\mathbb{S}^2} f(\omega)\overline{Y_l^m(\omega)}d\omega. \tag{1}$$

The inverse spherical harmonic transform reconstructs the function as:

$$f(\omega) = \mathcal{F}_{\mathbb{S}^2}^{-1}[\hat{f}](\omega) = \sum_{l=0}^{L}\sum_{m=-l}^{l} \hat{f}_l^m Y_l^m(\omega). \tag{2}$$

The spherical harmonics satisfy the orthogonality relation: $\int_{\mathbb{S}^2} Y_l^m(\omega)\overline{Y_{l'}^{m'}(\omega)}d\omega = \delta_{ll'}\delta_{mm'}$. Under a rotation $R \in SO(3)$, the spherical harmonics transform according to:

$$Y_l^m(R\omega) = \sum_{k=-l}^{l} D_{k,m}^{(l)}(R^{-1})Y_l^k(\omega), \tag{3}$$

where $D^{(l)}$ denotes the Wigner-D matrix (Driscoll & Healy, 1994). The Wigner-D matrix is related to spherical harmonics through:

$$\overline{D_{0,m}^{(l)}(R^{-1})} = \sqrt{\frac{4\pi}{2l+1}}\overline{Y_l^m(Rn)}, \tag{4}$$

where n is the north pole.

3.2 Fiber Orientation Distribution and Spherical Deconvolution

In the CSD formulation (Tournier et al., 2007; Khan et al., 2020; Jeurissen et al., 2014), the diffusion-weighted signal attenuation $S(\omega)$ measured over a set of gradient directions is modeled as the spherical convolution between a response function R and the fiber orientation distribution function F: $S(\omega) = (R * F)(\omega)$. Within this formulation, FOD is expressed using even-degree spherical harmonics:

$$F(\omega) = \sum_{l=0}^{L} \sum_{m=-2l}^{2l} \hat{F}_{2l}^{m} Y_{2l}^{m}(\omega), \tag{5}$$

where Y_{2l}^{m} is the spherical harmonic basis of even degree $2l$ and order m, and $\hat{F}_{2l}^{m}$ is the corresponding SH coefficient. Because FODs are antipodally symmetric, only even SH degrees are used.

3.3 Spherical Neural Operators

Neural operators aim to learn mappings between function spaces. Given an input function f and a target solution function g, a neural operator learns: $\mathcal{G}_\theta : f(\omega) \mapsto g(\omega)$, where both f and g may be discretized on different grids. On the sphere, existing spherical neural operator (Bonev et al., 2023; Tang et al., 2026) can be written as:

$$g(\omega) = \mathcal{F}_{\mathbb{S}^2}^{-1} \left[W_\theta \cdot \mathcal{F}_{\mathbb{S}^2}[f](l, m) \right](\omega), \tag{6}$$

where W_θ is a learnable spectral kernel. This formulation performs global angular mixing in the spherical harmonic domain while preserving the geometry of $\mathbb{S}^2$.

4 Methodology

Problem Formulation. Let $\Omega \subset \mathbb{R}^3$ denote the brain volume and $x \in \Omega$ denote a voxel location. At each voxel, the low-angular-resolution FOD is a spherical function: $f_x : \mathbb{S}^2 \to \mathbb{R}$, and the target high-angular-resolution FOD is: $g_x : \mathbb{S}^2 \to \mathbb{R}$. The goal of FOD angular super-resolution is to learn an operator on the sphere:

$$\mathcal{G}_\theta : f_x(\omega) \mapsto g_x(\omega), \quad \omega \in \mathbb{S}^2 \tag{7}$$

Unlike conventional patch-based regression, we treat each f_x and g_x as functions on $\mathbb{S}^2$. This allows the angular super-resolution task to be written as a spherical solution-operator learning problem.

4.1 CSD-inspired spherical fiber operator system

CSD models the measured diffusion signal as a spherical convolution (Tournier et al., 2007): $S_x(\omega) = (R_x * F_x)(\omega)$, where R_x is the response function and F_x is the underlying FOD. In clinical low-angular-resolution protocols, the recovered FOD f_x is an imperfect approximation of the desired high-angular-resolution FOD g_x. We therefore regard f_x as an observed low-bandwidth source function and g_x as the solution of an underlying spherical operator system $\mathcal{D}_x$:

$$\mathcal{D}_x\big(g_x(\omega)\big) = f_x(\omega), \quad \omega \in \mathbb{S}^2, \tag{8}$$

where $\mathcal{D}_x$ is a voxel-dependent spherical differential or pseudo-differential operator. This operator represents the degradation from high-angular-resolution FOD to low-angular-resolution FOD, including angular blurring, limited SH bandwidth, noise, and tissue-dependent uncertainty. The objective is to learn the inverse solution process:

$$g_x(\omega) = \mathcal{D}_x^{-1} f_x(\omega). \tag{9}$$

We next derive the solution to this spherical fiber operator system $\mathcal{D}_x$.

4.2 Operator Solution for spherical fiber system

Unlike general spherical signals, fiber orientation distributions (FODs) are antipodally symmetric due to the orientation ambiguity of white matter fibers:

$$f_x(\omega) = f_x(-\omega), \quad \omega \in \mathbb{S}^2. \tag{10}$$

According to the parity property of spherical harmonics $Y_l^m(-\omega) = (-1)^l Y_l^m(\omega)$, only even-degree spherical harmonics preserve antipodal symmetry. The input low-angular-resolution FOD $f_x(\omega)$ therefore inherits the antipodal symmetry of diffusion signals. Consequently, it admits an expansion over the even-degree spherical harmonic basis:

$$f_x(\omega) = \sum_{l=0}^{L} \sum_{m=-2l}^{2l} \hat{f}_{x,2l}^m Y_{2l}^m(\omega). \tag{11}$$

Based on the known input function $f_x(\omega)$, **our objective is to solve a solution target $g_x(\omega)$, which maintains the strict antipodal symmetry.** For solving this spherical fiber system as the Equation 8, we define the Green's function $G_x(\omega, \eta)$ satisfying:

$$\mathcal{D}_{x[w]} G_x(\omega, \eta) = \delta(\omega, \eta) = \begin{cases} \infty, & \omega = \eta, \\ 0, & \omega \neq \eta. \end{cases} \tag{12}$$

where $G_x(\omega, \eta)$ describes the response observed at ω generated by a unit impulse applied at orientation η. δ is the corresponding Dirac distribution on the sphere. Based on Eq. 12, the solution of the spherical fiber system is given by:

$$g_x(\omega) = \int_{\mathbb{S}^2} G_x(\omega, \eta) f_x(\eta) d\eta. \tag{13}$$

This solution is verified by applying $\mathcal{D}_x$ to both sides of Eq. 13 (Details in Appendix A). Therefore, Eq. 13 provides a valid Green's function solution on the antipodally symmetric spherical fiber space. With the SO(3) group equivariant inductive bias (Bonev et al., 2023), we obtain the Green's function on the antipodally symmetric basis:

$$G_x(\omega, \eta) = \sum_{l=0}^{L} \sum_{m=-2l}^{2l} \hat{G}_{x,2l}^m Y_{2l}^m(\eta^{-1}\omega). \tag{14}$$

Substituting into the Green's function solution in the Eq. 13 gives:

$$g_x(\omega) = \int_{\mathbb{S}^2} \left(\sum_{l=0}^{L} \sum_{m=-2l}^{2l} \hat{G}_{x,2l}^m Y_{2l}^m(\eta^{-1}\omega) \right) f_x(\eta) d\eta. \tag{15}$$

Based on spherical harmonic transform ($\mathcal{F}_{\mathbb{S}^2}$), we utilize the rotational property and orthogonality of spherical harmonics to yield:

$$\hat{g}_x = \mathcal{F}_{\mathbb{S}^2}[g_x(\omega)] = \sum_{m'=-2l}^{2l} \hat{G}_{x,2l}^{m'} \int_{SO(3)} f_x(Rn) \overline{D_{m',m}^{(2l)}(R^{-1})} dR. \tag{16}$$

To obtain the further solution, we let $I_{m'} = \int_{SO(3)} f_x(Rn) \overline{D_{m',m}^{(2l)}(R^{-1})} dR$. For an arbitrary rotation $R_n(\psi)$ around the reference axis n, we have $R_n(\psi)n = n$. Using the right invariance of the Haar measure, we obtain $I_{m'} = e^{-im'\psi} I_{m'}$ (Driscoll & Healy, 1994). Since this equality holds for every ψ, it follows that $I_{m'} = 0, \quad m' \neq 0$, and therefore only the zonal Green response survives:

$$\hat{g}_x = \hat{G}_{x,2l}^0 \int_{SO(3)} f_x(Rn) \overline{D_{0,m}^{(2l)}(R^{-1})} dR. \tag{17}$$

Using the relation between the Wigner-D matrix and spherical harmonics (Eq. 4), together with the decomposition of the Haar measure on $SO(3)$, the remaining integral reduces to the spherical harmonic coefficient of the input FOD:

$$\int_{SO(3)} f_x(Rn) \overline{D_{0,m}^{(2l)}(R^{-1})} dR = 2\pi \sqrt{\frac{4\pi}{4l+1}} \hat{f}_{x,2l}^m. \tag{18}$$

Consequently, Eq. 17 becomes:

$$\hat{g}_x = 2\pi\sqrt{\frac{4\pi}{4l+1}}\hat{G}^0_{x,2l}\hat{f}^m_{x,2l}. \tag{19}$$

Since the input FOD contains only even-degree spherical harmonic coefficients $\hat{f}^m_{x,2l}$ and the learned Green response is parameterized on the same even-degree harmonic subspace, Eq. (19) shows that the solution $\hat{g}_x$ also contains only even-degree coefficients. Consequently, the recovered spherical function is naturally expressed as $g_x(\omega) = \sum_{l=0}^{L}\sum_{m=-2l}^{2l}\hat{g}^m_{x,2l}Y^m_{2l}(\omega)$, which automatically satisfies the antipodal symmetry required by the physical definition of FODs. Therefore, the proposed NFO preserves antipodal symmetry as an explicit constraint. Finally, by defining the learnable parameter: $W_\theta(2l) = 2\pi\sqrt{\frac{4\pi}{2l\times 2+1}}\hat{G}^0_{x,2l}$, and utilizing the inverse spherical harmonic transform ($\mathcal{F}^{-1}_{\mathbb{S}^2}$), the neural fiber operator solution is finally derived as:

$$g_x(\omega) = \mathcal{F}^{-1}_{\mathbb{S}^2}(\hat{g}_x) = \mathcal{F}^{-1}_{\mathbb{S}^2}\left[W_\theta(2l)\hat{f}^m_{x,2l}\right](\omega). \tag{20}$$

Consequently, NFO operates entirely within the antipodally symmetric fiber harmonic space, preserving the physical structure of FODs.

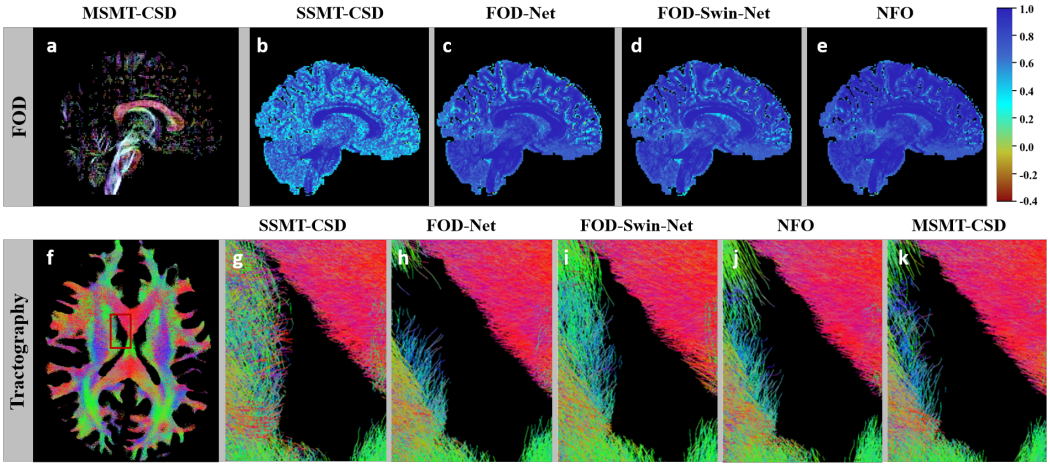


Figure 2: Held-out reconstruction comparison on the HCP-YA dataset. (a–e) show the sagittal FOD reference and ACC heatmaps for different methods. (f–k) show axial tractography results and corresponding zoomed-in views.

4.3 Neural Fiber Operators and Full Model Architecture

Neural Fiber Operator. Eq. 20 provides the analytical solution of the proposed fiber operator. Based on this formulation, we implement Neural Fiber Operator (NFO), illustrated in Fig. 1. Each NFO block directly realizes Eq. 20. Given an input FOD function $f_x(\omega)$, the block first utilize spherical harmonic transform to obtain the spherical harmonic coefficients $\hat{f}^m_{x,2l}$, then applies the learnable kernel $W_\theta(2l)$ on the even-degree coefficients, and finally reconstructs the learned fiber representation through the inverse spherical harmonic transform. Besides, a lightweight 1×1 convolution layer is introduced as the residual layer to enhance channel interaction, whose output is concatenated with the operator branch before nonlinear activation. Since every NFO block operates exclusively on the even-degree spherical harmonic subspace and the full network consists of multiple NFO blocks, the antipodal symmetry of FODs is preserved throughout the entire network.

Full NFO Model Architecture. To perform whole-brain or partial-brain FOD reconstruction with a large number of brain voxels, we adopt the same patch-learning strategy as FOD-Net (Zeng et al.,

Table 1: Held-out evaluation of different methods on HCP-YA dataset (ACC↑ and RMSE↓ on multiple tissues, including whole brain, white matter (WM), cortical gray matter (CGM) and subcortical gray matter (SGM)). Bold is the best result.

Methods		SSMT-CSD (Khan et al., 2020)	FOD-Net (Zeng et al., 2022)	FOD-Swin-Net (Da Silva et al., 2024)	**NFO**
Parameters		–	19.59 M	48.17 M	**2.74 M**
Reasoning Speed		–	161.36 s/sample	**7.14 s/sample**	282.86 s/sample
Whole Brain	ACC	0.5053 ± 0.0147	0.6833 ± 0.0082	0.6652 ± 0.0097	$\mathbf{0.7533 \pm 0.0106}$
	RMSE	0.0369 ± 0.0023	0.0169 ± 0.0010	0.0201 ± 0.0007	$\mathbf{0.0149 \pm 0.0009}$
White Matter	ACC	0.6684 ± 0.0211	0.8435 ± 0.0159	0.8223 ± 0.0154	$\mathbf{0.8871 \pm 0.0145}$
	RMSE	0.0409 ± 0.0022	0.0190 ± 0.0011	0.0248 ± 0.0006	$\mathbf{0.0171 \pm 0.0010}$
WM & CGM	ACC	0.5313 ± 0.0164	0.7413 ± 0.0138	0.7347 ± 0.0152	$\mathbf{0.8163 \pm 0.0157}$
	RMSE	0.0430 ± 0.0027	0.0180 ± 0.0012	0.0208 ± 0.0009	$\mathbf{0.0157 \pm 0.0011}$
WM & SGM	ACC	0.5641 ± 0.0226	0.7322 ± 0.0294	0.7004 ± 0.0305	$\mathbf{0.7588 \pm 0.0295}$
	RMSE	0.0429 ± 0.0023	0.0238 ± 0.0016	0.0285 ± 0.0011	$\mathbf{0.0223 \pm 0.0016}$

2022). As shown in the Fig. 1 (top row), for each valid center voxel, a local $9 \times 9 \times 9$ FOD patch is extracted, containing $9^3 = 729$ neighboring fiber functions. The full NFO network progressively aggregate these representations and reduce to a single fiber representation corresponding to the target voxel $g(\omega)$ as the final operator solution. Therefore, the full operator model architecture can be viewed as learning a local fiber operator: $\mathcal{T}_\theta : \{f_1(\omega), \ldots, f_{729}(\omega)\} \longrightarrow g(\omega)$.

5 Experiments and Results

We designed the experiments to evaluate both the in-domain reconstruction accuracy of NFO and the generalizability of the learned fiber solution operator. Specifically, all trainable models are optimized exclusively on the Human Connectome Project Young Adult dataset (HCP-YA) (Van Essen et al., 2013; Sotiropoulos et al., 2013), and are then evaluated under progressively more challenging settings: (1). large-scale whole-brain population modeling on held-out HCP-YA (Van Essen et al., 2013); (2). zero-shot transfer to the MGH–USC Human Connectome Project (MGH), assessing generalization across age groups and acquisition protocols (Fan et al., 2016); (3). transfer to microscopy-correlated ex-vivo corpus callosum (CC) brain specimens (Mollink et al., 2017), representing a major change in tissue condition, field strength, spatial resolution, diffusion weighting, and anatomical coverage (Tendler et al., 2022). Details about datasets, evaluation methods, metrics and model implementation are in Appendix B.

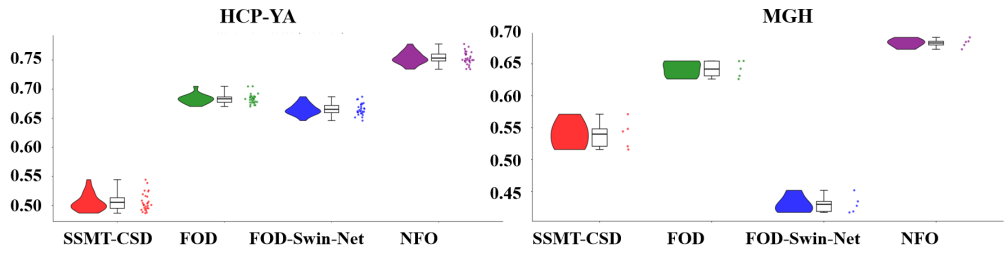


Figure 3: Whole brain ACC comparison of different methods by held-out reconstruction on the HCP-YA dataset and zero-shot reconstruction on the MGH dataset.

5.1 Large-Scale Whole-Brain Fiber Reconstruction

We first evaluate the proposed NFO on the large-scale HCP-YA benchmark, where both training and testing are performed on healthy young subjects under the same acquisition protocol. NFO consistently achieves the best performance across all evaluation settings (Table 1 and Fig. 2, details

Table 2: Zero-shot performance comparison of different methods on MGH and CC dataset (trained on HCP-YA dataset) (ACC↑ and RMSE↓). Bold represents the best result.

	Methods		SSMT-CSD (Khan et al., 2020)	FOD-Net (Zeng et al., 2022)	FOD-Swin-Net (Da Silva et al., 2024)	**NFO**
	Parameters		–	19.59 M	48.17 M	**2.74 M**
MGH	Reasoning Speed		–	124.10 s/sample	**2.51 s/sample**	244.31 s/sample
	Whole Brain	ACC	0.5399 ± 0.0224	0.6417 ± 0.0130	0.4298 ± 0.0141	**0.6827 ± 0.0070**
		RMSE	0.0471 ± 0.0028	0.0326 ± 0.0020	0.0381 ± 0.0010	**0.0323 ± 0.0019**
CC	Reasoning Speed		–	8.36 s/sample	**1.31 s/sample**	18.33 s/sample
	Corpus Callosum	ACC	0.7841 ± 0.0220	0.7894 ± 0.0142	0.4089 ± 0.0157	**0.8333 ± 0.0179**
		RMSE	0.0315 ± 0.0069	0.0273 ± 0.0050	0.0959 ± 0.0131	**0.0253 ± 0.0045**

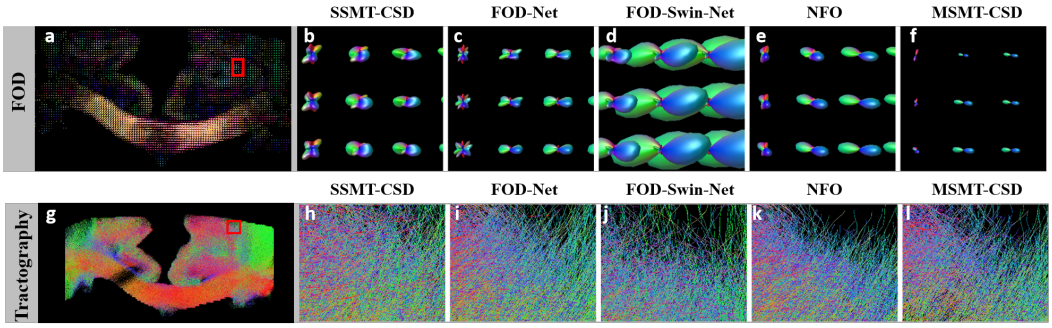


Figure 4: Zero-shot reconstruction comparison on the CC dataset. FOD reconstructions (b-f) and tractography results (h-l) are shown for zoomed-in regions indicated by the red boxes in the corpus callosum specimens.

shown in Table 6 and Fig. 6 in Appendix), showing its effectiveness across different tissue compositions, especially in regions with mixed tissue types. NFO is also highly parameter-efficient, with only 2.74M trainable parameters, significantly less than FOD-Net and FOD-Swin-Net, while consistently achieving the highest reconstruction accuracy.

To assess whether improvements in FOD reconstruction translate into fiber-level benefits, we further examine ACC heatmaps and anatomically constrained tractography (ACT). The pipeline of generating ACT following (Smith et al., 2012) is described in Appendix B.1.3. Fig. 2 qualitatively compares the ACC heatmaps and corresponding tractography. Consistent with the quantitative results, NFO reconstructs FOD profiles that more closely resemble the MSMT-CSD reference, with better-preserved principal orientations, clearer separation of crossing fiber populations, and fewer spurious angular responses.

5.2 Zero-Shot Generalization on Independent Clinical Cohorts

To evaluate whether the learned operator generalizes beyond the training distribution, we directly apply the model trained on HCP-YA to MGH datasets without any fine-tuning. Compared with HCP-YA, MGH introduces simultaneous changes in subject population, scanner, spatial resolution, diffusion shells and acquisition protocol, representing a substantial distribution shift from the training data. Table 2 (details in Table 7 in Appendix) reports the zero-shot performance on MGH. NFO achieves the best performance, demonstrating that the proposed framework preserves its advantage under both anatomical and acquisition variations. The fiber-level comparison of different methods further verify its advantage (Figure 7 in Appendix).

5.3 Cross-Domain Biological Validation

We next evaluate NFO on the microscopy-correlated ex-vivo CC dataset, representing the genuine biological domain shift from in-vivo to ex-vivo diffusion MRI. Compared with the in-vivo HCP datasets, CC differs simultaneously in tissue condition, magnetic field strength, spatial resolution,

Table 3: Connectome Comparison of different methods on HCP-YA by mean connectome disparity (MCD) and significant edges with different significance thresholds α. Bold is the best result.

Methods		SSMT-CSD (Khan et al., 2020)	FOD-Net (Zeng et al., 2022)	FOD-Swin-Net (Da Silva et al., 2024)	**NFO**
MCD ($\times 10^{-3}$)↓ (MEAN ± SD)		7.12 ± 0.94	3.94 ± 0.67	6.49 ± 0.90	$\mathbf{3.73 \pm 0.45}$
Significant Edges↓	$\alpha = 0.01$	4458	612	2993	**328**
	$\alpha = 0.05$	5633	1285	4203	**803**
	$\alpha = 0.10$	6302	1809	4956	**1222**

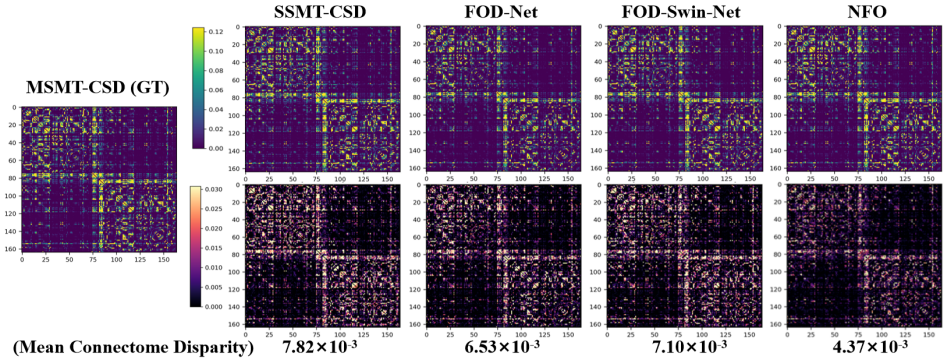


Figure 5: Reconstructed connectomes from different methods on HCP-YA dataset (first row) and the corresponding error maps (second row).

diffusion protocol and imaging domain. NFO consistently achieves the highest ACC on all three individual specimens, demonstrating stable zero-shot transfer across independent ex-vivo samples (Table 2, details in Table 8 in Appendix). Despite being trained exclusively on in-vivo HCP-YA data, NFO reconstructs FOD configurations that better preserve the complex local fiber organization observed in the reference, leading to improved tractography that more closely follows the underlying fiber pathways (Fig. 4). Together, these results support that NFO is a transferable and physically consistent operator for fiber reconstruction.

5.4 Structural Connectome Reconstruction

Although FOD reconstruction is the primary objective, its practical value also includes supporting reliable downstream structural connectivity analysis. We therefore investigate whether improvement in fiber reconstruction enhance the connectome reconstruction by comparing connectomes reconstructed from different methods against those obtained from the reference MSMT-CSD. The pipeline of generating connectomes is in Appendix B.1.3. Table 3 (details in Table 9 in Appendix) indicates that the connectomes of reconstruction results of NFO have the lowest mean connectome disparity and smallest amount of significant edges under different thresholds. As shown in Fig. 5, NFO produces connectivity patterns that more closely agree with the reference and achieves the lowest mean connectome disparity, indicating the advantage of NFO extends to preserveing anatomically meaningful fiber organization.

6 Conclusion

In this work, we introduced Neural Fiber Operators (NFO), a neural operator framework for FOD angular super-resolution. NFO formulates the prediction of FODs as learning a operator between fiber functions. Based on the CSD formulation, we derived the corresponding operator solution in the spherical harmonic domain while preserving the antipodal symmetry of FOD and implemented it as a deep neural operator architecture for whole-brain or partial-brain FOD reconstruction. Our experiments demonstrated that NFO consistently outperforms existing approaches while using substantially fewer parameters. Besides, NFO achieves generalization across different protocols and domains, indicating that the learned operator captures transferable spherical fiber function mappings. More broadly, this work provides a new perspective on diffusion MRI reconstruction by explicitly incorporating the intrinsic geometry of FOD into neural operator learning.

AI USE STATEMENT

In this work, we used generative AI tools for language polishing and writing assistance to improve the clarity, grammar, and readability of the manuscript. All core research contributions, including the research ideas, methodological development, theoretical derivations, experimental design and implementation, data analysis, and scientific conclusions, were developed and conducted by the authors. All AI-assisted text was carefully reviewed and verified by the authors. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This study uses existing, publicly available diffusion MRI datasets, including HCP-YA, MGH-HCP, and ex-vivo CC data, in accordance with their respective data-use agreements. No new human participants were recruited, and no additional human-subject experiments were conducted. All analyses were performed on de-identified research data, and no attempt was made to identify individual participants. The proposed method is intended for research purposes and has not been validated for clinical diagnosis or decision-making.

REPRODUCIBILITY STATEMENT

Appendix provides detailed methodological and implementation descriptions, experimental configurations, dataset analyses, and data source links.

The complete implementation `code` of the proposed method is provided in the **Supplementary Material** to further strengthen the reproducibility.

REFERENCES

Joseph J Bartlett, Catherine E Davey, Leigh A Johnston, and Jinming Duan. Recovering high-quality fiber orientation distributions from a reduced number of diffusion-weighted images using a model-driven deep learning architecture. *Magnetic Resonance in Medicine*, 92(5):2193–2206, 2024.

Peter J Basser, James Mattiello, and Denis LeBihan. Mr diffusion tensor spectroscopy and imaging. *Biophysical journal*, 66(1):259–267, 1994.

Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):289–300, 1995.

Kaushik Bhattacharya, Bamdad Hosseini, Nikola B Kovachki, and Andrew M Stuart. Model reduction and neural networks for parametric pdes. *The SMAI journal of computational mathematics*, 7:121–157, 2021.

Boris Bonev, Thorsten Kurth, Christian Hundt, Jaideep Pathak, Maximilian Baust, Karthik Kashinath, and Anima Anandkumar. Spherical fourier neural operators: Learning stable dynamics on the sphere. In *International conference on machine learning*, pp. 2806–2823. PMLR, 2023.

Mateus Oliveira Da Silva, Caio Pinheiro Santana, Diedre Santos Do Carmo, and Letícia Rittner. Fod-swin-net: angular super resolution of fiber orientation distribution using a transformer-based deep model. In *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 1–5. IEEE, 2024.

Flavio Dell'Acqua and J-Donald Tournier. Modelling white matter with spherical deconvolution: How and why? *NMR in Biomedicine*, 32(4):e3945, 2019.

Christophe Destrieux, Bruce Fischl, Anders Dale, and Eric Halgren. Automatic parcellation of human cortical gyri and sulci using standard anatomical nomenclature. *Neuroimage*, 53(1):1–15, 2010.

James R Driscoll and Dennis M Healy. Computing fourier transforms and convolutions on the 2-sphere. *Advances in applied mathematics*, 15(2):202–250, 1994.

Qiuyun Fan, Thomas Witzel, Aapo Nummenmaa, Koene RA Van Dijk, John D Van Horn, Michelle K Drews, Leah H Somerville, Margaret A Sheridan, Rosario M Santillana, Jenna Snyder, et al. Mgh–usc human connectome project datasets with ultra-high b-value diffusion mri. *Neuroimage*, 124:1108–1114, 2016.

Bruce Fischl. Freesurfer. *Neuroimage*, 62(2):774–781, 2012.

Matthew F Glasser, Stamatios N Sotiropoulos, J Anthony Wilson, Timothy S Coalson, Bruce Fischl, Jesper L Andersson, Junqian Xu, Saad Jbabdi, Matthew Webster, Jonathan R Polimeni, et al. The minimal preprocessing pipelines for the human connectome project. *Neuroimage*, 80:105–124, 2013.

Yifan Hu, Fukang Yin, Weimin Zhang, Kaijun Ren, Junqiang Song, Kefeng Deng, and Di Zhang. Spherical multigrid neural operator for improving autoregressive global weather forecasting. *Scientific Reports*, 15(1):11522, 2025.

Ben Jeurissen, Jacques-Donald Tournier, Thijs Dhollander, Alan Connelly, and Jan Sijbers. Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion mri data. *NeuroImage*, 103:411–426, 2014.

Ben Jeurissen, Maxime Descoteaux, Susumu Mori, and Alexander Leemans. Diffusion mri fiber tractography of the brain. *NMR in Biomedicine*, 32(4):e3785, 2019.

Ranjeet Ranjan Jha, Aditya Nigam, Arnav Bhavsar, Sudhir K Pathak, Walter Schneider, and K Rathish. Multi-shell d-mri reconstruction via residual learning utilizing encoder-decoder network with attention (msr-net). In *2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pp. 1709–1713. IEEE, 2020.

Ranjeet Ranjan Jha, Sudhir K Pathak, Vishwesh Nath, Walter Schneider, BV Rathish Kumar, Arnav Bhavsar, and Aditya Nigam. Vrfrnet: Volumetric roi fodf reconstruction network for estimation of multi-tissue constrained spherical deconvolution with only single shell dmri. *Magnetic Resonance Imaging*, 90:1–16, 2022.

Davood Karimi, Lana Vasung, Camilo Jaimes, Fedel Machado-Rivas, Simon K Warfield, and Ali Gholipour. Learning to estimate the fiber orientation distribution function from diffusion-weighted mri. *NeuroImage*, 239:118316, 2021.

Wasim Khan, Natalia Egorova, Mohamed Salah Khlif, Remika Mito, Thijs Dhollander, and Amy Brodtmann. Three-tissue compositional analysis reveals in-vivo microstructural heterogeneity of white matter hyperintensities following stroke. *Neuroimage*, 218:116869, 2020.

Nikola B Kovachki, Samuel Lanthaler, and Andrew M Stuart. Operator learning: Algorithms and analysis. *Handbook of Numerical Analysis*, 25:419–467, 2024.

Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.

Kenghong Lin, Xutao Li, Yunming Ye, Shanshan Feng, Baoquan Zhang, Guangning Xu, and Ziye Wang. Spherical neural operator network for global weather prediction. *IEEE Transactions on Circuits and Systems for Video Technology*, 34(6):4899–4913, 2023.

Lu Lu, Pengzhan Jin, and George Em Karniadakis. Deeponet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *arXiv preprint arXiv:1910.03193*, 2019.

Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature machine intelligence*, 3(3):218–229, 2021.

Ankur Mahesh, William Collins, Boris Bonev, Noah Brenowitz, Yair Cohen, Joshua Elms, Peter Harrington, Karthik Kashinath, Thorsten Kurth, Joshua North, et al. Huge ensembles part i: Design of ensemble weather forecasts using spherical fourier neural operators. *arXiv preprint arXiv:2408.03100*, 2024a.

Ankur Mahesh, William Collins, Boris Bonev, Noah Brenowitz, Yair Cohen, Peter Harrington, Karthik Kashinath, Thorsten Kurth, Joshua North, Travis OBrien, et al. Huge ensembles part ii: Properties of a huge ensemble of hindcasts generated with spherical fourier neural operators. *arXiv preprint arXiv:2408.01581*, 2024b.

Jeroen Mollink, Michiel Kleinnijenhuis, Anne-Marie van Cappellen van Walsum, Stamatios N Sotiropoulos, Michiel Cottaar, Christopher Mirfin, Mattias P Heinrich, Mark Jenkinson, Menuka Pallebage-Gamarallage, Olaf Ansorge, et al. Evaluating fibre orientation dispersion in white matter: Comparison of diffusion mri, histology and polarized light imaging. *Neuroimage*, 157: 561–574, 2017.

Claus Müller. *Spherical harmonics*, volume 17. Springer, 2006.

Vishwesh Nath, Sudhir K Pathak, Kurt G Schilling, Walt Schneider, and Bennett A Landman. Deep learning estimation of multi-tissue constrained spherical deconvolution with limited single shell dw-mri. In *Proceedings of SPIE–the International Society for Optical Engineering*, volume 11313, pp. 113130S, 2020.

Sara Sedlar, Théodore Papadopoulo, Rachid Deriche, and Samuel Deslauriers-Gauthier. Diffusion mri fiber orientation distribution function estimation using voxel-wise spherical u-net. In *Computational Diffusion MRI: International MICCAI Workshop, Lima, Peru, October 2020*, pp. 95–106. Springer, 2021.

Ittai Shamir and Yaniv Assaf. Tutorial: a guide to diffusion mri and structural connectomics. *Nature Protocols*, 20(2):317–335, 2025.

Robert E Smith, Jacques-Donald Tournier, Fernando Calamante, and Alan Connelly. Anatomically-constrained tractography: improved diffusion mri streamlines tractography through effective use of anatomical information. *Neuroimage*, 62(3):1924–1938, 2012.

Haykel Snoussi and Davood Karimi. Equivariant spherical cnns for accurate fiber orientation distribution estimation in neonatal diffusion mri with reduced acquisition time. *Frontiers in Neuroscience*, 19:1604545, 2025.

Stamatios N Sotiropoulos, Saad Jbabdi, Junqian Xu, Jesper L Andersson, Steen Moeller, Edward J Auerbach, Matthew F Glasser, Moises Hernandez, Guillermo Sapiro, Mark Jenkinson, et al. Advances in diffusion mri acquisition and processing in the human connectome project. *Neuroimage*, 80:125–143, 2013.

Hao Tang, Hao Chen, and Chao Li. Generalized spherical neural operators: Green's function formulation. In *International Conference on Learning Representations*, volume 2026, pp. 11180–11201, 2026.

Benjamin C Tendler, Taylor Hanayik, Olaf Ansorge, Sarah Bangerter-Christensen, Gregory S Berns, Mads F Bertelsen, Katherine L Bryant, Sean Foxley, Martijn P van den Heuvel, Amy FD Howard, et al. The digital brain bank, an open access platform for post-mortem imaging datasets. *Elife*, 11:e73153, 2022.

J-Donald Tournier, Fernando Calamante, David G Gadian, and Alan Connelly. Direct estimation of the fiber orientation density function from diffusion-weighted mri data using spherical deconvolution. *Neuroimage*, 23(3):1176–1185, 2004.

J-Donald Tournier, Fernando Calamante, and Alan Connelly. Robust determination of the fibre orientation distribution in diffusion mri: non-negativity constrained super-resolved spherical deconvolution. *Neuroimage*, 35(4):1459–1472, 2007.

J-Donald Tournier, Robert Smith, David Raffelt, Rami Tabbara, Thijs Dhollander, Maximilian Pietsch, Daan Christiaens, Ben Jeurissen, Chun-Hung Yeh, and Alan Connelly. Mrtrix3: A fast, flexible and open software framework for medical image processing and visualisation. *Neuroimage*, 202:116137, 2019.

David S Tuch. Q-ball imaging. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 52(6):1358–1372, 2004.

David C Van Essen, Stephen M Smith, Deanna M Barch, Timothy EJ Behrens, Essa Yacoub, Kamil Ugurbil, Wu-Minn HCP Consortium, et al. The wu-minn human connectome project: an overview. *Neuroimage*, 80:62–79, 2013.

Anderson M. Winkler, Gerard R. Ridgway, Matthew A. Webster, Stephen M. Smith, and Thomas E. Nichols. Permutation inference for the general linear model. *NeuroImage*, 92:381–397, 2014.

Tianyuan Yao, Francois Rheault, Leon Y Cai, Vishwesh Nath, Zuhayr Asad, Nancy Newlin, Can Cui, Ruining Deng, Karthik Ramadass, Andrea Shafer, 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):014005–014005, 2024.

Rui Zeng, Jinglei Lv, He Wang, Luping Zhou, Michael Barnett, Fernando Calamante, and Chenyu Wang. Fod-net: A deep learning method for fiber orientation distribution angular super resolution. *Medical Image Analysis*, 79:102431, 2022.

A Detailed Derivation and Analysis of Neural Fiber Operator

Based on Eq. 12, the solution of the spherical fiber system is given by

$$g_x(\omega) = \int_{\mathbb{S}^2} G_x(\omega, \eta) f_x(\eta)\, d\eta. \tag{21}$$

To verify this solution, we apply $\mathcal{D}_x$ to both sides:

$$\begin{aligned}
\mathcal{D}_{x[\omega]} g_x(\omega) &= \mathcal{D}_{x[\omega]} \left[\int_{\mathbb{S}^2} G_x(\omega, \eta) f_x(\eta)\, d\eta \right] \\
&= \int_{\mathbb{S}^2} \mathcal{D}_{x,\omega} G_x(\omega, \eta) f_x(\eta)\, d\eta \\
&= \int_{\mathbb{S}^2} \delta(\omega, \eta) f_x(\eta)\, d\eta \\
&= f_x(\omega).
\end{aligned} \tag{22}$$

Therefore, Eq. 21 provides a valid Green's function representation of the spherical fiber system. Let $n \in \mathbb{S}^2$ denote a fixed reference direction and, for each $\eta \in \mathbb{S}^2$, let $R_\eta \in SO(3)$ be a rotation satisfying $R_\eta n = \eta$. Under the $SO(3)$-equivariant inductive bias (Bonev et al., 2023), the rotational Green response can be expanded in spherical harmonics as

$$G_x(R_\eta^{-1}\omega) = \sum_{l=0}^{L} \sum_{m'=-2l}^{2l} \hat{G}_{x,2l}^{m'} Y_{2l}^{m'}(R_\eta^{-1}\omega). \tag{23}$$

Accordingly, the Green's function solution becomes

$$g_x(\omega) = \int_{\mathbb{S}^2} \left[\sum_{l=0}^{L} \sum_{m'=-2l}^{2l} \hat{G}_{x,2l}^{m'} Y_{2l}^{m'}(R_\eta^{-1}\omega) \right] f_x(\eta)\, d\eta. \tag{24}$$

Spherical Harmonic Transformation. We now derive the spectral form of Eq. 24. For a fixed even degree $2l$ and order m, the spherical harmonic coefficient of the output is

$$\hat{g}_{x,2l}^{m} = \int_{\mathbb{S}^2} g_x(\omega) \overline{Y_{2l}^{m}(\omega)}\, d\omega. \tag{25}$$

Substituting Eq. 24 into Eq. 25 yields

$$\begin{aligned}
\hat{g}_{x,2l}^{m} &= \int_{\mathbb{S}^2} \int_{\mathbb{S}^2} \left[\sum_{l'=0}^{L} \sum_{m'=-2l'}^{2l'} \hat{G}_{x,2l'}^{m'} Y_{2l'}^{m'}(R_\eta^{-1}\omega) \right] f_x(\eta) \overline{Y_{2l}^{m}(\omega)}\, d\eta\, d\omega \\
&= \sum_{l'=0}^{L} \sum_{m'=-2l'}^{2l'} \hat{G}_{x,2l'}^{m'} \int_{\mathbb{S}^2} f_x(\eta) \left[\int_{\mathbb{S}^2} Y_{2l'}^{m'}(R_\eta^{-1}\omega) \overline{Y_{2l}^{m}(\omega)}\, d\omega \right] d\eta.
\end{aligned} \tag{26}$$

Using the rotation law of spherical harmonics,

$$Y_{2l'}^{m'}(R_\eta^{-1}\omega) = \sum_{q=-2l'}^{2l'} D_{q,m'}^{(2l')}(R_\eta) Y_{2l'}^{q}(\omega), \tag{27}$$

the inner integral in Eq. 26 becomes

$$\begin{aligned}
&\int_{\mathbb{S}^2} Y_{2l'}^{m'}(R_\eta^{-1}\omega) \overline{Y_{2l}^{m}(\omega)}\, d\omega \\
&= \sum_{q=-2l'}^{2l'} D_{q,m'}^{(2l')}(R_\eta) \int_{\mathbb{S}^2} Y_{2l'}^{q}(\omega) \overline{Y_{2l}^{m}(\omega)}\, d\omega.
\end{aligned} \tag{28}$$

By the orthogonality of spherical harmonics (Müller, 2006),

$$\int_{\mathbb{S}^2} Y_{2l'}^{q}(\omega)\overline{Y_{2l}^{m}(\omega)}\,d\omega = \delta_{ll'}\delta_{qm}, \tag{29}$$

and hence

$$\int_{\mathbb{S}^2} Y_{2l'}^{m'}(R_\eta^{-1}\omega)\overline{Y_{2l}^{m}(\omega)}\,d\omega = \delta_{ll'}D_{m,m'}^{(2l)}(R_\eta). \tag{30}$$

Therefore, all degrees $l' \neq l$ vanish and Eq. 26 reduces to

$$\hat{g}_{x,2l}^{m} = \sum_{m'=-2l}^{2l} \hat{G}_{x,2l}^{m'} \int_{\mathbb{S}^2} f_x(\eta)D_{m,m'}^{(2l)}(R_\eta)\,d\eta. \tag{31}$$

Reduction to the Zonal Green Response. The sphere can be identified with the homogeneous space $\mathbb{S}^2 \simeq SO(3)/SO(2)$. We therefore lift the spherical function to $SO(3)$ by defining $\widetilde{f}_x(R) = f_x(Rn)$. Since any rotation $R_n(\psi)$ about the reference axis satisfies $R_n(\psi)n = n$, the lifted function is invariant under the corresponding right action (Driscoll & Healy, 1994):

$$\widetilde{f}_x(RR_n(\psi)) = f_x(RR_n(\psi)n) = f_x(Rn) = \widetilde{f}_x(R). \tag{32}$$

Consider

$$I_{m'} = \int_{SO(3)} f_x(Rn)\overline{D_{m',m}^{(2l)}(R^{-1})}\,dR. \tag{33}$$

Using the right invariance of the Haar measure and the change of variables $R \mapsto RR_n(\psi)$ gives

$$\begin{aligned} I_{m'} &= \int_{SO(3)} f_x(RR_n(\psi)n)\overline{D_{m',m}^{(2l)}\left((RR_n(\psi))^{-1}\right)}\,dR \\ &= \int_{SO(3)} f_x(Rn)\overline{D_{m',m}^{(2l)}\left(R_n(-\psi)R^{-1}\right)}\,dR. \end{aligned} \tag{34}$$

Based on the representation property $D^{(2l)}(R_n(-\psi)R^{-1}) = D^{(2l)}(R_n(-\psi))D^{(2l)}(R^{-1})$ and the principle that rotations around the reference axis are diagonal in the Wigner-D basis, $D_{m',q}^{(2l)}(R_n(-\psi)) = e^{im'\psi}\delta_{m'q}$ (Driscoll & Healy, 1994), we can obtain:

$$I_{m'} = e^{-im'\psi}I_{m'}. \tag{35}$$

Since Eq. 35 must hold for every $\psi \in [0, 2\pi)$,

$$I_{m'} = 0, \qquad m' \neq 0. \tag{36}$$

Thus, only the $m' = 0$ component of the rotational Green response contributes to the spherical output, yielding:

$$\hat{g}_{x,2l}^{m} = \hat{G}_{x,2l}^{0} \int_{SO(3)} f_x(Rn)\overline{D_{0,m}^{(2l)}(R^{-1})}\,dR. \tag{37}$$

Reduction to the Input Spherical Harmonic Coefficient. We next evaluate the remaining integral explicitly. Using unitarity of the Wigner-D representation (Driscoll & Healy, 1994),

$$\overline{D_{0,m}^{(2l)}(R^{-1})} = D_{m,0}^{(2l)}(R), \tag{38}$$

and the standard relation between spherical harmonics and Wigner-D matrices (Driscoll & Healy, 1994),

$$Y_{2l}^{m}(Rn) = \sqrt{\frac{4l+1}{4\pi}}\,\overline{D_{m,0}^{(2l)}(R)}, \tag{39}$$

we obtain:

$$D_{m,0}^{(2l)}(R) = \sqrt{\frac{4\pi}{4l+1}}\,\overline{Y_{2l}^{m}(Rn)}. \tag{40}$$

Substituting Eq. 40 into Eq. 37 gives

$$\hat{g}^m_{x,2l} = \hat{G}^0_{x,2l}\sqrt{\frac{4\pi}{4l+1}}\int_{SO(3)} f_x(Rn)\overline{Y^m_{2l}(Rn)}\,dR. \tag{41}$$

A rotation $R \in SO(3)$ can be decomposed into a direction $\eta = Rn \in \mathbb{S}^2$ and a residual rotation $R_n(\psi) \in SO(2)$ about the reference axis. With the corresponding Haar-measure decomposition, $dR = d\eta\, d\psi$, $\psi \in [0, 2\pi)$, and because both $f_x(Rn)$ and $Y^m_{2l}(Rn)$ depend only on $\eta = Rn$ (Driscoll & Healy, 1994), we obtain:

$$\begin{aligned}\int_{SO(3)} f_x(Rn)\overline{Y^m_{2l}(Rn)}\,dR &= \int_{\mathbb{S}^2}\int_0^{2\pi} f_x(\eta)\overline{Y^m_{2l}(\eta)}\,d\psi\,d\eta \\ &= 2\pi\int_{\mathbb{S}^2} f_x(\eta)\overline{Y^m_{2l}(\eta)}\,d\eta \\ &= 2\pi\hat{f}^m_{x,2l}.\end{aligned} \tag{42}$$

Substituting Eq. 42 into Eq. 41 therefore yields

$$\hat{g}^m_{x,2l} = 2\pi\sqrt{\frac{4\pi}{4l+1}}\hat{G}^0_{x,2l}\hat{f}^m_{x,2l}. \tag{43}$$

Antipodal Fiber Solution. Equation 43 establishes the spectral solution of the spherical fiber system. The input FOD satisfies antipodal symmetry, $f_x(-\omega) = f_x(\omega)$, and therefore contains only even-degree spherical harmonic coefficients:

$$\hat{f}^m_{x,k} = 0, \qquad k \text{ odd}. \tag{44}$$

Equation 43 acts independently on each surviving even degree $k = 2l$ and does not generate coefficients at odd degrees. Hence

$$\hat{g}^m_{x,k} = 0, \qquad k \text{ odd}, \tag{45}$$

and the recovered solution is

$$g_x(\omega) = \sum_{l=0}^{L}\sum_{m=-2l}^{2l}\hat{g}^m_{x,2l}Y^m_{2l}(\omega). \tag{46}$$

Using the parity property of spherical harmonics,

$$Y^m_k(-\omega) = (-1)^k Y^m_k(\omega), \tag{47}$$

we can directly verify the antipodal symmetry of the target solution g by:

$$\begin{aligned}g_x(-\omega) &= \sum_{l=0}^{L}\sum_{m=-2l}^{2l}\hat{g}^m_{x,2l}Y^m_{2l}(-\omega) \\ &= \sum_{l=0}^{L}\sum_{m=-2l}^{2l}(-1)^{2l}\hat{g}^m_{x,2l}Y^m_{2l}(\omega) \\ &= g_x(\omega).\end{aligned} \tag{48}$$

Thus, the derived operator maps an antipodally symmetric input fiber function to an antipodally symmetric output fiber function. Finally, we parameterize the effective Green response by:

$$W_\theta(2l) = 2\pi\sqrt{\frac{4\pi}{4l+1}}\hat{G}^0_{\theta,2l}. \tag{49}$$

Equation 43 becomes:

$$\hat{g}^m_{x,2l} = W_\theta(2l)\hat{f}^m_{x,2l}. \tag{50}$$

Applying the inverse spherical harmonic transform gives the final Neural Fiber Operator solution:

$$g_x(\omega) = \mathcal{F}^{-1}_{\mathbb{S}^2}\left[W_\theta(2l)\hat{f}^m_{x,2l}\right](\omega). \tag{51}$$

Consequently, NFO operates within the even-degree spherical harmonic subspace induced by the antipodal structure of the input FOD, and the resulting operator solution preserves this physical structure throughout the fiber-function mapping.

B EXPERIMENTAL DETAILS

B.1 EXPERIMENTAL SETUP

B.1.1 DATASETS AND EVALUATION OBJECTIVES

We evaluate NFO on four complementary diffusion MRI datasets, whose acquisition properties and roles in the generalization hierarchy are summarized in Tables 4 and 5. HCP-YA (Van Essen et al., 2013; Sotiropoulos et al., 2013) is the only dataset used for model optimization, whereas MGH and the microscopy-correlated ex-vivo CC dataset are used exclusively for zero-shot evaluation.

Following the protocol of FOD-Net (Zeng et al., 2022), we randomly select 110 HCP-YA subjects, including 50 for training, 10 for validation and 50 for testing. The target HAR-FOD is constructed from the complete multi-shell acquisition ($b = 0, 1000, 2000, 3000$), while the input LAR-FOD is estimated from one $b0$ image and 32 diffusion directions at $b = 1000$, simulating a clinically feasible single-shell acquisition.

MGH further evaluates transferability under an independent acquisition protocol (Fan et al., 2016). HAR-FOD is reconstructed from all available $b0$, $b = 1000$, $b = 3000$ and $b = 5000$ images, while the input again uses one $b0$ image and 32 directions at $b = 1000$. Consequently, the model must generalize to an unseen high-b acquisition without any parameter adaptation.

Finally, CC provides an extreme cross-domain benchmark consisting of three microscopy-correlated ex-vivo human brain specimens (Tendler et al., 2022; Mollink et al., 2017). Compared with HCP-YA, it introduces substantial shifts in tissue condition, magnetic field strength, spatial resolution and diffusion acquisition, allowing us to evaluate biological generalization beyond conventional in-vivo datasets. No external dataset is used for model selection, normalization estimation or fine-tuning.

Data Sources and Accessibility. All datasets used in this study are obtained from publicly accessible research repositories. The HCP-YA and MGH-HCP datasets are available through the BALSA database. The microscopy-correlated ex-vivo corpus callosum (CC) dataset is available through the Digital Brain Bank.

Table 4: Overview of the diffusion MRI datasets used in this study. The evaluation protocol is summarized separately in Table 5.

Property	HCP-YA (Van Essen et al., 2013)	MGH (Fan et al., 2016)	CC (Mollink et al., 2017)
Population	Healthy young adults	Healthy adults	Post-mortem brains
Age	22–35 yr	20–59 yr	51 / 65 / 91 yr
Tissue	In vivo	In vivo	Ex vivo
Scanner	Customized 3T Connectom	MGH Connectome	9.4T preclinical MRI
Resolution	1.25 mm	1.5 mm	0.4 mm
Target dMRI	$b = 0, 1000, 2000, 3000$	$b = 0, 1000, 3000, 5000$	$b \approx 0, 2500, 5000$
Clinical input	$b0$ & $32 \times b1000$	$b0$ & $32 \times b1000$	$b0$ & $32 \times b2500$

B.1.2 PREPROCESSING FOR FOD CONSTRUCTION

For every voxel, paired low- and high-angular-resolution FODs are constructed as

$$f_x(\omega) = \mathrm{CSD}(\mathcal{S}_{\mathrm{low},x}), \quad g_x^*(\omega) = \mathrm{CSD}(\mathcal{S}_{\mathrm{full},x}), \tag{52}$$

where $\mathcal{S}_{\mathrm{low}}$ and $\mathcal{S}_{\mathrm{full}}$ denote the clinically subsampled and complete diffusion measurements, respectively. HAR-FOD is estimated using MSMT-CSD, while LAR-FOD is reconstructed using SS3T-CSD. All FODs are represented by real antipodally symmetric spherical harmonics with $l_{\max} = 8$, retaining only even degrees and resulting in 45 SH coefficients per voxel.

Training samples are extracted as $9 \times 9 \times 9$ patches centred at valid brain voxels. The network predicts the HAR-FOD of the center voxel, and whole-brain reconstruction is obtained by sliding the trained model over all valid voxels. Unless otherwise specified, external datasets are evaluated at their native spatial resolution without resampling.

Table 5: Progressive operator generalization benchmarks. Only HCP-YA is used for training, while all remaining datasets are evaluated without fine-tuning.

Evaluation Property	HCP-YA	MGH	CC
Training	✓	–	–
Held-out testing	✓	–	–
Zero-shot evaluation	–	✓	✓
Cross-aging	–	✓	✓
Cross-scanner	–	✓	✓
Cross-acquisition protocol	–	✓	✓
Cross diffusion shells	–	✓	✓
Ultra-high diffusion weighting	–	✓	✓
In vivo → Ex vivo	–	–	✓

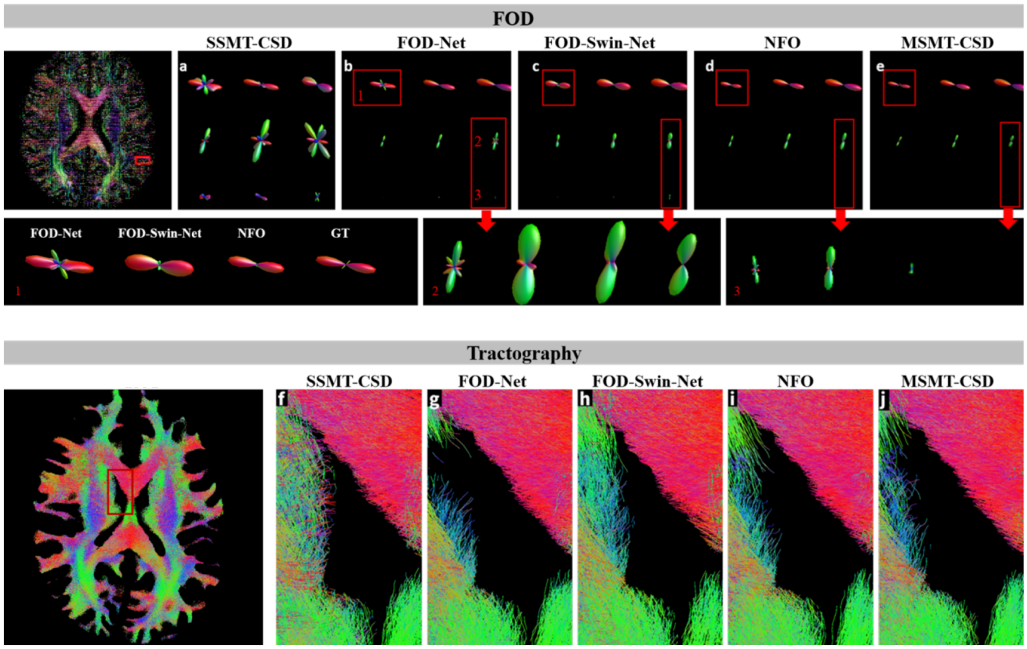


Figure 6: Held-out reconstruction comparison on the HCP-YA dataset. FOD reconstructions (a-e) and tractography results (f-j) are shown, both on axial cross-section. Zoomed-in display of voxels is also included.

B.1.3 Comparison Methods and Evaluation

We compare NFO with SSMT-CSD (Khan et al., 2020), the convolution-based FOD-Net (Zeng et al., 2022), and the Transformer-based FOD-Swin-Net (Da Silva et al., 2024). All learning-based methods use identical subject splits, preprocessing, patch size and evaluation masks whenever applicable. The FOD is compared on the reconstruction of all three datasets. For HCP-YA dataset with five-tissue-type (5TT) mask, our evaluation is also conducted using tractography and connectome, whose construction pipelines will be provided in the following content. The tractography is used for visualization of the angular errors of FOD and generating connectomes. The connectomes are compared by quantifiable metrics which will also be given.

Construction of Tractography To evaluate the practical impact of FOD reconstruction, tractography files of predicted and reference FODs are generated for qualitative and quantitative comparisons. As tractography propagates streamlines according to local fiber orientations, errors in FOD reconstruction can be directly indicated by fiber pathways. Whole-brain tractography is performed using `tckgen` command in MRtrix3 with the second-order integration over fiber orientation distributions

Table 6: Held-out evaluation of different methods on HCP-YA dataset for brain nerve fiber modeling (ACC↑ and RMSE↓, with UQ and LQ, on multiple tissues, including whole brain, WM, cortical gray matter CGM) and subcortical gray matter (SGM)). Bold represents the best result.

Methods		SSMT-CSD (Khan et al., 2020)	FOD-Net (Zeng et al., 2022)	FOD-Swin-Net (Da Silva et al., 2024)	**NFO**
Parameters		–	19.59 M	48.17 M	**2.74 M**
Reasoning Speed		–	161.36 s/sample	**7.14 s/sample**	282.86 s/sample
Whole Brain ACC	MEAN ± SD	0.5053 ± 0.0147	0.6833 ± 0.0082	0.6652 ± 0.0097	**0.7533 ± 0.0106**
	UQ	0.5134	0.6866	0.6720	**0.7603**
	LQ	0.4952	0.6774	0.6595	**0.7484**
Whole Brain RMSE	MEAN ± SD	0.0369 ± 0.0023	0.0169 ± 0.0010	0.0201 ± 0.0007	**0.0149 ± 0.0009**
	UQ	0.0412	0.0176	0.0207	**0.0155**
	LQ	0.0379	0.0159	0.0196	**0.0141**
White Matter ACC	MEAN ± SD	0.6684 ± 0.0211	0.8435 ± 0.0159	0.8223 ± 0.0154	**0.8871 ± 0.0145**
	UQ	0.6865	0.8554	0.8327	**0.8964**
	LQ	0.6524	0.8309	0.8108	**0.8772**
White Matter RMSE	MEAN ± SD	0.0409 ± 0.0022	0.0190 ± 0.0011	0.0248 ± 0.0006	**0.0171 ± 0.0010**
	UQ	0.0424	0.0197	0.0253	**0.0177**
	LQ	0.0394	0.0181	0.0244	**0.0162**
WM & CGM ACC	MEAN ± SD	0.5313 ± 0.0164	0.7413 ± 0.0138	0.7347 ± 0.0152	**0.8163 ± 0.0157**
	UQ	0.5413	0.7497	0.7459	**0.8284**
	LQ	0.5169	0.7342	0.7227	**0.8070**
WM & CGM RMSE	MEAN ± SD	0.0430 ± 0.0027	0.0180 ± 0.0012	0.0208 ± 0.0009	**0.0157 ± 0.0011**
	UQ	0.0446	0.0189	0.0215	**0.0165**
	LQ	0.0413	0.0172	0.0201	**0.0150**
WM & SGM ACC	MEAN ± SD	0.5641 ± 0.0226	0.7322 ± 0.0294	0.7004 ± 0.0305	**0.7588 ± 0.0295**
	UQ	0.5833	0.7462	0.7140	**0.7726**
	LQ	0.5490	0.7112	0.6809	**0.7409**
WM & SGM RMSE	MEAN ± SD	0.0429 ± 0.0023	0.0238 ± 0.0016	0.0285 ± 0.0011	**0.0223 ± 0.0016**
	UQ	0.0442	0.0253	0.0292	**0.0237**
	LQ	0.0412	0.0225	0.0276	**0.0209**

Table 7: Zero-shot performance comparison of different methods on MGH dataset (trained on HCP-YA dataset) for brain nerve fiber modeling (ACC↑ and RMSE↓,with UQ and LQ, on whole brain). Bold represents the best result.

Methods		SSMT-CSD (Khan et al., 2020)	FOD-Net (Zeng et al., 2022)	FOD-Swin-Net (Da Silva et al., 2024)	**NFO**
Parameters		–	19.59 M	48.17 M	**2.74 M**
Reasoning Speed		–	124.10 s/sample	**2.51 s/sample**	244.31 s/sample
Whole Brain ACC	MEAN ± SD	0.5399 ± 0.0224	0.6417 ± 0.0130	0.4298 ± 0.0141	**0.6827 ± 0.0070**
	UQ	0.5481	0.6543	0.4344	**0.6853**
	LQ	0.5208	0.6312	0.4185	**0.6797**
Whole Brain RMSE	MEAN ± SD	0.0471 ± 0.0028	0.0326 ± 0.0020	0.0381 ± 0.0010	**0.0323 ± 0.0019**
	UQ	0.0486	**0.0331**	0.0386	0.0333
	LQ	0.0457	0.0321	0.0375	**0.0320**

(iFOD2) algorithm. Anatomically constrained tractography (ACT) (Smith et al., 2012) is incorporated using the corresponding 5TT segmentation. The backtracking option is enabled, and streamline endpoints are cropped at the gray matter–white matter interface. Streamlines are randomly seeded within the whole-brain mask. The FOD amplitude cutoff is set to 0.05, the maximum curvature angle to 45°, and the minimum and maximum streamline lengths to 20 and 200 mm, respectively. Tracking is continued until 1 million valid streamlines were generated for each subject. On the CC dataset, microscopy registration is additionally used to assess biological consistency between reconstructed FODs and histological measurements.

Construction of Connectome Connectomes are also constructed to evaluate whether the improvement in FOD reconstruction supports reliable downstream structural connectivity analysis. As a direct method to verify the accuracy of predicted fiber, connectomes show the connectivity strength between regions guided by the fibers in tractography. In one connectome, each row and column

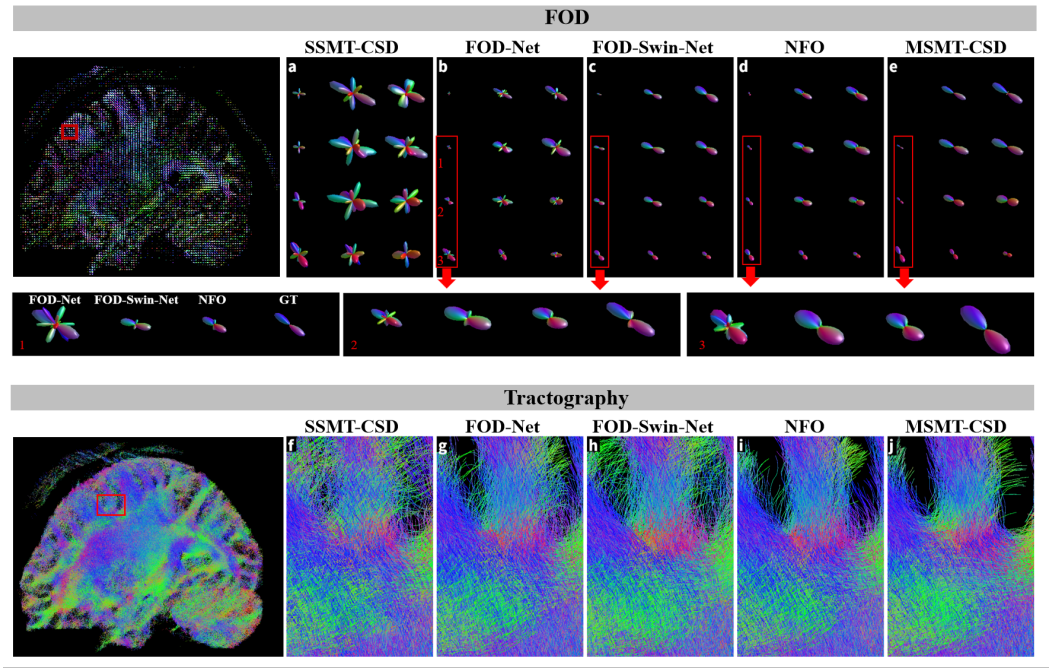


Figure 7: Zero-shot reconstruction comparison on the MGH dataset. FOD reconstructions (a-e) and tractography results (f-j) are shown, both on sagittal cross-section. Zoomed-in display of voxels is also included.

Table 8: Zero-shot performance comparison of different methods on CC dataset (trained on HCP-YA dataset) for brain nerve fiber modeling (ACC↑ and RMSE↓ on regions of corpus callosum). Bold represents the best result.

Methods		SSMT-CSD (Khan et al., 2020)	FOD-Net (Zeng et al., 2022)	FOD-Swin-Net (Da Silva et al., 2024)	**NFO**
Parameters		–	19.59 M	48.17 M	**2.74 M**
Reasoning Speed		–	8.36 s/sample	**1.31 s/sample**	18.33 s/sample
Corpus Callosum ACC	Sample 1	0.7594	0.7820	0.4177	**0.8127**
	Sample 2	0.8015	0.8057	0.4222	**0.8430**
	Sample 3	0.7915	0.7805	0.3868	**0.8442**
	MEAN ± SD	0.7841 ± 0.0220	0.7894 ± 0.0142	0.4089 ± 0.0157	**0.8333 ± 0.0179**
Corpus Callosum RMSE	Sample 1	0.0382	0.0310	0.1052	**0.0296**
	Sample 2	0.0244	0.0216	0.1014	**0.0206**
	Sample 3	0.0319	0.0292	0.0809	**0.0257**
	MEAN ± SD	0.0315 ± 0.0069	0.0273 ± 0.0050	0.0959 ± 0.0131	**0.0253 ± 0.0045**

represents an anatomical region and each off-diagonal element quantifies the structural connectivity between a pair of regions based on the streamlines connecting the corresponding nodes. Firstly, anatomical information is obtained from the T1-weighted image following the HCP preprocessing pipeline (Glasser et al., 2013). Cortical and subcortical regions are segmented and parcellated using FreeSurfer (Fischl, 2012) based on the Destrieux cortical atlas (Destrieux et al., 2010). The resulting labels are converted into a connectome-compatible node image using MRtrix3 and transformed into diffusion space (Tournier et al., 2019). Meanwhile, whole-brain tractography is reconstructed using ACT (Smith et al., 2012), with the corresponding 5TT mask image providing anatomical constraints. Streamline endpoints are then assigned to the anatomical nodes, and structural connectomes are finally generated using `tck2connectome` in MRtrix3 (Tournier et al., 2019).

Metric of FOD The primary evaluation metric is the angular correlation coefficient (ACC) (Zeng et al., 2022), which measures the similarity between predicted and target antipodally symmetric

Table 9: Comparison of different methods on the connectome of HCP-YA dataset by mean connectome disparity, thresholded edges (threshold of relative error are 10%, 20% and 50%) and significant edges with significance threshold α equals 0.01, 0.05 and 0.10. Bold represents the best result.

Methods		SSMT-CSD (Khan et al., 2020)	FOD-Net (Zeng et al., 2022)	FOD-Swin-Net (Da Silva et al., 2024)	**NFO**
Mean Connectome Disparity↓	MEAN ± SD ($\times 10^{-3}$)	7.12 ± 0.94	3.94 ± 0.67	6.49 ± 0.90	$\mathbf{3.73 \pm 0.45}$
	UQ ($\times 10^{-3}$)	7.79	4.23	7.05	**3.97**
	LQ ($\times 10^{-3}$)	6.15	3.58	5.94	**3.42**
Significant Edges↓	$\alpha = 0.01$	4458	612	2993	**328**
	$\alpha = 0.05$	5633	1285	4203	**803**
	$\alpha = 0.10$	6302	1809	4956	**1222**
Thresholded Edges (10%)↓	MEAN ± SD	8275.97 ± 372.45	7386.10 ± 372.42	8015.77 ± 391.79	$\mathbf{7322.93 \pm 346.07}$
	UQ	8537.25	7605.00	8336.25	**7521.00**
	LQ	7981.50	7120.50	7721.75	**7061.25**
Thresholded Edges (20%)↓	MEAN ± SD	7481.80 ± 342.18	6280.10 ± 328.95	7193.30 ± 359.06	$\mathbf{6171.60 \pm 278.98}$
	UQ	7724.75	6462.75	7482.00	**6317.25**
	LQ	7245.00	6063.50	6896.25	**5964.00**
Thresholded Edges (50%)↓	MEAN ± SD	5292.17 ± 289.43	3911.67 ± 152.89	4888.20 ± 266.88	$\mathbf{3797.47 \pm 137.32}$
	UQ	5466.75	4011.00	5111.50	**3897.00**
	LQ	5060.50	3784.25	4670.50	**3709.50**

FOD,

$$\mathrm{ACC}_x = \frac{\sum_{l,m} \widehat{g}^m_{x,l} \widehat{g}^{m,*}_{x,l}}{\sqrt{\sum_{l,m} (\widehat{g}^m_{x,l})^2} \sqrt{\sum_{l,m} (\widehat{g}^{m,*}_{x,l})^2} + \epsilon}, \tag{53}$$

where $\widehat{g}^m_{x,l}$ and $\widehat{g}^{m,*}_{x,l}$ are SH coefficients of the predicted FOD and the target FOD, respectively. ACC is the average ACC_x over all valid voxels within each subject. We additionally report the normalized root mean square error (RMSE) in the SH domain: $\mathrm{RMSE} = \|\widehat{\mathbf{g}} - \widehat{\mathbf{g}}^*\|_2 / (\|\widehat{\mathbf{g}}^*\|_2 + \epsilon)$.

Metrics of Connectomes Following (Zeng et al., 2022), we first compute the absolute connectome disparity between the estimated and reference connectivity matrices. We further summarize the disparity by averaging over all off-diagonal upper-triangular edges, yielding the mean connectome disparity (MCD). For one edge $C^m_s(i,j)$ in the connectome of sample s, the MCD of image generated by method m is defined as:

$$\mathrm{MCD}^m_s = \frac{1}{N_e} \sum_{i<j} |C^m_s(i,j) - C^{GT}_s(i,j)|, \tag{54}$$

where N_e is the number of edges calculated. Following the connectome-level evaluation strategy of (Zeng et al., 2022), statistical testing is performed independently for each undirected connectome edge to determine whether the model-derived connectivity differed significantly from the corresponding ground-truth connectivity. A paired non-parametric permutation test based on sign flipping is firstly employed (Winkler et al., 2014). For one edge, the subject-wise difference is defined as $d^{(s)}_{ij} = C^m_s(i,j) - C^{GT}_s(i,j)$ and the observed statistic for a dataset with N samples is: $T^{obs}_{ij} = \frac{1}{N} \sum_{s=1}^{N} d^{(s)}_{ij}$ and $T^{(b)}_{ij} = \frac{1}{N} \sum_{s=1}^{N} r^{(b)}_s d^{(s)}_{ij}$, where $r^{(b)}_s$ is randomly selected from $\{-1, +1\}$ in b th permutation. The two-sided Monte Carlo permutation p-value is calculated as:

$$p_{ij} = \frac{1 + \sum_{b=1}^{B} \mathbf{1}(|T^{(b)}_{ij}| \geqslant |T^{obs}_{ij}|)}{B + 1}. \tag{55}$$

To account for multiple comparisons across connectome edges, the resulting p-values were corrected using the Benjamini–Hochberg false discovery rate procedure (Benjamini & Hochberg, 1995). Under significance threshold α, an edge was considered significantly different from the ground truth when $p_{ij} < \alpha$.

For each sample, we also compared the thresholded edges which are connectomes elements of one method with a difference larger than a given ratio of GT, i.e., the relative difference. One significant point $C(i,j)$ satisfies:

$$|C_m(i,j) - C_{GT}(i,j)| > \tau |C_{GT}(i,j)|. \tag{56}$$

In our experiments, number of significant points with $\tau = 10\%$, $\tau = 20\%$ and $\tau = 50\%$ are calculated to evaluate the robustness of methods under different tractography complexity.

Evaluation on HCP-YA Dataset During the reconstruction of HCP-YA dataset, Table 6 shows that compared with regression-based baseline FOD-Net, NFO improves the whole-brain, white matter (WM) and WM&CGM ACC to 0.7533, 0.8871 and 0.8163. Besides the mean performance, NFO also achieves the best upper- and lower-quartile (UQ/LQ) statistics for nearly all metrics, suggesting that the improvement is consistently observed across the entire test population rather than being dominated by a small subset of subjects. To further understand how the quantitative improvements in FOD reconstruction translate into fiber-level benefits, we qualitatively examine both local FOD geometry and the resulting tractography.

In FOD images, the density of fibers are shown by the size of ball-stick-model. Both FOD and tractography images represent the principal diffusion directions using direction-encoded colors: red corresponds to left-right (x-axis), green to anterior-posterior (y-axis), and blue to superior-inferior (z-axis). This color-coding scheme facilitates the visualization of anatomical fiber orientations in the brain. Fig. 6 qualitatively compares the reconstructed FODs and corresponding tractography. Consistent with the quantitative results, NFO reconstructs FOD profiles that more closely resemble the MSMT-CSD reference, with better-preserved principal orientations, clearer separation of crossing fiber populations, and fewer spurious angular responses. These improvements are subsequently reflected in tractography, where NFO produces more coherent fiber trajectories and better preserves fiber branches through regions with complex orientations. In contrast, inaccurate or missing FOD peaks in the competing methods can lead to local propagation errors that accumulate along the tracking trajectory, resulting in deviated or missing streamlines. This correspondence between local FOD geometry and downstream fiber trajectories demonstrates that the reconstruction improvements of NFO benefit the fiber level.

Comparison in connectome are shown in Table. 9 and Fig. 5. NFO produces connectivity patterns that more closely agree with the reference connectomes, having the lowest MCD, significant edges and threshold edges with different threshold level, which shows NFO is capable of precise and stable connectivity prediction. Fig. 5 indicates an overall lower absolute error of NFO and the improvement is particularly evident for connections involving nodes 120–160.

These results provide downstream evidence that the benefit of NFO extends beyond voxel-wise FOD accuracy and propagates to large-scale structural connectivity, supporting the ability of the learned fiber operator to preserve anatomically meaningful fiber organization.

Evaluation on MGH Dataset In MGH dataset, Table 7 indicates that the whole-brain ACC increases to 0.6827. Shown in Fig. 7 NFO improves the local FOD configurations and substantially decreases false positive fibers in the reconstructed tractography, whereas larger angular discrepancies in the competing predictions propagate into worse streamline trajectories.

Evaluation on CC Dataset CC dataset contains the largest domain shift among three datasets. The most substantial domain shift is presented by the ex-vivo CC specimens in Fig. 4. Despite being trained exclusively on in-vivo HCP-YA data, NFO improves the mean ACC to 0.8333 while simultaneously reducing the RMSE to 0.0253, which is shown in Table 8. NFO reconstructs FOD configurations that better preserve the complex local fiber organization observed in the reference, which in turn leads to tractography with improved correspondence to the underlying fiber pathways.

B.1.4 Implementation Details

NFO is optimized using the mean squared error between predicted and target SH coefficients. Adam is adopted as the optimizer with an initial learning rate of 1×10^{-3}. Training is performed on HCP-YA using $9 \times 9 \times 9$ patches, 500 epochs and the checkpoint is used for all subsequent experiments. All implementations are based on PyTorch and trained on NVIDIA A5000 GPUs (24 GB). Other learning-based methods follow their official implementation (Zeng et al., 2022; Da Silva et al., 2024).

In the inference stage for all test datasets, FOD-Net and NFO follow the same patch-learning strategy (Zeng et al., 2022) and conduct voxel-wise reconstruction, in which voxels are predicted individually over the entire image volume. FOD-Swin-Net conducts patch-wise sliding-window reconstruction following the official implementation (96^3 sliding-window and 25% overlap) (Da Silva et al., 2024), in which overlapping patches are reconstructed separately and combined by averaging the predictions in overlapping regions.

For the reconstruction of FOD-Net and NFO, the patches are stacked along the first geometric dimension as one batch to reduce the time consumption while keeping the number of model reasoning identical in constructing one image. Besides reconstruction accuracy, we also report the number of trainable parameters and whole-brain inference time for all the learning-based approaches.