Equivariant spatio-hemispherical networks for diffusion MRI deconvolution

Axel Elaldi New York University axel.elaldi@nyu.edu Guido Gerig New York University gerig@nyu.edu Neel Dey MIT CSAIL dey@csail.mit.edu

Abstract

Each voxel in a diffusion MRI (dMRI) image contains a spherical signal corresponding to the direction and strength of water diffusion in the brain. This paper advances the analysis of such spatio-spherical data by developing convolutional network layers that are equivariant to the $\mathbf{E(3)} \times \mathbf{SO(3)}$ group and account for the physical symmetries of dMRI including rotations, translations, and reflections of space alongside voxel-wise rotations. Further, neuronal fibers are typically antipodally symmetric, a fact we leverage to construct highly efficient spatio-hemispherical graph convolutions to accelerate the analysis of high-dimensional dMRI data. In the context of sparse spherical fiber deconvolution to recover white matter microstructure, our proposed equivariant network layers yield substantial performance and efficiency gains, leading to better and more practical resolution of crossing neuronal fibers and fiber tractography. These gains are experimentally consistent across both simulation and in vivo human datasets.

1 Introduction

Instead of scalar intensities, each voxel of a diffusion MR image (dMRI) contains spatio-angular measurements of local water diffusion [64]. As in Fig. 1, this yields *spatio-spherical* images living on $\mathbb{R}^3 \times \mathcal{S}^2$ that are used to map neuronal organization *in vivo* [7, 44] alongside several other biomedical use cases [35, 37, 56, 65, 71]. Despite its potential, deriving neuronal fiber pathways using dMRI is hampered by significant partial voluming in both spatial and spherical domains due to limited clinical scanner resolutions and low SNR [1]. This paper presents a state-of-the-art dMRI deconvolution method to recover neuronal pathways in practical timeframes by developing highly efficient equivariant neural networks that account for dMRI's spatio-spherical data geometry and the voxel-level antipodal symmetry of neuronal fibers, all while demonstrating robustness to clinical dMRI resolutions.

Need for deconvolution. The diffusivity at a voxel reflects the underlying local tissue microstructure [7, 8, 44]. As multiple neuronal fibers can cross within a given resolution-limited voxel, recovering these fibers necessitates solving a blind fiber deconvolution problem at each voxel. Several voxel-level fiber models have been proposed [5, 45, 67, 69, 73, 79] and this work focuses on the widely used fiber Orientation Distribution Function (fODF) model, which represents the fiber configuration at a voxel as a sparse non-negative spherical function [66]. fODFs are the first step of common dMRI applications such as fiber tractography [37]. However, due to noise, subject motion, and clinically viable sampling resolutions (e.g., \leq 30 spherical samples per voxel), fODF recovery is highly ill-posed and requires significant regularization.

Iterative solutions. Constrained spherical deconvolution (CSD) [66] is the workhorse algorithm for dMRI deconvolution, solving the per-voxel fODF deconvolution problem iteratively, subject to non-negativity constraints. Several extensions of CSD further regularize fODFs to be sparse and/or spatially smooth [11, 12, 29, 57, 59, 74]. Further, as fODFs can be assumed to be antipodally symmetric [40, 67], many iterative algorithms optimize only over the hemisphere to accelerate

38th Conference on Neural Information Processing Systems (NeurIPS 2024).




Figure 1: A diffusion MRI (columns 1–3) and a T1w MRI (column 5) derived from a subject in the HCP Young Adult dataset [70]. The inset (column 4) visualizes a region's spatio-spherical diffusion signal $(b-1000mm/s^2)$, highlighting crossing-fiber patterns and the grey/white matter interface.

per-voxel optimization. For example, CSD only optimizes for the even-order spherical harmonics of the fODF and RUMBA [11] uses approximately hemispherical sampling to represent the fODF. Despite significant progress using iterative solutions, these methods still require high angular sampling resolutions that are not clinically feasible to resolve crossing fibers within a voxel.

Deep deconvolution networks. Recent deep learning frameworks for dMRI deconvolution incorporate supervision and/or additional inductive biases to improve results. Supervised approaches train U-Nets to regress the solutions of iterative models or regress ground-truth fODFs estimated from ex vivo animal histology. These methods are consequently upper-bounded by the quality of iterative solutions and the generalizability limitations of small animal datasets. Further extensions incorporate the underlying spherical geometry of the per-voxel estimation by using SO(3)-equivariant network layers but retain the supervised strategy of previous models [62]. ESD [19] instead proposes to use SO(3)-equivariant network layers and *unsupervised* and regularized reconstruction-based deconvolution losses to surpass supervised solutions.

However, the above methods operate entirely at the voxel level and do not model the strong correlation between neighboring fODFs. To this end, RT-ESD [20] proposed $\mathbf{E}(3) \times \mathbf{SO}(3)$ -equivariant network layers for dMRI that are equivariant to spatial rotations, translations, and reflections alongside voxel-wise spherical rotations to achieve high-quality deconvolution due to the joint spatio-spherical modeling. However, its high computational requirements limit its use in clinical or large-scale deployment. For instance, while subject-specific iterative fits with CSD require 3-5 minutes for the whole brain, RT-ESD requires about a day on an A100 GPU to converge when trained on a single subject.

Contributions. We present an efficient equivariant neural network framework for dMRI deconvolution that respects the data geometry of dMRI while ensuring computational practicality. Our spatio-hemispherical deconvolution (SHD) method addresses key redundancies and weaknesses of previous deep networks for dMRI deconvolution in three ways: (1) As fODFs are approximately antipodally symmetric at clinical resolutions, we replace RT-ESD's full spherical sampling with hemispherical sampling and find substantial efficiency gains by reducing the graph Laplacian of the voxel-wise SO(3) convolutions used in RT-ESD; (2) We then exploit the dense structure of dMRI data to further introduce optimized implementations and pre-computations, achieving a cumulative 65% reduction in processing time for an E(3) × SO(3)-equivariant U-Net; (3) Finally, we use explicit smoothness-promoting spatial regularization, leading to further improved fODF recovery.

Experimentally, we achieve state-of-the-art deconvolution results on two widely used simulated dMRI benchmarks with known ground truth. On real *in vivo* human dMRI, our method yields more spatially coherent fODF fields and higher robustness to changes in resolution from research-grade to clinical standards of single-shell low-angular protocols. Lastly, as the achieved efficiency gains enable training on a large set of human datasets, we can now train a single network for amortized inference on new human dMRI, as opposed to the subject-specific optimization of RT-ESD. Our code is available at https://github.com/AxelElaldi/fast-equivariant-deconv.

2 Related work

Deep learning for dMRI. Learned networks operating on dMRI have been extensively used to improve fiber tractography [46, 63], super-resolve diffusion signals or fODFs [48, 58, 78], regress




Figure 2: A deconvolution visualization comparing recovered fiber orientation distribution functions (fODFs) produced by the widely-used iterative CSD [66] model (top row) and our proposed SHD-TV model (bottom row) with high-resolution / clinically-infeasible (left) and low-resolution / clinically-feasible (right) spherical sampling. At high-resolutions (left), SHD-TV demonstrates enhanced localization of fiber orientations, heightened sensitivity to small-angle crossing fibers, and improved spatial consistency in the recovered fibers. At clinical low-resolutions (right), CSD struggles with the loss of input information, whereas our approach exhibits greater robustness to resolution losses and single-shell imaging protocols, yielding higher fidelity and spatially coherent fODFs. Appendix Fig. 9 visualizes comparisons with additional baselines.

microstructral indices in undersampled settings [13, 14, 27, 30, 31, 61], denoise artifacts [22, 23, 76], segment lesions [50], and more. Of these, most employ supervision via training on high-angular resolution research datasets to generalize to low-resolution clinical datasets. Our work is most related to networks that directly operate on $\mathbb{R}^3 \times \mathcal{S}^2$ images either through spherical networks applied voxelwise or approximate parameterizations of the spatio-spherical space [6, 9, 10, 20, 19, 47, 50, 62].

For spherical data, SO(3)-equivariant convolutions are often achieved via spherical harmonics-parameterized convolutions [16, 21, 42] or isotropic convolutions on spherical graphs [53]. These models have been extended to voxel-wise fODF estimation [19, 62] but do not use the spatial correlation between fODFs. [6, 10] develop convolutions for dMRI classification and super-resolution that exhibit voxel-wise SO(3)-equivariance and incorporate manifold-valued spatial averaging. To model spatio-spherical dependence explicitly, SE(3)-equivariance for improved dMRI segmentation has been achieved via tensor field networks [50] and separable kernels [9, 47] at the cost of high-memory usage. In contrast, RT-ESD [20] uses isotropic spatio-spherical kernels and focuses on deconvolution with equivariance to both joint and independent voxel and grid rotations. More recently, PONITA [9] proposes efficient SE(3)-equivariant convolutions for spatio-spherical molecular graphs, but as yet, does not scale to high-angular resolutions needed for dMRI deconvolution.

Spherical deconvolution methods. Constrained spherical deconvolution (CSD) [38, 66] and its sparse [12] and/or spatially regularized extensions [11, 29, 57, 59, 74] are established but rely on lengthy and data-dependent iterative optimization and struggle to resolve small-angle crossing fibers in the low-resolution setting (see Fig. 2). More recently, several trainable models have been proposed for direct fODF estimation including pattern matching methods [24] that use a dictionary to match diffusion signal to known tissue microstructure and deep neural networks [19, 20, 39, 43, 46, 51, 52, 61, 62, 78] that learn a mapping from diffusion signals to fODFs. Trainable models have the advantage of potentially decreasing the need for high-angular resolution dMRI input. ESD [19] further introduced an unsupervised learning framework using only a diffusion signal reconstruction loss and additional fODF regularization, removing the need for ground-truth fODFs for training.

Spatially informed deconvolution. Traditionally, spherical deconvolution is optimized voxel-wise. However, the underlying tissue microstructure has long-range spatial correlations. Iterative methods account for this through spatial regularization such as total-variation [11] or fiber continuity [29, 57, 59, 74], at the cost of increased optimization complexity. Current neural network frameworks

have overlooked explicit fODF spatial regularization and rely solely on high-quality fODF training data and implicit spatial regularization by mean of spatial-weight sharing, either employing grid-wise 3D convolution [20, 46, 78], or channel-wise concatenation of neighboring voxels [19, 61, 62].

3 Methods

3.1 Background

dMRI deconvolution. A dMRI image is a spatio-spherical function $S: \mathbb{R}^3 \times \mathcal{S}^2 \to \mathbb{R}^B$ with B features, called shells in the dMRI literature. The fODF model describes the tissue microstructure as a non-negative spatio-spherical function F, providing information on the local tissue composition and orientations, and links the dMRI and the fODFs $S = \mathcal{C}(F)$. dMRI deconvolution is interested in recovering the fODF from the dMRI by minimizing a constrained reconstruction loss $F^* = \operatorname{argmin}_{F \geq 0} ||S - \mathcal{C}(F)||_2^2$. We focus our study on learnable spatio-spherical deconvolution operators trained to reconstruct high-angular resolution fODFs from sparse dMRI measurements $\mathcal{N}_{\theta}(S) = F$.

Spatio-spherical convolutions. $\mathbf{E}(\mathbf{3}) \times \mathbf{SO}(\mathbf{3})$ is the group of independent $\mathbf{E}(\mathbf{3})$ and $\mathbf{SO}(\mathbf{3})$ transformations, where $\mathbf{E}(\mathbf{3})$ acts on the spatial grid \mathbb{R}^3 and $\mathbf{SO}(\mathbf{3})$ acts on the voxel sphere \mathcal{S}^2 . Let $\psi_\theta: \mathbb{R}^3 \times \mathcal{S}^2 \to \mathbb{R}$ be a learnable filter, such that convolving f and ψ_θ yields $f_{\mathrm{out}}(T,R) = \int_{y \in \mathbb{R}^3} \int_{p \in \mathcal{S}^2} \psi_\theta(T^{-1}y,R^{-1}p)f(y,p)\,dp\,dy$ where $(T,R) \in \mathbf{E}(\mathbf{3}) \times \mathbf{SO}(\mathbf{3})$. RT-ESD [20] implements these convolutions by only considering isotropic filters to limit computational complexity. In addition, the filter ψ_θ is expressed as a separable kernel $\psi_\theta(z_1,z_2) = \phi_{\theta_1}(z_1)\phi_{\theta_2}(z_2)$, turning the convolution into a two-step spherical-spatial convolution:

$$f_{\text{out}}(x,q) = \int_{y \in \mathbb{R}^3} \phi_{\theta_1}(||x - y||) \int_{p \in \mathcal{S}^2} \phi_{\theta_2}(qp^T) f(y,p) \, dp \, dy. \tag{1}$$

We build on RT-ESD [20], presented in Fig. 12, which uses DeepSphere [53]'s graph filtering based $\mathbf{SO(3)}$ -equivariant spherical convolution. At every spatial location $y \in \mathbb{R}^3$, f(y,.) is discretized on a set of spherical vertices $\mathcal{V} = \{q_i \in \mathcal{S}^2\}_{i \in [1,..,N_{\mathcal{V}}]}$, $\mathbf{f}(y) \in \mathbb{R}^{N_{\mathcal{V}}}$. The graph $\mathcal{G} = (\mathcal{V}, \mathbf{A})$ and its normalized Laplacian \mathbf{L} are constructed from the set \mathcal{V} , where $\mathbf{A} \in \mathbb{R}^{N_{\mathcal{V}} \times N_{\mathcal{V}}}$ is the spherical adjacency matrix. The voxel-wise spherical filtering output is then used to construct a spatio-spherical graph $\hat{\mathbf{f}}(y) = [T^0(\mathbf{L})\mathbf{f}(y),...,T^{K-1}(\mathbf{L})\mathbf{f}(y)] \in \mathbb{R}^{N_{\mathcal{V}} \times K}$ with K features, where K is the polynomial degree of the filtering and $T^k(\mathbf{L})$ is the Chebyshev polynomial of degree k. A learnable isotropic spatial filter $\alpha_k : \mathbb{R}^+ \to \mathbb{R}$ is then defined for each spherical filtered map $\hat{\mathbf{f}}_k$, resulting in the overall $\mathbf{E}(\mathbf{3}) \times \mathbf{SO(3})$ -convolutional operation expressed as:

$$\mathbf{f}_{out}(x) = \sum_{y \in \mathcal{P}(x)} \sum_{k=0}^{K-1} \alpha_k(||x - y||) T^k(\mathbf{L}) \mathbf{f}(y), \tag{2}$$


where $\mathcal{P}(x)$ is a spatial neighborhood around the position $x \in \mathbb{R}^3$. Spherical $\mathbf{SO(3)}$ -equivariance is achieved by using a uniform and symmetric \mathcal{V} sampling such as the HEALPix [32] grid, providing a uniform coverage of the spherical domain at different levels of resolution, and graph weight $p, q \in \mathcal{V}$, $\mathbf{A}(p,q) = e^{-||p-q||/\sigma}$ with σ 's value selected to minimize empirical equivariance error.

3.2 Efficient dMRI spatio-hemispherical convolutions

Spatio-Hemispherical Equivariant convolution. fODFs and dMRIs are approximately antipodally symmetric, which makes fully spherical convolutions redundant. We therefore improve the time and space efficiency of the RT-ESD convolution by reducing the spherical graph $\mathcal G$ to a hemispherical graph $\mathcal H$. Because f is antipodal symmetric, we assume, without loss of generality, that $\mathcal V$ is a symmetric sampling, i.e. $\forall p \in \mathcal V$, $\exists (-p) \in \mathcal V$. We construct the hemispherical sampling $\mathcal V^+ = \{p \in \mathcal V \text{ and } p_z \geq 0\}$ from the spherical sampling $\mathcal V$, with p_z chosen to be the 3rd spatial coordinate of p. We compute the hemispherical Laplacian weight as,

$$p, q \in \mathcal{V}^+, \ \mathbf{L}^+(p, q) = \mathbf{L}(p, q) + \mathbf{L}(p, -q).$$
 (3)

Thus, the spherical operation $\mathbf{L}f$ is entirely explained by the hemisphere graph Laplacian \mathbf{L}^+f^+ , i.e. $\forall p \in \mathcal{V}, \exists (p^+) \in \mathcal{V}^+$ s.t. $(\mathbf{L}^+f)(p^+) = (\mathbf{L}f)(p)$, where f^+ is the f sampling on \mathcal{V}^+ (proof






Figure 3: Contribution overview. A. We reduce the spherical graph $(\mathcal{G}, \mathbf{L})$ to an hemispherical graph $(\mathcal{H}, \mathbf{L}^+)$. B. The SHD deconvolution framework operates on a grid of spherical signals and reduces computation complexity while improving neuronal fiber deconvolution.

in appendix A.1). For spherical samplings \mathcal{V} such as HEALPix, our proposed sampling reduces the number of vertices sampled to $\frac{1}{2}|\mathcal{V}|$ and the hemispherical Laplacian size to $\frac{1}{4}|\mathcal{V}|^2$. Our reduction is both theoretically and empirically equivariant to $\mathbf{E}(3) \times \mathbf{SO}(3)$, see appendix A.2 for more details.

Dense Matrix Multiplication. DeepSphere's [53] used sparse matrix multiplication to compute $T^k(\mathbf{L})\mathbf{f}(x)$. However, for dMRI, our Laplacians are dense and sparse matrix multiplication adds significant computational overhead on dense matrices. To address this, we substitute the sparse matrix multiplication used in [53] with standard matrix multiplications and find substantial efficiency improvements when applied to dMRI data.

Pre-computed Chebyshev Polynomials. DeepSphere [53] assumes the spherical sampling \mathcal{V} , and thus the Laplacian \mathbf{L} , to be dependent on the input spherical signal \mathbf{f} . Thus, for efficiency, the Chebyshev polynomials are computed iteratively as $T^{k+1}(\mathbf{L})\mathbf{f}(x) = (2\mathbf{L}T^k(\mathbf{L}) - T^{k+1}(\mathbf{L}))\mathbf{f}(x)$ for each new \mathbf{f} and \mathbf{L} . In our setting, we use the same spherical sampling \mathcal{V} for every spherical signal \mathbf{f} , making the Laplacian \mathbf{L} of every convolutional layer fixed. We therefore precompute and store these polynomials before training to further eliminate redundancy.

3.3 Spatial Hemispherical Deconvolution (SHD) Network

Fig.3 presents an overview of our proposed spatial-hemispherical deconvolution (SHD) framework. Below, we first describe how our inputs are preprocessed and then detail our network and the losses used to train it.

Data normalization. The dMRI signal acquisition yields a raw spatio-spherical signal $S_{\mathcal{G}}$ sampled on a set of spherical coordinates \mathcal{G} , called shell sampling. The signal intensity range and shell sampling are scan-dependent and need to be normalized before being fed to our proposed neural network. Following [55], we normalize the dMRI scans such that the white matter B0-diffusion signal is comparable between scans. Then, following [19], the normalized data is interpolated onto a hemispherical HEALPix sampling \mathcal{V}^+ . We provide details in appendix C.2.

 $\mathbf{E}(3) \times \mathbf{SO}(3)$ -Unet. The interpolated diffusion signal is then deconvolved to recover the spatio-spherical fODFs $\mathbf{F}_{\mathcal{V}^+}$ on the same input HEALPix sampling \mathcal{V}^+ . We use a U-Net, with 4 depth levels, with layers acting on a spatial grid of hemispherical samples. We use our proposed layers in the U-Net architecture presented in Fig. 3. Besides the last, every block comprises a convolution, batch-norm, and a ReLU activation. The last block consists of a convolution and a Softplus activation function. The up/downsampling layers first involve mean spatial upsampling/pooling on \mathbb{R}^3 , followed by mean spherical upsampling/pooling on \mathcal{S}^2 , both introducing minor numerical equivariance error. Following DeepSphere [53], to minimize the equivariance error from the spherical upsampling/pooling, we use the hierarchical structure of the HEALPix grid.

Network training. Our proposed deconvolution network \mathcal{N}_{θ} takes as input the shell-sampled dMRI signal $\mathbf{S}_{\mathcal{G}}$ to produce fODFs $\mathbf{F}_{\mathcal{V}^+}^{\theta} = \mathcal{N}_{\theta}(\mathbf{S}_{\mathcal{G}})$. We train the network in an unsupervised manner by reconstructing the dMRI signal $\mathbf{S}_{\mathcal{D}}^{\theta} = \mathcal{C}(\mathbf{F}_{\mathcal{V}^+}^{\theta})$ from the estimated fODFs, with details on the




Figure 4: **Efficiency analysis.** Runtime (**A** & **C**) and GPU memory usage (**B** & **D**) expressed as the percentage of the baseline [20], for both: (**line plots**) a convolutional layer applied to increasing angular resolution samplings and (**bar plots**) a U-Net applied to high-angular resolution. The proposed convolution is more efficient than existing equivariant spatio-spherical convolutions.

reconstruction function C in appendix C.2. We note that the input and reconstruction shell samplings G and D need not be the same, allowing for angular super-resolution reconstruction during training if required. The parameters θ are estimated by minimizing a reconstruction loss subject to non-negativity [19, 66] and sparsity-promoting regularization [12, 19]:

$$\mathcal{L} = ||\mathbf{S}_{\mathcal{D}} - \mathbf{S}_{\mathcal{D}}^{\theta}||_{2}^{2} + \lambda_{\text{nn.}}||\mathbf{A}.\mathbf{F}_{\mathcal{V}^{+}}^{\theta}||_{2}^{2} + \lambda_{\text{spa.}}||\log\left(1 + \frac{\mathbf{F}_{\mathcal{V}^{+}}^{\theta^{-2}}}{\sigma^{2}}\right)||_{2}^{2} + \lambda_{\text{tv}}\mathcal{L}_{\text{tv}},\tag{4}$$

where the terms correspond to reconstruction, non-negativity, sparsity, and smoothness, respectively. We define **A** as a vector setting any positive component of $\mathbf{F}_{\mathcal{V}^+}^{\theta}$ to 0. We set $\sigma=10^{-5}$ following [12]. \mathcal{L}_{TV} is described below.

Spatial regularization. To add *explicit* spatial regularization, we propose to extend the spatial total-variation regularization from [11] to the trainable model framework $\mathcal{L}_{tv} = ||\nabla_x \mathbf{F}_{\mathcal{V}^+}^{\theta}||_2^2$ where ∇_x is the spatial gradient operator. The spatial regularizer \mathcal{L}_{tv} seeks to promote smoother reconstructions in ill-posed settings, especially at clinical angular resolutions used later in our experiments.


4 Experiments

We first quantify the runtime and memory efficiency gains produced by our contributions. We then analyze their use across a diversity of dMRI deconvolution settings on both real and synthetic datasets.

4.1 Runtime and memory efficiency improvements

Experimental details. We first compare the efficiency of our proposed convolution against two spatiospherical equivariant convolutions: RT-ESD [20] and PONITA [9], as measured by inference runtime and memory consumption. We then conduct an ablation analysis of our three proposed improvements. All analyses are conducted for both a single spatio-spherical layer applied to increasing spherical graph resolution and also for an entire spatio-spherical U-Net architecture detailed in Fig. 3. As input, we use a random $8\times8\times8\times V$ spatio-spherical volume, with V depending on the spherical HEALPix grid resolution. We fix the U-Net input HEALPix resolution to $8 \ (V = 754)$ and we vary the single layer HEALPix resolution from $1 \ (V = 12)$ to 8. We take the RT-ESD convolutions as a reference baseline and present runtime and memory usage results as a percentage of the baseline averaged across 50 inference steps. All comparisons are performed using an Nvidia RTX-8000 GPU.

[A] Dataset I and II: fODF estimation Generalization to unseen volume



[B] Dataset III: fODF and Tractography Unsupervised fine-tuning on a single volume




Figure 5: Overview of the diffusion MRI experiments in Section 4.2. [A] We perform superresolved fODF estimation experiments on two datasets, DiSCo and HCP, respectively. Here, we study the impact of using either high-angular or low-angular resolution as input. [B] We perform quantitative fODF and tractography estimation experiments on Tractometer. We extract fODFs and tractograms from the dMRI with both input and output having low-angular resolution.

Results. Fig. 4 presents the results of the efficiency analysis. For a single layer, our convolutions are 2 to 5 times faster than the RT-ESD baseline while decreasing memory footprint by 2.5 times on large spherical sampling. PONITA has a similar runtime and memory footprint as SHD on sparse spherical sampling but underperforms on large spherical graphs (that are common in dMRI deconvolution [12]), being 10 times slower and more memory intensive. Consequently, a U-Net designed for high-angular resolution sampling is 2.5 and 3.5 faster using our proposed convolutions over RT-ESD and PONITA, while using 4 and 20 times less memory, respectively. Our ablations indicate that using our hemispherical convolutions over the spherical convolutions, as well as using precomputed Laplacian polynomials, brings the most improvement for both runtime and memory on large spherical sampling while leveraging dense matrix multiplication reduces runtime on low-angular resolutions.

4.2 Diffusion MRI deconvolution experiments

Experimental details. Validation of fODF estimation methods is confounded by the absence of known and unbiased ground truth fODFs in *in vivo* human dMRI. To address this limitation, the literature commonly benchmarks methods on synthetic datasets with known ground truth tissue microstructure. We conduct both a quantitative (synthetic dataset) and qualitative (*in vivo* human dMRI) analysis of our deconvolution framework (Section 4.2.1). We then perform a quantitative analysis on the tractography downstream task on a benchmark synthetic dataset (Section 4.2.2). For quantitative evaluation, we first estimate fODFs, then detect fiber peaks, and report the peak angular error and false positive rate (FPR) following [17]. The angular error is the average angle between the ground truth peaks and their closest predicted peaks and the FPR is the number of predicted peaks that do not match any ground-truth peaks divided by the number of ground-truth peaks. See Appendix A.3 for more details. All deep network experiments were performed using a single RTX8000 GPU and used less than 16GB of system RAM.

Baselines. We compare the proposed SHD method against conventional iterative methods including CSD [66], RUMBA and RUMBA-TV [11], and state-of-the-art learning based methods leveraging different convolutional layer, such as the non-equivariant voxel-wise MLP in [52] and the volume-wise CNN in [46], that we also extend to have spatial regularization (CNN-TV). We also compare against equivariant voxel-wise spherical convolutions ESD [61, 19], and more recent equivariant spatio-spherical convolutions RT-ESD [20] and PONITA [9]. To focus on the inductive bias introduced by each method's convolution layer, we adopt a standardized network architecture and training framework, presented in section 3.3. We provide details about these different baselines in appendix C.3.

Data. We leverage three public dMRI datasets for validation. DiSCo [54] is a synthetic dMRI dataset with three volumes, split into training, validation, and testing volumes, with high-angular resolution sampling (4 shells each with 90 gradients) provided with ground truth fODF. We then use 100 in-vivo unrelated dMRI scans from the HCP young adult dataset [70], split into 65 training, 15 validation, and 20 testing volumes, with high-angular resolution (3 shells, each with 90 gradients), without any ground truth. As we aim to benchmark performance on both high-angular resolution data and more commonly acquired low-angular resolution data, we further simulate low-angular




Figure 6: **DiSCo fiber detection performances** on high (**left col.**) and low (**right col.**) angular resolutions. We first present fODF estimation results on high-angular [A] and low-angular resolution [B] input (**closer to bottom-left** is better). [C-D] then present a qualitative example of a two-crossing fiber estimation. Our faster implementation SHD does not negatively impact results in comparison to RT-ESD, while our improved model SHD-TV outperforms other methods by providing higher angular precision and less spurious fibers, especially at clinically-viable low-angular resolution.

resolution acquisitions on DiSCo and HCP by randomly selecting 29 gradients from only the lowest shell. The last dataset is Tractometer [49], a single-volume simulation of a real human brain with a low-angular resolution protocol (1 shell, 32 gradients), provided with ground truth tractography alongside a tractography scoring algorithm [60].

4.2.1 Super-resolved fODF estimation: synthetic DiSCo and real human HCP datasets

We perform two experiments, depicted in Fig.5A). First, we measure performance when all baselines are trained and tested on high-angular resolution data. Second, we give all methods high-angular resolution data only during training and test them on low-angular resolution data. In this setting, the networks are trained to regress the high-angular resolution data from low-angular resolution input. For quantitative evaluation, we extract fiber directions from the estimated fODF and we report the peak angular error and false positive rate as described in the previous section. All results are averaged over five random training seeds. Finally, as HCP lacks fODF and tractogram ground truth, we perform a qualitative analysis. Further validation details are in appendix A.3 and A.4.

Results. Fig. 6 presents DiSCo deconvolution results for both the low and high-angular resolution settings. We visualize qualitative results on a random HCP test subject in Fig.2, focusing on two areas with diverse tissue microstructures, such as crossing fibers, fiber bending, and multi-tissue compartments. Our proposed efficiency contributions retain all of the performance of previous methods as our proposed SHD model (without spatial regularization) achieves similar results to the much slower and more memory-intensive RT-ESD and PONITA methods. Moreover, non-equivariant methods, while having competitive angular errors, are negatively impacted by a high false positive rate and qualitatively high number of spurious fibers.

Finally, on research-grade high-angular resolution data, incorporating spatial information in the network input does not improve results. However, adding total-variation regularization on the output of the model as in our proposed SHD-TV model shows significantly improved results. Further, in a clinical low-angular setting, our spatially informed models (SHD and SHD-TV) widely outperform voxel-wise methods, showcasing the importance of contextual information for clinical dMRI.

4.2.2 Unsupervised fODF estimation and tractography w/ known ground truth: Tractometer

A high-level overview of this experiment is described in Fig.5B). As Tractometer consists of a single volume with a low-angular resolution protocol with no held-out testing data, we benchmark methods on unsupervised fODF estimation similarly to previous work [19]. We first train the models in an




Figure 7: Tractometer fODF estimation and tractography performance. Top: Unsupervised fODF estimation (A, closer to bottom left is better) and tractography (B, closer to top right is better) results. Bottom: In [C], we visualize ground-truth and estimated fibers projecting out from the brainstem into the right hemisphere. Overall, SHD and SHD-TV demonstrate more faithful fiber and streamline recovery as compared to the voxel-wise RUMBA and ESD methods. In particular, SHD-TV yields fewer invalid streamlines and increases spatial coherence.

unsupervised manner to estimate fODFs and then use them to investigate the effect of improved local fODF estimation on white matter tractography accuracy. Validation of the unsupervised estimated fODFs involves using the same local fiber detection evaluation as for the DiSCo experiments. Subsequently, the validation of the estimated tractographs is conducted by comparing them against 21 ground truth fiber bundles provided by the dataset [60]. The white matter streamlines of the tractograph are initially segmented into 21 target fiber bundles. The correctly identified bundles are counted as Valid Bundles and the streamlines belonging to one of the ground truth bundles are counted as Valid Streamlines. Further, the volume of each estimated bundle is compared against the volume of its corresponding ground truth bundle, leading to the computation of the Overlap score (analogous to a True Positive value) and the Overreach Score (similar to a False Positive value).

Results. We present Tractometer results in Fig.7. In this setting where all methods are trained and tested on the same volume, non-equivariant methods (CNN and CNN-TV) outperform most unregularized equivariant deconvolution methods. We speculate that when generalization is not required, non-equivariant models enjoy a higher overfitting capability in this problem, thereby diminishing the advantage of employing inductive priors in neural networks.

However, the proposed SHD-TV, which incorporates both spatio-spherical inductive biases and spatial regularization, significantly outperforms all other methods, regardless of whether they are equivariant or spatially regularized. This finding supports the compounding benefits of both our architectural modifications and our proposed regularization strategy. In turn, these enhanced fODF recovery performances positively impact tractography accuracy, reflected in a quantitatively higher number of valid streamlines and bundle overlap (albeit with an increase in bundle overreach) and more qualitatively coherent streamlines.

5 Discussion

Limitations and future work. Our work has certain limitations. First, while we focus on the key dMRI task of fiber deconvolution, our proposed layers are task-agnostic and can also be deployed to other tasks such as denoising and segmentation in future studies. Second, while our layers can be used in many fields with spatio-spherical data outside of neuroimaging, such as robotics [72], neural rendering [26], and molecular dynamics [9], our assumption of antipodal symmetry (widespread in dMRI) may need to be relaxed for many applications within these fields. Third, angularly

52103

undersampled reconstruction results have an inherent risk of missing fibers or fiber hallucination. While our work partially mitigates these unrealistic reconstructions by adding geometric and spatial priors to the reconstruction, future validation studies could analyze cohorts with paired scans on clinical low-resolution and research-grade high-resolution dMRI protocols. Finally, our methods need further testing on additional clinical challenges such as subject motion, imaging artifacts, and abnormal structures. However, as our work is entirely self-supervised and only based on regularized reconstruction of input data, we anticipate robustness to such imaging variation. Preliminary results on a pathological brain sample are provided in appendix A.5.

Conclusions. The analysis of *in vivo* white matter neuronal organization within the brain depends on the key task of dMRI deconvolution and current methods are either not accurate or too slow. To that end, we developed highly time and memory-efficient equivariant convolutional network layers that respect the physical symmetries of dMRI. Our contributions led to $3.5\times$ faster runtimes and up to $20\times$ lower memory consumption as compared to existing spatio-spherical layers, while exceeding or maintaining their performance on dMRI studies. These layers were then used to construct networks that achieved strong dMRI deconvolution performance in practical timeframes across multiple commonly used synthetic benchmark datasets and real *in vivo* human datasets. Finally, our methods are robust to the dMRI angular resolutions typically used in clinical practice outside of research settings, potentially enabling robust analyses of large-scale retrospective dMRI studies and direct application to the clinic.

Acknowledgements

The authors gratefully acknowledge support from NIH grants 1R01HD088125-01A1, R01HD055741, 1R01MH118362-01, 2R01EB021391-06 A1, R01ES032294, R01MH122447, and 1R34DA050287.

References

- [1] Andrew L Alexander, Khader M Hasan, Mariana Lazar, Jay S Tsuruda, and Dennis L Parker. Analysis of partial volume effects in diffusion-tensor mri. *Magnetic Resonance in Medicine*, 45(5):770–780, 2001.
- [2] Jesper LR Andersson, Stefan Skare, and John Ashburner. How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. *Neuroimage*, 20(2):870–888, 2003.
- [3] Jesper LR Andersson and Stamatios N Sotiropoulos. Non-parametric representation and prediction of single-and multi-shell diffusion-weighted mri data using gaussian processes. *Neuroimage*, 122:166–176, 2015.
- [4] Jesper LR Andersson and Stamatios N Sotiropoulos. An integrated approach to correction for off-resonance effects and subject movement in diffusion mr imaging. *Neuroimage*, 125:1063– 1078, 2016.
- [5] Yaniv Assaf and Peter J Basser. Composite hindered and restricted model of diffusion (charmed) mr imaging of the human brain. *Neuroimage*, 27(1):48–58, 2005.
- [6] Monami Banerjee, Rudrasis Chakraborty, Jose Bouza, and Baba C Vemuri. Volterranet: A higher order convolutionalnetwork with group equivariance forhomogeneous manifolds. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [7] Peter J Basser, James Mattiello, and Denis LeBihan. Estimation of the effective self-diffusion tensor from the nmr spin echo. *Journal of Magnetic Resonance, Series B*, 103(3):247–254, 1994.
- [8] Peter J Basser and Carlo Pierpaoli. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor mri. *Journal of magnetic resonance*, 213(2):560– 570, 2011.
- [9] Erik J Bekkers, Sharvaree Vadgama, Rob D Hesselink, Putri A van der Linden, and David W Romero. Fast, expressive se (n) equivariant networks through weight-sharing in position-orientation space. arXiv preprint arXiv:2310.02970, 2023.

- [10] Jose J Bouza, Chun-Hao Yang, David Vaillancourt, and Baba C Vemuri. A higher order manifold-valued convolutional neural network with applications to diffusion mri processing. In *International Conference on Information Processing in Medical Imaging*, pages 304–317. Springer, 2021.
- [11] Erick J Canales-Rodríguez, Alessandro Daducci, Stamatios N Sotiropoulos, Emmanuel Caruyer, Santiago Aja-Fernández, Joaquim Radua, Jesús M Yurramendi Mendizabal, Yasser Iturria-Medina, Lester Melie-García, Yasser Alemán-Gómez, et al. Spherical deconvolution of multi-channel diffusion mri data with non-gaussian noise models and spatial regularization. *PloS one*, 10(10):e0138910, 2015.
- [12] Erick Jorge Canales-Rodríguez, Jon Haitz Legarreta, Marco Pizzolato, Gaëtan Rensonnet, Gabriel Girard, et al. Sparse wars: A survey and comparative study of spherical deconvolution algorithms for diffusion mri. *NeuroImage*, 184:140–160, 2019.
- [13] Geng Chen, Bin Dong, Yong Zhang, Weili Lin, Dinggang Shen, and Pew-Thian Yap. Xq-sr: joint xq space super-resolution with application to infant diffusion mri. *Medical image analysis*, 57:44–55, 2019.
- [14] Geng Chen, Yoonmi Hong, Yongqin Zhang, Jaeil Kim, Khoi Minh Huynh, Jiquan Ma, Weili Lin, Dinggang Shen, Pew-Thian Yap, and UNC/UMN Baby Connectome Project Consortium. Estimating tissue microstructure with undersampled diffusion data via graph convolutional neural networks. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 280–290. Springer, 2020.
- [15] Taco S Cohen, Mario Geiger, Jonas Köhler, and Max Welling. Spherical cnns. arXiv preprint arXiv:1801.10130, 2018.
- [16] Benjamin Coors, Alexandru Paul Condurache, and Andreas Geiger. Spherenet: Learning spherical representations for detection and classification in omnidirectional images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 518–533, 2018.
- [17] Alessandro Daducci, Erick Jorge Canales-Rodriguez, Maxime Descoteaux, Eleftherios Garyfallidis, Yaniv Gur, et al. Quantitative comparison of reconstruction methods for intra-voxel fiber recovery from diffusion mri. *IEEE transactions on medical imaging*, 2013.
- [18] Thijs Dhollander, Remika Mito, David Raffelt, and Alan Connelly. Improved white matter response function estimation for 3-tissue constrained spherical deconvolution. In *Proc. Intl. Soc. Mag. Reson. Med*, volume 555, 2019.
- [19] Axel Elaldi, Neel Dey, Heejong Kim, and Guido Gerig. Equivariant spherical deconvolution: Learning sparse orientation distribution functions from spherical data. In *International Conference on Information Processing in Medical Imaging*, pages 267–278. Springer, 2021.
- [20] Axel Elaldi, Guido Gerig, and Neel Dey. E (3) x so (3)-equivariant networks for spherical deconvolution in diffusion mri. In *Medical Imaging with Deep Learning*, 2023.
- [21] Carlos Esteves, Christine Allen-Blanchette, Ameesh Makadia, and Kostas Daniilidis. Learning so (3) equivariant representations with spherical cnns. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 52–68, 2018.
- [22] Shreyas Fadnavis, Joshua Batson, and Eleftherios Garyfallidis. Patch2self: Denoising diffusion mri with self-supervised learning. Advances in Neural Information Processing Systems, 33:16293–16303, 2020.
- [23] Shreyas Fadnavis, Agniva Chowdhury, Joshua Batson, Petros Drineas, and Eleftherios Garyfallidis. Patch2self2: Self-supervised denoising on coresets via matrix sketching. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 27641–27651, 2024.
- [24] Patryk Filipiak, Timothy Shepherd, Ying-Chia Lin, Dimitris G Placantonakis, Fernando E Boada, and Steven H Baete. Performance of orientation distribution function-fingerprinting with a biophysical multicompartment diffusion model. *Magnetic Resonance in Medicine*, 88(1):418–435, 2022.

- [25] Bruce Fischl. Freesurfer. *Neuroimage*, 62(2):774–781, 2012.
- [26] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5501–5510, 2022.
- [27] Eric K Gibbons, Kyler K Hodgson, Akshay S Chaudhari, Lorie G Richards, Jennifer J Majersik, Ganesh Adluru, and Edward VR DiBella. Simultaneous noddi and gfa parameter map generation from subsampled q-space imaging using deep learning. *Magnetic resonance in medicine*, 81(4):2399–2411, 2019.
- [28] 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.
- [29] Alvina Goh, Christophe Lenglet, Paul M Thompson, and René Vidal. Estimating orientation distribution functions with probability density constraints and spatial regularity. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 877–885. Springer, 2009.
- [30] Vladimir Golkov, Alexey Dosovitskiy, Jonathan I Sperl, Marion I Menzel, Michael Czisch, Philipp Sämann, Thomas Brox, and Daniel Cremers. Q-space deep learning: twelve-fold shorter and model-free diffusion mri scans. *IEEE transactions on medical imaging*, 35(5):1344–1351, 2016.
- [31] Tobias Goodwin-Allcock, Jason McEwen, Robert Gray, Parashkev Nachev, and Hui Zhang. How can spherical cnns benefit ml-based diffusion mri parameter estimation? In *International Workshop on Computational Diffusion MRI*, pages 101–112. Springer, 2022.
- [32] Krzysztof M Gorski, Benjamin D Wandelt, Frode K Hansen, Eric Hivon, and Anthony J Banday. The healpix primer. *arXiv preprint astro-ph/9905275*, 1999.
- [33] Nate Gruver, Marc Finzi, Micah Goldblum, and Andrew Gordon Wilson. The lie derivative for measuring learned equivariance. *arXiv preprint arXiv:2210.02984*, 2022.
- [34] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.
- [35] Thierry AGM Huisman, Lee H Schwamm, Pamela W Schaefer, Walter J Koroshetz, Neetha Shetty-Alva, Yelda Ozsunar, Ona Wu, and A Gregory Sorensen. Diffusion tensor imaging as potential biomarker of white matter injury in diffuse axonal injury. *American Journal of Neuroradiology*, 25(3):370–376, 2004.
- [36] Mark Jenkinson, Christian F Beckmann, Timothy EJ Behrens, Mark W Woolrich, and Stephen M Smith. Fsl. Neuroimage, 62(2):782–790, 2012.
- [37] Ben Jeurissen, Maxime Descoteaux, Susumu Mori, and Alexander Leemans. Diffusion mri fiber tractography of the brain. *NMR in Biomedicine*, 32(4):e3785, 2019.
- [38] 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.
- [39] 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.
- [40] Suheyla Cetin Karayumak, Evren Özarslan, and Gozde Unal. Asymmetric orientation distribution functions (aodfs) revealing intravoxel geometry in diffusion mri. *Magnetic Resonance Imaging*, 49:145–158, 2018.

- [41] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint* arXiv:1412.6980, 2014.
- [42] Risi Kondor, Zhen Lin, and Shubhendu Trivedi. Clebsch–gordan nets: a fully fourier space spherical convolutional neural network. *Advances in Neural Information Processing Systems*, 31, 2018.
- [43] Simon Koppers and Dorit Merhof. Direct estimation of fiber orientations using deep learning in diffusion imaging. In *International Workshop on Machine Learning in Medical Imaging*, pages 53–60. Springer, 2016.
- [44] Denis Le Bihan, Eric Breton, Denis Lallemand, Philippe Grenier, Emmanuel Cabanis, and Maurice Laval-Jeantet. Mr imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders. *Radiology*, 161(2):401–407, 1986.
- [45] Denis Le Bihan, Jean-François Mangin, Cyril Poupon, Chris A Clark, Sabina Pappata, Nicolas Molko, and Hughes Chabriat. Diffusion tensor imaging: concepts and applications. *Journal of Magnetic Resonance Imaging*, 13(4):534–546, 2001.
- [46] Zhichao Lin, Ting Gong, Kewen Wang, Zhiwei Li, Hongjian He, Qiqi Tong, Feng Yu, and Jianhui Zhong. Fast learning of fiber orientation distribution function for mr tractography using convolutional neural network. *Medical physics*, 46(7):3101–3116, 2019.
- [47] Renfei Liu, Francois Bernard Lauze, Erik J Bekkers, Kenny Erleben, and Sune Darkner. Group convolutional neural networks for DWI segmentation. In *Geometric Deep Learning in Medical Image Analysis*, 2022.
- [48] Matthew Lyon, Paul Armitage, and Mauricio A Álvarez. Spatio-angular convolutions for super-resolution in diffusion mri. Advances in Neural Information Processing Systems, 36, 2024.
- [49] Klaus H Maier-Hein, Peter F Neher, Jean-Christophe Houde, Marc-Alexandre Côté, Eleftherios Garyfallidis, et al. The challenge of mapping the human connectome based on diffusion tractography. *Nature communications*, 8(1):1–13, 2017.
- [50] Philip Müller, Vladimir Golkov, Valentina Tomassini, and Daniel Cremers. Rotation-equivariant deep learning for diffusion mri. *arXiv preprint arXiv:2102.06942*, 2021.
- [51] Vishwesh Nath, Sudhir K Pathak, Kurt G Schilling, Walter Schneider, and Bennett A Landman. Deep learning estimation of multi-tissue constrained spherical deconvolution with limited single shell dw-mri. In *Medical Imaging 2020: Image Processing*, 2020.
- [52] Vishwesh Nath, Kurt G Schilling, Prasanna Parvathaneni, Colin B Hansen, Allison E Hainline, et al. Deep learning reveals untapped information for local white-matter fiber reconstruction in diffusion-weighted mri. *Magnetic resonance imaging*, 62:220–227, 2019.
- [53] Nathanaël Perraudin, Michaël Defferrard, Tomasz Kacprzak, and Raphael Sgier. Deepsphere: Efficient spherical convolutional neural network with healpix sampling for cosmological applications. *Astronomy and Computing*, 27:130–146, 2019.
- [54] Jonathan Rafael-Patino, Gabriel Girard, Raphaël Truffet, Marco Pizzolato, Emmanuel Caruyer, and Jean-Philippe Thiran. The diffusion-simulated connectivity (disco) dataset. *Data in Brief*, 38:107429, 2021.
- [55] David Raffelt, J-Donald Tournier, Stephen Rose, Gerard R Ridgway, Robert Henderson, Stuart Crozier, Olivier Salvado, and Alan Connelly. Apparent fibre density: a novel measure for the analysis of diffusion-weighted magnetic resonance images. *Neuroimage*, 59(4):3976–3994, 2012.
- [56] David A Raffelt, J-Donald Tournier, Robert E Smith, David N Vaughan, Graeme Jackson, Gerard R Ridgway, and Alan Connelly. Investigating white matter fibre density and morphology using fixel-based analysis. *Neuroimage*, 144:58–73, 2017.

- [57] Alonso Ramirez-Manzanares, Mariano Rivera, Baba C Vemuri, Paul Carney, and Thomas Mareci. Diffusion basis functions decomposition for estimating white matter intravoxel fiber geometry. *IEEE transactions on medical imaging*, 26(8):1091–1102, 2007.
- [58] Amelie Rauland and Dorit Merhof. Using synthetic training data in neural networks for the estimation of fiber orientation distribution functions from single shell data. In 2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI), pages 1–5. IEEE, 2023.
- [59] Marco Reisert and Valerij G Kiselev. Fiber continuity: an anisotropic prior for odf estimation. *IEEE transactions on medical imaging*, 30(6):1274–1283, 2011.
- [60] Emmanuelle Renauld, Antoine Théberge, Laurent Petit, Jean-Christophe Houde, and Maxime Descoteaux. Validate your white matter tractography algorithms with a reappraised ismrm 2015 tractography challenge scoring system. 2023.
- [61] Sara Sedlar, Abib Alimi, Théodore Papadopoulo, Rachid Deriche, and Samuel Deslauriers-Gauthier. A spherical convolutional neural network for white matter structure imaging via dmri. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 529–539. Springer, 2021.
- [62] 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, MICCAI Workshop*, 2020.
- [63] Fabian Leander Sinzinger and Rodrigo Moreno. Reinforcement learning based tractography with so (3) equivariant agents. In *Geometric Deep Learning in Medical Image Analysis (Extended abstracts)*, 2022.
- [64] Edward O Stejskal and John E Tanner. Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient. *The journal of chemical physics*, 42(1):288–292, 1965.
- [65] Lei Tang and Xiaohong Joe Zhou. Diffusion mri of cancer: From low to high b-values. *Journal of Magnetic Resonance Imaging*, 49(1):23–40, 2019.
- [66] 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.
- [67] 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.
- [68] J-Donald Tournier, Robert Smith, David Raffelt, Rami Tabbara, Thijs Dhollander, et al. Mrtrix3: A fast, flexible and open software framework for medical image processing and visualisation. *NeuroImage*, 202:116137, 2019.
- [69] David S Tuch. Q-ball imaging. Magnetic Resonance in Medicine, 52(6):1358–1372, 2004.
- [70] 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.
- [71] Steven Warach, D Chien, W Li, M Ronthal, and RR Edelman. Fast magnetic resonance diffusion-weighted imaging of acute human stroke. *Neurology*, 42(9):1717–1717, 1992.
- [72] Michael Watterson, Sikang Liu, Ke Sun, Trey Smith, and Vijay Kumar. Trajectory optimization on manifolds with applications to so (3) and r3xs2. In *Robotics: Science and Systems*, volume 20, 2018.
- [73] VJ Wedeen, TG Reese, DS Tuch, MR Weigel, JG Dou, RM Weiskoff, and D Chessler. Mapping fiber orientation spectra in cerebral white matter with fourier-transform diffusion mri. In *Proceedings of the 8th Annual Meeting of ISMRM*, 2000.

- [74] Ye Wu, Yoonmi Hong, Sahar Ahmad, Wei-Tang Chang, Weili Lin, Dinggang Shen, and Pew-Thian Yap. Globally optimized super-resolution of diffusion mri data via fiber continuity. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part VII 23*, pages 260–269. Springer, 2020.
- [75] Ye Wu, Xiaoming Liu, Yunzhi Huang, Tao Zhou, and Fan Zhang. An open relaxation-diffusion mri dataset in neurosurgical studies. *Scientific Data*, 11(1):177, 2024.
- [76] Tiange Xiang, Mahmut Yurt, Ali B Syed, Kawin Setsompop, and Akshay Chaudhari. DDM\$^2\$: Self-supervised diffusion MRI denoising with generative diffusion models. In *The Eleventh International Conference on Learning Representations*, 2023.
- [77] Fang-Cheng Yeh, Andrei Irimia, Dhiego Chaves de Almeida Bastos, and Alexandra J Golby. Tractography methods and findings in brain tumors and traumatic brain injury. *Neuroimage*, 245:118651, 2021.
- [78] 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.
- [79] Hui Zhang, Torben Schneider, Claudia A Wheeler-Kingshott, and Daniel C Alexander. Noddi: practical in vivo neurite orientation dispersion and density imaging of the human brain. *Neuroimage*, 61(4):1000–1016, 2012.

A Additional results

A.1 Spherical and hemisphere convolution

Hemispherical sampling. Consider the symmetric spherical sampling $\mathcal{V}=\{p\in\mathbb{S}^2\}$ such that $\forall p\in\mathcal{V}, \exists (-p)\in\mathcal{V}$, and the Euclidean coordinate $p=(p_x,p_y,p_z)\in\mathcal{S}^2$. The sphere can be divided into a north and south hemisphere $\mathcal{V}=\mathcal{V}^+\cup\mathcal{V}^-$, with $\mathcal{V}^+=\{p\in\mathcal{V},$ such that $p_z>0)\}$ and $\mathcal{V}^-=\{p\in\mathcal{V},$ such that $p_z<0)\}$. For point p lying on the circle $p_z=0$, we further add to \mathcal{V}^+ every point such that $p_y>0$ and to \mathcal{V}^- every point such that $p_y<0$. Finally, for the two points p such that $p_z=p_y=0$, we add p such that $p_x>0$ to \mathcal{V}^+ and p such that $p_x<0$ to \mathcal{V}^- .

Hemispherical restriction. Consider a spherical function $f: \mathcal{S}^2 \to \mathbb{R}$ and a function $\mathbf{L}: \mathcal{S}^2 \times \mathcal{S}^2 \to \mathbb{R}$, both sampled on \mathcal{V} such that $f \in \mathbb{R}^V$ and $\mathbf{L} \in \mathbb{R}^{V \times V}$. Consider the function $\mathbf{L}^+: \mathcal{S}^2 \times \mathcal{S}^2 \to \mathbb{R}$, $\mathbf{L}^+(p,q) = \mathbf{L}(p,q) + \mathbf{L}(p,-q)$, sampled on \mathcal{V}^+ such that $\mathbf{L}^+ \in \mathbb{R}^{V/2 \times V/2}$ and $f^+ \in \mathbb{R}^{V/2}$ the sampling of f on \mathcal{V}^+ .

Theorem 1 Consider an antipodally symmetric spherical function f and \mathbf{L} , such that for all $p, q \in \mathcal{V}$, f(p) = f(-p) and $\mathbf{L}(p,q) = \mathbf{L}(-p,-q)$, then $\forall p \in \mathcal{V}$, $\exists p^+ \in \mathcal{V}^+$, such that $(\mathbf{L}f)(p) = (\mathbf{L}^+f^+)(p^+)$.

Proof Let $p \in \mathcal{V}$. Then, define $p^+ = p$ if $p \in \mathcal{V}^+$, -p otherwise. Thus, we know that $p^+ \in \mathcal{V}^+$.

$$(\mathbf{L}f)(p) = \sum_{q \in \mathcal{V}} \mathbf{L}(p, q) f(q)$$
 (5)

$$= \sum_{q \in \mathcal{V}^{+}} \mathbf{L}(p,q) f(q) + \sum_{q \in \mathcal{V}^{-}} \mathbf{L}(p,q) f(q)$$
 (6)

$$= \sum_{q \in \mathcal{V}^+} \mathbf{L}(p, q) f(q) + \mathbf{L}(p, -q) f(-q)$$
(7)

$$= \sum_{q \in \mathcal{V}^+} (\mathbf{L}(p,q) + \mathbf{L}(p,-q)) f(q)$$
(8)

$$= \sum_{q \in \mathcal{V}^{+}} (\mathbf{L}(p^{+}, q) + \mathbf{L}(p^{+}, -q)) f(q)$$
(9)

$$= \sum_{q \in \mathcal{V}^+} \mathbf{L}(p^+, q) f(q) \tag{10}$$

$$= (\mathbf{L}^+ f^+)(p^+) \tag{11}$$

A.2 Numerical equivariance error analysis

Motivation. Due to discretization and implementation practicalities, equivariant networks are subject to aliasing and are only approximately equivariant [33]. Prior work in this space [53, 19, 20] has demonstrated low equivariance errors for SO(3) and $E(3) \times SO(3)$ convolutions and their U-Net counterparts when operating on full spherical sampling. We claim that the proposed accelerated implementations of these convolutions do not introduce additional equivariance error and provide empirical evidence below.

Model details. The hemispherical $E(3) \times SO(3)$ convolution is compared to a grid-wise E(3)-equivariant convolution with spherical information, E(3)-SH, and a voxel-wise SO(3)-equivariant convolution with spatial information, Concat-SO(3). The E(3)-SH convolution is a 3D convolution with isotropic kernels processing spherical harmonic coefficients as input features [52]. The Concat-SO(3) convolution is a spherical convolution processing spatial neighbors as input features [62, 20]. The spatial and spherical kernels are initialized randomly following [34].

Evaluation. We measure numerical equivariance error as the deviation from the equivariance equality, using:

$$\nabla \operatorname{Equiv}(\mathcal{N}, G, f) = \frac{||[G\mathcal{N}(f)] - \mathcal{N}([Gf])||_2^2}{||\mathcal{N}(f)||_2^2}$$
(13)

where \mathcal{N} is the tested operator, f is a spatio-spherical signal, and $G=(T,R)\in\mathbf{E}(3)\times\mathbf{SO}(3)$ is a composition of grid-wise and voxel-wise transformation. In practice we only test rotation-equivariance, and we limit T and R to be rotations. This quantity is measured over 1000 randomly sampled spatio-spherical signals f, random rotations (T,R), and randomly initialized \mathcal{N} . f is randomly generated from a uniform $\mathcal{U}[0,1]$ distribution on a $8\times 8\times 8\times V$ spatio-spherical volume with HEALPix hemispherical grid of resolution 8 (V=384 vectors). Moreover, we filter high spherical frequency by only keeping the first 8 spherical harmonic degrees.

Results. Fig.8 illustrates estimated equivariance errors. As expected, we find that the E(3)-SH convolution exhibits a lack of equivariance to voxel-wise rotations, and the Concat-SO(3) convolution exhibits a lack of equivariance to grid-wise rotations. The accelerated implementation of the $E(3) \times SO(3)$ convolution, taking into account both inductive bias of the spatial and spherical domain, maintains low voxel-wise rotation equivariance errors.




Figure 8: Quantitative evaluation of equivariance error, depending on the convolution equivariance group, and the applied rotation group.

A.3 DiSCo quantitative results

Experiment details. We provide more details on the DiSCo dataset, which we used to evaluate both fODF estimation performance and robustness to dMRI protocol angular resolution. The DiSCo dataset comprises three $40 \times 40 \times 40$ synthetic volumes with variable amounts of synthetic noise added. Our experiments use the volumes with a noise level of SNR=30. All three volumes share the same synthetic generation process and have four B0 images and four shells each with 90 gradients. We do not apply any pre-processing step. For quantitative analysis, we train the different models on the first volume, validate on the second, and test on the third. Given that the synthetic generation process of the DiSCo dataset is similar to a white matter/CSF tissue simulation, we limit the different models to a 2 tissue decomposition, and we analyze the white-matter fODFs.

Validation details. All results for deep network models are averaged over five random seeds. For quantitative evaluation, we first estimate fODFs, then detect peaks, and finally threshold them based on their voxel-wise relative fiber partial volume applying 19 different threshold values ranging from 0.05 to 0.95. For each threshold, we calculate the false positive rate (FP), false negative rate (FN), and angular precision of detected fibers following the evaluation framework proposed by [17]; utilizing a rejection cone of 25° . Precision and recall are computed at different thresholds using FP, FN, and true positive fibers (TP). Subsequently, we calculate the Precision-Recall F1 score. We report the angular error and false positive rate for the threshold that maximizes the F1 score on the DiSCo validation volume. We report raw average scores and 95% confidence interval (CI) average over five random seeds in Table 1. Notice that all scores have a CI < 0.01, but the angular error.

A.4 Qualitative in vivo human results

Experiment details. We provide more details on the HCP dataset, which we used for qualitative evaluation of the fODF estimation. The HCP young adult release contains 100 unrelated subjects and

Table 1: DiSCo fiber detection performances @ noise level SNR=30 on high and low-resolution data. Models requiring training are trained on the first volume, validated on the second, and tested on the third. Results average over 5 random initialized models. Confidence interval at 95% given if CI is greater than 0.01.

Model	PR AUC (†)		F1 score (†)		Angle @F1 (↓)		FNR @F1 (↓)		FPR @F1 (↓)	
Protocol resolution	High	Low	High	Low	High	Low	High	Low	High	Low
CSD	0.64	0.50	0.56	0.45	23.0	29.8	0.65	0.74	0.10	0.18
RUMBA	0.61	0.54	0.49	0.42	27.1	33.5	0.74	0.80	0.05	0.08
RUMBA-TV	0.49	0.55	0.49	0.48	26.4	27.5	0.71	0.74	0.12	0.09
MLP	0.59	0.48	0.53	0.42	23.8 ± 0.1	31.8 ± 0.2	0.68	0.78	0.12	0.15
CNN	0.56	0.51	0.51	0.47	24.9 ± 0.3	27.4 ± 0.3	0.70	0.74	0.11	0.12
S ² U-net	0.63	0.50	0.58	0.42	22.2 ± 0.1	32.4 ± 0.1	0.62	0.79	0.11	0.13
ESD	0.64	0.47	0.61	0.47	20.4 ± 0.2	27.9 ± 0.8	0.56	0.68	0.17	0.32
Concat-ESD	0.65	0.59	0.61	0.56	20.3 ± 0.1	22.6 ± 0.1	0.56	0.61	0.15	0.20
RT-ESD	0.64	0.60	0.60	0.57	20.4 ± 0.3	22.0 ± 0.2	0.58	0.61	0.14	0.16
SHD (ours)	0.64	0.61	0.60	0.57	20.0 ± 0.1	21.5 ± 0.5	0.57	0.60	0.17	0.19
SHD-TV (ours)	0.70	0.62	0.65	0.58	17.5 ± 0.2	20.7 ± 0.2	0.49	0.57	0.19	0.22

ensures the exclusion of twin subjects to prevent any data leakage between dataset splits. The dMRI volumes had undergone pre-processing using FSL [36] and FreeSurfer [25], and further details on the preprocessing pipeline can be found in the HCP dataset documentation [28, 2, 3, 4]. The 100 dMRI volumes adhered to a high-resolution diffusion protocol, featuring 18 B0 volumes per subject and three diffusion shells, each comprising 90 diffusion gradients. We apply similar downsampling to the HCP dataset, randomly selecting 29 diffusion gradients from the lower b-value shell (1000, s/mm²) while retaining the B0 volumes. For qualitative analysis on the HCP dataset, we train on 65 randomly selected subjects, monitor training on 15 validation subjects, and test on 20 subjects. In contrast with the DiSCo dataset, we estimated 3 group-wise tissue response functions, accounting for white matter, gray matter, and cerebrospinal fluid present in in-vivo human brains.

Validation details. In-vivo datasets lack microstructure ground truth. Some approaches [62] suggest using the CSD estimated fODFs from high-quality and high-resolution dMRI as a surrogate ground truth. However, we argue that this introduces a bias towards CSD and hinders the assessment of performance improvements in new methodologies. A more recent alternative proposed by [19, 20] involves using partial volume estimation (PVE) from fODFs as a validation metric. Nevertheless, the PVE scalar is rotation invariant, posing limitations on its ability to discern the impact of rotation-equivariant methods compared to their non-equivariant counterparts. Non-equivariant models may demonstrate high accuracy in invariant scalar estimation while failing to capture the underlying geometric structure adequately. Moreover, PVE validation relies heavily on the accuracy of third-party tissue segmentation on T1/T2 images, which may not consistently correlate with partial volume estimates derived from more information-rich but low-resolution dMRI data. While we acknowledge that a comprehensive solution for quantitative analysis of in-vivo fODF estimation is yet to be established, we advocate for a visual qualitative analysis of our proposed method. This approach aims to provide insights into the performance of our model without relying on potentially biased or limited validation metrics.




Figure 9: Qualitative illustration comparing the proposed equivariant dMRI deconvolution framework against all other baselines.

A.5 Preliminary analysis on an abnormal brain

Here, we demonstrate the preliminary applicability of our method on abnormal brains. In Fig. 10, we estimate fODFs on a brain with a glioma using the recently released dMRI dataset from [75]. We find that the additional priors from our method help fODF estimation and fiber tracking substantially, filling in a hole in the fiber tracks recovered using the baseline CSD method. However, we note that this analysis is highly preliminary and that tractography on brains with lesions is a highly active area of research that requires substantial modifications to tractography algorithms [77].




Figure 10: **fODF** and tractography estimation in a glioma-affected brain (gray arrow). In [A], we compare the conventional CSD method ([A.1]) with our proposed SHD-TV model ([A.2]). Our approach retrieves more spatially coherent fODFs with better fiber angular separation in voxels containing crossing fibers. Notably, the CSD method does not detect fODFs in the crossing area indicated by the red arrow, leading to an inadequate representation of microstructures by the tractography algorithm. Additionally, our model does not reveal any abnormal fODFs near or within the tumorous tissue.

B Equivariance vs. Data Augmentation: $\mathbb{R}^3 \times S^2$ MNIST

Motivation. We study the advantages of using spatio-spherical inductive biases for convolutional layers, as opposed to learning them from the data. We create a synthetic spatio-spherical MNIST dataset (inspired by the spherical MNIST experiments of [15]) to conduct controlled experiments.

Dataset. The $\mathbb{R}^3 \times \mathcal{S}^2$ MNIST data generation process is presented in Fig.11. Importantly, the dataset characterizes a spatio-spherical segmentation task, where voxel-wise information alone is insufficient for correct voxel classification. We construct two versions of this dataset to test generalization from either equivariance or data augmentation. The first version lacks any form of data augmentation, whereas the second version incorporates random grid and voxel-wise rotations. Because on-the-fly data augmentation of spatio-spherical signals is computationally expensive, we pre-generate a total of 1000 volumes per dataset, with a training-validation-test split of 716/142/142.

Dataset generation. We randomly position eight non-overlapping 4×4 squares on a 2D 16×16 slice. Subsequently, each square is assigned a random digit between 1 and 9, designating 0 as the background digit. This 2D slice is duplicated along the z-axis, yielding the final 3D segmentation volume. Notably, this volume comprises eight non-overlapping $4\times 4\times 16$ structures, termed tubes, aligned along the z-axis, each featuring the same classification digit across its voxels. Optionally, we apply a random grid rotation to the segmentation volume. In the second step, we randomly select an MNIST image corresponding to the classification digit assigned to each square. These digit images are projected onto a sphere using the methodology presented in [15]. We utilize HEALPix spherical sampling with a resolution of 4, equivalent to a spherical resolution of V=192 vertices. Given the straightforward nature of voxelwise spherical digit classification, we reduce the amount of information at each voxel. Instead of projecting the entire MNIST image onto the sphere, we randomly crop it to one-quarter of its original size before projecting the cropped digit onto the sphere. As a result, any segmentation network has to rely on both the spatial neighborhood and the spherical information within a voxel. Optionally, we apply a random voxel rotation to the spherically projected digit.




Figure 11: Top row: Spatio-Spherical MNIST generation process. Bottom row: Results from testing generalization from equivariance vs. data augmentation on the synthetic $\mathbb{R}^3 \times \mathcal{S}^2$ MNIST classification task. [A-C] Dataset generation process. [D-E] Segmentation dice score of models trained without (D) and with (E) data augmentation. Overall, our proposed $\mathbf{E}(3) \times \mathbf{SO}(3)$ convolution has a higher generalization power to unseen transformation than its non-equivariant counterparts, and the right inductive bias increases segmentation performance against data augmented-only models.

Model details. We compare four convolutional architectures, each exhibiting different equivariance properties while sharing the same U-Net structure presented in Fig.3. We compare the three previously presented spatio-spherical $\mathbf{E(3)} \times \mathbf{SO(3)}$, $\mathbf{E(3)}$ -SH, and Concat-SO(3) U-Nets, and, to highlight the importance of incorporating spatial information, we compare them with a voxel-wise $\mathbf{SO(3)}$ U-Net. Inputs to the U-Nets consist of $16 \times 16 \times 192$ $\mathbb{R}^3 \times \mathcal{S}^2$ MNIST volumes, and the outputs,

post-Softmax activation, are $16 \times 16 \times 16 \times 10$ volumes representing the classification probabilities of each voxel for the 10 digits.

Training details. Our training regimen uses a batch size of 8 and trains for 50 epochs with an initial learning rate of 3×10^{-3} , which is halved after the 25, 35, and 45 epochs. The optimized objective function encompasses a combined Dice and Cross-Entropy loss with class-dependent weight to address imbalanced labels.

Results. We first assess how equivariance influences the generalization of models to unseen data transformations. The four models are first trained on the dataset without data augmentation and subsequently tested on the dataset incorporating out-of-distribution random rotations. We then train and test the models on the dataset augmented with random rotations. The results are presented in Fig.11.D) and E). To quantify segmentation performance, we compute the dice score across every volume in the testing set.

In Fig.11.D), we observe that all four models, regardless of the embedded equivariance, perform well when trained and tested on a transformation-free (i.e., in-distribution) dataset. Importantly, the spatially informed networks outperform the voxel-wise $\mathbf{SO}(3)$ network, underscoring the significance of incorporating spatial context into the neural network. Furthermore, we emphasize the importance of incorporating inductive bias into the neural network to enhance generalization to unseen data transformations during training. The $\mathbf{E}(3) \times \mathbf{SO}(3)$ model, trained on the rotation-free dataset, achieves a relatively better dice score of 0.38 on the dataset with unseen rotations than both other models with only equivariance to $\mathbf{E}(3)$ or $\mathbf{SO}(3)$.

Fig.11.E) further illustrates the advantage of inductive bias compared to data augmentation during training. Data augmentation dramatically enhances the performance of the $E(3) \times SO(3)$ model on randomly rotated data, surpassing its non-equivariant counterparts. Meanwhile, data augmentation applied to E(3) or SO(3) models lacking spherical or spatial inductive bias improves segmentation performance but falls short of matching the performance of the $E(3) \times SO(3)$ equivariant model.

C Additional experimental details

C.1 Spatio-Hemispherical Convolution overview




Figure 12: Overview of the Hemispherical $E(3) \times SO(3)$ Convolution computation. This figure is adapted from [20]. [A-C] The input of the convolution is a 3D grid of hemispherical graphs with V vertices per voxel. The input is first processed by voxel-wise spherical filtering using the proposed hemispherical Laplacian and efficient implementation. [D-F] The 3D volume is then processed by a 3D isotropic convolution with weight-sharing across the hemispherical graph vertex.

C.2 Details of the deconvolution framework

fODF model. The fODF model assumes the voxel-wise diffusion signal to be produced by a composition of T tissue types, with each tissue further composed of an unknown number of cells/fibers, each one described by its local partial volume and orientation. At the voxel level, the fiber information is aggregated into one fiber orientation distribution function (fODF) per tissue, providing information on the local tissue composition and orientations. The model relates the dMRI signal and tissue microstructure by combining T tissue-specific Response Functions (RFs) and T fiber Orientation Distribution Functions (fODFs). The fODFs $F=\{F_t\}$ are tissue-dependent antipodally symmetric spatio-spherical functions. The response functions $RF=\{RF_t^b\}$ are tissue-dependent and shell-dependent spherical functions, assumed constant across the spatial domain, and to be antipodal and axial symmetric. The total b-shell dMRI signal is modeled as a sum over tissue-specific signals $S^b = \sum_{t=1}^T S_t^b$, where tissue-specific diffusion signals are modeled as a voxel-wise spherical convolution $S_t^b = \mathcal{C}(F_t, RF_t^b)$ between the fODF and the response function. The convolution is defined as $\mathcal{C}(F, RF)(x,q) = \int_{p \in \mathcal{S}^2} F(x,p)RF(q^Tp)\mathrm{d}p$.

Diffusion sampling. A dMRI image S is a function $S(x,q):\mathbb{R}^3\times\mathcal{S}^2\to\mathbb{R}^B$. In practice, the dMRI signal is discretized on a set of shells $\mathcal{B}=\{b_i\in\mathbb{R}^+\}_{i\in[1,..,B]}$ where B is the number of shells, such that for every $b\in\mathcal{B}$, the b-shell function S^b is a spatio-spherical signal. Furthermore, the b-shell function is sampled on a shell-dependent set of spherical gradient direction $\mathcal{V}_b=\{q_i^b\in\mathcal{S}^2\}_{i\in[1,..,N_b]}$ with N_b the number of b-shell gradient sampling. We note $\mathbf{S}_{\mathbf{b},\mathcal{V}_\mathbf{b}}(x)=[S(x,q_1^b,b),...,S(x,q_{N_b}^b,b)]\in\mathbb{R}^{N_b}$ the sampled b-shell signal on \mathcal{V}_b . Moreover, we note $\mathcal{D}=\{\mathcal{B},\{\mathcal{V}_b\}_{b\in\mathcal{B}}\}$ the full diffusion sampling set, and $\mathbf{S}_{\mathcal{D}}(x)=[\mathbf{S}_{\mathbf{b}_1,\mathcal{V}_{\mathbf{b}_1}}(x),...,\mathbf{S}_{\mathbf{b}_{\mathbf{B}},\mathcal{V}_{\mathbf{b}_{\mathbf{B}}}}(x)]\in\mathbb{R}^{N_{\mathcal{D}}}$ the sampled dMRI signal, with $N_{\mathcal{D}}=\sum_{b\in\mathcal{B}}N_b$ the total number of gradient sampling.

Data normalization. For every new dMRI scan $S^{ori.}$, we compute the scan-specific white matter B0 response function $RF_{wm}^{0,ori.}$ using MRtrix [68], and normalize the dMRI signal $S = S^{ori.}/RF_{wm}^{0,ori.}$. The input of our framework is the normalized shell-sampled diffusion MRI signal $\mathbf{S}_{\mathcal{G}}(x) = [\mathbf{S_{b_1}}, \nu_{\mathbf{b_1}}(x), ..., \mathbf{S_{b_B}}, \nu_{\mathbf{b_B}}(x)] \in \mathbb{R}^{N_{\mathcal{D}}}$ for $x \in \mathbb{R}^3$. Following [19], we first normalize the protocoldependent spherical sampling by first interpolating it, for each diffusion shell, to a fixed hemisphere HEALPix sampling \mathcal{V}^+ as defined in section 3.2. We implement the interpolation using spherical harmonics. For this, we first compute the spherical harmonic coefficients $\hat{\mathbf{S}_{\mathbf{b}}}(x) \in \mathbb{R}^{N_L}$ for every

 $b \in \mathcal{B}$ shell, where N_L is the number of harmonic coefficient for a maximal harmonic bandwidth of L. Because the diffusion signal is antipodal-symmetric, the odd-order spherical harmonic coefficients are equal to 0 and we only compute the even-order coefficients, making the total number of coefficient $N_L = (L/2+1)(L+1)$. Importantly, getting the coefficients $\hat{\mathbf{S}}_{\mathbf{b}}(x)$ from $\mathbf{S}_{\mathbf{b},\mathcal{V}_{\mathbf{b}}}(x)$ requires at least $|\mathcal{V}_b| \geq (L/2+1)(L+1)$ samples. We limit the maximum interpolation bandwidth to L=8, which requires at least 45 samples per shell. In case the shell sampling \mathcal{V}_b does not have enough samples, we use the theoretically maximum spherical harmonic degree estimated from the available gradients. We then interpolate the harmonic coefficients on the HEALPix grid \mathcal{V}^+ , $\mathbf{S}_{\mathbf{b},\mathcal{V}^+}(x) = \hat{\mathbf{S}}_{\mathbf{b}}(x)\mathbf{Y}_{\mathcal{V}^+}^L$, where $\mathbf{Y}_{\mathcal{V}^+}^L$ relates the even L bandwidth harmonic coefficients to the HEALPix spherical sampling. We get the B feature input spatio-spherical signal $\mathbf{S}_{\mathcal{V}^+}(x) \in \mathbb{R}^{V \times B}$.

Group-average response function. We estimate a population-based response function $\hat{\mathbf{RF}}$ following [55]. For N_s normalized dMRI scans in our training dataset, we first compute T tissue response functions $\hat{\mathbf{RF}}_{t,i}$ per dMRI image, where i is the image index, using the Dhollander algorithm, implemented in MRITIN [68] [18]. We then compute the per-tissue average response function $\hat{\mathbf{RF}}_t = 1/N_s \sum_{i=1}^{N_s} \hat{\mathbf{RF}}_{t,i}$.

Signal reconstruction. From the fODFs $\mathbf{F}_{\mathcal{V}^+}$, we reconstruct the diffusion signal $\mathbf{S}_{\mathcal{D}}$ on the reconstruction shell-sampling \mathcal{D} using a spherical convolution with response functions $\mathbf{R}\mathbf{F}$. For every voxel $x \in \mathbf{R}^3$, we compute the spherical harmonic coefficients $\hat{\mathbf{F}}(x) = \{f_{l,t}^m(x)\} \in \mathbb{R}^{N_L \times T}$ from the estimated $\mathbf{F}_{\mathcal{V}^+}(x)$. The fODFs being antipodal-symmetric, we only compute the even-degree coefficients. The maximum spherical harmonic degree L depends on the resolution of the HEALPix sampling \mathcal{V}^+ . We use a sampling resolution with 384 vertices per hemisphere and a fODF maximum spherical harmonic degree of L=18. Moreover, the response functions are also represented by their spherical harmonic coefficients $\hat{\mathbf{R}}\hat{\mathbf{F}}=\{r_{l,t}^{m,b}\}\in\mathbb{R}^{M_L\times T\times B}$. The RFs are antipodal and z-axis symmetric, setting to 0 every non-even degree and non-zero order coefficients. We then compute the reconstruction dMRI spherical harmonic coefficients $\hat{\mathbf{S}}(x)=\{s_l^{m,b}(x)\}\in\mathbb{R}^{N_L\times B}$, $s_l^{m,b}(x)=\sum_t \sqrt{2\pi/(2l+1)}f_{l,t}^m(x)r_{l,t}^{0,b}$. Finally, we can interpolate the reconstructed signal $\hat{\mathbf{S}}$ on any shell-sampling \mathcal{D} .

C.3 Details of compared models

We give an overview of details of the compared methods in Fig.13 and separate the different baseline methodologies into two groups.

Conventional Models. Conventional methods for fODF estimation generally solve the inverse problem iteratively. Constrained Spherical Deconvolution [66, 68] (CSD) and RUMBA [11] optimize voxel-wise dMRI reconstruction losses subject to a non-negativity constraint on the fODF. RUMBA-TV further enhances fODF smoothness by incorporating spatial total variation regularization. Importantly, conventional methods do not need prior training and are limited by the input data quality.

Deep-Learning Models. We evaluate our proposed framework against existing deep-learning methodologies for fODF estimation. While diverse network architectures, such as MLP, 2D/3D CNN, or spherical CNN, have been proposed, we adopt a standardized network architecture to focus our comparison on the inductive bias and training framework introduced by each method's convolution layer. We employ an MLP-based network that operates voxel-wise [52] and a CNN-based network operating on spatial patches [46]. Adding spherical inductive biases to the MLP network, we employ a voxel-wise S^2 -based network ESD [61, 19]. The SE(3)-equivariance P0NITA method adds spatial information and inductive bias, equivariant to joint transformation on the spatio-spherical domain. Finally, the RT-ESD model is a $E(3) \times SO(3)$ -equivariance spatio-spherical convolution network, equivariant to independent transformation on the spatial and spherical domain. We extend RT-ESD to an efficient implementation SHD and spatially regularized SHD-TV.

C.4 Details of architecture and training

Network Details. A high-level network architecture is illustrated in Fig. 3. For $E(3) \times SO(3)$ -equivariant models, we use spatio-spherical pooling and unpooling operations. The SO(3)-equivariant models employ spherical pooling/unpooling, the CNN models use only spatial pool-

ing/unpooling, and no pooling is used by the MLP model. The first layer maps the input features to $F_s = 32$ features for equivariant models and $F_s = 128$ features for non-equivariant models.

The input and output feature number of the first and last layers depend on the convolution inductive bias, the number of estimated tissue components T, the number of input diffusion shells B, and the diffusion gradients angular resolution. The MLP and CNN models have at each voxel a feature vector of size $F_i = B \times 45$ (high angular resolution) or $F_i = B \times 28$ (low angular resolution) consisting of the 45 or 28 spherical harmonic coefficients of degree 8 or 6, respectively. At each voxel, they output a feature vector of size $F_o = T \times 45$, representing the 8-degree spherical harmonic coefficients of the T tissue fODF. The S² U-Net, ESD, and RT-ESD models use an input spherical graph of 768 vertices (HEALPix grid resolution of 8) with $F_i = B$ spherical feature maps. The Concat-ESD model uses the same graph resolution but the number of feature maps increases to $F_i = B \times 3^3$ due to neighborhood concatenation. All four models output $F_o = T$ spherical maps (with the same resolution), corresponding to the T tissue fODFs.




Figure 13: **Model description.** We compare conventional and learnable methods. Learnable methods have increasing embedded geometric prior through equivariance.

Training implementation. All models are trained for 50 epochs with a batch size of 16 patches using the Adam optimizer [41] with an initial learning rate set at 1.7×10^{-2} . The learning rate is decayed by a factor of ten after the 30, 40, and 45 epochs. For unsupervised models, we tuned the regularization weights using the ESD model on the DiSCo validation volume, and further tune λ_{tv} for SHD: $\lambda_{nn} = 10^{-1}$, $\lambda_{sparse} = 5 \times 10^{-5}$, and $\lambda_{tv} = 5 \times 10^{-1}$. The spatially-informed models operate on $3 \times 3 \times 3$ spatial patches, whereas the spherical-only models process $1 \times 1 \times 1$ spatial patches.

fODF ground truth generation. Quantitative validation of the model requires access to ground truth fODF. The two benchmark datasets either do not directly provide this information or provide reference fODFs estimated with the CSD model on noise-free high-angular resolution dMRI volumes. However, both datasets provide ground-truth white matter streamlines. To avoid bias towards the CSD model and extract ground-truth fODF for the Tractometer dataset, we propose an alternative unbiased fODF estimation approach leveraging the ground-truth tractograph. We approximate the voxel-wise ground-truth fODFs by aggregating every streamline passing through a voxel and computing its spherical fiber density function. A white matter streamline can be represented as a set of vectors $\{v_j^i\}$ where $v \in \mathbf{S}^2$ is the local streamline direction, j is the streamline index and i is the vector index. For every voxel p, we find every local direction $\{v_{jk}^{ik}\}$ going through the voxel. We then apply a spherical kernel density estimation on the set of unit vectors $\{v_{jk}^{ik}\}$ using a uniform spherical kernel with an angular size of 15° . We then extract the ground-truth peak directions from the fiber density function using the MRtrix peak detection algorithm [68] with a maximum number of crossing fibers per voxel set to 10.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We provide empirical speed and memory efficiency in section 4.1, as well as improved deconvolution results in section 4.2.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discuss limitations and future work in Section 5.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We provide theoretical assumptions and proofs in Section 3.2 and Appendix A.1.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and crossreferenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide extensive technical details of the proposed convolution and deconvolution framework, with pre-processing and post-processing steps. We also provide the code of the proposed spatio-hemispherical convolution, as well as the deconvolution architecture, benchmark, training, and testing scripts in supplementary material, with provided hyperparameters. The DiSCo and HCP datasets are publicly available.

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The code, training, and testing scripts are provided available at https://github.com/AxelElaldi/fast-equivariant-deconv/ and all of our experimental datasets are publicly available.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be
 possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not
 including code, unless this is central to the contribution (e.g., for a new open-source
 benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how
 to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Training and testing details can be found in the main text and in appendix, as well as in the supplementary materials.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental
 material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Statistical significance given as 95% confidence interval in Appendix A.3.

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide memory and speed analysis of our method, as well as GPU and RAM resources.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: We reviewed the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a
 deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Our work pertains to the task of deconvolution of spatio-spherical images in neuroimaging. It has potential use in clinical workflows outside of research settings and this is discussed in Section 5.

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The released model does not pose a high risk of misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with
 necessary safeguards to allow for controlled use of the model, for example by requiring
 that users adhere to usage guidelines or restrictions to access the model or implementing
 safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Dataset, pre and post-processing algorithms, as well as benchmarked models, have been cited according to their licenses.

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.

- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the
 package should be provided. For popular datasets, paperswithcode.com/datasets
 has curated licenses for some datasets. Their licensing guide can help determine the
 license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: No new assets introduced.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: We retrospectively analyze publicly available datasets.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: All analyses were performed retrospectively on publicly available datasets that were IRB-approved at their original institutions.

Guidelines:

• The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

52125

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.