PRISM: Differentiable Analysis-by-Synthesis for Fixel Recovery in Diffusion MRI

Mohamed Abouagour(✉), Atharva Jaydeep Shah, and Eleftherios Garyfallidis

Indiana University Bloomington, Bloomington, IN 47405, USA
moabouag@iu.edu
(✉) Corresponding author

Abstract. Diffusion MRI microstructure fitting is nonconvex and often performed voxelwise, which limits fiber peak recovery in narrow crossings. This work introduces PRISM, a differentiable analysis-by-synthesis framework that fits an explicit multi-compartment forward model end-to-end over spatial patches. The model combines cerebrospinal fluid (CSF), gray matter, up to K white-matter fiber compartments (stick-and-zeppelin), and a restricted compartment, with explicit fiber directions and soft model selection via repulsion and sparsity priors. PRISM supports a fast MSE objective and a Rician negative log-likelihood (NLL) that jointly learns σ without oracle information. A lightweight nuisance calibration module (smooth bias field and per-measurement scale/offset) is included for robustness and regularized to identity in clean-data tests. On synthetic crossing-fiber data (SNR = 30; five methods, 16 crossing angles), PRISM achieves 3.5° best-match angular error with 95% recall, which is 1.9× lower than the best baseline (MSMT-CSD, 6.8°, 83% recall); in NLL mode with learned σ, error drops to 2.3° with 99% recall, resolving crossings down to 20°. On the DiSCo1 phantom (NLL mode), PRISM improves connectivity correlation over CSD baselines at all four tracking angles (best r=.934 at 25° vs. .920 for MSMT-CSD). Whole-brain HCP fitting (∼741k voxels, MSE mode) completes in ∼12 min on a single GPU (hardware in Sec. 3) with near-identical results across random seeds.

Keywords: Diffusion MRI · Microstructure · Analysis-by-synthesis · Differentiable physics · Inverse problems.

1 Introduction

Diffusion MRI (dMRI) encodes tissue microstructure in direction- and b-value-dependent signal attenuation. In practice, recovering multiple fiber populations from finite, noisy measurements is strongly nonconvex, and standard voxelwise fitting discards spatial coherence. Intensity variations (bias field and per-measurement affine correction) and Rician magnitude noise further destabilize narrow-crossing recovery when handled in separate preprocessing stages.

This work proposes PRISM, a *differentiable analysis-by-synthesis* framework for fitting microstructure and nuisance calibration parameters in one end-to-end optimization. PRISM initializes tissue and calibration parameters, synthesizes

the expected measurement through a differentiable forward model (Fig. 1), and iteratively minimizes reconstruction error:

$$\hat{\mathbf{y}} = \underbrace{\mathcal{A}\big(\overbrace{\mathcal{S}(\boldsymbol{\theta}_{\text{micro}})}^{\text{tissue signal}};\ \boldsymbol{\theta}_{\text{cal}}\big)}_{\text{nuisance calibration}}, \qquad \min_{\boldsymbol{\theta}} \mathcal{L}(\hat{\mathbf{y}}, \mathbf{y}) + \mathcal{R}(\boldsymbol{\theta}), \tag{1}$$

where $\mathcal{L}$ is either an MSE or Rician NLL data-fidelity term (Sec. 2.5). Because every component is differentiable, gradients flow from reconstruction error through the calibration module to every tissue parameter, enabling the optimizer to focus on biological structure while absorbing residual nuisance. Spatial priors further enforce coherence across voxels.

Multi-compartment analytic fitters such as NODDI [23], SANDI [14], Ball-and-Sticks [3], and multi-tissue CSD [9] fit each voxel independently on preprocessed data. Global formulations couple measurements across space: COMMIT [4] reweights a precomputed tractogram (complementary to local fixel recovery), and diffusion basis-function decomposition [17] adds spatial regularization to intra-voxel fiber estimation. PRISM instead combines interpretable biophysics with end-to-end optimization over volumes.

Related work. Scalar microstructure models (DTI [2], DKI [8]) and orientation estimators (CSD [21]) assume artifact-free input and fit each voxel independently. Deep-learning approaches learn voxelwise mappings from signals to microstructure [12, 13, 20] (see [10] for a survey), enabling fast inference but requiring labeled, protocol-matched training; PRISM is training-free. ReMiDi [11] builds a differentiable dMRI simulator for finite-element meshes and iterates a latent mesh representation to match observed signals. Forward-simulation methods such as ODF-FP [1] and FORCE [19] precompute libraries of plausible fiber configurations and identify each voxel by nearest-neighbor lookup in signal or ODF space. PRISM targets a different operating regime: the multi-compartment inverse problem on clinical-style acquisitions. It uses an analytic signal model, optimizes over 3D spatial patches with neighbor-aware regularization, and jointly calibrates bias field, per-measurement affine correction, and noise level.

Contributions. (1) A differentiable multi-compartment microstructure fitter with explicit K-fiber parameterization, repulsion+sparsity-based soft model selection, and a restricted compartment. **(2)** Two data-fidelity modes: MSE for fast convergence, and Rician NLL with *self-supervised* noise-level recovery, which lowers overall angular error from 3.5° to 2.3°. **(3)** 1.9× lower best-match angular error than the best baseline on a five-method crossing-fiber benchmark (3.5° vs. 6.8° at SNR = 30, MSE mode) with 95% recall; NLL mode with learned σ reduces error to 2.3° with 99% recall; +1.4–1.9 pp in tractography correlation on DiSCo1 (NLL). **(4)** Whole-brain fitting of ∼741k voxels in ∼12 min on a single GPU with near-identical Dice/correlation across random seeds; optional nuisance calibration halves angular error under gain perturbations (σ_g=0.20) and self-regularizes to identity on unperturbed data.

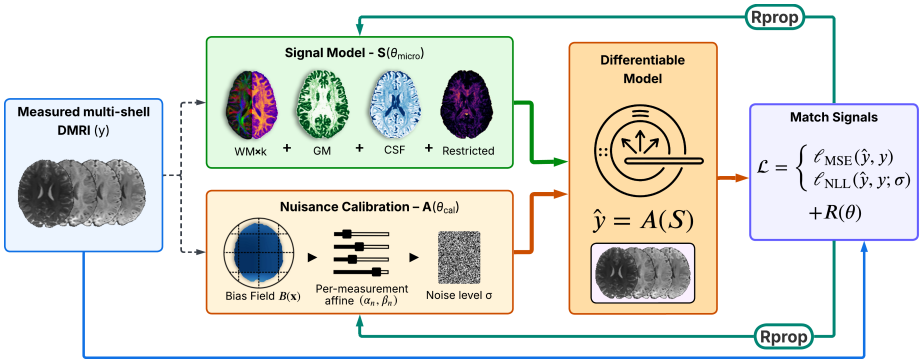

Fig. 1. PRISM analysis-by-synthesis pipeline. The multi-compartment tissue signal model $\mathcal{S}(\boldsymbol{\theta}_{\text{micro}})$ (top, green) combines up to K white-matter fiber populations with gray matter, CSF, and a restricted compartment. A nuisance calibration module $\mathcal{A}(\boldsymbol{\theta}_{\text{cal}})$ (bottom, orange) applies a smooth bias field $B(\mathbf{x})$, per-measurement affine correction (α_n, β_n), and noise level σ. The differentiable forward model synthesizes the expected measurement $\hat{y}{=}\mathcal{A}(\mathcal{S})$, which is compared to observed multi-shell dMRI via an MSE or NLL loss plus regularization $\mathcal{R}(\boldsymbol{\theta})$. Rprop updates tissue and calibration parameters jointly through gradient-based optimization.

2 Methods

2.1 Analysis-by-Synthesis Framework

PRISM fits a differentiable forward model to multi-shell dMRI data $\mathbf{y}$ by jointly optimizing tissue parameters $\boldsymbol{\theta}_{\text{micro}}$ and nuisance calibration parameters $\boldsymbol{\theta}_{\text{cal}}$ until the synthesized measurement matches the observed data. At voxel $\mathbf{x}$ and measurement n, the forward model composes a tissue signal with a lightweight calibration cascade:

$$\hat{y}(\mathbf{x}, n) = \mathcal{A}\big(\mathcal{S}(\mathbf{x}, n;\, \boldsymbol{\theta}_{\text{micro}});\, \boldsymbol{\theta}_{\text{cal}}\big) \tag{2}$$

where $\mathcal{S}$ generates the clean tissue signal and $\mathcal{A}$ applies nuisance intensity/noise calibration (Sec. 2.3). The forward model is closed-form analytic (Sec. 2.2); gradients come from PyTorch automatic differentiation, not a custom numerical or finite-element solver, distinguishing PRISM from mesh-based simulators [11]. Gradient-based optimization updates tissue and calibration parameters jointly.

2.2 Tissue Signal Model

The multi-compartment signal combines cerebrospinal fluid (CSF), gray matter (GM), white matter (WM, up to K fibers), and a restricted compartment, with fractions constrained via softmax:

$$\mathcal{S}(\mathbf{x}, n) = S_0(\mathbf{x}) \left(f_{\text{csf}}\, E_{\text{csf}} + f_{\text{gm}}\, E_{\text{gm}} + \sum_{k=1}^{K} f_{\text{wm},k}\, E_{\text{wm},k} + f_{\text{res}}\, E_{\text{res}} \right) \tag{3}$$

Isotropic compartments follow a mono-exponential decay with fixed diffusivities ($\times 10^{-3}\,\mathrm{mm}^2/\mathrm{s}$):

$$E_{\mathrm{csf}} = e^{-bD_{\mathrm{csf}}}, \quad E_{\mathrm{gm}} = e^{-bD_{\mathrm{gm}}}, \quad E_{\mathrm{res}} = e^{-bD_{\mathrm{res}}}, \tag{4}$$

where D_{csf}=3.0, D_{gm}=0.9, and D_{res}=0.2. Each WM fiber uses a stick-and-zeppelin model [15]:

$$E_{\mathrm{wm}} = f_{\mathrm{intra}}\, e^{-bD_{\parallel}\cos^2\theta} + (1-f_{\mathrm{intra}})\, e^{-b(D_{\parallel}\cos^2\theta + D_{\perp}\sin^2\theta)} \tag{5}$$

where θ is the angle between the fiber direction $\mathbf{d}_k$ and gradient direction $\mathbf{g}_n$, with fixed axial and radial diffusivities $D_{\parallel}$=1.7, $D_{\perp}$=0.4 ($\times 10^{-3}\,\mathrm{mm}^2/\mathrm{s}$), and learnable intra-axonal fraction f_{intra} (initialized at 0.5 via sigmoid).

Restricted compartment. A restricted isotropic compartment (D_{res}=0.2) captures slowly decaying high-b signal in small cellular structures (e.g., soma, astrocytes) [14] unexplained by extra-cellular or free-water diffusion.

Per-voxel parameters: S_0 (softplus), (K+3) tissue fractions (softmax over CSF, GM, K WM fibers, and restricted), $K\times 3$ fiber directions (ℓ_2-normalized), and f_{intra} (sigmoid).

2.3 Nuisance Intensity and Noise Calibration

The calibration module $\mathcal{A}$ applies three operators sequentially to the clean tissue signal $\mathcal{S}$: a spatial bias field, per-measurement intensity correction, and a learned noise level.

1. Spatial bias field $B(\mathbf{x})$: models multiplicative intensity variation from coil sensitivity and B_1 inhomogeneity. A low-resolution 8^3 control grid $\mathbf{b}_{\mathrm{lr}}$ (initialized at zero) is trilinearly upsampled to the data resolution and exponentiated: $B(\mathbf{x}) = \exp\big(\mathrm{upsample}(\mathbf{b}_{\mathrm{lr}})\big)$, enforcing low-frequency structure by construction.

2. Per-measurement intensity correction: A per-measurement affine model absorbs residual gradient-direction-dependent intensity variation: $\hat{y}_n \mapsto \exp(\alpha_n)\cdot \hat{y}_n + \beta_n$, where $\hat{y}_n$ is the predicted signal for measurement n, and α_n, β_n are initialized at zero (identity transform) and regularized toward identity ($\alpha_n \to 0$, $\beta_n \to 0$) via L_2 penalty. This captures multiplicative and additive per-volume drift that survives standard preprocessing.

3. Noise level $\sigma = \exp(\sigma_{\log})$ [6]: a learnable parameter recovered jointly with all others, serving as the noise scale for the Rician-aware loss mode (Sec. 2.5).

Composing the operators, the calibration chain for measurement n at voxel $\mathbf{x}$ is:

$$\hat{y}_n(\mathbf{x}) = \exp(\alpha_n)\, B(\mathbf{x})\, \mathcal{S}_n(\mathbf{x}) + \beta_n. \tag{6}$$

Scale control. The product $B(\mathbf{x}) \cdot S_0(\mathbf{x})$ is scale-ambiguous; the 8^3 control grid restricts B to low spatial frequencies while S_0 captures per-voxel intensity, giving a well-posed decomposition. L_2 penalties on $\mathbf{b}_{\mathrm{lr}}$ keep B near unity, and the per-measurement affine (α_n, β_n) is similarly regularized toward identity.

2.4 Regularization

The joint tissue–calibration model is inherently underdetermined: multiple configurations can explain the same observations. PRISM resolves this with lightweight regularizers that steer the solution toward anatomically plausible configurations without sacrificing data-driven flexibility. Weights are fixed globally and not tuned per dataset: because signals are b_0-normalized, the same values transfer across datasets (benchmark subsets in Sec. 3); the ablation quantifies each block's contribution, and a systematic sensitivity sweep is future work.

Spatial smoothness (local Huber-Laplacian). A local Laplacian prior over tissue fractions encourages spatial coherence across neighboring voxels:

$$\mathcal{R}_{\text{spatial}} = \frac{\lambda_{\text{sp}}}{|\mathcal{M}|} \sum_{\mathbf{x}\in\mathcal{M}} \rho\left(\mathbf{f}(\mathbf{x}) - \frac{1}{|\mathcal{N}(\mathbf{x})|} \sum_{\mathbf{x}'\in\mathcal{N}(\mathbf{x})} \mathbf{f}(\mathbf{x}') \right) \tag{7}$$

where $\mathbf{f}$ denotes the vector of tissue fractions, $\mathcal{N}(\mathbf{x})$ is a local neighborhood (6- or 26-connected), and $\rho(\cdot)$ is the Huber loss with transition δ=0.05. In ablation, 26-connected smoothing is not consistently better than 6-connected smoothing. Default: λ_{sp}=0.01.

Topology priors. The core model-selection term is direction repulsion with weight λ_{rep}=0.01: $\mathcal{R}_{\text{rep}} = \sum_{i<j} f_i f_j |\mathbf{d}_i \cdot \mathbf{d}_j|$. An L_1 sparsity penalty on minor fiber fractions (τ=0.15, λ_{sparse}=0.02) is also applied. Additional stabilization terms include orphan-WM suppression (penalizes WM fractions in voxels with near-zero S_0; λ=0.01) and fiber ordering (sorts fibers by descending fraction to break permutation symmetry; λ=0.01). A direction-continuity prior (λ_{dc}=0.005) discourages abrupt changes in each fiber's volume fraction *along its own orientation*, encoding that a fiber persists across neighboring voxels:

$$\mathcal{R}_{\text{dc}} = \frac{1}{K} \sum_{k=1}^{K} \mathbb{E}_{\mathbf{x}}\Big[f_k(\mathbf{x}) \big(|f_k(\mathbf{x}) - f_k(\mathbf{x}+\varepsilon\mathbf{d}_k)| + |f_k(\mathbf{x}) - f_k(\mathbf{x}-\varepsilon\mathbf{d}_k)| \big) \Big], \tag{8}$$

where $f_k \equiv f_{\text{wm},k}$ is the WM fraction of fiber k, sampled by trilinear interpolation at $\pm\varepsilon$ (ε=0.5 voxel) along its own direction $\mathbf{d}_k$. The f_k weighting localizes the penalty to fiber-bearing voxels, the $\pm$ pair respects antipodal symmetry, and the per-fiber sum is permutation-invariant, needing no cross-voxel peak matching.

Calibration regularizers. L_2 penalties on per-measurement log-scales (α_n) and offsets (β_n), and bias-field coefficients ($\mathbf{b}_{\text{lr}}$), plus total-variation smoothness on the upsampled bias field. These anchor all calibration parameters near identity, preventing them from absorbing biological signal.

2.5 Optimization

PRISM optimizes within a brain mask $\mathcal{M}$ on b_0-normalized signals (both $\hat{y}$ and y divided by the measured b_0). Two data-fidelity modes are available:

MSE mode treats magnitude noise as Gaussian:

$$\mathcal{L}_{\mathrm{MSE}} = \frac{1}{|\mathcal{M}|} \sum_{\mathbf{x},n} \big(\hat{y}(\mathbf{x},n) - y(\mathbf{x},n)\big)^2. \tag{9}$$

This is fast and stable, requiring no noise estimate.
Rician NLL mode uses the correct likelihood for magnitude MRI [6]:

$$\mathcal{L}_{\mathrm{NLL}} = \sum_{\mathbf{x},n} \Big[\log \sigma^2 + \tfrac{y^2+\hat{y}^2}{2\sigma^2} - \log y - \log I_0\Big(\tfrac{y\,\hat{y}}{\sigma^2}\Big) \Big], \tag{10}$$

where $\hat{y}$ is the predicted noise-free signal from Eq. (2), I_0 is the modified Bessel function of the first kind, and σ is the noise standard deviation. The $-\log y$ term depends only on the observed magnitude and is constant with respect to all optimized parameters; it does not affect the gradients or fitted values. Crucially, σ need not be known *a priori*: it is jointly optimized as part of the calibration module (Sec. 2.3, parameter $\sigma_{\log}$).

Both modes share the same regularization and solver. The total objective is $\mathcal{L}_{\mathrm{data}} + \sum_i \lambda_i \mathcal{R}_i + \mathcal{R}_{\mathrm{cal}}$.

Constraints and solver. Fractions via softmax; f_{intra} via sigmoid; S_0 via softplus; directions via ℓ_2-normalization. Rprop [18] is chosen for its sign-based updates, which naturally handle heterogeneous parameter scales (orientations, fractions, calibration fields) without a global learning rate; among eight optimizers on 20 axial slices it reaches the lowest final MSE (21.2×10^{-3} vs. $\geq$21.5) in the fewest iterations to threshold (22 vs. $\geq$39). Convergence requires 50–300 iterations depending on dataset and slab size.

Implementation details. Whole-brain fitting uses overlapping 30-slice slabs (5-slice overlap) stitched via weighted averaging. All λ values are fixed across every experiment (λ_{sp}=λ_{rep}=0.01, $\lambda_{\mathrm{sparse}}$=0.02). HCP whole-brain fitting uses 100 iterations per slab on a single GPU; DiSCo1 uses 300 iterations (NLL mode). Code and configurations are available at `https://github.com/m-agour/Prism`.

3 Experiments

Datasets. Evaluation spans three complementary benchmarks: (i) *Synthetic crossing-fibers* [5]: 3400 voxels (16 crossing angles from 15° to 90° in 5° steps, plus single-fiber controls); 3 shells ($b\in\{1{,}000, 2{,}000, 3{,}000\}$ s/mm^2, 64 directions each, N=193), SNR = 30 (also 10 and 50). A subset ($\{30, 45, 60, 90\}^\circ$, 200 voxels each) is reused for the artifact-robustness experiment; (ii) *DiSCo1 phantom* [16]: digital phantom with ground-truth connectivity (40^3, 16 ROIs, 120 ROI pairs), 4 shells (b up to 13,192 s/mm^2, N=364), SNR = 50; (iii) *HCP in-vivo* [22]: subject 100307, 145×174×145 (∼741k brain voxels), 3 shells (N=288).

Configuration. PRISM uses K=2 fibers for synthetic data, K=3 for HCP, and K=5 for DiSCo1, on a single NVIDIA RTX A6000 GPU (48 GB VRAM). Synthetic experiments evaluate both MSE and NLL modes with direction repulsion only (no spatial priors); NLL uses the self-supervised σ (no oracle). All

synthetic results are seeded for deterministic reproducibility in the reported software/hardware environment. DiSCo1 adds topology priors (300 iterations, NLL mode); HCP uses the full regularization suite (MSE mode, 100 iterations per 30-slice slab with 5-slice overlap). For spatial smoothness, both 6-connected and 26-connected neighborhoods are evaluated in ablation. Fiber directions are used directly as tracking peaks.

Baselines. CSD [21], MSMT-CSD [9], and ODF-FP [1]. CSD estimates its response function from ground-truth single-fiber voxels; MSMT-CSD receives oracle response functions matching the simulation eigenvalues. ODF-FP uses a 1,000,000-element library with max 3 peaks, $D_a \in [1.2, 2.2]$, $D_e, D_r \in [0.1, 0.6]$ ($\times 10^{-3}\,\mathrm{mm}^2/\mathrm{s}$), and DSI 8-fold tessellation. CSD receives a single shell (b=3,000 $\mathrm{s/mm}^2$); all other methods receive the full multi-shell input. ODF-FP is evaluated on the synthetic benchmark only; DiSCo1 uses CSD and MSMT-CSD.

Tracking & metrics. Identical deterministic tracking [5] (seed density 2, step 0.2 mm, fixed seed). Connectivity scored via Pearson r; angular error reported as best-match (each ground-truth fiber matched to its closest predicted fiber). F1 and recall quantify fiber detection.

4 Results

Robustness to Intensity Nuisance Factors. Nuisance calibration is evaluated under per-measurement gain perturbations (σ_g=0.02–0.20, SNR = 30, crossing angles $\{30, 45, 60, 90\}^\circ$, 200 voxels each). At σ_g=0.20, calibration halves angular error ($4.8^\circ \to 2.4^\circ$, -50%) and removes 85% of reconstruction MSE ($7.6 \to 1.1 \times 10^{-3}$); on clean or well-preprocessed data, all calibration parameters self-regularize to identity (see ablation below).

Crossing-Fiber Benchmark. Table 1 reports fiber recovery on the 16-angle synthetic benchmark (Sec. 3). $\mathrm{PRISM}_{\mathrm{MSE}}$ achieves 3.5° overall best-match angular error with 95% recall, 1.9× lower than the best baseline (MSMT-CSD, 6.8°, 83% recall). In NLL mode (learned σ), error drops to 2.3° with 99% recall and 99% F1. At narrow crossings ($\leq 30^\circ$), CSD and MSMT-CSD largely detect only a single fiber (overall recall 82–83%), while ODF-FP partially resolves the second peak and achieves the lowest baseline error in this regime (7.4°–13.7° vs. 7.8°–15.2° for MSMT-CSD); however, ODF-FP's dictionary discretization degrades at wider angles ($\geq 45^\circ$, Table 1). $\mathrm{PRISM}_{\mathrm{NLL}}$ resolves both fibers across all angles with 3.1°–5.8° best-match error at $\leq 30^\circ$. PRISM uses a deliberately mismatched forward model ($D_\perp$=0.4 vs. true $0.3 \times 10^{-3}\,\mathrm{mm}^2/\mathrm{s}$), while baselines receive oracle or ground-truth-derived response functions (Sec. 3).

DiSCo1 Tractography. On the DiSCo1 phantom (16 ROIs, 120 ROI pairs, 11,528 ground-truth strands, SNR = 50), we run identical deterministic tracking for CSD, MSMT-CSD, and PRISM (with Rician NLL) and evaluate connectivity correlation at tracking angles 15–30°. PRISM outperforms both baselines at *every* evaluated angle: at their optimal angles, CSD reaches r=.917 (30°), MSMT-CSD reaches

Table 1. Best-match angular error (°) at SNR = 30 (16 angles, 15–90°; representative subset). F1 and recall (%) over all angles.

Method	1-fib	15°	20°	25°	30°	45°	60°	75°	90°	Overall	F1	Rec.
$\text{PRISM}_{\text{MSE}}$	**0.7**	7.5	9.9	11.7	8.2	1.9	1.5	1.5	**1.3**	3.5	96	95
$\text{PRISM}_{\text{NLL}}$	1.4	**5.8**	**4.9**	**3.9**	**3.1**	**1.9**	**1.5**	**1.3**	**1.3**	**2.3**	**99**	**99**
CSD	3.1	7.9	10.3	12.7	15.2	8.8	3.9	4.0	3.5	7.7	90	82
MSMT-CSD	3.0	7.8	10.2	12.7	15.2	5.3	3.4	3.2	3.1	6.8	91	83
ODF-FP	3.1	7.4	9.7	11.9	13.7	16.9	10.9	9.4	10.9	11.6	87	86

r=.920 (30°), and PRISM achieves r=.934 (25°). Across the full range, PRISM scores r=.925/.930/.934/.928 versus .906/.911/.918/.920 for MSMT-CSD, with the largest per-angle gains at 15° and 20° (+1.9 pp). PRISM's optimum shifts to a tighter tracking angle (25° vs. 30°), indicating that its fiber orientations support narrower angular thresholds while maintaining higher overall connectivity.

In-Vivo Whole-Brain Fitting. On HCP data (741k voxels, 288 measurements, 3 shells), PRISM reaches MSE = 6.3×10^{-3} in ∼12 min (MSE mode; hardware in Sec. 3). The recovered tissue fractions are spatially consistent with known anatomy: cortical GM ribbons, deep WM tracts, and ventricular CSF.

Compared against FreeSurfer T_1 segmentations, PRISM achieves Dice = 0.79 (WM) and 0.80 (GM) with Pearson r=0.75/0.67; all summary metrics vary by <0.003 across three random seeds.

Ablation. We ablate PRISM on DiSCo1 as a cumulative build-up (25°): a vanilla 5-fiber model (MSE; no restricted compartment, priors, or calibration) reaches r=.840, already below MSMT-CSD (.920). Adding the restricted compartment is by far the largest single gain (r=.927, +8.7 pp)—it absorbs the high-b (b=13,192) signal—while nuisance calibration and the spatial/topology priors with Rician NLL (full, r=.934) add smaller refinements, consistent with calibration self-regularizing toward identity. PRISM's contribution is the framework that admits this compartment alongside nuisance calibration in one optimization.

5 Discussion and Conclusion

PRISM reframes microstructure estimation as a differentiable fitting problem, coupling an explicit multi-compartment forward model with end-to-end optimization rather than isolated voxelwise fitting stages; this explicit fiber parameterization yields a 1.9× angular-error reduction over the strongest baseline on the 16-angle benchmark (Table 1), because each fiber is treated as a first-class parameter rather than an implicit SH peak. The ablation shows the restricted compartment provides by far the largest single gain (vanilla r=.840 → +restricted .927, +8.7 pp; without it PRISM trails MSMT-CSD), with calibration and priors adding

smaller refinements toward the full r=.934—part of this gain is added model capacity the baselines lack. The intensity perturbation study confirms nuisance calibration activates under gain corruption while self-regularizing to identity on unperturbed data. On HCP, dMRI-derived tissue fractions agree with FreeSurfer (Dice$\approx$ 0.79/0.80 WM/GM) without T1 input.

Limitations. PRISM currently models intensity-domain nuisance factors; geometric preprocessing (motion, susceptibility) is still assumed. Diffusivities are fixed (NODDI/SANDI practice) to avoid ill-conditioning at clinical SNR; large mismatches can bias peaks [7], though PRISM is already evaluated under a deliberate $D_\perp$ mismatch (0.4 vs. 0.3) and still beats oracle-calibrated baselines. Convergence could be accelerated by warm-starting from a self-supervised encoder or a dictionary-based method [1, 19], reducing sensitivity to local minima.

In summary, PRISM demonstrates that differentiable analysis-by-synthesis with explicit fibers, restricted diffusion, and learned-σ Rician fitting yields substantial gains in crossing resolution, tractography agreement, and anatomically consistent whole-brain maps on a single GPU.

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

References

1. Baete, S.H., Cloos, M.A., Lin, Y.C., Placantonakis, D.G., Shepherd, T., Boada, F.E.: Fingerprinting orientation distribution functions in diffusion MRI detects smaller crossing angles. NeuroImage **198**, 231–241 (2019). https://doi.org/10.1016/j.neuroimage.2019.05.024
2. Basser, P.J., Mattiello, J., LeBihan, D.: MR diffusion tensor spectroscopy and imaging. Biophysical Journal **66**(1), 259–267 (1994)
3. Behrens, T.E.J., Woolrich, M.W., Jenkinson, M., Johansen-Berg, H., Nunes, R.G., Clare, S., Matthews, P.M., Brady, J.M., Smith, S.M.: Characterization and propagation of uncertainty in diffusion-weighted MR imaging. Magnetic Resonance in Medicine **50**(5), 1077–1088 (2003)
4. Daducci, A., Dal Palù, A., Lemkaddem, A., Thiran, J.P.: COMMIT: Convex optimization modeling for microstructure informed tractography. IEEE Transactions on Medical Imaging **34**(1), 246–257 (2015)
5. Garyfallidis, E., et al.: Dipy, a library for the analysis of diffusion MRI data. Frontiers in Neuroinformatics **8**, 8 (2014)
6. Gudbjartsson, H., Patz, S.: The Rician distribution of noisy MRI data. Magnetic Resonance in Medicine **34**(6), 910–914 (1995)
7. Guo, F., Tax, C.M.W., De Luca, A., Viergever, M.A., Heemskerk, A., Leemans, A.: Fiber orientation distribution from diffusion MRI: Effects of inaccurate response function calibration. Journal of Neuroimaging **31**(6), 1082–1098 (2021)
8. Jensen, J.H., Helpern, J.A., Ramani, A., Lu, H., Kaczynski, K.: Diffusional kurtosis imaging: The quantification of non-Gaussian water diffusion by means of magnetic resonance imaging. Magnetic Resonance in Medicine **53**(6), 1432–1440 (2005)
9. Jeurissen, B., Tournier, J.D., Dhollander, T., Connelly, A., Sijbers, J.: Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion MRI data. NeuroImage **103**, 411–426 (2014)

10. Karimi, D., Warfield, S.K.: Diffusion MRI with machine learning. Imaging Neuroscience **2**, 1–55 (2024)
11. Khole, P.P., Petiwala, Z.K., Magesh, S.P., Mirafzali, E., Gupta, U., Li, J.R., Ianus, A., Marinescu, R.: ReMiDi: Reconstruction of microstructure using a differentiable diffusion MRI simulator. arXiv preprint arXiv:2502.01988 (2025)
12. Koppers, S., Haarburger, C., Merhof, D.: Diffusion MRI signal augmentation: From single shell to multi shell with deep learning. In: Computational Diffusion MRI (MICCAI Workshop). pp. 61–70 (2017)
13. Nath, V., Schilling, K.G., Parvathaneni, P., Hansen, C.B., Hainline, A.E., Huo, Y., Blaber, J.A., Lyu, I., Janve, V., Gao, Y., Stepniewska, I., Anderson, A.W., Landman, B.A.: Deep learning reveals untapped information for local white-matter fiber reconstruction in diffusion-weighted MRI. Magnetic Resonance Imaging **62**, 220–227 (2019)
14. Palombo, M., Ianus, A., Guerreri, M., Nunes, D., Alexander, D.C., Shemesh, N., Zhang, H.: SANDI: A compartment-based model for non-invasive apparent soma and neurite imaging by diffusion MRI. NeuroImage **215**, 116835 (2020). https://doi.org/10.1016/j.neuroimage.2020.116835
15. Panagiotaki, E., Schneider, T., Siow, B., Hall, M.G., Lythgoe, M.F., Alexander, D.C.: Compartment models of the diffusion MR signal in brain white matter: A taxonomy and comparison. NeuroImage **59**(3), 2241–2254 (2012)
16. Rafael-Patino, J., Girard, G., Truffet, R., Pizzolato, M., Caruyer, E., Thiran, J.P.: The diffusion-simulated connectivity (DiSCo) dataset. Data in Brief **38**, 107429 (2021)
17. Ramirez-Manzanares, A., Rivera, M., Vemuri, B.C., Carney, P., Mareci, T.: Diffusion basis functions decomposition for estimating white matter intravoxel fiber geometry. IEEE Transactions on Medical Imaging **26**(8), 1091–1102 (2007)
18. Riedmiller, M., Braun, H.: A direct adaptive method for faster backpropagation learning: The RPROP algorithm. In: Proceedings of the IEEE International Conference on Neural Networks. pp. 586–591 (1993)
19. Shah, A.J., Henriques, R.N., Ramirez-Manzanares, A., Filipiak, P., Baete, S., Deka, K., Gor, M., Koudoro, S., Garyfallidis, E.: FORCE: FORward modeling for Complex microstructure Estimation. Research Square preprint (2025). https://doi.org/10.21203/rs.3.rs-8151109/v1, version 1, 20 November 2025
20. Tian, Q., Bilgic, B., Fan, Q., Liao, C., Ngamsombat, C., Hu, Y., Witzel, T., Setsompop, K., Polimeni, J.R., Huang, S.Y.: DeepDTI: High-fidelity six-direction diffusion tensor imaging using deep learning. NeuroImage **219**, 117017 (2020)
21. Tournier, J.D., Calamante, F., Connelly, A.: Robust determination of the fibre orientation distribution in diffusion MRI: Non-negativity constrained super-resolved spherical deconvolution. NeuroImage **35**(4), 1459–1472 (2007)
22. Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E.J., Yacoub, E., Ugurbil, K., WU-Minn HCP Consortium: The WU-Minn Human Connectome Project: An overview. NeuroImage **80**, 62–79 (2013)
23. Zhang, H., Schneider, T., Wheeler-Kingshott, C.A.M., Alexander, D.C.: NODDI: Practical in vivo neurite orientation dispersion and density imaging of the human brain. NeuroImage **61**(4), 1000–1016 (2012)