

Spinverse: Differentiable Physics for Permeability-Aware Microstructure Reconstruction from Diffusion MRI

Prathamesh Pradeep Khole^{1*}, Mario M. Brenes^{1*}, Zahra Kais Petiwala², Ehsan Mirafzali¹, Utkarsh Gupta¹, Jing-Rebecca Li³, Andrada Ianus⁴, and Razvan Marinescu¹

¹ University of California Santa Cruz, Santa Cruz, CA, USA
pkhole@ucsc.edu

² University of California San Diego, San Diego, CA, USA

³ Inria-Saclay, France

⁴ Kings College London, London, UK

Abstract. Diffusion MRI (dMRI) is sensitive to microstructural barriers, yet most existing methods either assume impermeable boundaries or estimate voxel-level parameters without recovering explicit interfaces. We present *Spinverse*, a permeability-aware reconstruction method that inverts dMRI measurements through a fully differentiable Bloch–Torrey simulator. *Spinverse* represents tissue on a fixed tetrahedral grid and treats each interior face permeability as a learnable parameter; low-permeability faces act as diffusion barriers, so microstructural boundaries whose topology is not fixed a priori (up to the resolution of the ambient mesh) emerge without changing mesh connectivity or vertex positions. Given a target signal, we optimize face permeabilities by backpropagating a signal-matching loss through the PDE forward model, and recover an interface by thresholding the learned permeability field. To mitigate the ill-posedness of permeability inversion, we use mesh-based geometric priors; to avoid local minima, we use a staged multi-sequence optimization curriculum. Across a collection of synthetic voxel meshes, *Spinverse* reconstructs diverse geometries and demonstrates that sequence scheduling and regularization are critical to avoid outline-only solutions while improving both boundary accuracy and structural validity. Our source code is available at <https://github.com/BioMedAI-UCSC/spinverse>.

Keywords: Differentiable physics · Diffusion MRI · permeability · microstructure reconstruction · inverse problems · finite element method.

1 Introduction

Diffusion MRI (dMRI) is sensitive to micron-scale barriers that restrict diffusion and enable exchange, offering a noninvasive window into tissue microstructure [1,14]. Yet, *reconstructing explicit microstructural boundaries* from the measured

* Equal contribution

signal remains challenging: the signal is indirect, integrates over sub-voxel geometries, and multiple configurations can produce similar attenuation profiles. Most approaches therefore estimate *voxel-level* parameters via simplified compartment models such as NODDI [23] and SANDI [15], which provide millimeter-scale biomarkers but do not recover *where* barriers are, and typically treat permeability as negligible.

Physics-based simulators [11] and simulation-driven pipelines [6] offer high-fidelity forward prediction in complex geometries, yet their outputs are signals or aggregate parameter estimates rather than explicit barrier reconstructions. Recent differentiable solvers [8,9] enable inverse mesh reconstruction from dMRI but remain ill-posed without strong priors; those works optimized vertex coordinates in a low-dimensional auto-encoded spectral space under an arbitrary but fixed topology. A method that optimizes meshes of variable topologies would be a significant step towards the reconstruction of accurate diffusion MRI geometries.

In this work, we introduce Spinverse, a method that reconstructs microstructural boundaries whose topology is not fixed a priori (up to the resolution of the ambient mesh) by optimizing face-wise permeabilities on a fixed tetrahedral grid. Our contributions are: **(1)** we formulate *permeability-aware boundary reconstruction* from dMRI as an inverse problem over a *face-wise permeability field* on a volumetric mesh, where boundaries emerge as low-permeability interfaces; **(2)** we implement an *end-to-end differentiable* Bloch–Torrey [21] simulator that treats each tetrahedron as an independent compartment with face permeabilities as optimizable parameters. Building on [11,9,8], we re-formulate the forward model in PyTorch [16] with parallelized per-element solves, enabling scalable gradient-based inversion via Autograd; and **(3)** we propose an inversion strategy combining *spatial priors* (MRF-based continuity [24,3,17] and manifold regularization [2]) with a *multi-sequence curriculum* to mitigate ill-posedness and improve interior boundary recovery.

2 Related Work

Compartment models (e.g., NODDI [23], SANDI [15]) estimate voxel-level parameters but do not recover explicit barriers or mesh permeabilities. FEM Bloch–Torrey simulators such as SpinDoctor [11] and [8] predict signals in complex geometries with permeable interfaces, and simulation-driven learning maps signals to aggregate parameters [6]. Differentiable inverse solvers such as ReMiDi [9] reconstruct mesh geometry under a Bloch–Torrey [21] forward model but assume fixed mesh topology; Spinverse instead infers boundaries by optimizing a face-wise permeability field on a fixed mesh grid, with interfaces emerging via the coupling operator $\mathbf{B}(\kappa)$ in (2), so the extracted interface topology is determined entirely by the optimization and is not constrained to match a pre-defined template, unlike vertex-based approaches that inherit the topology of their initial mesh. More broadly, differentiable mesh/topology relaxations such as DMesh++ [20] motivate soft-to-hard optimization, which we adopt under a Bloch–Torrey PDE [21] based forward model.




Fig. 1. Spinverse pipeline. *Top:* Given reference dMRI signals S_{ref} under multiple PGSE encodings, we optimize face permeabilities $\{\kappa_f\}$ on a fixed tetrahedral mesh by minimizing MSE loss with regularization. Gradients are backpropagated through a differentiable Bloch–Torrey simulator and used by Adam to update κ_f , recovering boundaries via thresholding. *Bottom (a–e):* Simulator internals: compartments and face permeabilities (a); coupling operator $B(\kappa)$ assembly (b); Bloch–Torrey FEM system (c); reduced-order eigensolve and matrix-exponential propagation (d); signal readout (e).

3 Methods

Problem setup: per-element compartments and learnable permeabilities. Let $\mathcal{M} = (\mathcal{V}, \mathcal{T})$ be an ambient tetrahedral mesh discretizing a bounded domain $\Omega \subset \mathbb{R}^3$. Rather than pre-segmenting $\mathcal{M}$ into tissue compartments, we treat *every tetrahedron as an independent compartment*. Each interior face $f \in \mathcal{F}_{\text{int}}$ (shared by exactly two tetrahedra) becomes a semi-permeable interface with learnable permeability $\kappa_f > 0$ (Fig. 1a). We optimize unconstrained parameters $\theta_f \in \mathbb{R}$ and obtain physical permeabilities via a tempered sigmoid reparametrization in log-space, $k_f = \log_{10} \kappa_f = k_{\min} + (k_{\max} - k_{\min}) \sigma(\theta_f / \tau_\sigma)$, where σ is the logistic sigmoid, $\tau_\sigma > 0$ a temperature, and $[k_{\min}, k_{\max}]$ bounds the log-permeability range. Vertex positions and mesh connectivity are fixed; microstructure emerges solely through $\{\kappa_f\}$, and a barrier set is recovered post-hoc as $\Gamma(\tau_b) = \{f \in \mathcal{F}_{\text{int}} : \kappa_f < \tau_b\}$ (Fig. 1, top).

Differentiable forward model. The transverse magnetization $m(\mathbf{x}, t)$ at position $\mathbf{x} \in \Omega$, time $t \in [0, T_E]$, satisfies the Bloch–Torrey equation [11] under diffusion encoding $a \in \mathcal{A}$:

$$\frac{\partial m}{\partial t} = \nabla \cdot (\mathbf{D} \nabla m) - \frac{1}{T_2} m - i\gamma \mathbf{G}_a(t) \cdot \mathbf{x} m, \quad m(\mathbf{x}, 0) = \rho, \quad (1)$$

where $\mathbf{D}$ is the diffusion tensor, T_2 the transverse relaxation time, γ the gyro-magnetic ratio, i the imaginary unit, ρ the proton density, and $\mathbf{G}_a(t)$ the PGSE

gradient waveform of acquisition a (encoding direction, δ , Δ , b -value). Under per-element compartmentalization, every interior face carries the interface condition $\mathbf{n}^\pm \cdot \mathbf{D}^\pm \nabla m^\pm = \kappa_f (m^\mp - m^\pm)$, where superscripts $\pm$ denote the two tetrahedra sharing face f and $\mathbf{n}^\pm$ their outward normals. This reduces to a reflecting barrier as $\kappa_f \rightarrow 0$ (Fig. 1a). The predicted signal is $S_a(\kappa) = \int_\Omega m(\mathbf{x}, T_E) d\mathbf{x}$.

Discretizing (1) with P1 (piecewise-linear) finite elements following [11] yields the semi-discrete system (Fig. 1b-c):

$$\mathbf{M}\dot{\mathbf{c}} = -(\mathbf{K} + \mathbf{R} + i\mathbf{Q}_a(t) + \mathbf{B}(\kappa))\mathbf{c}, \quad (2)$$

where $\mathbf{M}, \mathbf{K}, \mathbf{R}, \mathbf{Q}_a$ are mass, stiffness, relaxation, and gradient dephasing matrices. The coupling operator $\mathbf{B}(\kappa)$ is assembled on *all* interior faces from face-area-weighted Robin coupling terms; it depends linearly on $\{\kappa_f\}$ and is the *only* operator updated during optimization. For efficiency we project into a truncated generalized eigenbasis of $(\mathbf{K} + \mathbf{B}(\kappa))\mathbf{u} = \lambda\mathbf{M}\mathbf{u}$ and propagate in the reduced coordinates via matrix exponentials (Fig. 1d-e) [11,6,9], reading out the signal as $S_a = \mathbf{1}^\top \mathbf{M}\mathbf{U}\mathbf{y}(T_E)$. In practice, we refresh the truncated basis every N iterations and reuse it between refreshes; on intermediate iterations we update only the reduced interface operator $\mathbf{U}^\top \mathbf{B}(\kappa)\mathbf{U}$ while treating $\mathbf{U}$ as fixed, so gradients flow through the reduced projections and time propagation back to $\{\theta_f\}$.

Optimization: data fidelity, priors, and curriculum. We estimate θ^* by minimizing (Fig. 1, top)

$$\theta^* = \arg \min_{\theta} \sum_{a \in \mathcal{A}(t)} \lambda_{\text{data}} \mathcal{L}_{\text{data}}(S_a(\kappa(\theta)), S_a^*) + \lambda_{\text{cont}} \mathcal{R}_{\text{cont}}(\kappa(\theta)) + \lambda_{\text{man}} \mathcal{R}_{\text{man}}(\kappa(\theta)) \quad (3)$$

where $\kappa(\theta) = \{\kappa_f(\theta)\}_{f \in \mathcal{F}_{\text{int}}}$ are face permeabilities obtained from unconstrained parameters θ via a positive mapping (Sec. 3). We use a weighted squared error on normalized signals, $\mathcal{L}_{\text{data}}(S_a, S_a^*) = \|S_a/\hat{S}_0 - S_a^*/S_0^*\|_2^2$, where $\hat{S}_0$ and S_0^* are predicted and measured $b=0$ references.

The face-wise field admits many signal-equivalent configurations, so we regularize using mesh adjacency. *Weighted continuity* penalizes discrepancies between adjacent faces while allowing sharp transitions at *candidate interfaces* by downweighting face pairs likely to straddle a boundary. Specifically, $\mathcal{R}_{\text{cont}}(\kappa) = \sum_{(f,f') \in \mathcal{N}} w_{ff'} (\kappa_f - \kappa_{f'})^2$, with weights $w_{ff'} = 1 - |p_{\text{non}}(f) - p_{\text{non}}(f')|$. Here $\mathcal{N}$ is the set of adjacent interior-face pairs and $p_{\text{non}}(f)$ is a soft indicator of a non-permeable (barrier) face computed from κ_f .

We introduce a *soft manifold regularizer* that penalizes mesh edges with other than 0 or 2 incident boundary faces, discouraging non-manifold junctions. For an edge e with incident faces $\mathcal{F}_e$, let $k_e = \sum_{f \in \mathcal{F}_e} p_{\text{non}}(f)$; we penalize deviations from $\{0, 2\}$ via $\mathcal{R}_{\text{man}}(\kappa) = \sum_{e \in \mathcal{E}} \text{softmin}(k_e^2, (k_e - 2)^2)$, where $\text{softmin}(a, b) = -\tau_m \log(e^{-a/\tau_m} + e^{-b/\tau_m})$ with temperature $\tau_m > 0$.

Diffusion acquisitions probe complementary physical regimes: long diffusion times allow spins to encounter barriers and exchange across interfaces, making the signal sensitive to the coarse compartment layout whereas short diffusion times probe local restriction and fine boundary details. Naïvely optimizing all




Fig. 2. Sample data and reconstructions. (Left) Representative ground-truth meshes: single-axon (A) and two-axon (B) geometries embedded in the ambient tetrahedral grid. (Right) Four example Spinverse reconstructions (red) shown alongside ground truth meshes (blue). Each pair reports CD-L2 (CD) and BadEdge% (BE).

regimes jointly can induce gradient conflict. We therefore use a staged curriculum: we first optimize on long diffusion-time acquisitions (exchange-dominated) to establish a coarse interface configuration, then transition to short diffusion-time acquisitions (restriction-dominated) to refine fine-scale structure. Formally, we partition acquisitions into $\mathcal{A}_{\text{long}}$ and $\mathcal{A}_{\text{short}}$ and set $\mathcal{A}(t) = \mathcal{A}_{\text{long}}$ for $0 \leq t < T_{\text{switch}}$ and $\mathcal{A}(t) = \mathcal{A}_{\text{short}}$ for $T_{\text{switch}} \leq t \leq T$, where T is the total number of iterations and T_{switch} is fixed per experiment (Sec. 4).

4 Experiments and Results

We evaluate reconstructions with metrics for signal fidelity, geometric accuracy, and surface validity. The section is organized as qualitative reconstructions, ablations of priors and scheduling, and a comparison to learned predictors on multi-compartment axon-like scenes.

Setup. We evaluate our method on synthetic voxel-scale meshes of varying geometry (e.g., sphere, axons, torus) on a fixed cubic domain $x, y, z \in [-13.6, 13.6] \mu\text{m}$. Each ambient mesh contains 189 vertices, 1296 faces and 576 elements. Signals are simulated under PGSE with $b \in \{0, 1000, 2000, 3000, 4000, 5000\} \text{ s/mm}^2$, $\delta = 10 \text{ ms}$, and $\Delta \in \{20, 60\} \text{ ms}$ (short/long diffusion time). To accelerate the solver, we truncate the Laplace eigenbasis, retaining only $N_{\text{eig}} = 500$ eigenvalues and eigenmodes. We optimize a face-wise permeability field $\kappa_f \in [10^{-5}, 10^{-1}]$ and extract the interface by thresholding at $\tau_b = 10^{-3}$. Reconstructions are initialized at the threshold ($\kappa_f = 10^{-3}$), while reference meshes assign low permeability ($\kappa_f = 10^{-5}$) to boundary faces (all other faces set to $\kappa_f = 10^{-1}$). In our protocol, $\mathcal{A}_{\text{short}}$ denotes acquisitions with $\Delta = 20 \text{ ms}$ and $\mathcal{A}_{\text{long}}$ denotes acquisitions with $\Delta = 60 \text{ ms}$ (Sec. 3). We optimize θ with Adam [13,10] (base learning rate $\eta_0 = 0.75$) using a cyclic warmup+cosine schedule: each $C = 200$ -iteration cycle linearly warms up for $W = 50$ iterations, then cosine-decays to $\alpha\eta_0$ with $\alpha = 0.1$. We set $\lambda_{\text{data}} = 100$, $\lambda_{\text{cont}} = 2$, and linearly ramp λ_{man} from 0 to 2.0 over the first 400 iterations. We run for 400 iterations and switch acquisitions at $T_{\text{switch}} = 200$.





Fig. 3. Acquisition scheduling ablation (torus). Joint 50/50 optimization over $\mathcal{A}_{\text{long}} \cup \mathcal{A}_{\text{short}}$ (*Split*) versus a staged schedule $\mathcal{A}_{\text{long}} \rightarrow \mathcal{A}_{\text{short}}$ (*Staged*). Top: signal MSE and Chamfer (CD-L2) over iterations. Bottom: extracted interfaces at iterations 50/250/final (optimized, red) and ground truth (blue); staging yields a coherent torus boundary while joint optimization fails to recover the opening.

Metrics. We measure (i) *signal fidelity* using MSE on normalized signals S/S_0 , averaged across all b -values and gradient directions; (ii) *geometric accuracy* using symmetric Chamfer distance with squared Euclidean distances (CD-L2) [5], computed after centering each point set sampled from boundary faces (consistent with translation invariance of the diffusion MRI signal); and (iii) *structural validity* using edge-incidence statistics. Let $d(e)$ be the number of boundary faces incident to edge e (restricted to edges with $d(e) > 0$); we report BadEdge% (BE%) ($d(e) \neq 2$) to quantify deviations from a two-manifold interface.

Qualitative reconstructions. Fig. 2(i)–(iv) shows representative final reconstructions across geometries (optimized interface in red, ground truth in blue). Spinverse accurately recovers the sphere (i) and diagonal cylinder (ii) from a barrier-free initialization (CD-L2 0.00/1.37, BE% 0.00/20.15), and the additional examples in (iii)–(iv) exhibit similar cylindrical boundary recovery with comparable structural artifacts from thresholding. See the supplementary material for videos of these reconstructions, including the torus case.

Acquisition schedule ablation (torus). We focus on the torus as a challenging case with nontrivial topology and an interior opening (Fig. 3). Joint 50/50 optimization over $\mathcal{A}_{\text{long}} \cup \mathcal{A}_{\text{short}}$ reduces signal MSE but plateaus in Chamfer and converges to a fragmented interface that fails to recover the opening. In contrast, a staged schedule $\mathcal{A}_{\text{long}} \rightarrow \mathcal{A}_{\text{short}}$ progressively refines the boundary (iter 50 $\rightarrow$ 250 $\rightarrow$ final), yielding the correct topology (CD-L2 1.37, BE% 22.78). This ablation highlights an optimization failure mode when combining acquisition regimes that probe complementary diffusion physics: jointly fitting all regimes can produce gradient conflict and lead to local minima. Staging instead provides



Config	CD-L2 ↓	BE% ↓
Baseline	21.1 ± 2.8	66.3 ± 4.8
+ Cont.	5.0 ± 2.2	45.4 ± 8.8
+ Cont./Man.	3.9 ± 1.7	27.5 ± 7.7

Fig. 4. The left-hand plot shows average CD-L2 loss over 10 meshes for (i) baseline method (no regularization), (ii) baseline with continuity regularization and (iii) the full method with continuity/manifold regularization. The table shows final chamfer distance and BadEdge% values for each test configuration.




Fig. 5. Comparison with learned predictors on single-axon (A) and two-axon (B) meshes. Spinverse recovers compact boundaries; baselines produce many spurious faces.

a coarse-to-fine path that first localizes barriers under long diffusion time and then sharpens geometry under short diffusion time.

Ablations on regularization. Fig. 4 shows that adding continuity regularization substantially improves geometric recovery, reducing CD-L2 from 21.1 ± 2.8 to 5.0 ± 2.2 and decreasing BE% from 66.3 ± 4.8 to 45.4 ± 8.8 . Adding the manifold regularizer further improves interface structure, reducing BE% to 27.5 ± 7.7 while also improving CD-L2 to 3.9 ± 1.7 . Overall, continuity primarily drives the large gain in geometric accuracy, while the manifold term provides the largest additional reduction in non-manifold artifacts without sacrificing CD-L2. In absolute terms, BE% remains non-zero (e.g., $\approx 27.5\%$), reflecting that the thresholded face set can form open or branching interface complexes. In practice we observed a tradeoff: overly strong manifold regularization can suppress necessary boundary changes and degrade geometric recovery, whereas the chosen weight provides a favorable balance, improving interface validity without increasing CD-L2.

Comparison to learned predictors. We compare Spinverse against amortized baselines: an MLP [19] with learned encodings and two GNN variants (GraphSAGE [7] and GATv2 [4] with jumping knowledge [22]) that predict face-level barrier probabilities in a single forward pass. All networks are trained with Focal Loss [12] and FiLM conditioning [18] on 1000 synthetic bent-cylinder

Table 1. CD-L2 and BE% (mean $\pm$ std) across 100 test reconstructions each.

Geometry	Metric	MLP	GraphSAGE	GATv2+JK	Spinverse
1 Axon	CD-L2 $\downarrow$	10.81 $\pm$ 5.95	10.75 $\pm$ 6.26	8.91 $\pm$ 4.25	4.19$\pm$1.57
	BE% $\downarrow$	67.12 $\pm$ 6.47	68.18 $\pm$ 8.56	71.07 $\pm$ 8.91	23.40$\pm$9.10
2 Axon	CD-L2 $\downarrow$	8.28$\pm$2.42	9.03 $\pm$ 2.63	8.42 $\pm$ 2.52	9.79 $\pm$ 3.64
	BE% $\downarrow$	73.81 $\pm$ 4.25	75.76 $\pm$ 3.65	72.77 $\pm$ 5.82	35.62$\pm$5.64

meshes (single-axon and two-axon crossings, Fig. 2) and evaluated on 100 held-out meshes; Spinverse runs 400 iterations of per-instance optimization without any training data. Table 1 reports CD-L2 and BE%. On single-axon meshes, Spinverse roughly halves CD-L2 (4.19 vs. 8.91) and reduces BE% by $\sim 3\times$ (23.4% vs. 67–71%). On the harder two-axon setting, baselines achieve marginally lower CD-L2 (8.28–9.03 vs. 9.79) but at the cost of $\sim 2\times$ higher non-manifold artifacts (BE% 72–76% vs. 35.6%); the CD-L2 advantage is consistent with over-prediction of barrier faces, which improves point-cloud coverage but degrades structural validity (Fig. 5B). Fig. 5 confirms this qualitatively: baselines exhibit spurious faces and fragmented structure, whereas Spinverse recovers coherent boundaries.

Discussion. Spinverse performs most reliably on single-compartment geometries, where the long $\rightarrow$ short diffusion-time curriculum supports a coarse-to-fine refinement, recovering the dominant boundary structure before sharpening fine-scale barrier details. Continuity and manifold priors further regularize the ill-posed inverse problem, improving both geometric accuracy (CD-L2) and interface validity (BE%). BE% remains non-zero in several cases, and additional post-processing or stronger topology constraints may be required. We find that multi-compartment meshes remain challenging: reconstructions collapse toward symmetric envelopes that exploit the rotational degeneracy of the diffusion signal. This regime is inherently more ambiguous, as a larger set of interfaces must be inferred from the same limited acquisition set, exacerbating non-identifiability. Moreover, regularization strengths selected for single-compartment geometries may not transfer; stronger continuity/manifold priors may over-simplify the interface configuration. In terms of compute, each reconstruction takes ~ 33 minutes on a single 24GB NVIDIA L4 GPU and the solver’s $O(n^2)$ memory scaling per element currently limits the mesh resolution that can be handled. Finally, experiments are limited to controlled synthetic voxels with known forward-model parameters, and robustness to noise remains to be evaluated.

5 Conclusion

We presented Spinverse, a physics-informed method that reconstructs microstructural boundaries from diffusion MRI by optimizing face-wise permeabilities on a fixed tetrahedral mesh through a differentiable Bloch–Torrey simulator. Our experiments show that mesh-based geometric priors and staged multi-sequence curricula are critical for recovering coherent, structurally valid interfaces, and that

the physics-based inversion produces cleaner boundaries than learned baselines without requiring training data. Current evaluation is limited to synthetic geometries, and improving Spinverse on multi-compartment settings remains an open challenge. Future work will explore hybrid learned initialization with physics-informed refinement, more scalable solvers to enable finer mesh resolutions, task-driven sequence design that optimizes acquisition schedules to probe specific microstructural features, and validation on experimentally acquired dMRI data.

Acknowledgments. This work used computing resources provided by the National Research Platform (NRP) Nautilus cluster. We thank the NRP team for their support. We also thank Sanghyun Son and Yi Zhou, authors of [20] for helpful discussions during the early stages of this project. This work was supported by a UCSC Seed Grant, "An Open-Source Platform for Automated Culture and MRI Analysis of 3D Biological Tissues."

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

References

1. Alexander, D.C., Dyrby, T.B., Nilsson, M., Zhang, H.: Imaging brain microstructure with diffusion mri: practicality and applications. *NMR in Biomedicine* **32**(4), e3841 (2019). <https://doi.org/10.1002/nbm.3841>
2. Belkin, M., Niyogi, P., Sindhvani, V.: Manifold regularization: A geometric framework for learning from examples (07 2004)
3. Boykov, Y., Funka-Lea, G.: Graph cuts and efficient nd image segmentation. *International Journal of Computer Vision - IJCV* **70**, 109–131 (11 2006). <https://doi.org/10.1007/s11263-006-7934-5>
4. Brody, S., Alon, U., Yahav, E.: How attentive are graph attention networks? *CoRR* **abs/2105.14491** (2021), <https://arxiv.org/abs/2105.14491>
5. Fan, H., Su, H., Guibas, L.J.: A point set generation network for 3d object reconstruction from a single image. *CoRR* **abs/1612.00603** (2016), <http://arxiv.org/abs/1612.00603>
6. Fang, C., Yang, Z., Wassermann, D., Li, J.R.: A simulation-driven supervised learning framework to estimate brain microstructure using diffusion MRI. *Medical Image Analysis* **90**, 102979 (2023). <https://doi.org/10.1016/j.media.2023.102979>
7. Hamilton, W., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 30. Curran Associates, Inc. (2017), https://proceedings.neurips.cc/paper_files/paper/2017/file/5dd9db5e033da9c6fb5ba83c7a7e9-Paper.pdf
8. Khole, P.: Differentiable Diffusion Magnetic Resonance Imaging Simulator. Ph.D. thesis, University of California, Santa Cruz (2024), <https://escholarship.org/uc/item/77b250t7>, proQuest ID: Khole_ucsc_0036N_13189. Merritt ID: ark:/13030/m5f01jnx
9. 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 (2025), <https://arxiv.org/abs/2502.01988>

10. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization (2017), <https://arxiv.org/abs/1412.6980>
11. Li, J.R., Nguyen, V.D., Tran, T.N., Valdman, J., Trang, C.B., Nguyen, K.V., Vu, D.T.S., Tran, H.A., Tran, H.T.A., Nguyen, T.M.P.: Spinductor: A MATLAB toolbox for diffusion MRI simulation. *NeuroImage* **202**, 116120 (2019). <https://doi.org/10.1016/j.neuroimage.2019.116120>
12. Lin, T., Goyal, P., Girshick, R.B., He, K., Dollár, P.: Focal loss for dense object detection. *CoRR* **abs/1708.02002** (2017), <http://arxiv.org/abs/1708.02002>
13. Loshchilov, I., Hutter, F.: Fixing weight decay regularization in adam. *CoRR* **abs/1711.05101** (2017), <http://arxiv.org/abs/1711.05101>
14. Novikov, D.S., Fieremans, E., Jespersen, S.N., Kiselev, V.G.: Quantifying brain microstructure with diffusion mri: Theory and parameter estimation. *NMR in Biomedicine* **32**(4), e3998 (2019). <https://doi.org/10.1002/nbm.3998>
15. 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>
16. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library (2019), <https://arxiv.org/abs/1912.01703>
17. Paulsen, R., Bærentzen, A., Larsen, R.: Markov random field surface reconstruction. *IEEE transactions on visualization and computer graphics* **16**, 636–46 (09 2010). <https://doi.org/10.1109/TVCG.2009.208>
18. Perez, E., Strub, F., de Vries, H., Dumoulin, V., Courville, A.C.: Film: Visual reasoning with a general conditioning layer. *CoRR* **abs/1709.07871** (2017), <http://arxiv.org/abs/1709.07871>
19. Rumelhart, D.E., Hinton, G.E., Williams, R.J.: Learning representations by back-propagating errors. *Nature* **323**, 533–536 (1986). <https://doi.org/10.1038/323533a0>
20. Son, S., Gadelha, M., Zhou, Y., Fisher, M., Xu, Z., Qiao, Y.L., Lin, M.C., Zhou, Y.: Dmesh++: An efficient differentiable mesh for complex shapes (2025), <https://arxiv.org/abs/2412.16776>
21. Torrey, H.C.: Bloch equations with diffusion terms. *Phys. Rev.* **104**, 563–565 (Nov 1956). <https://doi.org/10.1103/PhysRev.104.563>, <https://link.aps.org/doi/10.1103/PhysRev.104.563>
22. Xu, K., Li, C., Tian, Y., Sonobe, T., Kawarabayashi, K., Jegelka, S.: Representation learning on graphs with jumping knowledge networks. *CoRR* **abs/1806.03536** (2018), <http://arxiv.org/abs/1806.03536>
23. Zhang, H., Schneider, T., Wheeler-Kingshott, C.A., Alexander, D.C.: Noddi: Practical in vivo neurite orientation dispersion and density imaging of the human brain. *NeuroImage* **61**(4), 1000–1016 (2012). <https://doi.org/10.1016/j.neuroimage.2012.03.072>
24. Zhang, Y., Brady, M., Smith, S.: Segmentation of brain mr images through a hidden markov random field model and the expectation maximization algorithm. *IEEE transactions on medical imaging* **20**, 45–57 (02 2001). <https://doi.org/10.1109/42.906424>