GPU Accelerated Modeling of Cortical Radial and Tangential Connectivity Changes in Neurodegeneration

Hongbo Zhang[1,2], Xinyu Nie[1,2], Jiaxin Yue[1,2], Yuan Li[1,2], John Ringman[3], and Yonggang Shi[1,2,3]

[1]Stevens Neuroimaging and Informatics Institute, Keck School of Medicine; [2]Ming Hsieh Department of Electrical and Computer Engineering, Viterbi School of Engineering; [3]Department of Neurology, Keck School of Medicine; University of Southern California, Los Angeles CA 90007, USA
`yshi@loni.usc.edu`

Abstract. The diffusion MRI signals in the human cerebral cortex are strongly associated with neurodegenerative diseases. Although models like NODDI have been extensively used to characterize cortical microstructure degeneration, they fall short in capturing detailed, orientation-specific connectivity changes within the cortex. In this study, we introduce a method to decompose cortical tissue diffusion signal to radial and tangential components. Our approach uses data from multi-shell diffusion imaging and combines it with anatomical information from brain surfaces. By applying a GPU accelerated probabilistic optimization framework, we can accurately and efficiently estimate these diffusion components while keeping the results smooth and consistent with the cortical anatomy. We test our method on data from HCP subjects and a clinical dataset of patients with autosomal dominant Alzheimer's Disease (ADAD) subjects. Our results demonstrate that the proposed method can more effectively reveal cortical gray matter connectivity changes related to tau pathology than metrics from the NODDI model. Our codebase is publicly available at https://github.com/Haibaobob/FOD-ctx-decomp.

Keywords: Diffusion MRI · Cerebral cortex · Radial fiber orientation.

1 Introduction

Traditional diffusion MRI (dMRI) research has primarily concentrated on deep white matter (DWM), where strong anisotropic signals facilitate fiber tracking, but the study of cortical gray matter (GM) remains challenging due to its inherently isotropic diffusion characteristics [1]. In recent years, several innovative techniques [2–6] have been developed to tackle the challenges posed by GM, where fiber orientations are predominantly radial or tangential relative to the cortical surface. Typically, radial fibers are considered extensions of myelinated axons from the DWM, whereas tangential fibers, which run within the cortical laminae, are thought to represent intracortical pathways [3, 4]. However, a comprehensive method for accurately decomposing diffusion signals into these distinct radial and tangential components is still lacking.

Although Neurite Orientation Dispersion and Density Imaging (NODDI) [7] offers detailed insights into brain microstructure, particularly in gray matter [5], it only studies the general density and orientation information, ignoring the radial and tangential differences. Earliest study on such differences employed the diffusion tensor model to analyze intra-cortical signals, effectively characterizing regional differences in radial and tangential diffusion components [3]. However, because the tensor model cannot resolve fiber crossings—a ubiquitous phenomenon in cortical gray matter [4]—its applicability is limited. To overcome this, [4] utilized the fiber orientation distribution (FOD) function, which can handle crossing fibers, and demonstrated distinct FOD structures across different cortical laminae via lamina-wise analysis. However, their radial component only relies on the voxel-wise normal direction and ignores the topographic regularity. Given that radial and tangential signals arise from different tissue types and have unique roles in neurodegeneration and neural development, accurately decomposing these components is essential for a comprehensive understanding of brain microstructure.

In this paper, we model the decomposition of radial and tangential connectivity as a discrete optimization problem. To address its computational complexity, we employ a probabilistic relaxation and develop a GPU-accelerated framework. Experiments on the HCP dataset show that our method efficiently and accurately decomposes the radial and tangential peak fields. Moreover, our results on an Alzheimer's Disease dataset demonstrate that these indicators offer distinct advantages for neurodegeneration detection compared to NODDI metrics.

2 Method

2.1 Free-water-aware FOD Reconstruction and Peak Extraction

Diffusion MRI aggregates various diffusion signals within each voxel, complicating the accurate characterization of distinct diffusion phenomena such as intra-axon water, free water, and other interference sources. Thus, we employ a fiber orientation distribution (FOD) model that incorporates compartment-specific information and leverages multi-shell dMRI data [8].

Since the model in [8] was originally developed for white matter, we have modified it to include a free-water compartment to account for the influence of cerebrospinal fluid (CSF) surrounding gray matter. With this modification, the diffusion signal $s(b, \mathbf{u})$ of a given b-value b and a gradient direction is modeled using spherical convolution:

$$s(b, \mathbf{u}) = (f *_s k)(\mathbf{u}) + \alpha e^{-b\lambda_{iso}} + \beta e^{-b\lambda_{FW}} + \eta, \tag{1}$$

where f represents the FOD to be deconvoluted, and k is a stick kernel function representing the intra-axonal diffusion. The parameters α and β represent the volume fractions of isotropic extra-axonal diffusion and free-water diffusion, respectively. The diffusivity of the isotropic extra-axonal component is denoted by an unknown parameter λ_{iso}, while λ_{FW} is fixed to 0.003 $\mu\text{m}^2/\text{s}$ to represent the free-water diffusivity. The final term η accounts for measurement noise. The sum of all compartment fractions satisfies the normalization condition [8].

We reconstruct the free-water-aware FODs using the same constrained energy minimization algorithm in [8]. The resulting FODs are represented as functions $f_{\text{FOD}} : \mathbb{S}^2 \to \mathbb{R}^+$, where $\mathbb{S}^2$ represents the unit sphere. To facilitate further analysis, we extract the local maxima of these FODs, which indicate the dominant fiber orientations within each voxel. Since the number of FOD peaks can vary across voxels, we standardize the representation by fitting all peaks into a 5D tensor $\boldsymbol{S}$ with a maximum of m peaks. The FOD peaks are sorted by magnitude, and if a voxel contains fewer than m peaks, the remaining positions are padded with zeros. In all our experiments, we set m=5.

2.2 Estimate Radial and Tangential Directions in Cortical Voxels

We aim to decompose the FOD function into tangential and radial components relative to the white matter (WM) and pial surfaces. In [9], the FOD function was projected onto the deformed WM surface to track U-fibers, but this approach did not yield the radial component. In [10], a principle FOD peak direction was selected in each voxel—by solving a Markov Random Fields (MRF) problem—to follow the fiber bundle's trajectory.

Problem Formulation Inspired by [10], we posit that the radial and tangential FOD peak directions vary smoothly across the cortex and can be effectively modeled using an MRF framework. Given voxel-wise radial peak directions $\boldsymbol{R}_i$ and principle peak tangential directions $\boldsymbol{T}_i$, and acknowledging that either $\boldsymbol{R}_i$ or $\boldsymbol{T}_i$ may be absent in some voxels, the optimal peak vectors are obtained by minimizing the following loss function:

$$\mathcal{L} = \sum_{i \in \Omega} (\mathcal{L}_u(i) + \lambda_p \sum_{j \in \mathcal{N}(i)} \mathcal{L}_p(i,j)), \tag{2}$$

where Ω is the set of cortical voxels, and $\mathcal{N}(i)$ denotes the set of voxels spatially adjacent to voxel i. The strength of the pairwise smoothness regularization is controlled by the positive weighting factor λ_p, which can be adjusted based on the specific dataset and the desired level of smoothness.

Let $\langle\ ,\ \rangle$ denote the cosine similarity between vectors and $\delta(\cdot)$ denote the indicator function. We further define the unary loss function $\mathcal{L}_u$ and the pairwise loss $\mathcal{L}_p$:

$$\mathcal{L}_u(i) = -\frac{\delta(\boldsymbol{R}_i \neq \boldsymbol{0})}{|\Omega_R|}\,|\langle \boldsymbol{R}_i, \boldsymbol{N}_i \rangle| + \frac{\delta(\boldsymbol{T}_i \neq \boldsymbol{0})}{|\Omega_T|}\,|\langle \boldsymbol{T}_i, \boldsymbol{N}_i \rangle| - (\frac{\|\boldsymbol{R}_i\|_2}{|\Omega_R|} + \frac{\|\boldsymbol{T}_i\|_2}{|\Omega_T|}), \tag{3}$$

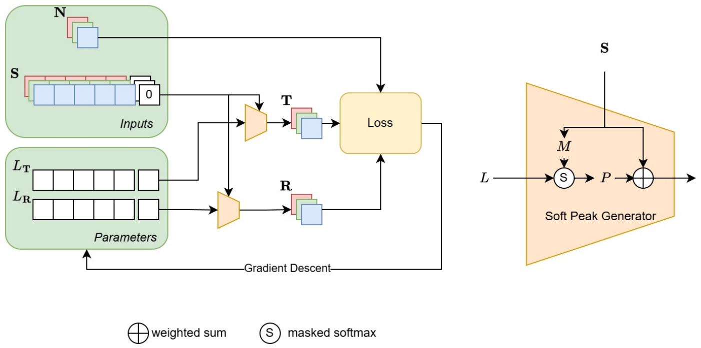

Fig. 1. GPU accelerated optimization framework. Each rectangle in the diagram represents a 3D tensor. The optimization pipeline operates as follows: During each forward pass, the **Soft Peak Generator** (orange trapezoid) generates soft peak tensors $\boldsymbol{R}$ (radial) and $\boldsymbol{T}$ (tangential) by applying a masked softmax to the logit tensors $L_{\boldsymbol{R}}$, $L_{\boldsymbol{T}}$, followed by a probability-weighted sum with the peak tensor S. These soft peaks, along with the interpolated normal direction $\boldsymbol{N}$, are used to compute the loss. The computed loss then backpropagates gradients to update the logit tensors $L_{\boldsymbol{R}}$ and $L_{\boldsymbol{T}}$ through gradient descent.

$$\begin{aligned}\mathcal{L}_p(i,j) = &- \frac{\delta(\boldsymbol{R}_i, \boldsymbol{R}_j \neq \mathbf{0})}{|\mathcal{E}_R|} \left|\langle \boldsymbol{R}_i, \boldsymbol{R}_j \rangle\right| \\ &- \frac{\delta(\boldsymbol{R}_i, \boldsymbol{R}_j \neq \mathbf{0})}{|\mathcal{E}_R|} \left|\langle \boldsymbol{R}_i - \operatorname{proj}_{\boldsymbol{T}_i} \boldsymbol{R}_i, \boldsymbol{R}_j - \operatorname{proj}_{\boldsymbol{T}_j} \boldsymbol{R}_j \rangle\right|,\end{aligned} \tag{4}$$

where $\boldsymbol{N}_i$ denotes the normal at the center of voxel i, obtained by interpolating the normals from the pial and WM surfaces. $|\Omega_R| = \sum_{i \in \Omega} \delta(\boldsymbol{R}_i \neq \mathbf{0})$ and $|\Omega_T| = \sum_{i \in \Omega} \delta(\boldsymbol{T}_i \neq \mathbf{0})$ counts non-zero vectors. $|\mathcal{E}_R| = \sum_{(i,j) \in \mathcal{E}} \delta(\boldsymbol{R}_i, \boldsymbol{R}_j \neq \mathbf{0})$ counts adjacent pairs with non-zero radial vectors. $\operatorname{proj}_{\boldsymbol{T}_i} \boldsymbol{R}_i$ denotes the projection of $\boldsymbol{R}_i$ onto $\boldsymbol{T}_i$.

The unary loss function $\mathcal{L}_u$ is designed to minimize the average cosine similarity between non-zero $\boldsymbol{R}_i$ and the corresponding normals $\boldsymbol{N}_i$, while maximize the average similarity between non-zero $\boldsymbol{T}_i$ and $\boldsymbol{N}_i$. The last term in $\mathcal{L}_u$ penalizes small magnitudes of $\boldsymbol{R}_i$ and $\boldsymbol{T}_i$ to suppress noise contributions. In contrast, the pairwise loss $\mathcal{L}_p$ regularizes the spatial smoothness by averaging the cosine similarity over adjacent voxel pairs, of which both vectors are non-zero.

However, this discrete modeling makes it computationally hard to iterate all possible combinations and find the global optimum. Although advanced algorithms such as belief propagation (BP) [11], as utilized in [10], can provide approximate solutions, both the memory and time requirements escalate rapidly as the graph scales to encompass the whole brain.

GPU Accelerated Optimization To address the computational inefficiencies inherent in discrete optimization, we reformulate the problem as a continuous probabilistic relaxation. More specifically, for each voxel i, we estimate the probability mass functions $P^i_{\boldsymbol{R}}$ and $P^i_{\boldsymbol{T}}$ over the candidate peaks $\mathcal{C} = \{\boldsymbol{S}^k_i | k \le m\} \cup \{\boldsymbol{0}\}$, ensuring that the selected peak minimizes the loss $\mathcal{L}$ in (2).

As of common practice in probability regression, instead of explicitly defining probability distributions, we introduce logits $L_{\boldsymbol{R}}$ and $L_{\boldsymbol{T}}$, where the last logit determines the probability of selecting zero peaks. We apply masked softmax operations with a mask tensor M to transform the logits to probability, where M mask out the padded zero peaks in $\boldsymbol{S}$. The softmax ensures that probabilities sum up to 1 while the probability of padded zeros to be 0.

The logits are first initialized based on the cosine similarity between peaks in $\mathcal{C}$ and $\boldsymbol{N}_i$. During each forward step, soft peaks $\boldsymbol{R}_i = \sum_{k \le m} P^{i,k}_{\boldsymbol{R}} \boldsymbol{S}^k_i$ and $\boldsymbol{T_i} = \sum_{k \le m} P^{i,k}_{\boldsymbol{T}} \boldsymbol{S}^k_i$ are generated by the probability-weighted sum of peaks. The soft peaks are not necessarily from $\mathcal{C}$, which may generate spurious peaks that well minimize the loss $\mathcal{L}$. Therefore, we regularize the entropy $H_{\boldsymbol{R}}$ and $H_{\boldsymbol{T}}$, which condense the probability distribution to a single candidate as optimization iterates. Moreover, the Jensen-Shannon divergence between $P_{\boldsymbol{R}}$ and $P_{\boldsymbol{T}}$ is maximized to avoid selecting the same peak as both radial and tangential direction. The overall optimization problem is reformulated as:

$$\min_{L_{\boldsymbol{R}}, L_{\boldsymbol{T}}} \mathcal{L} + \lambda_{\text{reg}} \sum_{i \in \Omega} (H^i_{\boldsymbol{R}} + H^i_{\boldsymbol{T}} - \text{JS-Div}(P^i_{\boldsymbol{R}}, P^i_{\boldsymbol{T}})), \tag{5}$$

where λ_{reg} is a regularization parameter, and $P^i_{\boldsymbol{R}}$ and $P^i_{\boldsymbol{T}}$ are the masked softmax of $L^i_{\boldsymbol{R}}$ and $L^i_{\boldsymbol{T}}$ respectively.

The entire optimization process—illustrated in Fig. 1—is implemented using differentiable tensor operations. This design enables us to leverage the PyTorch library for GPU acceleration. Once convergence is achieved, the candidate peaks with the highest probabilities are selected as the radial and tangential directions.

Given the estimated directions $\boldsymbol{R}_i$ and $\boldsymbol{T}_i$, we quantify the radial and tangential components $\mathcal{R}_i$ and $\mathcal{T}_i$ from the FOD functions through spherical integrations over a given threshold θ. For example, $\mathcal{R}_i$ can be computed as follows:

$$\mathcal{R}_i = \begin{cases} \displaystyle\int_{\{\langle \boldsymbol{A}, \boldsymbol{R}_i \rangle > \cos(\theta)\}} f^i_{\text{FOD}} \, d\boldsymbol{A}, & \text{if } \boldsymbol{R}_i \neq \boldsymbol{0} \\ 0, & \text{otherwise} \end{cases} \tag{6}$$

The tangential component $\mathcal{T}_i$ can be computed similarly.

2.3 Implementation Details

The optimization was implemented in PyTorch using the Adam optimizer with a learning rate of 0.1, $\beta_1 = 0.9$, and $\beta_2 = 0.999$. The smoothness regularization weight λ_p was set to 0.5 for the HCP dataset and 0.1 for the ADAD dataset. The entropy regularization weight λ_{reg} was fixed at 1×10^{-6} for all experiments. The number of optimization steps was set to 500 for HCP and 300 for ADAD.

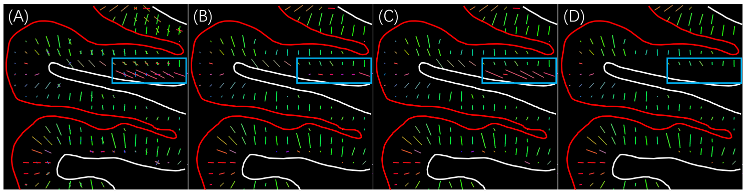

Fig. 2. Qualitative comparison of radial peaks of HCP subject 100408. (A) All extracted FOD peaks. (B) Radial directions from NN. (C) Radial directions from BP. (D) Radial directions by the proposed method. Red and white curves represent the pial and white matter surfaces, respectively. Voxels in blue boxes indicate that the proposed method identifies tangential peak field and excludes those peaks from the radial peaks.

3 Experiments and Analysis

3.1 Data Preprocessing

Data preprocessing begins with structural processing using Freesurfer's T1 recon-all pipeline [12]. Multi-shell diffusion data is then employed to reconstruct free-water-aware FODs. To mitigate distortions in cortical regions, the T1 image is nonlinearly registered to the dMRI space using ANTs software [13]. Finally, the cortical surfaces extracted from the T1 image are warped into the dMRI space using the computed deformation field and affine transformation, yielding accurate representations of the pial and white matter boundaries in dMRI space.

3.2 Method Validation

To evaluate our proposed method, we compare with two baseline approaches: the nearest neighbor (NN) method and the belief propagation (BP) based algorithm [10]. The experiments were conducted on a server equipped with an NVIDIA(R) A6000 GPU and an Intel(R) Xeon(R) Gold 5317 CPU.

In this experiment, we use the Human Connectome Project (HCP) dataset [14]. The HCP dMRI data was acquired from 270 gradient directions distributed over three different b values: 1000, 2000, and 3000 s/mm^2. The resolution of dMRI is $1.25mm$ isotropic.

The NN method selects the peak that is closest to the interpolated normal direction in each voxel; however, it often produces noisy and inconsistent estimates. Meanwhile, the BP algorithm yields smoother estimates, but it is hampered by high computational costs and memory demands. Moreover, both baseline methods struggle to distinguish tangential peaks from radial peaks in voxels where the radial peak is absent. In contrast, our proposed approach reduces false positives and delivers smoother, more consistent results, as demonstrated in Fig. 2.

We evaluate performance by quantifying the smoothness of the radial peak field and its radial component. Specifically, we compute the mean Cosine Similarity (mCS) between neighboring radial peaks and the mean Scale-Invariant Gradient (mSIG) of the radial component, given by $\frac{\|\nabla\mathcal{R}\|_2}{|\mathcal{R}|}$. The results on 72

Table 1. Quantitative Smoothness Evaluation

Method	**mCS**	**mSIG**	**Average Time (sec)**
NN	0.886 (0.005)	0.950 (0.083)	0.958 (0.090)
BP	0.898 (0.004)	0.770 (0.042)	3693.463 (868.632)
Ours	**0.916** (0.005)	**0.692** (0.047)	89.462 (1.040)

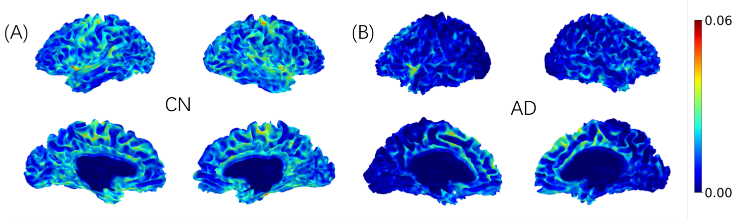

Fig. 3. The radial component profile on surfaces in dMRI space. The AD subject (B) shows widespread intensity reduction especially in temporal and parietal lobes compared to the CN subject (A).

HCP subjects, presented in Table 1, demonstrate that our method outperforms the baselines in terms of smoothness while maintaining acceptable computational efficiency.

3.3 Application to Neurodegeneration Analysis

We applied our method to a cohort of 30 autosomal dominant Alzheimer's disease (ADAD) subjects carrying the A431E mutation [15], including 19 cognitively normal (CN) individuals and 11 patients with Alzheimer's disease (AD). Following the HCP protocol, T1-weighted MRI and multi-shell dMRI were acquired on a 3T Siemens Prisma scanner.

For comparison, we computed the Neurite Density Index (NDI) and Orientation Dispersion Index (ODI) from the NODDI model for cortical gray matter following the method in [16]. To mitigate potential partial volume effects, both the proposed and NODDI volume features were interpolated onto surface profiles by averaging the interpolated volume intensities at three equally spaced surfaces (25%, 50%, and 75%) between white matter and pial surfaces. Fig. 3 presents the radial component projected onto the cortical surface, highlighting regional differences in radial tissue diffusivity between AD and CN subjects.

The analysis was conducted at both the ROI and vertex levels. In the ROI-wise analysis, surface image profiles were averaged within each cortical ROI as defined by Freesurfer, providing a region-level summary of diffusion characteristics and facilitating comparisons across anatomically meaningful regions. In contrast, the vertex-wise analysis involved interpolating the surface profiles onto a common template surface, which enabled a finer-grained assessment of localized diffusion abnormalities associated with neurodegeneration.

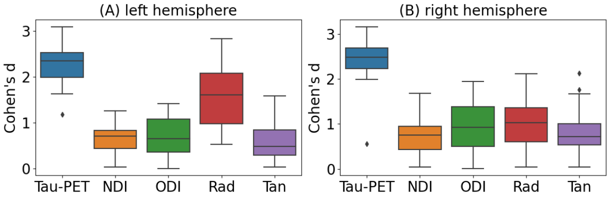

Fig. 4. Cohen's d across 36 ROIs of left hemisphere (A) and right hemisphere (B). The radial component (Rad) yields higher effect sizes, especially in left hemishpere. Paired one-tailed t-tests show Rad significantly outperforms NDI and ODI on the left ($p^* < 0.001$), and NDI on the right ($p^* = 0.034$); other differences were not significant. p^* denote Bonferroni correction.

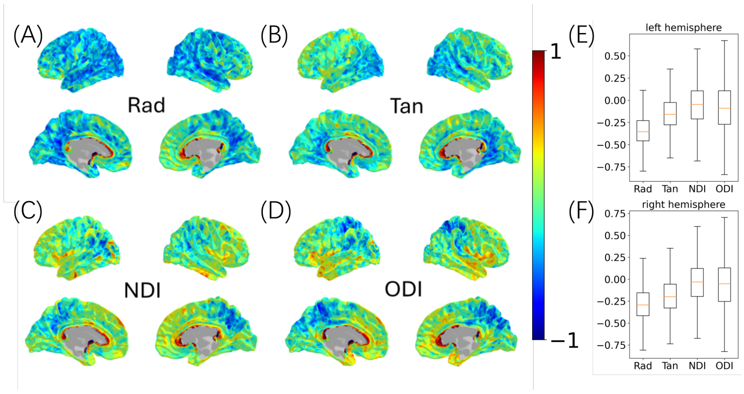

Fig. 5. Correlation coefficients (CC) between tau PET SUVR and cortical connectivity metrics. (A,B) CC of tau PET SUVR with the radial (Rad) and tangential (Tan) component of our method. (C, D) CC of tau PET SUVR with the NDI and ODI metrics of NODDI. (E,F): Boxplots of vertex-wise CCs on the left and right hemisphere.

Cohen's *d* between AD and CN In this experiment, we compared the Cohen's d effect size between AD and CN groups across 36 ROIs (Fig. 4). The radial component identified by our method exhibited a higher effect size than NODDI. Notably, the results revealed the hemispheric asymmetries in ADAD pathology: the left hemisphere showed more pronounced radial degeneration, while the right hemisphere exhibited more uniform degeneration. These nuanced observations were not attainable with the NODDI model alone.

Correlation with Tau-PET Imaging For 30 subjects, Tau PET scans were acquired using the 18F-AV-1451 tracer. We followed the PET-Surfer processing pipeline [17] to compute the standard uptake value ratios (SUVRs) of tau PET signals in cortical areas. Fig. 5 shows our radial component demonstrated a higher negative correlation with Tau-PET SUVR compared to NODDI($p^* < 0.001$, Bonferroni-corrected, paired t-tests), especially in the parietal, temporal and frontal lobes. This suggests that tau deposition may be more strongly associated with radial-wise degeneration in these regions, providing potential biomarkers for neurodegeneration detection.

4 Conclusions

In this work, we developed a GPU accelerated algorithm for decomposing the radial and tangential connectivity in cortical gray matter. Validated using data from HCP subjects and an ADAD cohort, we demonstrate that our method can more effectively detect cortical connectivity changes than the NODDI approach. We also showed that the radial components have much stronger association with tau pathology than the tangential and NODDI metrics, which suggests the potential of the radial connectivity in cortical gray matter as a novel biomarker for neurodegeneration.

Acknowledgments. This work was supported by the National Institute of Health (NIH) under grants R01EB022744, RF1AG077578, RF1AG064584, U19AG078109, and P30AG066530.

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

References

1. Dmitry S Novikov, Els Fieremans, Sune N Jespersen, and Valerij G Kiselev. Quantifying brain microstructure with diffusion mri: Theory and parameter estimation. *NMR in Biomedicine*, 32(4):e3998, 2019.
2. Alfred Anwander, André Pampel, and Thomas R Knösche. In vivo measurement of cortical anisotropy by diffusion-weighted imaging correlates with cortex type. In *ISMRM Joint Annual Meeting*, 2010.
3. Jennifer A McNab, Jonathan R Polimeni, Ruopeng Wang, Jean C Augustinack, Kyoko Fujimoto, Allison Stevens, Thomas Janssens, Reza Farivar, Rebecca D Folkerth, Wim Vanduffel, et al. Surface based analysis of diffusion orientation for identifying architectonic domains in the in vivo human cortex. *Neuroimage*, 69:87–100, 2013.
4. Christoph WU Leuze, Alfred Anwander, Pierre-Louis Bazin, Bibek Dhital, Carsten Stüber, Katja Reimann, Stefan Geyer, and Robert Turner. Layer-specific intracortical connectivity revealed with diffusion mri. *Cerebral cortex*, 24(2):328–339, 2014.
5. Hikaru Fukutomi, Matthew F Glasser, Hui Zhang, Joonas A Autio, Timothy S Coalson, Tomohisa Okada, Kaori Togashi, David C Van Essen, and Takuya Hayashi. Neurite imaging reveals microstructural variations in human cerebral cortical gray matter. *Neuroimage*, 182:488–499, 2018.
6. Yaniv Assaf. Imaging laminar structures in the gray matter with diffusion mri. *Neuroimage*, 197:677–688, 2019.
7. 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.

8. Giang Tran and Yonggang Shi. Fiber orientation and compartment parameter estimation from multi-shell diffusion imaging. *IEEE Transactions on Medical Imaging*, 34(11):2320–2332, 2015.
9. Xinyu Nie, Jialiang Ruan, Maria Concepción Garcia Otaduy, Lea Tenenholz Grinberg, John Ringman, and Yonggang Shi. Surface-based probabilistic fiber tracking in superficial white matter. *IEEE transactions on medical imaging*, 43(3):1113–1124, 2023.
10. Xinyu Nie and Yonggang Shi. Topographic filtering of tractograms as vector field flows. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part III 22*, pages 564–572. Springer, 2019.
11. Pedro F Felzenszwalb and Daniel P Huttenlocher. Efficient belief propagation for early vision. *International journal of computer vision*, 70:41–54, 2006.
12. Bruce Fischl. Freesurfer. *Neuroimage*, 62(2):774–781, 2012.
13. Brian B Avants, Nicholas J Tustison, Gang Song, Philip A Cook, Arno Klein, and James C Gee. A reproducible evaluation of ants similarity metric performance in brain image registration. *Neuroimage*, 54(3):2033–2044, 2011.
14. Stamatios N Sotiropoulos, Saad Jbabdi, Junqian Xu, Jesper L Andersson, Steen Moeller, Edward J Auerbach, Matthew F Glasser, Moises Hernandez, Guillermo Sapiro, Mark Jenkinson, et al. Advances in diffusion mri acquisition and processing in the human connectome project. *Neuroimage*, 80:125–143, 2013.
15. Jill Murrell, Bernardino Ghetti, Elizabeth Cochran, Miguel Angel Macias-Islas, Luis Medina, Arousiak Varpetian, Jeffrey L Cummings, Mario F Mendez, Claudia Kawas, Helena Chui, et al. The a431e mutation in psen1 causing familial alzheimer's disease originating in jalisco state, mexico: an additional fifteen families. *Neurogenetics*, 7:277–279, 2006.
16. Talia M Nir, Julio E Villalón-Reina, Lauren E Salminen, Elizabeth Haddad, Hong Zheng, Sophia I Thomopoulos, Clifford R Jack Jr, Michael W Weiner, Paul M Thompson, Neda Jahanshad, et al. Cortical microstructural associations with csf amyloid and ptau. *Molecular Psychiatry*, 29(2):257–268, 2024.
17. Douglas N Greve, Claus Svarer, Patrick M Fisher, Ling Feng, Adam E Hansen, William Baare, Bruce Rosen, Bruce Fischl, and Gitte M Knudsen. Cortical surface-based analysis reduces bias and variance in kinetic modeling of brain pet data. *Neuroimage*, 92:225–236, 2014.