

Diffusion-Informed Spatial Smoothing of fMRI Data in White Matter Using Spectral Graph Filters

David Abramian^{a,b,*}, Martin Larsson^c, Anders Eklund^{1a,b,g}, Iman Aganj^{d,e}, Carl-Fredrik Westin^f, Hamid Behjat^{h,f,d,*}

^aDepartment of Biomedical Engineering, Linköping University, Linköping, Sweden

^bCenter for Medical Image Science and Visualization, Linköping University, Linköping, Sweden

^cCentre of Mathematical Sciences, Lund University, Lund, Sweden

^dMartinos Center for Biomedical Imaging, Massachusetts General Hospital, Harvard Medical School, Boston, USA

^eComputer Science and Artificial Intelligence Lab, Massachusetts Institute of Technology, Cambridge, USA

^fDepartment of Radiology, Brigham and Women's Hospital, Harvard Medical School, Boston, USA

^gDepartment of Computer and Information Science, Linköping University, Linköping, Sweden

^hDepartment of Biomedical Engineering, Lund University, Lund, Sweden

Abstract

Brain activation mapping using functional magnetic resonance imaging (fMRI) has been extensively studied in brain gray matter (GM), whereas in large disregarded for probing white matter (WM). This unbalanced treatment has been in part due to controversies in relation to the nature of the blood oxygenation level-dependent (BOLD) contrast in WM and its detachability. However, an accumulating body of studies has provided solid evidence of the functional significance of the BOLD signal in WM and has revealed that it exhibits anisotropic spatio-temporal correlations and structure-specific fluctuations concomitant with those of the cortical BOLD signal. In this work, we present an anisotropic spatial filtering scheme for smoothing WM fMRI data that accounts for known spatial constraints on the BOLD signal in WM. In particular, the spatial correlation structure of the WM BOLD signal is highly anisotropic and closely linked to local axonal structure in terms of shape and orientation, suggesting that isotropic Gaussian filters conventionally used for smoothing fMRI data are inadequate for denoising the WM BOLD signal as they are incapable of adapting to the underlying domain of the BOLD signal in white matter. The fundamental element in the proposed method is a graph-based description of the WM that encodes the underlying anisotropy observed across WM, derived from diffusion MRI data. Based on this representation, and leveraging graph signal processing principles, we design subject-specific spatial filters that adapt to a subject's unique WM structure at each position in the WM that they are applied at. We use the proposed filters to spatially smooth WM fMRI data, as an alternative to the conventional practice of using isotropic Gaussian filters. We test the proposed filtering approach on two sets of phantoms, showcasing its greater sensitivity and specificity for the detection of slender anisotropic activations, compared to that achieved with isotropic Gaussian filters. We also present WM activation mapping results on the Human Connectome Project's 100-unrelated subject dataset, across seven functional tasks, showing the capacity of the proposed method for detecting streamline-like activations within axonal bundles.

Keywords: functional MRI, diffusion MRI, white matter, graph signal processing, anisotropy

1. Introduction

To date, reports on task-based functional magnetic resonance imaging (fMRI) activation mapping and resting-state functional connectivity have been overwhelmingly restricted to the gray matter (GM), whereas white matter (WM) functional data has been largely ignored or treated as a nuisance regressor. Such unbalanced treatment of fMRI data within GM and WM, due in part to controversies in relation to the source of the BOLD signal in WM, has led to a systematic underreporting of BOLD-related activity in WM (Mazerolle et al., 2019; Gawryluk et al., 2014).

Despite past controversies, evidence provided by an increasing body of recent studies, see e.g. Grajauskas et al. (2019) and Gore et al. (2019) and references therein, has led to more

widespread acceptance of the detectability and functional relevance of the BOLD signal in WM. For example, Ding et al. (2013) showed that resting-state BOLD signals in WM exhibit structure-specific temporal correlations along WM tracts, which coincide with fiber patterns revealed by diffusion tensor imaging (DTI), and which, under functional load, become more pronounced in functionally relevant structures (Ding et al., 2016). More specifically, Mishra et al. (2020) showed that varying experimental task parameters results in a coupled modulation of the BOLD signal in the visual cortex and relevant WM tracts, corroborating past findings of simultaneous BOLD activations in structurally-connected regions of GM and WM (Mazerolle et al., 2010). Furthermore, a growing number of recent studies have shown that low frequency BOLD fluctuations can be used to estimate the dynamic functioning of fiber tracts (Gore et al., 2019), in both health (Marussich et al., 2017; Huang et al., 2018b; Li et al., 2020b) and disease (Jiang et al.,

*Corresponding authors.

2019; Ji et al., 2019; Gao et al., 2020), providing a powerful means to study how information is transferred and integrated between functionally specialized cortices.

Due to the significantly lower vascularization density in WM compared to that in GM (Logothetis and Wandell, 2004; Jochimsen et al., 2010), the overall magnitude of the WM BOLD signal is substantially lower than that in GM (Yarkoni et al., 2009), which has been reported to be as low as 10% of that observed in GM and modulated as a function of distance from the cortical layer (Li et al., 2019b). In addition to being weak, the WM BOLD signal is affected by unique confounding factors, suggesting the need for WM-tailored acquisition and processing schemes. Broadly speaking, the BOLD contrast and its detection in WM can potentially be enhanced in three ways: i) development and use of MRI sequences optimal for WM fMRI (e.g. increased T2-weighting (Gawryluk et al., 2009) or tailored field strengths (Mazerolle et al., 2013)); ii) design of temporal models that account for the unique hemodynamic response function (HRF) in WM, which substantially differs from that in GM (Yarkoni et al., 2009; Fraser et al., 2012; Erdoğan et al., 2016)); and iii) design of spatial models that account for the unique spatial features of the BOLD contrast in WM, which is highly anisotropic (Ding et al., 2013, 2016). This paper focuses on the third category, presenting the case for the importance of filter design when handling fMRI data in WM, particularly in relation to the inherent differences between the spatial profiles of BOLD signal in WM relative to those in GM.

1.1. Spatial smoothing tailored to white matter fMRI data

Typical fMRI analysis pipelines rely on the assumption that the BOLD signal exhibits isotropic spatial profiles at focal activated regions (Carp, 2012). Isotropic Gaussian kernels applied to functional data, which is a staple of conventional fMRI analysis, is only justified under this assumption, and generally trades spatial specificity for increased sensitivity. In particular, by virtue of the matched filter argument, spatial filters are optimal only for detecting activations that conform to the size and shape of the filter kernel, and can otherwise result in loss of information regarding the spatial extent and shape of activation areas (Geissler et al., 2005; Mikl et al., 2008), obliterating all non-smooth singularities in the data.

In order to improve on the sensitivity-specificity trade-off afforded by conventional isotropic spatial smoothing, multiple smoothing methods that adapt to local spatial image features have been proposed. These include steerable filters (Knutsson et al., 1983), which enable directionally-adaptive spatial smoothing (Friman et al., 2003; Eklund et al., 2011; Zhuang et al., 2017; Abramian et al., 2020b), wavelet transforms (Mallat, 1989; Bullmore et al., 2004), which try to strike a balance between localization in space and frequency domain (Ruttimann et al., 1998; Van De Ville et al., 2004; Breakspear et al., 2006), and non-linear filters (e.g. bilateral filters) that locally adapt to various features of adjacent voxels (Smith and Brady, 1997; Rydell et al., 2008; Lohmann et al., 2018). While such methods have been successfully applied to GM, their adaptive properties rely on the spatial features manifested by the BOLD contrast. Since this contrast is substantially reduced

in WM, the effectiveness of these methods would likely be reduced when applied to fMRI data in WM.

Rather than adapting the smoothing operation to features present in the BOLD signal, alternative adaptive smoothing approaches incorporate information from the *domain* on which the data resides, typically provided by complementary anatomical images. One common approach is cortical surface smoothing, shown to provide increased sensitivity and specificity (Jo et al., 2007; Coalson et al., 2018). Such methods have also been used to formulate smoothing approaches that respect tissue boundaries (Behjat et al., 2019), preventing artifacts resulting from the mixing of signals from adjacent but differing tissue types during filtering. In both of these scenarios the anatomical information is provided by T1-weighted images.

An important distinguishing feature of the BOLD signal in WM is that it exhibits a spatial correlation structure grossly consistent with the directions of water diffusion, as measured by DTI (Ding et al., 2013), which is present during rest and becomes more pronounced under functional loading (Wu et al., 2017; Ding et al., 2018). The anatomical basis for this observation can be that up to half of the blood volume in WM resides in vessels that run in parallel to WM tracts (Doucette et al., 2019). As a consequence, conventional isotropic Gaussian filters may prove especially unsuited for the task of increasing the SNR of the BOLD signal in the highly anisotropic WM domain. Filtering methods adaptive to features of the BOLD signal may prove more effective, but the low BOLD contrast manifested in WM will potentially limit their usefulness. On the other hand, the strong anatomical dependence in the correlation structure of the BOLD signal in WM suggests that domain-informed smoothing methods can be particularly beneficial. Such methods can rely on T1-weighted images as well as diffusion-weighted MRI (DW-MRI) to adapt the filtering to the morphology and the axonal microstructure of WM, respectively. This paper presents the design and validation of such a filtering scheme.

1.2. Structure-informed processing of fMRI data through GSP

In the past five years, an increasing number of studies have showcased the use of principles from the recently emerged field of graph signal processing (GSP) within neuroimaging, in particular, in proposing intuitive methodologies for structure-informed processing of fMRI data. The fundamental idea in GSP is to analyze data recorded at a discrete set of positions in such way that the underlying structural relationship between those positions is accounted for, wherein this underlying structure can be represented in the form of a graph, i.e., a structure consisting of a set of nodes and edges. We refer the reader to Shuman et al. (2013) for an introduction to GSP and to Ortega et al. (2018) and Stanković et al. (2020) for an overview of recent developments, challenges, and applications.

An increasing number of studies have proposed the use of region of interest (ROI) based structural connectomes (Sporns et al., 2005), derived from tractography data, as underlying backbones for interpreting fMRI data (Atasoy et al., 2016; Abdelnour et al., 2018; Huang et al., 2018a). When structural connectomes are interpreted as graphs, a number of their Laplacian eigenvectors manifest spatial patterns that are reminiscent of

well-established functional networks, as shown by Atasoy et al. (2016). Under this framework, methods have been proposed for spatio-temporal deconvolution of fMRI data (Bolton et al., 2019), quantification of the coupling strength of resting-state fMRI data with underlying structure (Medaglia et al., 2018; Preti and Van De Ville, 2019), implementation of neural field models (Aqil et al., 2020), prediction of behaviorally relevant scores (Bolton and van De Ville, 2020), and for characterization of functional connectivity dynamics in health (Huang et al., 2018b), and its changes, for instance, due to concussion (Sihag et al., 2020), and under hallucinogenic drugs (Atasoy et al., 2017).

As alternatives to macro-scale ROI-based graphs, a number of voxel-wise brain graph designs have been proposed for analysis of fMRI data. Graphs encoding GM morphology have been proposed for enhanced activation mapping in GM, for both group-level (Behjat et al., 2015) and subject-level (Behjat et al., 2013, 2014) analyses, and for discriminative characterization of fMRI data across functional tasks (Behjat and Larsson, 2020). A closely related work to that presented here is by Tarun et al. (2020), in which DW-MRI data is used to encode the WM fiber structure, for the task of visualizing WM fiber pathways based on the functional activity observed at the cortical layer.

1.3. Aim and overview

To the best of our knowledge, no method has to date been presented to specifically account for the spatial features of the BOLD contrast in WM when it comes to processing WM fMRI data. The main objective of this work is to consider the importance of filter design when handling fMRI data in WM, particularly, in relation to the inherent difference between the spatial profiles of BOLD signal in WM relative to those in GM.

In this paper, we develop an adaptive spatial smoothing method tailored to the processing of fMRI data in WM. Using diffusion orientation distribution functions (ODF) obtained from high angular resolution diffusion imaging (HARDI) data, we construct subject-specific voxel-wise WM graphs. A spectral heat kernel filter is then defined on the spectrum of the resulting graphs, and implemented in a computationally efficient way for the task of fMRI data filtering, using principles from GSP. When instantiated at any position within the WM, the proposed filters adapt to the local axonal orientation, becoming consistent with the spatial correlation structure of the WM BOLD signal.

The remainder of this paper is organized as follows: in Section 2, we review relevant GSP principles and describe our proposed graph and filter designs, as well as the construction of phantoms. In Section 3, we examine the smoothing filters produced by the proposed design and evaluate their performance on two types and on real task fMRI data. We conclude the paper in Section 4 with a discussion on design consideration, limitations and future work.

2. Materials and Methods

2.1. Data and preprocessing

Data used in the preparation of this work were obtained from the WU-Minn Human Connectome Project (HCP) (Van Essen et al., 2013) database¹. We use the 100 unrelated adult subject sub-group (54% female, mean age = 29.11 ± 3.67 , age range = 22-36), which we denote as the HCP100 subject set. Five of the subjects were excluded due to incomplete WM coverage of the DW-MRI data, leaving a total of 95 subjects. The HCP data acquisition study was approved by the Washington University Institutional Review Board and informed consent was obtained from all subjects. We used the minimally preprocessed structural, task fMRI, and DW-MRI data. Task fMRI data for each subject consists of 1940 time frames across seven functional tasks: Emotion, Gambling, Language, Motor, Relational, Social, and Working Memory, comprising 23 experimental conditions in total. The method proposed in this paper heavily relies on the accurate co-registration between the structural and functional data, as provided by the minimally processed HCP data. The imaging parameters and image preprocessing steps have been fully described by Glasser et al. (2013). All data processing in this work was done using the Matlab software and the SPM12 toolbox². Diffusion ODFs were generated using the method presented by Yeh et al. (2010) and implemented in the DSI Studio software package³.

The HCP data are provided in a mixture of three spatial resolutions within two neurological spaces (ACPC, i.e., native subject space, and MNI): 0.7 mm isotropic ACPC for the structural data, 1.25 mm isotropic ACPC for the DW-MRI data, and 2 mm isotropic MNI for the fMRI data. A fundamental necessity for the proposed methodology is to reconcile the three datasets into a single set of working parameters. However, the resampling process and the nonlinear conversion between ACPC and MNI spaces have the potential of negatively affecting the data quality. The number of voxels is also a relevant parameter, since it determines to a great extent the memory usage and computation time of the various processing steps. Given the importance of axonal orientation information to the proposed method, we also prioritized minimizing the manipulations applied to the DW-MRI data.

Based on these considerations, we chose the ACPC space at the resolution of the diffusion data, i.e., 1.25 mm isotropic, as the working space. As such, the HCP preprocessed fMRI volumes were warped back into ACPC space and upsampled to the voxel resolution of the diffusion data. This mapping was done by leveraging the `mni2acpc.nii` displacement maps provided with the HCP preprocessed data, using first order splines as the basis for interpolation. In addition, the segmentation volume `aparc+aseg.nii`, computed via FreeSurfer (Fischl, 2012) and provided with the HCP data, was downsampled to the working resolution, from which voxels associated to WM were extracted.

¹<https://ida.loni.usc.edu/login.jsp>

²<https://www.fil.ion.ucl.ac.uk/spm/software/spm12/>

³<http://dsi-studio.labsolver.org>

2.2. GSP preliminaries

The fundamental idea in GSP is the application of signal processing procedures to data residing on the vertices of a graph, wherein the graph defines the underlying irregular domain of the data. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{A})$ denote an undirected, single-connected, weighted graph, defined by a vertex set \mathcal{V} of size N_g , denoting the size of the graph, an edge set \mathcal{E} consisting of connecting pairs (i, j) of vertices, and a symmetric adjacency matrix \mathbf{A} whose nonzero elements $a_{i,j}$ represent the weight of edges $(i, j) \in \mathcal{E}$. Let $\ell^2(\mathcal{G})$ denote the Hilbert space of all square-integrable graph signals $\mathbf{f} : \mathcal{V} \rightarrow \mathbb{R}$ defined on the vertex set \mathcal{V} . A graph signal $\mathbf{f} \in \ell^2(\mathcal{G})$ is in essence an $N_g \times 1$ vector, whose n -th component represents the signal value at the n -th vertex of \mathcal{G} .

The graph spectral domain, analogous to the Euclidean Fourier domain, can be defined using a graph's Laplacian matrix. In particular, the *normalized* Laplacian matrix of \mathcal{G} is defined as $\mathbf{L} = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$, where \mathbf{D} denotes the graph's degree matrix, which is diagonal with elements defined as $d_{i,i} = \sum_j a_{i,j}$. Since \mathbf{L} is real, symmetric, diagonally dominant, and with non-negative diagonal entries, it is positive semi-definite; i.e., all its N_g eigenvalues are real and non-negative, and they are also no larger than 2 due to the normalization used in the definition of \mathbf{L} . This set of eigenvalues defines the spectrum of \mathcal{G} , denoted as $\Lambda = \{0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_{N_g} \stackrel{\text{def}}{=} \lambda_{\max} \leq 2\}$. The associated eigenvectors, denoted $\{\mathbf{u}_l\}_{l=1, \dots, N_g}$, form an orthonormal basis spanning the $\ell^2(\mathcal{G})$ space. A graph signal \mathbf{f} can be transformed into a spectral representation through the use of the Laplacian eigenvectors as

$$\hat{\mathbf{f}}[l] = \sum_{n=1}^{N_g} \mathbf{u}_l[n] \mathbf{f}[n] \quad (1)$$

$$= \mathbf{u}_l^T \mathbf{f}, \quad l = 1, \dots, N_g. \quad (2)$$

This spectral representation enjoys a perfect reconstruction, that is, the signal can be recovered as $\mathbf{f} = \sum_{l=1}^{N_g} \hat{\mathbf{f}}[l] \mathbf{u}_l$.

In contrast to filters in classical signal processing, graph filters are shift-variant, adapting their shape to the underlying graph structure when localized at any given vertex. Consequently, individual filters defined on the spectral domain of a graph will become spatially-adaptive by the nature of GSP. This valuable property of graph filters enables the proposed methodology, but it also prevents the implementation of filtering operations as straightforward convolutions. Rather, in analogy to frequency-domain filtering in classical signal processing, graph signal filtering can be conveniently defined in the graph spectral domain. Given the spectral profile of a desired filter, $k(\lambda) : [0, 2] \rightarrow \mathbb{R}$, a graph signal \mathbf{f} can be filtered with $k(\lambda)$ as

$$\tilde{\mathbf{f}} = \sum_{l=1}^{N_g} k(\lambda_l) \hat{\mathbf{f}}[l] \mathbf{u}_l \quad (3)$$

$$\stackrel{(2)}{=} \sum_{l=1}^{N_g} k(\lambda_l) \mathbf{u}_l^T \mathbf{f} \mathbf{u}_l, \quad (4)$$

Implementing (4) requires the Laplacian eigenvectors, i.e., a full diagonalization of \mathbf{L} , which is impractical for large graphs,

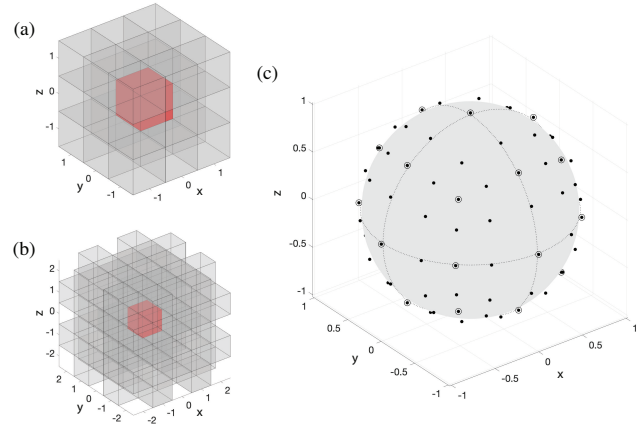


Figure 1: (a) 26 voxels within the $3 \times 3 \times 3$ neighborhood (gray) used to define edges to the focal voxel (red). (b) 98 voxels within the $5 \times 5 \times 5$ neighborhood (gray), used to define edges to the focal voxel (red). (c) Scattered dots on the unit sphere specify the 98 neighborhood directions encoded by the $5 \times 5 \times 5$ voxel neighborhood. Circled dots represent the subset of 26 directions encoded by the $3 \times 3 \times 3$ voxel neighborhood.

such as those presented in this work. An efficient alternative approach is to implement the filtering using polynomial approximations of $k(\lambda)$. We refer the interested reader to Appendix I for details on the implementation.

2.3. WM graph design

In order to take advantage of GSP tools, it is necessary to define graphs that encode relevant information in their vertices, edges, and weights. For the purpose of allowing diffusion-informed smoothing of the WM BOLD signal, we require graphs capable of encoding the subject's axonal microstructure. Filters defined on the spectral domain of such graphs will become locally adapted to this microstructure due to the shift-variant nature of graph filters.

We define a WM graph as a graph whose vertex set \mathcal{V} consists of all WM voxels, resulting in graphs with $240k \pm 60k$ vertices on the HCP100 subject set. The graph's edge set \mathcal{E} is defined on the basis of voxel adjacency, with pairs of vertices being connected to each other whenever their associated voxels are spatially neighboring. Two neighborhood definitions are considered, corresponding to cubic lattices of sizes $3 \times 3 \times 3$ (henceforth 3-conn) and $5 \times 5 \times 5$ (henceforth 5-conn), where the focal voxel is located in the center of the lattice. The 3-conn lattice specifies 26 voxels in the neighborhood of the focal voxel, whereas for the 5-conn lattice, voxels in the outer layer that fall in parallel to the voxels within the inner layer are excluded, resulting in 98 voxels in the neighborhood; see Figure 1.

The encoding of axonal microstructure by the graph is principally achieved through the edge-weighting scheme, inspired by the work of Iturria-Medina et al. (2007). The weights provide a discretization of the diffusion ODF at each point, and include information on the coherence of diffusion orientation among neighboring voxels. Let $O_i(\hat{r})$ denote the ODF associated to voxel v_i , with its coordinate origin at the voxel's center, and with \hat{r} denoting the unit direction vector. Let $\hat{r}_{i,j}$ denote

the unit vector pointing from the center of voxel v_i to the center of neighboring voxel v_j . A discretization of the ODF along direction $\hat{r}_{i,j}$ can be obtained as

$$p(i, \hat{r}_{i,j}) = \int_{\Omega_{i,j}} O_i(\hat{r}) d\Omega, \quad (5)$$

where $\Omega_{i,j}$ denotes the solid angle of $4\pi/26$ (for 3-conn) or $4\pi/98$ (for 5-conn) around $\hat{r}_{i,j}$ and $d\Omega$ denotes the infinitesimal solid angle element. This measure can be approximated by taking N_t samples of the ODF within the solid angle $\Omega_{i,j}$ as

$$p(i, \hat{r}_{i,j}) \approx \tilde{p}(i, \hat{r}_{i,j}) = \frac{1}{N_t} \sum_{k=1}^{N_t} O_i(\hat{r}_{i,j}^k), \quad (6)$$

where $\hat{r}_{i,j}^k$ denotes the k -th sampling direction within $\Omega_{i,j}$. Details of the sampling process are given in Appendix II. Furthermore, we normalize this metric as

$$q_{i,j} = \frac{\tilde{p}(i, \hat{r}_{i,j})}{2 \max_j \{\tilde{p}(i, \hat{r}_{i,j}) \mid (i, j) \in \mathcal{E}\}}, \quad (7)$$

which bounds it in the $[0, 0.5]$ range. The maximum value of 0.5 is reached if the ODF at v_i shows its maximal diffusion along $\hat{r}_{i,j}$, whereas otherwise $q_{i,j} < 0.5$.

The measure defined in (7) constitutes a normalized discretization of the diffusion ODF at voxel v_i . However, it does not guarantee symmetry, i.e., generally $q_{i,j} \neq q_{j,i}$, which makes it unsuitable for the edge weights in an undirected graph. Nevertheless, we can obtain a symmetric weight by considering a bidirectional measure of diffusion given by

$$w_{i,j} = w_{j,i} = q_{i,j} + q_{j,i}, \quad (8)$$

which is constrained to the $[0, 1]$ range. Consequently, we define the graph's edge weights as

$$a_{i,j} = a_{j,i} = h(w_{i,j}), \quad (9)$$

where $h(\cdot) : [0, 1] \rightarrow [0, 1]$ is a tunable sigmoid function defined as

$$h(x) = \frac{((1-\alpha)x)^\beta}{((1-\alpha)x)^\beta + ((1-x)\alpha)^\beta} \in [0, 1], \quad (10)$$

where parameters $\alpha \in (0, 1)$ and $\beta > 0$ control the threshold level and the steepness of the transition from 0 to 1, respectively; see Figure 2. Given that diffusion ODFs generally manifest non-zero magnitudes in all directions, with little contrast between directions of strong and weak diffusion, the thresholding step enables associating weights only to the main directions of diffusion, without the need to use sharpened ODFs as presented in our preliminary work (Abramian et al., 2020a). The choice of the sigmoid function over a Heaviside step ensures retaining a single connected structure in the graph; that is, any non-zero value is mapped to a non-zero value. In this work we use a fixed value of $\beta = 50$, but study the effect of varying the threshold point, in particular, for values $\alpha = 0.85, 0.9$ and 0.95 .

The expression for the edge weight between a pair of voxels (9) integrates information about the extent of diffusion along




Figure 2: Sigmoid function used for thresholding edge weights, for three different values of α and a fixed value $\beta = 50$.




Figure 3: Spectral graph heat kernels, defined within the bounds of the spectrum of a normalized graph Laplacian matrix, i.e., $[0, 2]$.

$\hat{r}_{i,j}$ from both v_i and v_j , amounting to a measure of orientational coherence of the diffusion ODFs at these voxels. In addition, the α parameter of the thresholding function provides added flexibility to this representation.

2.4. Spectral graph heat kernel filters

We design spatial smoothing filters with a heat kernel profile in the graph spectral domain, defined by

$$k(\lambda) = e^{-\tau\lambda}, \quad \forall \lambda \in [0, \lambda_{\max}], \quad (11)$$

where τ is a free parameter determining the spatial extent of the filter. Figure 3 shows several realizations of the heat kernel over a range of τ . When instantiated in the vertex domain, such filters are roughly similar in shape to the Gaussian filters typically used for fMRI analysis; however, there is no direct equivalence between the two filters given the irregular domain represented by the graph.

The filtering is implemented using the polynomial approximation scheme described in Appendix I. The polynomial order required to obtain a suitable approximation of the heat kernel varies depending on the choice of τ . For the range of τ investigated in this study, we used polynomial approximations of order 15, resulting in negligible approximation error in representing the filters.

2.5. Circular phantom construction

Due to the discrete nature of graphs, the set of orientations that can be perfectly captured by edges between voxels is limited by the neighborhood definition used. To evaluate the influence of angular resolution on denoising performance, we tested

the 3-conn and 5-conn neighborhood definitions on a set of circular phantoms of various orientations and radii. These phantoms aim to simulate a wide range of streamline orientations and curvatures, which could be encountered in practice.

Each phantom consisted of an activation profile in the shape of a circular streamline, accompanied by an ODF map oriented along its tangent, representing strong diffusion along the circle. The phantoms were constructed in 93 different orientations in 3D space, selected in a roughly uniform way by subdividing the faces of an icosahedron three times, and from the resulting polyhedron, selecting its subset of vertices that fall in the spherical sector of $0 \leq \theta, \phi \leq \pi/2$; see Figure 4(a). Due to symmetries in the phantoms and the neighborhood definitions, this set of phantom orientations provides a relatively exhaustive sampling of the effects of streamline orientation on smoothing performance. Additionally, to study the effects of curvature, we created the phantoms with three different radii for each orientation: 10, 20, and 30 voxels at 1.25 mm isotropic resolution.

2.6. Streamline-based phantom construction

Since the correlation structure of the BOLD signal in WM is highly anisotropic and resemblant of the diffusion tensor (see Section 1.1), activation patterns in this tissue are likely to have elongated shapes which locally follow the direction of diffusion. To validate the performance of the proposed filtering scheme at detecting such activation patterns, we performed tests on a set of semi-synthetic phantoms that simulate streamline-shaped activations. We denote the phantoms as semi-synthetic since the activation patterns were derived from real diffusion data from the HCP100 dataset. Each phantom consisted of a set of non-uniformly spread activation patterns diffusing along WM streamlines obtained through deterministic tractography of the HCP100 subject set; see Figures 4(b) and (c). Details of the construction of the phantoms are given in Appendix III.

Time-series versions of the streamline-based phantoms were also generated in order to evaluate the performance of the proposed method in the context of a typical general linear model (GLM) fMRI analysis. These were created by using each streamline-based phantom as the underlying ground-truth activity in a 100-volume fMRI time series with a block design alternating 20 volume stretches of rest and activity in an off-on-off paradigm.

3. Results

We validated the performance of the proposed diffusion-informed spatial smoothing (DSS) relative to isotropic Gaussian spatial smoothing (GSS) through a series of tests on synthetic phantoms—circular and streamline-based phantoms—as well as on real data from the HCP100 subject set.

3.1. Diffusion-informed filters

The adaptive properties of DSS filters are illustrated in Figure 5. The three filters shown were generated using identical parameters ($\alpha = 0.9$, $\tau = 7$), and differ only in the location within the WM where they were instantiated. The filters closely

follow the local diffusion orientation in WM described by the diffusion ODFs. For highly anisotropic WM regions this results in slender and strongly oriented filters, yet, as the third column shows, the filters are also capable of taking isotropic shapes in regions of low anisotropy or crossing fibers.

The shape of DSS filters can be controlled by setting the τ parameter of the graph spectral filter kernel (see (11)) and the α parameter of the weight thresholding function (see (10)). While the former controls the spatial extent of the filter in a manner akin to the full width at half maximum (FWHM) of isotropic Gaussian filters, the latter controls the minimum edge weights retained by the graph, which in turn, constrains filters to follow main directions of diffusion. Figure 6 presents a range of different filter shapes that can be achieved by varying these parameters. High values of α result in very narrow, streamline-like filters, whereas lower values result in more isotropic filters. Low enough values of α negate the diffusion-adaptive properties of DSS, with the resulting filters adapting solely to the morphology of the WM domain (see Supplementary Figure S1).

The choice of neighborhood definition plays a significant role in the shape of the resulting filters. In combination with the 5-conn neighborhood definition, higher α values can result in non-local averaging filters when the ODFs are oriented along a neighborhood direction in the outer shell of the neighborhood (see Figure 5 middle left, Figure 6 bottom row). This effect is not present in filters created using the 3-conn neighborhood definition (see Figure S2), which additionally show a more limited capacity to represent orientation due to the reduced angular resolution of the neighborhood definition. More exhaustive results for both 5-conn and 3-conn filters are presented in Supplementary Figures S1-S6.

3.2. Validations on circular phantoms

Circular phantoms of 93 different orientations and 3 different radii were created as described in Section 2.5. Each phantom was corrupted with 10 realizations of additive white Gaussian noise of standard deviation 1, and subsequently denoised by spatial filtering with GSS and DSS over a range of parameters. The FWHM of GSS and the τ parameter of DSS were varied over a range from 1 to 8 in unit steps. Both the 3-conn and 5-conn neighborhood definitions were tested for DSS, which we will refer to as DSS3 and DSS5, respectively. The α parameter of DSS was set to 0.9 throughout.

To assess the denoising performance of GSS, DSS3 and DSS5, we performed receiver operating characteristic (ROC) analyses. The filtered phantom volumes were each thresholded at 300 uniformly-spaced consecutive levels spanning the minimum and maximum value in each filtered volume. The resulting detections for each threshold level were compared with the ground truth of the phantom, yielding true positive rates (TPR) and false positive rates (FPR) that were collected in ROC curves. The area under the curve (AUC) of the ROC curves was then computed, resulting in an overall measure of performance.

Figures 7(a) and (b) show the overall performance of DSS3, DSS5 and GSS as characterized by the AUCs. Due to the lack of equivalence between DSS and GSS filters, there is no direct correspondence between individual values of FWHM and

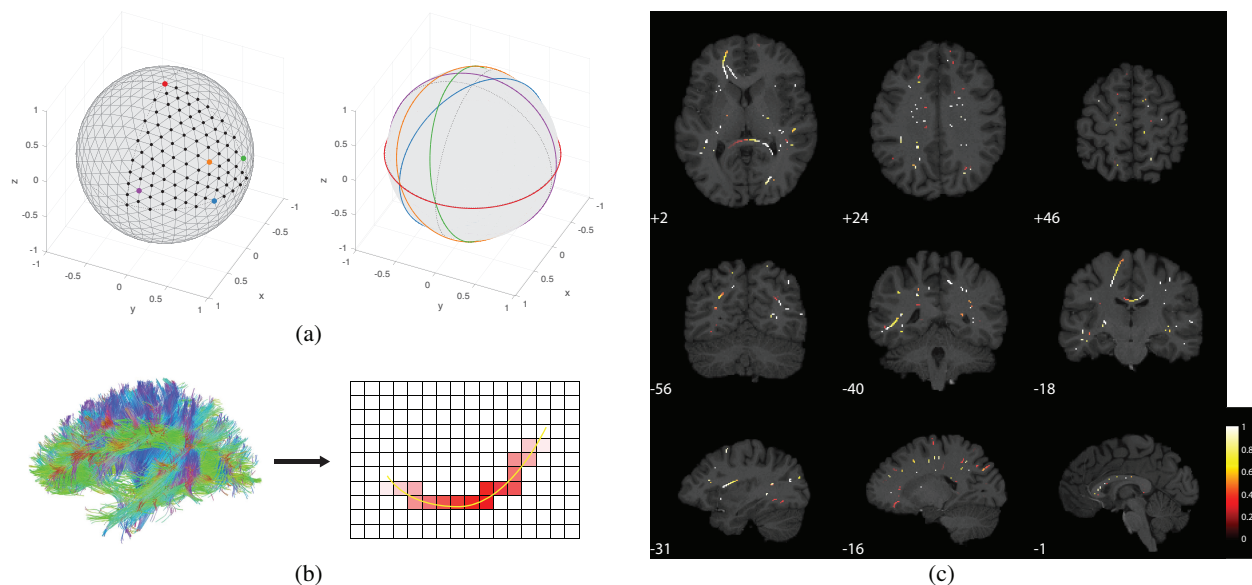


Figure 4: Phantom construction. (a) Circular phantom construction. Left: A subset of vertices of a 3-level subdivided icosahedron, 93 out of 642, were selected. Vectors pointing from the center of the sphere to these vertices constitute the normal vectors of the planes within which circular phantoms were realized. Right: Five representative unit circles with orientations corresponding to the vertices on the left of matching color. (b) Streamline-based phantom construction. A WM streamline constructed using tractography is randomly selected, a focal point along the streamline is randomly selected, and a diffused non-binary activation pattern is created around the focal point. (c) Axial, coronal and sagittal view of a representative streamline-based phantom with 100 streamline activations, overlaid on subject's T1-weighted image.

τ . However, it can be noted that the performance of GSS peaks at 2 mm FWHM, and diminishes for larger filter sizes. On the other hand, both DSS3 and DSS5 achieve substantially higher maximum performances, which are not negatively affected by increased filter size in the range of τ tested.

The median AUC of DSS5 consistently falls above that of DSS3 for $\tau \geq 2$ and all three phantom radii. The performance gap between DSS5 and DSS3 increases for larger τ , and slightly increases on circular phantoms with larger radii, i.e., smaller curvatures. These results corroborate the improvements in detection performance thanks to the increased angular resolution of the 5-conn neighborhood definition. This is further illustrated by Figure 7(c), which shows the performance improvement of DSS5 over DSS3 for individual phantom orientations. The wide range of performance gains is representative of the varying difficulty of representing specific spatial orientations in the discrete domain of graphs, highlighting the importance of angular resolution for the proposed filters.

Given the overall superior performance of DSS5 over DSS3, in the following DSS results are only presented for graphs using the 5-conn neighborhood definition.

3.3. Validations on streamline-based phantoms

A similar analysis was performed on streamline-based phantoms. A single phantom with $N_s = 50, 100$ and 200 streamline activations was created for each of the 95 subjects as described in Section 2.6. As before, each phantom was corrupted with 10 realizations of additive white Gaussian noise of standard deviation 1, and denoised by spatial filtering with GSS and DSS over the same range of parameters. The α parameter of DSS was set to 0.9, whereas values of 0.85 and 0.95 were also tested

on the 100-streamline phantoms. The denoising performance of both methods was assessed by applying the same ROC/AUC analysis described in Section 3.2.

Figures 8(a) and (b) show AUC results on all three types of phantoms for DSS and GSS, respectively. Due to the substantial amount of noise present in the phantoms, spatial smoothing using either GSS or DSS generally leads to better performance compared to no smoothing. DSS outperforms GSS across the range of τ and FWHM values tested, and across the different settings. As with the circular phantoms, the performance of GSS peaks at 2 mm FWHM, with increased size negatively affecting performance beyond that value. DSS shows a similar pattern, with peak performance achieved for τ of 3 and 4 for $\alpha = 0.9$. Both GSS and DSS show better performance on phantoms with a greater number of streamlines. Additional results show that DSS outperforms GSS in both sensitivity and specificity (see Supplementary Figure S7(a)), and across a range of SNR values (see Supplementary Figure S7(b)).

To assess the performance of DSS and GSS in combination with temporal modeling, i.e., as used within fMRI activation mapping studies, time-series version of the streamline-based phantoms were generated as described in Section 2.6. The phantoms were corrupted with additive white Gaussian noise of standard deviation 1 and subsequently spatially filtered with GSS and DSS with the same range of parameters used previously. The smoothed phantoms were subjected to a standard single-subject analysis in SPM, and the resulting t-maps were used in the ROC/AUC analysis.

Figures 8(c) and (d) show AUC results from the time-series phantoms. Due to the increased detection power afforded by temporal modeling, AUCs are higher for all scenarios in the

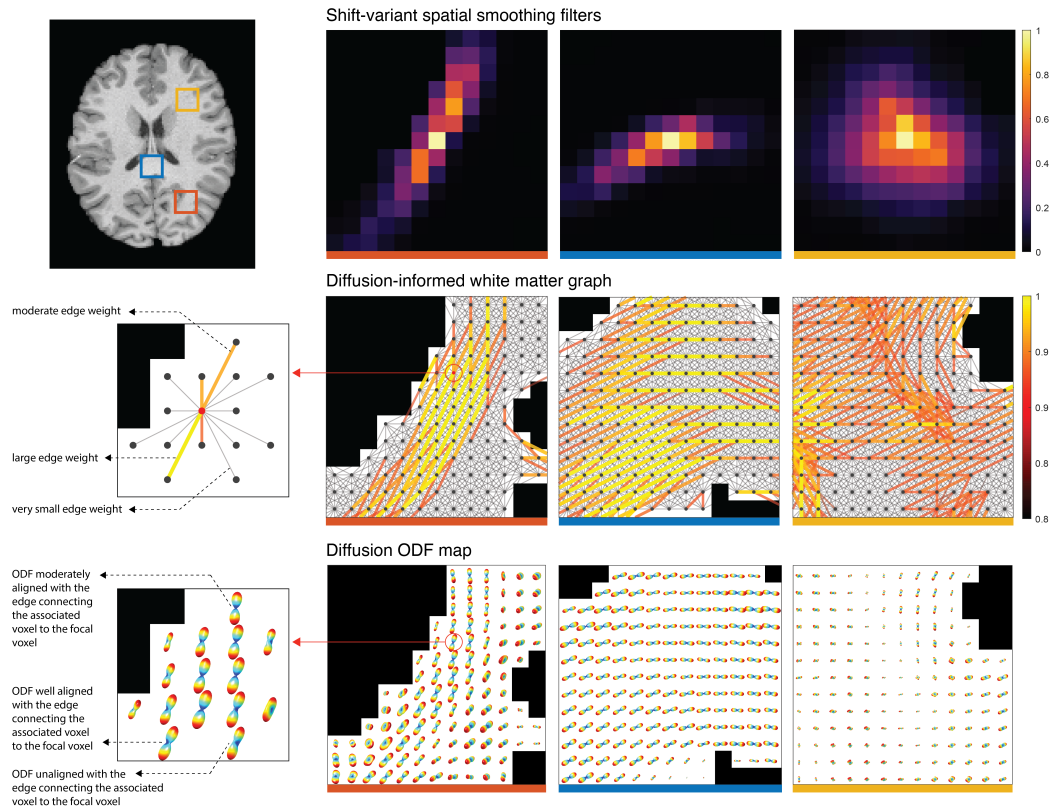


Figure 5: Generation of diffusion-informed smoothing filters. Diffusion ODFs (bottom row) serve as the basis for the creation of a WM graph (middle row). Every WM voxel corresponds to a vertex in the graph, with weighted connections to neighboring voxels (middle left). The edge weights are determined on the basis of coherence between the directions of diffusion and the orientation of the graph edges (bottom left). Using this WM graph definition, graph filters from a single spectral profile become adaptive to the local axonal microstructure when instantiated in different WM regions (top row). Note that both the edges connecting voxels and the graph filters extend in three dimensions, whereas their 2D axial intersection centered at the focal voxel are shown. Graph parameters: 5-connn neighborhood, $\alpha = 0.9$, $\beta = 50$; filter parameters: $\tau = 7$. Filters are shown normalized to the $[0, 1]$ range. ODF interpolation and visualization were performed using the public CSA-ODF package⁴.




Figure 6: Effects of parameters τ and α on the shape of DSS filters located at red ROI shown in Figure 5. Graph parameters: 5-connn neighborhood, $\beta = 50$. Filters are shown normalized to the $[0, 1]$ range.

time-series analysis compared to those in the single-volume analysis. Similarly to the single-volume phantom results, GSS achieves its best performance for 2 mm filters, and considerably deteriorates beyond that size. Notably, GSS only provides a distinct improvement over no smoothing for 2 mm filters. DSS results also show a negative correlation between filter size and performance for $\tau > 2$, but the overall performance is superior to GSS and provides a benefit over no smoothing in most tested cases, with best results achieved for τ between 2 and 4. After subjecting the t-maps to activation mapping with false discovery rate (FDR) correction at 5% (Genovese et al., 2002), the detection maps resulting from DSS showed substantially higher sensitivity and specificity than those from GSS (see Supplementary Figures S8-S10). These results also illustrate that the diminished performance of both methods on phantoms with a greater number of streamline activations is a consequence of increased FPR when using large filters.

Figures 8(a) and (c) also illustrate the effects of varying the α parameter of DSS in single-volume and time-series phantoms, respectively. For both types of phantoms higher values of α generally resulted in better performance. In the case of single-volume phantoms, filters with $\alpha=0.9$ outperformed the others

⁴<https://www.nitrc.org/projects/csaodf-hough>

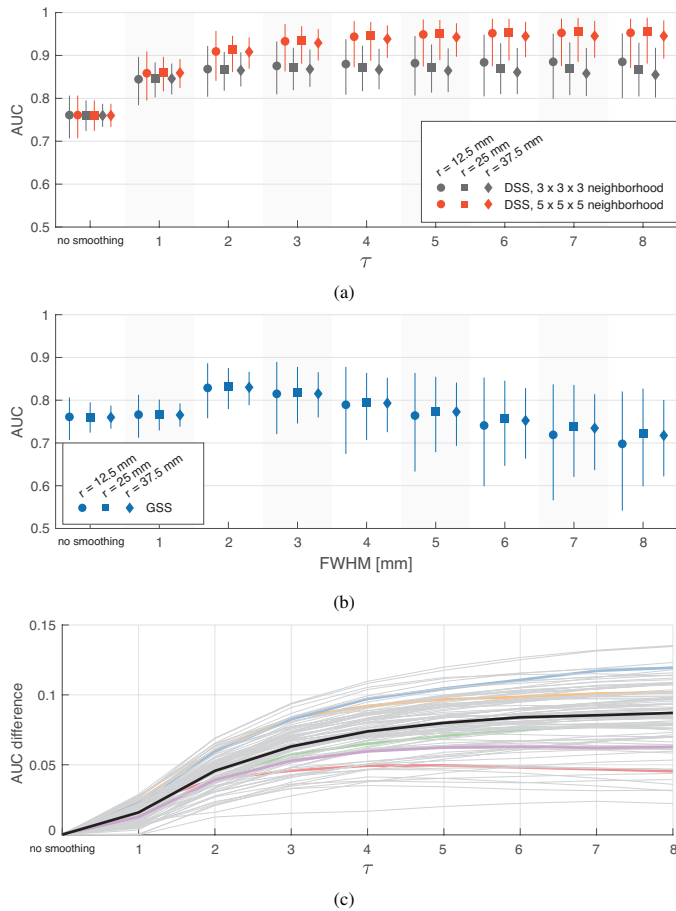


Figure 7: Validation of spatial smoothing on circular phantoms. (a)-(b) AUC of ROC curves obtained from volumes spatially smoothed with DSS and GSS, respectively. The markers show the median AUC over 930 ROCs (93 orientations \times 10 realizations), whereas the whiskers represent 5–95% percentiles. (c) Difference between AUC values for DSS5 and DSS3 for phantoms with 25 mm radius. The black curve shows the difference between the median performances shown in (a), whereas the remaining curves show the difference between the 10-realization medians for each of the 93 phantom orientations. The five colored curves correspond to the phantom orientations shown in Fig. 4(a).

for small filter sizes, while $\alpha=0.95$ is superior for larger filter sizes and across all sizes for time-series phantoms. In addition, filters with $\alpha = 0.95$ show minimal decay in performance as filter size increases for both kinds of phantoms. Filters with $\alpha = 0.85$ consistently performed worse than the others.

3.4. Validations on task fMRI data

In order to explore the effects of the proposed smoothing method on real task fMRI data, we used SPM12 to perform activation mapping on the HCP100 task fMRI data, comprising 23 experimental conditions across 7 tasks. Each GLM analysis included 12 motion regressors (raw and temporal derivative) in addition to regressors for 2 to 8 experimental conditions associated with each task. The canonical HRF model, corresponding to a double gamma, was used. Temporal noise modeling was done using a global AR(1) model. The fMRI data was smoothed using GSS and DSS with the same parameters used previously. For GSS, each fMRI volume was first multiplied with the WM mask, to avoid introducing signal from GM. This

step is not required for DSS, since the method by its nature functions only in WM. The resulting t-maps were then thresholded to determine significant active voxels after FDR correction at 5%. Our choice of FDR as the correction method was due to it only assuming the p -values to be uniformly distributed under the null hypothesis. Correction methods based on assumptions about the smoothness of the data, such as those based on Gaussian random field theory, would be difficult to justify for an adaptive smoothing approach.

Figure 9 shows t-maps and detections from a representative subject obtained using GSS and DSS across a range of filter sizes. Visual inspection of the t-maps reveals that GSS results in generally round features with little oriented structure. For larger Gaussian filters very little visible structure remains. In contrast, t-maps obtained using DSS present notable spatial detail, with linear features in the shape of streamlines discernible across filter sizes. These differences are also present in the detections from both methods. While GSS detections are generally large and rounded, with very few detections present for small filters, DSS is capable of detections showing remarkable spatial detail, and with considerable detection power for small filter sizes. Specifically, the presented detection maps highlight the capability of DSS in identifying separate streamline-shaped activations in two contiguous parallel axonal bundles, which remain distinct across the tested filter sizes. On the other hand, with GSS, these activations are combined into a single active region when large filters are used, or are not present when small filters are used. Notably, the case of FWHM = 3 mm shows activation foci being combined *across* rather than along axonal bundles, suggesting that these activations may not be separable with GSS. Additional activation mapping results are presented in Supplementary Figures S11 and S12.

In order to quantitatively investigate the degree to which spatial structure is present in t-maps obtained using the two smoothing methods, we analyzed the t-maps using *structure tensor* methods (Knutsson, 1989). While a thorough introduction to such methods falls outside the scope of this work, it is sufficient for our purposes to point out that the eigenvalues and eigenvectors of the structure tensor provide information on the presence and orientation of spatial structure, in the form of lines and edges, at a given point in an image or volume.

For each t-map we constructed a quantitative structure map by computing the sum of the structure tensor eigenvalues at every voxel (a measure of the amount of spatial structure in each voxel). The mean value of each structure map provides a global measure of the presence of spatial structure in the corresponding t-map. Figure 10(a) shows a comparison of this global structure measure for DSS and GSS. For both methods the amount of structure present in the t-maps diminishes as the filter size increases, which is consistent with the loss of spatial detail resulting from smoothing the data. This effect is very pronounced for GSS, while t-maps generated using DSS exhibit a more consistent amount of spatial structure across the tested filter sizes.

To determine the extent to which the structure present in the t-maps is influenced by the diffusion information introduced by DSS, we computed the Pearson correlation coefficient between

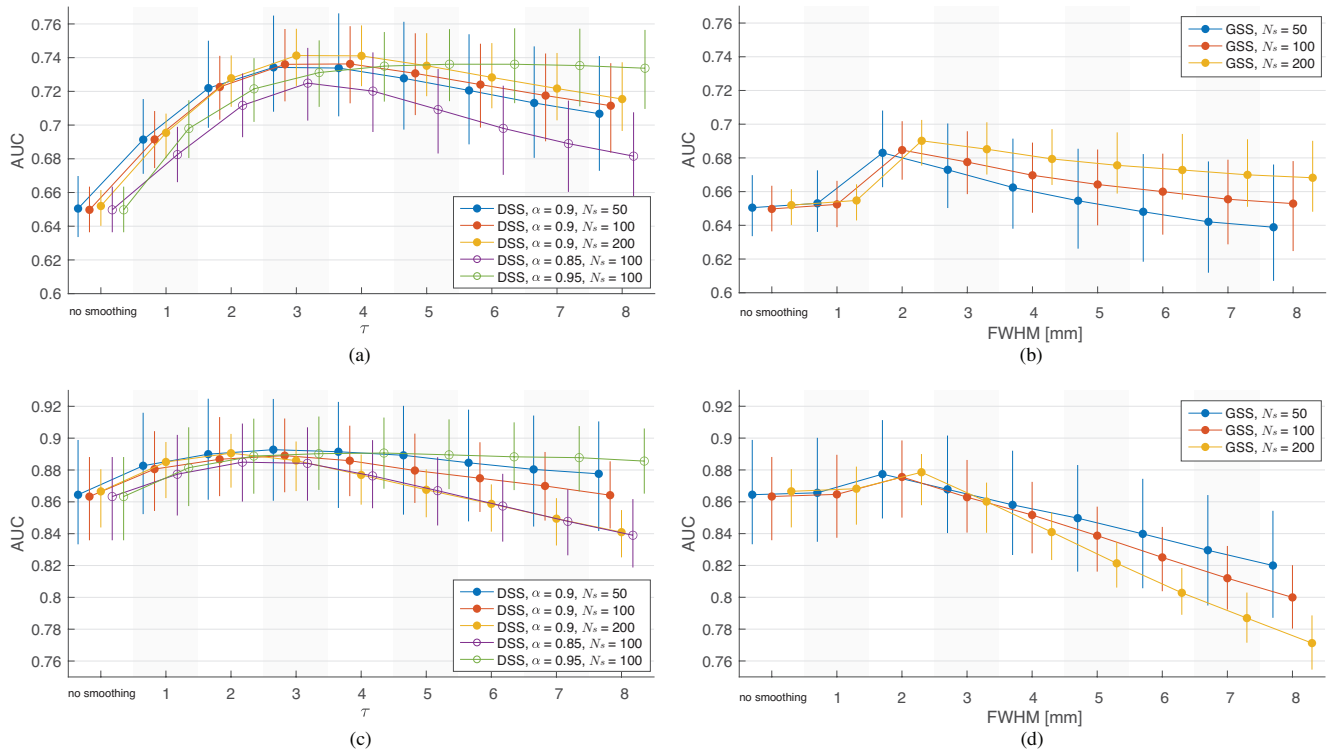


Figure 8: Validation of spatial smoothing on streamline-based phantoms. (a)-(b) AUC of ROC curves obtained from volumes spatially smoothed with DSS and GSS, respectively. (c)-(d) AUC of ROC curves obtained from activation mapping t-maps of time-series streamline-based phantoms smoothed with DSS and GSS, respectively. The markers show the median AUC over 950 ROCs (95 subjects \times 10 realizations), whereas the whiskers represent 5 – 95% percentiles.

the quantitative structure maps and the quantitative anisotropy (QA) map (Yeh et al., 2013) of the associated subject; see Figure 10(b). For DSS this correlation is close to zero at $\tau = 1$, and steadily increases for increased filter sizes. In contrast, the structure manifested in t-maps obtained through GSS shows a slight negative correlation with QA for small filter sizes, which becomes more positive as the filter size increases, but stays substantially below the correlation obtained with DSS. These results suggest that DSS is successful at informing the smoothing process with the local diffusion properties of the underlying WM, with larger values of τ resulting in stronger diffusion encoding.

Figure 11 compares the number of detections obtained from DSS and GSS. To prevent bias due to differences in brain size, we present the fraction of each subject’s WM mask being declared as active. Overall, the detection rates for both methods increase as a function of filter size, with DSS exhibiting a more linear increase than GSS. Moreover, the number of detections on t-maps obtained from volumes smoothed with GSS using large FWHM values, 5 to 8 mm, is on average larger than for maps detected using DSS with any of the tested τ values, as manifested by comparing the average values of corresponding tasks.

In the absence of ground truth it is not possible to make definitive statements on the relationship between differences in the number of voxels deemed active by each method and potential differences in their sensitivity and specificity. However, it can be insightful to quantify the difference between the detec-

tion maps generated with DSS and GSS. To quantify the similarity between a pair of detection maps we computed the Dice coefficient between them, defined as

$$d_{\tau, \text{fwhm}} = \frac{2|M_{\tau} \cap M_{\text{fwhm}}|}{|M_{\tau}| + |M_{\text{fwhm}}|}, \quad (12)$$

where M_{τ} denotes the set of detected voxels using DSS with a given τ , M_{fwhm} denotes the set of detected voxels using GSS with a given FWHM, and $|\cdot|$ denotes set cardinality. The Dice coefficient is constrained to the $[0, 1]$ range, where a value of 1 signifies perfect overlap between the detection maps and a value of 0 represents no overlap.

For every subject and experimental condition we calculated Dice coefficients between detection maps obtained with GSS and DSS of all filter sizes, and arranged them into 8×8 Dice matrices. Additionally, we calculated the maximum Dice coefficient between each DSS filter size and every GSS filter size for each subject and condition. Figure 12 shows Dice results for several representative experimental conditions. The overall similarity between the detection maps obtained with DSS and GSS is relatively low. The highest ensemble Dice is achieved for $\tau = 1$ and FWHM = 2 mm, where it reaches a value of 0.6, and diminishes to less than 0.4 for the largest filters tested. The relationship between the τ and FWHM values that result in the highest similarity in the detection maps is also shown to be nonlinear, with values of τ between 5 and 8 giving maps most similar to those from FWHM = 4 mm. The similarity between the detection maps also shows considerable variation

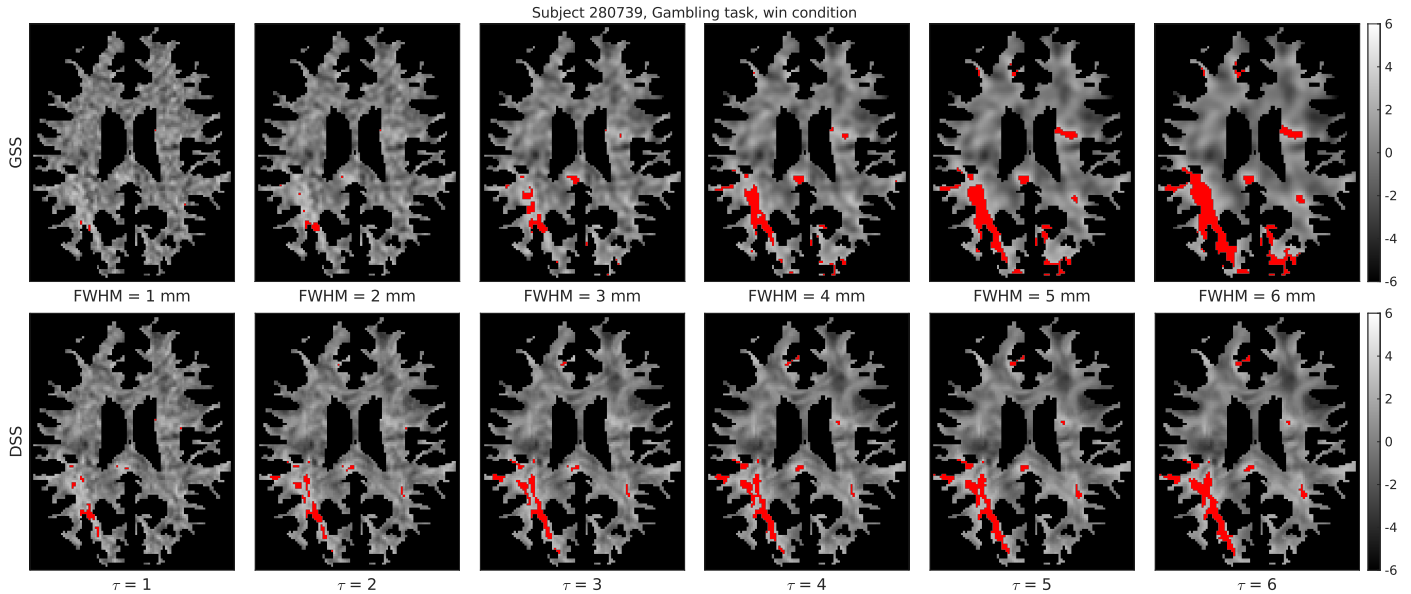


Figure 9: Comparison of representative activation mapping results for a single subject generated with GSS (top) and DSS (bottom), with t-maps shown in grayscale and detections overlaid in red (FDR-corrected at 5%).

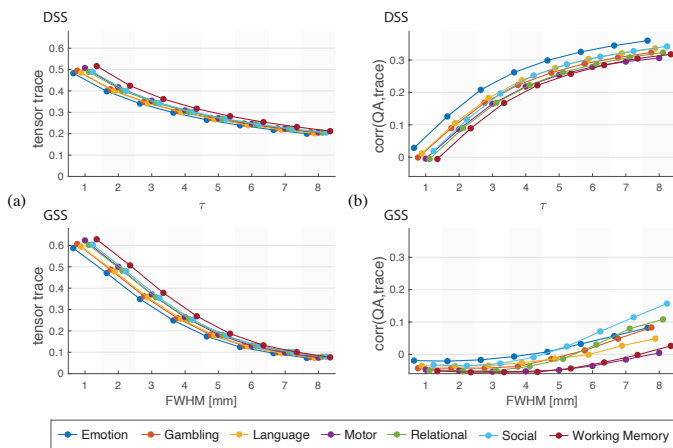


Figure 10: Structural analysis of task fMRI t-maps, obtained using local structure tensor analysis (Knutsson, 1989) where the eigenvalues of the structure tensor denote the amount of spatial structure. (a) Quantification of the amount of anisotropic structure observed in t-maps, specified by the mean structure map value, averaged across the task's experimental conditions. (b) Correlation between subjects' QA maps and structure maps, averaged across each task's experimental conditions. Markers shows the median value across the 95 subjects.

across tasks and individual experimental conditions (see results for all experimental conditions in Supplementary Figure S13), with below-average similarity in the Language and Motor tasks and above-average in the Gambling and Relational tasks.

In order to determine whether the detections generated by either method are a subset of the detections from the other, we examined the number of common and unique detections produced by DSS and GSS. For all subjects and experimental conditions, the detection maps produced by DSS were compared with the most similar maps produced by GSS. Figure 12, bottom right, shows the average number of voxel detections common to both methods, as well as those unique to each method, for the tested

values of τ . These results show that, across filter sizes, both DSS and GSS produce a considerable number of detections that are not produced by the other method. This, together with the generally low Dice results, suggest the presence of substantial differences in the localization and spatial extent of activations detected using DSS and GSS.

3.5. Processing time

Although the proposed methodology requires additional MRI scanning time for the acquisition of DW-MRI data, it does not impose a dramatic increase in processing time over conventional approaches. Using a workstation with an Intel Core i7-7700K processor and 64 GB of RAM, the generation of diffusion ODFs from DW-MRI data required approximately 90 seconds. The graph and its Laplacian matrix could then be calculated from the ODFs in under 15 seconds. Both of these operations need only be performed once per subject.

In our implementation, the average filtering time of a single volume with GSS was 10.3 ms using the `imgaussfilt3` Matlab function (the same operation required about 450 ms when using the smoothing implemented in SPM). On the other hand, DSS filtering scales efficiently with the number of filter kernels used. Average single-volume DSS filtering times for a single kernel were 115 ms for the 5-conn neighborhood and 56 for 3-conn, and became reduced to 17.7 ms and 11.0 ms, respectively, when using 8 filter kernels at once. With worst case performance, the proposed method gave filtering times of around 45 seconds for a 405-volume series (the longest of those available in HCP data, corresponding to the Working Memory task).

4. Discussion

4.1. Interpretation of results from simulated data

Previous implementations of voxel-wise graphs on GM (Behjat et al., 2015; Maghsadagh et al., 2019; Behjat and Larsson,




Figure 11: Fraction of voxels within WM mask detected as being significant using DSS (top left) and GSS (bottom left) across 7 functional tasks, over 95 subjects. Significant voxels were determined after FDR correction at 5%. In the plots on the left, each dot corresponds to one subject, whereas ■ shows the average value across the 95 subjects. The plots on the right show the trend of the average value as a function of filter parameters τ and FWHM for GSS and DSS respectively.

2020) have used the 3-conn neighborhood in defining graph edges. However, given the different nature of the proposed encoding for WM graphs, which needs to represent axonal orientations rather than the morphology of GM, we considered the potential advantages of using larger neighborhood definitions. To this end, we compared the denoising performance obtained with graphs using the 3-conn and 5-conn neighborhood definitions on circular phantoms of multiple orientations and radii. Such phantoms were used because, barring discretization artifacts, they offer an exhaustive sampling of all possible orientations in which data can appear in three dimensions. The results show a clear improvement from using the larger neighborhood definition (see Figures 7(a) and (c)), which can be attributed to its superior angular resolution of 98 neighborhood directions, against the 26 of the 3-conn definition. Furthermore, comparing performances obtained on phantoms of different radii shows that the larger neighborhood definition provides more stable performance across spatial curvatures than the smaller neighborhood, which performs worse for smaller curvatures, particularly for larger filters. Compared to isotropic Gaussian smoothing (see Figure 7(b)), both the 3-conn and 5-conn neighborhood definitions used in DSS showed enhanced denoising performance on circular phantoms. In particular, while the performance of GSS deteriorates for larger filter sizes, the performance of DSS reaches a plateau instead, suggesting that the diffusion-informed nature of DSS filters is capable of minimizing the introduction of spurious signal even for larger filter sizes.

To better mimic spatial activation patterns manifested as BOLD contrast in WM, we designed and studied semi-synthetic streamline-based phantoms, whose diffuse activation patterns are representative of WM fiber structures, along which corre-

lated BOLD activity is expected (Ding et al., 2013, 2016, 2018). The phantoms were studied in two settings. In the first setting the denoising performance was studied in the absence of temporal modeling, in the same way as the circular phantoms. In this setting, both methods provided an improvement over no smoothing, but DSS outperformed GSS for all tested filter sizes (Figure 8(a)). In the second setting the phantoms were studied within the context of GLM activation mapping, i.e. with temporal modeling. In this scenario, GSS provided only minimal improvements over no smoothing, whereas DSS provided a notable improvement (Figure 8(b)). These phantoms were also used to study the influence of the α parameter of DSS, which sets a lower bound on the weight of connections allowed in the WM graph. Due to the narrower and more directional filters resulting from higher α values (Figure 5), the increased performance on the streamline-based phantoms would be expected (Figures 8(a) and (c)). However, this result may not be readily extensible to real fMRI data, as the spread of real activation patterns is not known.

4.2. Interpretation of results from real data

We compared activation mapping results from DSS and GSS on task fMRI data from the HCP100 subject set. Structure tensor analysis of the resulting t-maps revealed that the overall amount of structure present diminished for larger filter sizes, an effect which is more pronounced for GSS (Figure 10(a)). Such results reflect the loss of spatial details that happens as a result of lowpass filtering. However, due to the highly anisotropic shapes that DSS filters take within the WM (Figure 5), features in the shape of lines and edges can be present in t-maps even for larger filter sizes (Figure 9). In addition, the spatial structure present in the t-maps obtained with DSS is correlated with

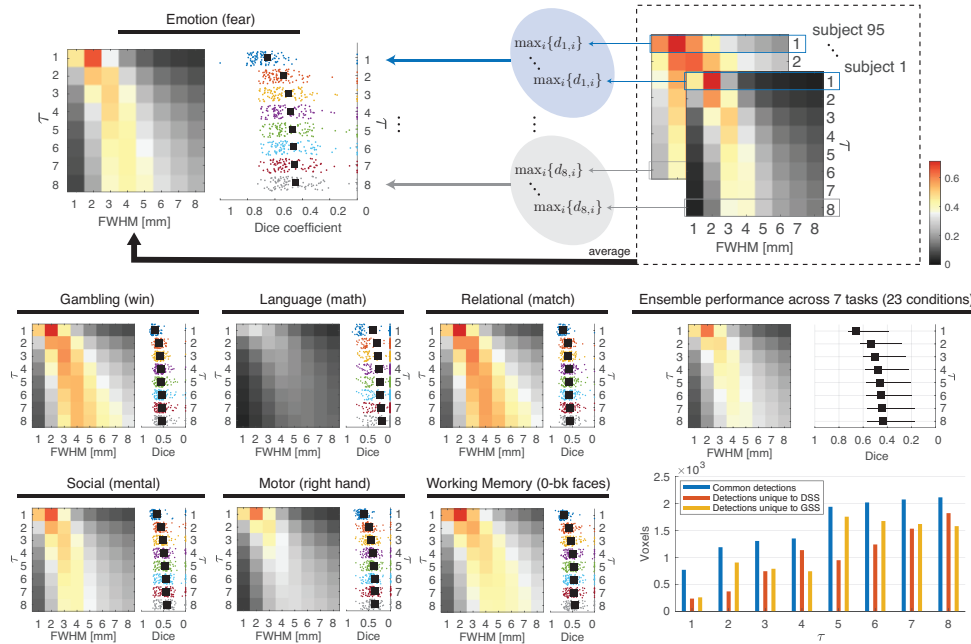


Figure 12: Dice similarity between detection maps generated with DSS and GSS. For each subject and condition, an 8×8 Dice matrix was computed, where each element represented $d_{\tau, \text{fwhm}}$, see (12). For a given subject, if neither DSS nor GSS led to any detections for a given combination of τ and FWHM, the corresponding element was excluded from further analysis. The schematic on top explains how the results were ensembled across subjects, resulting in two plots for each experimental condition; in the plots on the right, the mean of the scattered values is indicated by \blacksquare . Results are presented for a representative experimental condition in each task—see results across the 23 conditions in Supplementary Figure S13, as well as ensembled across 23 conditions; the ensemble plot on the left shows the average across conditions, whereas the one on the right shows the median and range of the mean maximum Dice values across conditions. The plot on the bottom right shows the average number of common and unique detections generated by DSS and GSS across all subjects and conditions, wherein every value of τ was compared with the FWHM that resulted in the maximum ensemble Dice coefficient.

regions of high diffusion anisotropy (Figure 10(b)), indicating that DSS successfully adapts its smoothing to the underlying WM microstructure.

Due to the differences in their definitions, as well as the adaptive nature of DSS, there is no direct correspondence between GSS and DSS filters. This is corroborated by the low Dice coefficients between detection maps resulting from both methods (Figure 12). While GSS with $\text{FWHM} \geq 4$ mm generally results in more detections than DSS with any of the tested filter sizes (Figure 11), this result is consistent with the lower specificity that can be expected from GSS due to its lack of domain adaptation (see Supplementary Figures S8-S10). On the other hand, example detection maps corroborate that DSS is capable of resolving slender, streamline-like activation patterns along axonal pathways across multiple filter sizes (Figure 9, Supplementary Figures S11 and S12) by leveraging information about the spatial correlation structure of the BOLD signal in WM.

4.3. Limitations

We used a sigmoid function, see (10), as a means of boosting orientation encoding, allowing diffusion only along main directions of diffusion coherence. We studied three threshold values, $\alpha = 0.85, 0.9$ and 0.95 , all of which yielded better performance than GSS on phantom data, with noticeable variations in performance among the three values. However, the general choice of the thresholding function and its associated parameters is rather ad-hoc, which is a complication of similar nature as that encountered in connectomic studies. Future work should consider a more rigorous validation of the thresholding scheme for obtaining optimal performance, especially on real fMRI data.

Accurate multi-modal co-registration of functional, structural, and diffusion MRI data is a cornerstone of the proposed methodology. Within this study, we used preprocessed HCP

data, which have been diligently motion-corrected, distortion-corrected, and co-registered (Glasser et al., 2013). However, conducting solid preprocessing steps may not be possible in some datasets, and if so, results obtained using the proposed method should be interpreted with care.

A recent body of studies has highlighted substantial differences of the HRF response in WM from that in GM (Li et al., 2019b). However, due to the ongoing nature of this research, we decided to use the standard HRF model that is conventionally used in fMRI activation mapping. Given that our work is comparative, and since the choice of the HRF model affects the methods in the same way, we do not believe that our conclusions would be substantially affected by the use of a more precise model. Nevertheless, future work aimed at investigating the WM BOLD signal can most likely benefit from combining a more appropriate HRF model with adaptive smoothing of the BOLD signal by DSS.

4.4. Outlook; potential extensions and other applications

Due to the limited degree to which diffusion ODFs can differentiate fiber orientation (Jones et al., 2013), we boost orientation encoding by means of a weight thresholding scheme. Alternatively, the proposed design can be extended to leverage fiber orientation distribution functions estimated from either the diffusion ODFs (Descoteaux et al., 2008) or the raw diffusion data (Tournier et al., 2007), to obviate the need for thresholding. In the absence of HARDI data but presence of DTI data, the proposed method can be readily extended to leverage diffusion tensors instead of diffusion ODFs, which can be of particular interest for reanalyzing the vast extent of currently available fMRI datasets that are accompanied by DTI data.

In the absence of any DW-MRI data, it would be possible to adapt the proposed method to use a structure ten-

representation (Knutsson, 1989) derived from T1-weighted MRI images as the complementary contrast (Abramian et al., 2020b), wherein the proposed filtering scheme could be extended to function across the entire brain mask. The resulting morphology-based spatial smoothing could then be seen as a GSP-based alternative to non-linear filtering algorithms which enable spatial smoothing within similar anatomical compartments (Smith and Brady, 1997; Weickert and Scharr, 2002; Ding et al., 2005; Rydell et al., 2008; Lohmann et al., 2018), but will not provide adaptation to WM fiber orientations.

In addition to performing denoising through heat kernel smoothing (i.e., lowpass filtering), the proposed WM graphs can be used to implement graph-wavelet denoising, similar to that implemented by Behjat et al. (2015) for GM graphs, using novel data-driven GSP denoising schemes (de Loynes et al., 2019) in combination with computationally efficient multi-scale spectral graph decomposition methods (Li et al., 2019c; Shuman, 2020) that can be tractably implemented on large graphs.

In the present study, we only explored spatial smoothing of task-based fMRI data within the context of activation mapping, whereas DSS can be readily applied to WM resting-state fMRI data, where recent studies of WM resting-state fMRI data have used Gaussian smoothing of the data as a pre-processing step. Such research appears particularly promising in light of studies reporting the existence of BOLD-like response in resting-state data (Liu and Duyn, 2013; Petridou et al., 2013; Karahanoğlu and Van De Ville, 2015), and the current growing interest in exploring functional dynamics of WM at rest (Peer et al., 2017; Ding et al., 2018; Li et al., 2019a; Wang et al., 2020; Li et al., 2020a).

Another research avenue that can benefit from the proposed WM graph design is structural studies. The eigenvalues of cortical surface graphs as well as their eigenmodes have been leveraged in multiple applications, namely, quantifying cortical folding patterns (Germanaud et al., 2012; Rabiei et al., 2016; Dubois et al., 2019), age prediction (Wachinger et al., 2015; Masoumi et al., 2019), and analysis of brain asymmetry in health (Wachinger et al., 2015; Maghsadgh et al., 2019) and in disease (Wachinger et al., 2016a,b; Masoumi et al., 2019). Such analyses can be extended to leverage the spectra of WM graphs. Analysis on similarly designed graphs using DW-MRI data—covering the entire brain rather than just the WM—has shown that an initial subset of the graph eigenmodes provides informative features to distinguish between subjects (Tarun et al., 2019). Lastly, ODF-based WM graphs may be found beneficial in deriving structural connectivity measures that account for direct as well as indirect pathways, for example, as an alternative to the DTI-based conductance model proposed by Frau-Pascual et al. (2019).

5. Conclusion

The development of methods geared specifically towards WM can prove substantially helpful in investigating the functional significance of the WM BOLD signal. Notwithstanding the repository of sophisticated smoothing techniques found in

the literature, to date, studies on WM fMRI data have mainly resorted to isotropic Gaussian smoothing. An apparent reason is the ease in implementing Gaussian smoothing and its availability in widely used open-access software packages, which facilitate its routine application. The proposed diffusion-informed spatial filtering method, in conjunction with the use of WM specific HRF models and MR sequences, holds promise to aid better understanding of the functional role of WM.

Code and data availability

An implementation of the methods proposed in this work will be made available as a MATLAB package on <https://www.nitrc.org/>—customized link to will added in the final version of the manuscript. The simulated circular and streamline-based phantoms used in this work will be made available <https://openneuro.org/>—customized link will be added in the final version of the manuscript.

Acknowledgements

Data used in this work were provided by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 Institutes and Centers of the National Institutes of Health (NIH) that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

This work, and HB, were supported by the Swedish Research Council under Grant 2018-06689, and in part by the Royal Physiographic Society of Lund, the Thorsten and Elsa Segerfalk Foundation, and the Hans Werthén Foundation. DA and AE were supported by the Swedish Research Council under Grant 2017-04889, by the ITEA3 / VINNOVA funded project Intelligence based iMprovement of Personalized treatment And Clinical workflow support (IMPACT), and by the Center for Industrial Information Technology (CENIIT) at Linköping University. IA was supported by the BrightFocus Foundation (A2016172S) and the NIH, specifically the National Institute of Diabetes and Digestive and Kidney Diseases (K01DK101631) and the National Institute on Aging (R56AG068261).

A preliminary version of this work has been presented (Abramian et al., 2020a).

Appendix I: Spectral graph filtering through polynomial approximation

Spectral graph filtering can be efficiently implemented using polynomial approximation schemes, mitigating the need to diagonalize large \mathbf{L} matrices as those used in the present work. Using this approach, a spectral kernel $k(\lambda)$ is first approximated using a polynomial of suitable order, denoted $p(\lambda) : [0, 2] \rightarrow \mathbb{R}$, and filtering of signal \mathbf{f} is then implemented as

$$\tilde{\mathbf{f}} \stackrel{(3)}{=} \sum_{l=1}^{N_g} p(\lambda_l) \hat{\mathbf{f}}[l] \mathbf{u}_l, \quad (13)$$




Figure 13: Uniform sampling within solid angles along different orientations. (a) An icosahedron with five levels of subdivision, wherein the subset of its vertices that fall within the solid angle $4\pi/98$ around the z-axis direction, marked with black dots, are treated as a template sampling pattern. (b) The template sampling pattern (black) is then rotated towards other neighborhood directions; two directions shown here, in red and blue.

where the vectorized form of (3) is invoked. Noting that $\mathbf{L}\mathbf{u}_l = \lambda_l\mathbf{u}_l \Rightarrow p(\mathbf{L})\mathbf{u}_l = p(\lambda_l)\mathbf{u}_l$, (13) can then be simplified as

$$\tilde{\mathbf{f}} = p(\mathbf{L}) \sum_{l=1}^{N_g} \hat{\mathbf{f}}[l]\mathbf{u}_l = p(\mathbf{L})\mathbf{f}, \quad (14)$$

where in the last equality we used $\mathbf{f} = \sum_{l=1}^{N_g} \hat{\mathbf{f}}[l]\mathbf{u}_l$. Using this scheme, a filtering is performed through a series of polynomial matrix operations on \mathbf{L} , without the need to access the Laplacian eigenvalues. In this work we leveraged truncated Chebyshev polynomial approximations of spectral kernels as presented by Hammond et al. (2011), which have the benefit of approximating a minimax polynomial, minimizing an upper bound on the approximation error.

Appendix II: Uniform sampling of ODFs

We define a spherical sampling grid using the vertices of an icosahedron with five levels of subdivision, which results in a total of 10,242 vertices on the unit sphere. Due to non-uniformity in the spatial spread of the vertices, the number and distribution of vertices that fall within the solid angles $\Omega_{i,j}$ subtended along the 26/98 different $\hat{r}_{i,j}$ neighborhood directions vary. To overcome this bias, we treat the vertices that fall within $\Omega_{i,j}$ around the z-axis as a sampling template, resulting in $N_t = 389$ and 105 template directions for the 3-conn and 5-conn neighborhood definitions, respectively. The sampling template is then rotated and centered around each neighborhood direction $\hat{r}_{i,j}$, resulting in a set of sampling directions $\{\hat{r}_{i,j}^k | k = 1, \dots, N_t\}$ (see Figure 13).

Appendix III: Streamline-based phantom construction

For each subject, 10 thousand streamlines, denoted $\{s_i(x) \in \mathbb{R}^3\}_{i=1 \dots 10000}$, were generated through deterministic tractography using the method presented by Yeh et al. (2013) as implemented in DSI Studio. A subset of S streamlines from a single subject

was randomly selected and used as the basis to produce a phantom. Each streamline s_i was first voxelized, resulting in a vector \mathbf{s}_i containing the indices of the voxels through which it passes. A random source point for the activation was then selected, represented by an indicator vector \mathbf{d}_i of the same length as \mathbf{s}_i with a single 1. An adjacency matrix \mathbf{A}_i was then defined, specifying that every voxel in \mathbf{s}_i is connected to itself and its neighbors within a $3 \times 3 \times 3$ neighborhood, with equal weights adding up to 1. The diffuse activation pattern, denoted \mathbf{p}_i , was then constructed as

$$\mathbf{p}_i = \frac{\mathbf{A}_i^n \mathbf{d}_i}{\max \mathbf{A}_i^n \mathbf{d}_i}, \quad (15)$$

where the exponent n is a parameter that controls the extent of spatial spread of the activation, and was set to 250 in the design of all the phantoms used in this work. Finally, the phantom was created by merging the various activation patterns $\{\mathbf{p}_i\}_{i=1 \dots S}$ into a single volume.

References

- Abdelnour, F., Dayan, M., Devinsky, O., Thesen, T., Raj, A., 2018. Functional brain connectivity is predictable from anatomic network's Laplacian eigenstructure. *NeuroImage* 172, 728–739.
- Abramian, D., Larsson, M., Eklund, A., Behjat, H., 2020a. Improved functional MRI activation mapping in white matter through diffusion-adapted spatial filtering, in: 2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI), IEEE. pp. 539–543.
- Abramian, D., Sidén, P., Knutsson, H., Villani, M., Eklund, A., 2020b. Anatomically informed Bayesian spatial priors for fMRI analysis, in: 2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI), IEEE. pp. 1026–1030.
- Aqil, M., Atasoy, S., Kringelbach, M.L., Hindriks, R., 2020. Graph neural fields: a framework for spatiotemporal dynamical models on the human connectome. *bioRxiv*.
- Atasoy, S., Donnelly, I., Pearson, J., 2016. Human brain networks function in connectome-specific harmonic waves. *Nat. Commun.* 7, 10340.
- Atasoy, S., Roseman, L., Kaelen, M., Kringelbach, M.L., Deco, G., Carhart-Harris, R.L., 2017. Connectome-harmonic decomposition of human brain activity reveals dynamical repertoire re-organization under LSD. *Scientific reports* 7, 17661.
- Behjat, H., Doğan, Z., Van De Ville, D., Sörnmo, L., 2019. Domain-informed spline interpolation. *IEEE Transactions on Signal Processing* 67, 3909–3921.
- Behjat, H., Larsson, M., 2020. Spectral characterization of functional MRI data on voxel-resolution cortical graphs, in: 2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI), IEEE. pp. 558–562.
- Behjat, H., Leonardi, N., Sörnmo, L., Van De Ville, D., 2014. Canonical cerebellar graph wavelets and their application to fMRI activation mapping, in: *Proc. IEEE Int. Conf. Eng. Med. Biol. Soc.*, pp. 1039–1042.
- Behjat, H., Leonardi, N., Sörnmo, L., Van De Ville, D., 2015. Anatomically-adapted graph wavelets for improved group-level fMRI activation mapping. *Neuroimage* 123, 185–199.
- Behjat, H., Leonardi, N., Van De Ville, D., 2013. Statistical parametric mapping of functional MRI data using wavelets adapted to the cerebral cortex, in: *Proc. IEEE Int. Symp. Biomed. Imaging*, pp. 1070–1073.
- Bolton, T.A., van De Ville, D., 2020. Dynamics of brain activity captured by graph signal processing of neuroimaging data to predict human behaviour, in: 2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI), IEEE. pp. 549–553.
- Bolton, T.A., Farouj, Y., Inan, M., Van De Ville, D., 2019. Structurally-informed deconvolution of functional magnetic resonance imaging data, in: 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019), IEEE. pp. 1545–1549.
- Breakspear, M., Bullmore, E.T., Aquino, K., Das, P., Williams, L.M., 2006. The multiscale character of evoked cortical activity. *Neuroimage* 30, 1230–1242.

- Bullmore, E., Fadili, J., Maxim, V., Sendur, L., Whitcher, B., Suckling, J., Brammer, M., Breakspear, M., 2004. Wavelets and functional magnetic resonance imaging of the human brain. *Neuroimage* 23 Suppl 1, S234–S249.
- Carp, J., 2012. The secret lives of experiments: methods reporting in the fmri literature. *Neuroimage* 63, 289–300.
- Coalson, T.S., Van Essen, D.C., Glasser, M.F., 2018. The impact of traditional neuroimaging methods on the spatial localization of cortical areas. *Proceedings of the National Academy of Sciences* 115, E6356–E6365.
- Descoteaux, M., Deriche, R., Knosche, T.R., Anwander, A., 2008. Deterministic and probabilistic tractography based on complex fibre orientation distributions. *IEEE transactions on medical imaging* 28, 269–286.
- Ding, Z., Gore, J.C., Anderson, A.W., 2005. Reduction of noise in diffusion tensor images using anisotropic smoothing. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 53, 485–490.
- Ding, Z., Huang, Y., Bailey, S.K., Gao, Y., Cutting, L.E., Rogers, B.P., Newton, A.T., Gore, J.C., 2018. Detection of synchronous brain activity in white matter tracts at rest and under functional loading. *Proc. Natl Acad. Sci.* 115, 595–600.
- Ding, Z., Newton, A.T., Xu, R., Anderson, A.W., Morgan, V.L., Gore, J.C., 2013. Spatio-temporal correlation tensors reveal functional structure in human brain. *PLoS one* 8, e82107.
- Ding, Z., Xu, R., Bailey, S.K., Wu, T.L., Morgan, V.L., Cutting, L.E., Anderson, A.W., Gore, J.C., 2016. Visualizing functional pathways in the human brain using correlation tensors and magnetic resonance imaging. *Magnetic resonance imaging* 34, 8–17.
- Doucette, J., Wei, L., Hernández-Torres, E., Kames, C., Forkert, N.D., Aamand, R., Lund, T.E., Hansen, B., Rauscher, A., 2019. Rapid solution of the Bloch-Torrey equation in anisotropic tissue: Application to dynamic susceptibility contrast MRI of cerebral white matter. *NeuroImage* 185, 198–207.
- Dubois, J., Lefèvre, J., Angleys, H., Leroy, F., Fischer, C., Lebenberg, J., Dehaene-Lambertz, G., Borradori-Tolsa, C., Lazeyras, F., Hertz-Pannier, L., et al., 2019. The dynamics of cortical folding waves and prematurity-related deviations revealed by spatial and spectral analysis of gyrification. *NeuroImage* 185, 934–946.
- Eklund, A., Andersson, M., Knutsson, H., 2011. Fast random permutation tests enable objective evaluation of methods for single-subject fMRI analysis. *International journal of biomedical imaging* 2011.
- Erdoğan, S.B., Tong, Y., Hocke, L.M., Lindsey, K.P., deB Frederick, B., 2016. Correcting for blood arrival time in global mean regression enhances functional connectivity analysis of resting state fMRI-BOLD signals. *Frontiers in human neuroscience* 10, 311.
- Fischl, B., 2012. *Freesurfer*. *Neuroimage* 62, 774–781.
- Fraser, L.M., Stevens, M.T., Beyea, S.D., D’Arcy, R.C., 2012. White versus gray matter: fMRI hemodynamic responses show similar characteristics, but differ in peak amplitude. *BMC neuroscience* 13, 91.
- Frau-Pascual, A., Fogarty, M., Fischl, B., Yendiki, A., Aganj, I., 2019. Quantification of structural brain connectivity via a conductance model. *NeuroImage* 189, 485–496.
- Friman, O., Borga, M., Lundberg, P., Knutsson, H., 2003. Adaptive analysis of fMRI data. *Neuroimage* 19, 837–845.
- Gao, Y., Li, M., Huang, A.S., Anderson, A.W., Ding, Z., Heckers, S.H., Woodward, N.D., Gore, J.C., 2020. Declined functional connectivity of white matter during rest and working memory tasks associates with cognitive impairments in schizophrenia. *medRxiv*.
- Gawryluk, J.R., Brewer, K.D., Beyea, S.D., D’arcy, R.C., 2009. Optimizing the detection of white matter fMRI using asymmetric spin echo spiral. *Neuroimage* 45, 83–88.
- Gawryluk, J.R., Mazerolle, E.L., D’Arcy, R.C., 2014. Does functional mri detect activation in white matter? a review of emerging evidence, issues, and future directions. *Front. Neurosci.* 8, 239.
- Geissler, A., Lanzenberger, R., Barth, M., Tahamtan, A.R., Milakara, D., Gartus, A., Beisteiner, R., 2005. Influence of fmri smoothing procedures on replicability of fine scale motor localization. *Neuroimage* 24, 323–331.
- Genovese, C.R., Lazar, N.A., Nichols, T., 2002. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage* 15, 870–878.
- Germanaud, D., Lefèvre, J., Toro, R., Fischer, C., Dubois, J., Hertz-Pannier, L., Mangin, J.F., 2012. Larger is twistier: spectral analysis of gyrification (SPANGY) applied to adult brain size polymorphism. *Neuroimage* 63, 1257–1272.
- Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., et al., 2013. The minimal preprocessing pipelines for the human connectome project. *Neuroimage* 80, 105–124.
- Gore, J.C., Li, M., Gao, Y., Wu, T.L., Schilling, K.G., Huang, Y., Mishra, A., Newton, A.T., Rogers, B.P., Chen, L.M., et al., 2019. Functional mri and resting state connectivity in white matter—a mini-review. *Magnetic resonance imaging* 63, 1–11.
- Grauskas, L.A., Frizzell, T., Song, X., D’Arcy, R.C., 2019. White matter fMRI activation cannot be treated as a nuisance regressor: Overcoming a historical blind spot. *Frontiers in neuroscience* 13, 1024.
- Hammond, D., Vandergheynst, P., Gribonval, R., 2011. Wavelets on graphs via spectral graph theory. *Appl. Comput. Harmon. Anal.* 30, 129–150.
- Huang, W., Bolton, T.A.W., Medaglia, J.D., Bassett, D.S., Ribeiro, A., Ville, D.V.D., 2018a. A graph signal processing perspective on functional brain imaging. *Proceedings of the IEEE* 106, 868–885. doi:10.1109/JPROC.2018.2798928.
- Huang, Y., Bailey, S.K., Wang, P., Cutting, L.E., Gore, J.C., Ding, Z., 2018b. Voxel-wise detection of functional networks in white matter. *Neuroimage* 183, 544–552.
- Iturria-Medina, Y., Canales-Rodriguez, E., Melie-Garcia, L., Valdes-Hernandez, P., Martinez-Montes, E., Alemán-Gómez, Y., Sánchez-Bornot, J., 2007. Characterizing brain anatomical connections using diffusion weighted MRI and graph theory. *Neuroimage* 36, 645–660.
- Ji, G.J., Ren, C., Li, Y., Sun, J., Liu, T., Gao, Y., Xue, D., Shen, L., Cheng, W., Zhu, C., et al., 2019. Regional and network properties of white matter function in Parkinson’s disease. *Human brain mapping* 40, 1253–1263.
- Jiang, Y., Luo, C., Li, X., Li, Y., Yang, H., Li, J., Chang, X., Li, H., Yang, H., Wang, J., et al., 2019. White-matter functional networks changes in patients with schizophrenia. *NeuroImage* 190, 172–181.
- Jo, H.J., Lee, J.M., Kim, J.H., Shin, Y.W., Kim, I.Y., Kwon, J.S., Kim, S.I., 2007. Spatial accuracy of fMRI activation influenced by volume- and surface-based spatial smoothing techniques. *Neuroimage* 34, 550–564.
- Jochimsen, T.H., Ivanov, D., Ott, D.V., Heinke, W., Turner, R., Möller, H.E., Reichenbach, J.R., 2010. Whole-brain mapping of venous vessel size in humans using the hypercapnia-induced BOLD effect. *Neuroimage* 51, 765–774.
- Jones, D.K., Knösche, T.R., Turner, R., 2013. White matter integrity, fiber count, and other fallacies: the do’s and don’ts of diffusion MRI. *Neuroimage* 73, 239–254.
- Karahanoğlu, F.I., Van De Ville, D., 2015. Transient brain activity disentangles fMRI resting-state dynamics in terms of spatially and temporally overlapping networks. *Nature communications* 6, 1–10.
- Knutsson, H., 1989. Representing local structure using tensors, in: 6th Scandinavian Conference on Image Analysis, Oulu, Finland, Linköping University Electronic Press. pp. 244–251.
- Knutsson, H., Wilson, R., Granlund, G., 1983. Anisotropic nonstationary image estimation and its applications: Part i—restoration of noisy images. *IEEE Trans. Commun.* 31, 388–397.
- Li, J., Biswal, B.B., Meng, Y., Yang, S., Duan, X., Cui, Q., Chen, H., Liao, W., 2020a. A neuromarker of individual general fluid intelligence from the white-matter functional connectome. *Translational Psychiatry* 10, 1–12.
- Li, J., Biswal, B.B., Wang, P., Duan, X., Cui, Q., Chen, H., Liao, W., 2019a. Exploring the functional connectome in white matter. *Human brain mapping* 40, 4331–4344.
- Li, M., Gao, Y., Gao, F., Anderson, A.W., Ding, Z., Gore, J.C., 2020b. Functional engagement of white matter in resting-state brain networks. *NeuroImage*, 117096.
- Li, M., Newton, A.T., Anderson, A.W., Ding, Z., Gore, J.C., 2019b. Characterization of the hemodynamic response function in white matter tracts for event-related fMRI. *Nature communications* 10, 1–11.
- Li, S., Jin, Y., Shuman, D.I., 2019c. Scalable m -channel critically sampled filter banks for graph signals. *IEEE Transactions on Signal Processing* 67, 3954–3969.
- Liu, X., Duyn, J.H., 2013. Time-varying functional network information extracted from brief instances of spontaneous brain activity. *Proceedings of the National Academy of Sciences* 110, 4392–4397.
- Logothetis, N., Wandell, B., 2004. Interpreting the BOLD signal. *Annu. Rev. Physiol.* 66, 735–769.
- Lohmann, G., Stelzer, J., Lacosse, E., Kumar, V.J., Mueller, K., Kuehn, E.,

- Grodd, W., Scheffler, K., 2018. LISA improves statistical analysis for fMRI. *Nature communications* 9, 1–9.
- de Loynes, B., Navarro, F., Olivier, B., 2019. Data-driven thresholding in denoising with spectral graph wavelet transform. *arXiv preprint arXiv:1906.01882*.
- Maghsadgh, S., Eklund, A., Behjat, H., 2019. Graph spectral characterization of brain cortical morphology, in: *Proc. IEEE Int. Conf. Eng. Med. Biol. Soc.*, pp. 458–462. doi:10.1109/EMBC.2019.8856468.
- Mallat, S.G., 1989. A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Trans. Pattern Anal. Mach. Intell.* 11, 674–693.
- Marussich, L., Lu, K.H., Wen, H., Liu, Z., 2017. Mapping white-matter functional organization at rest and during naturalistic visual perception. *Neuroimage* 146, 1128–1141.
- Masoumi, M., Toews, M., Lombaert, H., 2019. WaveletBrain: Characterization of human brain via spectral graph wavelets. *arXiv preprint arXiv:1906.06158*.
- Mazerolle, E., Ohlhauser, L., Mayo, C., Sheriff, A., Gawryluk, J., 2019. Evidence of underreporting of white matter fMRI activation. *Journal of magnetic resonance imaging: JMIR*.
- Mazerolle, E.L., Beyea, S.D., Gawryluk, J.R., Brewer, K.D., Bowen, C.V., D'Arcy, R.C., 2010. Confirming white matter fMRI activation in the corpus callosum: co-localization with DTI tractography. *Neuroimage* 50, 616–621.
- Mazerolle, E.L., Gawryluk, J.R., Dillen, K.N., Patterson, S.A., Feindel, K.W., Beyea, S.D., Stevens, M.T.R., Newman, A.J., Schmidt, M.H., D'Arcy, R.C., 2013. Sensitivity to white matter fMRI activation increases with field strength. *PLoS one* 8, e58130.
- Medaglia, J.D., Huang, W., Karuza, E.A., Kelkar, A., Thompson-Schill, S.L., Ribeiro, A., Bassett, D.S., 2018. Functional alignment with anatomical networks is associated with cognitive flexibility. *Nat. Hum. Behav.* 2, 156.
- Mikl, M., Mareček, R., Hlušík, P., Pavlicová, M., Drastich, A., Chlebus, P., Brázdil, M., Krupa, P., 2008. Effects of spatial smoothing on fMRI group inferences. *Magnetic resonance imaging* 26, 490–503.
- Mishra, A., Li, M., Anderson, A.W., Newton, A.T., Ding, Z., Gore, J.C., 2020. Concomitant modulation of BOLD responses in white matter pathways and cortex. *NeuroImage*, 116791.
- Ortega, A., Frossard, P., Kovačević, J., Moura, J.M.F., Vandergheynst, P., 2018. Graph signal processing: Overview, challenges, and applications. *Proc. IEEE* 106, 808–828. doi:10.1109/JPROC.2018.2820126.
- Peer, M., Nitzan, M., Bick, A.S., Levin, N., Arzy, S., 2017. Evidence for functional networks within the human brain's white matter. *J. Neurosci.* 37, 6394–6407.
- Petridou, N., Gaudes, C.C., Dryden, I.L., Francis, S.T., Gowland, P.A., 2013. Periods of rest in fMRI contain individual spontaneous events which are related to slowly fluctuating spontaneous activity. *Human brain mapping* 34, 1319–1329.
- Preti, M.G., Van De Ville, D., 2019. Decoupling of brain function from structure reveals regional behavioral specialization in humans. *Nature communications* 10, 1–7.
- Rabiei, H., Richard, F., Coulon, O., Lefèvre, J., 2016. Local spectral analysis of the cerebral cortex: New gyrification indices. *IEEE transactions on medical imaging* 36, 838–848.
- Ruttimann, U.E., Unser, M., Rawlings, R.R., Rio, D., Ramsey, N.F., Mattay, V. S. Hommer, D., Frank, J., Weinberger, D.R., 1998. Statistical analysis of functional MRI data in the wavelet domain. *IEEE Trans. Med. Imag.* 17, 142–154.
- Rydell, J., Knutsson, H., Borga, M., 2008. Bilateral filtering of fMRI data. *IEEE J. Sel. Top. Signal Process.* 2, 891–896.
- Shuman, D.I., 2020. Localized spectral graph filter frames: A unifying framework, survey of design considerations, and numerical comparison. *arXiv preprint arXiv:2006.11220*.
- Shuman, D.I., Narang, S.K., Frossard, P., Ortega, A., Vandergheynst, P., 2013. The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Process. Mag.* 30, 83–98.
- Sihag, S., Naze, S., Taghdiri, F., Tator, C., Wennberg, R., Mikulis, D., Green, R., Colella, B., Tartaglia, M.C., Kozloski, J.R., 2020. Multimodal dynamic brain connectivity analysis based on graph signal processing for former athletes with history of multiple concussions. *IEEE Transactions on Signal and Information Processing over Networks* 6, 284–299.
- Smith, S.M., Brady, J.M., 1997. SUSAN — a new approach to low level image processing. *International journal of computer vision* 23, 45–78.
- Sporns, O., Tononi, G., Kötter, R., 2005. The human connectome: a structural description of the human brain. *PLoS Comput. Biol.* 1, e42.
- Stanković, L., Mandić, D., Daković, M., Scalzo, B., Brajović, M., Sejdić, E., Constantinides, A.G., 2020. Vertex-frequency graph signal processing: A comprehensive review. *Digital Signal Processing*, 102802.
- Tarun, A., Abramian, D., Behjat, H., De Ville, D.V., 2019. Graph spectral analysis of voxel-wise brain graphs from diffusion-weighted mri, in: *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pp. 159–163. doi:10.1109/ISBI.2019.8759496.
- Tarun, A., Behjat, H., Bolton, T., Abramian, D., Van De Ville, D., 2020. Structural mediation of human brain activity revealed by white-matter interpolation of fMRI. *NeuroImage*, 116718.
- Tournier, J.D., Calamante, F., Connelly, A., 2007. Robust determination of the fibre orientation distribution in diffusion MRI: non-negativity constrained super-resolved spherical deconvolution. *Neuroimage* 35, 1459–1472.
- Van De Ville, D., Blu, T., Unser, M., 2004. Integrated wavelet processing and spatial statistical testing of fMRI data. *Neuroimage* 23, 1472–1485.
- Van Essen, D., Smith, S., Barch, D., Behrens, T., Yacoub, E., Ugurbil, K., Consortium, W.M.H., 2013. The WU-Minn human connectome project: an overview. *Neuroimage* 80, 62–79.
- Wachinger, C., Golland, P., Kremen, W., Fischl, B., Reuter, M., Initiative, A.D.N., et al., 2015. Brainprint: a discriminative characterization of brain morphology. *Neuroimage* 109, 232–248.
- Wachinger, C., Reuter, M., Initiative, A.D.N., et al., 2016a. Domain adaptation for Alzheimer's disease diagnostics. *Neuroimage* 139, 470–479.
- Wachinger, C., Salat, D., Weiner, M., Reuter, M., Initiative, A.D.N., 2016b. Whole-brain analysis reveals increased neuroanatomical asymmetries in dementia for hippocampus and amygdala. *Brain* 139, 3253–3266.
- Wang, P., Meng, C., Yuan, R., Wang, J., Yang, H., Zhang, T., Zaborszky, L., Alvarez, T.L., Liao, W., Luo, C., et al., 2020. The organization of the human corpus callosum estimated by intrinsic functional connectivity with white-matter functional networks. *Cerebral Cortex* 30, 3313–3324.
- Weickert, J., Schar, H., 2002. A scheme for coherence-enhancing diffusion filtering with optimized rotation invariance. *Journal of Visual Communication and Image Representation* 13, 103–118.
- Wu, X., Yang, Z., Bailey, S.K., Zhou, J., Cutting, L.E., Gore, J.C., Ding, Z., 2017. Functional connectivity and activity of white matter in somatosensory pathways under tactile stimulations. *Neuroimage* 152, 371–380.
- Yarkoni, T., Barch, D.M., Gray, J.R., Conturo, T.E., Braver, T.S., 2009. BOLD correlates of trial-by-trial reaction time variability in gray and white matter: a multi-study fMRI analysis. *PLoS One* 4, e4257.
- Yeh, F.C., Verstynen, T.D., Wang, Y., Fernández-Miranda, J.C., Tseng, W.Y.I., 2013. Deterministic diffusion fiber tracking improved by quantitative anisotropy. *PLOS ONE* 8. URL: <https://doi.org/10.1371/journal.pone.0080713>, doi:10.1371/journal.pone.0080713.
- Yeh, F.C., Wedeen, V.J., Tseng, W.Y.I., 2010. Generalized q-sampling imaging. *IEEE transactions on medical imaging* 29, 1626–1635.
- Zhuang, X., Yang, Z., Curran, T., Byrd, R., Nandy, R., Cordes, D., 2017. A family of locally constrained CCA models for detecting activation patterns in fMRI. *NeuroImage* 149, 63–84.