SLEPIANS: A GENERALIZED GRAPH FOURIER TRANSFORM TO RESOLVE LOCALISED FUNCTIONAL BRAIN INTERACTIONS

Thomas AW Bolton[⋆†]*, Younes Farouj*[⋆†]*, Dimitri Van De Ville*[⋆†]

[⋆] Institute of Bioengineering, Ecole Polytechnique Federale de Lausanne (EPFL), Switzerland
[†] Department of Radiology and Medical Informatics, University of Geneva (UNIGE), Switzerland

ABSTRACT

Neuroimaging techniques revealed that the brain function exhibits coordinated patterns that can also be reflected in the anatomical structure. From a data processing point of view, this finding can be harnassed to constrain the dynamics of functional time-series to the underlying structural pathways while performing basic operations such as filtering. In this work, following recent advances in Graph signal processing, we introduce linear Slepian estimators to uncover localised functional brain interactions. The method projects the functional Magnetic Resonance Imaging (fMRI) signal onto a collection of Slepian vectors that are defined on a graph extracted, for example, from structural and diffusion data. The Slepian decomposition allows an optimal multi-bandwidth description of signals that are maximally concentrated within a subset of nodes as it is often the case for neural activity. The estimator, itself, is constructed by keeping only the Slepian coefficents that are within a certain spatial frequency range. Through some numerical experiments on simulated data, we show the advantages of using this technique over classical Laplacian and localized Laplacian filtering. To complement these results, we present an illustration on real data from the Human Connectome project (HCP) that demonstrate the potential of Slepians at retrieving localised interaction patterns in the context of a visual stimulation task.

***Index Terms*—** Graph Fourier transform, Slepians, Graph signal processing

1. INTRODUCTION

Magnetic resonance imaging (MRI) has opened a number of avenues for the study of the brain. At the structural level, diffusion MRI enables to resolve the physical connectivity that exists between different regions. Viewed from the prism of graph analysis [1], this information is embedded in an adjacency matrix, which can then be decomposed into a set of canonical structural elements, commonly termed the *eigenmodes*.

In parallel, functional MRI offers a window on brain activity over time, and on its dynamics at rest or upon cognitive challenge. There is an exquisite relationship between brain function and the underlying structural scaffold [2]; for this reason, bimodal analytical approaches that can combine those two pieces of information are particularly tailored, and an emerging topic of interest.

In particular, graph signal processing (GSP) [3] has recently gained momentum for this purpose [4]. In this framework, functional information is regarded as temporal signals on a graph defined from structural measurements, and viewed as a linear combination of eigenmodes. This has the double advantage of enabling the study of structure/function relationships, and of permitting a wide set of signal processing operations in the spectral domain to improve the quality of functional signals.

In such approaches, a similar importance is typically assigned to each brain region (i.e., each node of the studied network) at the decomposition stage. However, in some settings, it may be desirable to enhance the accuracy of the analysis on a subset of particularly important areas (for example, the brain regions expected to respond to a given paradigm). At the same time, some flexibility is also desired when it comes to defining this subset, so that rough prior knowledge be enough information to provide.

Here, we show how this can be achieved through a generalised decomposition into Slepian vectors [5]. On simulated data, we show that if provided with a set of nodes of interest, noisy functional signals can be recovered more accurately within this set compared to a standard decomposition. In addition, we show how error estimates are robust to changes in the selected subset. On a real data example, we exemplify the potential of Slepians to reveal subtler, localised interaction patterns in the context of a visual stimulation task.

2. METHODS

2.1. Graph signal processing basics

Let a graph $\mathcal{G} = (\mathcal{V}, \mathbf{W})$ characterised by the set of N nodes $\mathcal{V}$, linked as described by the symmetrical adjacency matrix $\mathbf{W} \in \mathbb{R}^{N \times N}$. In the brain application considered here, $w_{i,j}$ will be large if brain regions i and j are strongly physically connected. The resulting Laplacian matrix admits an eigendecomposition as $\mathbf{L} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^\top$ [3], with $\mathbf{V} = [\mathbf{v_1}|\mathbf{v_2}|\dots|\mathbf{v_N}]$ containing the eigenmodes as its columns, arranged in ascending eigenvalue order $\lambda_1 < \lambda_2 < \dots < \lambda_N$. Because $\lambda_k = \mathbf{v_k}^\top \mathbf{L} \mathbf{v_k} = \sum_{i \neq j} W_{i,j}([v_k]_i - [v_k]_j)^2$, eigenmodes of smaller eigenvalues will represent *low frequency* structural

patterns on the graph (i.e., for which strongly connected nodes show similar values), while eigenmodes of larger eigenvalues will denote less organised patterns with respect to the graph structure.

From this description, a signal $\mathbf{x} \in \mathbb{R}^{N\times 1}$ can conveniently be expressed in the graph domain as $\hat{\mathbf{x}} = \mathbf{V}^\top \mathbf{x}$, which is known as the *graph Fourier transform* (GFT) [6]; conversely, we also have $\mathbf{x} = \mathbf{V}\hat{\mathbf{x}}$. Each element of $\hat{\mathbf{x}}$ then represents the strength with which an eigenmode is contributing to the signal at hand.

In addition, note that the eigenmodes are also solutions to the Laplacian embedding problem, where the goal is to find a mapping of the graph nodes on a line so that connected ones stay as close as possible, that is, to find $\mathbf{x}$ as:

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmin}}\, \mathbf{x}^\top \mathbf{L}\mathbf{x} = \underset{\mathbf{x}}{\operatorname{argmin}}\, \mathbf{x}^\top \mathbf{V}\mathbf{\Lambda}\mathbf{V}^\top \mathbf{x}, \qquad (1)$$

with $\mathbf{x}^\top\mathbf{x} = 1$ and $\mathbf{x}^\top\mathbf{1} = 0$ [7].

2.2. Slepian vectors

The goal is to derive an alternative set of basis vectors, under the constraints that (1) their energy should be localised within a predefined subset of nodes $\mathcal{S}$, and (2) they should be derived from a bandwidth-limited subset of original eigenvectors, with dimension $K \leq N$ [8]. To highlight the selected nodes, we define $\mathbf{M}$ as the diagonal matrix with $m_{i,i} = 1$ if a node is selected in the subset, and 0 otherwise. Further, we define $\mathbf{V_T} \in \mathbb{R}^{N\times K}$ as the trimmed set of eigenmodes.

The concentration to optimise is given by $\mu = \frac{\hat{\mathbf{x}}^\top \mathbf{C}\hat{\mathbf{x}}}{\hat{\mathbf{x}}^\top\hat{\mathbf{x}}}$, with $\mathbf{C} = \mathbf{V_T}^\top \mathbf{M}\mathbf{V_T}$. For the analogy with the classical eigenmodes, where λ_i represents a spatial frequency on the graph, we consider an alternative formulation where, using the equality $\mathbf{\Lambda} = \mathbf{\Lambda}^{1/2}\mathbf{V}^\top\mathbf{V}\mathbf{\Lambda}^{1/2}$, the solution vectors $\mathbf{s_k}$, $k = 1, \ldots, K$ satisfy a generalised Laplacian embedding formulation [8]:

$$\hat{\mathbf{s}}^* = \underset{\hat{\mathbf{s}}}{\operatorname{argmin}}\, \hat{\mathbf{s}}^\top {\mathbf{\Lambda_T}}^{1/2}\mathbf{C}{\mathbf{\Lambda_T}}^{1/2}\hat{\mathbf{s}}. \qquad (2)$$

The *local frequency* within $\mathcal{S}$ is denoted ξ_i, and the matrix $\mathbf{S} = [\mathbf{s_1}|\mathbf{s_2}|\ldots|\mathbf{s_K}]$ contains the solution vectors, arranged in ascending local frequency ($\xi_1 < \xi_2 < \ldots < \xi_M$). Note that ξ_i will be low either if $\mathbf{s_i}$ is not concentrated within $\mathcal{S}$, or if it displays a low local frequency.

2.3. Linear Slepian estimators

Let $\mathbf{X} \in \mathbb{R}^{N\times T}$ be a signal of length T defined on each node of the graph. At any given time point t, the current value of the signal, $\mathbf{X_t}$, can be projected onto the Slepian basis formed by the collection of vectors $\mathbf{s_k}$, $k = 1, \ldots, K$.

$$\mathbf{S}^\top\mathbf{X_t} = <\mathbf{X_t}, \mathbf{s_i}>_{i=1,\ldots,K}. \qquad (3)$$

Now, assuming that $\mathbf{X_t}$ is composed of a smooth component (i.e., the signal of interest) and of a noise component of larger frequency, it is possible to truncate the projection $\mathbf{S}^\top\mathbf{X_t}$ up to a certain range without loss of information, by discarding the Slepian vector coefficients associated to the largest ξ_i. At the same time, one can also focus on the subset $\mathcal{S}$ by discarding the non-concentrated Slepian vectors ($\mu_i < \epsilon$). Formally, we retrieve the output $\mathbf{Y_t^{(\bar{\xi})}} \in \mathbb{R}^{N\times 1}$ as:

$$\mathbf{Y_t^{(\bar{\xi})}} = \mathbf{S}\mathbf{H}^{(\bar{\xi})}\mathbf{S}^\top\mathbf{X_t}, \qquad (4)$$

where $\bar{\xi}$ is the cut-off frequency and $\mathbf{H}^{(\bar{\xi})}$ is a diagonal matrix with $H_{i,i}^{(\bar{\xi})} = 1$ if $\xi_i < \bar{\xi}$ and $\mu_i > \epsilon$, and 0 otherwise. We call this estimator *linear* in analogy with classical estimators on regular domains (DCT [9], wavelets [10]), as it depends only on the cut-off frequency and does not involve any point-wise thresholding procedure.

2.4. Evaluations on simulated data

We examined how well a simulated ground truth signal $\mathbf{X} \in \mathbb{R}^{N\times T}$ on $\mathcal{G}$, corrupted with noise, could be retrieved on a subset of nodes $\mathcal{S}$ through graph filtering. We considered simulated modular graphs of $N = 200$ nodes, with 4 communities, a minimal community number of 40 nodes, and world density $\frac{1}{N}$. We created simulated time courses of $T = 900$ time points, where the largest graph community was selected as $\mathcal{S}$. Within $\mathcal{S}$, we created two independent temporal paradigms of 180 time points, each occurring in half of the subset nodes. Each of the other 3 communities was also assigned a separate paradigm time course (see Figure 1A).

To add noise on top of the ground truth time courses, at each time point, a random subset of n_r nodes was chosen, and corrupted with noise of intensity $I \sim \mathcal{N}(0, \sigma_r^2)$ (see Figure 1B). We assessed the ability of (1) a classical basis of eigenmodes (GFT), (2) a localised eigenmodes decomposition on the subset $\mathcal{S}$ (LGFT), and (3) a Slepian basis focused on the same subset (SLEP) to retrieve the ground truth. We probed ranges of $n_r \in [0 : 20 : 200]$ and $\sigma_r \in [0 : 0.5 : 4]$, simulating 20 datasets in each case. We computed the average mean squared error (MSE) across all nodes within $\mathcal{S}$ for all possible filtering cutoffs, and selected the minimal MSE as our error measure.

In addition, we compared the LGFT and SLEP decompositions in their sensitivity to an imperfect subset selection, where for each of 10 simulated trial, $\mathcal{S}_i$ ($|\mathcal{S}_i| = 0.2|\mathcal{S}|$) was picked as the subset of interest, and the mean error across those trials was considered.

2.5. Exploratory application to real fMRI data

To extend our observations on simulated data to the fMRI setting, we consider one subject (ID:100307) from the Human Connectome Project initiative [11], for which we down-

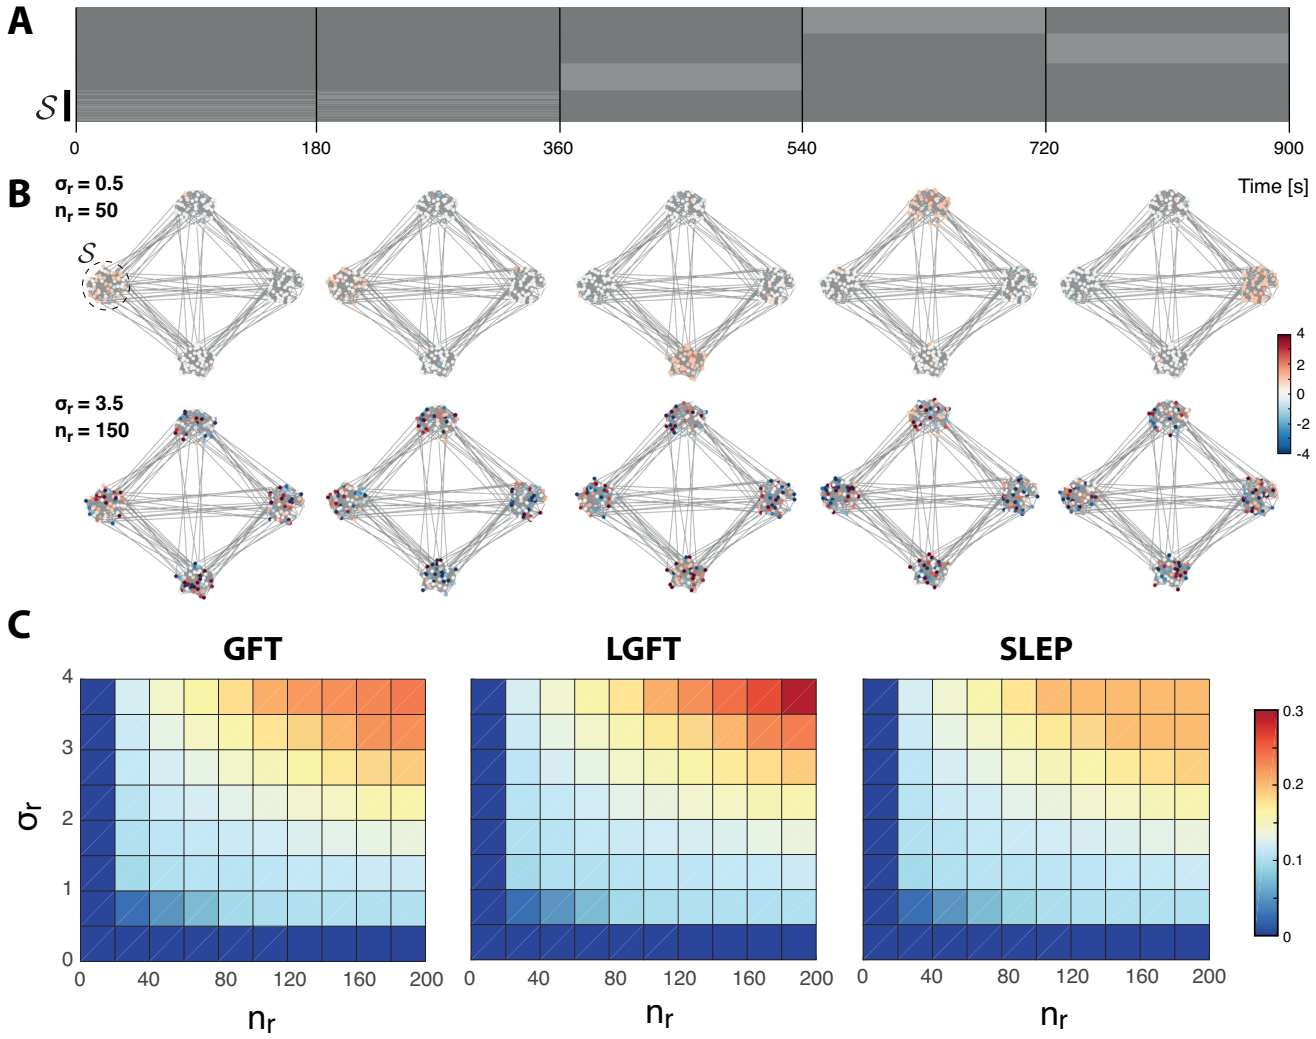

Fig. 1. (A) Simulated paradigm for all nodes (top to bottom) across time (left to right). Light gray denotes activation, and $\mathcal{S}$ reflects the selected subset. (B) In low noise (top) or hight noise (bottom) cases, example signals generated at different moments from the paradigm, on a simulated graph. (C) For GFT (left), LGFT (middle) or SLEP (right) cases, evolution of median MSE for increasing noise levels.

loaded diffusion MRI, structural MRI and task fMRI (*Working Memory*) data. We created $\mathbf{W}$ from the structural and diffusion data, using the MRtrix toolbox (multi-shell multi-tissue response function estimation, spherical deconvolution, tractogram generation with 10^7 output streamlines between 838 regions from the Craddock atlas [12]).

The considered working memory task fMRI recording included blocks of fixation, and of image presentation (faces, places, tools or body parts), in 0-back or 2-back fashion [13]. Here, we focused on the visual aspect of the paradigm, and selected a subset $\mathcal{S}$ of 3 nodes from the occipital brain (see Figure 2A). We generated Slepian graph coefficients at a bandwidth of 180, and truncated them with $\epsilon = 10^{-2}$, yielding 2 remaining coefficients at each time point. We compared this two-dimensional representation of functional brain activity to the one achieved from the coefficients of the 2 GFT or LGFT eigenmodes of lowest frequency.

3. RESULTS

3.1. Evaluations on simulated data

Across trials, MSE increased with larger n_r and larger σ_r, for all three evaluated methods (GFT, LGFT, SLEP; Figure 1C). Performance was equivalent at low noise levels, but in the noisiest cases (top right part of the matrices), a better ground truth recovery was achieved with Slepians. For instance, for $n_r = 200$ and $\sigma_r = 4$, median MSE values were $MSE_{GFT} = 0.26$, $MSE_{LGFT} = 0.36$, and $MSE_{SLEP} =$

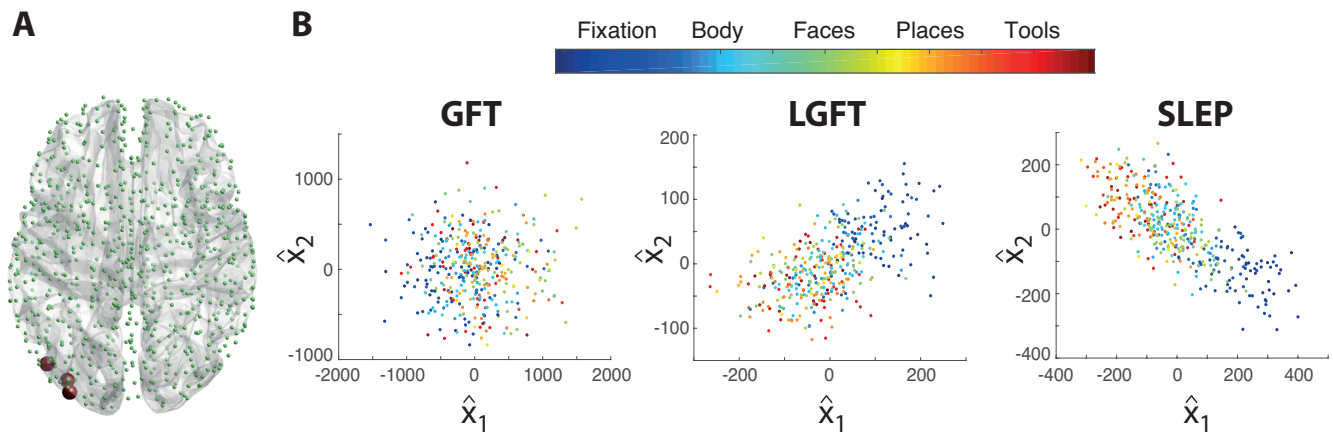

Fig. 2. (A) Brain depiction of the considered set of nodes, with the selected subset $\mathcal{S}$ highlighted in red. (B) For the GFT (left), LGFT (middle) and SLEP (right) cases, evolution over time of the signal projected on the two lowest frequency eigenmodes/Slepian vectors. Colour coding reflects the type of stimulus presented at each time point (note that the paradigm regressor was convolved with a hemodynamic response function estimate to retrieve an fMRI-compatible timing).

0.2.

In addition, the LGFT approach was less robust than the SLEP one to an incomplete $\mathcal{S}$ selection: for example, for still moderate noise levels of $n_r = 100$ and $\sigma_r = 2$, median MSE in the LGFT case was 0.29 (compared to 0.13 for a full subset selection), whereas it stayed at 0.14 in the SLEP case (against 0.13 for a full subset selection).

3.2. Exploratory application to real fMRI data

Plotting the spectral coefficients linked to the 2 lowest frequency GFT eigenmodes (Figure 2B, left plot), no clear discrimination could be achieved between the moments when different types of visual stimuli were presented. In the localised GFT case (middle plot), moments of fixation (dark blue) could be segregated from visual presentation. In the Slepians case (right plot), on top of distinguishing fixation from visual presentation, the presentation of body parts or faces (light blue or green data points) could be separated from moments when places or tools (orange or red data points) were shown.

4. DISCUSSION

GSP tools are gaining a lot of interest in the neuroimaging community, but the strong noise levels at play in the functional data at hand complicate their use. In this work, we demonstrated how a projection onto a collection of vectors called Slepians opposition to many research fields where GSP finds its use, brain gra

GSP tools are gaining a lot of interest in the neuroimaging community, but the strong noise levels at play in the functional data at hand complicate their use.

Discuss the fact that some types of noise (here, several nodes spanning the whole network activate together, but not following the graph structure; i.e., two very close nodes in the graph are equally likely to activate compared to two remote nodes) can impede the classical GSP tools. Slepians enable to focus the analysis on a limited subpart of the graph, and so, to an extent, get rid of this noise effect.

Discuss the fact that at the same time, the bandwidth parameter enables more robustness to imperfect node selection, which is nice because in neuroscience, we do not necessarily know all interesting foci of activation beforehand. The thing is that the lower bandwidth enables us to use *structural* information from the graph, in order to facilitate *functional* data processing. In itself, this is a really cool way to combine modalities, which is also an emerging trend.

Discuss possible follow-up extensions of the framework (e.g., have a set of nodes that we want to be active, and another that we want to be deactive, when defining $\mathbf{M}$. If possible, find other nice applications of those tools...

5. REFERENCES

[1] Y Iturria-Medina, EJ Canales-Rodriguez, L Melie-Garcia, PA Valdes-Hernandez, E Martinez-Montes, Y Alemán-Gómez, and JM Sánchez-Bornot, "Characterizing brain anatomical connections using diffusion weighted mri and graph theory," *Neuroimage*, vol. 36, no. 3, pp. 645–660, 2007.

[2] Ed Bullmore and Olaf Sporns, "Complex brain networks: graph theoretical analysis of structural and functional systems," *Nature reviews. Neuroscience*, vol. 10, no. 3, pp. 186, 2009.

[3] David I Shuman, Sunil K Narang, Pascal Frossard, Antonio Ortega, and Pierre Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *IEEE Signal Processing Magazine*, vol. 30, no. 3, pp. 83–98, 2013.

[4] Weiyu Huang, Leah Goldsberry, Nicholas F Wymbs, Scott T Grafton, Danielle S Bassett, and Alejandro Ribeiro, "Graph frequency analysis

of brain signals," *IEEE Journal of Selected Topics in Signal Processing*, vol. 10, no. 7, pp. 1189–1203, 2016.

[5] David Slepian and Henry O Pollak, "Prolate spheroidal wave functions, fourier analysis and uncertaintyi," *Bell Labs Technical Journal*, vol. 40, no. 1, pp. 43–63, 1961.

[6] Aliaksei Sandryhaila and José MF Moura, "Discrete signal processing on graphs," *IEEE transactions on signal processing*, vol. 61, no. 7, pp. 1644–1656, 2013.

[7] Mikhail Belkin and Partha Niyogi, "Laplacian eigenmaps for dimensionality reduction and data representation," *Neural computation*, vol. 15, no. 6, pp. 1373–1396, 2003.

[8] Dimitri Van De Ville, Robin Demesmaeker, and Maria Giulia Preti, "When slepian meets fiedler: Putting a focus on the graph spectrum," *IEEE Signal Processing Letters*, vol. 24, no. 7, pp. 1001–1004, 2017.

[9] Leonid P Yaroslavsky, Karen O Egiazarian, and Jaakko T Astola, "Transform domain image restoration methods: review, comparison, and interpretation," in *Photonics West 2001-Electronic Imaging*. International Society for Optics and Photonics, 2001, pp. 155–169.

[10] David L Donoho and Jain M Johnstone, "Ideal spatial adaptation by wavelet shrinkage," *biometrika*, vol. 81, no. 3, pp. 425–455, 1994.

[11] David C Van Essen, Stephen M Smith, Deanna M Barch, Timothy EJ Behrens, Essa Yacoub, Kamil Ugurbil, Wu-Minn HCP Consortium, et al., "The wu-minn human connectome project: an overview," *Neuroimage*, vol. 80, pp. 62–79, 2013.

[12] R Cameron Craddock, G Andrew James, Paul E Holtzheimer, Xiaoping P Hu, and Helen S Mayberg, "A whole brain fmri atlas generated via spatially constrained spectral clustering," *Human brain mapping*, vol. 33, no. 8, pp. 1914–1928, 2012.

[13] Deanna M Barch, Gregory C Burgess, Michael P Harms, Steven E Petersen, Bradley L Schlaggar, Maurizio Corbetta, Matthew F Glasser, Sandra Curtiss, Sachin Dixit, Cindy Feldt, et al., "Function in the human connectome: task-fmri and individual differences in behavior," *Neuroimage*, vol. 80, pp. 169–189, 2013.