Thesis Proposal: Using structured sparse regression to model functional brain connectivity with anatomical constraints

Brian Kent
Adviser: Alessandro Rinaldo
Committee: Christopher Genovese, Robert Kass,
Aarti Singh (Machine Learning), Timothy Verstynen (Psychology)

June 4, 2012

Abstract

Many common methods for estimating functional brain connectivity are limited because they are bivariate, descriptive, and non-sparse. Furthermore, the functional magnetic resonance imaging used to fit these models has poor temporal resolution, causing spurious connections to be included. The first set of problems can be overcome by learning functional network structure with penalized regression models, which are multivariate, sparse, and predictive. To address the problem of spurious connections we use anatomical networks—estimated with diffusion spectrum imaging and high-definition fiber tractography—as soft constraints on functional connectivity. In the first phase of this research we define anatomical and functional network nodes by clustering fiber tract terminal points. In the second, we investigate how high-dimensional regression methods can be used to estimate the edges in a functional network in a way that reflects the underlying anatomical structure.

1 Introduction

Many cognitive tasks are thought to arise from the characteristics and dynamics of distributed networks of brain regions. Modeling these networks accurately will allow us to better understand interesting brain functions and potentially to diagnose pathologies (McIntosh, 2000; Honey et al., 2009).

Functional brain networks are often modeled as graphs. First, a set of regions-of-interest (ROIs) are defined as the graph nodes and the activation at each ROI is measured over time, typically with functional magnetic resonance imaging (fMRI). The graph edges are most commonly estimated using a measure of pairwise statistical association like correlation, partial correlation, coherence, or mutual information (Bullmore and Sporns, 2009; Smith

et al., 2011). This standard modeling process has several limitations. Specifically, it is intrinsically descriptive and does not allow for statistical or scientific validation through predictive checks. Furthermore, it does not identify subtle graph structures like loops or hierarchies very well and the lack of sparsity makes scientific interpretation more difficult in general.

The biggest drawback to fMRI-based functional connectivity estimates, however, is the fMRI modality itself. fMRI typically has a temporal resolution of 0.5-1 Hz, while neural signals may take less than 100 ms to travel from one ROI to the next. In addition, fMRI measures the level of blood oxygenation, which is not a direct indicator of neural activity. As a result, fMRI-based functional network estimates likely include connections that have no anatomical basis (in other words, type I errors).

Incorporating anatomical connectivity. Our hypothesis is that the connections in a functional network are constrained by an underlying anatomical network, and that we can improve functional connectivity models by explicitly including anatomical information. To model anatomical networks we use diffusion spectrum imaging (DSI) (Hagmann et al., 2006; Wedeen et al., 2005) and high-definition fiber tractography (HDFT) (Verstynen et al., 2011; Jarbo et al., 2012). DSI indicates the directions of maximum water displacement in each voxel of a subject's brain; by connecting adjacent voxels according to these primary directions any desired number of *fiber tracts* can be constructed. Because water tends to diffuse more along the major axis of neurons, the fiber tracts indicate the location of major white matter fiber bundles. Figure 1 shows an example of fiber tracts constructed in the cortico-basal ganglia system.

The most direct way that anatomical networks can inform functional network estimation is through measures of the structural connection strength between two network nodes. This can be as simple as the number of fibers that connect a pair of ROIs, although more sophisticated statistics like quantitative anisotropy are typically available because they must be computed for the fiber tractography process (Yeh et al., 2010).

We can also use fiber tractography to influence functional connectivity estimates by using clusters of fiber terminal points to define the functional network nodes. Defining these nodes correctly is critical for accurate connectivity estimates (Smith et al., 2011), and there are several benefits to defining the nodes with HDFT data. First, fiber tractography has much better resolution than the most common node identification method of projecting a generic brain atlas onto each individual subject. ROIs can also be defined directly from the fMRI signal, but fMRI voxels are relatively large and may not capture complicated structural patterns. Fiber tract terminal points have a sub-voxel resolution, and we can represent this in functional network nodes by having each fMRI voxel contribute to graph nodes in proportion to the cluster representation in the voxel. In addition, fMRI measures the resting state blood oxygen level-dependent (rsBOLD) contrast, which is a noisy, indirect indicator of neural activity. DSI provides more direct insight into neural connections.

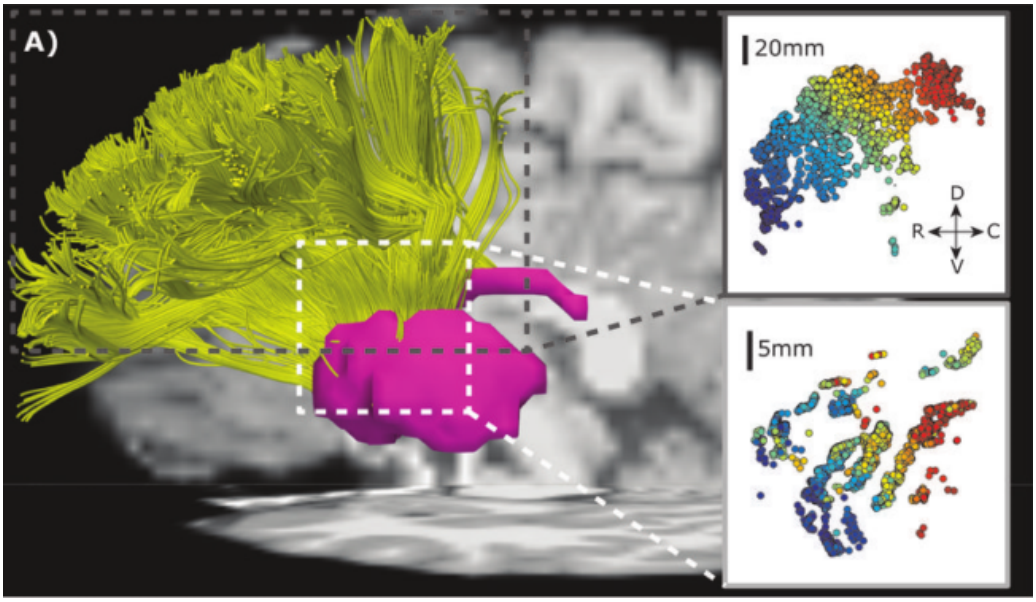

Figure 1: An example of HDFT fibers constructed in the cortico-basal ganglia system. Fibers are shown in yellow and the basal ganglia is purple. The top inset shows the spatial arrangement of the cortical terminal fiber points, and bottom inset shows the striatal terminal points (color indicates depth). *Source: Verstynen et al. (2012)*

There are drawbacks to defining graph nodes with structural networks. Most notably, deterministic HDFT is prone to type II errors; small clusters of fibers are often overlooked due to random noise and partial volume effects in the DSI process (Lemkaddem et al., 2012). This phenomenon may explain some spurious functional connections, which is why our proposed methods down-weight functional edges that lack an anatomical link but do not rule them out entirely.

Clustering fiber tract endpoints. Fiber tract endpoints are often arranged in highly irregular patterns, so defining graph nodes according to endpoint clusters is a challenging task. For example, fiber terminal points in the striatum lie along curved stripes (shown clearly in the bottom right panel of Figure 1), corresponding to projections from several different cortical regions. This complex anatomy suggests a complex arrangement of functional assignments, and it is critical that each stripe be identified as a unique cluster so that the stripes become vertices in our functional network (Verstynen et al., 2012). Simple clustering algorithms like k-means fail to capture these irregular endpoint groupings, but techniques like spectral clustering, diffusion maps, and density clustering show more promise.

Both anatomical and functional networks can be defined and modeled at different scales. Figure 2 from Jarbo et al. (2012) illustrates anatomical connectivity across the corpus callosum; there are clearly patterns at both the lobe and sub-lobe scales. Similarly, Verstynen et al. (2012) find that fiber tracts projecting into the basal ganglia can be segregated on at least two levels. Ideally we would like to estimate functional connectivity simultaneously at multiple scales, using a hierarchy of fiber tract clusters to define the functional network nodes. Each level of this hierarchy would correspond to a partition of the set of fibers, such that lower levels are refinements of the partitions on higher levels.

Penalized regression. We use penalized regression methods to estimate functional connectivity networks; this family of models offers several advantages for this particular application. First and foremost, the penalty function in many models can be modified to incorporate anatomical connection strength as a soft constraint on functional connectivity. In addition to encouraging functional models to conform to a subject's anatomy, the sparsity induced by the penalty functions make the models statistically valid and scientifically tractable.

The regression framework is also multivariate and predictive, overcoming two of the main pitfalls of the correlation-type methods. Having predictive models is especially useful because the connectivity estimates can be validated both by statistical methods like cross-validation and also through repeat scans of the same subjects.

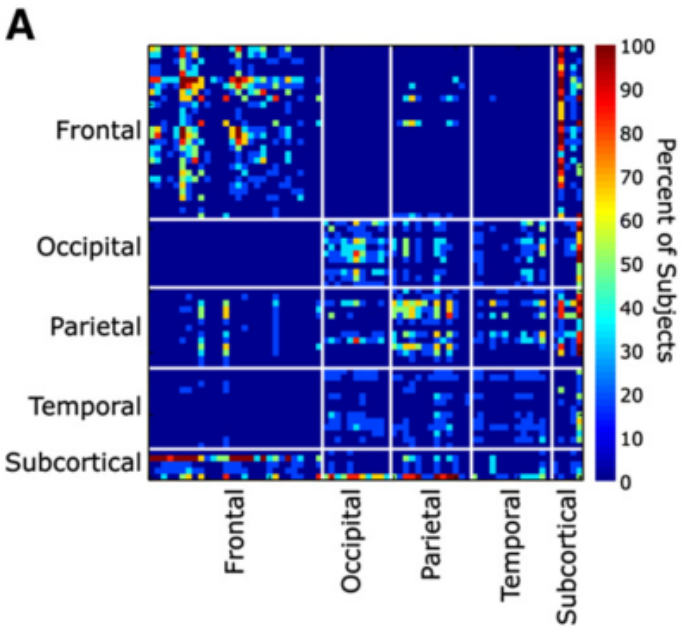

Figure 2: Anatomical connectivity across the corpus callosum for pairs of ROIs, measured as the fraction of subjects with a significant connection. Rows indicate left hemisphere ROIs and columns indicate right hemisphere ROIs. There is clearly structure at both the lobe and sub-lobe scales. *Source: Jarbo et al. (2012)*

2 Literature Review

Estimating the anatomical network. Numerous methods exist for clustering fiber tracts and estimating the connectivity of fiber bundles. The following examples are meant to illustrate the breadth of approaches. Cammoun et al. (2011) construct a hierarchical map of structural connectivity by defining about 1000 ROIs on a template brain and recursively combining these *by hand* into 66 clusters. They map these template groups onto fiber tracts for an individual subject and use the sum of inverse fiber lengths to estimate pairwise cluster connectivity. The authors note that a multi-scale approach is important for accurate structural mapping because of the large variation in the size of fiber bundles and the anatomical regions they connect.

Guevara et al. (2011) implement a hierarchical fiber tract clustering routine that avoids manual ROI labeling but requires a byzantine cascade of steps. A whole brain set of about 1.5 million fiber tracts is subdivided into major regions according to an atlas, then fiber groups are created within each region according to fiber length. These groups and their pairwise connectivities are inputs to a hierarchical clustering of white matter voxel space into parcels, but these parcels often have substantial overlap between different fiber bundles. To fix this, the endpoints of fibers are clustered within a group into fascicles. The combination of all of these subdivisions can lead to a sort of over-fitting, so fascicle centroids are computed and clustered hierarchically. This process is shown to work well in simulated fiber tract data, real fiber sets, and phantoms but it does not lend itself to statistical analysis in any way.

As a preprocessing step for functional network node localization, Zhu et al. (2011) use principal components analysis to transform fiber tract bundles into collections of unit normalized vectors that indicate the bundles' primary direction and shape. They use affinity propagation to group several hundred fiber bundles constructed in a local neighborhood of each proposed functional ROI into three clusters.

Estimating functional connectivity. Estimating fMRI-based functional networks requires definition of network nodes and selection of the edge set. Standard approaches to the first problem include expert manual labeling of ROIs, projection of a brain template onto individuals, statistical decomposition of fMRI signals into independent or variance-maximizing components, and task-based fMRI responses (Cammoun et al., 2011; Zhu et al., 2011). Methods for determining the edge set can be divided into undirected and directed graph structure learning techniques. The latter category includes Bayes nets, lag-based approaches such as Granger causality, and methods that look for asymmetries in pairwise conditional probabilities (Smith et al., 2011). The sparse regression methods that we will investigate produce undirected graphs and we focus on comparable methods.

The simplest way to measure pairwise relationships is the Pearson correlation coefficient

of the time series at two nodes (we refer to this as *full correlation*). This method has several major shortcomings, most notably that it often identifies indirect connections as direct (Smith et al., 2011). To overcome this problem, the *partial correlation* method finds the correlation between two vertices after the effects of all other nodes have been regressed out. Correlation can be replaced by any statistic that measures pairwise dependencies, such as mutual information.

Both full and partial correlations can fail to detect a true but lagged relationship. *Coherence* is a temporal generalization of correlation that can overcome this problem and can allow non-stationarity to be detected. Pairwise coherence can be computed using wavelet transforms or by estimating cross-spectral densities at various frequencies (Smith et al., 2011).

Dynamic causal modeling (DCM) and *structural equation modeling* (SEM) have also been used to estimate undirected graph structure. Smith et al. reject these methods because they do not allow for a computationally efficient search across all possible network topologies and are intended to be used to compare only a few hypotheses about network structure. DCM is particularly inappropriate because it assumes pre-specified input times, but we are estimating a network of resting state functional connectivity where there are no input times (Smith et al., 2011).

Smith et al. (2011) test the ability of each of these methods to detect graph edges and edge directions, and the robustness of the methods to data problems. Using simulated fMRI BOLD data, they find partial correlation and *regularized precision matrix* methods are the best for detecting true network connections, estimating about 90% of true edges to have higher weight than the 95th percentile strength of false positive edges. The full correlation method is not far behind, but mutual information- and coherence-based techniques perform poorly. Interestingly, the Bayes net approach is also effective for discovering graph edges, but not for identifying edge directions.

Relating anatomical and functional connectivity. The principle underlying this research is that anatomical connectivity and functional connectivity are related, and a handful of papers show that this relationship is indeed plausible. Greicius et al. (2009) use fRMI data to estimate functional connectivity and diffusion tensor imaging (DTI) to estimate anatomical connectivity, and find that the former "reflects" the latter in the default mode network across 23 subjects.

Honey et al. (2009) analyze the structure-function relationship with the same data types that we use—DSI and fMRI—but use only pairwise correlation of the fMRI to estimate connection strength between two functional graph nodes. They find that strong structural connections are predictive of strong functional connections, but that strong functional connections can occur without direct anatomical links.

Zhu et al. (2011) improve the accuracy of functional node localization by incorporating

HDFT information, although they do this by maximizing the inter-subject consistency of HDFT fiber bundles that intersect a proposed ROI. For this project we analyze data within individual subjects, so this method is not feasible.

Bunea et al. (2011) use the lasso and elastic net methods to add anatomical information to a model of brain output, but in a much different framework than our approach. The response variables are cognitive function test scores, which are regressed on a variety of covariates including ROI volumes and fiber tractography functions (such as fractional anisotropy and mean diffusivity), as well as clinical measurements and demographic categories.

Penalized regression for graph structure learning. Meinshausen and Bühlmann (2006) describe how penalized regression can be used to estimate the edges in an undirected graphical model. The neighborhood of node a in vertex set Γ (denoted ne_a) is the smallest set of nodes such that a is conditionally independent of the remaining nodes given a's neighborhood. These are precisely the nodes that are adjacent to a in a graphical model representation of random vector X, and when X is Gaussian these nodes are indicated by nonzero entries in the a'th row of the inverse covariance matrix (Lauritzen, 1996). Estimating the presence of these edges from n observations of X can be achieved with ℓ_1-penalized regression. Let X_a be the n-dimensional vector of observations of the a'th column of the data matrix $\mathbf{X}$, $\mathbf{X}_{(a)}$ be the matrix of n observations (rows) of the other $p-1$ random variables in X, and β^a be a length $p-1$ vector. Then let

$$\widehat{\beta}^{a,\lambda} = \arg\min_{\beta} \left\{ \|X_a - \mathbf{X}_{(a)}\beta\|_2^2 + \lambda\|\beta\|_1 \right\}.$$

The neighborhood estimate for node a is

$$\widehat{\text{ne}}_a^{\lambda} = \left\{ b \in \Gamma : \beta_b^{a,\lambda} \neq 0 \right\}$$

and edge set estimate is

$$\widehat{E}^{\lambda} = \left\{ (a,b) : a \in \widehat{ne}_b^{\lambda} \text{ OR } b \in \widehat{ne}_a^{\lambda} \right\}.$$

We refer to this method as *parallel lasso*. Meinshausen and Bühlmann show that under a set of assumptions about dimensionality ($p >> n$), non-singularity of the covariance matrix, sparsity, effect size, and neighborhood stability, then both the neighborhood selection and the full edge set estimation procedures are consistent. The tuning parameter λ can be chosen to control the probability of falsely connecting two distinct connected components in the true graph.

The strong assumptions required by Meinshausen and Bühlmann are a drawback to this method. More general research on the lasso method establishes that the *irrepresentible conditions*—which preclude irrelevant covariates from being correlated with relevant

ones—are sufficient to for ℓ_1 regularization to recover a sparse support set (Zhao and Yu, 2006; Meinshausen and Yu, 2009; Zhang, 2011b). Unfortunately, these conditions are not generally verifiable, and given the large impact of hemodynamic noise on fMRI data, they are likely not satisfied in our application.

Zhao et al. (2009) develop a family of linear regression penalty functions called composite absolute penalties (CAP) that incorporates prior knowledge about covariate structure. In particular, penalty functions can be created that represent hierarchical structures, with the most natural example being ANOVA regression where interaction terms should not enter the model until lower-order terms have entered.

Suppose there are K groups of covariates, and g_k is a subset of the covariate indices for $k \in 1, \ldots, K$. Let γ_k be the desired norm for group k and γ_0 be an "overall norm". The coefficient vector for group k is $\beta_{g_k} = (\beta_j)_{j \in g_k}$. The CAP penalty is

$$T_{g,\gamma}(\beta) = \sum_{k=1}^{K} \left(\|\beta_{g_k}\|_{\gamma_k} \right)^{\gamma_0},$$

and the full penalized regression problem is

$$\widehat{\beta}_{g,\gamma}(\lambda) = \arg\min_{\beta} \left\{ L(X, Y, \beta) + \lambda T_{g,\gamma}(\beta) \right\}.$$

If $\gamma_i \geq 1$ for all i then the penalty function is convex which allows for efficient computation of solutions (under typical loss functions). The authors demonstrate that CAP penalties induce hierarchical sparsity if the overall norm γ_0 is set to 1 and γ_k is set to be greater than 1 for each group. If a set of variables g_1 should enter the model before another set g_2, then g_2 appears in every penalty term where g_1 appears, and there should be at least one penalty term where g_2 is included but g_1 is not. Under these conditions, once g_2 enters the model g_1 almost surely enters as well.

The hierarchical version of this family of penalty functions assumes the average correlation between predictors and the response increases monotonically when moving from the leaves to the root. This is a natural assumption for the ANOVA regression setting but it is the opposite of the brain anatomy structure we would like to characterize. That is, we would like small clusters of brain nodes to enter our model only if the larger aggregate clusters have already entered. Furthermore, Zhao et al. show in simulation studies that CAP-penalized regression performs worse than regular lasso when the interaction effects are very strong.

The CAP penalty is one of many structured sparsity-inducing regularizations that have been proposed in the past decade. Other prominent examples include the *adaptive lasso*, which uses weights on the elements of the ℓ_1 norm (Zou, 2006); *adaptive bridge* regression, which uses weights in the penalty function and allows any p-norm (Kawano, 2012); *group*

lasso, which has an ℓ_1 penalty on groups of covariates and an ℓ_2 penalty on individual features (Yuan and Lin, 2006); and *k-overlap norm* regression, which estimates a model whose support is a union of covariate groups (Argyriou et al., 2012; Jacob et al., 2009).

Stepwise procedures constitute another family of methods that can be used to obtain sparse estimates in high-dimensional settings. Zhang (2011a) proposes an adaptive forward-backward stepwise algorithm for high-dimensional regression and demonstrates that this method performs better than lasso under certain conditions. Percival et al. (2011) use a forward stepwise regression procedure to achieve structured sparsity in estimating the association between protein mutations and HIV drug resistance. In each iteration, the covariates that are not yet included in the model are weighted according to their proximity to included covariates and their correlation with the current model's residuals. The candidate covariate with the highest weight is then added to the model. This algorithm gives sparse estimates of the coefficient vector such that non-zero components of the coefficient vector tend to cluster together. Percival et al. show their algorithm is consistent in the sense that with high probability all selected predictors are nonzero in the true target coefficient vector and the number of false negatives is bounded by the number of small true nonzero predictors, but the authors note that stepwise regression is a greedy algorithm unlike the convex program of the lasso.

3 Preliminary Results

3.1 Data and model framework

Our data consist of DSI fiber tracts and fMRI time series for several subjects imaged in the past year. We analyze each subject on an individual basis. The number of fiber tracts varies depending on the particular brain region under investigation, but is typically about $10,000$. To estimate the anatomical network we focus primarily on fiber tract terminal points. Denote $x^s \in \mathbb{R}^3$ to be the subcortical terminus and $x^c \in \mathbb{R}^3$ to be the cortical end. Let X^s be the set of all n subcortical points and X^c be the set of all cortical points. The full HDFT data set is $\mathcal{X} = X^s \cup X^c$.

The functional data is a set of time series of the rsBOLD fMRI contrast, one for each voxel in a standardized $96 \times 96 \times 51$ whole-brain grid. We represent each voxel as a point $y \in \mathbb{N}^3$ so that $\{y_t\}_{t=1}^T$ is the rsBOLD time series at voxel y.

To map the physical space into voxel space we define a function $g : \mathbb{R}^3 \to \mathbb{N}^3$. In practice this mapping is done in the fiber tractography process, before we see the data. The set of endpoints in voxel y is $g^{-1}(y) = \{x \in \mathcal{X} : g(x) = y\}$. A clustering function is also defined that assigns a label to each point: $c : \mathbb{R}^3 \to \{1, \ldots, K\}$, where K—the number of clusters—can either be fixed in advance or chosen based on the data. Applying

this function to all the fiber endpoints in voxel y yields a cluster assignment vector for y: $c(g^{-1}(y))$.

Functional network nodes are formed using the voxel cluster vectors. In the simplest set-up each voxel is labeled according to its modal cluster and the voxels with the same label are grouped into a graph node. Suppose v is the vertex that corresponds to cluster k. Then

$$v = \big\{y : \text{mode}\{c(g^{-1}(y))\} = k\big\}.$$

Each node covers many voxels and their rsBOLD time series; these are averaged to create a single time series for the node:

$$v_t = \frac{1}{n_v} \sum_{y \in v} y_t.$$

where n_v is the number of voxels in node v.

Fiber tractography provides many ways to estimate the strength of the connection between two ROIs. The simplest methods are based only on the number of fibers that connect a pair of nodes. For example, let n_u and n_v be the number of fibers with endpoints in nodes u and v, and let $n_{uv} = n_{vu}$ be the number of fibers that have one end in node u and the other end in node v. The degree of connectivity is

$$\alpha_{uv} = \frac{n_{uv}}{n_u n_v}.$$

3.2 Clustering fiber endpoints

To define functional network nodes we cluster fiber tract terminal points. For our initial exploration of endpoint clustering we analyze three fiber data sets separately. The first has 5,000 fibers where x^s falls in the caudate nucleus, the second has 10,000 fibers such that x^s is in the putamen, and the third has 10,000 fibers with x^s located in the thalamus.

Because of the striatal pattern of the fiber endpoints, particularly in the putamen, simple clustering algorithms like k-means do not work well. Instead we use a spectral clustering algorithm, following the 2006 tutorial by von Luxburg (2006). A substantial portion of our preliminary work has been investigating and optimizing the performance of this technique. For each of the six endpoint groups (three cortical and three subcortical) we construct a k-nearest neighbor similarity graph and represent this graph as an adjacency matrix W. The edge weights are

$$w_{ij} = \begin{cases} \exp\left\{-\frac{\|x_i - x_j\|^2}{2c}\right\} & x_i \text{ and } x_j \text{ are k-neighbors} \\ 0 & \text{else.} \end{cases}$$

Two points are considered neighbors if either one is in the other's top 1% of similarities. This neighborhood definition is chosen empirically to create a connected but very sparse graph. The choice of c remains a problem; we set it arbitrarily to the 10'th percentile squared euclidean distance.

Outliers are known to be a problem for spectral clustering so we remove all points in the first percentile of similarity degree (i.e. the row sum of W). This is another arbitrary threshold, but in our experience this step dramatically improves the results. After the final clustering step each outlier can be assigned to the group with the closest cluster center, although it is often easier to visualize the data when the outliers are omitted for good.

We construct the diagonal degree matrix D and use it to normalize the adjacency matrix:

$$L = D^{-1}W.$$

Keeping only the first twenty (ordered) eigenvalues from a spectral decomposition of L, we choose the optimal number of clusters k to be the index of the largest weighted gap between consecutive values in this set. The weights are designed to bias k toward lower values. Figure 3 shows the eigenvalues for caudate points in a typical subject; it is not obvious which eigenvalue gap is the most significant. This is clearly a crude heuristic and is a major shortcoming of this method.

We assume the embedding dimension is equal to the optimal cluster number, and define V to be a matrix with columns containing the first k eigenvectors. The original points are now represented in eigen space by the rows of V, and these rows are clustered with the k-means++ algorithm. Figure 4 shows the putamen points for an example subject colored according to cluster. It is clear that spectral clustering is effective at identifying the natural endpoint groups.

Many of the parameters in this process are chosen empirically to produce consistently good clusterings, where the quality of a labeling is judged by a neuroimaging expert. Perhaps the most influential parameter choice is the fraction of points to label as outliers and remove from the graph Laplacian. Figure 5 shows a poor spectral clustering result for one subject, in that there are two groups of green points that clearly do not belong together (note this is a slightly different set of points than the baseline investigation).

Figures 6 and 7 illustrate a likely cause for this bad clustering result. The outlier points in orange push all of the other points to be very close to each other in the first three dimensions of the eigenspace, which causes dimensions 4-6 to be more heavily weighted in the clustering process. In these higher dimensions the green groups are actually adjacent to each other. Removing points in the first percentile of similarity degree seems to solve this problem.

3.3 fMRI pre-processing and simple connectivity models

For our preliminary investigation of functional connectivity we decouple this phase of the research from the process of clustering fiber endpoints. Instead we use a fixed set of 105 ROIs identified in each subject by the brain template matching algorithm in FreesurferFischl et al. (2002); Dale et al. (1999). Each ROI is associated with an rsBOLD time series with 210 observations, averaged over all of the voxels contained in that ROI. These time series are centered and scaled so the units are fractional deviations from the mean, and they are smoothed with a bandpass filter.

Each pair of ROIs has several associated measures of anatomical connectivity; Figure 8a illustrates the simplest of these in heatmap form. Each cell shows the log of the count of fibers between the ROIs at the cell's row and column indices. ROI pairs with no connecting fibers are blank (white space). We have not yet incorporated this anatomical connectivity network into our functional connectivity estimates, but we present it here for the sake of comparison.

As discussed in the literature review, the most common estimators for functional connectivity are pairwise full and partial correlation. The heatmap of these estimates are shown in Figures 8b and 8c; the more reddish a cell is, the more the row and column ROIs are associated. One interesting pattern that emerges are the two diagonal lines of strong association (ignoring the main diagonal, of course). These associations are a result of the fact that symmetric pairs of nodes—where nearly identical ROIs occur in each hemisphere of the brain—appear roughly 40 spots apart in the ROI list. This arrangement is used as a diagnostic check to ensure estimation methods perform correctly at a baseline level.

Results from a preliminary application of parallel lasso are shown in Figure 8d. Each row of this heatmap represents the index of the ROI used as the response node, and the red cells in a row indicate that ROI's estimated neighborhood. Once again the diagonal pattern of hemispheric symmetry appears very strongly.

4 Research Plan

Comparing fiber clustering methods. The primary goal of this dissertation is to explore and analyze high-dimensional regression methods for learning functional connectivity networks in the presence of prior knowledge about the underlying physical structure of those networks. The first step in this process is to identify which clustering algorithms work best for grouping fiber tract terminal points into network nodes. We know naïve k-means clustering does not work well because of the irregular spatial distribution of the endpoints in the target brain regions, especially the striatum, and our preliminary results show that spectral clustering can be effective but requires a substantial amount of tuning.

We plan to compare a number of clustering algorithms, including expectation maximization, hierarchical methods, density-based techniques, and variants of spectral clustering such as diffusion maps. This phase of the research is expected to take about two months.

Multi-scale clustering. Ultimately we want to simultaneously estimate the functional connectivity at multiple scales, and to define the nodes of the connectivity graph at each scale according to a hierarchical clustering of fiber tract terminal points. Spectral clustering works well to identify network nodes for the 5,000-10,000 points constructed in small brain regions like the caudate and thalamus, but more global networks consist of hundreds of thousands of fiber endpoints. Computing and storing all of the pairwise distances between these points is not computationally feasible so we plan to use recently-developed active clustering algorithms.

Active clustering assigns clusters hierarchically in a coarse-to-fine direction. At each level of the hierarchy a small sample of seed points is drawn and clustered with a traditional method. Each of the remaining points is then compared *only to the seed clusters* and is assigned to the most similar group. The $\mathcal{O}(n \log^2 n)$ pairwise distance calculations required for this process are a substantial improvement over the n^2 calculations required for typical clustering methods, and for active spectral clustering there are additional computational savings because the spectral decomposition is done only on the small sample of seed points. Assuming the *tight clustering condition*, which requires the similarities between points within a true cluster to exceed the similarities between points in two different true clusters, Eriksson et al. (2011) and Krishnamurthy et al. (2012) show the hierarchical version of this algorithm has good statistical properties. The authors show the method still recovers the true clusters with high probability even if a small subset of pairwise similarities is noisy or anomalous. Implementing this method and exploring its performance under the more severe violations of the tight clustering condition that likely occur in our data is expected to take two to three months.

Learning functional connectivity. The second major goal of this project is to design a regression-based strategy for estimating functional network structure in the presence of anatomical constraints. When modeling the functional connectivity at only one resolution the myriad sparse regression methods can be applied in a straightforward way. For example, the weights in the adaptive lasso penalty function can be used to represent the weakness of the anatomical connection between a target node i and all of the other nodes. Let the weakness measure be γ_{ij}. Then the neighborhood estimate for node i is the set of nonzero entries in

$$\widehat{\beta}^{i,\lambda} = \arg\min_{\beta} \left\{ \|X_i - \mathbf{X}_{(i)}\beta\|_2^2 + \lambda \sum_{j \neq i} \gamma_{ij} |\beta_j| \right\},$$

where X_i is the fMRI time series for node i and $\mathbf{X}_{(i)}$ is an $n \times (p-1)$ matrix with the other nodes' time series. The full edge set is estimated to be either the union or intersection of the node neighborhoods.

There are several questions that should be addressed for methods that model connectivity even at a single resolution. Do the consistency results proved by Meinshausen and Bühlmann (2006) continue to hold when different penalty functions are used in a parallel setting instead of lasso? Which of the three broad classes of regression approach—parallel lasso, stepwise regression, or regularized estimation of the inverse covariance matrix—performs best for our application, and how is the performance affected by incorporating anatomical information as with the adaptive lasso example above? If there are substantial correlations between the covariates, does a variant like the elastic net become optimal? To answer these performance questions we will use cross-validation and eventually a repeat scan of several of the subjects.

The more difficult aspect of learning functional connectivity is to estimate the network structure for an entire hierarchy of nodes using a structured sparsity-inducing regression penalty. To our knowledge none of the extant structured sparsity methods has good statistical and computational properties when covariate groups overlap completely, as they must in a tree structure. Investigation of structured sparse regression is expected to take four to six months.

Software. Statistical programming is a key component of this application-driven thesis. Many of the clustering and structured sparse algorithms that have been proposed do not have publicly available implementations, leading us to write much of our own software from scratch. As part of this project, our code will be disseminated publicly and maintained regularly.

A Figures

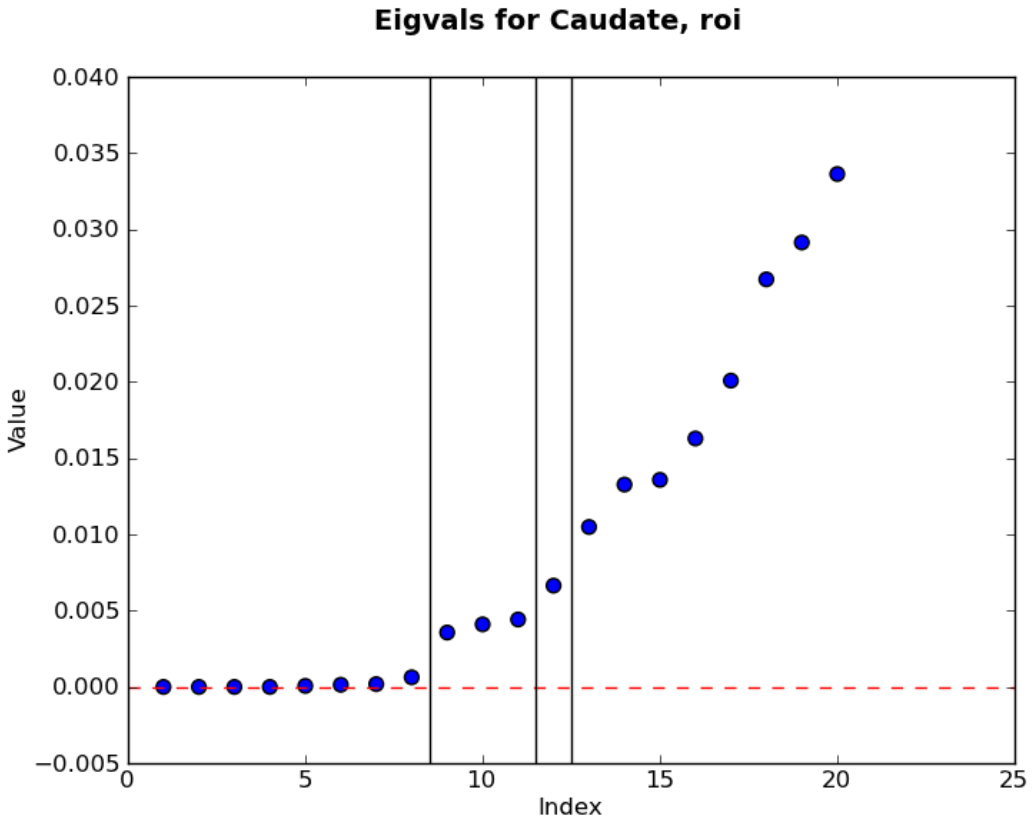

Figure 3: Eigenvalues of the normalized graph Laplacian of caudal fiber endpoints for a typical subject. Vertical black lines indicate the largest three *weighted* gaps.

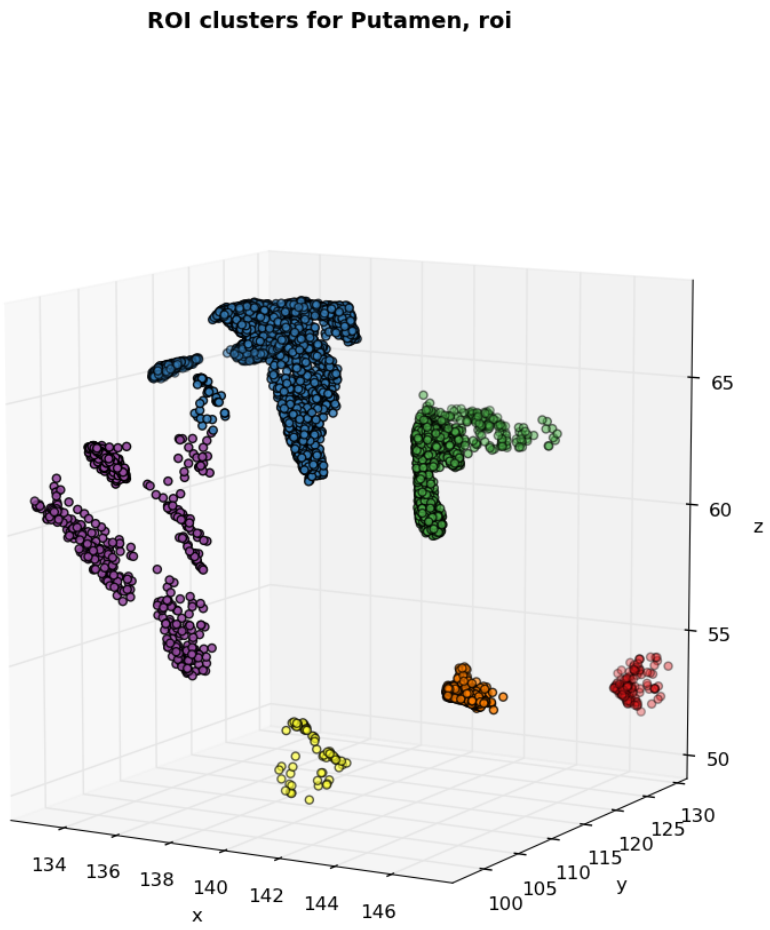

Figure 4: Putamen fiber endpoints (excluding outliers) for one subject. The six clusters are represented by color. Units are millimeters; X is the right-left axis and Y is the anterior-posterior axis.

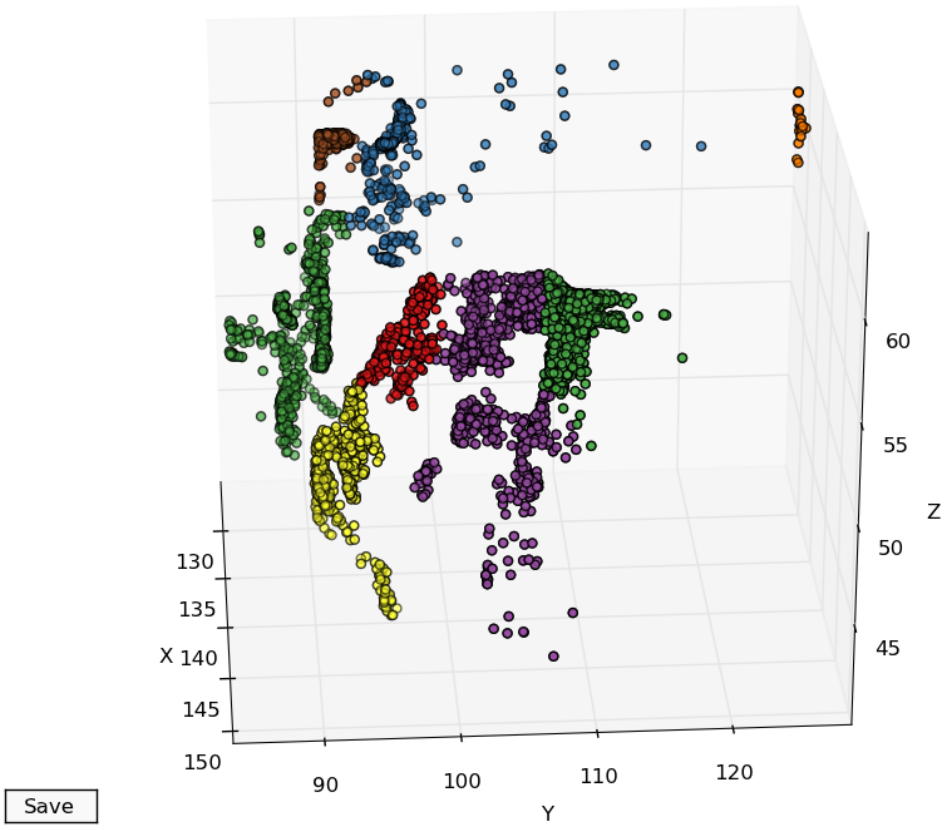

Figure 5: Dorsolateral prefrontal cortex (DLPFC) fiber endpoints for one subject, colored according to spectral clustering results. This clustering is of poor quality because it fails to group the main stripes. Note in particular the two groups of green endpoints that should not be considered one cluster.

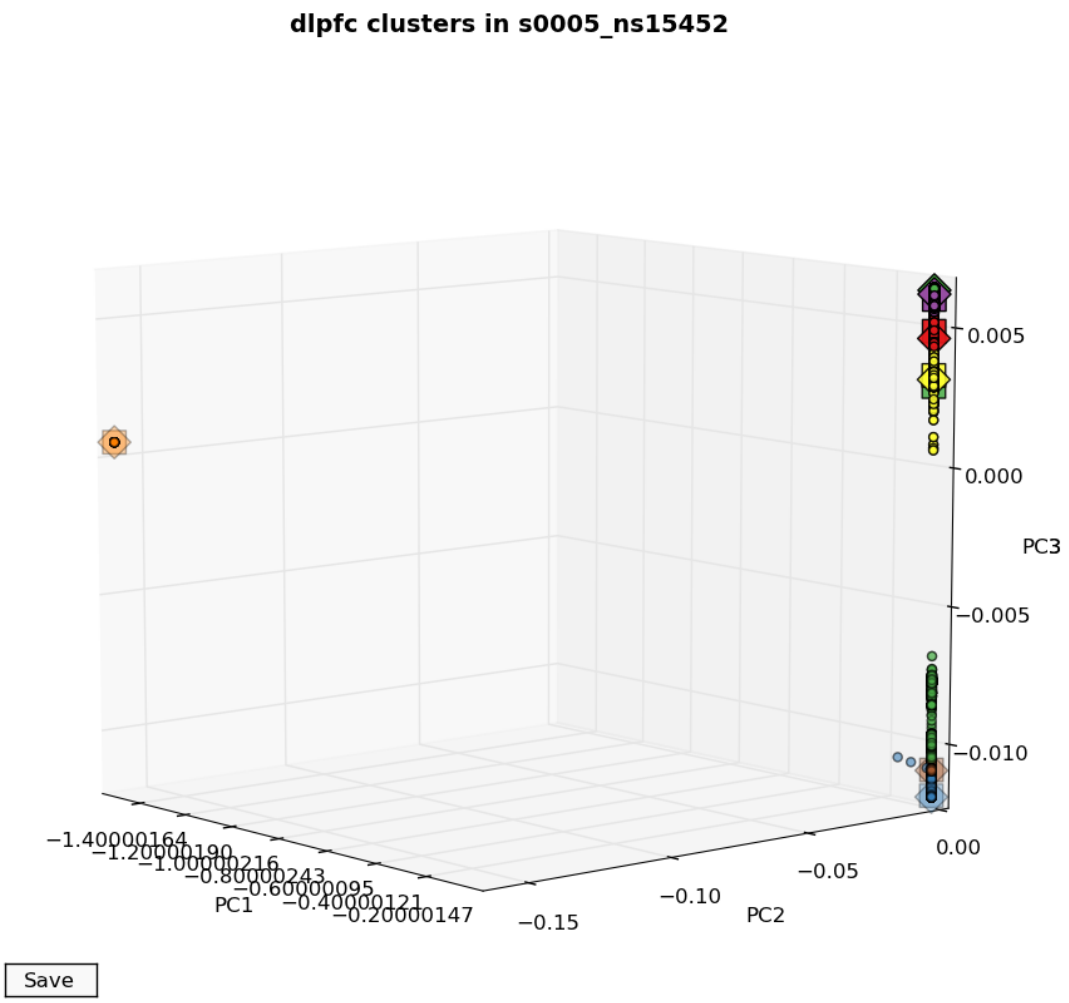

Figure 6: DLPFC fiber endpoints from Figure 5 represented in the first three dimensions of the eigenspace derived from a spectral decomposition. Note how the orange outlier points push all remaining points relatively close together.

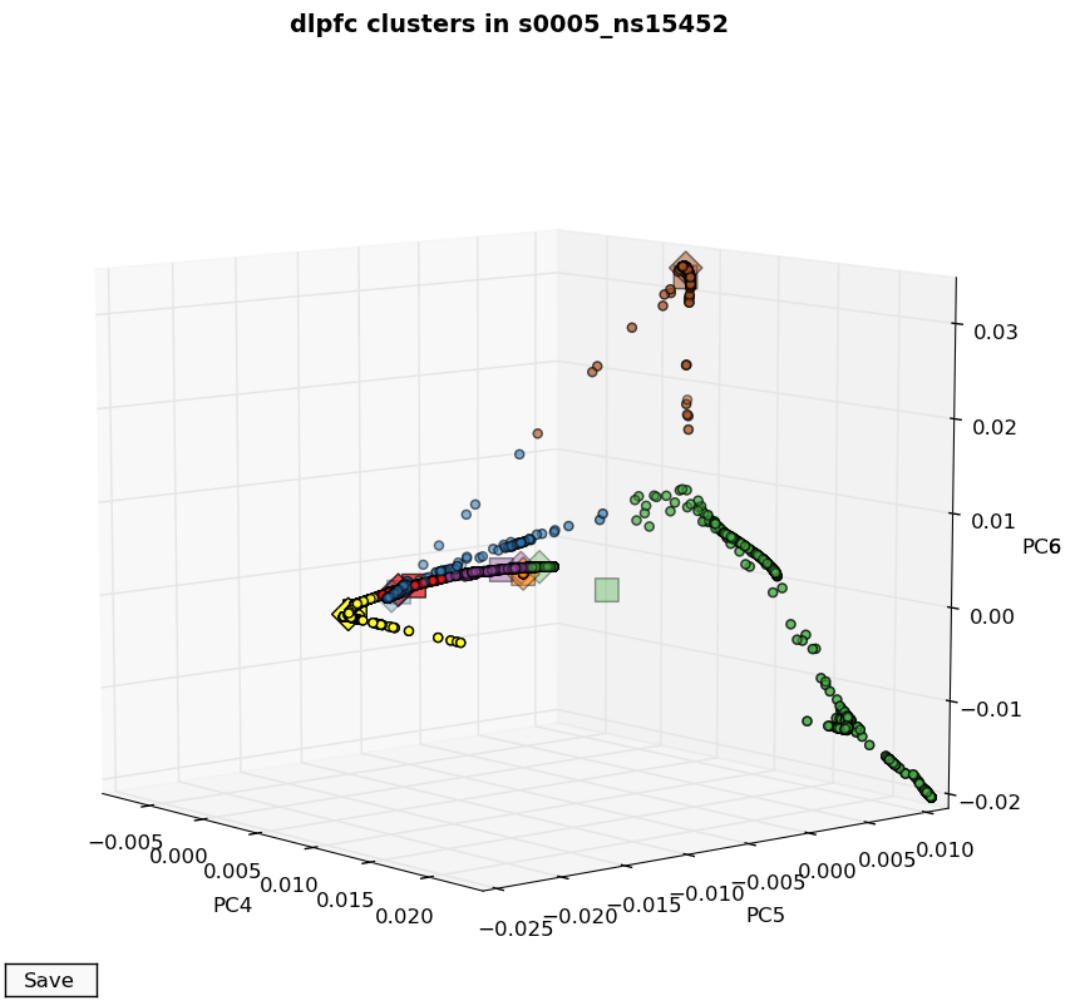

Figure 7: DLPFC fiber endpoints from Figure 5 represented in dimensions 4-6 of the spectral decomposition eigen space. Note the two groups of green points are close together in these dimensions.

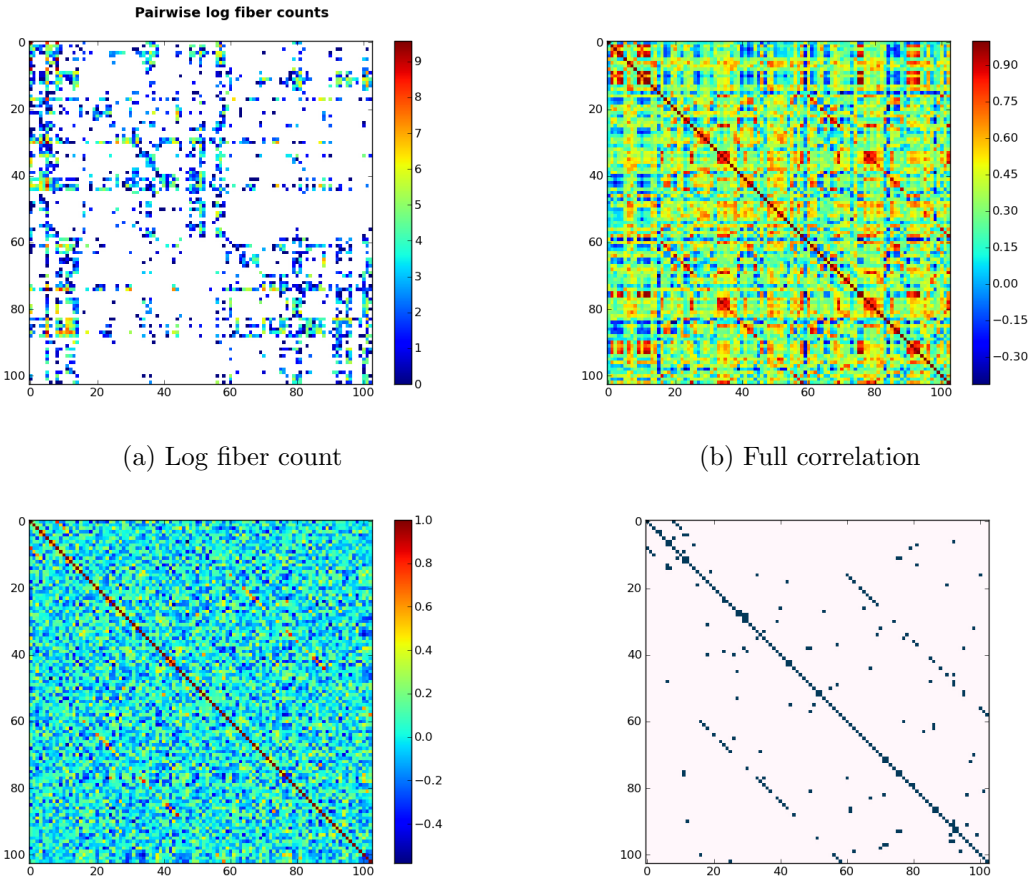

(a) Log fiber count

(b) Full correlation

(c) Partial correlation

(d) Parallel lasso neighborhoods

Figure 8: Various estimates of network connectivity for one subject. Panel (a) shows the log of the fiber count between pairs of ROIs, which is a measure of anatomical connectivity. The remaining three panels show estimates of functional connectivity. In panel (d), the red cells indicate the estimated neighborhood for the nodes listed by row. These estimates were made without using any structural information.

References

Andreas Argyriou, Rina Foygel, and Nathan Srebro. Sparse Prediction with the k-Overlap Norm. *arXiv*, 1204.5043v:1–9, 2012.

Ed Bullmore and Olaf Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature reviews. Neuroscience*, 10 (3):186–98, March 2009. ISSN 1471-0048. doi: 10.1038/nrn2575. URL `http://www.ncbi.nlm.nih.gov/pubmed/19190637`.

Florentina Bunea, Yiyuan She, Hernando Ombao, Assawin Gongvatana, Kate Devlin, and Ronald Cohen. Penalized least squares regression methods and applications to neuroimaging. *NeuroImage*, 55(4):1519–27, April 2011. ISSN 1095-9572. doi: 10.1016/j.neuroimage.2010.12.028. URL `http://www.ncbi.nlm.nih.gov/pubmed/21167288`.

Leila Cammoun, Xavier Gigandet, Djalel Meskaldji, Jean Philippe Thiran, Olaf Sporns, Kim Q. Do, Philippe Maeder, Reto Meuli, and Patric Hagmann. Mapping the human connectome at multiple scales with diffusion spectrum MRI. *Journal of Neuroscience Methods*, 203(2):386–397, October 2011. ISSN 01650270. doi: 10.1016/j.jneumeth.2011.09.031. URL `http://linkinghub.elsevier.com/retrieve/pii/S0165027011005991`.

Anders M Dale, Bruce Fischl, and Martin I Sereno. Cortical Surface-Based Analysis. *NeuroImage*, 9:179–194, 1999.

Brian Eriksson, Gautam Dasarathy, Aarti Singh, and Robert Nowak. Active Clustering : Robust and Efficient Hierarchical Clustering using Adaptively Selected Similarities. In Geoffrey Gordon, David Dunson, and Miroslav Dudik, editors, *JMLR Workshop and Conference Proceedings: AISTATS 2011*, volume 15, pages 260–268, Fort Lauderdale, FL, 2011.

Bruce Fischl, David H Salat, Evelina Busa, Marilyn Albert, Megan Dieterich, Christian Haselgrove, Andre van der Kouwe, Ron Killiany, David Kennedy, Shuna Klaveness, Albert Montillo, Nikos Makris, Bruce Rosen, and Anders M Dale. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron*, 33(3):341–55, January 2002. ISSN 0896-6273. URL `http://www.ncbi.nlm.nih.gov/pubmed/11832223`.

Michael D Greicius, Kaustubh Supekar, Vinod Menon, and Robert F Dougherty. Resting-state functional connectivity reflects structural connectivity in the default mode network. *Cerebral Cortex*, 19(1):72–80, January 2009. ISSN 1460-2199. doi: 10.1093/cercor/bhn059. URL `http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2605172&tool=pmcentrez&renderty`

P Guevara, C Poupon, D Rivière, Y Cointepas, M Descoteaux, B Thirion, and J Mangin. Robust clustering of massive tractography datasets. *NeuroImage*, 54(3): 1975–1993, 2011. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2010.10.028. URL `http://dx.doi.org/10.1016/j.neuroimage.2010.10.028`.

Patric Hagmann, Lisa Jonasson, Philippe Maeder, Jean-philippe Thiran, Van J Wedeen, and Reto Meuli. Understanding Diffusion MR Imaging Techniques: From Scalar Diffusion-weighted Imaging to Diffusion Tensor Imaging and Beyond. *RadioGraphics*, 26:205–224, 2006. doi: 10.1148/rg.26si065510.

C J Honey, O Sporns, L Cammoun, X Gigandet, J P Thiran, R Meuli, and P Hagmann. Predicting human resting-state functional connectivity from structural connectivity. *PNAS*, 106(6):2035–2040, 2009. doi: 10.1073/pnas.0811168106.

Laurent Jacob, Guillaume Obozinski, and Jean-Philippe Vert. Group Lasso with Overlap and Graph Lasso. *Proceedings of the 26th International Conference on Machine Learning*, pages 433–440, 2009. doi: 10.1145/1553374.1553431.

Kevin Jarbo, Timothy Verstynen, and Walter Schneider. In vivo quantification of global connectivity in the human corpus callosum. *NeuroImage*, 59(3):1988–1996, October 2012. ISSN 1095-9572. doi: 10.1016/j.neuroimage.2011.09.056. URL `http://www.ncbi.nlm.nih.gov/pubmed/21985906`.

Shuichi Kawano. Adaptive bridge regression modeling with model selection criteria. *arXiv*, 1204.3130v:1–14, 2012.

Akshay Krishnamurthy, Sivaraman Balakrishnan, Min Xu, and Aarti Singh. Efficient Active Algorithms for Hierarchical Clustering. *Proceedings of the International Conference on Machine Learning*, In Press, 2012.

Steffen Lauritzen. *Graphical Models*. Oxford University Press, Oxford, UK, 1996.

A Lemkaddem, A Daducci, S Vulliemoz, K O'Brien, F Lazeyras, M Hauf, R Wiest, R Meuli, M Seeck, G Krueger, and J-P Thiran. A multi-center study: Intra-scan and inter-scan variability of diffusion spectrum imaging. *NeuroImage*, May 2012. ISSN 1095-9572. doi: 10.1016/j.neuroimage.2012.04.045. URL `http://www.ncbi.nlm.nih.gov/pubmed/22569062`.

Ulrike Von Luxburg. A Tutorial on Spectral Clustering. Technical Report August, Max Planck Institute for Biological Cybernetics, Tubingen, Germany, 2006.

A.R. McIntosh. Towards a network theory of cognition. *Neural Networks*, 13(8-9):861–70, 2000. ISSN 0893-6080. URL `http://www.ncbi.nlm.nih.gov/pubmed/11156197`.

Nicolai Meinshausen and Peter Buhlmann. High-Dimensional Graphs and Vari-

able Selection with the Lasso. *Annals of Statistics*, 34(3):1436–1462, 2006. doi: 10.1214/009053606000000281.

Nicolai Meinshausen and Bin Yu. Lasso-type recovery of sparse representations for high-dimensional data. *The Annals of Statistics*, 37(1):246–270, February 2009. ISSN 0090-5364. doi: 10.1214/07-AOS582. URL http://projecteuclid.org/euclid.aos/1232115934.

Daniel Percival, Kathryn Roeder, Roni Rosenfeld, and Larry Wasserman. Structured, Sparse Regression With Application To Hiv Drug Resistance. *The annals of applied statistics*, 5(2A):628–644, June 2011. ISSN 1941-7330. doi: 10.1214/10-AOAS428. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3165199&tool=pmcentrez&renderty

Stephen M Smith, Karla L Miller, Gholamreza Salimi-khorshidi, Matthew Webster, Christian F Beckmann, Thomas E Nichols, Joseph D Ramsey, and Mark W Woolrich. Network modelling methods for FMRI. *NeuroImage*, 54(2):875–91, January 2011. ISSN 1095-9572. doi: 10.1016/j.neuroimage.2010.08.063. URL http://dx.doi.org/10.1016/j.neuroimage.2010.08.063 http://www.ncbi.nlm.nih.gov/pubmed/20817103.

Timothy Verstynen, Kevin Jarbo, Sudhir Pathak, and Walter Schneider. In vivo mapping of microstructural somatotopies in the human corticospinal pathways. *Journal of neurophysiology*, 105(1):336–46, January 2011. ISSN 1522-1598. doi: 10.1152/jn.00698.2010. URL http://www.ncbi.nlm.nih.gov/pubmed/21068263.

Timothy D Verstynen, David Badre, Kevin Jarbo, and Walter Schneider. Microstructural organizational patterns in the human corticostriatal system. *Journal of neurophysiology*, February 2012. ISSN 1522-1598. doi: 10.1152/jn.00995.2011. URL http://www.ncbi.nlm.nih.gov/pubmed/22378170.

Van J Wedeen, Patric Hagmann, Wen-yih Isaac Tseng, Timothy G Reese, and Robert M Weisskoff. Mapping Complex Tissue Architecture With Diffusion Spectrum Magnetic Resonance Imaging. *Magnetic Resonance in Medicine*, 54:1377–1386, 2005. doi: 10.1002/mrm.20642.

Fang-Cheng Yeh, Van Jay Wedeen, and Wen-Yih Isaac Tseng. Generalized q-sampling imaging. *IEEE transactions on medical imaging*, 29(9):1626–35, September 2010. ISSN 1558-254X. doi: 10.1109/TMI.2010.2045126. URL http://www.ncbi.nlm.nih.gov/pubmed/20940832.

Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B*, 68(1):49–67, 2006. doi: 10.1111/j.1467-9868.2005.00532.x.

Tong Zhang. Adaptive Forward-Backward Greedy Algorithm for Learning Sparse

Representations. *IEEE Transactions on Information Theory*, 57(7):4689–4708, July 2011a. ISSN 0018-9448. doi: 10.1109/TIT.2011.2146690. URL http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5895111.

Tong Zhang. Multi-stage Convex Relaxation for Feature Selection. *arXiv*, 1106.0565v(2), 2011b.

Peng Zhao and Bin Yu. On Model Selection Consistency of Lasso. *Journal of Machine Learning Research*, 7:2541–2563, 2006. URL http://dl.acm.org/citation.cfm?id=1248547.1248637.

Peng Zhao, Guilherme Rocha, and Bin Yu. The composite absolute penalties family for grouped and hierarchical variable selection. *The Annals of Statistics*, 37(6):3468–3497, 2009. doi: 10.1214/07-AOS584.

Dajiang Zhu, Kaiming Li, Carlos Cesar Faraco, Fan Deng, Degang Zhang, Lei Guo, L. Stephen Miller, and Tianming Liu. Optimization of Functional Brain ROIs via Maximization of Consistency of Structural Connectivity Profiles. *NeuroImage*, 59(2): 1382–1393, August 2011. ISSN 10538119. doi: 10.1016/j.neuroimage.2011.08.037. URL http://linkinghub.elsevier.com/retrieve/pii/S105381191100930X.

Hui Zou. The Adaptive Lasso and Its Oracle Properties. *Journal of the American Statistical Association*, 101(476):1418–1429, December 2006. ISSN 0162-1459. doi: 10.1198/016214506000000735. URL http://pubs.amstat.org/doi/abs/10.1198/016214506000000735.