11-1-2015

Partial covariance based functional connectivity computation using Ledoit-Wolf covariance regularization

Matthew R. Brier

Anish Mitra

John E. McCarthy
Washington University in St Louis, mccarthy@wustl.edu

Beau M. Ances

Abraham Z. Snyder

Partial covariance based functional connectivity
computation using Ledoit-Wolf covariance

Matthew R. Brier[1], Anish Mitra[2], John E. McCarthy[3],
Beau M. Ances[1,2,4], Abraham Z. Snyder[2,1]

[1]Department of Neurology, Washington University in St Louis.
[2]Department of Radiology, Washington University in St Louis.
[3]Department of Mathematics, Washington University in St Louis.
[4]Department of Biomedical Engineering, Washington University in St Louis.

Total Word Count: 4,528
Abstract Word Count: 222
Number of Figures: 6
Number of Tables: 2
Supplemental Figures: 1

Correspondence to:
Matthew Brier
Department of Neurology
Washington University in St Louis
660 S Euclid
St. Louis, MO 63110
brierm@wusm.wustl.edu

Abstract

Functional connectivity refers to shared signals among brain regions and is typically assessed in a task free state. Functional connectivity commonly is quantified between signal pairs using Pearson correlation. However, resting-state fMRI is a multivariate process exhibiting a complicated covariance structure. Partial covariance assesses the unique variance shared between two brain regions excluding any widely shared variance, hence is appropriate for the analysis of multivariate datasets as exemplified by fMRI. However, calculation of partial covariance requires inversion of the covariance matrix, which, in most functional connectivity studies, is not invertible owing to rank deficiency. Here we apply Ledoit-Wolf shrinkage (L_2 regularization) to invert the high dimensional BOLD covariance matrix. We investigate the network organization and brain-state dependence of partial covariance-based functional connectivity. Although RSNs are conventionally defined in terms of shared variance, removal of widely shared variance, surprisingly, improved the separation of RSNs in a spring embedded graphical model. This result suggests that pair-wise unique shared variance plays a heretofore unrecognized role in RSN covariance organization. In addition, application of partial correlation to fMRI data acquired in the eyes open vs. eyes closed states revealed focal changes in uniquely shared variance between the thalamus and visual cortices. This result suggests that partial correlation of resting state BOLD time series reflect functional processes in addition to structural connectivity.

2

Introduction

The brain is highly active even in the absence of observable behavior (Raichle, 2011). Intrinsic brain activity is not random, but rather exhibits a stereotypical correlation structure (Biswal et al., 2010). Spontaneous fluctuations in the blood oxygen level dependent (BOLD) signal (Ogawa et al., 1993) have been used to investigate the organization of intrinsic activity in the resting-state, i.e., in the absence of explicit task performance (Biswal et al., 1995). This phenomenon commonly is referred to as functional connectivity; the associated topographies define resting-state networks (RSNs). Seed based correlation mapping (Biswal et al., 2010) and spatial independent component analysis (sICA) (Beckmann et al., 2005) together account for the vast majority of functional connectivity studies. These techniques are able to define RSN topographies but do not isolate variance unique to brain region pairs. Partial covariance techniques offer the possibility of improved understanding of brain organization by more precisely attributing widely and uniquely shared variance in different brain states. However, partial covariance analysis of fMRI datasets has been difficult owing to the rank deficiency, hence non-invertibility, of high dimensional fMRI datasets (Schafer and Strimmer, 2005). Extant strategies for dealing with this problem in the context of resting-state fMRI are listed in Table 1.

Covariance Matrix Conditioning via Shrinkage

BOLD time-series give rise to an empirical covariance matrix of the form: $\hat{\Sigma} = \frac{1}{T}(X - \bar{X})^t(X - \bar{X})$, where $X \in \mathbb{R}^{T \times M}$ (T time points and M regions of interest [ROI])

with the overbar indicating the column mean. When M is large, $\hat{\Sigma}$ typically is ill-conditioned, i.e., not invertible. Ledoit-Wolf shrinkage describes a process wherein rank deficient covariance matrices can be made invertible by shrinkage (Ledoit and Wolf, 2003). Defining $\tilde{\Sigma} = (1 - \alpha)\hat{\Sigma} + \alpha\Delta$, where α is a tuning parameter and Δ is the shrinkage target (Ledoit and Wolf, 2004), yields an invertible matrix. Shrinkage targets vary in form but generally have some favorable property. In this case, the shrinkage target is full rank. The value of α that results in $\tilde{\Sigma}$ that most closely approximates the theoretical value of Σ (the true covariance matrix given infinite data) has in the past been determined using cross-validation e.g., (Efron and Morris, 1975). However, it has been shown that the optimal value of α can be calculated in closed form under weak assumptions (Ledoit and Wolf, 2003)[1]. Matrix shrinkage is a low bias strategy in the sense that $\tilde{\Sigma}$ is close to the true theoretical covariance matrix. Given $\tilde{\Sigma}$, partial covariance and related quantities in high dimensional covariance matrices can be calculated.

Ledoit-Wolf regularization has previously been used to condition resting-state BOLD covariance matrices (Varoquaux et al., 2012, Deligianni et al., 2014) but the properties and consequences of this approach to functional connectivity have not been thoroughly investigated. In order to introduce this approach to the broader neuroimaging community, we first demonstrate the favorable properties of the partial covariance quantity. Specifically, partial covariance matrices calculated using

---

[1] The weak assumptions are that the first four moments of the true covariance matrix are defined.

Ledoit-Wolf regularization can be calculated with large numbers of ROIs and have improved test-retest reliability compared to full covariance matrices. We use the partial covariance quantity to investigate RSN organization. RSNs are defined by their widely shared variance; the partial covariance quantity removes this variance. It is unclear what, if any, RSN organization remains in the partial covariance matrix. Partial covariance functional connectivity has been reported to be topographicallysimilar to structural connectivity (i.e., constant over short time-scales). We next investigate the brain-state dependence of the partial correlation quantity using a well-established contrast: eyes open vs. eyes closed.

Methods

<u>Subject Characteristics</u>

Two datasets are used in this study. Dataset A consisted of 57 adults (mean age: 30 years, range: 18-45 years). All subjects were neurologically and psychiatrically normal on examination and neuropsychological testing. These subjects were enrolled as controls in ongoing imaging studies and two short resting state fMRI scans with the eyes open and fixated were collected. Dataset B consisted of 10 adults (mean age: 25 years, range: 22 – 31 years). Two long resting state fMRI scans were collected, one with eyes open and the other with eyes closed. All subjects provided written informed consent in accordance with the Washington University in St Louis Institutional Review Board.

<u>Scanning Parameters</u>

Commented [JM1]: They are calculated by inversion; you don't want to invert the matrix of them.

Commented [JM2]: What does reported to be mean?

5

Image acquisition was performed using a 3T Siemens Trio scanner (Erlangen, Germany) equipped with a standard 12-channel head coil. A high-resolution structural scan was acquired using a 3-dimensional sagittal T1-weighted magnetization-prepared rapid gradient echo (MPRAGE, echo time [TE] = 16 msec, repetition time [TR] = 2,400 msec, inversion time [TI] = 1,000 msec, flip angle = 8°, 256×256 acquisition matrix, $1 \times 1 \times 1$mm voxels). This scan was used for atlas registration. High-resolution 2D multi-slice oblique axial spin density/T2-weighted fast spin echo (FSE) structural images were also acquired using slice tilts and positions computed by slice preregistration (TE=455 msec, TR = 3,200 msec, 256×256 acquisition matrix, 1 acquisition, $1 \times 1 \times 1$mm voxels). These T2-weighted FSE data were used for fMRI atlas registration. Resting state fMRI scans were collected using a gradient spin-echo sequence (TE = 30 msec, TR = 2200 msec, field of view = 256 mm, flip angle = 90°, 4mm isotropic voxels) sensitive to blood oxygen level dependent (BOLD) contrast (Ogawa et al., 1993). In dataset A, two six-minute resting state fMRI runs (164 volumes per run) were acquired during which participants were asked to fixate on a visual cross-hair and not fall asleep. In dataset B, two forty-five minute resting state fMRI runs (1225 volumes per run) were acquired. During one run the participants were asked to fixate on a visual cross-hair (Eyes open; EO) and during the other run the participants were asked to keep their eyes closed (EC) but not to fall asleep. The order of these scans were counterbalanced across subjects.

Preprocessing

6

Initial preprocessing of resting state fMRI data followed conventional methods as previously described (Brier et al., 2014). Generic preprocessing of fMRI data included correction for slice-dependent intensity differences related to interleaved acquisition (debanding) (Hacker et al., 2013), rigid body correction for head movement within and between fMRI runs, and atlas transformation. Volumes strongly contaminated by head movement (Power et al., 2012) were removed and voxelwise replaced with linearly interpolated values (Power et al., 2013). Frame censoring was computed using the DVARS measure (Smyser et al., 2010) modified to include a 10mm FWHM Gaussian spatial pre-blur[2]. The frame exclusion threshold was set at 0.7% rms BOLD signal change over successive frames, counting only voxels within the brain (Brier et al., 2014). Excluded frames were replaced with linear interpolations only for the purposes of facilitating preprocessing. Only subjects with fewer than 40% of frames excluded were passed on to the next stage of processing. Signals of non-interest were extracted from white matter, ventricles, and the global signal averaged over the whole brain (Fox et al., 2009). These signals of non-interest along with movement time-series and their first temporal derivatives were regressed from the voxelwise BOLD time-series. In order to investigate the effects of global signal regression (GSR) on the calculation of partial covariance matrices, a subset of the following analyses were performed without GSR. The residual BOLD time-series was then low-pass filtered to retain frequencies below 0.1Hz and spatially smoothed with Gaussian blur (6mm FWHM in each direction). The linearly interpolated volumes were excluded in all subsequent

_____

[2] This pre-blur step was used only to calculate DVARS and was not carried forward.

analyses. Following preprocessing, an average of 308 (stdev = 30) frames were retained for each subject. These BOLD time-series were autocorrelated across approximately 3 frames (6.6 seconds); thus each subject had approximately 103 degrees of freedom.

Region of Interest Definition

In Dataset A, we used two sets of regions of interest (ROI) to examine the effects of matrix conditioning under different degrees of rank deficiency. The first set contained 36 ROIs representing 5 RSNs including the default mode, dorsal attention, executive control, salience, and sensorimotor network (Brier et al., 2012). The second set contained 264 ROIs representing 12 RSNs including somatomotor, cingulo-opercular, auditory, default mode, parietal encoding/retrieval, visual, fronto-parietal control, salience, subcortical, ventral attention, dorsal attention, and cerebellum networks (Power et al., 2011). Each ROI was a 6mm radius sphere centered on previously published coordinates.

Covariance Calculations and Ledoit-Wolf Shrinkage

For each subject, the BOLD time-series extracted from $M = 36$ or $M = 264$ ROIs had length T, where T is the number of BOLD frames that were not contaminated with movement. Thus, define $X \in \mathbb{R}^{T \times M}$. The sample covariance matrix is then defined as $\hat{\Sigma} = \frac{1}{T}(X - \bar{X})^t(X - \bar{X})$. Because large subspaces in the covariance matrix are highly related and because the condition $T \gg M$ is not satisfied, $\hat{\Sigma}$ is rank deficient and therefore uninvertible. We define the shrunken covariance matrix as $\tilde{\Sigma} = (1 - \alpha)\hat{\Sigma} +$

$\alpha\Delta$ (Ledoit and Wolf, 2004, Schafer and Strimmer, 2005). α was calculated in closed-form as described in (Ledoit and Wolf, 2003) and Δ is the identity matrix (Ledoit and Wolf, 2004). The details of the calculation of α are described in Appendix 1.

Full and Partial Covariance

For each subject, the sample and conditioned covariance matrix have the form:

$$\hat{\Sigma} = \begin{bmatrix} \hat{\sigma}_1 & \cdots & \hat{\sigma}_{i,j} \\ \vdots & \ddots & \vdots \\ \hat{\sigma}_{j,i} & \cdots & \hat{\sigma}_i \end{bmatrix} \qquad \tilde{\Sigma} = \begin{bmatrix} \tilde{\sigma}_1 & \cdots & \tilde{\sigma}_{i,j} \\ \vdots & \ddots & \vdots \\ \tilde{\sigma}_{j,i} & \cdots & \tilde{\sigma}_i \end{bmatrix}$$

where $i, j \in [1, \dots, 36]$ or $i, j \in [1, \dots, 264]$ and σ_i denotes the variance of the i-th time-series and $\sigma_{i,j}$ denotes the full covariance between the i-th and j-th time-series. The off-diagonal values in $\hat{\Sigma}^{-1}$ or $\tilde{\Sigma}^{-1}$ quantify the partial covariance between the i-th and j-th time-series accounting for all other time-series. To conform with the definition of partial correlations (Weatherburn, 1961) and resolve a sign flip associated with inverting the covariance matrix we always scaled $\hat{\Sigma}^{-1}$ and $\tilde{\Sigma}^{-1}$ by a factor of -1. This sign manipulation ensures that the reported partial covariances have the same sign as the corresponding partial correlations. For the purpose of demonstrating the favorable properties of Ledoit-Wolf regularization as applied to fMRI data we use covariance as it is the general case (Dataset A). However, the application of matrix shrinkage is equally valid for correlation matrices (see Supplemental Material) and is used in the investigation of Dataset B in order to conform to the extant literature.

Graphical LASSO Comparison

In Dataset A, we compared the results of partial covariance values calculated using Ledoit-Wolf shrinkage with an extant L$_1$-based method, namely, the graphical LASSO (Friedman et al., 2008). The inverse covariance matrix was estimated for each subject using a range of *a priori* defined penalty value. It is also possible to determine a penalty value by cross-validation. We calculated the ideal penalty coefficient using a popular cross-validation scheme (Bien and Tibshirani, 2011). To compare these results with the previous results, we calculated the correlation across all ROI pairs and subjects.

Identification of Topography Sensitive to Eyes Open vs. Eyes Closed

In order to investigate the changes in widely and uniquely shared variance in Dataset B we first identified a set of ROIs that were most affected by the eyes open vs. eyes closed contrast. We divided the entire grey matter into 9mm isotropic ROIs and calculated the correlation between each ROI and all voxels for the EO and EC condition separately. The resulting parametric volumes were averaged across subjects to yield $C^{EO}(M \times V)$ and $C^{EC}(M \times V)$ where M is the number of 9mm isotropic ROIs and V is the number of voxels. Given M ROIs, there exist $\binom{M}{2}$ functional connectivity pairs; these data require data reduction. To achieve concise summary topographies we adopted a PCA approach. Let $X = C^{EO} - C^{EC}$. PCA is performed on X such that $X(M, V \in GM) = USW^T$ where GM is the grey matter mask. The number of significant principal components is determined by an information criteria (Minka, 2000). Let w_i be the ith column of W with length M. The topography of the effect of Eyes Open vs. Eyes Closed is defined as:

10

$$T_i = \frac{1}{\sqrt{M}} \sum_{m=0}^{M} w_i(m) \cdot X(m)$$

T_i has length V and represents a topography. Large values (positive and negative) were extracted to form ROIs ($|T_i| > 0.05$; cluter size > 100 voxels).

Intersubject Variability as a Funciton of Brain State

To quantify the level of inter-subject variability we calculated the distance between a single subject's full/partial correlation matrix in the EO or EC condition and the group average without that single subject. For some subject *n*, either full or partial, EO or EC, the inter-subject variability for generic matrix C is defined as:

$$\sigma_n = \frac{1}{N} \sum_{n=1}^{N} \left\| C_n - \left(\frac{1}{N-1} \sum_{\eta \neq n} C_\eta \right) \right\|_2$$

Thus, four distributions (full/partial, EO/EC) are then subjected to a matrix type (full, partial) by condition (EO, EC) repeated measures ANOVA.

Results

Matrix shrinkage accomplished significant matrix conditioning

We calculated $\hat{\Sigma}$ in the 36 and 264 ROI sets with and without GSR (Figure 1A). The covariance matrices exhibited block organization commonly observed in resting-state BOLD fMRI correlation matrices. With GSR, positive blocks along the diagonal correspond to within RSN correlations and off-diagonal blocks of negative correlations indicate anti-correlations (e.g., between the default-mode network (DMN) and dorsal attention network (DAN)). Without GSR, the same block

11

organization was evident, but all values were shifted towards positive values. $\tilde{\Sigma}$, computed as described in section 2.5, is shown in Figure 1B. The difference between $\hat{\Sigma}$ and $\tilde{\Sigma}$ is not apparent to the eye (Figure 1A vs 1B). However, the effect of matrix conditioning is clearly evident in the eigenspectra (Figure 1C). The critical feature evident in Figure 1C is the exponential decrease in eigenvalues of $\hat{\Sigma}$ (blue line), especially in the larger ROI set. In contrast, the eigenvalues of $\tilde{\Sigma}$ showed a less steep decline in the smaller ROI set, and a plateau in the larger ROI set (α is the lower bound).

Matrix invertibility is frequently quantified as the condition number, defined as the ratio of the largest to smallest eigenvalue. In practice, matrices with condition numbers less than $\sim 10^4$ are invertible using standard double precision arithmetic. In the larger ROI set the condition number was intolerably large but was made significantly smaller with shrinkage (Figure 1D). For the 36 ROI set matrix shrinkage improved the condition number by about a factor of 2. However, for the 264 ROI set matrix conditioning improved the condition number from a value representing the limits of machine precision ($\sim 10^{17}$) to an invertible condition ($\sim 10^3$). This improvement in condition number was achieved with α less than 0.05 for all conditions (Figure 1E). By way of comparison, the value of α obtained in a highly cited study of gene expression data using Ledoit-Wolf regularization with 100 variables was 0.20 (Schafer and Strimmer, 2005). In the present data, significant conditioning was achieved with a relatively small amount of shrinkage. Similar results were achieved for correlation matrices (Supplemental Figure 1).

12

Matrix conditioning stabilized matrix inversion and improved test-retest reliability

The results shown in Figure 1 demonstrate a dramatic improvement in condition number following modest conditioning. The error in matrix inversion is bounded by the condition number (Conte and de Boor, 1980). To investigate this important effect, we report empirical results using the present data and a split sample approach. More specifically, the full and partial covariance matrices derived from the first and second fMRI runs were compared with and without conditioning. Assuming stationarity, the results obtained from both runs should be the same. Non-stationarity of BOLD time-series has been reported, e.g., (Chang and Glover, 2010, Allen et al., 2012), but recent results suggest that this effect is small when appropriate models are used (Lindquist et al., 2014). Accepting this view, the difference between the first and second runs is interpretable in terms of test-retest reliability. To quantify the difference between the matrices resulting from the first and second fMRI runs, we calculated the Euclidean distance (Frobenius norm) between the corresponding matrices. We subjected these log transformed Euclidean distance measures to a linear model investigating the effect of matrix regularization correcting for the α value. Matrix conditioning had little effect on the full covariance matrices, but reduced the distance between the inverse covariance matrices by more than 10 orders of magnitude in the case of the larger ROI set (Figure 2A) and by a factor of 10 in the smaller ROI set. We also observed a negative relationship between the distance between the inverse covariance matrices prior to regularization and α. This effect was approximately the same with and without GSR.

13

Thus, matrix conditioning, as implemented here, was sufficient to stabilize the inversion step even for a large number of ROIs. This result follows from the relationship between condition number and inversion precision and has practical utility in improving test-retest reliability of high-dimensional partial correlations.

We next investigated the structure of the first vs second run differences, which we define as error for the present purposes. Unstructured error may be reduced by increasing numbers of subjects, increasing the amount of data acquired in a single subject, or by averaging across ROIs. However, structured error is more difficult to remove. Thus, reduction of structured error is an ideal characteristic of an analytic approach. The mean (across subjects) error for each of the matrix types was investigated. As noted above, the magnitude of the error differs by orders of magnitude depending on conditioning and ROI number. We investigated the presence or absence of structured error. Accordingly, the error was Z-transformed (removing the mean and normalizing by the standard deviation across ROI pairs), thereby scaling all results to approximately the same range. The full covariance matrices are presented above the diagonal in Figure 2B,C. The first vs. second run differences were clustered roughly by RSN; critically these differences were not affected by conditioning. The first vs. second run differences in partial covariance matrices (below the diagonal, Figure 2B,C) exhibited no structure in the 36 ROI set. In contrast, in the 264 ROI set, the differences without conditioning were structured (red arrows); conditioning removed this structure. To quantify the effect of conditioning on error structure in inverse covariance matrices, we calculated the

14

first eigenvalue of each difference matrix (Table 2). Large and small first eigenvalues correspond, respectively, to structured and unstructured error. A large first eigenvalue indicates similar error structure across some set of ROIs; averaging or other data reduction techniques will not remove this error. On the other hand, a small first eigenvalue indicates random error across ROIs that is removable via averaging or other approaches. Matrix conditioning significantly reduced this metric of structured error in the inverse covariance matrices, particularly in the large ROI set (all $p<10^{-6}$). We emphasize that without removing the overall effect of magnitude (Z-scoring), the effect of conditioning would have been much more marked.

The relationship between partial covariance analysis and global signal regression is shown in the Supplemental Material (Supplemental Figure 2).

Ledoit-Wolf shrinkage and LASSO (L$_1$) based methods yielded similar results
We have demonstrated the utility of Ledoit-Wolf shrinkage in the computation of partial covariance matrices. A frequently used alternative, namely, the graphical LASSO approach, enforces an L$_1$ penalty on the inverse solution (Friedman et al., 2008). We compared the calculated partial covariance values resulting from the graphical LASSO approach vs. the present Ledoit-Wolf based approach. As there exists no closed form method to select the weight of the L$_1$ penalty (β), we calculated partial covariance matrices over a range of β values [0.10 0.25 0.50 0.75 0.90]. The 264 ROI case failed to converge after 1 week of computation. Accordingly, we report

15

a comparison between the graphical LASSO approach and the Ledoit-Wolf approach

for the 36 ROI case (Figure 3A). To simulate the thresholding effect of the L_1 penalty

in the Ledoit-Wolf based result, we thresholded the absolute value of the data over a

range of ρ values [0.00 0.01 ... 0.25]. For all values of β and ρ, L_1:L_2 correlation was

high (r > 0.50). However, there was obvious structure in the L_1:L_2 correlation

results. The highest correlation was found with no L_2 thresholding (ρ=0) and

minimal L_1 sparsity (β=0.10) (Figure 3B). Similarly high correlation was found

between the penalty determined by cross-validation (β=0.12) (Figure 3C). Sparsity

is a requirement for inversion in L_1 based methods, but increasing sparsity reduces

test-retest reliability (Varoquaux et al., 2012). The present approach avoids sparsity

altogether and is computationally feasible for larger ROI sets.

Partial covariance exhibits strong homotopy and reveals novel features of RSN

organization

Having demonstrated the practical utility of Ledoit-Wolf regularization for the

computation of partial covariance functional connectivity, we now investigate the

brain organization that is revealed by this approach. We analyzed the 36 ROI set as

the larger ROI set does not explicitly include homotopic ROIs. Full covariance

matrices (Figure 4A, above the diagonal) demonstrate the familiar block

organization reflecting RSN organization. This feature is notably absent in the

partial covariance matrix (Figure 4A, below the diagonal). Instead, the partial

covariance matrix is dominated by homotopic functional connectivity, although this

feature is not immediately obvious on inspection of Figure 4A. To more clearly

demonstrate this feature, we display the values of Figure 4A in histogram format

16

(Figure 4B). In both the full and partial covariance results, homotopic values (shown in blue) were concentrated in the right tail of the distribution. However, the separation between homotopic versus other covariance values was more complete using partial covariance (Figure 4B, enlargement).

The topography of homotopic functional connectivity differed in the full versus partial covariance results. Repeated measures ANOVA revealed that both full ($F_{14,786}$=3.44, p=0.000018) and partial ($F_{14,786}$=23.7, p<10^{-8}) homotopic covariance values varied across brain regions (Figure 4C). However, in detail, the topographies of homotopic full and partial covariance were not correlated (r = 0.088, p = 0.76). This result suggests that partial covariance reveals features in the functional organization of the brain that are not captured in the full covariance values.

Full covariance functional connectivity is dominated by clusters of widely shared variance (RSNs); partial covariance calculation removes that variance. We next investigated if RSN organization still existed after removing the variance that defines RSNs using a semi-quantitative graphical approach. Full and partial covariance matrices were thresholded using a one-sample t-test (across subjects) against the null-hypothesis of no covariance. t-statistics corresponding to suprathreshold (p<0.001, uncorrected) ROI pairs were used as edge weights in a graph. The resultant graph was displayed using a force directed algorithm (Hu, 2005). In the force directed results obtained with the 36 ROI set, nodes belonging to the same RSN were clustered for both mean $\tilde{\Sigma}$ (Figure 4D) and mean $\tilde{\Sigma}^{-1}$ (Figure

17

4E). RSNs are conventionally defined in terms of widely shared variance between multiple ROIs; partial covariance removes that shared variance. Therefore one might expect RSN structure to fall apart the force directed graphical model derived from partial covariance analysis. However, contrary to this expectation, the separation between RSNs was more complete in the result obtained with mean $\tilde{\Sigma}^{-1}$. Comparable results were obtained with the 264 ROI set (Supplemental Figure 3). The critical observation here is that clustering of nodes by RSN affiliation persisted in the partial covariance results even though shared signals were removed. In full covariance matrices, RSN affiliation is reflected in block organization. This feature is not apparent in partial covariance matrices (Figure 4A, below the diagonal). Nevertheless, nodes in the same RSN remained clustered. Inspection of the force directed graph result obtained with the 36 ROI set suggests that RSNs are clustered in the partial covariance results on the basis of conditionally dependent pair-wise connections. Thus, RSNs can be defined in two ways: 1) on the basis of widely shared variance across a set of ROIs and 2) on the basis of uniquely shared variance between pairs of member ROIs.

Partial Correlation is Focally Dependent on Changes in Brain State

Partial correlation functional connectivity has been proposed to more closely approximate structural connectivity (e.g., white matter tractography). If partial correlation functional connectivity were a simple reflection of structural connectivity (Smith et al., 2013) it would be insensitive to brain state. To test this hypothesis, we investigated the partial correlation organization in two commonly

investigated brain states: eyes open (EO) and eyes closed (EC). We began by identifying a topography of brain regions sensitive to the EO vs. EC full correlation contrast (see Methods). Two significant principal components were identified (Figure 5A). From these topographies, eight ROIs were extracted. Full and partial correlation functional connectivity was calculated (Figure 5B). Full correlation functional connectivity was dramatically affected by the change in brain state (Figure 5C, above diagonal) whereas partial correlation functional connectivity was relatively unchanged (Figure 5C, below diagonal). However, partial correlations between the thalamus and visual areas and the thalamus and frontal areas were affected. Thus, partial correlation organization is modifiable by brain state and those perturbations are relatively focal compared to full correlation changes.

Next we investigated the inter-subject variability in the full and partial correlation matrices in the EO and EC condition. The inter-subject variability was calculated as the distance (Euclidean norm) between a single subject's matrix and the group mean (less that subject)'s matrix (Table 3). These values were subjected to a correlation type (full, partial) by condition (EO, EC) repeated measures ANOVA. There was a significant effect of correlation type ($F_{(1,9)}=11.11$, $p=0.0088$) wherein the partial correlation matrices showed less inter-subject variability. There was no effect of condition ($F_{(1,9)}=0.0046$, $p=0.95$) nor an interaction ($F_{(1,9)}=0.27$, $p=0.62$).

Discussion

Summary

19

This report utilizes an approach to covariance matrix conditioning that facilitates calculation of high dimensional partial covariance matrices for the assessment of functional connectivity. This approach yields partial covariance matrices that are stable within subjects and reproducible across subject groups. Application of this computational strategy provided insight into the brain's functional organization. For example, RSNs are defined by both their widely shared variance and by pair-wise unique shared variance. Finally, partial correlations are sensitive to brain state suggesting they contain dynamic representations of brain connectivity and not simple 1:1 reflections of structural connectivity.

<u>Matrix shrinkage circumvents limitations in extant approaches to partial covariance calculation</u>

The favorable properties of partial correlation and covariance computation in the context of resting state fMRI have been frequently recognized (Table 1). The primary challenge has been overcoming rank deficiency of high-dimensional BOLD fMRI covariance matrices. The simplest approach is to limit the dimensionality of the system (Marrelec et al., 2006, Fransson and Marrelec, 2008, Zhang et al., 2008). This approach has the advantage of avoiding ill conditioned matrices but limits the scope of the scientific inquiry. The present matrix shrinkage strategy provides computational stability and enables study of large partial covariance organizations.

The second extant approach to overcoming the rank deficiency problem is dimensionality reduction using PCA followed by ICA (Liang et al., 2011, Yu et al.,

2011, Smith et al., 2013). The covariance matrix calculated over independent components is full rank and therefore invertible. However, this approach includes an arbitrary parameter (the number of components retained in the PCA step) and does not allow for investigation of partial correlations between individual regions, although, partial correlations between RSNs can be computed. Ledoit-Wolf shrinkage involves no arbitrary parameters and supports the computation of partial correlations between *a priori* selected ROIs.

The third approach in the extant literature for calculating partial correlations is imposing an L_1 penalty to constrain the inversion step (Varoquaux et al., 2010, Fiecas et al., 2013, Smith et al., 2013). The L_1 norm forces small values in the inverse matrix to 0. This procedure effectively reduces the number of free parameters in the inversion computation. This approach includes a tunable parameter determined by simulation but allows for the investigation of large number of ROIs. One of the consequences of the L_1 constraint is that the inverse covariance matrix is sparse (i.e., has many off-diagonal 0s). This feature has been described as an advantage (Varoquaux et al., 2010) but a recent report has argued that the L_1 constraint is too strict (Ryali et al., 2012). Ryali and colleagues used the elastic net (linear combination of the L_1 and L_2 penalty (Hastie et al., 2009)) to constrain the inversion step at the cost of an additional free parameter. The L_2 norm penalizes large values in the inverse but does not force small values to zero. The present approach is analogous to imposing a pure L_2 penalty without any tunable parameters.

21

Thus, the present approach for calculating partial correlation/covariance values has three primary advantages over alternative approaches. First, the calculation of Ledoit-Wolf shrinkage is possible for a large number of ROIs and computation time does not dramatically increase as in non-closed form solutions (e.g., graphical LASSO). Second, this approach does not assume a model (e.g., sparsity). Sparsity may be appropriate for some forms of connectivity (e.g., white matter tractography) but has not been conclusively demonstrated for functional measures. Finally, the arbitrary parameter is calculated in closed form, precludes simulations, reduces computational burden, and stabilizes interpretation.

Previous applications of matrix shrinkage to functional connectivity

Covariance matrix shrinkage has been recently used in the investigation of functional connectivity (Varoquaux et al., 2012, Deligianni et al., 2014, Shou et al., 2014). Stein's paradox (Stein, 1956) asserts that a shrunken estimate of the mean outperforms the sample mean in predicting the true mean of a multivariate distribution. Shou and colleagues used knowledge of the group functional connectivity matrix to predict the retest outcome in subjects scanned twice. Shrinkage of individual functional connectivity matrices towards the group mean improved test-retest reliability. This result is consistent with the known behavior of multivariate distributions under shrinkage (Efron and Morris, 1975) but represents an application of matrix shrinkage very different from the present use. We used shrinkage to improve matrix invertibility. Nevertheless, we observed that partial covariance estimates were more reliable than full covariance estimates both at the

22

individual and group level. Other work has used Ledoit-Wolf shrinkage to calculate partial correlation values (Varoquaux et al., 2012, Deligianni et al., 2014), but these studies did not systematically investigate the properties of the Ledoit-Wolf approach.

Homotopy dominates partial covariance organization.

Large partial covariance values identify variance uniquely shared between two variables. In the present results (Figure 4), the largest partial covariance values were observed between homotopic brain regions. Strong homotopic functional connectivity is well documented (Salvador et al., 2008) and is thought to be supported by callosal and other commissural fibers (O'Reilly et al., 2013, Shen et al., 2015). It has been suggested that that unique shared variance implies direct communication between brain regions (Smith et al., 2013). This is an appealing notion but it is not entirely supported by the present results. Over all ROI pairs, the largest partial covariance value was observed between left and right primary visual cortex (Figure 4C). The only known direct inter-hemispheric connections between primary visual cortices connect the representations of the vertical meridian (Van Essen et al., 1982). It seems unlikely that this anatomy supports the observed high degree of unique shared variance. Hence, some other mechanism may be responsible for identification of highly reliable inter-hemispheric partial correlations in primary visual cortex. One hypothesis that potentially explains the present observations is that high degrees of shared variance may arise from common efferents or afferents (Adachi et al., 2012).

23

RSN organization is preserved in partial covariance matrices.

Direct comparison of full and partial covariance matrices reveals several striking differences. Partial covariance matrices lacked the familiar RSN-related block organization that dominates full covariance matrices. Loss of prominent RSN organization follows from the definition of both RSNs and partial covariance. RSNs are defined by shared variance between member regions; this shared variance was removed in partial covariance matrices. What remains is unique shared variance between regions. Surprisingly, this organization is sufficient to separate RSNs in force-directed graphs (Figure 4D). This result may reflect some underlying skeletonized organization of the brain (van den Heuvel et al., 2012). We note that partial covariance matrices were reliable, which suggests that they reflect some biologically important aspect of brain organization.

Partial Correlation Functional Connectivity is Sensitive to Brain State

The similar organization of partial correlation functional connectivity and structural connectivity has been previously observed (Smith et al., 2013). The extent to which partial correlations represent structural connectivity alone or dynamic brain organization is unclear. The contrast of brain organization in the eyes open vs. eyes closed state is an operationally simple contrast that is known to elicit functional connectivity changes. Closing of the eyes elicits a change in mean BOLD signal in primary visual (Bianciardi et al., 2009) as well as non-visual cortices (Marx et al., 2003, McAvoy et al., 2008). Functional connectivity is also affected. Consistent with

24

previous literature we found broad changes in full correlation functional connectivity (Figure 5) (Jao et al., 2013, Liang et al., 2014, Xu et al., 2014). In particular, we observed an increasingly strong negative correlation between visual cortices and the thalamus (Zou et al., 2009, McAvoy et al., 2012). In the partial correlations, functional connectivity between the thalamus and visual regions is also made more negative by eye closure. This demonstrates that the change in functional connectivity between the thalamus and visual cortex due to eye state is mediated by uniquely shared variance while the other, more broad, modulations are due to widely shared variance. Speculatively, we hypothesize that the changes in broadly shared variance may be due to arousal (Tagliazucchi and Laufs, 2014) while the changes in uniquely shared variance may be more attributable to the eye condition-related brain-state change. Regardless of the exact mechanism, this result demonstrates that partial correlation based functional connectivity is modulated by brain state. This modulation indicates that uniquely shared variance reveals functional brain organization, not simply structural connectivity.

Appendix 1

In this appendix we briefly outline the calculation of the Ledoit-Wolf shrinkage parameter α. This approach was originally described in (Ledoit and Wolf, 2003) and code is generously made available by those authors[3]. We outline the process briefly. Capital letters denote matrices, lower case letters denote elements in a matrix, and t indicates transpose.

Given data in the form

$$X \in \mathbb{R}^{T \times M}$$

the empirical covariance matrix is calculated as

$$\hat{\Sigma} = \frac{1}{T}(X - \bar{X})^t (X - \bar{X})$$

where the overbar indicates the column mean. The shrinkage target is computed as

$$\Delta = \bar{\sigma} \cdot I$$

where $\bar{\sigma} = \langle diag(\hat{\Sigma}) \rangle_M$ and I is the identity matrix of $\mathbb{R}^{M \times M}$. The shrinkage parameter is calculated as follows. Define:

$$p_{ij} = \frac{1}{T} \sum_{t=1}^{T} \left\{ (x_{it} - \bar{x}_i)(x_{jt} - \bar{x}_j) - \hat{\Sigma}_{ij} \right\}^2$$

$$c = \left\| \hat{\Sigma} - \Delta \right\|_2^2$$

From this, the shrinkage constant can be calculated:

$$\alpha = \frac{(\Sigma_{ij} \, p/c)}{T}$$

---

[3] http://www.ledoit.net/cov1para.m

Table/Figure Captions

Table 1: Literature Review. Table summarizing previous reports utilizing partial correlation/ covariance approaches. Table shows the reference, solution adopted to allow partial correlation estimation, number of ROIs and number of subject. The general area of scientific inquiry is also noted.

Table 2: 1st eigenvalue of Z transformed difference of 1st and 2nd half scan session. Table shows the mean (standard deviation) of the largest eigenvalue of the Z transformed difference between the first and second fMRI runs. This is computed separately for the 36 and 264 ROI set, with and without GSR, and for the covariance and inverse covariance matrices. The far right column shows the t and p values for the comparison of the sample (unconditioned) and conditioned condition.

Table 3: Inter-subject Variability. Mean (standard deviation) values of σ as calculated in the methods. Low values correspond to low levels of inter-subject variability.

Figure 1: Matrix conditioning significantly improves the condition number of BOLD covariance matrices. A: BOLD covariance matrices (Σ) calculated in the 36 and 264 ROI set with and without global signal regression (GSR). Matrices with GSR exhibit familiar block organization representing resting-state networks (RSNs). Matrices without GSR exhibit similar organization biased towards positive values. **B:** Conditioned BOLD covariance matrices ($\hat{\Sigma}$) in the same configuration as A. Notably,

27

there is no visual difference between matrices in A and B. **C:** Mean eigenspectra of matrices in A and B. The blue line indicates the eigenspectra corresponding to $\mathbf{\Sigma}$ and the red lines indicates the eigenspectra corresponding to $\widehat{\mathbf{\Sigma}}$. The eigenvalues of $\mathbf{\Sigma}$ decay towards zero but the eigenvalues of $\widehat{\mathbf{\Sigma}}$ plateau at a non-zero constant. **D:** Quantification of matrix invertibility by mean condition number. The condition number is defined as the ratio of the largest eigenvalue to the smallest. **E:** mean shrinkage coefficient (α) in each ROI set with and without global signal regression. Error bars indicate standard error of the mean estimated across subjects.

Figure 2: Matrix conditioning is sufficient to constrain the covariance inverse.
A: Mean Euclidean distance between covariance and inverse covariance matrices calculated from the first and second fMRI sessions. Blue bars display results without conditioning; red bars display results with conditioning. Error bars represent standard deviation across subjects. Asterisks indicate significant differences; in this case $p < 0.001$. Conditioning had no significant effect on distance for full covariance results, but significantly reduced the distance (error) between the inverse covariance matrices. **B:** Z-transformed (removal of mean and normalization by standard deviation) mean difference between the first and second sessions without conditioning. Full covariance results are above the diagonal; partial covariance results are below the diagonal. Red arrow heads indicate organized pattern of large values in the inverse matrix. **C:** Same as B, but with conditioning. Structured noise is absent in the conditioned result.

28

Figure 3: Ledoit-Wolf shrinkage and L$_1$ based methods yield similar results. A: The correlation (using ROI pairs and subjects as samples) between the partial covariance quantities resulting from the present Ledoit-Wolf shrinkage based approach and the graphical LASSO approach are presented for varying L1 penalty weights (β) and varying absolute thresholds applied to the Ledoit-Wolf based approach (ρ). Only representative β values are shown as LASSO is computationally expensive. **B:** Scatter plot illustrating the maximal correlation result ($\rho=0$ and $\beta=0.10$) between the two techniques. **C:** Same as B, but $\beta=0.12$ which was determined by cross-validation.

Figure 4: The partial covariance matrix is dominated by homotopic functional connectivity and contains RSN organization. A: Full (above diagonal) and partial (below diagonal) covariance matrices shown for the 36 ROI set with GSR. Black boxes denote RSN membership. **B:** Histogram of full and partial covariance values with homotopic covariance values in blue. Inset magnifies right side of distribution. Homotopic covariance is large in both distributions, but is better separated in the partial covariance result. **C:** Mean full and partial covariance between homotopic brain regions (error bars are standard error. **D:** The 36 ROI covariance and partial covariance matrix were transformed into a graph representations. The position of each node optimized using an automated algorithm (Hu, 2005). Colors denote *a priori* RSN membership. Lines indicating connections are omitted from the 264 ROI result for visual clarity.

Figure 5: Partial Correlation Organization is Modulated by Brain-state A: First two principal components of the EO vs. EC full correlation contrast. Strongly affected regions (positive and negative) were extracted to form ROIs. ROIs are superimposed as colored regions. **B:** Full (above diagonal) and partial (below diagonal) correlation matrices using extracted ROIs in the EO and EC condition. **C:** ROI-pair-wise contrasts of the EO and EC condition. Cell color indicated t-statistic value. White X indicated p < 0.001, uncorrected for multiple comparisons.

Supplemental Figure 1: Matrix conditioning significantly improves the condition number of BOLD correlation matrices. A: BOLD correlation matrices ($\widehat{C}$) calculated in the 36 and 264 ROI set with and without global signal regression (GSR). Matrices with GSR exhibit familiar block organization representing resting-state networks (RSNs). Matrices without GSR exhibit similar organization biased towards positive values. **B:** Conditioned BOLD correlation matrices ($\widetilde{C}$) in the same configuration as A. Notably, there is no visual difference between matrices in A and B. **C:** Mean eigenspectra of matrices in A and B. The blue line indicates the eigenspectra corresponding to $\widehat{C}$ and the red lines indicates the eigenspectra corresponding to $\widetilde{C}$. The eigenvalues of $\widehat{C}$ decay towards zero but the eigenvalues of $\widetilde{C}$ plateau at a non-zero constant. **D:** Quantification of matrix invertibility by mean condition number. The condition number is defined as the ratio of the largest eigenvalue to the smallest. **E:** mean shrinkage coefficient (α) in each ROI set with and without global signal regression.

Supplemental Figure 2: Partial covariance matrices do not depend on global signal regression. A: Mean Z transformed covariance (above the diagonal) and partial covariance (below the diagonal) matrices. **B:** Scatter plots show the relationship between the result derived with and without GSR. Red line indicates the line of identity. In full covariance matrices, GSR represents roughly an additive constant, however the partial covariance results are nearly invariant to GSR. Inset in each plot is the slope (m) and intercept (b) or the best fit linear regression.

Supplemental Figure 3: RSN organization in 264 ROI set. Figure in same style as Figure 4D but examined in the 264 ROI set.

References

Adachi Y, Osada T, Sporns O, Watanabe T, Matsui T, Miyamoto K, Miyashita Y (2012) Functional connectivity between anatomically unconnected areas is shaped by collective network-level effects in the macaque cortex. Cerebral cortex 22:1586-1592.

Allen EA, Damaraju E, Plis SM, Erhardt EB, Eichele T, Calhoun VD (2012) Tracking Whole-Brain Connectivity Dynamics in the Resting State. Cerebral cortex.

Beckmann CF, DeLuca M, Devlin JT, Smith SM (2005) Investigations into resting-state connectivity using independent component analysis. Philosophical transactions of the Royal Society of London Series B, Biological sciences 360:1001-1013.

Bianciardi M, Fukunaga M, van Gelderen P, Horovitz SG, de Zwart JA, Duyn JH (2009) Modulation of spontaneous fMRI activity in human visual cortex by behavioral state. NeuroImage 45:160-168.

Bien J, Tibshirani RJ (2011) Sparse estimation of a covariance matrix. Biometrika 98:807-820.

Biswal B, Yetkin Z, Haughton VM, Hyde JS (1995) Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. Magnetic Resonance in Medicine 34:537-541.

Biswal BB, Mennes M, Zuo XN, Gohel S, Kelly C, Smith SM, Beckmann CF, Adelstein JS, Buckner RL, Colcombe S, Dogonowski AM, Ernst M, Fair D, Hampson M, Hoptman MJ, Hyde JS, Kiviniemi VJ, Kotter R, Li SJ, Lin CP, Lowe MJ, Mackay C, Madden DJ, Madsen KH, Margulies DS, Mayberg HS, McMahon K, Monk CS, Mostofsky SH, Nagel BJ, Pekar JJ, Peltier SJ, Petersen SE, Riedl V, Rombouts SA, Rypma B, Schlaggar BL, Schmidt S, Seidler RD, Siegle GJ, Sorg C, Teng GJ, Veijola J, Villringer A, Walter M, Wang L, Weng XC, Whitfield-Gabrieli S, Williamson P, Windischberger C, Zang YF, Zhang HY, Castellanos FX, Milham MP (2010) Toward discovery science of human brain function. Proceedings of the National Academy of Sciences of the United States of America 107:4734-4739.

Brier MR, Thomas JB, Snyder AZ, Benzinger TL, Zhang D, Raichle ME, Holtzman DM, Morris JC, Ances BM (2012) Loss of intranetwork and internetwork resting state functional connections with Alzheimer's disease progression. The Journal of neuroscience : the official journal of the Society for Neuroscience 32:8890-8899.

Brier MR, Thomas JB, Snyder AZ, Wang L, Fagan AM, Benzinger T, Morris JC, Ances BM (2014) Unrecognized preclinical Alzheimer disease confounds rs-fcMRI studies of normal aging. Neurology.

Chang C, Glover GH (2010) Time-frequency dynamics of resting-state brain connectivity measured with fMRI. NeuroImage 50:81-98.

Conte S, de Boor C (1980) Elementary Numerical Analysis: McGraw-Hill.

Deligianni F, Centeno M, Carmichael DW, Clayden JD (2014) Relating resting-state fMRI and EEG whole-brain connectomes across frequency bands. Frontiers in neuroscience 8:258.

Efron B, Morris C (1975) Data analysis using Stein's estimator and its generalizations. Journal of the American Statistical Association 70:311-319.

Fiecas M, Ombao H, van Lunen D, Baumgartner R, Coimbra A, Feng D (2013) Quantifying temporal correlations: a test-retest evaluation of functional connectivity in resting-state fMRI. NeuroImage 65:231-241.

Fox MD, Zhang D, Snyder AZ, Raichle ME (2009) The global signal and observed anticorrelated resting state brain networks. Journal of neurophysiology 101:3270-3283.

Fransson P, Marrelec G (2008) The precuneus/posterior cingulate cortex plays a pivotal role in the default mode network: Evidence from a partial correlation network analysis. NeuroImage 42:1178-1184.

Friedman J, Hastie T, Tibshirani R (2008) Sparse inverse covariance estimation with the graphical lasso. Biostatistics 9:432-441.

Hacker CD, Laumann TO, Szrama NP, Baldassarre A, Snyder AZ, Leuthardt EC, Corbetta M (2013) Resting state network estimation in individual subjects. NeuroImage 82:616-633.

Hastie T, Tibshirani R, Friedman JH (2009) The elements of statistical learning : data mining, inference, and prediction. New York, NY: Springer.

Hu Y (2005) Efficient, high-quality force-directed graph drawing. The Mathematica Journal 10:37-71.

Jao T, Vertes PE, Alexander-Bloch AF, Tang IN, Yu YC, Chen JH, Bullmore ET (2013) Volitional eyes opening perturbs brain dynamics and functional connectivity regardless of light input. NeuroImage 69:21-34.

Ledoit O, Wolf M (2003) Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. Journal of Empirical Finance 10:603-621.

Ledoit O, Wolf M (2004) A well-conditioned estimator for large-dimensional covariance matrices. Journal of Multivariate Analysis 88:365-411.

Liang B, Zhang D, Wen X, Xu P, Peng X, Huang X, Liu M, Huang R (2014) Brain spontaneous fluctuations in sensorimotor regions were directly related to eyes open and eyes closed: evidences from a machine learning approach. Frontiers in human neuroscience 8:645.

Liang Z, King J, Zhang N (2011) Uncovering intrinsic connectional architecture of functional networks in awake rat brain. The Journal of neuroscience : the official journal of the Society for Neuroscience 31:3776-3783.

Lindquist MA, Xu Y, Nebel MB, Caffo BS (2014) Evaluating dynamic bivariate correlations in resting-state fMRI: A comparison study and a new approach. NeuroImage.

Marrelec G, Krainik A, Duffau H, Pelegrini-Issac M, Lehericy S, Doyon J, Benali H (2006) Partial correlation for functional brain interactivity investigation in functional MRI. NeuroImage 32:228-237.

Marx E, Stephan T, Nolte A, Deutschlander A, Seelos KC, Dieterich M, Brandt T (2003) Eye closure in darkness animates sensory systems. NeuroImage 19:924-934.

McAvoy M, Larson-Prior L, Ludwikow M, Zhang D, Snyder AZ, Gusnard DL, Raichle ME, d'Avossa G (2012) Dissociated mean and functional connectivity BOLD signals in visual cortex during eyes closed and fixation. Journal of neurophysiology 108:2363-2372.

McAvoy M, Larson-Prior L, Nolan TS, Vaishnavi SN, Raichle ME, d'Avossa G (2008) Resting states affect spontaneous BOLD oscillations in sensory and paralimbic cortex. Journal of neurophysiology 100:922-931.

Minka T (2000) Automatic choice of dimensionality for PCA. NIPS 13:598-604.

O'Reilly JX, Croxson PL, Jbabdi S, Sallet J, Noonan MP, Mars RB, Browning PG, Wilson CR, Mitchell AS, Miller KL, Rushworth MF, Baxter MG (2013) Causal effect of disconnection lesions on interhemispheric functional connectivity in rhesus monkeys. Proceedings of the National Academy of Sciences of the United States of America 110:13982-13987.

Ogawa S, Menon RS, Tank DW, Kim SG, Merkle H, Ellermann JM, Ugurbil K (1993) Functional brain mapping by blood oxygenation level-dependent contrast magnetic resonance imaging. A comparison of signal characteristics with a biophysical model. Biophys J 64:803-812.

Power JD, Barnes KA, Snyder AZ, Schlaggar BL, Petersen SE (2012) Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. NeuroImage 59:2142-2154.

Power JD, Cohen AL, Nelson SM, Wig GS, Barnes KA, Church JA, Vogel AC, Laumann TO, Miezin FM, Schlaggar BL, Petersen SE (2011) Functional network organization of the human brain. Neuron 72:665-678.

Power JD, Mitra A, Laumann TO, Snyder AZ, Schlaggar BL, Petersen SE (2013) Methods to detect, characterize, and remove motion artifact in resting state fMRI. NeuroImage 84C:320-341.

Raichle ME (2011) The restless brain. Brain connectivity 1:3-12.

Ryali S, Chen T, Supekar K, Menon V (2012) Estimation of functional connectivity in fMRI data using stability selection-based sparse partial correlation with elastic net penalty. NeuroImage 59:3852-3861.

Salvador R, Martinez A, Pomarol-Clotet E, Gomar J, Vila F, Sarro S, Capdevila A, Bullmore E (2008) A simple view of the brain through a frequency-specific functional connectivity measure. NeuroImage 39:279-289.

Schafer J, Strimmer K (2005) A shrinkage approach to large-scale covariance matrix estimation and implications for functional genomics. Statistical applications in genetics and molecular biology 4:Article32.

Shen K, Misic B, Cipollini BN, Bezgin G, Buschkuehl M, Hutchison RM, Jaeggi SM, Kross E, Peltier SJ, Everling S, Jonides J, McIntosh AR, Berman MG (2015) Stable long-range interhemispheric coordination is supported by direct anatomical projections. Proceedings of the National Academy of Sciences of the United States of America.

Shou H, Eloyan A, Nebel MB, Mejia A, Pekar JJ, Mostofsky S, Caffo B, Lindquist MA, Crainiceanu CM (2014) Shrinkage prediction of seed-voxel brain connectivity using resting state fMRI. NeuroImage.

Smith SM, Beckmann CF, Andersson J, Auerbach EJ, Bijsterbosch J, Douaud G, Duff E, Feinberg DA, Griffanti L, Harms MP, Kelly M, Laumann T, Miller KL, Moeller S, Petersen S, Power J, Salimi-Khorshidi G, Snyder AZ, Vu AT, Woolrich MW, Xu J, Yacoub E, Ugurbil K, Van Essen DC, Glasser MF, Consortium WU-MH (2013) Resting-state fMRI in the Human Connectome Project. NeuroImage 80:144-168.

34

Smyser CD, Inder TE, Shimony JS, Hill JE, Degnan AJ, Snyder AZ, Neil JJ (2010) Longitudinal analysis of neural network development in preterm infants. Cerebral cortex 20:2852-2862.

Stein C (1956) Inadmissibility of the usual estimator for the mean of a multivariate distribution. Proceedings of the Third Berkely Symposium on Mathematics, Statistics, and Probability 1.

Tagliazucchi E, Laufs H (2014) Decoding wakefulness levels from typical fMRI resting-state data reveals reliable drifts between wakefulness and sleep. Neuron 82:695-708.

van den Heuvel MP, Kahn RS, Goni J, Sporns O (2012) High-cost, high-capacity backbone for global brain communication. Proceedings of the National Academy of Sciences of the United States of America 109:11372-11377.

Van Essen DC, Newsome WT, Bixby JL (1982) The pattern of interhemispheric connections and its relationship to extrastriate visual areas in the macaque monkey. The Journal of neuroscience : the official journal of the Society for Neuroscience 2:265-283.

Varoquaux G, Gramfort A, Poline JB, Thirion B (2010) Brain covariance selection: better individual functional connectivity models using population prior. Advances in Neural Information Processing Systems 23.

Varoquaux G, Gramfort A, Poline JB, Thirion B (2012) Markov models for fMRI correlation structure: Is brain functional connectivity small world, or decomposable into networks? J Physiol Paris 106:212-221.

Weatherburn CE (1961) A first course in mathematical statistics. Cambridge Eng.: University Press.

Xu P, Huang R, Wang J, Van Dam NT, Xie T, Dong Z, Chen C, Gu R, Zang YF, He Y, Fan J, Luo YJ (2014) Different topological organization of human brain functional networks with eyes open versus eyes closed. NeuroImage 90:246-255.

Yu Q, Sui J, Rachakonda S, He H, Gruner W, Pearlson G, Kiehl KA, Calhoun VD (2011) Altered topological properties of functional network connectivity in schizophrenia during resting state: a small-world brain network study. PloS one 6:e25423.

Zhang D, Snyder AZ, Fox MD, Sansbury MW, Shimony JS, Raichle ME (2008) Intrinsic functional relations between human cerebral cortex and thalamus. Journal of neurophysiology 100:1740-1748.

Zou Q, Long X, Zuo X, Yan C, Zhu C, Yang Y, Liu D, He Y, Zang Y (2009) Functional connectivity between the thalamus and visual cortex under eyes closed and eyes open conditions: a resting-state fMRI study. Human brain mapping 30:3066-3078.

Table 1

Reference	Conditioning Strategy	# of ROIs/ components	# of Subjects	Scientific Topic
Liu et al., 2008	No conditioning	90	31 controls, 31 schizophrenics	Small world structure in schizophrenia
Balenzuela et al., 2010	No conditioning	90	12 controls, 12 back pain patients	Graph analysis of functional connectivity in lower back pain
Zhang et al., 2011	No conditioning	90	63 controls, 30 MDD	Small world structure in major depressive disorder
Marrelec et al., 2006	Limit # of ROIs	6	7	Motor system organization
Fransson and Marrelec, 2008	Limit # of ROIs	9	17	DMN organization
Zhang et al., 2008	Limit # of ROIs	5	17	Thalamic organization
Liang et al., 2011	ICA dimensionality reduction	38	6 rats	Network organization in the awake rate

Yu et al., 2011	ICA dimensionality reduction	57	19 controls, 19 schizophrenics	Small world structure in schizophrenia
Smith et al., 2013*	ICA dimensionality reduction	22/78	20	Methods development for the connectome project
Varoquaux et al., 2010	L_1 Normalization	122	20	Covariance model selection
Fiecas et al., 2013	L_1 Normalization	90	25	Test-retest reliability assessment
Smith et al., 2013*	L_1 Normalization	22/78	20	Methods development for the connectome project
Ryali et al., 2012	Elastic Net Normalization	90	22	Stability of different regularization techniques
Varoquaux et al., 2012	Markov Structure (L_1 and Ledoit-Wolf)	105	12	RSN vs small-world brain organization
Deligianni et al., 2014	Ledoit-Wolf shrinkage	82	17	Relationship between BOLD and EEG functional connectivity

* Smith et al., 2013 compared multiple approaches so is presented twice

Table 2

			Sample	Conditioned	t,p
\multicolumn 1st eigenvalue of Z transformed difference of 1st and 2nd half scan session					
36 ROI Set	w/ GSR	Covariance	17.8 (6.4)	17.8 (6.6)	0.7, 0.49
		Inverse	16.7 (3.4)	13.0 (2.0)	7.3, <10^{-8}
	w/o GSR	Covariance	19.0 (4.7)	19.0 (4.8)	0.5, 0.64
		Inverse	17.2 (3.5)	13.0 (2.1)	8.5, <10^{-10}
264 ROI Set	w/ GSR	Covariance	108.2 (34.2)	107.6 (35.4)	1.3, 0.19
		Inverse	138.9 (20.9)	35.2 (5.6)	35.5, <10^{-39}
	w/o GSR	Covariance	128.4 (30.8)	128.3 (31.0)	1.0, 0.28
		Inverse	146.6 (22.8)	35.0 (4.5)	38.1,<10^{-40}

Table 3

	Full Correlation	Partial Correlation
EO	1.33 (0.46)	0.91 (0.26)
EC	1.30 (0.31)	0.95 (0.25)