

NETWORK NEURO SCIENCE

an open access  journal



Citation: Wodeyar, A., & Srinivasan, R. (2022). Structural connectome constrained graphical lasso for MEG partial coherence. *Network Neuroscience*, 6(4), 1219–1242. https://doi.org/10.1162/netn_a_00267

DOI: https://doi.org/10.1162/netn_a_00267

Supporting Information: <https://github.com/wodeyara/AdaptiveGraphicalLassoForParCoh>

Received: 13 April 2022
Accepted: 6 July 2022

Competing Interests: The authors have declared that no competing interests exist.

Corresponding Author:
Anirudh Wodeyar
wodeyar@bu.edu

Handling Editor:
Vince Calhoun

Copyright: © 2022
Massachusetts Institute of Technology
Published under a Creative Commons
Attribution 4.0 International
(CC BY 4.0) license



METHODS

Structural connectome constrained graphical lasso for MEG partial coherence

Anirudh Wodeyar^{1,2,3,4} and Ramesh Srinivasan²

¹Department of Cognitive Sciences, University of California, Irvine, California, USA

²Department of Statistics, University of California, Irvine, California, USA

³Department of Biomedical Engineering, University of California, Irvine, California, USA

⁴Department of Mathematics and Statistics, Boston University, Boston, Massachusetts, USA

Keywords: MEG, Gaussian graphical model, Coherence, Structural connectivity, Functional connectivity

ABSTRACT

Structural connectivity provides the backbone for communication between neural populations. Since axonal transmission occurs on a millisecond time scale, measures of M/EEG functional connectivity sensitive to phase synchronization, such as coherence, are expected to reflect structural connectivity. We develop a model of MEG functional connectivity whose edges are constrained by the structural connectome. The edge strengths are defined by partial coherence, a measure of conditional dependence. We build a new method—the adaptive graphical lasso (AGL)—to fit the partial coherence to perform inference on the hypothesis that the structural connectome is reflected in MEG functional connectivity. In simulations, we demonstrate that the structural connectivity’s influence on the partial coherence can be inferred using the AGL. Further, we show that fitting the partial coherence is superior to alternative methods at recovering the structural connectome, even after the source localization estimates required to map MEG from sensors to the cortex. Finally, we show how partial coherence can be used to explore how distinct parts of the structural connectome contribute to MEG functional connectivity in different frequency bands. Partial coherence offers better estimates of the strength of direct functional connections and consequently a potentially better estimate of network structure.

INTRODUCTION

Electrophysiological signals are sampled on a millisecond time scale capturing aggregate synaptic activity from populations of neurons. These neuro-physiological signals have intrinsic time scales, organized in frequency bands; and intrinsic spatial organization, organized by functional localization and integrated by the anatomical connectivity (Nunez & Srinivasan, 2006). Functional connectivity (FC) (Friston, 2011) refers to statistical dependence between signals recorded from two different areas of the brain, usually measured in a predefined frequency band. This broad definition encompasses different preprocessing methods and statistical models that emphasize different temporal and spatial scales of the underlying brain activity.

Coherence is a widely used measure of electroencephalography and magnetoencephalography (M/EEG) functional connectivity (Nunez & Srinivasan, 2006). Coherence is modulated

across different cognitive tasks and clinical disease states (Baillet, 2017; Gross et al., 2013; Rouhinen, Panula, Palva, & Palva, 2013; Roux & Uhlhaas, 2014; Siebenhühner, Wang, Palva, & Palva, 2016; Stam, 2014). Coherence is expected to reflect delayed signal transmission along white-matter tracts, that is, structural connections (Abdelnour, Voss, & Raj, 2014; Chu et al., 2015; Meier et al., 2016; Nunez & Srinivasan, 2006; Schneider, Dann, Sheshadri, Scherberger, & Vinck, 2020; Srinivasan, Winter, Ding, & Nunez, 2007) and is thus used to characterize network structure.

Leakage effects:
Artifactual shared activity across
brain regions after M/EEG source
localization.

However, mapping coherence to the anatomy is difficult due to its susceptibility to inflation from leakage effects. Leakage effects are the shared activity across brain sources caused by the limited resolution of source localization methods (Baillet, Mosher, & Leahy, 2001; Brookes et al., 2011; Gross et al., 2001; Hämäläinen & Ilmoniemi, 1994; Wipf & Nagarajan, 2009) to spatially separate source activity mixed by EEG volume conduction and MEG field spread (Nunez & Srinivasan, 2006; Srinivasan et al., 2007). Leakage effects result in common signals with zero phase difference between sources. One approach suggested to reduce leakage effects, the imaginary coherence (Nolte et al., 2004), is based on using only the projection of signals onto a phase difference of ± 90 degrees. However, this distorts the interpretation of the strength of functional connectivity, by weighting toward signals with preselected phase differences. Moreover, this approach is still susceptible to spurious connections when genuine long-range connections exist at a delay and this activity is leaked to neighboring regions (Palva et al., 2018).

We can use coherence or imaginary coherence to define the network edge weights, a critical first step for analyzing network structure, for example, using graph theoretical methods (De Vico Fallani, Richiardi, Chavez, & Achard, 2014; Maldjian, Davenport, & Whitlow, 2014; Niso et al., 2015; Schoonheim et al., 2013). However, both coherence and imaginary coherence reflect activity over single and multistep structural connectivity (Abdelnour et al., 2014; Chu et al., 2015; Meier et al., 2016). This distorts the definition of a path (Avena-Koenigsberger, Misić, & Sporns, 2018; Blinowska & Kaminski, 2013; Kaminski & Blinowska, 2018) over an undirected network and thus raises questions about the validity of network structure analyses using networks defined by the strength of coherence.

Partial coherence:
Transformation of the cross-spectral
density that allows inference of
conditional dependencies.

In contrast to the coherence or imaginary coherence, partial coherence accounts for both instantaneous and lagged shared information across multiple areas (Dahlhaus, 2000; Reid et al., 2019; Sanchez-Romero & Cole, 2021). Partial coherence has a long history in neuroscience: initially applied to spike trains (Rosenberg, Halliday, Breeze, & Conway, 1998) and generalized in Baccalá and Sameshima (2001) to the partial directed coherence. The real-valued analogue, partial correlation, has been applied to fMRI data across many studies (Hinne, Janssen, Heskes, & van Gerven, 2015; Huang et al., 2010; Ng, Varoquaux, Poline, & Thirion, 2012; Ryali, Chen, Supekar, & Menon, 2012; Smith et al., 2011; Varoquaux, Gramfort, Poline, & Thirion, 2010; Wodeyar, Cassidy, Cramer, & Srinivasan, 2020). Partial coherence represents the strength of linear relationships between a pair of brain areas when accounting for their relationships with all other brain areas (Dahlhaus, 2000; Epskamp & Fried, 2018; Whittaker, 2009). It reduces false positive detection of direct connections resulting from activity over indirect connections, as would result from leakage effects and multistep paths. Thus, we can better interpret partial coherence as connection strength to define a functional network. However, partial coherence estimation can be challenging. Most previous studies using partial coherence have focused on cases where there are only a few nodes in the network or used the L2-norm penalization for regularization (Baccalá & Sameshima, 2001; Colclough et al., 2016; Dahlhaus, 2000; Medkour, Walden, & Burgess, 2009; Ter Wal et al., 2018), without obvious justification. The use of the L2-norm is counterintuitive, as

L2-norm penalization:
Penalizing the sum of squares of
precision weights, which we
normalize to get the partial
coherence.

Penalization structure:
The L1-norm penalties chosen for the
edges and non-edges of the network
constraint provided to the adaptive
graphical lasso.

L1-norm penalization:
Estimating the precision while penalizing the sum of the absolute values of the precision.

Adaptive graphical lasso:
Extension of L1-norm penalization that permits different penalties to be applied to different edges of the precision.

the structural connectivity of the brain is known to be sparse, and there is little reason to minimize the edge strengths. In the fMRI literature, when estimating partial correlation, several studies have experimented with alternative regularization approaches: L1-norm (Huang et al., 2010), elastic net (Ryali et al., 2012), group-based penalization approaches (Varoquaux et al., 2010), edge-specific penalization (Ng et al., 2012), as well as Bayesian approaches to estimation (Hinne et al., 2015). However, these alternative regularization approaches have not been attempted in partial coherence estimation, in part because of the difficulty in implementing them.

We expect that functional connectivity is constrained by the structural connectome. In this article, we make explicit use of the structural connectome to facilitate regularization of partial coherence estimates. We use a graphical lasso technique modified to use the structural connectome to guide the L1 penalization, a method we call the adaptive graphical lasso (AGL). To our knowledge, this is the first time that the graphical lasso (L1-norm), and further the graphical lasso using a constraint-based penalization, has been used to estimate partial coherence for neural signals (Colclough et al., 2016; Ter Wal et al., 2018). We select the lasso penalization through a novel cross-validation technique that separately identifies the optimal penalization on and off the structural connectome. If the penalization is lower for edges in the structural connectome, we have clearly identified that the pattern of connectivity is influenced by the structural connectome. Note that the entire structural connectome need not be estimated in the partial coherence, a subset may be estimated as a function of the data. Through simulations, we aim to demonstrate that (1) the partial coherence can be estimated accurately using the AGL, (2) we can directly test whether the structural connectome is a useful constraint in network identification, and (3) the partial coherence serves as a better functional connectivity metric than the coherence or imaginary coherence. Finally, we use the AGL-estimated partial coherence to demonstrate distinct contributions of the structural connectome to MEG signals in different frequency bands.

METHODS

Overview

This work is guided by the intuition that the statistics of neural activity data collected at the mesoscale (intracranial electrocorticography - ECoG) and macroscale (M/EEG) are constrained by structural connectivity of the axon fiber systems of the cortex. As such, we have built a minimal generative computational model, representing the partial coherence, that is derived from estimates of structural connectivity and we have developed a method to infer model parameters. We allowed the structural connectivity to potentially guide the estimation of the partial coherence and developed new simulations to link this work with M/EEG and ECoG data.

Structural Connectome

We built a template of the structural connectome (SC) from a probabilistic atlas. We used streamlines generated with deterministic tractography by Yeh et al. (2018) using the HCP842 dataset (Van Essen et al., 2013) transformed to the MNI152 template brain obtained from the FMRIB Software Library (FSL). In this dataset experts vet the streamlines to remove potentially noisy estimates of axonal fibers. We applied the Lausanne parcellation (Cammoun et al., 2012) of 114 cortical and 15 subcortical regions of interest (ROIs) to the MNI152 template brain and generated a volumetric representation for each region of interest using the *easy_lausanne* toolbox (Cieslak, 2015). Each streamline was approximated by a single 100-point cubic spline using code adapted from the *along-tract-stats* toolbox (Colby et al.,

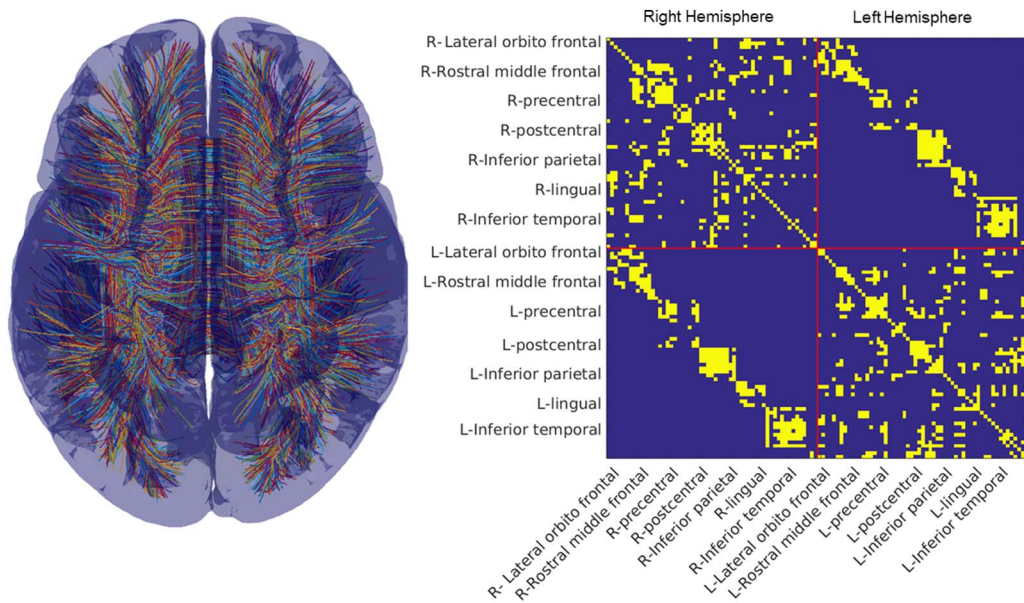


Figure 1. Structural connectome. (Left) We show streamlines derived from the work of Yeh et al. (2018) using the HCP-842 dataset. (Right) The structural connectome for the 114 areas of the Lausanne parcellation. We have labeled a subset of areas each with one to three subdivisions (see Cammoun et al., 2012, for all subdivisions of the Lausanne parcellation). We show the undirected and unweighted SC, with any nonzero edge being shown in yellow.

2012). By identifying the streamlines which terminated in a pair of ROIs, we were able to create the SC for the Lausanne parcellation. Each streamline only connected a single pair of ROIs. An edge W_{ij} for ROIs i and j existed if there was a streamline connecting the pair.

From this process, we built the 129×129 undirected and unweighted structural connectome with 1,132 edges. We reduced this matrix to 114×114 with 720 edges (see Figure 1) after removing all the subcortical structures and limiting interhemispheric connections to homologous white-matter tracts. This latter step helped remove potentially noisy estimates of connections (while potentially increasing false negatives) where streamlines intersected and passed outside the cortical surface before reaching the terminal point in a brain region. The resulting template of structural connectivity shown in Figure 1 is referred to as the structural connectome (SC). This template is incomplete in that it does not include subcortical to cortical projections. Thus, functional connectivity resulting from structural connections not captured by this template may exist in the data. Our estimation procedure for the graphical models of functional connectivity described in the next section allows for such connections, if needed, to account for the statistical structure in the data.

Generative Model

Complex-valued Gaussian graphical model. We assume that a vector of activity (\mathbf{Z}) in one frequency band is a sample drawn from a complex-valued multivariate Gaussian. Here Φ is the precision—the unnormalized partial coherence—and is determined by the SC:

$$\mathbf{Z} \sim N(\mathbf{0}, \Phi) \tag{1}$$

In the frequency domain, a signal can be characterized by samples of amplitude and phase, or equivalently, by complex-valued coefficients with real and imaginary parts corresponding to sine and cosine components of the signal.

The complex-valued multivariate Gaussian for a zero-mean (where $E(\mathbf{Z}) = 0 + 0i$) process (Schreier & Scharf, 2010) is defined as:

$$\rho(\mathbf{Z}) = \frac{1}{\pi^n \det^{\frac{1}{2}}(\Sigma)} \exp\left(-\frac{1}{2} \mathbf{Z} \Sigma^{-1} \mathbf{Z}^H\right) \quad (2)$$

where

$$\Sigma = \begin{bmatrix} \mathbf{R}_{\mathbf{z}\mathbf{z}} & \tilde{\mathbf{R}}_{\mathbf{z}\mathbf{z}} \\ \tilde{\mathbf{R}}_{\mathbf{z}\mathbf{z}}^H & \mathbf{R}_{\mathbf{z}\mathbf{z}}^H \end{bmatrix} \quad (3)$$

and

$$\mathbf{R}_{\mathbf{z}\mathbf{z}} = E[\mathbf{z}\mathbf{z}^H]; \tilde{\mathbf{R}}_{\mathbf{z}\mathbf{z}} = E[\mathbf{z}\mathbf{z}^T] \quad (4)$$

The key parameter in this model is the covariance matrix Θ and its inverse, the *precision* matrix $\Phi = \Theta^{-1}$. As defined in Equations 3 and 4, the covariance matrix for complex-valued data is composed of the familiar cross-spectrum $\mathbf{R}_{\mathbf{z}\mathbf{z}}$ and the complementary cross-spectrum $\tilde{\mathbf{R}}_{\mathbf{z}\mathbf{z}}$. Most spectral analysis methods only make use of $\mathbf{R}_{\mathbf{z}\mathbf{z}}$ and implicitly assume circular symmetry, that is, $\tilde{\mathbf{R}}_{\mathbf{z}\mathbf{z}} = 0$ (Schreier & Scharf, 2010). In this case, the complex-valued data is labeled as *proper*. With the assumption of circular symmetry, we can parameterize the complex-valued Gaussian using the precision as:

$$\Phi = \mathbf{R}_{\mathbf{z}\mathbf{z}}^{-1} \quad (5)$$

Each value in the precision matrix Φ is the conditional covariance between any two variables (here, sources representing two ROIs) given the other variables (all other ROIs). The precision represents a model of functional connectivity—the *conditional dependence* between sources. The strength of the conditional dependence represents the linear relationship between any pair of sources when linear effects from all other sources are removed (see Section 2.2.2 of Pourahmadi, 2011 for an intuitive explanation in terms of multivariate linear regression). For any pair of sources, if the precision is zero, there is no need for a relationship between the sources to account for observed coherence. Such apparent coherences arise from connections mediated via other sources in the model. Note that the precision directly represents a complex-valued Gaussian graphical model (Whittaker, 2009).

In the generative model, we choose to set up the precision matrix Φ to have a nonzero entry only at edges that have a connection in the SC. We are assuming that in each frequency band, coherence represents the result of joint random fluctuations of a set of oscillators whose connections are determined by the SC. The precision values are estimated using the graphical lasso in a cross-validated procedure that allows potentially using the SC as a guide for the L1 penalization. In this way the nonzero locations and values of the precision are determined by the data.

Adaptive graphical lasso. The graphical lasso (Friedman, Hastie, & Tibshirani, 2008) is a method that has been applied in multiple fields in the past decade, from genomics (Menéndez, Kourmpetis, ter Braak, & van Eeuwijk, 2010) to fMRI functional connectivity (Ng et al., 2012; Ryali et al., 2012; Varoquaux et al., 2010; Wodeyar et al., 2020) and climate models (Zerener, Friederichs, Lehnertz, & Hense, 2014). It is used to identify a sparse approximation to the regularized precision matrix while solving problems arising from rank deficiency and small numbers of samples. To apply the lasso, we optimize the penalized likelihood function for a

multivariate Gaussian (Meinshausen & Bühlmann, 2006) to estimate the precision—where Θ (R_{zz} in Equation 4) is the cross-spectral density (CSD):

$$\hat{\Phi} = \arg \min_{\Phi > 0} \left(-\log(\det \Phi) + \text{tr}(\Theta \Phi) + \lambda \sum_{j < k} |\Phi_{jk}| \right) \quad (6)$$

The penalization parameter λ in the graphical lasso determines the nonzero set of precision values. The output of the lasso from Equation 6 is the precision matrix $\hat{\Phi}$.

We made use of the lasso to estimate the precision while taking advantage of the knowledge of the SC to hypothesize the likely locations of nonzero precision values. We made use of the lasso optimization from quadratic approximation for sparse inverse covariance or QUIC (Hsieh, Dhillon, Ravikumar, & Sustik, 2011) using a matrix penalty term (this process is also called the adaptive lasso (Zou, 2006) determined by the SC with edges \mathbf{W} (and $\lambda_1 = \lambda_2$):

$$\hat{\Phi} = \arg \min_{\Phi > 0} \left(-\log(\det \Phi) + \text{tr}(\Theta \Phi) + \lambda_1 * \sum_{j < k; \mathbf{W}_{jk} \in SC} |\Phi_{jk}| + \lambda_2 * \sum_{j < k; \mathbf{W}_{jk} \notin SC} |\Phi_{jk}| \right) \quad (7)$$

Note that in the limiting case of $\lambda_1 = \lambda_2$, the likelihood function is the same as it is for the graphical lasso. We determine the λ_1 and λ_2 using cross-validation. This crucial setup simultaneously provides (1) a measure of the usefulness of the SC as a hypothesis on MEG functional connectivity and (2) serves as a principled thresholding mechanism for weak connections.

By optimizing the penalized likelihood, we leveraged the information in the SC as a hypothesis for our lasso estimate. We derive the graph G with vertices $V = 1, 2, \dots, 114$ and edges $W_{est} = G_{ij} = 1, i, j \in V$ from the precision based on the nonzero values in $\hat{\Phi}$. The final precision matrix $\tilde{\Phi}$ is estimated under the unpenalized Gaussian likelihood for the set of edges \mathbf{W}_{est} defined by the graphical model using the function *ggmFitHtf* (PMTK3 toolbox; Murphy & Dunham, 2008) which optimizes (unpenalized Gaussian log-likelihood):

$$\tilde{\Phi} = \arg \min_{\Phi > 0; |\Phi| > 0 = G} - (\log(\det \Phi) + \text{tr}((\Theta + \delta * I)\Phi)) \quad (8)$$

Since Θ (covariance) is usually rank deficient, we add a small value (δ) along the diagonal to make it full rank. We fixed δ as 0.001 times the maximum value along the upper triangle of the covariance.

Cross-validation. We test whether the AGL produced estimates of the precision that show reduced error relative to applying the graphical lasso using cross-validation. Note that applying the graphical lasso would be equivalent to having the penalization inside and outside the SC be equal, that is, $\lambda_1 = \lambda_2$. We estimated the appropriate value for λ_1 and λ_2 using cross-validation.

We split data into four ensembles, and repeated the following analysis with each ensemble. We estimated the precision $\tilde{\Phi}_i$ on one ensemble of the data (i) and estimated the deviance when using this precision as the inverse for the covariance Θ_j for all the other ensembles j of the data (and vice versa). Deviance was estimated as:

$$Dev = \sum_{i=1:4} \sum_{j=1:4; j \neq i} (-\log(\det \tilde{\Phi}_i) + \text{tr}(\Theta_j \tilde{\Phi}_i)) \quad (9)$$

Partial coherence. In every frequency band, or for each iteration of our simulation, we estimated the precision for complex-valued data incorporating amplitude and phase for a

Deviance:
Estimate of the goodness-of-fit between the estimated precision and out-of-sample estimated cross-spectral density.

frequency band. The normalization of the precision (Φ) yields the partial coherence (PC) (Dahlhaus, 2000), estimated using:

$$PC_{z_1 z_2} = \left| \frac{\Phi_{z_1 z_2}}{\sqrt{\Phi_{z_1 z_1} * \Phi_{z_2 z_2}}} \right|^2 \quad (10)$$

Contemporary Methods for Functional Connectivity

We considered three alternative methods to compare the partial coherence model estimated from AGL: coherence, imaginary coherence, and the partial coherence estimated when regularizing using the L2 norm. We estimate coherence C from the cross-spectral density Θ , where z_1, z_2 are the amplitude and phase information in one frequency band from two sources, as:

$$C_{z_1 z_2} = \left| \frac{\Theta_{z_1 z_2}}{\sqrt{\Theta_{z_1 z_1} * \Theta_{z_2 z_2}}} \right|^2 \quad (11)$$

Imaginary coherence is believed to reduce the influence of volume conduction and zero phase lag connectivity (such as would exist from source leakage). The idea is to minimize this effect by estimating the consistency of the imaginary part of the cross-spectral density between two sources. We measure it using (where *imag* refers to the imaginary component of the complex value from the cross-spectral density):

$$IC_{z_1 z_2} = \left| \frac{\text{imag}(\Theta_{z_1 z_2})}{\sqrt{\Theta_{z_1 z_1} * \Theta_{z_2 z_2}}} \right|^2 \quad (12)$$

Coherence and imaginary coherence networks are defined using a threshold derived using bootstrapping (Zoubir & Boashash, 1998). We define a population distribution by resampling 1,000 times with replacement. We kept C or IC edges with distributions that did not cover 0 at an alpha value of 0.05.

Finally, we consider an alternative regularization to estimate the partial coherence—an L2-norm penalization. This style of regularization does not force precision values to zero but instead minimizes them to optimize the likelihood. The penalized likelihood for the L2 norm inverse is:

$$\hat{\Phi} = \arg \min_{\Phi > 0} \left(-\log(\det \Phi) + \text{tr}(\Theta \Phi) + \eta * \sum_{j < k} |\Phi_{jk}|^2 \right) \quad (13)$$

We need to identify a threshold for inference on the edges of the precision. Using a novel cross-validation procedure that mirrors the approach we applied under the AGL (using the likelihood function to estimate deviance), we optimize for the L2-norm penalization (η) and the threshold. The threshold to be applied is determined as a percentile—between 5 to 95—of the weights whose optimal value is identified using cross-validation.

Simulations

Overview. We wished to test the accuracy of the AGL to estimate the precision matrix. To do so we simulate from a generative model and attempt to recover the parameters. The generative model we use is a complex-valued multivariate normal where the nonzero values in the precision define an undirected network (as specified in Equation 4). For each simulation, and each iteration, we generated new networks with random weights for edges. While the edge locations are kept consistent within a simulation, we randomized the weights on the edges.

The internal variability of each area/node changes across simulation iterations therefore changing the signal-to-noise ratio for each edge. We examined each simulation under two (or more) sampling scenarios—one where the number of samples is comparable to the number of nodes and one where there are many more samples than the number of nodes. For each simulation, where we always possess ground truth information, we assessed whether the AGL inferred (1) the usefulness of the network constraint, (2) recovered the true edges, (3) controlled the false positives, and (4) correctly estimated the edge weights of the partial coherence.

Simulation 1: Structural connectome simulation. In all three simulations, to generate novel precision matrices, we retained the edge locations from the original SC but simulated random weights for the edges sampled from a normal distribution, $N(100, 30)$. Finally, each edge is assigned a random phase (μ) based on sampling from a Gaussian distribution (mean = $\frac{\pi}{2}$, $SD = 0.25$). After multiplying each edge weight with the phase, we can generate the precision. This represents the complex-valued, circularly symmetric precision matrix (Φ) for a frequency band. We tested whether the precision is positive-definite by attempting to generate the Cholesky factorization of the matrix using the MATLAB function *chol*. If not, we continuously added the summed absolute value of the rows to the diagonal until the matrix was positive-definite.

Using the precision, we determined the cross-spectral density as its inverse ($\Theta = \Phi^{-1}$). The cross-spectral density has a real-valued equivalent representation (Schreier & Scharf, 2010). We can treat the real and imaginary components of the CSD as separate variables governed by a joint covariance structure. Complex-valued Gaussian values were sampled using the MATLAB function *mvnrand* operating on the real-valued CSD.

Simulation 2: Fake network constraint. In the second simulation we examined whether the AGL permits inference about the hypothesized network, that is, can we use the penalizations chosen under cross-validation to judge accuracy of the hypothesized network. We began with the same approach as in the first simulation, generating precision matrices and samples from the true structural connectome. However, we changed how we applied the AGL. Rather than use the true network, we provided a fake network generated by shuffling the nodes of the structural connectome, thus allowing us to preserve the degree distribution of the original network. We shuffle nodes using the *randperm* function in MATLAB to generate 114 integers between 1 and 114 without repetition. Every iteration of the simulation, we shuffled the nodes of the SC so that the number of edges and general connectome structure are retained while the actual node identities are altered. The penalization structure under a fake network is expected to revert to the vanilla graphical lasso, with constant penalization across the entire matrix. We collapsed results across all iterations to assess if this occurred.

Simulation 3: Forward solution and source localization simulation. In the third simulation, we generated pseudo-MEG data. This simulation tested the ability of different methods to overcome the spatial blurring induced by the process of source localization—leakage effects and incomplete demixing of source signals. For a visual depiction of this simulation, please see Figure 2.

We first built an MEG forward model—an estimate of the magnetic field measured at MEG sensors above the scalp generated by current sources located in the brain. We built the forward model for the Neuromag MEG system consisting of 306 MEG coils at 102 locations above the scalp (shown in Figure 2). At each location, there are 3 sensors—one magnetometer that measures the component of the magnetic field passing through the coil and two planar gradiometers that measure the gradient of this magnetic field in two orthogonal directions. We made use only of the orthogonal pair of planar gradiometer coils (102 pairs of sensors at 102

Forward model:
Estimate of the mixing of the magnetic field at the MEG sensors from brain source of activity.

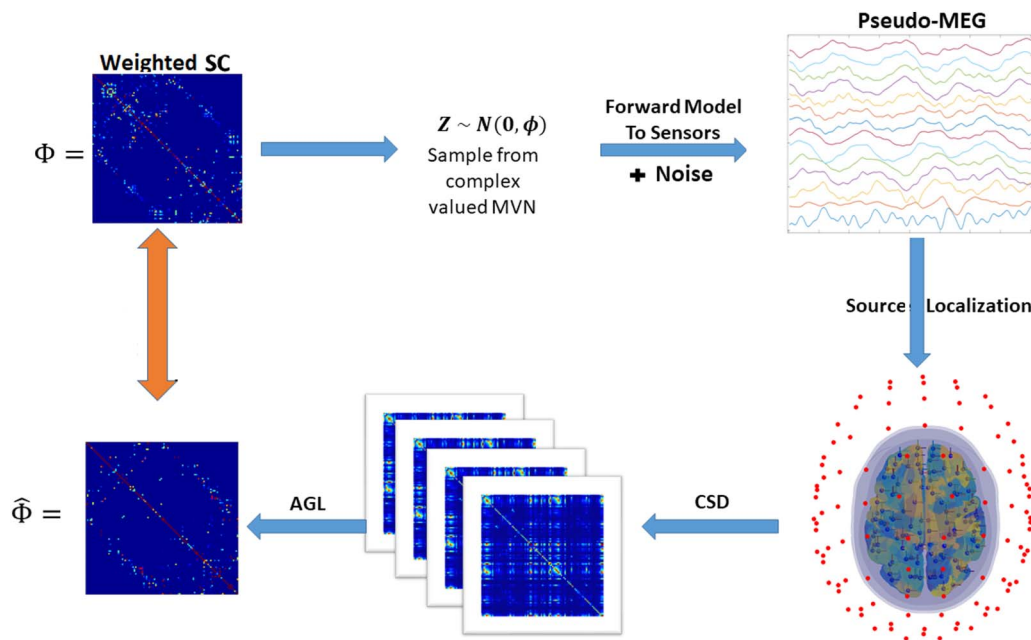


Figure 2. Source localization simulation, describing each figure above in a clockwise manner from top left. First, we used the SC edge locations to constrain the precision on each iteration of the simulation. We generated random weights and phases for each edge. Second, we sampled from a complex-valued multivariate normal distribution using the precision generated in the first step. Third, we used the MEG forward matrix to forward model the samples to the sensors. Fourth, we applied an inverse solution to source localize data. Fifth, we split the data into four ensembles of 120 samples (represented are the four covariance matrices from these ensembles of data). Finally, these four ensembles served as the input for the adaptive graphical lasso. The estimated precision from this procedure was compared to the original precision (orange arrow) by examining the penalizations applied, the edges recovered and edge weights recovered.

locations), as planar gradiometer coils have better spatial resolution than magnetometer coils thus facilitating source localization (Malmivuo, Malmivuo, & Plonsey, 1995).

The forward model is built for a specific head model, which we developed here from the *fsaverage* MRI image from the Freesurfer toolbox (Fischl, 2012). The tessellated cortical surfaces for right and left hemisphere were extracted using the *recon-all* pipeline in Freesurfer and then downsampled to 81,000 (81k) vertices (*mris_decimate* from Freesurfer). We used this surface to constrain dipole orientation and define the volume of the model corresponding to the cortex. We generated the inner skull, outer skull, and scalp surfaces approximated with 2,562 vertices from the *fsaverage* head generated using the *mri_watershed* function. Using these surfaces, and with the conductivities of the scalp, CSF and brain set at 1 S/m and the skull at 0.025 S/m (i.e., 40 times lower conductivity), we applied the OpenMEEG toolbox (Gramfort, Papadopoulos, Olivi, & Clerc, 2010) to compute a boundary element model (BEM). Each row of the MEG forward matrix from the BEM is the magnetic field gradient detected across all 204 gradiometers from a unit current density source at one of the 81k cortical surface vertices.

Using the Lausanne parcellation for 114 cortical ROIs (Cammoun et al., 2012), we subdivided the cortical surface and identified vertices belonging to each ROI using the volumetric parcellation of the *fsaverage* brain. Using this organization of vertices we then reduced the representation of the current source for each ROI down to a set of three dipoles in the *x*, *y*, and *z* directions at a single location. The location of the source for each ROI was selected by taking a weighted average of vertex locations where the weight of each location was determined by the magnitude (L2 norm) of the field generated at the gradiometers. In this way, we

reduced our source model to 114 source locations, with three sources at each location in the canonical x , y , and z directions. We computed a new MEG forward matrix (\mathbf{M}) of dimension 204×342 using OpenMEEG which approximates the linear mixing of source activity at the gradiometers to generate the measured MEG signals.

We simulate source activity \mathbf{S} across 114 areas using the precision with edges determined by the structural connectome, that is, one sample from the real-valued equivalent of the inverse of the precision is a 114×1 vector. To this source activity, we added independent noise with variance set such that the ratio of the trace of the noise to the CSD was controlled at 25 dB. We forward modeled the data to the MEG sensors. A sample of the MEG data is represented as a complex-valued vector \mathbf{B} of length equal to the number of MEG sensors (204 sensors). The set of samples of \mathbf{B} relates to source activity \mathbf{S} by:

$$\mathbf{B} = \mathbf{M}\mathbf{S} \quad (14)$$

where \mathbf{M} is the MEG forward matrix.

We localize activity to the 342 sources (three directions, along x , y , and z axes at 114 locations) by inverting the reduced lead field using regularized weighted minimum norm estimation (weighted L2 norm; Dale & Sereno, 1993) and applying it to data at the scalp. We estimated the inverse \mathbf{M}^- using (where ν is a penalization term):

$$\text{diag}(\mathbf{W}) = \|\mathbf{M}\|_2^5 \quad (15)$$

$$\mathbf{M}^- = (\mathbf{M}\mathbf{W})^T \left((\mathbf{M}\mathbf{W})^T * (\mathbf{M}\mathbf{W}) + \nu \mathbf{I} \right)^{-1} \quad (16)$$

We defined ν as the 10th percentile of the weights of $\mathbf{M}\mathbf{W}$. The estimated source activity is then $\mathbf{S} = \mathbf{M}^- \mathbf{B}$. We identify the time series for the three dipoles along the x , y , and z directions. Using a singular value decomposition at each ROI, we identified the optimal orientation of the dipole as the first singular vector. Using the first singular vector at each ROI, we reduced the source data from 342×1 to 114×1 for each sample. We used the source localized data as the input to the AGL to estimate partial coherence. We also estimated the coherence, imaginary coherence, and partial coherence under the L2 norm.

Metrics for the accuracy of the functional connectivity estimates. Across all simulations we used the ground truth to help us understand the performance of different algorithms. To understand whether the AGL is better than the vanilla graphical lasso, we examined the penalization applied on the edges and nonedges of the network provided as a constraint in simulations 1, 2, and 3. Across all methods in simulations 1 and 3, we looked at the number of true edges recovered, the number of false positives estimated, and the accuracy of estimated edge weights. To ascertain the accuracy of estimated edge weights, we calculated the Pearson correlation between the Fisher r -to- z transformed edge weights across the set of true edges, that is, all edges in the ground truth model.

Application to MEG Data

MEG data. The MEG data we analyzed was shared by the Cambridge Centre for Ageing and Neuroscience (CamCAN). CamCAN funding was provided by the UK Biotechnology and Biological Sciences Research Council (grant number BB/H008217/1), together with support from the UK Medical Research Council and University of Cambridge, UK. This data was obtained from the CamCAN repository (available at <https://www.mrc-cbu.cam.ac.uk/datasets/camcan/>; Shafto et al., 2014; Taylor et al., 2017) and was conducted in accordance with the Helsinki

declaration and approved by the Cambridgeshire 2 Research Ethics Committee (reference: 10/H0308/50).

MEG data was collected using a 306 sensor VectorView MEG system (Electa Neuromag, Helsinki). The 306 sensors consisted of 102 magnetometers and 204 planar gradiometers. The data were sampled at 1000 Hz and highpass filtered at 0.3 Hz. This data was run through temporal signal space separation (tSSS; Taulu et al., 2005; MaxFilter 2.2, Elekta Neuromag Oy, Helsinki, Finland) to remove noise from external sources and to help correct for head movements (location of the head was continuously estimated using Head Position Indicator coils). MaxFilter was also used to remove the 50 Hz line noise and also to automatically detect and reconstruct noisy channels.

Spectral analysis. We extracted 480 seconds of resting-state gradiometer data for a single individual. We first applied a band-pass filter between 0.5 and 100 Hz and a notch filter at 50 Hz to remove line noise. We 97 built elliptic filters (designed using *fdesign.bandpass* function in MATLAB) with stop band set to 0.5 Hz below and above pass band, stopband attenuation set to 100 dB, and passband ripple set to 0.02. Band-pass filtering was then done using the *filtfilt* function in MATLAB to minimize phase distortion. We analyzed five frequency bands: delta (1–3 Hz), theta (4–7 Hz), alpha (8–13 Hz), beta (14–29 Hz), and gamma (30–80 Hz). Within each band we optimized the dipole orientation across 114 ROIs as described in the section describing *Simulation 3*. Using the band-pass filtered data we were able to estimate adaptively source localized data and within each frequency band. Source localized broadband data, using band-specific source dipole orientations, was multitapered and Fourier transformed in 1-second windows. We used the all frequencies in every band, to avoid averaging over frequencies, to generate a 480×114 complex-valued matrix used for estimating the cross-spectral density.

Using the complex-valued data within each frequency band, we have a 480×114 matrix which served as the input for estimating the partial coherence. We split the 480 samples from 114 sources into four continuous ensembles of 120 samples each based on the expectation that we would have robust, stationary networks estimable with 120 seconds (Chu et al., 2012). Further, having four ensembles allowed for four-fold cross-validation. Within each ensemble we estimated the cross-spectral density and, using the AGL, the precision. We then followed the same procedure as described earlier in the section on cross-validation. Thus, we had at the end of the analysis for each subject, partial coherence across all five frequency bands.

RESULTS

Simple Five Node Network

As a proof-of-concept simulation, we examined network recovery of a sparse five node network with five edges (see *Figure 3A*) representing the precision. We sampled data for each node from the inverse of the precision, the cross-spectral density. We apply the AGL to the observed data to extract the partial coherence; a network with weighted edges. We considered two cases, one where we have a small number of samples (24 independent samples) and, second, when we have a large number of samples (240 independent samples). Each simulation (24 and 240 samples) is repeated 200 times.

The cross-validation process allows the AGL to place the same penalization everywhere, thus the penalization values assess the usefulness of a network constraint. We see from the penalization distribution (*Figure 3*) that there is reduced penalization for true edges relative to nonedges, as we expected, both with 24 and 240 samples. The second metric of interest

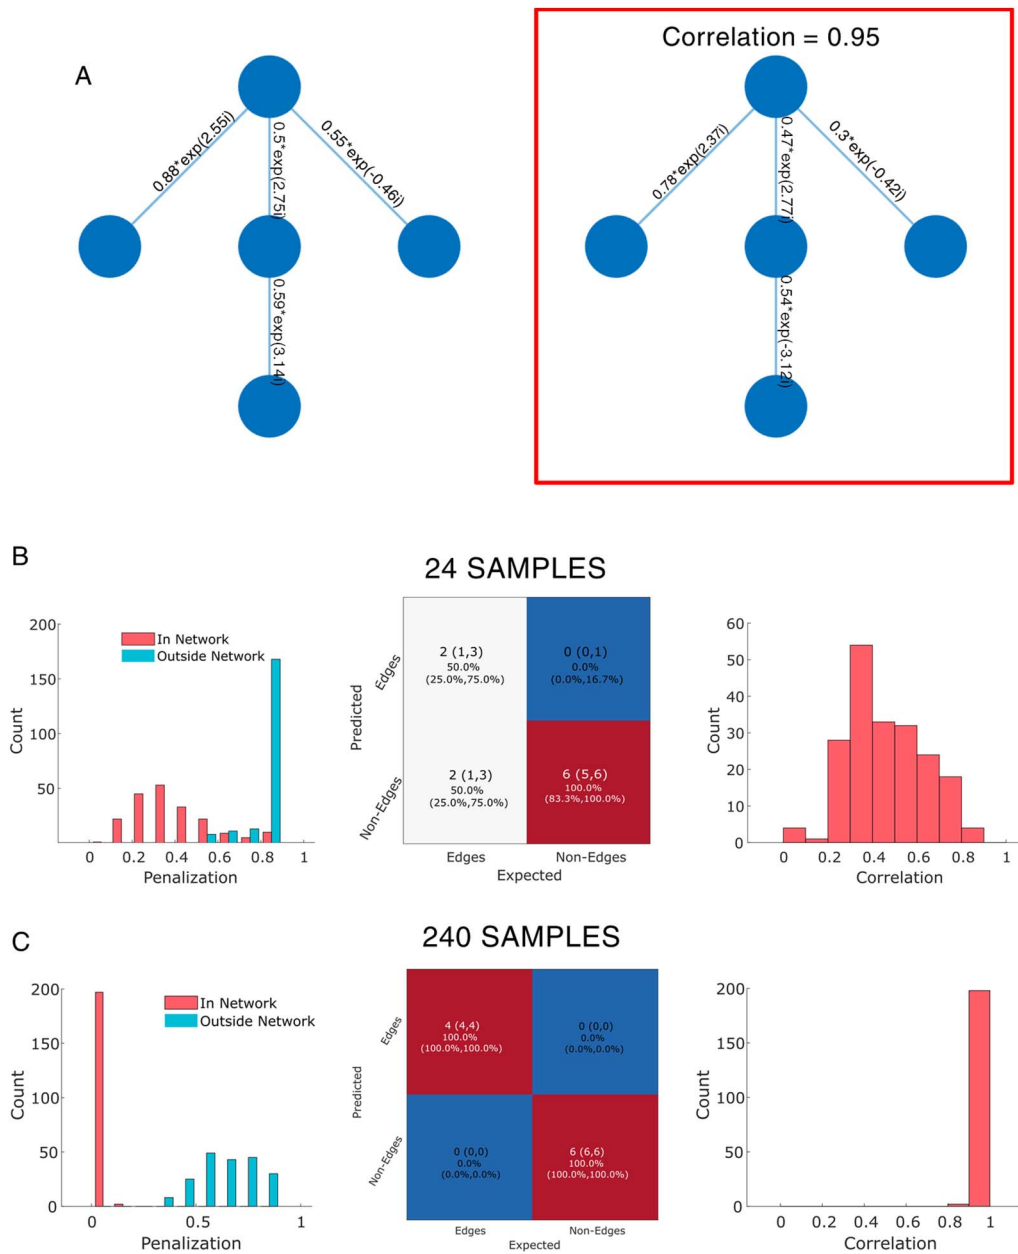


Figure 3. AGL provides accurate five node complex-valued network inference and controls false positives. (A) (Left) We simulate from a simple five node network with complex-valued edges. (Right) An example of a recovered network with accurately reconstructed weights when using 240 samples. (B, C) Left column shows penalization distributions for the constraint network provided (red) and outside the constraint network (blue) when using 24 and 240 samples across 200 simulations. Y-axis shows count of simulations and X-axis the penalization multiplier applied. Middle column shows confusion matrices with median (25th, 75th quartiles). Right column shows the correlation of edge weights across all true edges in reconstructed networks.

is edge recovery. In Figure 3B (middle column) we can see that the false positives are well controlled (with the distribution concentrated at zero edges) while we recover between two to all five of the true edges present despite only 24 available samples. With 240 samples (Figure 3C, middle column), we recovered all true edges in all 200 simulations and avoided any false positives in 95% of simulations. The final test is the recovery of the actual edge weights—the complex values representing connection strength and relative phase. We

estimate this correspondence using a correlation between the true edges and the recovered edges. A high correlation implies that the complex-valued vectors tend to align with the orientations and strengths of the original complex-valued vectors and a correlation close to 0 indicates incorrect weight and orientation (an orthogonal vector or a zero vector). From Figures 3B and 3C we can see that correlation is 0.5 with 24 samples while it is nearly 1 with 240 samples. We conclude that we are able to recover the weights and edges of the precision even when we have only 24 samples, but with (an order of magnitude) more samples, we are almost able to recover the precision perfectly.

Recovering the Structural Connectome

In the second simulation, we considered an order of magnitude increase in the number of nodes and edges. We used the structural connectome across 114 areas. The network is sparse, with 720 weighted edges out of a total possible of 6,441 edges. The inverse of the precision determined from the SC could represent the cross-spectral density estimated from intracranial electrocorticography (ECoG). Similar to the first simulation, we examine the performance of the AGL to estimate the correct partial coherence when we have 480, 960, 1,440, 1,920 and 2,400 samples. Since we simulate from a covariance structure with nonzero intra-ROI variance, the signal-to-noise ratio of each individual edge is modulated in every simulation iteration.

When simulating data from the structural connectome in a low samples case (480 samples), AGL identifies the correct penalty structure (Figure 4A, left column) and controls false positives (Figure 4A, middle column). Network recovery under AGL in a high sampling situation (2,400 samples) is nearly perfect (Figure 4B). The penalization structure consistently (across all sampling scenarios) indicated lower penalization on SC edges relative to non-SC edges, the false positives are controlled (also across all sampling scenarios) and real edges identified (≥ 500 of 720) and finally, the edge weights, that is, the partial coherence, are well recovered (correlation ≥ 0.7). This showed that the AGL is able to infer a penalization structure that uses the structural connectome. Even when we simulated only 480 samples the AGL minimized false positives, showed the usefulness of knowledge of the SC and reasonably recovered the network weights. We conclude that low numbers of samples do not pose an impossible hurdle in judging the usefulness of the structural connectome, recovering the structural connectome and controlling false positives.

Inferring an Inaccurate Structural Connectome Constraint

We forced model misspecification onto the AGL and examined the results. Model misspecification is done by altering the constraint provided to the AGL relative to the generative network model. We expect that the penalization structure of the AGL will reflect when we have used an incorrect network as a potential constraint: a shift in penalization toward the graphical lasso, that is, a uniform penalization. An alternative hypothesis is that the AGL always uses any constraint provided: the penalization can never approach the graphical lasso. We test these hypotheses by shuffling node identities for the SC network constraint provided to the AGL. However, we generate data from the structural connectome determined precision.

Examining the penalization structure (Figure 5B and 5C, left column), we find that the AGL does not place lower penalization values on the fake network edges, instead approaching the flat penalization of the graphical lasso. However, this does not imply network recovery in either the fake or the true networks (Figure 5B and 5C, middle and right column), with both the false positives and the true edges suppressed in both networks at 480 samples. Penalization

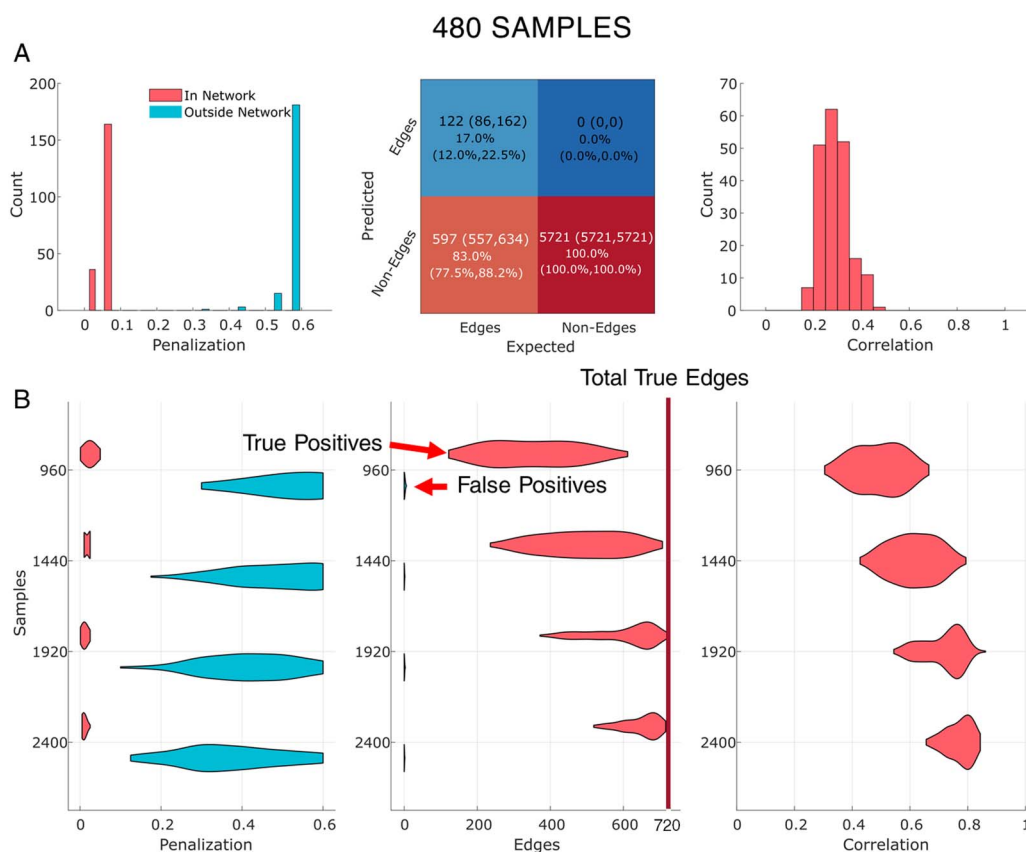


Figure 4. AGL recovers the partial coherence derived from the structural connectome with 114 nodes. (A) We ran 200 iterations of a simulation where we sampled from the weighted structural connectome as the precision. The left column shows the distribution of penalization on the structural connectome, and outside the structural connectome when we use 480 samples. Middle column shows the confusion matrix of edges recovered. Right column shows the correlation of all true edges with edges recovered. (B) We show changes in performance as a function of the sample size. Distributions are across all simulations. The left and right columns are similar to panel A, while the middle column now shows only the distributions of false and true positive edges recovered across 200 simulations. As the number of samples increases, AGL improves significantly, with network recovery almost perfect at 2,400 samples. In the middle column, the red vertical line represents the total number of true edges present (720).

at 480 samples is placed uniformly high across the whole network (on and off the fake network edges). However, at 2,400 samples, the AGL places a small uniform penalization across the whole network allowing more true edges to be estimated (Figure 5C, right column), while the false positives driven by the fake network continue to be controlled (Figure 5C, middle). This suggests that while the AGL remained constrained by the network provided, (1) an incorrect network constraint can be inferred from the penalization structure and (2) with sufficient samples the true network can be partially recovered despite an incorrect constraint.

Comparing AGL to Contemporary Network Recovery Approaches

Many contemporary algorithms aim to estimate the networks scaffolding EEG/MEG/ECoG data. We compared three methods that make similar assumptions about the data as using the AGL-estimated partial coherence : coherence (Bendat & Piersol, 2011), imaginary coherence (Nolte et al., 2004), and the partial coherence estimated under an L2-norm inverse (Colclough et al., 2016; Ter Wal et al., 2018). We first compared these methods when




Figure 5. Using a fake network constraint provides no useful information. (A) On the left is the true network from which the data is simulated. On the right we show an example fake network constraint where we shuffle nodes in the true network, thus retaining the degree distribution of the original network. (B and C) The left column shows the penalization applied inside and outside the fake network constraint. Middle column shows the confusion matrix when predicting the fake network edges. Right column shows the confusion matrix when predicting the original network edges. Note that when we have 2,400 samples, more true edges from the original network are estimated and very few in the fake network.

recovering a network with structural connectome edges using 480 samples. We compared the methods on the true positives, false positives, and the network weight recovery (Figure 6).

We found that, at 480 samples, all methods were able to recover the true SC edges; however, they also estimated a considerable number of false positive edges. The AGL was considerably better at controlling false positive edges than all other methods, with the imaginary coherence performing better than coherence and the L2 norm. When estimating network weights, the L2 norm inverse and the AGL did considerably better than the coherence and imaginary coherence. At 2,400 samples, we saw similar performance differences across




Figure 6. AGL outperforms alternative methods at SC recovery. We compare AGL to an L2 norm inverse, coherence, and imaginary coherence. We show two cases: 480 (A) and 2,400 samples (B). (A) (Left) For each method on the y-axis, there are two violin plots with the upper violin representing edges discovered in the SC (true positives) and the lower violin showing edges discovered outside the SC (false positives). No method controls the false positives as well as the AGL (see bottom violin); however, all alternative methods recover more edges in the SC. In the right column we show correlation of all true edges with edges estimated. The L2 norm inverse recovers edge weights comparably to the AGL. (B) At 2,400 samples, the AGL outperforms all other algorithms at controlling both false positives and at edge weight recovery. Note that since the L2 norm is thresholded based on percentiles of edge weights, the number of edges recovered outside the SC is discrete.

methods, with the AGL continuing to outperform all other methods at controlling false positives. Further, the AGL is better than the L2 norm inverse at estimating the network weights of the true network at 2,400 samples. We conclude that the AGL-estimated partial coherence outperforms contemporary algorithms at recovering the underlying network, both when we have limited independent samples and when we have large numbers of independent samples.

Network Recovery After Source Localization

In M/EEG research, we must both recover the network from limited samples and reduce the impact of signal leakage from source localization. We apply a commonly used source localization technique (weighted L2 norm inverse; see details in [Methods](#), section on [Simulation 3](#))




Figure 7. AGL outperforms alternative methods at network recovery under source localization. We generate data from a network in the brain with SC edges and forward model to MEG sensors. After applying source localization (weighted L2 norm inverse), we attempt to recover the original network. (A and B) In panel A we show results when using 480 samples and in panel B we show the results when using 2,400 samples. Left column shows the true and false positives discovered by each method. Right column shows the correlation of the true edges with estimated edges. AGL outperforms all algorithms at 2,400 samples at recovering true edges, controlling false positives and recovering edge weights.

to attempt to recover sources. We apply the AGL and other algorithms to this recovered source activity.

Examining the 480 sample case (Figure 7A), we see that the AGL continues to outperform all other methods at controlling false positives in the network. However, other network recovery techniques were comparable in recovering true SC edges, with the coherence recovering all edges but also including a large number of false positives. All methods were comparably poor at recovering the network weights. When we have more samples (2,400; see Figure 7B), we see that the AGL clearly outperforms all algorithms in all metrics measured, with the correlation with network weights reaching 0.58. This suggests that when more samples were available we were able to partially overcome the difficulties imposed by source localization by using the AGL-estimated partial coherence.

Application to MEG Data

We extracted 480 seconds of preprocessed resting-state MEG data from a single subject from the open-source CamCAN dataset. We source localized this data (using weighted L2 norm; Dale & Sereno, 1993) to the 114 areas of the Lausanne parcellation. After source localization, we used 1-second windows to get amplitude and phase samples at each frequency from 1 to 50 Hz using the multitaper method. We applied the AGL with our cross-validation procedure (see section Cross Validation) to estimate the partial coherence. Note that since we examined only a single subject, we intended this to only be a demonstration of how the AGL-estimated partial coherence could be used. Further, we do not have a ground truth in this situation so we focus on the penalization structure to infer if the structural connectome (SC) is useful information in modeling the coherence. We did find that the SC serves as a useful constraint in the delta (2/3 frequencies), theta (1/4 frequencies), and beta bands (11/15 frequencies), but not in the alpha (0/6 frequencies) or gamma (0/20 frequencies) bands. The null results in the alpha

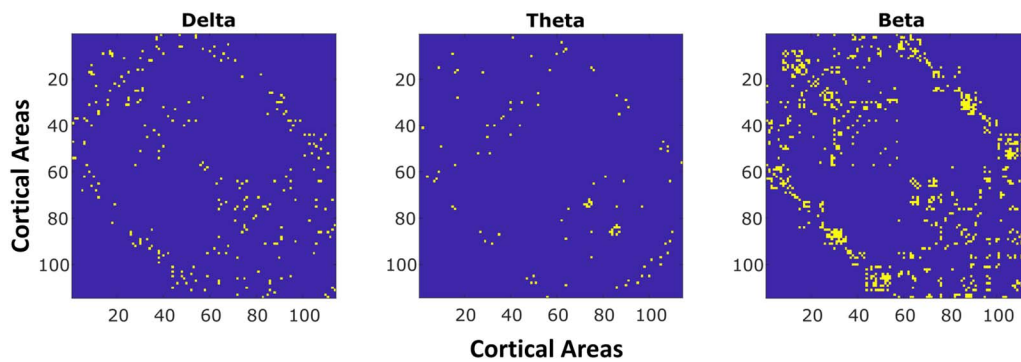


Figure 8. AGL recovers unique partial coherence networks in different frequency bands. We apply the AGL to delta, alpha, theta, beta, and gamma frequency bands. We show here results from the bands where there was a lower penalization applied to the SC. In the delta, theta, and beta bands we show the edges that were maximum in strength for that band. This selection shows how there exist different networks at each frequency with the beta band showing the greatest spread across the SC and the theta band network showing the greatest selectivity.

and gamma band indicate that the measured functional connectivity involves other connections, for example, thalamocortical or other subcortical projections not included in the structural connectome. Finally, when we applied a fake network—a shuffled SC—as a potential constraint, we found that none of the frequency bands use the constraint applied and the algorithm chose to use vanilla graphical lasso. This indicates that only the SC serves as a useful constraint in the delta, theta, and beta bands.

For the cases when the SC was a useful constraint, the partial coherence estimates the edges of the SC that are relevant for each frequency, which can be a subset of the structural connectome. We examined the edge weights in the delta, theta, and beta bands, looking for which band had the highest weight at each SC edge. We show this in Figure 8. Beta band networks tend to have connections distributed across the cortex, while theta and delta connections are more circumscribed. Delta band shows connections within frontal and cingulate regions and from frontal/cingulate to parietal regions. Theta band shows consistent connectivity across left and right hemispheres between temporal and parietal/occipital regions. Beta band connectivity dominates throughout the rest of the structural connectivity, with little specificity. The relevance of the structural connectivity to beta band functional connectivity is consistent with past research (Garcés et al., 2016). We conclude that the AGL can be applied to empirical data to discover networks in different frequency bands.

DISCUSSION

We developed a model of MEG coherence constrained by knowledge of anatomical connectivity in the structural connectome. We showed that we can accurately infer the weighted network connectivity by means of partial coherence, for the first time, using the AGL. This method can assess if the structural connectome is useful as a constraint for estimation of the partial coherence by comparing the penalization applied to the structural connectome to the penalization applied outside it. Finally, we demonstrated how, when the functional connectivity is simulated from the structural connectome, the AGL-estimated partial coherence outperforms coherence, imaginary coherence, and the L2-norm estimated partial coherence.

Functional Connectivity Network Using Partial Coherence

The AGL yields a new measure of functional connectivity that is based on the expectation that the structural connectivity scaffolds the functional connectivity. Critically, the method also

allows functional connections to exist that are not prescribed in the structural connectome. Given the ability of the partial coherence to reduce false positives and provide an accurate definition of a path (Avena-Koenigsberger et al., 2018), it serves as a useful electrophysiological functional connectivity measure for network analyses (Reid et al., 2019). Further, the precision can potentially be applied toward other analyses that attempt to decode the causal direction of connections (Baccalá & Sameshima, 2001; Sanchez-Romero & Cole, 2021).

Estimating Partial Coherence Using the Structural Connectome

Regularization is essential to estimate partial coherence for large networks. We argued that an L1 norm regularization is more intuitive than the L2 norm because the structural connectome is sparse. We can explicitly incorporate the structural connectome (SC) into the partial coherence estimate through the AGL. Past work applying a matrix penalty term to the graphical lasso (Pineda-Pardo et al., 2014)—using it to estimate the partial correlation—has directly forced the SC connection weights onto the penalization weighting. In contrast to Pineda-Pardo et al. (2014), we expected that the SC strengths are unlikely to map directly onto the strengths of the precision due to individual differences and variance within individuals across functional brain states. In addition, we expected that the SC can have different contributions across frequency bands yielding different connection weights. For these reasons we used the binarized SC to potentially organize the L1 penalization, that is, we allowed the penalization to entirely ignore the SC if appropriate.

Using Larger Numbers of Samples in Functional Connectivity Research

We found that the accuracy of network recovery is contingent on the number of samples used. While a subset of the network was recoverable when samples were comparable to nodes, from simulations it appeared that there was considerably improved performance with higher numbers of samples. While past work has suggested that for coherence there can be convergence within a few hundreds of samples (Chu et al., 2012), we saw that for the imaginary coherence and for partial coherence, larger numbers of samples significantly improved performance. This knowledge provides impetus to use longer recordings (10 minutes or more) to estimate resting-state electrophysiological functional connectivity, similar to recent work in functional-MRI research (Gordon et al., 2017).

Limitations

In the simulations, we assumed a generative model where brain areas show random oscillatory behavior linked by the structural connectome. This could be represented using a zero-mean complex multivariate normal with a circularly symmetric precision. More detailed mean-field models of neural activity may be more phenomenologically accurate (Cabral et al., 2014); although, past work suggests there is limited gain in using them when explaining empirical data (Finger et al., 2016; Messé, Rudrauf, Giron, & Marrelec, 2015). As such, there is value in having multiple models to explain the data as a function of the hypothesis being tested.

While partial coherence offers a clearer definition of a direct connection between areas, it is potentially susceptible to false positives, depending on the nature of causal direction of common effects in the network (Sanchez-Romero & Cole, 2021). Specifically, if two nodes *A* and *B* are directly influencing a third node *C*, and *A* and *B* are unassociated, then a false positive connection can appear between *A* and *B*. As such, the partial coherence potentially could be better used in concert with coherence as proposed by Sanchez-Romero and Cole (2021).

Mean-field models:
Mathematical model of the mean
neural population dynamics.

In humans, structural connectivity is only estimated from diffusion weighted imaging and is an imperfect measure, subject to its own limitations (Karnath, Sperber, & Rorden, 2018; Maier-Hein et al., 2017). There are difficulties in tractography linked to overlapping fiber bundles that make it hard to identify correct bundle endpoints, and strict correction of incorrect streamlines can rapidly lead to large numbers of false negatives (Maier-Hein et al., 2017). The decision to remove nonhomologous interhemispheric connectivity may also have introduced a few false negatives. Finally, we used a group-averaged SC template for all the subjects, and while individual variability in SC is low (Chamberland et al., 2017), better models may be built using an individualized SC estimate. An important future direction will be to examine the optimal structural connectivity estimate for MEG data.

Source localization can be formulated in several ways based on prior assumptions. While we used a weighted L_2 norm inverse, beamformer reconstruction approaches are also quite common in MEG (Brookes et al., 2011; Gross et al., 2001) and require investigation within this framework. Bayesian techniques accounting for priors more explicitly can afford better source reconstruction (Baillet et al., 2001; Wipf & Nagarajan, 2009). Examining these alternative approaches was beyond our scope, but the AGL is equally applicable under these alternatives. Additionally, we chose to limit our analysis to an SC with 114 nodes; a future extension to this work might examine cases with more (or fewer) sources. We also ignore for our purposes subcortical source activity and connectivity. This may have led to the large variation in the estimated results in the MEG data. In this example, the alpha and gamma rhythms may not have mapped onto the structural connectome because of strong thalamocortical contributions. Estimation of subcortical activity in MEG, while possible, is difficult without explicit prior knowledge (Krishnaswamy et al., 2017), and would also potentially benefit from including the magnetometer recordings and developing individual subject head models.

Conclusion

Understanding the relevance of different potential constraints, such as source modeling and the structural connectome, on a "big data" measurement technique such as MEG data improves our ability to infer genuine signal variability from noise. This work developed a simple model derived from the constraint of the structural connectome and demonstrated that we can recover the model parameters in simulations. This method is useful in clinical situations and in cognitive neuroscience for understanding network structure. For example, estimating a gamma band partial coherence network in a working memory task to understand which structural connections are most strongly activated. Another example, as we have recently demonstrated in fMRI research (Wodeyar et al., 2020), is to examine the influence of lesions and concomitant structural disconnection on MEG or EEG functional connectivity. Interpreting M/EEG coherence is contingent on building and comparing different models of the data, and we believe our work takes us a significant step in this direction.

SUPPORTING INFORMATION

Supporting information for this article is available at <https://github.com/wodeyar/AdaptiveGraphicalLassoForParCoh> (Wodeyar, 2022).

AUTHOR CONTRIBUTIONS

Anirudh Wodeyar: Conceptualization; Formal analysis; Investigation; Methodology; Software; Visualization; Writing – original draft; Writing – review & editing. Ramesh Srinivasan:

Conceptualization; Funding acquisition; Methodology; Resources; Supervision; Writing – review & editing.

REFERENCES

- Abdelnour, F., Voss, H. U., & Raj, A. (2014). Network diffusion accurately models the relationship between structural and functional brain connectivity networks. *NeuroImage*, *90*, 335–347. <https://doi.org/10.1016/j.neuroimage.2013.12.039>, PubMed: 24384152
- Avena-Koenigsberger, A., Misisic, B., & Sporns, O. (2018). Communication dynamics in complex brain networks. *Nature Reviews Neuroscience*, *19*(1), 17–33. <https://doi.org/10.1038/nrn.2017.149>, PubMed: 29238085
- Baccalá, L. A., & Sameshima, K. (2001). Partial directed coherence: A new concept in neural structure determination. *Biological Cybernetics*, *84*(6), 463–474. <https://doi.org/10.1007/PL00007990>, PubMed: 11417058
- Baillet, S. (2017). Magnetoencephalography for brain electrophysiology and imaging. *Nature Neuroscience*, *20*(3), 327–339. <https://doi.org/10.1038/nn.4504>, PubMed: 28230841
- Baillet, S., Mosher, J. C., & Leahy, R. M. (2001). Electromagnetic brain mapping. *IEEE Signal Processing Magazine*, *18*(6), 14–30. <https://doi.org/10.1109/79.962275>
- Bendat, J. S., & Piersol, A. G. (2011). *Random data: Analysis and measurement procedures* (Vol. 729). Hoboken, NJ: John Wiley & Sons. <https://doi.org/10.1002/9781118032428>
- Blinowska, K. J., & Kaminski, M. (2013). Functional brain networks: random, “small world” or deterministic? *PLoS One*, *8*(10), e78763. <https://doi.org/10.1371/journal.pone.0078763>, PubMed: 24205313
- Brookes, M. J., Woolrich, M., Luckhoo, H., Price, D., Hale, J. R., Stephenson, M. C., ... Morris, P. G. (2011). Investigating the electrophysiological basis of resting state networks using magnetoencephalography. *Proceedings of the National Academy of Sciences*, *108*(40), 16783–16788. <https://doi.org/10.1073/pnas.1112685108>, PubMed: 21930901
- Cabral, J., Luckhoo, H., Woolrich, M., Joansson, M., Mohseni, H., Baker, A., ... Deco, G. (2014). Exploring mechanisms of spontaneous functional connectivity in MEG: How delayed network interactions lead to structured amplitude envelopes of band-pass filtered oscillations. *NeuroImage*, *90*, 423–435. <https://doi.org/10.1016/j.neuroimage.2013.11.047>, PubMed: 24321555
- Cammoun, L., Gigandet, X., Meskaldji, D., Thiran, J. P., Sporns, O., Do, K. Q., ... Hagmann, P. (2012). Mapping the human connectome at multiple scales with diffusion spectrum MRI. *Journal of Neuroscience Methods*, *203*(2), 386–397. <https://doi.org/10.1016/j.jneumeth.2011.09.031>, PubMed: 22001222
- Chamberland, M., Girard, G., Bernier, M., Fortin, D., Descoteaux, M., & Whittingstall, K. (2017). On the origin of individual functional connectivity variability: The role of white matter architecture. *Brain Connectivity*, *7*(8), 491–503. <https://doi.org/10.1089/brain.2017.0539>, PubMed: 28825322
- Chu, C. J., Kramer, M. A., Pathmanathan, J., Bianchi, M. T., Westover, M. B., Wison, L., & Cash, S. S. (2012). Emergence of stable functional networks in long-term human electroencephalography. *Journal of Neuroscience*, *32*(8), 2703–2713. <https://doi.org/10.1523/JNEUROSCI.5669-11.2012>, PubMed: 22357854
- Chu, C. J., Tanaka, N., Diaz, J., Edlow, B. L., Wu, O., Hämäläinen, M., ... Kramer, M. A. (2015). EEG functional connectivity is partially predicted by underlying white matter connectivity. *NeuroImage*, *108*, 23–33. <https://doi.org/10.1016/j.neuroimage.2014.12.033>, PubMed: 25534110
- Cieslak, M. (2015). *Easy lausanne*, GitHub, https://github.com/mattcieslak/easy_lausanne
- Colby, J. B., Soderberg, L., Lebel, C., Dinov, I. D., Thompson, P. M., & Sowell, E. R. (2012). Along-tract statistics allow for enhanced tractography analysis. *NeuroImage*, *59*(4), 3227–3242. <https://doi.org/10.1016/j.neuroimage.2011.11.004>, PubMed: 22094644
- Colclough, G. L., Woolrich, M. W., Tewarie, P. K., Brookes, M. J., Quinn, A. J., & Smith, S. M. (2016). How reliable are MEG resting-state connectivity metrics? *NeuroImage*, *138*, 284–293. <https://doi.org/10.1016/j.neuroimage.2016.05.070>, PubMed: 27262239
- Dahlhaus, R. (2000). Graphical interaction models for multivariate time series. *Metrika*, *51*(2), 157–172. <https://doi.org/10.1007/s001840000055>
- Dale, A. M., & Sereno, M. I. (1993). Improved localization of cortical activity by combining EEG and MEG with MRI cortical surface reconstruction: A linear approach. *Journal of Cognitive Neuroscience*, *5*(2), 162–176. <https://doi.org/10.1162/jocn.1993.5.2.162>, PubMed: 23972151
- De Vico Fallani, F., Richiardi, J., Chavez, M., & Achard, S. (2014). Graph analysis of functional brain networks: Practical issues in translational neuroscience. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *369*(1653), 20130521. <https://doi.org/10.1098/rstb.2013.0521>, PubMed: 25180301
- Epskamp, S., & Fried, E. I. (2018). A tutorial on regularized partial correlation networks. *Psychological Methods*, *23*(4), 617–634. <https://doi.org/10.1037/met0000167>, PubMed: 29595293
- Finger, H., Bönstrup, M., Cheng, B., Messé, A., Hilgetag, C., Thomalla, G., ... König, P. (2016). Modeling of large-scale functional brain networks based on structural connectivity from DTI: Comparison with EEG derived phase coupling networks and evaluation of alternative methods along the modeling path. *PLoS Computational Biology*, *12*(8), e1005025. <https://doi.org/10.1371/journal.pcbi.1005025>, PubMed: 27504629
- Fischl, B. (2012). FreeSurfer. *NeuroImage*, *62*(2), 774–781. <https://doi.org/10.1016/j.neuroimage.2012.01.021>, PubMed: 22248573
- Friedman, J., Hastie, T., & Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, *9*(3), 432–441. <https://doi.org/10.1093/biostatistics/kxm045>, PubMed: 18079126
- Friston, K. J. (2011). Functional and effective connectivity: A review. *Brain Connectivity*, *1*(1), 13–36. <https://doi.org/10.1089/brain.2011.0008>, PubMed: 22432952

- Garcés, P., Pereda, E., Hernández-Tamames, J. A., Del-Pozo, F., Maestú, F., & Ángel Pineda-Pardo, J. (2016). Multimodal description of whole brain connectivity: A comparison of resting state MEG, fMRI, and DWI. *Human Brain Mapping, 37*(1), 20–34. <https://doi.org/10.1002/hbm.22995>, PubMed: 26503502
- Gordon, E. M., Laumann, T. O., Gilmore, A. W., Newbold, D. J., Greene, D. J., Berg, J. J., ... Dosenbach, N. U. F. (2017). Precision functional mapping of individual human brains. *Neuron, 95*(4), 791–807. <https://doi.org/10.1016/j.neuron.2017.07.011>, PubMed: 28757305
- Gramfort, A., Papadopoulos, T., Olivi, E., & Clerc, M. (2010). Open-MEEG: Opensource software for quasistatic bioelectromagnetics. *Biomedical Engineering Online, 9*(1), 45. <https://doi.org/10.1186/1475-925X-9-45>, PubMed: 20819204
- Gross, J., Hoogenboom, N., Thut, G., Schyns, P., Panzeri, S., Belin, P., & Garrod, S. (2013). Speech rhythms and multiplexed oscillatory sensory coding in the human brain. *PLoS Biology, 11*(12), e1001752. <https://doi.org/10.1371/journal.pbio.1001752>, PubMed: 24391472
- Gross, J., Kujala, J., Hämäläinen, M., Timmermann, L., Schnitzler, A., & Salmelin, R. (2001). Dynamic imaging of coherent sources: Studying neural interactions in the human brain. *Proceedings of the National Academy of Sciences, 98*(2), 694–699. <https://doi.org/10.1073/pnas.98.2.694>, PubMed: 11209067
- Hämäläinen, M. S., & Ilmoniemi, R. J. (1994). Interpreting magnetic fields of the brain: Minimum norm estimates. *Medical & Biological Engineering & Computing, 32*(1), 35–42. <https://doi.org/10.1007/BF02512476>, PubMed: 8182960
- Hinne, M., Janssen, R. J., Heskens, T., & van Gerven, M. A. (2015). Bayesian estimation of conditional independence graphs improves functional connectivity estimates. *PLoS Computational Biology, 11*(11), e1004534. <https://doi.org/10.1371/journal.pcbi.1004534>, PubMed: 26540089
- Hsieh, C.-J., Dhillon, I. S., Ravikumar, P. K., & Sustik, M. A. (2011). Sparse inverse covariance matrix estimation using quadratic approximation. In *Advances in neural information processing systems* (pp. 2330–2338). <https://proceedings.neurips.cc/paper/2011/hash/2ba8698b79439589fdd2b0f7218d8b07-Abstract.html>
- Huang, S., Li, J., Sun, L., Ye, J., Fleisher, A., Wu, T., ... Alzheimer's Disease Neuroimaging Initiative. (2010). Learning brain connectivity of Alzheimer's disease by sparse inverse covariance estimation. *NeuroImage, 50*(3), 935–949. <https://doi.org/10.1016/j.neuroimage.2009.12.120>, PubMed: 20079441
- Kaminski, M., & Blinowska, K. J. (2018). Is graph theoretical analysis a useful tool for quantification of connectivity obtained by means of EEG/MEG techniques? *Frontiers in Neural Circuits, 12*, 76. <https://doi.org/10.3389/fncir.2018.00076>, PubMed: 30319364
- Karnath, H.-O., Sperber, C., & Rorden, C. (2018). Mapping human brain lesions and their functional consequences. *NeuroImage, 165*, 180–189. <https://doi.org/10.1016/j.neuroimage.2017.10.028>, PubMed: 29042216
- Krishnaswamy, P., Obregon-Henao, G., Ahveninen, J., Khan, S., Babadi, B., Iglesias, J. E., ... Purdon, P. L. (2017). Sparsity enables estimation of both subcortical and cortical activity from MEG and EEG. *Proceedings of the National Academy of Sciences, 114*(48), E10465–E10474. <https://doi.org/10.1073/pnas.1705414114>, PubMed: 29138310
- Maier-Hein, K. H., Neher, P. F., Houde, J.-C., Côté, M.-A., Garyfallidis, E., Zhong, J., ... Descoteaux, M. (2017). The challenge of mapping the human connectome based on diffusion tractography. *Nature Communications, 8*(1), 1349. <https://doi.org/10.1038/s41467-017-01285-x>, PubMed: 29116093
- Maldjian, J. A., Davenport, E. M., & Whitlow, C. T. (2014). Graph theoretical analysis of resting-state MEG data: Identifying inter-hemispheric connectivity and the default mode. *NeuroImage, 96*, 88–94. <https://doi.org/10.1016/j.neuroimage.2014.03.065>, PubMed: 24699016
- Malmivuo, P., Malmivuo, J., & Plonsey, R. (1995). *Bioelectromagnetism: Principles and applications of bioelectric and biomagnetic fields*. New York, NY: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195058239.001.0001>
- Medkour, T., Walden, A. T., & Burgess, A. (2009). Graphical modeling for brain connectivity via partial coherence. *Journal of Neuroscience Methods, 180*(2), 374–383. <https://doi.org/10.1016/j.jneumeth.2009.04.003>, PubMed: 19375455
- Meier, J., Tewarie, P., Hillebrand, A., Douw, L., van Dijk, B. W., Stufflebeam, S. M., & Van Mieghem, P. (2016). A mapping between structural and functional brain networks. *Brain Connectivity, 6*(4), 298–311. <https://doi.org/10.1089/brain.2015.0408>, PubMed: 26860437
- Meinshausen, N., & Bühlmann, P. (2006). High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics, 34*(3), 1436–1462. <https://doi.org/10.1214/009053606000000281>
- Menéndez, P., Koumpetis, Y. A. I., ter Braak, C. J. F., & van Eeuwijk, F. A. (2010). Gene regulatory networks from multifactorial perturbations using graphical lasso: Application to the DREAM4 challenge. *PLoS One, 5*(12), e14147. <https://doi.org/10.1371/journal.pone.0014147>, PubMed: 21188141
- Messé, A., Rudrauf, D., Giron, A., & Marrelec, G. (2015). Predicting functional connectivity from structural connectivity via computational models using MRI: An extensive comparison study. *NeuroImage, 111*, 65–75. <https://doi.org/10.1016/j.neuroimage.2015.02.001>, PubMed: 25682944
- Murphy, K., & Dunham, M. (2008). PMTK: Probabilistic modeling toolkit. In *Neural information processing systems (NIPS) workshop on probabilistic programming*. <https://www.cs.ubc.ca/~murphyk/pmtk/pmtk1/pmtk-nips08-slides.pdf>
- Ng, B., Varoquaux, G., Poline, J. B., & Thirion, B. (2012). A novel sparse graphical approach for multimodal brain connectivity inference. In *International conference on medical image computing and computer-assisted intervention* (pp. 707–714). Berlin, Germany: Springer. https://doi.org/10.1007/978-3-642-33415-3_87, PubMed: 23285614
- Niso, G., Carrasco, S., Gudín, M., Maestú, F., Del-Pozo, F., & Pereda, E. (2015). What graph theory actually tells us about resting state interictal MEG epileptic activity. *NeuroImage: Clinical, 8*, 503–515. <https://doi.org/10.1016/j.nicl.2015.05.008>, PubMed: 26106575
- Nolte, G., Bai, O., Wheaton, L., Mari, Z., Vorbach, S., & Hallett, M. (2004). Identifying true brain interaction from EEG data using the imaginary part of coherency. *Clinical Neurophysiology, 115*(10), 2292–2307. <https://doi.org/10.1016/j.clinph.2004.04.029>, PubMed: 15351371
- Nunez, P. L., & Srinivasan, R. (2006). *Electric fields of the brain: The neurophysics of EEG*. New York, NY: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195050387.001.0001>

- Palva, J. M., Wang, S. H., Palva, S., Zhigalov, A., Monto, S., Brookes, M. J., ... Jerbi, K. (2018). Ghost interactions in MEG/EEG source space: A note of caution on inter-areal coupling measures. *NeuroImage*, *173*, 632–643. <https://doi.org/10.1016/j.neuroimage.2018.02.032>, PubMed: 29477441
- Pineda-Pardo, J. A., Bruña, R., Woolrich, M., Marcos, A., Nobre, A. C., Maestú, F., & Vidaurre, D. (2014). Guiding functional connectivity estimation by structural connectivity in MEG: An application to discrimination of conditions of mild cognitive impairment. *NeuroImage*, *101*, 765–777. <https://doi.org/10.1016/j.neuroimage.2014.08.002>, PubMed: 25111472
- Pourahmadi, M. (2011). Covariance estimation: The GLM and regularization perspectives. *Statistical Science*, *26*(3), 369–387. <https://doi.org/10.1214/11-STS358>
- Reid, A. T., Headley, D. B., Mill, R. D., Sanchez-Romero, R., Uddin, L. Q., Marinazzo, D., ... Cole, M. W. (2019). Advancing functional connectivity research from association to causation. *Nature Neuroscience*, *22*(11), 1751–1760. <https://doi.org/10.1038/s41593-019-0510-4>, PubMed: 31611705
- Rosenberg, J. R., Halliday, D. M., Breeze, P., & Conway, B. A. (1998). Identification of patterns of neuronal connectivity—Partial spectra, partial coherence, and neuronal interactions. *Journal of Neuroscience Methods*, *83*(1), 57–72. [https://doi.org/10.1016/S0165-0270\(98\)00061-2](https://doi.org/10.1016/S0165-0270(98)00061-2), PubMed: 9765051
- Rouhinen, S., Panula, J., Palva, J. M., & Palva, S. (2013). Load dependence of β and γ oscillations predicts individual capacity of visual attention. *Journal of Neuroscience*, *33*(48), 19023–19033. <https://doi.org/10.1523/JNEUROSCI.1666-13.2013>, PubMed: 24285906
- Roux, F., & Uhlhaas, P. J. (2014). Working memory and neural oscillations: Alpha–gamma versus theta–gamma codes for distinct wm information? *Trends in Cognitive Sciences*, *18*(1), 16–25. <https://doi.org/10.1016/j.tics.2013.10.010>, PubMed: 24268290
- 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*(4), 3852–3861. <https://doi.org/10.1016/j.neuroimage.2011.11.054>, PubMed: 22155039
- Sanchez-Romero, R., & Cole, M. W. (2021). Combining multiple functional connectivity methods to improve causal inferences. *Journal of Cognitive Neuroscience*, *33*(2), 180–194. https://doi.org/10.1162/jocn_a_01580, PubMed: 32427070
- Schneider, M., Dann, B., Sheshadri, S., Scherberger, H. G., & Vinck, M. (2020). A general theory of coherence between brain areas. *bioRxiv*. <https://doi.org/10.1101/2020.06.17.156190>
- Schoonheim, M. M., Geurts, J. J., Landi, D., Douw, L., van der Meer, M. L., Vrenken, H., ... Stam, C. J. (2013). Functional connectivity changes in multiple sclerosis patients: A graph analytical study of MEG resting state data. *Human Brain Mapping*, *34*(1), 52–61. <https://doi.org/10.1002/hbm.21424>, PubMed: 21954106
- Schreier, P. J., & Scharf, L. L. (2010). *Statistical signal processing of complex-valued data: The theory of improper and noncircular signals*. Cambridge, UK: Cambridge University Press. <https://doi.org/10.1017/CBO9780511815911>
- Shafiq, M. A., Tyler, L. K., Dixon, M., Taylor, J. R., Rowe, J. B., Cusack, R., ... Cam-CAN. (2014). The Cambridge Centre for Ageing and Neuroscience (Cam-CAN) study protocol: A cross-sectional, lifespan, multidisciplinary examination of healthy cognitive ageing. *BMC Neurology*, *14*(1), 204. <https://doi.org/10.1186/s12883-014-0204-1>, PubMed: 25412575
- Siebenhühner, F., Wang, S. H., Palva, J. M., & Palva, S. (2016). Cross-frequency synchronization connects networks of fast and slow oscillations during visual working memory maintenance. *Elife*, *5*, e13451. <https://doi.org/10.7554/eLife.13451>, PubMed: 27669146
- Smith, S. M., Miller, K. L., Salimi-Khorshidi, G., Webster, M., Beckmann, C. F., Nichols, T. E., ... Woolrich, M. W. (2011). Network modelling methods for fMRI. *NeuroImage*, *54*(2), 875–891. <https://doi.org/10.1016/j.neuroimage.2010.08.063>, PubMed: 20817103
- Srinivasan, R., Winter, W. R., Ding, J., & Nunez, P. L. (2007). EEG and MEG coherence: Measures of functional connectivity at distinct spatial scales of neocortical dynamics. *Journal of Neuroscience Methods*, *166*(1), 41–52. <https://doi.org/10.1016/j.jneumeth.2007.06.026>, PubMed: 17698205
- Stam, C. J. (2014). Modern network science of neurological disorders. *Nature Reviews Neuroscience*, *15*(10), 683–695. <https://doi.org/10.1038/nrn3801>, PubMed: 25186238
- Taulu, S., Simola, J., & Kajola, M. (2005). Applications of the signal space separation method. *IEEE Transactions on Signal Processing*, *53*(9), 3359–3372. <https://doi.org/10.1109/TSP.2005.853302>
- Taylor, J. R., Williams, N., Cusack, R., Auer, T., Shafiq, M. A., Dixon, M., ... Henson, R. N. (2017). The Cambridge Centre for Ageing and Neuroscience (Cam-CAN) data repository: Structural and functional MRI, MEG, and cognitive data from a cross-sectional adult lifespan sample. *NeuroImage*, *144*, 262–269. <https://doi.org/10.1016/j.neuroimage.2015.09.018>, PubMed: 26375206
- Ter Wal, M., Cardellicchio, P., LoRusso, G., Pelliccia, V., Avanzini, P., Orban, G. A., & Tiesinga, P. H. (2018). Characterization of network structure in stereoeeg data using consensus-based partial coherence. *NeuroImage*, *179*, 385–402. <https://doi.org/10.1016/j.neuroimage.2018.06.011>, PubMed: 29885486
- Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E. J., Yacoub, E., Uğurbil, K., ... WU-Minn HCP Consortium. (2013). The WU-Minn Human Connectome Project: An overview. *NeuroImage*, *80*, 62–79. <https://doi.org/10.1016/j.neuroimage.2013.05.041>, PubMed: 23684880
- Varoquaux, G., Gramfort, A., Poline, J.-B., & Thirion, B. (2010). Brain covariance selection: Better individual functional connectivity models using population prior. In *Advances in neural information processing systems* (pp. 2334–2342). <https://proceedings.neurips.cc/paper/2010/file/db576a7d2453575f29eab4bac787b919-Paper.pdf>
- Whittaker, J. (2009). *Graphical models in applied multivariate statistics*. Chichester, UK: Wiley Publishing.
- Wipf, D., & Nagarajan, S. (2009). A unified bayesian framework for MEG/EEG source imaging. *NeuroImage*, *44*(3), 947–966. <https://doi.org/10.1016/j.neuroimage.2008.02.059>, PubMed: 18602278
- Wodeyar, A. (2022). *AdaptiveGraphicalLassoForParCoh*, GitHub, <https://github.com/wodeyara/AdaptiveGraphicalLassoForParCoh>
- Wodeyar, A., Cassidy, J. M., Cramer, S. C., & Srinivasan, R. (2020). Damage to the structural connectome reflected in resting-state fMRI

- functional connectivity. *Network Neuroscience*, 4(4), 1197–1218. https://doi.org/10.1162/netn_a_00160, PubMed: 33409436
- Yeh, F.-C., Panesar, S., Fernandes, D., Meola, A., Yoshino, M., Fernandez-Miranda, J. C., ... Verstynen, T. (2018). Population-averaged atlas of the macroscale human structural connectome and its network topology. *NeuroImage*, 178, 57–68. <https://doi.org/10.1016/j.neuroimage.2018.05.027>, PubMed: 29758339
- Zerenner, T., Friederichs, P., Lehnertz, K., & Hense, A. (2014). A Gaussian graphical model approach to climate networks. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 24(2), 023103. <https://doi.org/10.1063/1.4870402>, PubMed: 24985417
- Zou, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American Statistical Association*, 101(476), 1418–1429. <https://doi.org/10.1198/016214506000000735>
- Zoubir, A. M., & Boashash, B. (1998). The bootstrap and its application in signal processing. *IEEE Signal Processing Magazine*, 15(1), 56–76. <https://doi.org/10.1109/79.647043>