

Learning Interpretable Models for Coupled Networks Under Domain Constraints

Hongyuan You, Sikun Lin, Ambuj K. Singh

Department of Computer Science
University of California, Santa Barbara
2120A Harold Frank Hall
Santa Barbara, California 93106
{hyou, sikun, ambuj}@ucsb.edu

Abstract

Modeling the behavior of coupled networks is challenging due to their intricate dynamics. For example in neuroscience, it is of critical importance to understand the relationship between the functional neural processes and anatomical connectivities. Modern neuroimaging techniques allow us to separately measure functional connectivity through fMRI imaging and the underlying white matter wiring through diffusion imaging. Previous studies have shown that structural edges in brain networks improve the inference of functional edges and vice versa. In this paper, we investigate the idea of coupled networks through an optimization framework by focusing on interactions between structural edges and functional edges of brain networks. We consider both types of edges as observed instances of random variables that represent different underlying network processes. The proposed framework does not depend on Gaussian assumptions and achieves a more robust performance on general data compared with existing approaches. To incorporate existing domain knowledge into such studies, we propose a novel formulation to place hard network constraints on the noise term while estimating interactions. This not only leads to a cleaner way of applying network constraints but also provides a more scalable solution when network connectivity is sparse. We validate our method on multishell diffusion and task-evoked fMRI datasets from the Human Connectome Project, leading to both important insights on structural backbones that support various types of task activities as well as general solutions to the study of coupled networks.

Introduction

Recently, there has been an effort to move research from the investigation of single networks to the more realistic scenario of multiple coupled networks. In this paper, we consider the case of pairs of networks, $(\mathcal{G}_1, \mathcal{G}_2)$, that are categorized into different modalities over a population. Such coupled network systems can be found in infrastructures of modern society (energy-communication), financial systems (ownership-trade), or even human brains (anatomical substrate-cortical activation). Our goal is to reconstruct one network from information on the other, and during such a process obtaining a concise interpretation of how one network affects the other.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

To achieve the above goal, we consider an edge-by-edge formulation. We treat one set of edges in \mathcal{G}_1 as predictors and the other set of edges in \mathcal{G}_2 as response variables in a multivariate linear regression model. Past research for the above problem relies on the restrictive Gaussian assumption, which simplifies the problem but is difficult to justify, especially in the domain of brain architectures. Adopting Gaussian assumption on non-Gaussian data can significantly prevent the detection of conditional dependencies and may lead to incorrectly inferred relationships among variables.

The learning of relationships between two different modalities can be difficult without sufficient data. As a result, in sparser data settings, the ability to specify constraints based on domain knowledge can be beneficial. For example, in the case of brain data, functional edges have mainly local influences, and structural edges are more responsible for long-distance influences (Rubinov and Sporns 2010; Batista-García-Ramó and Fernández-Verdecia 2018). We want preferences encoded in domain knowledge to guide the selection of partial correlations of unexplained noise terms in the constructed model.

Based on the above motivations, we propose a flexible and efficient framework CC-MRCE (Convex-set Constrained Multivariate Regression with Covariance Estimation) that simultaneously learns both regression coefficients between two coupled networks and the correlation structure of noise terms. In a departure from existing methods, our framework encodes domain knowledge as a set of convex constraints and adopts a pseudolikelihood-based neighborhood-selection objective in partial correlation estimation, which has been shown to be more robust to non-Gaussian data. Because of the CC-MRCE objective's bi-convex nature, we alternately solve a regression sub-problem and a constrained partial correlation sub-problem until convergence. The latter sub-problem requires feasible solutions under given domain constraints that we render tractable via a modified two-stage proximal gradient descent method.

We illustrate the use of our method in the context of the human brain. Brain data presents one of the greatest technical challenges in analysis and modeling due to a network-based characterization (Bassett and Sporns 2017; Bassett, Zurn, and Gold 2018), non-Gaussian nature of data, high dimensionality, a small number of samples, and the need to incorporate domain knowledge. We apply the pro-

posed framework on the Human Connectome Project (HCP) dataset (Van Essen et al. 2013), where two coupled networks are constructed from fMRI scans (representing cortical activation) and diffusion scans (representing the anatomical substrate). We successfully predict a brain functional network from the given structural network; our method outperforms previous state-of-art methods, and our obtained models are easier to interpret. We investigate the structure-function coupling for seven different tasks. Our findings agree with the nature of fMRI task and brain region functions in existing literature, thus validating our model’s ability to discover meaningful couplings.

Our main contributions are as follows:

- We propose a regularized multiple regression approach that adapts to non-Gaussian data. (Sections 2.1, 2.2)
- We incorporate prior domain knowledge to model estimation by formulating constraints into an optimization problem. (Section 2.3)
- We develop a fast method based on nested FISTA for solving the proposed optimization problem. (Section 3)
- We show the effectiveness of our model on HCP brain data using quantitative comparisons with existing approaches as well as a qualitative analysis. (Section 4.2)

Problem Formulation

In this section, we first introduce existing works on multiple regression under Gaussian assumptions and then motivate our approach under non-Gaussian settings and domain constraints.

Multiple-output Regression Problem

Let \mathcal{D} be an n -subject sample set in which all subjects share the same coupling $(\mathcal{G}_1, \mathcal{G}_2)$ but have different edge values. For subject i in \mathcal{D} , let $\mathbf{x}^{(i)} = (x_1^{(i)}, \dots, x_p^{(i)})$ be p -dimensional inputs that represent edge values in the first modality network \mathcal{G}_1 , and $\mathbf{y}^{(i)} = (y_1^{(i)}, \dots, y_p^{(i)})$ be p -dimensional outputs¹ that stand for edge values in the second modality network \mathcal{G}_2 . We assume that the inputs x_i and outputs y_i are correlated through a multivariate linear regression model:

$$\mathbf{y}^{(i)} = \mathbf{x}^{(i)} \mathbf{B} + \boldsymbol{\epsilon}^{(i)}, \quad \text{for } i = 1, \dots, n \quad (1)$$

where \mathbf{B} is the $p \times p$ regression coefficient matrix and its element β_{jk} is the regression coefficient that measures the cross-modality impact of edge x_j to edge y_k , and $\boldsymbol{\epsilon}^{(i)}$ is the noise vector of subject i . The model can be expressed in the matrix form:

$$\mathbf{Y} = \mathbf{X} \mathbf{B} + \mathbf{E} \quad (2)$$

where row i of $\mathbf{X} \in \mathbb{R}^{n \times p}$ and $\mathbf{Y} \in \mathbb{R}^{n \times p}$ are the structural and functional edge vectors $\mathbf{x}^{(i)}$ and $\mathbf{y}^{(i)}$ of subject i .

A straightforward approach to estimating \mathbf{B} is to solve p separate regression problems, assuming noise terms are independent and uncorrelated. Recently, advanced methods have been proposed to exploit the correlation in noise terms

¹In general, models do not require the same dimensions for inputs and outputs. We use the equality setting only for simplicity.

to improve the modeling. They accomplish the goal by introducing an assumption that noise terms $\boldsymbol{\epsilon}^{(1)}, \dots, \boldsymbol{\epsilon}^{(n)}$ are all *i.i.d.* Gaussian $\mathcal{N}(\mathbf{0}, \boldsymbol{\Omega}^{-1})$ and then simultaneously estimating regression coefficients \mathbf{B} and inverse covariance matrix $\boldsymbol{\Omega}$ of the noise terms. Two popular methods along this direction are MRCE (Rothman, Levina, and Zhu 2010) and CGGM (Sohn and Kim 2012; Wytock and Kolter 2013; Yuan and Zhang 2014). The MRCE method considers the conditional distribution $\mathbf{Y}|\mathbf{X} \sim \mathcal{N}(\mathbf{X}\mathbf{B}, \boldsymbol{\Omega}^{-1})$ and estimates both \mathbf{B} and $\boldsymbol{\Omega}$ by alternately minimizing the negative conditional Gaussian likelihood, with the ℓ_1 lasso penalty applied on the entries of \mathbf{B} and $\boldsymbol{\Omega}$. The other method, CGGM, further assumes that \mathbf{X} and \mathbf{Y} are jointly Gaussian. Under such formulation, the conditional distribution of $\mathbf{Y}|\mathbf{X}$ is given by $\mathcal{N}(-\mathbf{X}\boldsymbol{\Omega}_{XY}\boldsymbol{\Omega}^{-1}, \boldsymbol{\Omega}^{-1})$, which reparameterizes the regression coefficient \mathbf{B} as $-\boldsymbol{\Omega}_{XY}\boldsymbol{\Omega}^{-1}$. Compared with MRCE, the objective of CGGM is based on the negative conditional Gaussian likelihood as well, but is jointly convex for $\boldsymbol{\Omega}_{XY}$ and $\boldsymbol{\Omega}$, and therefore more friendly to computation.

Relaxing Gaussian Assumptions

Although MRCE and CGGM have received significant attention in solving multi-output regression problems, one drawback of these two approaches is the Gaussian assumption, especially for applications to brain data (Freyer et al. 2009; Hlinka et al. 2011; Eklund, Nichols, and Knutsson 2016). Recall that MRCE assumes the Gaussian noise, and CGGM further assumes joint Gaussian distribution over both inputs and outputs. We tested whether the HCP structural and functional data is Gaussian with a significance level of 0.05. The test rejects the Gaussian null hypothesis for 97.5% of structural edges and 36.3% of functional ones. Since our sample size is small, false negatives are more likely to occur (Columb and Atkinson 2016), namely failing to reject the Gaussian hypothesis when the underlying data is non-Gaussian. Therefore, the proportion of non-Gaussian data in brain networks is expected to be even higher. Thus, relying on Gaussian assumptions is likely to affect the constructed models negatively.

To avoid a Gaussian assumption, we propose a pseudolikelihood approach for learning multi-output regression models by optimizing the following objective function:

$$\min_{\{\mathbf{B}_k\}, \{\omega_{jk}\}} \left[-n \sum_{j=1}^p \log \omega_{jj} + \frac{1}{2} \sum_{j=1}^p \sum_{i=1}^n \left(\omega_{jj} (\mathbf{y}_j^{(i)} - \mathbf{x}^{(i)} \mathbf{B}_j) + \sum_{k \neq j} \omega_{jk} (\mathbf{y}_k^{(i)} - \mathbf{x}^{(i)} \mathbf{B}_k) \right)^2 + \lambda_1 \sum_{j < k} |\omega_{jk}| + \lambda_2 \sum_{j < k} |\beta_{jk}| \right]$$

or in a neat matrix notation:

$$\min_{\mathbf{B}, \boldsymbol{\Omega}} -n \log |\boldsymbol{\Omega}_D| + \frac{1}{2} \text{tr} \left((\mathbf{Y} - \mathbf{X} \mathbf{B})^T (\mathbf{Y} - \mathbf{X} \mathbf{B}) \boldsymbol{\Omega}^2 \right) + \lambda_1 \|\mathbf{B}\|_1 + \lambda_2 \|\boldsymbol{\Omega}_X\|_1 \quad (3)$$

where $\boldsymbol{\Omega} = \{\omega_{jk}\}$ denotes the inverse covariance matrix, $\mathbf{B} = \{\beta_{jk}\}$ denotes the coefficient matrix, and $\boldsymbol{\Omega}_D$ and $\boldsymbol{\Omega}_X$ denote the diagonal and off-diagonal parts of $\boldsymbol{\Omega}$. The proposed objective can be considered as a reparameterization of the Gaussian likelihood with $\boldsymbol{\Omega}^2$ and an approximation to the

log-determinant term. It has been proven that under mild singularity conditions, such reparameterization can guarantee estimation consistency for distributions with sub-Gaussian tails (Peng et al. 2009; Khare, Oh, and Rajaratnam 2015).

Next, we develop an optimization algorithm to minimize the objective. The objective function itself is not jointly convex for both variables \mathbf{B} and \mathbf{X} , but remains convex with respect to each of them while keeping the other fixed. Therefore, we adopt the alternating minimization idea. In the t -th iteration, we first fix \mathbf{B} as the estimated $\hat{\mathbf{B}}^{(t-1)}$ from the previous $(t-1)$ -th iteration, and calculate the empirical covariance matrix \mathbf{S} of noise terms:

$$\mathbf{S}^{(t-1)} = \frac{1}{n}(\mathbf{Y} - \mathbf{X}\mathbf{B}^{(t-1)})^T(\mathbf{Y} - \mathbf{X}\mathbf{B}^{(t-1)}) \quad (4)$$

Next, we estimate the inverse covariance matrix: $\Omega^{(t)}$ with the given $\mathbf{S}^{(t-1)}$ as a constant:

$$\Omega^{(t)} = \arg \min_{\Omega} -\log |\Omega_D| + \frac{1}{2} \text{tr}(\mathbf{S}^{(t-1)}\Omega^2) + \lambda_2 \|\Omega_X\|_1 \quad (5)$$

Observe that the above subproblem follows CONCORD's original form, which is more robust to heavy-tailed data (Khare, Oh, and Rajaratnam 2015; Koanantakool et al. 2018) than the conventional Gaussian likelihood approach and can be efficiently solved using proximal gradient methods with a convergence rate of $O(1/t^2)$ (Oh et al. 2014).

Lastly, we keep Ω fixed at $\Omega^{(t)}$ and optimize the regression coefficients \mathbf{B} :

$$\mathbf{B}^{(t)} = \arg \min_{\mathbf{B}} \frac{1}{2} \text{tr}((\mathbf{Y} - \mathbf{X}\mathbf{B})^T(\mathbf{Y} - \mathbf{X}\mathbf{B})\Omega^{(t)}) + \lambda_1 \|\mathbf{B}\|_1 \quad (6)$$

Note that subproblem (6) is convex when $\Omega^{(t)}$ is positive semi-definite. We present the above regression-based approach as CONCORD-MRCE in **Algorithm 1** (pseudocode, see appendix).

Imposing Domain Constraints

Due to limited sample size of real-world datasets and their high dimensionality, incorporating accurate domain constraints can reduce the search space and avoid over-fitting.

Under a linear mapping assumption, the partial correlation of response variables arises only from correlations in the noise terms. Therefore, Ω not only represents the inverse covariance of noise terms, but also equals the conditional inverse covariance of $\mathbf{Y}|\mathbf{X}$. The nonzero entries of Ω encode direct relationships among the target modality outputs \mathbf{Y} that cannot be explained by weighted inputs $\mathbf{X}\mathbf{B}$ of the source modality. It will be beneficial if the zero-vs-nonzero structure of Ω is partially given by domain experts and used as hard constraints in model estimation.

More formally, let M be a binary matrix that has the same dimensions as Ω . We can define a convex matrix set \mathbb{S}_M , containing all matrices that share the same set of zero entries with M . We can then improve the previous regression-based approach to estimate Ω under the domain constraint that takes the form of $\Omega \in \mathbb{S}_M$, written in an equivalent unconstrained convex form:

$$\hat{\Omega} = \arg \min_{\mathbf{B}, \Omega} -\frac{n}{2} \log |\Omega_D^2| + \frac{1}{2} \text{tr}((\mathbf{Y} - \mathbf{X}\mathbf{B})^T(\mathbf{Y} - \mathbf{X}\mathbf{B})\Omega^2) + \lambda_1 \|\mathbf{B}\|_1 + \lambda_2 \|\Omega_X\|_1 + \mathbb{I}\{\Omega \in \mathbb{S}_M\} \quad (7)$$

where $\mathbb{I}\{\Omega \in \mathbb{S}_M\}$ is an indicator function. This formulation can be extended to $\Omega \in C$ whenever C is a closed convex set of positive definite matrices.

Solution

In this section, we show how to adapt the previous solution when the inverse covariance Ω is constrained during estimation. Notice that the ideas of **Algorithm 1** can be mostly used to solve Eq (7), except for the Ω -update step, which is now affected by the added constraints. The new Ω -update step needs to solve the following sub-problem:

$$\Omega^{(t)} = \arg \min_{\Omega} -\log |\Omega_D^2| + \text{tr}(\mathbf{S}^{(t-1)}\Omega^2) + \lambda_2 \|\Omega_X\|_1 + \mathbb{I}\{\Omega \in \mathbb{S}_M\} \quad (8)$$

We follow the FISTA (Fast Iterative Soft-Thresholding Algorithm (Beck and Teboulle 2009)) approach that is used in CONCORD (Oh et al. 2014). This method utilizes an accelerated gradient algorithm using soft-thresholding as its proximal operator for the L_1 norm and achieves a fast $O(1/t^2)$ convergence rate. Previous work (Oh et al. 2014) has also applied FISTA for partial correlation estimation and proved its efficiency. To adapt our constrained problem into the FISTA framework, we split our objective function (8) into a smooth part and a non-smooth part:

$$\begin{aligned} h_1(\Omega) &= -\log |\Omega_D^2| + \text{tr}(\mathbf{S}\Omega^2) \\ h_2(\Omega) &= \lambda_2 \|\Omega_X\|_1 + \mathbb{I}\{\Omega \in \mathbb{S}_M\} \end{aligned}$$

For any symmetric matrix Ω , the gradient of the smooth function can be easily calculated as: $\nabla h_1(\Omega) = -2\Omega_D^{-1} + 2\Omega\mathbf{S}$. With this formulation, we now adapt the FISTA iterative scheme to solve our network-constrained problem (8):

$$\alpha_{t+1} = (1 + \sqrt{1 + 4\alpha_t^2})/2 \quad (9)$$

$$\Theta^{(t+1)} = \Omega^{(t)} + \frac{\alpha_t - 1}{\alpha_{t+1}}(\Omega^{(t)} - \Omega^{(t-1)}) \quad (10)$$

$$\Omega^{(t+1)} = \text{prox}_{\gamma h_2}[\Theta^{(t+1)} - (n\tau_t/2)\nabla h_1(\Theta^{(t+1)})] \quad (11)$$

where τ_t is the step length and t denotes the iteration number. γ is a trade-off parameter that controls the extent to which the proximal operator maps points towards the minimum of $h_2(\Omega)$, with larger values of γ associated with larger movement near the minimum.

In these iterative steps, Θ_{t+1} is an expected position, updated purely by momentum. Within each loop, the algorithm first takes a gradient step of the estimated future position (Eq.10) and then applies the proximal mapping of a closed convex function $h_2(\Omega)$.

In contrast to the standard FISTA approach, the composite function $h_2(\Omega)$ consists of a sparsity penalty and a network-constrained indicator function. More specifically, we can write down the explicit form of Eq. (11) according to the proximal operator definition:

$$\begin{aligned} \hat{\Omega} &= \hat{\Omega}_X + \mathbf{A}_D \\ \hat{\Omega}_X &= \text{prox}_{\gamma h_2}(\mathbf{A}_X) \\ &= \arg \min_{\Omega_X \in C} \frac{1}{2\gamma} \|\Omega_X - \mathbf{A}_X\|_F^2 + \lambda_2 \|\Omega_X\|_1 \quad (12) \end{aligned}$$

where $\mathbf{A} = \Theta^{(t+1)} - (n\tau_t/2)\nabla h(\Theta^{(t+1)})$. Instead of directly solving the original problem (12), we consider its dual problem as follows. Let matrix \mathbf{H} be the dual variable of matrix Ω . We have:

$$\begin{aligned} & \min_{\Omega_X \in C} \left(\frac{1}{2\gamma} \|\Omega_X - \mathbf{A}_X\|_F^2 + \lambda_2 \max_{\|\mathbf{H}_X\|_\infty \leq 1} \text{vec}(\mathbf{H}_X)^T \text{vec}(\Omega_X) \right) \\ &= \max_{\|\mathbf{H}_X\|_\infty \leq 1} \min_{\Omega_X \in C} \frac{1}{2\gamma} \left(\|\Omega_X - (\mathbf{A}_X - \gamma\lambda_2 \mathbf{H}_X)\|_F^2 \right. \\ & \quad \left. - \|\mathbf{A}_X - \gamma\lambda_2 \mathbf{H}_X\|_F^2 + \|\mathbf{A}_X\|_2^2 \right) \end{aligned} \quad (13)$$

where \mathbf{A}_D and \mathbf{A}_X denote the diagonal and off-diagonal part of \mathbf{A} . Since the initial objective function above is convex in Ω_X and concave in \mathbf{H}_X , we exchange the order of the minimum and maximum operator in which the inner minimization problem has an obvious solution through orthogonal projection theorem (Rockafellar 1970), written as

$$\Omega_X = \mathbb{P}_C(\mathbf{A}_X - \gamma\lambda_2 \mathbf{H}_X) \quad (14)$$

where \mathbb{P}_C is defined as an projection operator: $\mathbb{P}_C(\Gamma) = \text{argmin}_{\mathbf{R} \in C} \|\mathbf{R} - \Gamma\|_F^2$ and its orthogonal projection operator \mathbb{P}_{C^\perp} is defined as $I - \mathbb{P}_C$. In the special case that $C = \mathbb{S}_M$, projection $\mathbb{P}_C(\Gamma)$ is equivalent to removing invalid nonzero entries of the input matrix Γ .

Inserting the optimal Ω_X back into objective (13), we now obtain the final dual form of the problem (12):

$$\begin{aligned} \hat{\mathbf{H}}_X &= \arg \min_{\|\mathbf{H}_X\|_\infty \leq 1} \|\mathbf{A}_X - \gamma\lambda_2 \mathbf{H}_X\|_2^2 \\ & \quad - \|\mathbb{P}_{C^\perp}(\mathbf{A}_X - \gamma\lambda_2 \mathbf{H}_X)\|_2^2 \end{aligned} \quad (15)$$

where any solution $\hat{\mathbf{H}}_X$ to the dual problem corresponds to a primal solution through Eq.(14). Since the dual objective is continuously differentiable and constraints on l_∞ -norm are convex, we can again efficiently solve it with additional inner FISTA iterations, which is to minimize an equivalent composite objective $\min g_1(\mathbf{H}_X) + g_2(\mathbf{H}_X)$, where $g_1(\mathbf{H}_X)$ is smooth and $g_2(\mathbf{H}_X)$ is non-smooth:

$$\begin{aligned} g_1(\mathbf{H}_X) &= \|\mathbf{A}_X - \gamma\lambda_2 \mathbf{H}_X\|_2^2 - \|\mathbb{P}_{C^\perp}(\mathbf{A}_X - \gamma\lambda_2 \mathbf{H}_X)\|_2^2 \\ g_2(\mathbf{H}_X) &= I_{\{\|\mathbf{H}_X\|_\infty \leq 1\}}(\mathbf{H}_X). \end{aligned}$$

To adapt FISTA to the problem (15), the only thing left is to obtain the gradient of smooth function $g_1(\mathbf{H}_X)$, for which we need the lemma below:

Lemma 1. (Moreau 1965) *If g is a closed proper convex function, and for any positive t , define a proximal operator $g_t(\mathbf{x}) := \inf_{\mathbf{u}} [g(\mathbf{u}) + \frac{1}{2t}\|\mathbf{u} - \mathbf{x}\|^2]$, then its infimum is attained at the unique point $\text{prox}_t(g)(\mathbf{x})$. Further, g_t is continuously differentiable on \mathbb{E} with a $1/t$ -Lipschitz gradient given by $\nabla g_t(\mathbf{x}) = (\mathbf{x} - \text{prox}_t(g)(\mathbf{x}))/t$.*

According to Lemma 1, we simply plug $g(\mathbf{u}) = \delta(\mathbf{u} \in C)$, and obtain that $\text{prox}_t(g)(\mathbf{x}) = \mathbb{P}_C(\mathbf{x})$. Therefore the gradient ∇g_1 is calculated as:

$$\nabla g_1(\mathbf{H}_X) = -2\gamma\lambda_2 \cdot \mathbb{P}_C(\mathbf{A}_X - \gamma\lambda_2 \mathbf{H}_X). \quad (16)$$

and the proximal mapping of function $g_2(\Omega)$ becomes a projection of \mathbf{H}_X into the L_∞ -ball:

$$(\text{prox}_{g_2}(\mathbf{H}))_{ij} = \text{sign}(\mathbf{H}_X) \min\{|\mathbf{H}_X|, \mathbf{1}_X\} \quad (17)$$

This completes the modified Ω -update step in the constrained setting (referred to as Constrained-CONCORD). Its corresponding pseudocode is presented in **Algorithm 3** (pseudocode, see appendix). The overall framework of simultaneously estimating Ω and \mathbf{B} is presented in **Algorithm 2** (pseudocode, see appendix), which we name as CC-MRCE.

Experiments

We present two sets of experiments. First, to understand our method's strengths and limitations, we utilize simulated datasets that allow us to inspect both reconstruction performance and model selection performance. We compare the proposed method CC-MRCE with other baselines when the underlying data distribution does not follow the Gaussian assumption. We show that both the non-Gaussian assumption and the network constraints contribute to the improvement of performance. Second, we conduct experiments on the Human Connectome Project (HCP) data (Van Essen et al. 2013). Our model offers a quantitative advantage over baseline methods for predicting functional networks from structural ones; at the same time, our results agree with existing neuroscience literature.

Application to Simulated Data

Data Generation. Using a similar approach to existing works (Peng et al. 2009; Khare, Oh, and Rajaratnam 2015; Rothman, Levina, and Zhu 2010), we generate our simulated dataset by first synthesizing two key model parameter matrices Ω_0 and \mathbf{B}_0 , and then construct input, output, and noise terms (*i.e.* $\mathbf{x}^{(i)}$, $\mathbf{y}^{(i)}$ and $\epsilon^{(i)}$'s).

Note that every positive-definite matrix has a Cholesky decomposition that takes the form of $\mathbf{L}\mathbf{L}^T$, where \mathbf{L} is a lower triangular matrix \mathbf{L} , and if \mathbf{L} is sparse enough then $\mathbf{L}\mathbf{L}^T$ is sparse as well. Therefore, we first sample a sparse lower triangular \mathbf{L} with real and positive diagonal entries, and then generate our inverse covariance matrix with $\Omega_0 = \mathbf{L}\mathbf{L}^T$. The generated $p \times p$ positive definite matrix Ω_0 has 10% nonzeros entries and a condition number of 4.3. To demonstrate the robustness of proposed method CC-MRCE on non-Gaussian data, we sample the noise terms $\{\epsilon^{(i)}\}_{i=1}^n$ according to a multivariate t -distribution with zero mean and covariance matrix $\Sigma = \Omega_0^{-1}$.

Next, we generate a sparse coefficient matrix \mathbf{B}_0 using the matrix element-wise product trick, $\mathbf{B}_0 = \mathbf{W} \circ \mathbf{K} \circ \mathbf{Q}$. In this construction approach, \mathbf{W} has entries with independent draws from standard normal distribution $\mathcal{N}(0, 1)$. Each entry in \mathbf{K} is drawn independently from a Bernoulli distribution that takes value 1 with probability s_1 . \mathbf{Q} has rows that are either all one or all zero, which are determined by independent Bernoulli draws with success probability s_2 . Generating the sparse \mathbf{B}_0 in this manner, we not only control its sparsity level, but also forcibly make $(1 - s_2)p$ predictors to be irrelevant for p responses, and guarantee that each relevant predictor is associated with $s_1 p$ response variables.

In the following experiments, the probabilities s_1 and s_2 are chosen to be 0.15 and 0.8, the sample size n is fixed at 50, and the input and output dimensions p and q are both

Model	Reconstruction MSE percentage (100%)	Pearson's r -score	min p-value	max p-value	relative AUC w.r.t. Ω	relative AUC w.r.t. B
CC-MRCE (unconstrained)	50.88	0.709	5.146E-09	0.101	0.520	0.536
CC-MRCE (SNR:1.0)	43.24	0.763	2.265E-11	0.083	0.855	0.640
CC-MRCE (SNR:2.0)	43.18	0.764	2.109E-11	0.073	0.925	0.662
CC-MRCE (perfect)	42.98	0.764	2.692E-11	0.064	1.000	0.671
MRCE	60.22	0.653	8.686E-09	0.174	0.375	0.540
CGGM	76.21	0.552	2.141E-07	0.965	0.330	0.195

Table 1: Reconstruction performance and model selection performance of models on simulated dataset. CC-MRCE variations uniformly outperform MRCE and CGGM. CC-MRCE variations with more strict constraints perform better.

set to 20. The input $x^{(i)}$'s are sampled from a multivariate normal distribution $\mathcal{N}(0, \Sigma_X)$ where $(\Sigma_X)_{jk} = 0.5^{|j-k|}$, following previous works (Yuan and Lin 2007; Peng et al. 2010). The output $y^{(i)}$'s are calculated as the linear model assumption in Eq.(1). We replicate the above process for independently generating a validation dataset of the same sample size. All penalty parameters are selected simultaneously and tuned according to the validation error.

Methods. In this heavy-tailed setting, we compare the performance of CC-MRCE to CGGM and MRCE, which are developed under Gaussian settings. The CGGM implementation that we used in this experiment is provided by (McCarter and Kim 2016) and has been optimized for large-scale problems and limited memory. The MRCE implementation is provided by (Rothman, Levina, and Zhu 2010). Both CGGM and MRCE do not adapt to the network constraints we impose here.

In order to inspect the effectiveness of network constraints, we apply multiple variations of constraint sets to the proposed CC-MRCE method, designated as CC-MRCE (unconstrained), CC-MRCE (SNR: 2.0), CC-MRCE (SNR: 1.0) and CC-MRCE (perfect). *Network constraints* in each variation are defined as follows. For CC-MRCE (perfect), we choose $\mathbb{S}_E = \{\Omega : \Omega(j, k) = 0 \text{ if } \Omega_0(j, k) = 0\}$, which forces selected nonzeros in solution Ω to completely fall into ground-truth nonzeros. For CC-MRCE (SNR: 2.0) and CC-MRCE (SNR: 1.0), we loosen the feasible set by adding $50\% \|\Omega_0\|_0$ and $100\% \|\Omega_0\|_0$ spurious nonzero positions, which are randomly sampled from positions of zero entries in Ω_0 . For CC-MRCE (unconstrained), we remove all constraints so that the method only relies on regularized multi-regression.

Performance Evaluation. We evaluate the reconstruction performance of models using the conventional MSE error (in percentage). Correlation coefficients between predicted and ground-truth outputs are also provided along with their corresponding p-values. In addition, we use the relative area under the curve (AUC) of receiver operating characteristic (ROC) curves (Fawcett 2006; Friedman, Hastie, and Tibshirani 2010), with regards to Ω and B , as key measures to compare model selection performance of all these methods. The AUC of a perfect ROC curve, which would be 1, indicates an ideal recovery of ground truth zero-vs-nonzero structure in Ω (or B). However, models with large false-positive rates (FPR) are barely meaningful in real scenarios. To focus on the initial portion of ROC curves, we control the




Figure 1: (a) ROC curve for Ω estimation, (b) ROC curve for B estimation. Benefit from domain constraints, CC-MRCE obtains better ROC curves when uncovering nonzero entries of B and Ω .

FPRs simultaneously to be smaller than 0.2 for both Ω and B estimation. Thus, the maximum AUC value that a model can reach is just 0.2. For ease of comparison, we provide relative AUC values, divided by 0.2 to normalize to 1. For each method, we run the algorithm with at least 25 appropriate parameter pairs (λ_1, λ_2) to get its ROC curve. Recall that all methods in this section are required to estimate 800 parameters given $n = 50$ samples.

Test Results. Table 1 displays the test results of six different settings in all measures mentioned above. As can be seen, all four CC-MRCE variations (with different network constraints) obtain significantly smaller reconstruction MSE percentages, higher correlation coefficients, and smaller p-values. Note that CGGM behaves the worst in all measures since it is deeply rooted in the Gaussian setting and is consequently misled by these assumptions. We also see that the performance of CC-MRCE gradually improves when the applied network constraints are more informative.

We also plot ROC curves for Ω and B estimation in Figure 1.(a) and Figure 1.(b), respectively. It is clear that CC-MRCE performs better than MRCE and CGGM, across different choices of network constraints. For Ω estimation, as expected, CC-MRCE with more strict constraints has steeper curves, suggesting that it recovers mostly correct partial relationships between variables with very few spurious connections, and therefore achieves higher AUC scores. CC-MRCE (perfect) behaves perfectly for Ω -ROC, by its definition. For the estimation of matrix B , a similar phe-

nomena demonstrates that network constraints improve the learning of regression coefficients and lead to a better reconstruction performance. Without imposing network constraints, unconstrained formulation of CC-MRCE is likely to generate a biased estimate of $\hat{\Omega}$ on small datasets and can not recover ground truth features in B .

Application to Human Connectome Data

Problem Formulation. Many works in literature report on coupling between brain structural connectivities (SCs) and functional connectivities (FCs) for both resting state (Hagmann et al. 2008; Hermundstad et al. 2013; Honey et al. 2009, 2007) and task-evoked states (Hermundstad et al. 2013; Raichle 2015; Cole et al. 2014; Davison et al. 2015). One such recent work (Becker et al. 2018) maps SC to resting-state FC by aligning both the eigenvalues and eigenvectors of a subject’s SC and FC matrices, and evaluates the mapping by reconstructing resting FCs from SCs.

Inspired by the domain observation (Batista-García-Ramó and Fernández-Verdecia 2018) that two functional connections sharing the same node are more likely to form a meaningful pathway or functional activation pattern, we constrain the partial correlations between non neighboring edges to be zero, i.e. $\Omega_{jk} = 0$ if e_j and e_k are not incident edges in the fMRI-constructed network.

We conducted experiments by reconstructing task FCs from SCs of 51 subjects with their fMRI data from HCP dataset under seven task states. We used the parcellation scheme in (Hagmann et al. 2008) with a spatial scale of 33, resulting in 83 brain regions and 3403 possible edges. Columns in predictor matrix, Equation (2), $\mathbf{X} \in \mathbb{R}^{51 \times 3403}$ represent SCs, whose entries are numbers of white matter streamlines intersecting pairs of brain regions, and columns in response matrix $\mathbf{Y} \in \mathbb{R}^{51 \times 3403}$ represent FCs, whose entries are functional correlations between cortical activities of brain regions. B denotes the mapping (or coupling) between SCs and FCs, and E denotes the part of FCs that cannot be explained by SCs. We also normalized SC values to (0, 1], a range comparable to FC. We ran a 10-fold cross-validation of our model for each task, splitting data in a 9-1 train-validation ratio (46-5 split for the 51 subjects). We selected the optimal hyperparameters λ_1 and λ_2 by a 5×5 grid search in the log-scale between $10^{-1.6}$ to $10^{-0.4}$, keeping the models with smallest Mean Squared Error (MSE) percentage on the validation sets, averaged across 10 folds. In our case, both λ_1 and λ_2 have optimal values around 0.1 across tasks. Aside from MSE, we also tested the Pearson correlation coefficient between predicted FCs and ground truth FCs, (referred to as Pearson’s r -score, listed in Table 2) and minimum, maximum p -values. The reconstruction MSE percentage is below 1% for the training data and around 8% for the validation data. Strong and significant positive correlations are shown for both training (r -score around 0.6 to 0.8) and validation (r -score around 0.5) data. These results indicate our model’s effectiveness in FC reconstruction by exploiting cross-modal coupling between SC-FC and domain prior knowledge on FC-FC relationships.

Performance Evaluation. Regarding the ability of finding accurate mappings between SCs and FCs, we com-

Tasks		Reconstruction MSE percentage (100%)	Pearson’s r -score
EMOTION	1	89.78 \pm 21.95	-0.0309 \pm 0.0285
	2	81.57 \pm 2.38	0.0777 \pm 0.0070
	3	61.54 \pm 33.84	0.4228 \pm 0.0349
	4	17.50 \pm 1.84	-0.0014 \pm 0.0086
	5	8.84 \pm 0.84	0.4575 \pm 0.0402
LANGUAGE	1	41.38 \pm 7.10	0.0270 \pm 0.0448
	2	72.68 \pm 6.33	0.0815 \pm 0.0081
	3	54.23 \pm 30.18	0.4764 \pm 0.0383
	4	35.02 \pm 58.49	0.0020 \pm 0.0052
	5	7.95 \pm 0.87	0.4988 \pm 0.0205
MOTOR	1	117.85 \pm 27.32	-0.0016 \pm 0.0456
	2	77.30 \pm 4.24	0.0782 \pm 0.0110
	3	57.26 \pm 29.63	0.4156 \pm 0.0548
	4	21.02 \pm 7.75	0.0023 \pm 0.0090
	5	7.80 \pm 0.66	0.4807 \pm 0.0480
GAMBLING	1	108.99 \pm 32.83	-0.0211 \pm 0.0795
	2	79.14 \pm 3.65	0.0804 \pm 0.0071
	3	54.73 \pm 30.70	0.4781 \pm 0.0380
	4	22.63 \pm 13.15	0.0033 \pm 0.0072
	5	7.86 \pm 1.72	0.5014 \pm 0.0301
SOCIAL	1	112.82 \pm 36.07	-0.0064 \pm 0.076
	2	79.42 \pm 3.72	0.0772 \pm 0.0088
	3	47.89 \pm 28.09	0.4912 \pm 0.0353
	4	18.68 \pm 3.74	0.0025 \pm 0.0071
	5	6.81 \pm 0.95	0.5578 \pm 0.0404
RELATIONAL	1	147.61 \pm 49.24	-0.0265 \pm 0.0754
	2	81.11 \pm 2.98	0.0706 \pm 0.0081
	3	54.86 \pm 30.68	0.4758 \pm 0.0412
	4	31.77 \pm 18.30	-0.0019 \pm 0.012
	5	8.43 \pm 1.09	0.4858 \pm 0.0460
WM (Working Memory)	1	81.46 \pm 29.72	-0.0447 \pm 0.0566
	2	77.99 \pm 2.89	0.0799 \pm 0.0109
	3	56.25 \pm 32.85	0.4767 \pm 0.0579
	4	103.96 \pm 111.99	-0.0041 \pm 0.0089
	5	7.69 \pm 1.76	0.4968 \pm 0.0543

Table 2: Functional connectivity reconstruction performance of seven tasks with different models. Models are numbered as follows. 1: CGGM; 2: VAE; 3: Spectral Mapping; 4: Random B ; 5: CC-MRCE (Ours).

pared our model with the optimization approach CGGM, a deep learning approach Variational AutoEncoder (VAE) (Kingma and Welling 2013), and the Spectral Mapping method (Becker et al. 2018). For CGGM, the time required to run a 10-fold cross-validation with one set of hyperparameters on a single task ranges from one day to a week. So we ran three sets of hyperparameters on EMOTION and LANGUAGE, selected the best parameter pair, and fixed it for the rest of the five tasks. In particular, we chose penalty terms on B and E to be both 0.01. For VAE, both encoder and decoder consist of two fully connected layers, with latent variable dimensions being two. We use MSE as the training loss. In our experiments, increasing the number of layers or latent dimensions of VAE did not improve the final performance. For the Spectral Mapping method, we follow the setup of the original paper, setting maximum path length k to seven. The 10-fold cross validation results of CGGM, VAE and Spectral

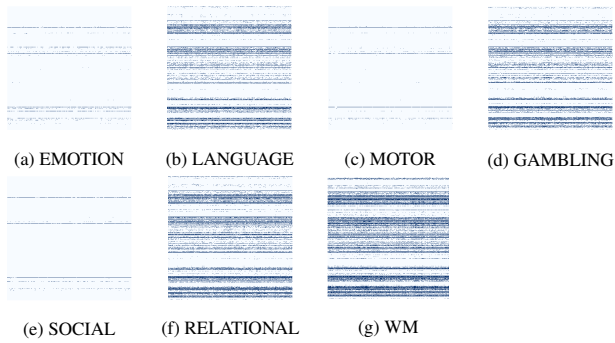


Figure 2: Visualizations of the B matrix for seven tasks. For each task, each fold in the 10-fold cross validation may lead to different models. Here, we only show those entries that are nonzero more than five times.

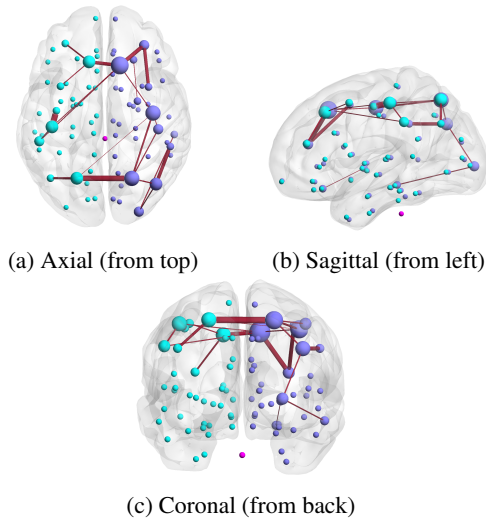


Figure 3: Visualization of the edges contributing to *all* seven tasks. Node size denotes the degree and edge width denotes its importance, as in the mapping B .

Mapping method are reported in Table 2.

The assumption of Gaussian noise weakens CGGM’s performance on the HCP dataset: it has unstable and large average MSE percentages across tasks, and the correlations between predictions and ground truth are very small and even negative, which are also not statistically significant with regards to p -values. On the other hand, VAE models have stable MSE percentage (around 80%) and average Pearson’s r -scores with small standard deviations across all tasks. The correlations of VAE models are weak (all around 0.08), yet statistically significant with p -values constantly smaller than 0.0025 for all tasks. This shows VAE learns a slightly meaningful mapping, but with such a small sample size, deep learning models are unlikely to perform well. Lastly, the Spectral Mapping method is designed for maximizing the correlation between fMRI prediction and ground truth for brain data, so it performs well as for correlation, however the prediction values are off, resulting in high MSE percent-

ages. In all, our regression-based model performs better in both correlation and value reconstruction, showing its superiority in prediction on non-Gaussian data with small sample sizes.

Result Interpretation. Apart from better reconstruction performance, our model also has the advantage of result interpretability. We can explore the SC-FC mapping through the resulting coefficient matrix B s. Since our problem definition is $FC = SC \cdot B + E$, the i^{th} row in B corresponds to i^{th} edge pair in the SC vector. In the following, we say an edge i exists if row i of B has nonzero entries. As the experiment is run under the 10-fold cross validation setting and each training partition may generate a different mapping B , we consider an entry in the common B to be nonzero if it is nonzero more than five times in these 10 trials. The results are shown in Figure 2. From the figure, we can see for every task, several rows in B have many more nonzero entries than the others. This indicates the existence of several significant structural edges being responsible for most of the functional activities. To test if this assumption is valid, we compared B s from our model to randomly generated B s with the same levels of sparsity. The resulting MSE percentages, although having large variances, often have smaller means than that of CGGM and VAE, implying the importance of sparsity level of the coupling. However, the resulting r -scores of predicted FCs using a random B is the lowest among all methods. Together with very large p -values, the results predicted by random coupling show no correlation between predictions and ground truth. This indicates that our models learned meaningful mapping information from SCs to FCs, and that *structured sparsity* of B is important for getting predictions besides the level of sparsity alone.

We now analyze the B matrices for different cognitive tasks. During fMRI data acquisitions of all seven tasks, participants are presented with visual cues, either as images or videos, and they need to use motions such as pressing buttons to complete the tasks (Barch et al. 2013). Interestingly, apart from the LANGUAGE task, the mappings learned by our model predicts the strongest contribution of left precentral and left postcentral connection, which is on the motor cortex responsible for right-side body movement. From this, we assume most participants use their right hands to conduct the required finger movements for these tasks. All mappings also contain edges in the occipital lobe, complying with the visual nature of these tasks. The visualization of common edges that exist in all seven tasks is shown in Figure 3. This “backbone” roughly resembles the Default Mode Network (Greicius et al. 2009; Raichle 2015).

We then examined which structural edges contribute significantly to the functional activity under different tasks. For this, we plot the “high-contributing” edges in LANGUAGE, GAMBLING, RELATIONAL and WM tasks in Figure 4. An edge is considered as high-contributing if the number of nonzero entries of its corresponding row in B is more than half of the maximum number of nonzero entries of any row in B . From Figure 4, we notice although a common backbone exists, structural connections in different brain regions are responsible for specific tasks (e.g. SCs in and around hippocampus area appear to be highly contributing to the

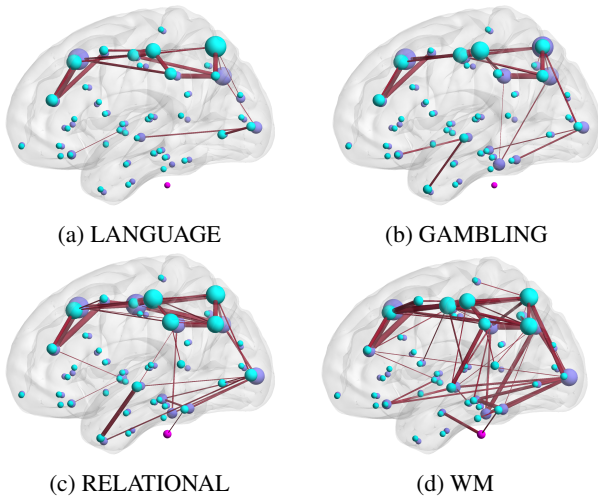


Figure 4: Task-specific visualizations for high-contributing structural edges. Assuming that the maximum number of nonzero entries of a row in \mathbf{B} is m , we only show the edges correspond to rows containing more than $m/2$ nonzero entries.

WM FCs but not so for the LANGUAGE ones, which is consistent with the literature (Baddeley, Jarrold, and Vargha-Khadem 2011)). We also plot both entry-wise and edge-wise overlap ratio for the mapping of seven tasks in Figure 5.

Another interesting phenomenon in Figure 2 is that the number of nonzero entries of \mathbf{B} for LANGUAGE, GAMBLING, RELATIONAL and WM are much larger than the other three tasks: EMOTION, MOTOR and SOCIAL, although the final model for each task have a similar level of prediction performance and similar hyperparameters. This is largely caused by the nature of non-overlapping \mathbf{B} s for EMOTION, MOTOR and SOCIAL tasks: their \mathbf{B} overlap ratios are significantly smaller than the other four tasks as shown in Table 3. Here we define the overlap ratio as the number of nonzero entries in the common \mathbf{B} (entry ij being nonzero if it's nonzero more than five times) over the number of nonzero entries in $\mathbf{B}_{\cup} = \mathbf{B}_1 \cup \dots \cup \mathbf{B}_{10}$ with \mathbf{B}_k being the predicted \mathbf{B} using the k^{th} split in 10-fold cross validation. This is also why we omit these three tasks for Figure 4, as the predicted \mathbf{B} s are not stable across the population and only a few common connections show significant contributions. We assume this results from group heterogeneity when carrying out these tasks. Further studies with heterogeneous models for the population will be useful to verify this assumption.

EMOTION	LANGUAGE	MOTOR	GAMBLING
5.34	25.58	3.60	35.36
SOCIAL	RELATIONAL	WM	
4.12	37.77	32.53	

Table 3: Overlap ratios (%) of predicted \mathbf{B} (SCs-FCs mapping) across 10 folds for seven tasks.

EMOTION	1	0.06	0.38	0.06	0.44	0.05	0.03
LANGUAGE	0.06	1	0.06	0.22	0.06	0.22	0.2
MOTOR	0.38	0.06	1	0.06	0.49	0.05	0.03
GAMBLING	0.06	0.22	0.06	1	0.06	0.24	0.23
SOCIAL	0.44	0.06	0.49	0.06	1	0.04	0.03
RELATIONAL	0.05	0.22	0.05	0.24	0.04	1	0.24
WM	0.03	0.2	0.03	0.23	0.03	0.24	1

(a)

EMOTION	1	0.14	0.91	0.13	1	0.11	0.09
LANGUAGE	1	1	1	0.9	1	0.79	0.64
MOTOR	0.7	0.1	1	0.1	1	0.08	0.07
GAMBLING	1	0.93	1	1	1	0.82	0.66
SOCIAL	0.35	0.05	0.46	0.05	1	0.04	0.03
RELATIONAL	1	0.99	1	0.99	1	1	0.79
WM	1	1	1	1	1	0.99	1

(b)

Figure 5: Entry-wise and edge-wise overlap ratio for the mapping of seven tasks. 5a considers entry-wise overlap of predicted \mathbf{B} of different tasks. The value on position (task i , task j) is the entry-wise IoU (Intersection over Union) of task i 's \mathbf{B} and task j 's \mathbf{B} , i.e. number of nonzero entries in $\mathbf{B}_i \cap \mathbf{B}_j$ over number of nonzero entries in $\mathbf{B}_i \cup \mathbf{B}_j$. 5b considers the the SC edges responsible for different tasks. An edge is considered to exist when its corresponding row in \mathbf{B} has nonzero entries. The value in position (task i , task j) is the number of common SC edges of task i and task j over the number of SC edges of task j .

Conclusions

In this paper, we proposed a regularized regression-based approach (CC-MRCE) for jointly learning linear models and partial correlations among variables under domain constraints. Motivated by the neuroscience application of predicting functional brain activities from structural connections, the CC-MRCE method discards the Gaussian assumption and incorporates domain constraints into model estimation. We further developed a fast algorithm based on nested FISTA to solve the optimization problem. With synthetic data analysis, we demonstrated that both domain constraints and assumption of non-Gaussian data contribute to the performance improvement of CC-MRCE. Our experimental results on Human Connectome Project data show that CC-MRCE outperforms existing methods on prediction tasks and uncovers couplings that agree with existing neuroscience literature.

Algorithm 1: CONCORD-MRCE

Input: penalty parameter λ_1 and λ_2
Initialize $t = 0$, $\hat{B}^{(0)} = \mathbf{0}$ and $\hat{\Omega}^{(0)} = \hat{\Omega}(\hat{B}^{(0)})$.
while not converged do
 step 1: Compute $\hat{S}^{(t-1)} = \hat{S}(\hat{B}^{(t-1)})$ as Eq.(4);
 step 2: Update $\hat{\Omega}^{(t)} = \hat{\Omega}(\hat{S}^{(t-1)})$ as Eq. (5) by
 calling CONCORD($\hat{S}^{(t-1)}$);
 step 3: Update $\hat{B}^{(t)} = \hat{B}(\hat{\Omega}^{(t)})$ as Eq.(6);
return B and Ω

Algorithm 2: CC-MRCE

Input: penalty parameter λ_1 and λ_2 , convex
constraint set C .
Initialize $\hat{B}^{(0)} = \mathbf{0}$ and $\hat{\Omega}^{(0)} = \hat{\Omega}(\hat{B}^{(0)})$.
while not converged do
 step 1: Compute $\hat{S}^{(t-1)} = \hat{S}(\hat{B}^{(t-1)})$ as Eq.(4);
 step 2: Update $\hat{\Omega}^{(t)} = \hat{\Omega}(\hat{S}^{(t-1)})$ by calling
 Constrained-CONCORD($\hat{S}^{(t-1)}, E, \lambda_1$);
 step 3: Update $\hat{B}^{(t)} = \hat{B}(\hat{\Omega}^{(t)})$ as Eq.(6);
return B and Ω

Acknowledgments

This project was partially supported by funding from the National Science Foundation under grant IIS-1817046.

Appendix A: Pseudocodes

CC-MRCE (**Algorithm 2**) improves CONCORD-MRCE by imposing hard constraints on the solution space.

In every step of the while loop, CC-MRCE calls sub-function Constrained-CONCORD to estimate the partial correlation matrix $\Omega^{(t)}$ under hard constraints, i.e. the solution $\Omega^{(t)}$ shares the same zero and nonzero pattern as matrix E (shown in **Algorithm 3**). Such constraints can be replaced by any set of convex constraints on $\Omega^{(t)}$. Note that superscript t' and t in **Algorithm 3** are iteration counters of inner and outer stages in the Constrained-CONCORD algorithm, respectively.

Lemma 2. *Let $L(g_1)$ be the Lipschitz constant of the gradient of objective function $g_1(\mathbf{H}_X)$, then $L(g_1) \leq 2\lambda_2^2\gamma^2$.*

Although we use line-search to pick a proper step length $\kappa_{t'}$ in the inner-loop of **Algorithm 3**, it can be replaced with a constant step length $\kappa_{t'} = 2\lambda_2^2\gamma^2$ according to the above lemma.

References

Baddeley, A.; Jarrold, C.; and Vargha-Khadem, F. 2011. Working memory and the hippocampus. *Journal of Cognitive Neuroscience* 23(12): 3855–3861.

Barch, D. M.; Burgess, G. C.; Harms, M. P.; Petersen, S. E.; Schlaggar, B. L.; Corbetta, M.; Glasser, M. F.; Curtiss, S.; Dixit, S.; Feldt, C.; et al. 2013. Function in the human connectome: task-fMRI and individual differences in behavior. *NeuroImage* 80: 169–189.

Algorithm 3: Constrained-CONCORD

Input: sample covariance matrix S , sparsity pattern E , penalty parameter λ_2
Output: partial correlation matrix Ω
set $\Theta^{(1)} = \Omega^{(0)} \in \mathbb{S}_M^p$, $\alpha_1 = 1$, $\tau_{(0,0)} \leq 1$, $c < 1$
while not converged do
 $\mathbf{G}^{(t)} = \nabla h_1(\Theta^{(t)})$
 search largest $\tau_t \in \{c^j \tau_{(t,0)}\}_{j=0,1,\dots}$
 $\mathbf{A}^{(t)} = \Theta^{(t)} - (n\tau_t/2)\mathbf{G}^{(t)}$
 set $\tilde{\Theta}^{(1)} = \mathbf{H}^{(0)} \in \mathbb{S}_{d \times d}$, $\tilde{\alpha}_1 = 1$, $\kappa_{(0,0)} \leq 1$,
 $\tilde{c} < 1$
 while not converged do
 $\tilde{\mathbf{G}}_X^{(t')} = \nabla g_1(\mathbf{H}_X^{(t')})$
 search largest $\kappa_{t'} \in \{\tilde{c}^{j'} \kappa_{(t',0)}\}_{j'=0,1,\dots}$
 $\mathbf{H}_X^{(t')} = \text{prox}_{g_2}(\tilde{\Theta}_X^{(t')} - \kappa_{t'} \tilde{\mathbf{G}}_X^{(t')})$
 check backtrack line search criterion
 $\tilde{\alpha}_{t'+1} = (1 + \sqrt{1 + 4\tilde{\alpha}_{t'}^2})/2$
 $\tilde{\Theta}^{(t'+1)} =$
 $\mathbf{H}_X^{(t')} + \frac{\alpha_{t'-1}}{\alpha_{t'+1}}(\mathbf{H}_X^{(t')} - \mathbf{H}_X^{(t'-1)})$
 compute $\kappa_{(t'+1,0)}$
 $\Omega^{(t)} = \mathbb{P}_C(\mathbf{A}_X^{(t)} - \gamma\lambda_2\mathbf{H}_X^{(t)}) + \mathbf{A}_D^{(t)}$
 check backtrack line search criterion
 $\alpha_{t+1} = (1 + \sqrt{1 + 4\alpha_t^2})/2$
 $\Theta^{(t+1)} = \Omega^{(t)} + \frac{\alpha_t - 1}{\alpha_{t+1}}(\Omega^{(t)} - \Omega^{(t-1)})$
 compute $\tau_{(t+1,0)}$

Bassett, D. S.; and Sporns, O. 2017. Network neuroscience. *Nature Neuroscience* 20(3): 353.

Bassett, D. S.; Zurn, P.; and Gold, J. I. 2018. On the nature and use of models in network neuroscience. *Nature Reviews Neuroscience* 19(9): 566–578.

Batista-García-Ramó, K.; and Fernández-Verdecia, C. I. 2018. What we know about the brain structure–function relationship. *Behavioral Sciences* 8(4): 39.

Beck, A.; and Teboulle, M. 2009. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences* 2(1): 183–202.

Becker, C. O.; Pequito, S.; Pappas, G. J.; Miller, M. B.; Grafton, S. T.; Bassett, D. S.; and Preciado, V. M. 2018. Spectral mapping of brain functional connectivity from diffusion imaging. *Scientific Reports* 8(1): 1–15.

Cole, M. W.; Bassett, D. S.; Power, J. D.; Braver, T. S.; and Petersen, S. E. 2014. Intrinsic and task-evoked network architectures of the human brain. *Neuron* 83(1): 238–251.

Columb, M.; and Atkinson, M. 2016. Statistical analysis: sample size and power estimations. *Bja Education* 16(5).

Davison, E. N.; Schlesinger, K. J.; Bassett, D. S.; Lynall, M.-E.; Miller, M. B.; Grafton, S. T.; and Carlson, J. M. 2015. Brain network adaptability across task states. *PLoS Computational Biology* 11(1).

- Eklund, A.; Nichols, T. E.; and Knutsson, H. 2016. Cluster failure: Why fMRI inferences for spatial extent have inflated false-positive rates. *Proceedings of the National Academy of Sciences* 113(28): 7900–7905.
- Fawcett, T. 2006. An introduction to ROC analysis. *Pattern Recognition Letters* 27(8): 861–874.
- Freyer, F.; Aquino, K.; Robinson, P. A.; Ritter, P.; and Breakspear, M. 2009. Bistability and non-Gaussian fluctuations in spontaneous cortical activity. *Journal of Neuroscience* 29(26): 8512–8524.
- Friedman, J.; Hastie, T.; and Tibshirani, R. 2010. Applications of the lasso and grouped lasso to the estimation of sparse graphical models. Technical report, Stanford University.
- Greicius, M. D.; Supekar, K.; Menon, V.; and Dougherty, R. F. 2009. Resting-state functional connectivity reflects structural connectivity in the default mode network. *Cerebral Cortex* 19(1): 72–78.
- Hagmann, P.; Cammoun, L.; Gigandet, X.; Meuli, R.; Honey, C. J.; Wedeen, V. J.; and Sporns, O. 2008. Mapping the structural core of human cerebral cortex. *PLoS Biology*.
- Hermundstad, A. M.; Bassett, D. S.; Brown, K. S.; Aminoff, E. M.; Clewett, D.; Freeman, S.; Frithsen, A.; Johnson, A.; Tipper, C. M.; Miller, M. B.; et al. 2013. Structural foundations of resting-state and task-based functional connectivity in the human brain. *Proceedings of the National Academy of Sciences* 110(15): 6169–6174.
- Hlinka, J.; Paluš, M.; Vejmelka, M.; Mantini, D.; and Corbetta, M. 2011. Functional connectivity in resting-state fMRI: is linear correlation sufficient? *NeuroImage* 54(3).
- Honey, C.; Sporns, O.; Cammoun, L.; Gigandet, X.; Thiran, J.-P.; Meuli, R.; and Hagmann, P. 2009. Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences* 106(6): 2035–2040.
- Honey, C. J.; Kötter, R.; Breakspear, M.; and Sporns, O. 2007. Network structure of cerebral cortex shapes functional connectivity on multiple time scales. *Proceedings of the National Academy of Sciences* 104(24): 10240–10245.
- Khare, K.; Oh, S.-Y.; and Rajaratnam, B. 2015. A convex pseudolikelihood framework for high dimensional partial correlation estimation with convergence guarantees. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 77(4): 803–825.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *ArXiv Preprint arXiv:1312.6114*.
- Koanantakool, P.; Ali, A.; Azad, A.; Buluc, A.; Morozov, D.; Olikier, L.; Yelick, K.; and Oh, S.-Y. 2018. Communication-avoiding optimization methods for distributed massive-scale sparse inverse covariance estimation. In *International Conference on Artificial Intelligence and Statistics*, 1376–1386. PMLR.
- McCarter, C.; and Kim, S. 2016. Large-scale optimization algorithms for sparse conditional gaussian graphical models. In *Artificial Intelligence and Statistics*, 528–537.
- Moreau, J.-J. 1965. Proximité et dualité dans un espace hilbertien. *Bulletin de la Société mathématique de France* 93: 273–299.
- Oh, S.; Dalal, O.; Khare, K.; and Rajaratnam, B. 2014. Optimization methods for sparse pseudo-likelihood graphical model selection. In *Advances in Neural Information Processing Systems*, 667–675.
- Peng, J.; Wang, P.; Zhou, N.; and Zhu, J. 2009. Partial correlation estimation by joint sparse regression models. *Journal of the American Statistical Association* 104(486): 735–746.
- Peng, J.; Zhu, J.; Bergamaschi, A.; Han, W.; Noh, D.-Y.; Pollack, J. R.; and Wang, P. 2010. Regularized multivariate regression for identifying master predictors with application to integrative genomics study of breast cancer. *The Annals of Applied Statistics* 4(1): 53.
- Raichle, M. E. 2015. The brain’s default mode network. *Annual Review of Neuroscience* 38: 433–447.
- Rockafellar, R. T. 1970. *Convex analysis*, volume 28. Princeton University Press.
- Rothman, A. J.; Levina, E.; and Zhu, J. 2010. Sparse multivariate regression with covariance estimation. *Journal of Computational and Graphical Statistics* 19(4): 947–962.
- Rubinov, M.; and Sporns, O. 2010. Complex network measures of brain connectivity: uses and interpretations. *NeuroImage* 52(3): 1059–1069.
- Sohn, K.-A.; and Kim, S. 2012. Joint estimation of structured sparsity and output structure in multiple-output regression via inverse-covariance regularization. In *Artificial Intelligence and Statistics*, 1081–1089.
- Van Essen, D. C.; Smith, S. M.; Barch, D. M.; Behrens, T. E.; Yacoub, E.; Ugurbil, K.; Consortium, W.-M. H.; et al. 2013. The WU-Minn human connectome project: an overview. *NeuroImage* 80: 62–79.
- Wytock, M.; and Kolter, Z. 2013. Sparse Gaussian conditional random fields: Algorithms, theory, and application to energy forecasting. In *International Conference on Machine Learning*, 1265–1273.
- Yuan, M.; and Lin, Y. 2007. Model selection and estimation in the Gaussian graphical model. *Biometrika* 94(1): 19–35.
- Yuan, X.-T.; and Zhang, T. 2014. Partial gaussian graphical model estimation. *IEEE Transactions on Information Theory* 60(3): 1673–1687.