The Sylvester Graphical Lasso (SyGlasso)

Yu Wang
University of Michigan
wayneyw@umich.edu

Byoungwook Jang
University of Michigan
bwjang@umich.edu

Alfred Hero
University of Michigan
hero@umich.edu

Abstract

This paper introduces *the Sylvester graphical lasso* (SyGlasso) that captures multiway dependencies present in tensor-valued data. The model is based on the Sylvester equation that defines a generative model. The proposed model complements the tensor graphical lasso (Greenewald et al., 2019) that imposes a Kronecker sum model for the inverse covariance matrix by providing an alternative Kronecker sum model that is generative and interpretable. A nodewise regression approach is adopted for estimating the conditional independence relationships among variables. The statistical convergence of the method is established, and empirical studies are provided to demonstrate the recovery of meaningful conditional dependency graphs. We apply the SyGlasso to an electroencephalography (EEG) study to compare the brain connectivity of alcoholic and non-alcoholic subjects. We demonstrate that our model can simultaneously estimate both the brain connectivity and its temporal dependencies.

1 Introduction

Estimating conditional independence patterns of multivariate data has long been a topic of interest for statisticians. In the past decade, researchers have focused on imposing sparsity on the precision matrix (inverse covariance matrix) to develop efficient estimators in the high-dimensional statistics regime where $n \ll p$. The success of the ℓ_1-penalized method for estimating multivariate dependencies was demonstrated in Meinshausen and Bühlmann (2006) and Friedman et al.

(2008) for the multivariate setting. This has naturally led researchers to generalize these methods to multiway tensor-valued data. Such generalizations are of benefit for many applications, including the estimation of brain connectivity in neuroscience, reconstruction of molecular networks, and detecting anomalies in social networks over time.

The first generalizations of multivariate analysis to the tensor-variate settings were presented by Dawid (1981), where the matrix-variate (a.k.a. two-dimensional tensor) distribution was first introduced to model the dependency structures among both rows and columns. Dawid (1981) extended the multivariate setting by rewriting the tensor-variate data as a vectorized (vec) representation of the tensor samples $\boldsymbol{\mathcal{X}} \in \mathbb{R}^{m_1 \times \cdots \times m_k}$ and analyzing the overall precision matrix $\boldsymbol{\Omega} = \mathbb{E}\big(\text{vec}(\boldsymbol{\mathcal{X}})\text{vec}(\boldsymbol{\mathcal{X}})^T\big) \in \mathbb{R}^{m \times m}$, where $m = \prod_{k=1}^{K} m_k$. Even for a two-dimensional tensor $\boldsymbol{\mathcal{X}} \in \mathbb{R}^{m_1 \times m_2}$, the computation complexity and sample complexity is high since the number of parameters in the precision matrix grows quadratically as m^2. Therefore, in the regime of tensor-variate data, unstructured precision matrix estimation has posed challenges due to the large number of samples needed for accurate structure recovery.

To address the sample complexity challenges, sparsity can be imposed on the precision matrix $\boldsymbol{\Omega}$ by using a sparse Kronecker product (KP) or Kronecker sum (KS) decompositions of $\boldsymbol{\Omega}$. The earliest and most popular form of sparse structured precision matrix estimation represents $\boldsymbol{\Omega}$ as the Kronecker product of smaller precision matrices. Tsiligkaridis et al. (2013) and Zhou (2014) proposed to model the precision matrix as a sparse Kronecker product of the covariance matrices along each mode of the tensor in the form $\boldsymbol{\Omega} = \boldsymbol{\Psi}_1 \otimes \cdots \otimes \boldsymbol{\Psi}_K$. The KP structure on the precision matrix has the nice property that the corresponding covariance matrix is also a KP. Zhou (2014) provides a theoretical framework for estimating the $\boldsymbol{\Omega}$ under KP structure and showed that the precision matrices can be estimated from a single instance under the matrix-variate normal distribution. Recently, Lyu

et al. (2019) extended the KP structured model to tensor-valued data, and provided new theoretical insights into the KP model. An alternative, called the Bigraphical Lasso, was proposed by Kalaitzis et al. (2013) to model conditional dependency structures of precision matrices by using a Kronecker sum representation $\mathbf{\Omega} = \mathbf{\Psi}_1 \oplus \mathbf{\Psi}_2 = (\mathbf{\Psi}_1 \otimes \mathbf{I}) + (\mathbf{I} \otimes \mathbf{\Psi}_2)$. On the other hand, Rudelson and Zhou (2017) and Park et al. (2017) studied the KS structure on the covariance matrix $\mathbf{\Sigma} = \mathbf{A} \oplus \mathbf{B}$ which corresponds to errors-in-variables models. More recently, Greenewald et al. (2019) proposed a model that generalized the KS structure to model tensor-valued data, called the TeraLasso. As shown in their paper, compared to the KP structure, KS structure on the precision matrix leads to a non-separable covariance matrix that provides a richer model than the KP structure.

KP vs KS: The KP model admits a simple stochastic representation as $\mathbf{X} = \mathbf{C}^{-1}\mathbf{Z}\mathbf{D}^{-1}$, where $\mathbf{A} = \mathbf{C}\mathbf{C}^T, \mathbf{B} = \mathbf{D}\mathbf{D}^T$, and $\mathbf{Z}$ is white Gaussian. It can be shown using properties of KP that $X \sim \mathcal{N}(0, (\mathbf{A} \otimes \mathbf{B})^{-1})$. Unlike the KP model, the KS model does not have a simple stochastic representation. From another perspective, the Kronecker structures can be characterized by the product graphs of the individual components. Kalaitzis et al. (2013) first motivated the KS structure on the precision matrix by relating $(\mathbf{\Psi}_1 \oplus \cdots \oplus \mathbf{\Psi}_K)$ to the associated Cartesian product graph. Thus, the overall structure of $\mathbf{\Omega}$ naturally leads to an interpretable model that brings the individual components together. The KP, however, corresponds to the direct tensor product of the individual graphs and leads to a denser dependency structure in the precision matrix Greenewald et al. (2019).

The Sylvester Graphical Lasso (SyGlasso): We propose a *Sylvester-structured graphical model* to estimate precision matrices associated with tensor data. Similar to the KP- and KS-structured graphical models, we simultaneously learn K graphs along each mode of the tensor data. However, instead of a KS or KP model for the precision matrix, the Sylvester structured graphical model uses a KS model for the square root factor of the precision matrix. The model is estimated by joint sparse regression models that impose sparsity on the individual components $\mathbf{\Psi}_k$ for $k = 1, \ldots, K$. The Sylvester model reduces to a squared KS representation for the precision matrix $\mathbf{\Omega} = (\mathbf{\Psi}_1 \oplus \cdots \oplus \mathbf{\Psi}_K)^2$, which is motivated by a stochastic representation of multivariate data with such a precision matrix. SyGlasso is the first KS-based graphical lasso model that admits a stochastic representation (i.e., Sylvester). Thus, our proposed SyGlasso puts the KS representations on similar ground as the KP representations in terms of interpretablility.

Notations

We adopt the notations used by Kolda and Bader (2009). A K-th order tensor is denoted by boldface Euler script letters, e.g., $\boldsymbol{\mathcal{X}} \in \mathbb{R}^{m_1 \times \cdots \times m_K}$. $\boldsymbol{\mathcal{X}}$ reduces to a vector for $K = 1$ and to a matrix for $K = 2$. The $(i_1, \ldots, i_K)$-th element of $\boldsymbol{\mathcal{X}}$ is denoted by $\boldsymbol{\mathcal{X}}_{i_1,\ldots,i_K}$, and we define the vectorization of $\boldsymbol{\mathcal{X}}$ to be $\text{vec}(\boldsymbol{\mathcal{X}}) := (\boldsymbol{\mathcal{X}}_{1,1,\ldots,1}, \boldsymbol{\mathcal{X}}_{2,1,\ldots,1}, \ldots, \boldsymbol{\mathcal{X}}_{m_1,1,\ldots,1}, \boldsymbol{\mathcal{X}}_{1,2,\ldots,1}, \ldots, \boldsymbol{\mathcal{X}}_{m_1,m_2,\ldots,m_k})^T \in \mathbb{R}^m$ with $m = \prod_{k=1}^{K} m_k$.

There are several tensor algebra concepts that we recall. A fiber is the higher order analogue of the row and column of matrices. It is obtained by fixing all but one of the indices of the tensor, e.g., the mode-k fiber of $\boldsymbol{\mathcal{X}}$ is $\boldsymbol{\mathcal{X}}_{i_1,\ldots,i_{k-1},:,i_{k+1},\ldots,i_K}$. Matricization, also known as unfolding, is the process of transforming a tensor into a matrix. The mode-k matricization of a tensor $\boldsymbol{\mathcal{X}}$, denoted by $\boldsymbol{\mathcal{X}}_{(k)}$, arranges the mode-$k$ fibers to be the columns of the resulting matrix. It is possible to multiply a tensor by a matrix – the k-mode product of a tensor $\boldsymbol{\mathcal{X}} \in \mathbb{R}^{m_1 \times \cdots \times m_K}$ and a matrix $\mathbf{A} \in \mathbb{R}^{J \times m_k}$, denoted as $\boldsymbol{\mathcal{X}} \times_k \mathbf{A}$, is of size $m_1 \times \cdots \times m_{k-1} \times J \times m_{k+1} \times \ldots m_k$. Its entry is defined as $(\boldsymbol{\mathcal{X}} \times_k \mathbf{A})_{i_1,\ldots,i_{k-1},j,i_{k+1},\ldots,i_K} := \sum_{i_k=1}^{m_k} \boldsymbol{\mathcal{X}}_{i_1,\ldots,i_K} A_{j,i_k}$. In addition, for a list of matrices $\{\mathbf{A}_1, \ldots, \mathbf{A}_K\}$ with $\mathbf{A}_k \in \mathbb{R}^{m_k \times m_k}$, $k = 1, \ldots, K$, we define $\boldsymbol{\mathcal{X}} \times \{\mathbf{A}_1, \ldots, \mathbf{A}_K\} := \boldsymbol{\mathcal{X}} \times_1 \mathbf{A}_1 \times_2 \cdots \times_K \mathbf{A}_K$. Lastly, we define the K-way Kronecker product as $\bigotimes_{k=1}^{K} \mathbf{\Psi}_k = \mathbf{\Psi}_1 \otimes \cdots \otimes \mathbf{\Psi}_K$, and the equivalent notation for the Kronecker sum as $\bigoplus_{k=1}^{K} \mathbf{\Psi}_k = \mathbf{\Psi}_1 \oplus \cdots \oplus \mathbf{\Psi}_K = \sum_{k=1}^{K} \mathbf{I}_{[d_{1:k-1}]} \otimes \mathbf{\Psi}_k \otimes \mathbf{I}_{[d_{k+1:K}]}$, where $\mathbf{I}_{[d_{k:\ell}]} = \mathbf{I}_{d_k} \otimes \cdots \otimes \mathbf{I}_{d_\ell}$.

2 Sylvester Graphical Lasso

Let a random tensor $\boldsymbol{\mathcal{X}} \in \mathbb{R}^{m_1 \times \cdots \times m_K}$ be generated by the following representation:

$$\boldsymbol{\mathcal{X}} \times_1 \mathbf{\Psi}_1 + \cdots + \boldsymbol{\mathcal{X}} \times_K \mathbf{\Psi}_K = \boldsymbol{\mathcal{T}}, \qquad (1)$$

where $\mathbf{\Psi}_k \in \mathbb{R}^{m_k \times m_k}, k = 1, \ldots, K$ are sparse symmetric positive definite matrices and $\boldsymbol{\mathcal{T}}$ is a random tensor of the same order as $\boldsymbol{\mathcal{X}}$. Equation (1) is known as the Sylvester tensor equation. The equation often arises in finite difference discretization of linear partial equations in high dimension (Bai et al., 2003) and discretization of separable PDEs (Kressner and Tobler, 2010; Grasedyck, 2004). When $K = 2$ it reduces to the Sylvester matrix equation $\mathbf{\Psi}_1 \mathbf{X} + \mathbf{X}\mathbf{\Psi}_2^T = \mathbf{T}$ which has wide application in control theory, signal processing and system identification (see, for example Datta and Zou (2017) and references therein).

It is not difficult to verify that the Sylvester representation (1) is equivalent to the following system of

linear equations:

$$\left(\bigoplus_{k=1}^{K} \boldsymbol{\Psi}_k\right) \text{vec}(\boldsymbol{\mathcal{X}}) = \text{vec}(\boldsymbol{\mathcal{T}}), \tag{2}$$

If $\boldsymbol{\mathcal{T}}$ is a random tensor such that $\text{vec}(\boldsymbol{\mathcal{T}})$ has zero mean and identity covariance, it follows from (2) that any $\boldsymbol{\mathcal{X}}$ generated from the stochastic relation (1) satisfies $\mathbb{E}\,\text{vec}(\boldsymbol{\mathcal{X}}) = \mathbf{0}$ and $\boldsymbol{\Sigma} = \boldsymbol{\Omega}^{-1} := \mathbb{E}\,\text{vec}(\boldsymbol{\mathcal{X}})\,\text{vec}(\boldsymbol{\mathcal{X}})^T = \left(\bigoplus_{k=1}^{K} \boldsymbol{\Psi}_k\right)^{-2}$. In particular, when $\text{vec}(\boldsymbol{\mathcal{T}}) \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_m)$, we have that $\text{vec}(\boldsymbol{\mathcal{X}}) \sim \mathcal{N}\left(\mathbf{0}, \left(\bigoplus_{k=1}^{K} \boldsymbol{\Psi}_k\right)^{-2}\right)$.

This paper proposes a procedure for estimating $\boldsymbol{\Omega}$ with N independent copies of the tensor data $\{\boldsymbol{\mathcal{X}}^i\}_{i=1}^{N}$ that are generated from (1). For the rest of the paper, we assume that the last mode of the data tensor corresponds to the observations mode. For example, when $K = 2$, $\boldsymbol{\mathcal{X}} \in \mathbb{R}^{m_1 \times m_2 \times N}$ is the matrix-variate data with N observations. Our goal is to estimate the K precision matrices $\{\boldsymbol{\Psi_k}\}_{k=1}^{K}$ each of which describes the conditional independence of k-th data dimension. The resulting precision matrix is $\boldsymbol{\Omega} = \left(\bigoplus_{k=1}^{K} \boldsymbol{\Psi}_k\right)^2$. By rewriting (2) element-wise, we first observe that

$$\left(\sum_{k=1}^{K} (\boldsymbol{\Psi}_k)_{i_k, i_k}\right) \boldsymbol{\mathcal{X}}_{i_{[1:K]}}$$
$$= -\sum_{k=1}^{K} \sum_{j_k \neq i_k} (\boldsymbol{\Psi}_k)_{i_k, j_k} \boldsymbol{\mathcal{X}}_{i_{[1:k]}, j_k, i_{[k+1:K]}} + \boldsymbol{\mathcal{T}}_{i_{[1:K]}}. \tag{3}$$

Note that the left-hand side of (3) involves only the summation of the diagonals of the $\boldsymbol{\Psi}$'s and the right-hand side is composed of columns of $\boldsymbol{\Psi}$'s that exclude the diagonal terms. Equation (3) can be interpreted as an autogregressive model relating the $(i_1, \ldots, i_K)$-th element of the data tensor (scaled by the sum of diagonals) to other elements in the fibers of the data tensor. The columns of $\boldsymbol{\Psi}'s$ act as regression coefficients. The formulation in (3) naturally leads us to consider a pseudolikelihood-based estimation procedure (Besag, 1977) for estimating $\boldsymbol{\Omega}$. It is known that inference using pseudo-likelihood is consistent and enjoys the same $\sqrt{N}$ convergence rate as the MLE in general (Varin et al., 2011). This procedure can also be more robust to model misspecification. Specifically, we define the sparse estimate of the underlying precision matrices along each axis of the data as the solution of the following convex optimization problem:

$$\min_{\substack{\boldsymbol{\Psi}_k \in \mathbb{R}^{m_k \times m_k} \\ k=1,\ldots K}} -N \sum_{i_1,\ldots,i_K} \log \boldsymbol{\mathcal{W}}_{i_{[1:K]}}$$
$$+ \frac{1}{2} \sum_{i_1,\ldots,i_K} \|(I) + (II)\|_2^2 + \sum_{k=1}^{K} P_{\lambda_k}(\boldsymbol{\Psi}_k). \tag{4}$$

where $P_{\lambda_k}(\cdot)$ is a penalty function indexed by the tuning parameter λ_k and

$$(I) = \boldsymbol{\mathcal{W}}_{i_{[1:K]}} \boldsymbol{\mathcal{X}}_{i_{[1:K]}}$$
$$(II) = \sum_{k=1}^{K} \sum_{j_k \neq i_k} (\boldsymbol{\Psi}_k)_{i_k, j_k} \boldsymbol{\mathcal{X}}_{i_{[1:k]}, j_k, i_{[k+1:K]}},$$

with $\boldsymbol{\mathcal{W}}_{i_{[1:K]}} := \sum_{k=1}^{K} (\boldsymbol{\Psi}_k)_{i_k, i_k}$. Here we focus on the ℓ_1-norm penalty, i.e., $P_{\lambda_k}(\boldsymbol{\Psi}_k) = \lambda_k \|\boldsymbol{\Psi}_k\|_{1,\text{off}}$.

The optimization problem (4) can be put into the following matrix form:

$$\min_{\substack{\boldsymbol{\Psi}_k \in \mathbb{R}^{m_k \times m_k} \\ k=1,\ldots K}} -\frac{N}{2} \log \left|(\text{diag}(\boldsymbol{\Psi}_1) \oplus \cdots \oplus \text{diag}(\boldsymbol{\Psi}_K))^2\right|$$
$$+ \frac{N}{2} \text{tr}(\mathbf{S}(\boldsymbol{\Psi}_1 \oplus \cdots \oplus \boldsymbol{\Psi}_K)^2) + \sum_{k=1}^{K} P_{\lambda_k}(\boldsymbol{\Psi}_k)$$

where $\text{diag}(\boldsymbol{\Psi}_k) \in \mathbb{R}^{m_k \times m_k}$ is a matrix of the diagonal entries of $\boldsymbol{\Psi}_k$ and $\mathbf{S} \in \mathbb{R}^{m \times m}$ is the sample covariance matrix, i.e., $\mathbf{S} = \frac{1}{N} \text{vec}(\boldsymbol{\mathcal{X}})^T \text{vec}(\boldsymbol{\mathcal{X}})$. Note that the pseudolikelihood above approximates the ℓ_1-penalized Gaussian negative log-likelihood in the log-determinant term by including only the Kronecker sum of the diagonal matrices instead of the Kronecker sum of the full matrices. Further discussion of pseudolikelihood- and likelihood-based approaches for (inverse) covariance estimations can be found in Khare et al. (2015).

We also note that when $K = 1$ the objective (4) reduces to the objective of the CONCORD estimator (Khare et al., 2015), and is similar to those of SPACE (Peng et al., 2009) and Symmetric lasso (Friedman et al., 2010). Our framework is a generalization of these methods to higher order tensor-valued data, when the Sylvester representation (1) holds.

Remark: In our formulation $\boldsymbol{\Omega} = (\bigoplus_{k=1}^{K} \boldsymbol{\Psi}_k)^2$ does not uniquely determine $\{\boldsymbol{\Psi}_k\}_{k=1}^{K}$ due to the trace ambiguity: scaled identity factors can be added to/subtracted from the $\boldsymbol{\Psi}'_k s$ without changing the matrix $\boldsymbol{\Omega}$. To address this non-identifiability, we rewrite the overall precision matrix $\boldsymbol{\Omega}$ as

$$\boldsymbol{\Omega} = \left(\bigoplus_{k=1}^{K} \boldsymbol{\Psi}_k\right)^2 = \left(\bigoplus_{k=1}^{K} \boldsymbol{\Psi}_k^{\text{off}} + \bigoplus_{k=1}^{K} \text{diag}(\boldsymbol{\Psi}_k)\right)^2,$$

where $\boldsymbol{\Psi}_k^{\text{off}} = \boldsymbol{\Psi}_k - \text{diag}(\boldsymbol{\Psi}_k)$, and estimate the off-diagonal entries $\boldsymbol{\Psi}_k^{\text{off}}$ and $\bigoplus_{k=1}^{K} \text{diag}(\boldsymbol{\Psi}_k)$ separately. This allows us to reconstruct the overall precision matrix $\boldsymbol{\Omega}$ when $\boldsymbol{\Psi}_k^{\text{off}}$ is penalized with an ℓ_1 penalty.

2.1 Estimation of the graphical model

Let $Q_N(\mathcal{W}, \{\boldsymbol{\Psi}_k^{\text{off}}\}_{k=1}^K)$ denote the objective function in (4), where $\mathcal{W} = \bigoplus_{k=1}^K \text{diag}(\boldsymbol{\Psi}_k)$. We adopt a convergent alternating minimization approach (Khare and Rajaratnam, 2014) that cycles between optimizing $\boldsymbol{\Psi}_k$ and $\mathcal{W}$ while fixing other parameters. In particular, for $1 \le k \le K$, $1 \le i_k < j_k \le m_k$, define

$$
\begin{aligned}
T_{i_k j_k}(\boldsymbol{\Psi}_k^{\text{off}}) &= \underset{\substack{(\tilde{\boldsymbol{\Psi}}_l)_{m,n} = (\boldsymbol{\Psi}_l)_{m,n} \\ \forall (l,m,n) \ne (k,i_k,j_k)}}{\arg\min} Q_N(\tilde{\mathcal{W}}, \{\tilde{\boldsymbol{\Psi}}_k^{\text{off}}\}_{k=1}^K) \\
T(\mathcal{W}) &= \underset{\substack{\tilde{\boldsymbol{\Psi}}_k^{\text{off}} = \boldsymbol{\Psi}_k^{\text{off}} \\ \forall k}}{\arg\min} Q_N(\tilde{\mathcal{W}}, \{\tilde{\boldsymbol{\Psi}}_k^{\text{off}}\}_{k=1}^K).
\end{aligned}
$$
(5)

For each (k, i_k, j_k), $T_{i_k j_k}(\boldsymbol{\Psi}_k^{\text{off}})$ updates the (i_k, j_k)-th entry with the minimizer of $Q_N(\mathcal{W}, \{\boldsymbol{\Psi}\}_{k=1}^K)$ with respect to $(\boldsymbol{\Psi}_k)_{i_k j_k}^{\text{off}}$ holding all other variables constant. Similarly, $T(\mathcal{W})$ updates $\mathcal{W}_{i_{[1:K]}}$ with the solution of $\min Q_N(\mathcal{W}, \{\boldsymbol{\Psi}\}_{k=1}^K)$ with respect to $\mathcal{W}_{i_{[1:K]}}$ holding all other variables constant. The closed form updates $T_{i_k j_k}(\boldsymbol{\Psi}_k^{\text{off}})$ and $T(\mathcal{W})$ are detailed in Appendix A.

Algorithm 1: Nodewise SyGlasso

Input: Standardized data $\mathcal{X}$, penalty parameter λ_k

Output: $\{\hat{\boldsymbol{\Psi}}_k\}_{k=1}^K$, $\hat{\boldsymbol{\Omega}} = \left(\bigoplus_{k=1}^K \hat{\boldsymbol{\Psi}}_k\right)^2$

Initialize $\{\hat{\boldsymbol{\Psi}}_k^{(0)}\}_{k=1}^K$, $\hat{\boldsymbol{\Omega}}^{(0)} = \left(\bigoplus_{k=1}^K \hat{\boldsymbol{\Psi}}_k^{(0)}\right)^2$

while *not converged* **do**
  # *Update off-diagonal elements*;
  for $k \leftarrow 1, \ldots, K$ **do**
    for $i_k \leftarrow 1, \ldots, m_k - 1$ **do**
      for $j_k \leftarrow i_k + 1, \ldots, m_k$ **do**
        $(\hat{\boldsymbol{\Psi}}_k^{(t+1)})_{i_k, j_k} \leftarrow (T_{i_k, j_k}(\boldsymbol{\Psi}_k^{(t)}))_{i_k, j_k}$;
              from (9) in Appendix A.1
      end
    end
  end
  # *Update diagonal elements*;
  $\hat{\mathcal{W}}^{(t+1)} \leftarrow T(\mathcal{W}^{(t)})$ from (10) in Appendix A.2
end

---

3 Large Sample Properties

We show that under suitable conditions, the Sylvester graphical lasso (SyGlasso) estimator (Algorithm 1) achieves both model selection consistency and estimation consistency. As in other studies, for the convergence analysis we make standard assumptions that the diagonal of $\boldsymbol{\Omega}$ is known (Khare et al., 2015; Peng et al.,

2009)[1]. We analyze the theoretical properties of the SyGlasso under the assumption that $\mathcal{W}$ is given. In practice, we can estimate $\mathcal{W}$ using Algorithm 1, and if the diagonals of each individual $\boldsymbol{\Psi}_k$ are desired, we can incorporate any available prior knowledge of the variation along each data dimension.

We estimate $\{\boldsymbol{\Psi}_k^{\text{off}}\}_{k=1}^K$ by solving the following ℓ_1 penalized problem:

$$
\min_{\boldsymbol{\beta}} L_N\left(\mathcal{W}, \boldsymbol{\beta}, \mathcal{X}\right) + \sum_{k=1}^K \lambda_k \|\boldsymbol{\Psi}_k\|_{1,\text{off}},
$$
(6)

where $L_N\left(\mathcal{W}, \boldsymbol{\beta}, \mathcal{X}\right) := \frac{1}{N} \sum_{s=1}^N L\left(\mathcal{W}, \boldsymbol{\beta}, \mathcal{X}^s\right)$, with

$$
\begin{aligned}
L\left(\mathcal{W}, \boldsymbol{\beta}, \mathcal{X}^s\right) = &-N \sum_{i_{[1:K]}} \log \mathcal{W}_{i_{[1:K]}} \\
&+ \frac{1}{2} \sum_{i_1, \ldots, i_K} ((I) + (II))^2.
\end{aligned}
$$
(7)

where

$(I) = \mathcal{W}_{i_{[1:K]}} \mathcal{X}_{i_{[1:K]}}$

$(II) = \sum_{k=1}^K \sum_{j_k \ne i_k} (\boldsymbol{\Psi}_k)_{i_k, j_k} \mathcal{X}_{i_{[1:k-1]}, j_k, i_{[k+1:K]}}$

$\boldsymbol{\beta} = ((\boldsymbol{\Psi}_1)_{1,2}, (\boldsymbol{\Psi}_1)_{1,3}, \ldots, (\boldsymbol{\Psi}_1)_{1,m_1}, \ldots, (\boldsymbol{\Psi}_k)_{m_k-1,m_k})^T$

and $\boldsymbol{\beta}$ denotes the off-diagonal entries of all $\boldsymbol{\Psi}_k's$.

We first state the regularity conditions needed for establishing convergence of the SyGlasso estimator. Let $\mathcal{A}_k := \{(i,j) : (\boldsymbol{\Psi}_k)_{i,j} \ne 0, i \ne j\}$ and $q_k := |\mathcal{A}_k|$ for $k = 1, \ldots, K$ be the true edge set and the number of edges, respectively. Let $\mathcal{A} = \cup_{k=1}^K \mathcal{A}_k$. We use $\bar{\boldsymbol{\beta}}, \bar{\boldsymbol{\Omega}}, \bar{\mathcal{W}}$ to emphasize that they are the true values of the corresponding parameters.

(A1 - Subgaussianity) The data $\mathcal{X}^1, \ldots, \mathcal{X}^N$ are i.i.d subgaussian random tensors, that is, $\text{vec}(\mathcal{X}^i) \sim \mathbf{x}$, where $\mathbf{x}$ is a subgaussian random vector in $\mathbb{R}^p$, i.e., there exist a constant $c > 0$, such that for every $\mathbf{a} \in \mathbb{R}^p$, $\mathbb{E}e^{\mathbf{a}^T x} \le e^{c\mathbf{a}^T \bar{\boldsymbol{\Sigma}} \mathbf{a}}$, and there exist $\rho_j > 0$ such that $\mathbb{E}e^{tx_j^2} \le K$ whenever $|t| < \rho_j$, for $1 \le j \le p$.

(A2 - Bounded eigenvalues) There exist constants $0 < \Lambda_{\min} \le \Lambda_{\max} < \infty$, such that the minimum and maximum eigenvalues of $\boldsymbol{\Omega}$ are bounded with $\lambda_{\min}(\bar{\boldsymbol{\Omega}}) = (\sum_{k=1}^K \lambda_{\max}(\boldsymbol{\Psi}_k))^{-2} \ge \Lambda_{\min}$ and $\lambda_{\max}(\bar{\boldsymbol{\Omega}}) = (\sum_{k=1}^K \lambda_{\min}(\boldsymbol{\Psi}_k))^{-2} \le \Lambda_{\max}$.

(A3 - Incoherence condition) There exists a constant $\delta < 1$ such that for $k = 1, \ldots, K$ and all $(i,j) \in \mathcal{A}_k$

$$|\bar{L}''_{ij,\mathcal{A}_k}(\bar{\mathcal{W}}, \bar{\boldsymbol{\beta}})[\bar{L}''_{\mathcal{A}_k,\mathcal{A}_k}(\bar{\mathcal{W}}, \bar{\boldsymbol{\beta}})]^{-1}\text{sign}(\bar{\boldsymbol{\beta}}_{\mathcal{A}_k})| \le \delta,$$

---

[1] When $K = 1$ it is possible to relax this assumption to require only accurate estimates of the diagonals, see Khare et al. (2015); Peng et al. (2009) for details.

where for each k and $1 \le i < j \le m_k$, $1 \le k < l \le m_k$,

$$\bar{L}''_{ij,kl}(\bar{\boldsymbol{\mathcal{W}}}, \bar{\boldsymbol{\beta}}) := E_{\bar{\boldsymbol{\mathcal{W}}}, \bar{\boldsymbol{\beta}}} \left(\frac{\partial^2 L(\boldsymbol{\mathcal{W}}, \boldsymbol{\beta}, \boldsymbol{\mathcal{X}})}{\partial(\boldsymbol{\Psi}_k)_{i,j} \partial(\boldsymbol{\Psi}_k)_{k,l}} |_{\boldsymbol{\mathcal{W}}=\bar{\boldsymbol{\mathcal{W}}}, \boldsymbol{\beta}=\bar{\boldsymbol{\beta}}} \right).$$

Note that conditions analogous to (A3) have been used in Meinshausen and Bühlmann (2006) and Peng et al. (2009) to establish high-dimensional model selection consistency of the nodewise graphical lasso in the case of $K = 1$. Zhao and Yu (2006) show that such a condition is almost necessary and sufficient for model selection consistency in lasso regression, and they provide some examples when this condition is satisfied.

Inspired by Meinshausen and Bühlmann (2006) and Peng et al. (2009) we prove the following properties:

1. Theorem 3.1 establishes estimation consistency and sign consistency for the nodewise SyGlasso restricted to the true support, i.e., $\boldsymbol{\beta}_{\mathcal{A}^c} = 0$,

2. Theorem 3.2 shows that no wrong edge is selected with probability tending to one,

3. Theorem 3.3 establishes consistency result of the nodewise SyGlasso.

Theorem 3.1. Suppose that conditions (A1-A2) are satisfied. Suppose further that $\lambda_{N,k} = O(\sqrt{\frac{m_k \log p}{N}})$ for all k and $N > O(\max_k q_k m_k \log p)$ as $N \to \infty$. Then there exists a constant $C(\bar{\boldsymbol{\beta}})$, such that for any $\eta > 0$, the following hold with probability at least $1 - O(\exp(-\eta \log p))$:

- There exists a global minimizer $\hat{\boldsymbol{\beta}}_{\mathcal{A}}$ of the restricted SyGlasso problem:

$$\min_{\boldsymbol{\beta}:\boldsymbol{\beta}_{\mathcal{A}^c}=0} L_N(\bar{\boldsymbol{\mathcal{W}}}, \boldsymbol{\beta}, \boldsymbol{\mathcal{X}}) + \sum_{k=1}^{K} \lambda_k \|\boldsymbol{\Psi}_k\|_{1,\text{off}}. \quad (8)$$

- (Estimation consistency) Any solution $\hat{\boldsymbol{\beta}}_{\mathcal{A}}$ of (8) satisfies:

$$\|\hat{\boldsymbol{\beta}}_{\mathcal{A}} - \boldsymbol{\beta}_{\mathcal{A}}\|_2 \le C(\bar{\boldsymbol{\beta}})\sqrt{K} \max_k \sqrt{q_k} \lambda_{N,k}.$$

- (Sign consistency) If further the minimal signal strength satisfies: $\min_{(i,j)\in\mathcal{A}_k} |(\boldsymbol{\Psi}_k)_{i,j}| \ge 2C(\bar{\boldsymbol{\beta}})\sqrt{K} \max_k \sqrt{q_k}\lambda_{N,k}$ for each k, then $\text{sign}(\hat{\boldsymbol{\beta}}_{\mathcal{A}_k})=\text{sign}(\bar{\boldsymbol{\beta}}_{\mathcal{A}_k})$.

Theorem 3.2. Suppose that the conditions of Theorem of 3,1 and (A3) are satisfied. Suppose further that $p = O(N^\kappa)$ for some $\kappa \ge 0$. Then for $\eta > 0$, for N sufficiently large, the solution of (8) satisfies:

$$P_{\bar{\boldsymbol{\mathcal{W}}}, \bar{\boldsymbol{\beta}}} \left(\max_{(i,j)\in\mathcal{A}_k^c} |L'_{N,ij}(\bar{\boldsymbol{\mathcal{W}}}, \hat{\boldsymbol{\beta}}_{\mathcal{A}_k}, \boldsymbol{\mathcal{X}})| < \lambda_{N,k} \right)$$
$$\ge 1 - O(\exp(-\eta \log p))$$

for each k, where $L'_{N,ij} := \partial L_N / \partial(\boldsymbol{\Psi}_k)_{ij}$.

Theorem 3.3. Assume the conditions of Theorem 3.2. Then there exists a constant $C(\bar{\boldsymbol{\beta}}) > 0$ such that for any $\eta > 0$ the following events hold with probability at least $1 - O(\exp(-\eta \log p))$:

- There exists a global minimizer $\hat{\boldsymbol{\beta}}$ to problem (4).

- (Estimation consistency) Any minimizer $\hat{\boldsymbol{\beta}}$ of (4) satisfies:

$$\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2 \le C(\bar{\boldsymbol{\beta}})\sqrt{K} \max_k \sqrt{q_k}\lambda_{N,k}.$$

- (Sign consistency) If $\min_{(i,j)\in\mathcal{A}_k} |(\boldsymbol{\Psi}_k)_{i,j}| \ge 2C(\bar{\boldsymbol{\beta}}) \max_k \sqrt{q_k}\lambda_{N,k}$ for each k, then $\text{sign}(\hat{\boldsymbol{\beta}})=\text{sign}(\bar{\boldsymbol{\beta}})$.

Proofs of the above theorems are given in Appendix B.

4 Numerical Illustrations

We evaluate the proposed SyGlasso estimator (Algorithm 1) in terms of optimization and graph recovery accuracy. We also compare the graph recovery performance with other models recently proposed for matrix- and tensor-variate precision matrices. We first illustrate the differences among these models by investigating the sparsity pattern of $\boldsymbol{\Omega}$ with $K = 3$ modes and $m_k = 4, \forall k$. For simplicity, we generate $\boldsymbol{\Psi}_k$ for $k = 1, 2, 3$ as identical 4×4 precision matrices that follow a one dimensional autoregressive-1 (AR1) process. We recall the KP and KS models:

Kronecker Product (KP): The KP model restricts the precision matrix and the covariance matrix to be separable across the K data dimensions and suffers from a multiplicative explosion in the number of edges. As they are separable models and the constructed $\boldsymbol{\Omega}$ corresponds to the direct product of the K graphs, KP is unable to capture more complex nested patterns captured by the KS and SyGlasso models as shown in Figure 1 (c) and (d).

Kronecker Sum (KS): The covariance matrix under the KS precision matrix assumption is nonseparable across K data dimensions, and the KS-structured models can be motivated from a maximum entropy point of view. Contrary to the KP structure, the number of edges in the KS structure grows as the sum of the edges of the individual graphs (as a result of Cartesian product of the associated graphs), which leads to a more controllable number of edges in $\boldsymbol{\Omega}$.

We compare these methods under different model assumptions to explore the flexibility of the proposed SyGlasso model under model mismatch. To empirically assess the efficiency of the proposed model, we
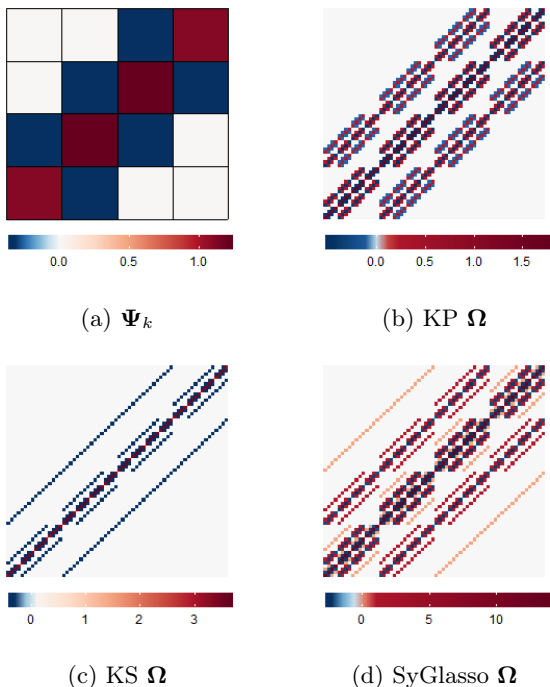
(a) $\mathbf{\Psi}_k$

(b) KP $\mathbf{\Omega}$

(c) KS $\mathbf{\Omega}$

(d) SyGlasso $\mathbf{\Omega}$

Figure 1: Comparison of SyGlasso to Kronecker sum (KS) and product (KP) structures. All models are composed of the same components $\mathbf{\Psi}_k$ for $k = 1, 2, 3$ generated as an AR(1) model with $m_k = 4$ as shown in (a). The AR(1) components are brought together to create the final 64×64 precision matrix $\mathbf{\Omega}$ following (b) the KP structure with $\mathbf{\Omega} = \bigotimes_{k=1}^{3} \mathbf{\Psi}_k$, (c) the KS structure with $\mathbf{\Omega} = \bigoplus_{k=1}^{3} \mathbf{\Psi}_k$, and (d) the proposed Sylvester model with $\mathbf{\Omega} = \left(\bigoplus_{k=1}^{3} \mathbf{\Psi}_k \right)^2$. The KP does not capture nested structures as it simply replicates the individual component with different multiplicative scales. The SyGlasso model admits a precision matrix structure that strikes a balance between KS and KP.

generate tensor-valued data based on three different precision matrices. The $\mathbf{\Psi}_k$'s are generated from one of 1) AR1(ρ), 2) Star-Block (SB), or 3) Erdos-Renyi (ER) random graph models described in Appendix C.

We test SyGlasso with $K = 2$ under: 1) SB with $\rho = 0.6$ and sub-blocks of size 16 and AR1($\rho = 0.6$); 2) SB with $\rho = 0.6$ and sub-blocks of size 16 and ER with 256 randomly selected edges. In both scenarios we set $m_1 = 128$ and $m_2 = 256$ with 10 samples. Figure 2 shows the iterative optimization performance of Algorithm 1. All the plots for the various scenarios exhibit iterative optimization approximation errors that quickly converge to values below the statistical errors. Note that these plots also suggest that our algorithm can attain linear convergence rates. We also test our method for model selection accuracy over a range of
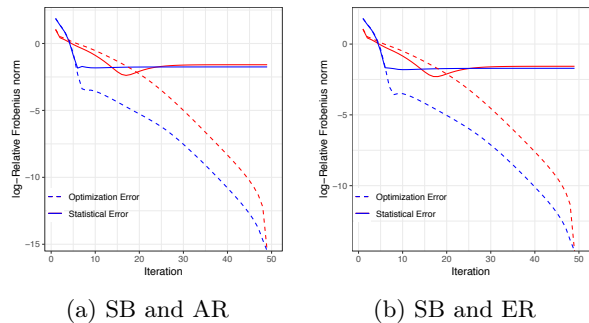

(a) SB and AR

(b) SB and ER

Figure 2: Performance of the SyGlasso estimator against the number of iterations under different topologies of $\mathbf{\Psi}_k$'s. The solid line shows the statistical error $\log(\|\hat{\mathbf{\Psi}}_k^{(t)} - \mathbf{\Psi}_k\|_F \backslash \|\mathbf{\Psi}_k\|_F)$, and the dotted line shows the optimization error $\log(\|\hat{\mathbf{\Psi}}_k^{(t)} - \hat{\mathbf{\Psi}}_k\|_F \backslash \|\hat{\mathbf{\Psi}}_k\|_F)$, where $\hat{\mathbf{\Psi}}_k$ is the final SyGlasso estimator. The performances of $\mathbf{\Psi}_1$ and $\mathbf{\Psi}_2$ are represented by red and blue lines, respectively.

penalty parameters (we set $\lambda_k = \lambda, \forall k$). Figure 3 displays the sum of false positive rate and false negative rate (FPR+FNR), it suggests that the nodewise SyGlasso estimator is able to fully recover the graph structures for each mode of the tensor data.
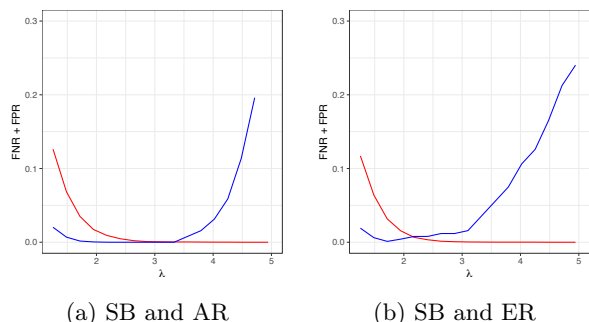

(a) SB and AR

(b) SB and ER

Figure 3: The performance of model selection measured by FPR + FNR. The performances of $\mathbf{\Psi}_1$ and $\mathbf{\Psi}_2$ are represented by red and blue lines, respectively. With an appropriate choice of λ, the SyGlasso recovers the dependency structures encoded in each $\mathbf{\Psi}_k$.

We compare the proposed SyGlasso to the TeraLasso estimator (Greenewald et al., 2019), and to the Tlasso estimator proposed by Lyu et al. (2019) for KP, on data generated using precision matrices $(\mathbf{\Psi}_1 \oplus \mathbf{\Psi}_2 \oplus \mathbf{\Psi}_3)^2$, $\mathbf{\Psi}_1 \oplus \mathbf{\Psi}_2 \oplus \mathbf{\Psi}_3$, and $\mathbf{\Psi}_1 \otimes \mathbf{\Psi}_2 \otimes \mathbf{\Psi}_3$, where $\mathbf{\Psi}$'s are each 16×16 ER graphs with 16 nonzero edges. We use the Matthews correlation coefficient (MCC) to compare model selection performances. The MCC is defined as (Matthews, 1975)

$$\text{MCC} = \frac{\text{TP} \times \text{TN} - \text{FP} \times \text{FN}}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}},$$

where we follow Greenewald et al. (2019) to consider each nonzero off-diagonal element of $\boldsymbol{\Psi}_k$ as a single edge.

The results shown in Figure 4 indicate that all three estimators perform well when $N = 5$, even under model misspecification. In the single sample scenario, the graph recovery performance of each estimator does well under each true underlying data generating process. Note that for data generated using KP, the SyGlasso performs surprisingly well and is comparable to Tlasso. These results seem to indicate that SyGlasso is very robust under model misspecification. The superior performance of SyGlasso under KP model, even with one sample, suggests again that SyGlasso structure has a flavor of both KS and KP structures, as seen in Figure 1. This follows from the observation that $(\boldsymbol{\Psi}_1 \oplus \boldsymbol{\Psi}_2)^2 = \mathbf{I}_{m_1} \otimes \boldsymbol{\Psi}_1^2 + \boldsymbol{\Psi}_2^2 \otimes \mathbf{I}_{m_2} + 2\boldsymbol{\Psi}_1 \otimes \boldsymbol{\Psi}_2 = \boldsymbol{\Psi}_1^2 \oplus \boldsymbol{\Psi}_2^2 + 2\boldsymbol{\Psi}_1 \otimes \boldsymbol{\Psi}_2$.

5 EEG Analysis

We revisit the alcoholism study conducted by Zhang et al. (1995) to explore multiway relationships in EEG measurements of alcoholic and control subjects. Each of 77 alcoholic subjects and 45 control subjects was visually stimulated by either a single picture or a pair of pictures on a computer monitor. Following the analyses of Zhu et al. (2016) and Qiao et al. (2019), we focus on the α frequency band (8 - 13 Hz) that is known to be responsible for the inhibitory control of the subjects (see Knyazev (2007) for more details). The EEG signals were bandpass filtered with the cosine-tapered window to extract α-band signals. Previous Gaussian graphical models applied to such α frequency band filtered EEG data could only estimate the connectivity of the electrodes as they cannot be generalized to tensor valued data. The SyGlasso reveals similar dependency structure as reported in Zhu et al. (2016) and Qiao et al. (2019) while recovering the chain structure of the temporal relationship.

Specifically, after the band-pass filter was applied, we work with the tensor data $\boldsymbol{\mathcal{X}}_{alcoholic}, \boldsymbol{\mathcal{X}}_{control} \in \mathbb{R}^{m_{nodes} \times m_{time} \times m_{trial}}$ corresponding to an alcoholic subject and a control subject. We simultaneously estimate $\boldsymbol{\Psi}_{node} \in \mathbb{R}^{m_{node} \times m_{node}}$ that encodes the dependency structure among electrodes and $\boldsymbol{\Psi}_{time} \in \mathbb{R}^{m_{time} \times m_{time}}$ that shows the relationship among time points that span the duration of each trial. Previous studies consider the average of all trials, for each subject and use the number of subjects as observations to estimate the dependency structures among 64 electrodes. Instead, we look at one subject at a time and consider different experimental trials as observations. Our analysis focuses on recovering the precision ma-
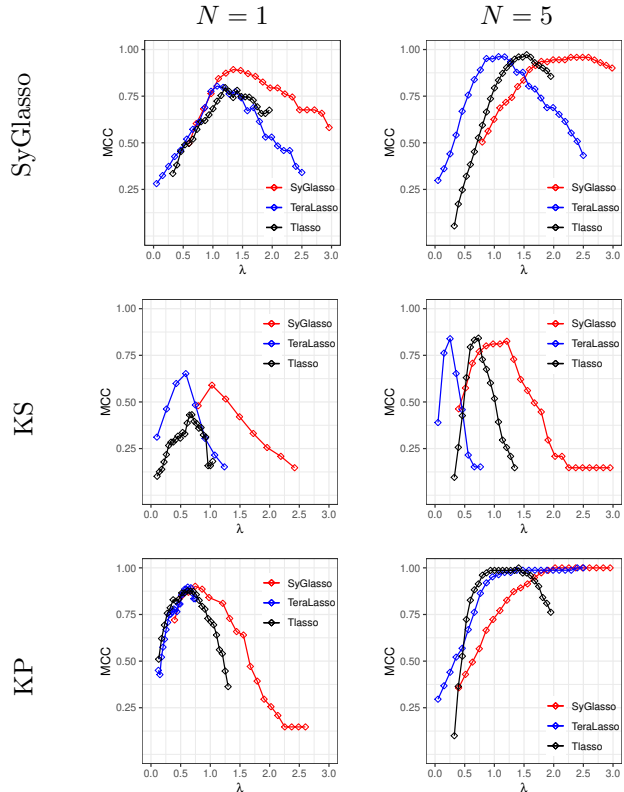

Figure 4: Performance of SyGlasso, TeraLasso (KS), and Tlasso (KP) measured by MCC under model misspecification. MCC of 1 represents a perfect recovery of the sparsity pattern in $\boldsymbol{\Omega}$, and MCC of 0 corresponds to random guess. From top to bottom, the synthetic data were generated with the precision matrices from SyGlasso, KS, and KP models. The left column shows the results for a single sample ($N = 1$), and the right column shows the results for $N = 5$ observations. Note that the SyGlasso has better performance for a single sample (left column) when data is generated from the matched Kronecker model and as good performance for the mismatched Kronecker models.

trices of electrodes and time points, but it can be easily generalized to estimate the dependency structure among trials as well.

Figure 5 shows the result of the SyGlasso estimated network of electrodes. For comparison, both graphs were thresholded to match 5% sparsity level. Similar to the findings of Qiao et al. (2019), our estimated graph $\boldsymbol{\Psi}_{node}$ for the alcoholic group shows the asymmetry between the left and the right side of the brain compared to the more balanced control group. Our finding is also consistent with the result in Hayden et al. (2006) and Zhu et al. (2016) that showed frontal asymmetry of the alcoholic subjects.

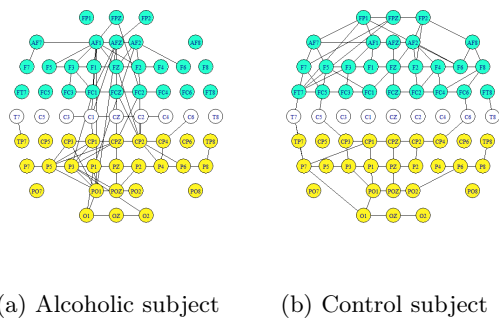(a) Alcoholic subject   (b) Control subject

Figure 5: Estimated brain connectivity results from SyGlasso for (a) the alcoholic subject and (b) the control subject. The blue nodes correspond to the frontal region, and the yellow nodes correspond to the parietal and occipital regions. The alcoholic subject has asymmetric brain connections in the frontal region compared to the control subject.
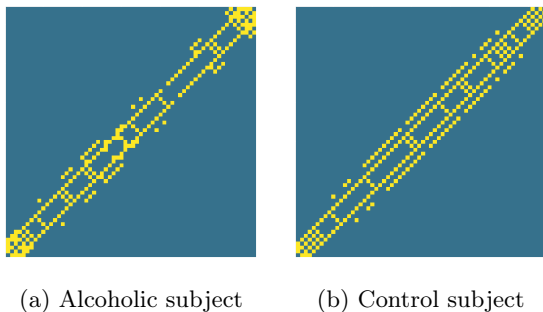

(a) Alcoholic subject   (b) Control subject

Figure 6: Support (off-diagonals) of SyGlasso-estimated temporal Sylvester factor $\hat{\mathbf{\Psi}}_{time}$ of the precision matrix for (a) the alcoholic subject and (b) the control subject. Both subjects exhibit banded conditional dependency structures over time.

While previous analyses on this EEG data using graphical models only focused on the precision matrix of the electrodes, here we exhibit in Figure 6 the second precision matrix that encodes temporal dependency. Overall both subjects exhibit banded dependency structures over time, since adjacent timepoints are conditionally dependent. However, note that the conditional dependency structure of the timepoints for the alcoholic subject appears to be more chaotic.

6 Discussion

This paper proposed Sylvester-structured graphical model and an inference algorithm, the SyGlasso, that can be applied to tensor-valued data. The current tools available for researchers are limited to Kronecker product and Kronecker sum models on either the covariance or the precision matrix. Our model is moti-

vated by a generative stochastic representation based on the Sylvester equation. We showed that the resulting precision matrix corresponds to the squared Kronecker sum of the precision matrices $\mathbf{\Psi}_k$ along each mode. The individual components $\mathbf{\Psi}_k$'s are estimated by the nodewise regression based approach.

There are several promising future directions to take the proposed SyGlasso. First is to relax the assumption that the diagonals of the factors are fixed - an assumption that is standard among the Kronecker structured models for theoretical analysis. Practically, Sy-Glasso is able to recover the off-diagonals of the individual $\mathbf{\Psi}_k$ and the diagonal of $\mathbf{\Omega}$, which only requires to estimating $\bigoplus_{k=1}^{K} \mathrm{diag}(\mathbf{\Psi}_k)$ instead of all diagonal entries $\mathrm{diag}(\mathbf{\Psi_k})$ for all k. However, we believe that analyzing the sparsity pattern of the squared Kronecker sum matrix would help us estimate the diagonal entries of the individual components $\mathbf{\Psi}_k$'s. In addition, it would be worthwhile to study optimization procedures that perform matrix-wise estimation of $\mathbf{\Psi}_k$ (yielding simultaneous estimates for both off-diagonal and diagonal entries) and compare its empirical and theoretical properties with the approach proposed in this paper.

Secondly, in terms of the statistical properties, our theoretical results guarantee sparsistency of the individual graphs with a slower convergence rate than that is proposed in Greenewald et al. (2019), while empirical evidence suggests that a faster rate can be achieved. Improvement of this statistical convergence rate analysis will be worthwhile. Also, our results do not guarantee the statistical convergence of individual $\mathbf{\Psi}_k$'s nor $\mathbf{\Omega}$ with respect to the operator norm. Similar to the solution proposed in Zhou et al. (2011), we plan to adopt a two-step procedure using SyGlasso for variable selection followed by refitting the precision matrix $\mathbf{\Omega}$ using maximum likelihood estimation with edge constraint.

Lastly, an exciting future direction is to investigate the utility of SyGlasso as a tool to perform system identification for systems that can be modeled by Sylvester approximations to differential equations. Such dynamical systems include, for example, physical processes governed by separable elliptic PDEs as described in Grasedyck (2004), Kressner and Tobler (2010). It is likely that myriad physical processes can be well-modeled using a Sylvester equation, leading to sparse precision matrices (e.g., a space-time process satisfying the Poisson equation).

References

Bai, Z.-Z., Golub, G. H., and Ng, M. K. (2003). Hermitian and skew-hermitian splitting methods for non-hermitian positive definite linear systems. *SIAM Journal on Matrix Analysis and Applications*, 24(3):603–626.

Besag, J. (1977). Efficiency of pseudolikelihood estimation for simple gaussian fields. *Biometrika*, pages 616–618.

Datta, A. and Zou, H. (2017). Cocolasso for high-dimensional error-in-variables regression. *The Annals of Statistics*, 45(6):2400–2426.

Dawid, A. P. (1981). Some matrix-variate distribution theory: notational considerations and a bayesian application. *Biometrika*, 68(1):265–274.

Friedman, J., Hastie, T., and Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441.

Friedman, J., Hastie, T., and Tibshirani, R. (2010). Applications of the lasso and grouped lasso to the estimation of sparse graphical models. Technical report, Technical report, Stanford University.

Grasedyck, L. (2004). Existence and computation of low kronecker-rank approximations for large linear systems of tensor product structure. *Computing*, 72(3-4):247–265.

Greenewald, K., Zhou, S., and Hero III, A. (2019). Tensor graphical lasso (teralasso). *To appear in JRSS-B. arXiv preprint arXiv:1705.03983*.

Hayden, E. P., Wiegand, R. E., Meyer, E. T., Bauer, L. O., O'Connor, S. J., Nurnberger Jr, J. I., Chorlian, D. B., Porjesz, B., and Begleiter, H. (2006). Patterns of regional brain activity in alcohol-dependent subjects. *Alcoholism: Clinical and Experimental Research*, 30(12):1986–1991.

Kalaitzis, A., Lafferty, J., Lawrence, N. D., and Zhou, S. (2013). The bigraphical lasso. In *International Conference on Machine Learning*, pages 1229–1237.

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.

Khare, K. and Rajaratnam, B. (2014). Convergence of cyclic coordinatewise l1 minimization. *arXiv preprint arXiv:1404.5100*.

Knyazev, G. G. (2007). Motivation, emotion, and their inhibitory control mirrored in brain oscillations. *Neuroscience & Biobehavioral Reviews*, 31(3):377–395.

Kolda, T. G. and Bader, B. W. (2009). Tensor decompositions and applications. *SIAM review*, 51(3):455–500.

Kressner, D. and Tobler, C. (2010). Krylov subspace methods for linear systems with tensor product structure. *SIAM journal on matrix analysis and applications*, 31(4):1688–1714.

Lyu, X., Sun, W. W., Wang, Z., Liu, H., Yang, J., and Cheng, G. (2019). Tensor graphical model: Nonconvex optimization and statistical inference. *IEEE transactions on pattern analysis and machine intelligence*.

Matthews, B. W. (1975). Comparison of the predicted and observed secondary structure of t4 phage lysozyme. *Biochimica et Biophysica Acta (BBA)-Protein Structure*, 405(2):442–451.

Meinshausen, N. and Bühlmann, P. (2006). High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics*, pages 1436–1462.

Park, S., Shedden, K., and Zhou, S. (2017). Nonseparable covariance models for spatio-temporal data, with applications to neural encoding analysis. *arXiv preprint arXiv:1705.05265*.

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.

Qiao, X., Guo, S., and James, G. M. (2019). Functional graphical models. *Journal of the American Statistical Association*, 114(525):211–222.

Rudelson, M. and Zhou, S. (2017). Errors-in-variables models with dependent measurements. *Electronic Journal of Statistics*, 11(1):1699–1797.

Tsiligkaridis, T., Hero III, A. O., and Zhou, S. (2013). On convergence of kronecker graphical lasso algorithms. *IEEE transactions on signal processing*, 61(7):1743–1755.

Varin, C., Reid, N., and Firth, D. (2011). An overview of composite likelihood methods. *Statistica Sinica*, pages 5–42.

Zhang, X. L., Begleiter, H., Porjesz, B., Wang, W., and Litke, A. (1995). Event related potentials during object recognition tasks. *Brain Research Bulletin*, 38(6):531–538.

Zhao, P. and Yu, B. (2006). On model selection consistency of lasso. *Journal of Machine learning research*, 7(Nov):2541–2563.

Zhou, S. (2014). GEMINI: Graph estimation with matrix variate normal instances. *The Annals of Statistics*, 42(2):532–562.

Zhou, S., Rütimann, P., Xu, M., and Bühlmann, P. (2011). High-dimensional covariance estimation based on gaussian graphical models. *Journal of Machine Learning Research*, 12(Oct):2975–3026.

Zhu, H., Strawn, N., and Dunson, D. B. (2016). Bayesian graphical models for multivariate functional data. *The Journal of Machine Learning Research*, 17(1):7157–7183.