Unsupervised Network Discovery for Brain Imaging Data

Zilong Bai
University of California, Davis
Computer Science
zlbai@ucdavis.edu

LCDR Peter Walker
Naval Medical Research Center
peter.b.walker.mil@mail.mil

Anna Tschiffely
Naval Medical Research Center
anna.e.tschiffely.ctr@mail.mil

Fei Wang
Cornell University
Healthcare Policy and Research
feiwang03@gmail.com

Ian Davidson
University of California, Davis
Computer Science
davidson@cs.ucdavis.edu

ABSTRACT

A common problem with spatiotemporal data is how to simplify the data to discover an underlying network that consists of cohesive spatial regions (nodes) and relationships between those regions (edges). This network discovery problem naturally exists in a multitude of domains including climate data (dipoles), astronomical data (gravitational lensing) and the focus of this paper, fMRI scans of human subjects. Whereas previous work requires strong supervision, we propose an unsupervised matrix tri-factorization formulation with complex constraints and spatial regularization. We show that this formulation works well in controlled experiments with synthetic networks and is able to recover the underlying ground-truth network. We then show that for real fMRI data our approach can reproduce well known results in neurology regarding the default mode network in resting-state healthy and Alzheimer affected individuals.

CCS CONCEPTS

• **Information systems** → *Clustering*;

KEYWORDS

network discovery; spatial regularization; spatiotemporal data; fMRI; brain

1 INTRODUCTION

Network discovery is crucial in analyzing complex spatiotemporal data in a multitude of applications, such as ocean temperature monitoring[14], climate modeling[8], astronomical data[12], and the subject of this paper, fMRI scans of human subjects[4]. Simplifying such spatiotemporal data into a network of cohesive spatial regions (nodes) and relationships between those regions (edges) offers great insight. Consider the focus of this paper, discovering cognitive networks in fMRI data. It is well known that the 70,000 voxels in a typical fMRI scan can be grouped into 116 anatomical regions[16]; what is unknown is how those regions are coactive to form a functional network for different tasks. Our earlier work[4] addresses this problem by adding strong supervision in the form of the anatomical regions' boundaries.

In this paper we remove this requirement that the regions are given and instead require our method to discover their boundaries and the connectivity between them *without supervision*. Figure 1 shows a simplified example involving a single slice of brain activity. This problem is particularly difficult given that these studies rarely involve more than 100 patients, rendering methods such as deep learning inapplicable.

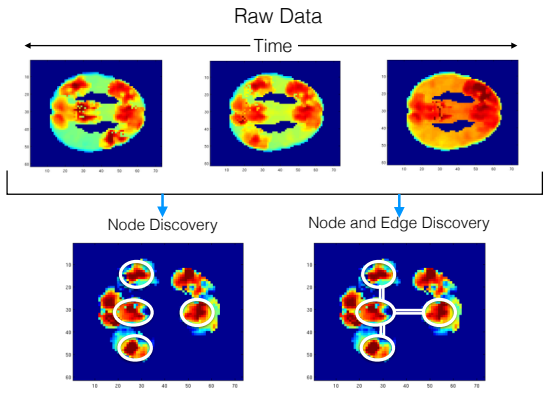

Figure 1: Our aim is to take spatiotemporal data and simultaneously discover nodes and edges to uncover the underlying network the person possesses without any supervision.

Though unsupervised matrix factorization methods such as the Orthogonal Nonnegative Matrix tri-Factorization (ONMtF) algorithms in [6] and [9] can be applied, they are not directly applicable to spatiotemporal problems. As shown in Figure 2a, when we set the rank of factorization to *six*, ONMtF cannot correctly discover *six* spatially continuous nodes; it can only discover *six* clusters of voxels with relatively high intra-cluster correlations. Even if we recursively apply this method by dividing the brain into two regions (using a factor 2, i.e. rank-2, ONMtF) and then divide each of those

KDD '17, August 13-17, 2017, Halifax, NS, Canada

ACM ISBN 978-1-4503-4887-4/17/08…$15.00
https://doi.org/10.1145/3097983.3098023

into two regions and so forth, the results are that recursive ONMtF cannot separate non-adjacent nodes even in the third recursion as shown in Figure 2b. To be interpreted as a cognitive network, the spatial regions must be continuous as these allow explanations in terms of the anatomical atlas as we have done in the captions of Figures 10, 11 and 12.

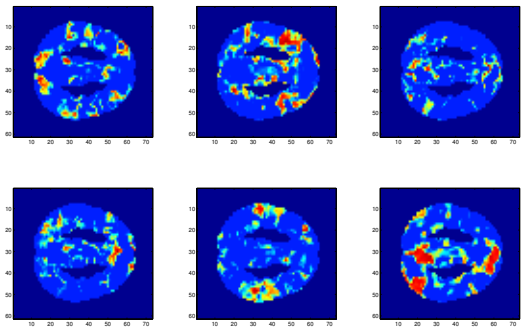

(a) A typical orthogonal non-negative symmetric matrix rank-6 tri-factorization by using ONMtF [6] .

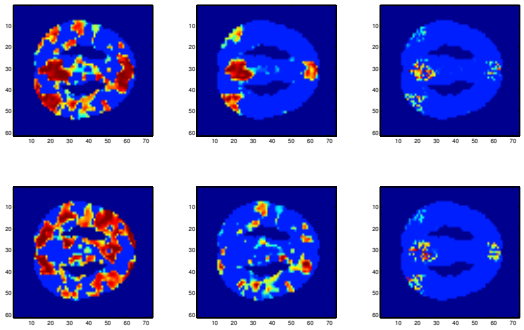

(b) Recursive foreground/background separation with ONMtF [6] for the typical scan sequence. Each column represents a 2 way factorization, one of which is then further refined.

Figure 2: The limitations of prior work (ONMtF [6]) for spatiotemporal data. Compare with Figure 10.

In this work, we extend the previous formulation of matrix tri-factorization (see Equation (1)) by adding a spatial regularizer (see Equation (2)). This is a non-trivial extension and produces a new type of solver involving new multiplicative update rules and a non-negative least squares solver. Our main contributions are as follows:

- We formulate a novel tri-factorization with spatial continuity regularization (ONMtF-SCR). This produces a novel non-negative least squares solver (Equation (17)) and multiplicative update rules (Equation (10)).
- The tri-factorization allows a natural representation where one matrix represents the nodes and another the associations between the nodes, thus simultaneously discovering nodes and edges in one framework.
- Our tri-factorization allows the discovery of networks without strong supervision. This produces results where each factor/node can be interpreted as an activated anatomical region and most importantly can generate new insights.
- We verify that our method discovers interpretable nodes that reproduce well known results in neurology regarding the default mode network in resting-state healthy and Alzheimer affected individuals for fMRI scans of human subjects.

We begin this paper by discussing related works about both the problem and the model. In the next section, we introduce and define the unsupervised network discovery problem. In the third section, we present our formulation of orthogonal non-negative matrix tri-factorization with continuity regularization; our solver is presented in the subsequent section. Extensive experiments were conducted to evaluate our approach first on synthetic networks where the ground truth is known and then on fMRI data. Finally, we conclude our work.

2 RELATED WORK

This work is related to two areas: the application domain of network discovery and underlying formulation of matrix tri-factorization.

2.1 Network Discovery

Most work on network discovery focuses on either node identification or edge identification but to our knowledge not simultaneously discovering both as we propose. Furthermore, finding a network with heterogeneous edges to our knowledge has not been addressed. A variety of current data mining tools have been used to determine, in an unsupervised setting, clusterings of voxels (what we refer to as nodes/regions) that appear to exhibit relatively high connectivity or covary with each other in systematic ways [2, 17]. Another stream of research focuses on just identifying connectivity between given brain regions using dynamic belief networks [3], and our own work using sparse inverse covariance estimation [15]. Some supervised learning work could be used to predict the existence of a network. Such work performs dimension reduction, followed by classification or regression, on the 4D array provided by the Blood Oxygenation Level Dependence (BOLD) fMRI [5]. However, probing functional connectivity goes far beyond a yes or no question of whether two regions are functionally connected; activity in one region can lag, excite another region, or coordinate only in certain circumstances [11].

Early work [1] used projection methods to uncover the subset of voxels that are co-active. For example principal component analysis can be used to determine which subset of the voxels have similar co-activation patterns. Here the data to project consists of a matrix $Cohort \times Voxel$. The projection vector's sign can be interpreted as an indicator vector for those voxels that are coactive. This work only provides a coarse measurement of the network as it divides the voxels into "in network" and "out of network". Later work is more refined in that it discovers the individual parts of the network. This work can naturally be modeled as a tensor decomposition problem into k factors using PARAFAC where each factor corresponds to an anatomical region. Consider a slice of activity modeled as an order 3 tensor: $x \times y \times time$. Then the decomposition produces factor

matrices $A \times B \times T$ where A is $x \times k$, B is $y \times k$ and T is $t \times k$. The outer-product of each column of the matrices (i.e. A_i and B_i) produces a spatial pattern which can be interpreted as a combination of voxels that are co-active according to the vector T_i. However, such a formulation does not discover spatial patterns that correspond to anatomical regions. To overcome this, strong supervision must used as we did in our earlier work [4]. However, use of such strong supervision means making a strong prior assumption that this anatomical region information is known and accurate which is often not the case if the study involves complex understudied settings where neuroplasticity (brain reconfiguration) occurs such as TBI (traumatic brain injury), PTSD (post traumatic stress disorder) and Alzheimer affected individuals.

Finally, finding consensus networks in populations is primarily done by visual inspection of ICA components of rasterized data; little work focuses on finding similarities and differences between two cohorts (sub-populations).

2.2 Matrix tri-Factorization

Matrix tri-factorization with non-negative constraints has been applied to simplify traffic data [10]. As we focus on discovering interpretable cognitive networks, the discovered nodes should be non-overlapped clusters of voxels. Among various models for matrix tri-factorization, the model most related to our work is [6], in which multiplicative algorithms are proposed for solving matrix tri-factorization with orthogonality constraints on clustering indicator matrices to ensure non-overlapped clusters. An Orthogonal Non-negative Matrix tri-Factorization (ONMtF) model for symmetric matrix is formulated in Equation (1). However, the orthogonality constraints are *incapable* of discovering nodes with spatial continuity as shown in empirical experiments on fMRI scans (See Figure 2).

$$\underset{\mathbf{F}\geq 0, \mathbf{M}\geq 0}{Minimize} \|\mathbf{X} - \mathbf{F}\mathbf{M}\mathbf{F}^T\|_F^2 \quad s.t. \quad \mathbf{F}^T\mathbf{F} = \mathbf{I} \tag{1}$$

To our knowledge, the use of *non-negative matrix tri-factorization with spatial continuity regularization* to discover cognitive networks is novel and is particularly challenging when *orthogonality* is incorporated on top of these requirements. More specifically, our work differs from [6] in our novelty of incorporating spatial continuity into the multiplicative update rule for $\mathbf{F}$ and using nonnegative least squares solver to solve for symmetric $\mathbf{M}$.

3 PROBLEM OVERVIEW AND DEFINITION

We now discuss our proposed contributions at the functional level and go into more details with respect to formulations and solvers in the next sections. fMRI is a predominant method for capturing brain activity as it processes information [7, 18] with a natural question being to simplify the activity into the underlying cognitive network. The example in Figure 1 shows that there are multiple groups of voxels which coactivate together even though they are spatially far apart.

Problem 1 ***Unsupervised Network Discovery Problem***
Input: $\mathcal{D} \in \mathbb{R}^{Spatial \times Temporal}$, *k the number of nodes to discover.*
Output: A collection of k non-overlapping continuous nodes $\mathbf{F}_1 \ldots \mathbf{F}_k$ *where* $\mathbf{F}_i \in [0,1]^{Spatial}, \forall i \neq j, \mathbf{F}_i^T\mathbf{F}_j = 0,$
$\mathbf{M} \in \mathbb{R}^{k \times k}$ *a measure of association between the nodes.*

In this paper we take $\mathcal{D}$ and convert it to an affinity/similarity matrix $\mathbf{X}$ between the voxels.

3.1 Problem Formulation: ONMtF-SCR

In this section we formulate the network discovery problem as a non-negative symmetric matrix tri-factorization model with orthogonality constraint and spatial continuity regularization (ONMtF-SCR) as shown in Equation (2).[1]

$$\begin{aligned} &\underset{\mathbf{F}\geq 0, \mathbf{M}\geq 0}{Minimize} \|\mathbf{X} - \mathbf{F}\mathbf{M}\mathbf{F}^T\|_F^2 + \beta tr(\mathbf{F}^T\Theta\mathbf{F}) \\ &s.t. \quad \mathbf{F}^T\mathbf{F} = \mathbf{I} \end{aligned} \tag{2}$$

The Modeling Assumptions Made. We first discuss how Equation (2) simplifies the underlying behavior in the fMRI data. The tri-factorization effectively rewrites the fMRI image as a linear combination of subnetworks in the form $\sum_{i,j}(\mathbf{M})_{ij}\mathbf{F}_{*i}\mathbf{F}_{*j}^T$ as diagrammatically shown in Figure 3. Here the *weighted* **outer** product of $\mathbf{F}_{*i}$ and $\mathbf{F}_{*j}$ represents a subgraph consisting of node i and j, and the weight $\mathbf{M}_{ij}$ thus indicates the amount of inter-node (or intra-node if $i = j$) associations between node i and j. In this way the network is modeled as a series of node pairs with a strength of association of each pair.

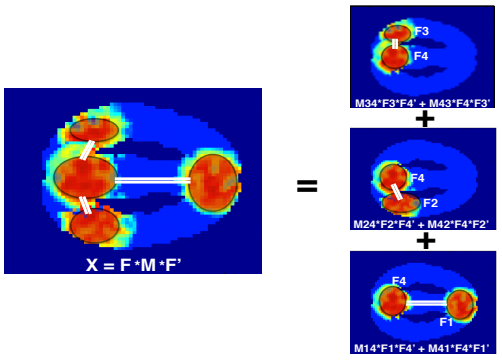

Figure 3: How our formulation models a complex network as a linear combination of simpler parts. The example 4-node underlying network is decomposed into a linear combination of three subnetworks.

Interpretation of ONMtF-SCR for Medical Imaging. Here $\mathbf{X}$ can be the correlations between each pair of voxels over time according to the BOLD values taken from the fMRI scans.

$\mathbf{F}$ is the node indicator matrix. Each column of $\mathbf{F}$ is a nonnegative ($\mathbf{F}_{i,j} \geq 0, \forall i,j$) and normalized ($\|\mathbf{F}_{*j}\|_2 = 1, \forall j$) indicator vector where $\mathbf{F}_{i,j}$ denotes the membership weight of voxel i in node j. The columns of $\mathbf{F}$ are orthogonal to each other ($\mathbf{F}_{*i}^T\mathbf{F}_{*j} = 0, \forall i \neq j$) to guarantee that the discovered nodes are **non-overlapping**.

$\mathbf{M}$ is a symmetric non-negative matrix that describes the associations/interactions between the nodes indicated by $\mathbf{F}$. The entries of $\mathbf{M}$ can be used to construct the edges between the nodes of the functional network shown in Figure 1.

[1] In this paper superscript " T " denotes matrix transpose, and "tr()" represents matrix trace calculation.

The Spatial Regularizer. Spatial continuity within each discovered node is enforced by the regulation term $\beta tr(\mathbf{F}^T\Theta\mathbf{F})$ where Θ is a penalty matrix that discourages voxels located far apart spatially from being clustered into the same node. Note our regularizer is very general and does not simply state the adjacent columns (or rows) in the matrix $\mathbf{X}$ should be together. Any forms of spatial properties can be used. Specifically, Θ is a reciprocal Gaussian Kernel matrix for each pair of voxels within brain area as shown in Equation (3). Vectors $\mathbf{v}_i$ and $\mathbf{v}_j$ indicate the spatial $x-y$ locations of voxels i and j on the slice of the brain image this paper evaluates. The diagonal entries of Θ *enforce L-2 norm sparsity regularization* on the region indicator vectors.

$$(\Theta)_{ij} = e^{\frac{\|\mathbf{v}_i-\mathbf{v}_j\|_2^2}{2\sigma^2}} \tag{3}$$

The spatial continuity regularization in Equation (2) is incorporated to improve the interpretability of discovered networks since the orthogonality constraint *does not* guarantee each discovered node to be spatially continuous.

Solving Equation (2) Directly is Challenging. One common approach to solve constrained optimization problems is the Lagrangian multiplier method and accordingly the Lagrangian function of Equation (2) is derived below.

$$\begin{aligned} L(\mathbf{F},\mathbf{M},\lambda) &= tr(\mathbf{XX}) - 2tr(\mathbf{XFMF}^T) + \beta tr(\mathbf{F}^T\Theta\mathbf{F}) \\ &+ tr(\mathbf{FMF}^T\mathbf{FMF}^T) + tr(\lambda(\mathbf{F}^T\mathbf{F})) - tr(\lambda) \end{aligned} \tag{4}$$

The presence of the 4th-order term $tr(\mathbf{FMF}^T\mathbf{FMF}^T)$ in Equation (4) makes directly solving it extremely challenging and it *cannot* be reduced to $tr(\mathbf{FM}^2\mathbf{F}^T)$ unless $\mathbf{F}^T\mathbf{F} \equiv \mathbf{I}$ is guaranteed through each iteration of the solver.

4 ALGORITHM FOR SOLVING ONMTF-SCR

In this section we present the solver for ONMtF-SCR in Algorithm 1. Given the difficulty of directly solving Equation (2) there are two possible directions: *1)* Simplify the formulation or *2)* Approximately solve the complex formulation. The formulation cannot be simplified, as each part of it is essential to formulate the problem of unsupervised network discovery, and it is our desire to address the problem directly. Thus, the latter option is taken: instead of directly solving Equation (2), the problem is reformulated into two subproblems, which are then solved within one framework. Particularly, multiplicative update rules are developed for solving the node discovery subproblem, and a non-negative least squares subproblem is solved for edge discovery. Our experiments suggest the solutions found are effective on both synthetic and real-world data.

4.1 Unsupervised Node Discovery Subproblem

To solve the node discovery subproblem, Equation (2) is simplified to Equation (5) by substituting the term $\mathbf{FM}^T$ with a single optimization matrix variable $\mathbf{G}$ while keeping both the orthogonality constraints and spatial regularization on $\mathbf{F}$.

$$\begin{aligned} &\underset{\mathbf{F}\geq 0,\mathbf{G}\geq 0}{Minimize} \|\mathbf{X}-\mathbf{FG}^T\|_F^2 + \beta tr(\mathbf{F}^T\Theta\mathbf{F}) \\ &s.t. \quad \mathbf{F}^T\mathbf{F} = \mathbf{I} \end{aligned} \tag{5}$$

This simplification makes it plausible to derive multiplicative update rules that can be efficiently applied to obtain an effective solution

Algorithm 1 Algorithm for Network Discovery.

Require: i. Θ: Spatial continuity penalty matrix (constructed with spatial bandwidth σ)
ii. β: Weight of spatial continuity regularization
iii. r: Rank of factorization
Output: $\mathbf{F}$, $\mathbf{M}$
Initialization: Randomly initialize $\mathbf{F}_{s\times r}$ and $\mathbf{G}_{s\times r}$;
Solve Equation (5) and set $\mathbf{F}$ for node discovery:
while Termination condition not satisfied **do**
  Update $\mathbf{F}$ with Equation (10), λ calculated with Equation (14);
  Update $\mathbf{G}$ with Equation (11);
end while
Solve Equation (17) and set $\mathbf{M}$ for edge discovery;
Rescale $\mathbf{M}$ with Equation (18);
Rescale $\mathbf{F}$ with Equation (19);
return: $\mathbf{F}$, $\mathbf{M}$

for $\mathbf{F}$ and $\mathbf{G}$. The Lagrangian of Equation (5) is given by Equation (6) with the $r\times r$ symmetric matrix λ as Lagrangian multipliers.

$$L(\mathbf{F},\mathbf{G},\lambda) = \|\mathbf{X}-\mathbf{FG}^T\|_F^2 + \beta tr(\mathbf{F}^T\Theta\mathbf{F}) + tr(\lambda(\mathbf{F}^T\mathbf{F}-\mathbf{I})) \tag{6}$$

Since $\|\mathbf{X}-\mathbf{FG}^T\|_F^2 = tr((\mathbf{X}-\mathbf{FG}^T)(\mathbf{X}-\mathbf{FG}^T)^T)$, partial differentiation over $\mathbf{F}$ on $L(\mathbf{F},\mathbf{G},\lambda)$ is given by Equation (7)

$$\frac{\partial L(\mathbf{F},\mathbf{G},\lambda)}{\partial \mathbf{F}} = -2\mathbf{XG} + 2\mathbf{FG}^T\mathbf{G} + 2\beta\Theta\mathbf{F} + 2\mathbf{F}\lambda \tag{7}$$

A local optima is obtained when Equation (7) is set to 0 and by using *KKT* (*Karush-Kuhn-Tucker*) complementary condition [6] yields:

$$(-2\mathbf{XG} + 2\mathbf{FG}^T\mathbf{G} + 2\beta\Theta\mathbf{F} + 2\mathbf{F}\lambda)_{ik}\mathbf{F}_{ik} = 0 \tag{8}$$

This is equivalent to Equation (9) if either $\mathbf{F}_{ik} = 0$ or the left factor is equal to 0.

$$(-2\mathbf{XG} + 2\mathbf{FG}^T\mathbf{G} + 2\beta\Theta\mathbf{F} + 2\mathbf{F}\lambda)_{ik}\mathbf{F}_{ik}^2 = 0 \tag{9}$$

Hence,

Multiplicative Update Rule for F

$$\mathbf{F}_{i,k} \leftarrow \mathbf{F}_{i,k}\sqrt{\frac{(\mathbf{XG})_{i,k}}{(\mathbf{FG}^T\mathbf{G} + \beta\Theta\mathbf{F} + \mathbf{F}\lambda)_{i,k}}} \tag{10}$$

By partial differentiation over $\mathbf{G}$ on $L(\mathbf{F},\mathbf{G},\lambda)$,

Multiplicative Update Rule for G

$$\mathbf{G}_{i,k} \leftarrow \mathbf{G}_{i,k}\frac{(\mathbf{X}^T\mathbf{F})_{i,k}}{(\mathbf{G}(\mathbf{F}^T\mathbf{F}))_{i,k}} \tag{11}$$

An approximate compact expression for λ can be derived by summing over i in Equation (8) (a technique also used by [6]):

$$\begin{aligned} &(-\mathbf{F}^T\mathbf{XG} + (\mathbf{F}^T\mathbf{F})(\mathbf{G}^T\mathbf{G}) + (\mathbf{F}^T\mathbf{F})\lambda + \beta\mathbf{F}^T\Theta\mathbf{F})_{k,k} = 0 \\ &(-\mathbf{F}^T\mathbf{XG} + (\mathbf{G}^T\mathbf{G}) + \lambda + \beta\mathbf{F}^T\Theta\mathbf{F})_{k,k} = 0 \end{aligned} \tag{12}$$

Hence,

$$\lambda_{k,k} = (\mathbf{F}^T\mathbf{XG} - \mathbf{G}^T\mathbf{G} - \beta\mathbf{F}^T\Theta\mathbf{F})_{k,k} \tag{13}$$

By ignoring the non-negativity of $\mathbf{F}$, approximate off-diagonal entries of $\lambda_{i,j}, i\neq j$ can be obtained, which lead to compact expression for λ:

$$\lambda = (\mathbf{F}^T\mathbf{XG} - \mathbf{G}^T\mathbf{G} - \beta\mathbf{F}^T\Theta\mathbf{F}) \tag{14}$$

In contrast to the ONMF with only orthogonality constraints in [6], the update rule of **F** in our model is novel in incorporating spatial continuity regularization with Θ and β.

4.2 Edge Discovery Subproblem

In the second stage we aim to recover the association matrix **M**. As suggested by the simplification in the first stage, a straightforward way looks for the solution of the nonnegative least squares problem in Equation (15).

$$\underset{\mathbf{M}\geq 0}{Minimize}\|\mathbf{G}-\mathbf{F}\mathbf{M}^T\|_F^2 \tag{15}$$

However, it is crucial to note that the solution **M** to subproblem 15 need *not* be symmetric (as **G** is not definitely within the space spanned by the column vectors of **F** due to the *non-negativity* constraint) which is a requirement for the simpler networks we explore with the original formulation 2. Hence we use Equation (16) to solve for the symmetric edge discovery subproblem. Note Equation (16) can be equivalently rewritten as Equation (17) in the form of standard nonnegative least squares problem, which can be efficiently solved by various existing solvers. We make use of the result that the structure of the Kronecker product of two identical **F** matrices in Equation (17) guarantees the solution **M** reshaped from $vec(\mathbf{M})$ to be symmetric. Note our nonnegative least squares solver for the core matrix **M** is different from the multiplicative update rule used in [6] for tri-factorizations.

Symmetric Edge Discovery Solver

$$\underset{\mathbf{M}\geq 0}{Minimize}\|\mathbf{X}-\mathbf{F}\mathbf{M}\mathbf{F}^T\|_F^2 \equiv \tag{16}$$

$$\underset{\mathbf{M}\geq 0}{Minimize}\|(\mathbf{F}\otimes\mathbf{F})vec(\mathbf{M})-vec(\mathbf{X})\|_2^2 \tag{17}$$

where $\otimes$ is the Kronecker product and $vec(\cdot)$ is the vectorization of a matrix (i.e. stack columns altogether).

4.3 Rescaling for Reproducibility of M and F.

Since the columns of **F** solved from Equation (5) are normalized, **M** solved from Equation (17) is affected by the sizes of the clusters. **M** is rescaled with Equation (18) such that **M** only reflects weighted correlations within/between nodes. Correspondingly, **F** needs to be rescaled with Equation (19) after rescaling **M** to make the maximum entry of each column in **F** equal to 1.

$$\mathbf{M}_{i,j} \leftarrow \mathbf{M}_{i,j} * (max(\mathbf{F}_{*i}) * max(\mathbf{F}_{*j})), \forall i,j \tag{18}$$

$$\mathbf{F}_{i,j} \leftarrow \frac{\mathbf{F}_{i,j}}{\underset{i}{max}(\mathbf{F}_{i,j})}, \forall i,j \tag{19}$$

In the experimental results for showing the relative strength of the correlations between and within the discovered nodes, **M** is further rescaled with Equation (20).

$$\mathbf{M}_{i,j} \leftarrow \frac{\mathbf{M}_{i,j}}{\underset{i,j}{max}(\mathbf{M}_{i,j})} \tag{20}$$

5 EXPERIMENTS

Previously we showed the limitations of existing ONMtF formulation (see Equation (1)) [6] for functional network discovery (Figure 2). Here we evaluate the performance of our network discovery method with experiments on both synthetic networks and networks generated from real fMRI spatiotemporal data. For the synthetic network experiments we use anatomical region structures from an anatomical atlas to define the functional nodes and background regions (Figure 4) but create artificial ground truth interactions (Figure 5). For real world data we compare results of finding networks in young, elderly and demented individuals and present typical results.

Setting Parameters. The matrix Θ (see Equation (3)) determines the penalty for placing two non-adjacent voxels in the same cluster. To use our regularizer we need to set two parameters: σ (the bandwidth of the kernel) and β (the regularizer weight). The parameter σ effectively specifies the diameter of continuous regions that our method discovers. The parameter β controls the tradeoff between the continuity regularization on the discovered nodes and how well our decomposition model reconstructs the original graph.
Using cross-validation we set $\sigma = 7$, $\beta = 40$. This parameter setting achieves the following: *1)* one continuous node with strong intra-cluster correlation will not be partitioned into two nodes, while two spatially separated nodes with strong inter-cluster correlation will not be clustered into one node; *2)* our method guarantees the continuity within each discovered node while achieving a reasonable reconstruction of the original graph with the decomposition model.

All MATLAB code to reproduce the experimental results in this paper will be available on `github` [2].

5.1 Synthetic Network Experiments With The Brain Atlas

Here we conduct experiments on different synthetic networks. The underlying regions we use are from the *Automated Anatomical Labeling (AAL) Atlas* (see Figure 4) and we create artificial associations between the nodes/regions to answer the following questions regarding the potentials of our method:
(1) How our method performs on a *simple network*, where the functional nodes of underlying network are both strongly self-correlated and strongly inter-correlated (see Figure 5a).
(2) How our method performs on a *degrading network*, where there is strong correlation within each individual node, but the inter-cluster correlation gradually decreases. This network is common in demented individuals (see Figure 5b).
(3) How our method performs on a degraded network with *local noise* in the form of weak correlations within each background (remaining) region (see Figure 5c). This is a setting that mirrors patients with the PTSD pathology.
(4) How our method performs when there are networks with *global noise*. Here there are not only strong inter-correlations between the set of functional nodes but also weak correlations between functional nodes and the background regions (see Figure 5d).

The synthetic networks $\mathbf{X}_{syn}$ are constructed from $\mathbf{F}_{syn}\mathbf{M}_{syn}\mathbf{F}_{syn}^T$ where $\mathbf{F}_{syn}$ is the indicator matrix for the segments of the synthetic brain activation network and $\mathbf{M}_{syn}$ describes the average intra/inter-cluster correlations of the segments. We now describe how we create the synthetic regions and their interactions.

[2] https://github.com/ZilongBai/KDD2017.git

Label	Correlation Definition
C_f	Within the same functional node
C_{ff}	Between different functional nodes
C_b	Within the same background region
C_{fb}	Between functional node and background region
C_{bb}	Between different background regions

Table 1: Definitions of different types of correlations for synthetic networks.

	C_f	C_{ff}	C_b	C_{fb}	C_{bb}
1.	1.0	1.0	0.0	0.0	0.0
2.	1.0	[0.1, 0.9]	0.0	0.0	0.0
3.	1.0	0.6	[0.1, 0.8]	0.0	0.0
4.	1.0	0.8	0.5	[0.1, 0.9]	0.0

Table 2: Specification of $\mathbf{M}_{syn}$, the synthetic correlation matrix, for different settings. *1. Simple Network, 2. Degrading Network, 3. Local Noise, 4. Global Noise.*

Creating Realistic But Synthetic Regions. The spatial region of the entire synthetic brain network is partitioned into eight major segments (based on the anatomical atlas) to synthesize the real brain activation as shown in Figure 4. Each segment is comprised of multiple spatially adjacent anatomical regions. The four segments on the first row of Figure 4 are spatially separated nodes that comprise the underlying functional network; the four segments on the second row of Figure 4 are remaining parts that serve to synthesize background regions.

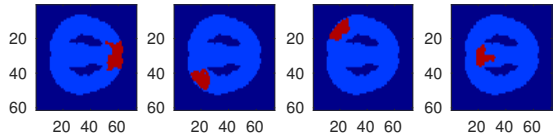

(a) Functional Nodes

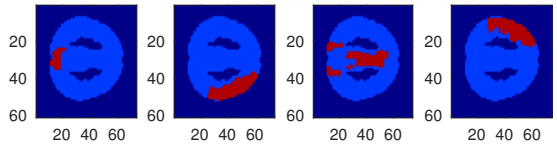

(b) Background Regions

Figure 4: Synthetic nodes/regions used in our experiments.

Creating Synthetic Activations. Figure 5 visually shows examples of synthetic node interaction matrices. Here we include precise details for reproduciblity of experiments, which are important as we show how we increase the challenge of the network discovery problem. To construct $\mathbf{M}_{syn}$ we define 5 major types of correlations which we label as $C_f, C_{ff}, C_b, C_{fb}, C_{bb}$. Here the suffix f means functional node and b means background region as shown in Figure 4. Details of their definitions are in Table 1. Specifically for the entires of $\mathbf{M}_{syn}$:

- $\mathbf{M}_{syn}(i,i) = C_f, 1 \le i \le 4$,
- $\mathbf{M}_{syn}(i,i) = C_b, 5 \le i \le 8$,
- $\mathbf{M}_{syn}(i,j) = C_{ff}, 1 \le i,j \le 4, i \ne j$,
- $\mathbf{M}_{syn}(i,j) = C_{bb}, 5 \le i,j \le 8, i \ne j$,
- $\mathbf{M}_{syn}(i,j) = C_{fb}, 1 \le i \le 4, 5 \le j \le 8$,
- $\mathbf{M}_{syn}(i,j) = C_{fb}, 5 \le i \le 8, 1 \le j \le 4$

A network is then defined by a combination of these different types of correlations and allows us to meaningfully test in a variety of situations. In our work we construct four synthetic settings defined formally in Table 2 and visually in Figure 5.

Results for Synthetic Networks With Our Method. Figure 6 and Figure 7 respectively show the discovered functional nodes $\mathbf{F}_{dis}$ and associations of nodes $\mathbf{M}_{dis}$ for each example ground truth network. $\mathbf{F}_{dis}$ should be compared with the ground truth functional nodes in Figure 4. Similarly, $\mathbf{M}_{dis}$ should be compared to the ground truth correlations between functional nodes in Figure 5.

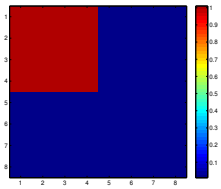

(a) Simple Network

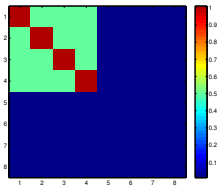

(b) Degrading Network

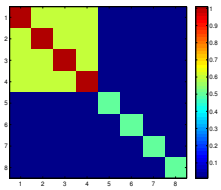

(c) Local Noise

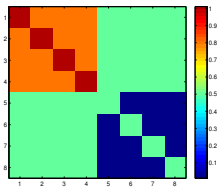

(d) Global Noise

Figure 5: Example synthetic correlation matrices $\mathbf{M}_{syn}$ of synthetic settings *1,2,3,4*. Each graph shows the edges of eight component networks (four functional nodes and four background regions from Figure 4). The *top left* 4×4 *patch of each graph shows the ground truth correlations of synthetic functional nodes our method attempts to recover.* The colors indicate the strength of each correlation.

5.2 Quantifying Node Discovery

Whilst in the previous section one can visually compare the ground truth (top row of Figure 4) with the recovered nodes in our four settings (Figure 6), here we quantify how good our method is at recovering the underlying nodes in various challenging situations.

To evaluate the performance of the proposed algorithm, we define a performance metric as follows:

Definition (Node Discovery Accuracy) Node Discovery Accuracy (NDA) of functional node i is defined as the maximum cosine similarity between all discovered nodes and the ground-truth functional

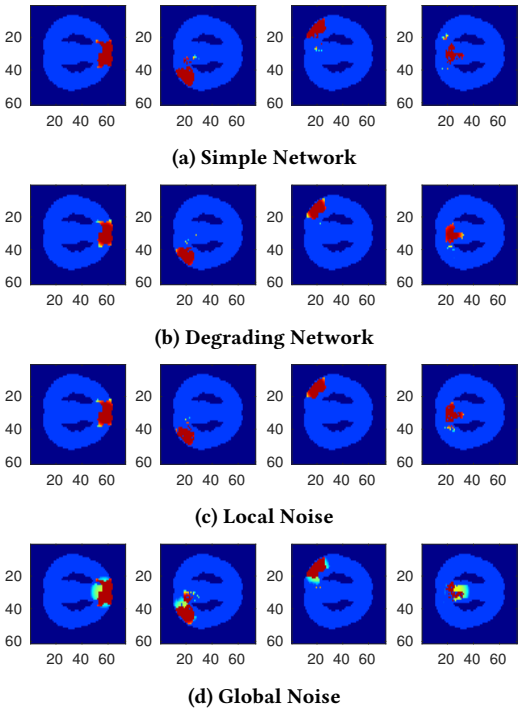

(a) Simple Network

(b) Degrading Network

(c) Local Noise

(d) Global Noise

Figure 6: Recovered functional nodes F_{dis} from the example synthetic networks in synthetic settings *1,2,3,4*. Compare with the top row of Figure 4, which is the ground truth. Our method recovers functional nodes with various complex shapes in the presence of different noises. The spatial properties of the ground-truth nodes are preserved accurately in the recovered nodes.

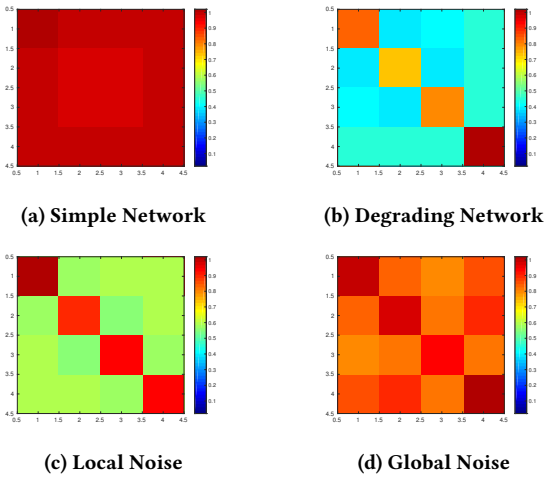

(a) Simple Network

(b) Degrading Network

(c) Local Noise

(d) Global Noise

Figure 7: Discovered interactions M_{dis} between functional nodes from the example synthetic networks of the four synthetic settings with our rank-4 factorization. *Compare with the top left* 4×4 *patches of the graphs from Figure 5 which are the ground truth edges between synthetic functional nodes. Our method performs well in recovering interactions between functional nodes in the face of different noises.*

node i (see Equation (21)).

$$accuracy(node_i) = \max_{1\le j\le r} \frac{\mathbf{F}_{dis(*,j)}^T \mathbf{F}_{syn(*,i)}}{\|\mathbf{F}_{dis(*,j)}\|_2 \|\mathbf{F}_{syn(*,i)}\|_2} \quad (21)$$

where F_{dis} denotes the discovered nodes of rank-r factorization; F_{syn} denotes the indicator matrix of ground-truth synthetic functional nodes/background regions as shown in Figure 4, with $F_{syn(*,i)}$, $1 \le i \le r$ for functional nodes. Cosine similarity enables NDA to provide fair comparison between different nodes without being affected by their sizes.

Figure 8 shows the NDA results of the four underlying functional nodes in *Degrading Network*, *Local Noise* and *Global Noise*. As one progresses along the x-axis in each graph the situations become more and more challenging. In *Degrading Network* (Figure 8a), the x-axis of this graph shows C_{ff} changing from 0.1 to 0.9. Whilst discovery of node 1,2, and 3 remain unaffected by the increasing strength of correlations between synthetic functional nodes, it becomes more challenging for our method to recover node 4 as its NDA drops significantly when C_{ff} increases to greater than 0.6. In the settings with *Local Noise* (Figure 8b), the x-axis represents C_b increasing from 0.1 to 0.8. It shows that our method is robust against the increasing correlation within each background region until $C_b \ge 0.8$ (note $C_{ff} = 0.6$ as specified in Table 2-*3*). It becomes notably challenging to accurately discover nodes in the presence of *Global Noise* (Figure 8c). The x-axis tells C_{fb} increasing from 0.1 to 0.9. Unlike Figures 8a and 8b, the NDA of node 3 drops to lower than 0.8 when $C_{fb} > 0.7$; the NDA of node 4 decreases to lower than 0.6 when $C_{fb} \ge 0.6$. Note node 4 completely fails to be discovered when $C_{fb} = 0.3$.

In summary, our approach is robust against various types of noise to certain extents. Comparing to other nodes, node 4 tends to be more susceptible to various noises and is particularly vulnerable to the *Global Noise*. This can be explained by the spatial properties of node 4: it has the most irregular shape amongst all synthetic functional nodes and is *surrounded* by background regions. The exceptional drop at $C_{fb} = 0.3$ in the presence of *Global Noise* (Figure 8c) can be explained by the difficulty of capturing the most appropriate solutions with 20 randomly initialized experiments for its particular shape, location, and size. We also witness that node 3 tends to suffer along with node 4 in the face of significant *Global Noise*. This can be explained by: *i)* the two synthetic functional nodes are closely located in spatial domain; *ii)* the space between them is filled up with some synthetic background regions.

5.3 Experiment with Real Brain Imaging Data

In these experiments we work on trying to discover the underlying networks in real fMRI data from ADNI [3]. The matrix $\mathbf{X}$ in Equation (5) is constructed by calculating the absolute Pearson correlation between each pair of voxels based on their temporal activations and hence represents a correlation/similarity matrix. In this paper for ease of visualization, we focus on the activation network on

[3] http://adni.loni.usc.edu/study-design/collaborative-studies/dod-adni/

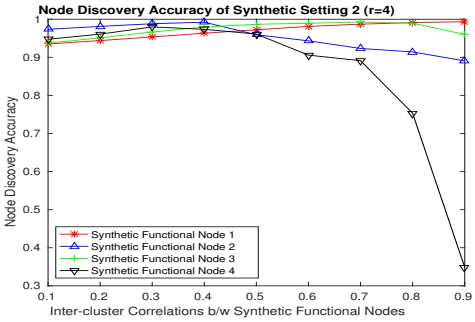

(a) Degrading Network

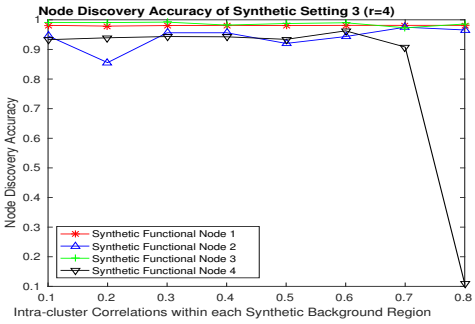

(b) Local Noise

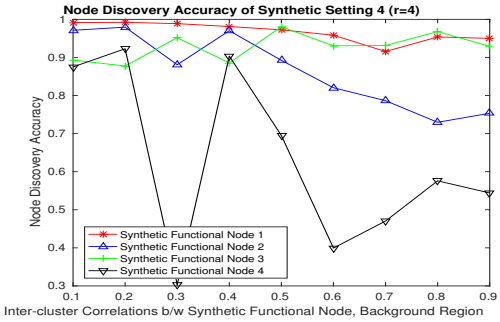

(c) Global Noise

Figure 8: Node discovery accuracy measured by Equation (21) for different synthetic networks. The higher the NDA, the better the node discovery performance. In each graph the network discovery difficulty increases along the x-axis from left to right. The interactions (correlations) of functional nodes and background regions of each graph are specified in Table 2. Each synthetic setting is tested with 20 randomly initialized experiments and the tri-factorization solution with the lowest relative reconstruction error is chosen.

the 36*th* slice of each 3-dimensional scan image. The data contains three cohorts of individuals: **y** young normal subjects, **o** old normal subjects and **d** old demented (Alzheimer affected) subjects at resting state. At resting state it is well known that even the most diseased individual will exhibit the well known default mode network (DMN) (Figure 9). We expect the strength of the DMN to diminish from **y** to **o** to **d** and for parts of the DMN to become missing for demented individuals. Figures 10, 11 and 12 illustrate the network

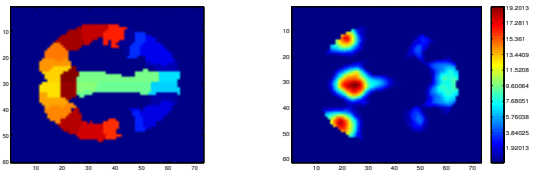

Figure 9: Left: Anatomical structures/nodes in the brain at slice 36. Different colors represent different nodes. Right: The subset of structures that form the default mode network (DMN). This network activates while a person is in resting state.

discovered by our method using a rank-6 factorization. For each of the cohorts we present the results of one typical subject and the average network discovery result. Similar to our synthetic experiments, we conduct multiple experiments with random initialization on the same subject and select the tri-factorization solution with the lowest relative reconstruction error. The results from experiments on real fMRI data demonstrate the capability of our method to reproduce the well known results in neurology [7] regarding the default mode network in resting-state subjects. Figure 10 shows a strong and completely intact DMN for a young healthy person as expected whilst Figure 11 shows for a healthy elderly subject a mostly intact but weakened network. Most importantly, Figure 12 shows that a demented individual's network is not intact in that the front cingular is missing from the network. This behavior has been observed previously [13]. The percentage amounts listed in part (c) of Figures 10, 11 and 12 are significant. What they indicate is that our method was able to simplify the network into regions

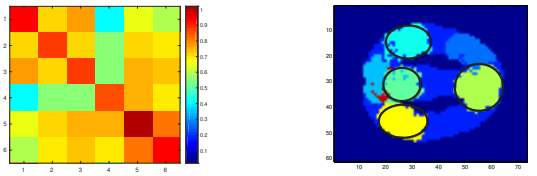

(a) Discovered edges (b) Discovered Network

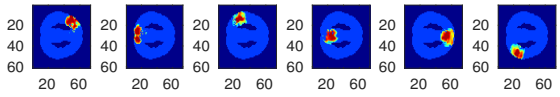

(c) Discovered nodes and corresponding largest covered anatomical regions (cosine similarity) from left to right: Frontal Inf. Oper. R. (59.12%), Cuneus L. (59.15%), Precentral L. (66.61%), Precuneus R. (48.45%), SupraMarginal L. (61.10%), Frontal Sup. Medial. L. (50.56%).

Figure 10: Network discovered from a typical young normal subject. Normal activities are discovered from the *Parietal and Occipital Lobes* (ice blue and lemon regions in (b)), the *Cingulum Region Posterior* (sea foam region in (b)), the *Frontal Cortex* (lime region in (b)).

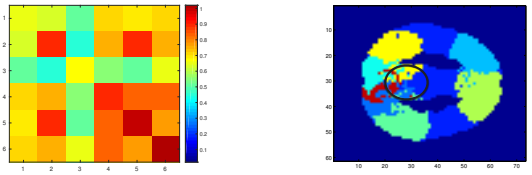

(a) Discovered edges (b) Discovered Network

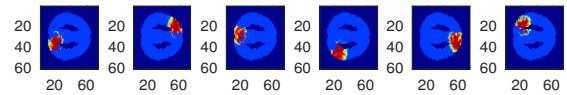

(c) Discovered nodes and corresponding largest covered anatomical regions (cosine similarity) from left to right: Precuneus R. (28.49%), Occipital Mid. R. (61.59%), Angular L. (56.42%), SupraMarginal R. (59.52%), Frontal Sup. Medial. L. (44.60%), Occipital Mid. L. (51.63%).

Figure 11: Network discovered from a typical elderly normal subject. The *Parietal and Occipital Lobes* and *Frontal Cortex* regions are discovered as seen in the typical young normal scan. The region within black oval in (b) shows some activity on the edge of *Cingulum Region Posterior* but not as much as in Alzheimer affected patient, which likely indicates normal slowed activity that comes with normal aging.

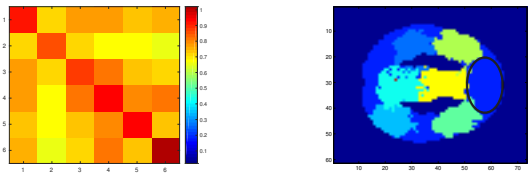

(a) Discovered edges (b) Discovered Network

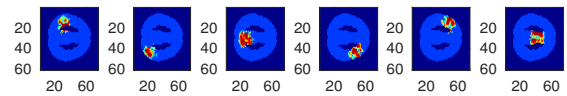

(c) Discovered nodes and corresponding largest covered anatomical regions (cosine similarity) left to right: SupraMarginal R. (23.28%), Angular L. (63.27%), Precuneus L. (51.63%), Frontal Inf. Oper. L. (48.40%), Frontal Inf. Oper. R. (54.92%), Cingulum Mid. L. (40.46%).

Figure 12: Network discovered from a typical elderly demented subject (Alzheimer affected individual). The *Occipital and Angular* are discovered with retarded activity levels likely due to Alzheimer's disease. The region within black oval in (b) shows that *Frontal Cortex* is not discovered as a functional node; hypo-activity of this region shows possible lack of decision making or executive function likely due to Alzheimer's disease. The *Cingulum Region Anterior* in lemon is discovered as a functional node, showing the relative overactivity of cingulum that may manifest as depression. Depression is common in Alzheimer's disease affected patients.

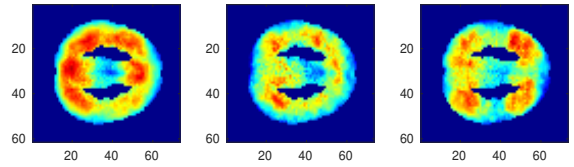

Figure 13: Average network discovery result of each cohort. Left: young normal, 40 fMRI scans. Middle: elderly normal, 19 fMRI scans. Right: elderly demented (Alzheimer affected), 21 fMRI scans.

that matched known anatomical regions **despite not** being given those regions, that is without any supervision.

6 CONCLUSION

In this paper, we pose the problem of unsupervised network discovery involving the discovery of regions with cohesive behavior (the nodes) and the relationship between those regions (the edges) from spatiotemporal data. Though previous work has solved the problem of *either* node discovery *or* edge discovery, to our knowledge no method has been proposed to simultaneously solve both without supervision. We formulate the unsupervised network discovery problem as a nonnegative matrix tri-factorization with orthogonality constraints and a novel spatial continuity regularizer to ensure easily interpretable and useful results. We develop effective solvers for the formulation with *i)* novel update rules for node discovery, *ii)* nonnegative least squares solver for edge discovery. We explored the strength of our methods with extensive experiments on a variety of synthetic but realistic networks and on real world datasets. On synthetic data our method demonstrates robustness in the presence of different kinds of background noise. For real world data, we are not able to present all individuals' results due to space, but we show an average result and the result of a typical subject for each cohort. These results are consistent with known earlier work [2] and demonstrate the significance of our method: it has the potential to effectively explore new networks in scenarios where specific domain knowledge is unavailable or insufficient. This model can be applied to spatiotemporal data from other domains if their node discovery requires geometrical continuity within each node.

ACKNOWLEDGMENTS

This research is supported by NSF Grant IIS-1422218 'Functional Network Discovery for Brain Connectivity' and The Henry M. Jackson Foundation Grant 'Small World and Other Graph Properties in Brains'. Peter B. Walker is a military service member. This work was prepared as part of his official duties. Title 17 U.S.C. 101 defines U.S. Government work as a work prepared by a military service member or employee of the U.S. Government as part of that person's official duties. The opinions of the authors do not necessarily reflect those of the United States.

REFERENCES

[1] Dan A Alcantara, Owen Carmichael, Will Harcourt-Smith, Kirstin Sterner, Stephen R Frost, Rebecca Dutton, Paul Thompson, Eric Delson, and Nina Amenta. 2009. Exploration of shape variation using localized components analysis. *IEEE transactions on pattern analysis and machine intelligence* 31, 8 (2009), 1510–1516.

[2] Jessica R Andrews-Hanna, Jay S Reidler, Jorge Sepulcre, Renee Poulin, and Randy L Buckner. 2010. Functional-anatomic fractionation of the brain's default network. *Neuron* 65, 4 (2010), 550–562.

[3] John Burge, Terran Lane, Hamilton Link, Shibin Qiu, and Vincent P Clark. 2009. Discrete dynamic Bayesian network analysis of fMRI data. *Human brain mapping* 30, 1 (2009), 122–137.

[4] Ian Davidson, Sean Gilpin, Owen Carmichael, and Peter Walker. 2013. Network Discovery via Constrained Tensor Analysis of fMRI Data *(KDD '13)*. ACM, New York, NY, USA, 194–202.

[5] Federico De Martino, Francesco Gentile, Fabrizio Esposito, Marco Balsi, Francesco Di Salle, Rainer Goebel, and Elia Formisano. 2007. Classification of fMRI independent components using IC-fingerprints and support vector machine classifiers. *Neuroimage* 34, 1 (2007), 177–194.

[6] Chris Ding, Tao Li, Wei Peng, and Haesun Park. 2006. Orthogonal nonnegative matrix t-factorizations for clustering *(KDD '06)*. ACM, 126–135.

[7] Christopher R Genovese, Nicole A Lazar, and Thomas Nichols. 2002. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage* 15, 4 (2002), 870–878.

[8] Forrest M Hoffman, William W Hargrove Jr, David J Erickson III, and Robert J Oglesby. 2005. Using clustered climate regimes to analyze and compare predictions from fully coupled general circulation models. *Earth Interactions* 9, 10 (2005), 1–27.

[9] Keigo Kimura, Yuzuru Tanaka, and Mineichi Kudo. 2014. A Fast Hierarchical Alternating Least Squares Algorithm for Orthogonal Nonnegative Matrix Factorization.. In *ACML*.

[10] Chia-Tung Kuo, James Bailey, and Ian Davidson. *A Framework for Simplifying Trip Data into Networks via Coupled Matrix Factorization.* Chapter 83, 739–747.

[11] Pierre-Jean Lahaye, Jean-Baptiste Poline, Guillaume Flandin, Silke Dodel, and Line Garnero. 2003. Functional connectivity: studying nonlinear, delayed interactions between BOLD signals. *Neuroimage* 20, 2 (2003), 962–974.

[12] Marc Postman and P Murdin. 2001. Distribution of Galaxies, Clusters, and Superclusters. *Encyclopedia of Astronomy and Astrophysics* (2001).

[13] Marcus E Raichle and Abraham Z Snyder. 2007. A default mode of brain function: a brief history of an evolving idea. *Neuroimage* 37, 4 (2007), 1083–1090.

[14] Marian Scott, Claire Miller, Francesco Finazzi, and Ruth Haggarty. 2013. Coherency in space of lake and river temperature and water quality records. (2013).

[15] Liang Sun, Rinkal Patel, Jun Liu, Kewei Chen, Teresa Wu, Jing Li, Eric Reiman, and Jieping Ye. 2009. Mining brain region connectivity for alzheimer's disease study via sparse inverse covariance estimation *(KDD '09)*. ACM, 1335–1344.

[16] Nathalie Tzourio-Mazoyer, Brigitte Landeau, Dimitri Papathanassiou, Fabrice Crivello, Olivier Etard, Nicolas Delcroix, Bernard Mazoyer, and Marc Joliot. 2002. Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *Neuroimage* 15, 1 (2002), 273–289.

[17] Vincent G van de Ven, Elia Formisano, David Prvulovic, Christian H Roeder, and David EJ Linden. 2004. Functional connectivity as revealed by spatial independent component analysis of fMRI measurements during rest. *Human brain mapping* 22, 3 (2004), 165–178.

[18] Peter B Walker and Ian N Davidson. 2011. Exploring new methodologies for the analysis of functional magnetic resonance imaging (fMRI) following closed-head injuries. In *Foundations of Augmented Cognition. Directing the Future of Adaptive Systems*. Springer, 120–128.