Applying joint graph embedding to study Alzheimer's neurodegeneration patterns in volumetric data

Rosemary He[1,2], Daniel Tward[2,3], and for the Alzheimer's Disease Neuroimaging Initiative*

[1]Department of Computer Science, University of California, Los Angeles, CA, 90095, United States
[2]Department of Computational Medicine, University of California, Los Angeles, CA, 90095, United States
[3]Ahmanson-Lovelace Brain Mapping Center, University of California, Los Angeles, CA, 90095, United States

Abstract

Neurodegeneration measured through volumetry in MRI is recognized as a potential Alzheimer's Disease (AD) biomarker, but its utility is limited by lack of specificity. Quantifying spatial patterns of neurodegeneration on a whole brain scale rather than locally may help address this. In this work, we turn to network based analyses and extend a graph embedding algorithm to study morphometric connectivity from volume-change correlations measured with structural MRI on the timescale of years. We model our data with the multiple random eigengraphs framework, as well as modify and implement a multigraph embedding algorithm proposed earlier to estimate a low dimensional embedding of the networks. Our version of the algorithm guarantees meaningful finite-sample results and estimates maximum likelihood edge probabilities from population-specific network modes and subject-specific loadings. Furthermore, we propose and implement a novel statistical testing procedure to analyze group differences after accounting for confounders and locate significant structures during AD neurodegeneration. Family-wise error rate (FWER) is controlled at 5% using permutation testing on the maximum statistic. We show that results from our analysis reveal networks dominated by known structures associated to AD neurodegeneration, indicating the framework has promise for studying AD. Furthermore, we find network-structure tuples that are not found with traditional methods in the field.

Keywords

Alzheimer's disease, structural MRI, graph embedding, network analysis, familywise error rate control

1 Introduction

Alzheimer's disease (AD) is a progressive mental disorder associated with neurodegeneration that generally occurs in old ages. It is one of the most common diseases in seniors, killing more than breast cancer

*Data used in preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database `adni.loni.usc.edu`). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: `http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf`

and prostate cancer combined[3]. However, AD can only be formally diagnosed through an autopsy after a patient is deceased, encouraging the research of alternative proxies for AD diagnosis. There are currently three classes of potential biomarkers that could offer useful alternatives to assess AD diagnosis: $\beta-$amyloid(A), tau(T) and biomarkers for neurodegeneration or neuronal injury(N)[13], with the last class being the focus of this paper. Neurodegeneration can be measured noninvasively, such as through structural MRI, but is not specific to AD as it can reflect other diseases[15]. For example, atrophy is a biomarker that is often associated with AD, but also occurs in a variety of disorders such as epilepsy and anoxia[15]. The lack of specificity of these neurodegeneration biomarkers poses a major limitation in their utility for early-stage clinical diagnosis of AD. A promising approach to improving specificity is to consider patterns of volumetric changes across the whole brain, rather than focusing on a small number of regions. For example, in discussing the diagnosis of Mild Cognitive Imparment due to AD, Albert et al. suggests the possibility of biomarkers describing "complex patterns of tissue loss" through "data driven statistical approaches in which many different brain regions are evaluated simultaneously" [2].

While rare in volumetric analysis, such patterns have been studied extensively in functional brain imaging. One popular approach to studying brain connectomics is to extract information about interactions between volumetric pixels (voxels) from time-series functional magnetic resonance imaging (fMRI) data collected over a timescale of weeks, months or years[9]. Functional connectivity, in particular, studies temporal dependencies among anatomically separated regions[24]. There are several approaches to study functional connectivity in fMRI data. As our interests lie within the whole brain instead of a single voxel, we discuss only multivoxel pattern analysis methods that study networks as a whole[19]. A traditional method in the field is seed-based analysis, in which a region of interest (ROI) is selected, and all voxels correlated to the ROI is identified[10]. For example, in an earlier work Biswal et al. studied the motor cortex to identify the sensorimotor network[6]. Unsupervised clustering methods including k-means, hierarchical and graph-based methods do not require a priori ROI and group voxels together by their similarities in time series data[16]. In their paper, Lee et al. uses fuzzy-c-means clustering to identify resting state networks[17].

Over the years, graph-based approaches have gained popularity in studying functional connectivity. To convert a brain into a graph, the regions are modeled as nodes and connections between regions as edges. Under this model, one can construct a matrix of all pairs of connections in the brain, known as the functional connectome[12], of which decomposition or embedding methods can be applied to uncover latent variables. Independent component analysis decomposes data into linearly independent components, grouping brain regions into networks based on their voxel activation correlations[8]. Non-negative matrix factorization is a dimensionality reduction method that forces non-negativity constraints on the components[16]. Some popular embedding methods include Adjacency Spectral Embedding (ASE)[23], which embeds a single symmetric adjacency matrix using eigenvectors corresponding to the largest eigenvalues, and Laplacian Eigenmap (LE)[4], which embeds a single graph-Laplacian matrix using its eigenvectors corresponding to the smallest nonzero eigenvalues. However, several limitations lie within these graph-based approaches. First, they embed one graph at a time, and combining individual embeddings across multiple graphs is not a straighforward task. Second, the results are difficult to interpret, and further analysis is required[28].

One technique for embedding multiple graphs at once is omnibus embedding, in which the matrices of multiple graphs are combined into one and embedding is done on the big matrix[18]. However, the combined matrix is usually very large and require lots of computational power. Dictionary learning is another framework for uncovering low-dimensional embeddings across multiple graphs, allowing for group comparison. Drawing upon clustering and linear decomposition methods, this method allows for additional constraints to achieve better formed solutions[1]. In the work of D'Souza et al.[11], they use a dictionary learning method to model interactions between resting state functional MRI and behavioral data in Autism Spectrum Disorder. Their method finds shared dictionary elements across multiple graphs and a subject specific loading onto the elements, which are then used as inputs to a neural network for disease prediction[11]. In the work of Wang et al.[25], they propose a joint graph embedding method

to estimate a low dimensional embedding across multiple graphs and each graph's projection onto that embedding, which we will discuss more in detail in section 2.2.

In this work, we shift our attention away from functional connectivity and propose a graph-based approach to study neurodegeneration using volumetric data. To our knowledge, this is the first time joint graph embeddings have been used to study volumetric brain data. Similar to methods above, we model the brain as a graph where a node represents a structure of interest and an edge represents a correlation in atrophy patterns between two structures[27]. We then use a multigraph embedding technique to try and understand these patterns and uncover potential biomarkers. By applying multigraph embedding to study neurodegeneration, which spans a much longer timescale than fMRI measures, we hope to increase specificity of our results. In addition, we hope to increase the specificity of neurodegeneration biomarkers by modeling the dataset with a more complex pattern than existing approaches. Rather than looking at volume changes in each structure individually in the traditional mass univariate method[22], we add complexity by characterizing pair-wise relationships between structures in the context of uncovered networks.

2 Material and methods

2.1 Data preprocessing

We obtain our data from Alzheimer's Disease Neuroimaging Initiative (ADNI) database (`adni.loni.usc.edu`). The ADNI was launched in 2003 as a public-private partnership, led by Principal Investigator Michael W. Weiner,MD. The primary goal of ADNI has been to test whether serial magnetic resonance imaging (MRI), positron emission tomography (PET), other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of mild cognitive impairment (MCI) and early Alzheimer's disease (AD). Specifically, we took the ADNI1 3Y1.5T Longitudinal FreeSurfer dataset[26] prepared by University of California, San Francisco, comprised of 699 individuals in total. We selected 108 regions of interest common in studying neurodegeneration, excluding non-brain and whole-brain structures as we are interested in structures on a smaller scale. We selected a cohort in which each individual has at least 3 visits during the span of 3 years. For subjects with missing volumes, we forward filled in time. We model each individual's brain as a graph, where each anatomical structure is a node and an edge exists between two nodes if they are highly correlated during neurodegeneration. For each individual in the cohort, we converted volumetric data into a correlation matrix of size 108 by 108 and then an adjacency matrix based on a threshold of 0.8 by absolute value, as shown in Figure 1. For defining disease groups, the Clinical Dementia Rating (CDR) is referenced, which consists of 5 levels: 0 (None), 0.5 (Questionable), 1 (Mild), 2 (Moderate), and 3 (Severe)[20]. We decided a threshold of $<=1$ to separate the group into none/mild and severe cognitive impairment. After filtering and pre-processing, the cohort contains 494 individuals, specifically the none/mild group with 322 and the severe group with 172.

2.2 Multiple Random Eigengraphs and Joint Graph Embedding

We first review the mathematical model and original joint graph embedding algorithm proposed by Wang et al.[25], of which our algorithm is based on. In this work, we refer to a random graph as a graph in which the edges are generated under a probability distribution. The Multiple Random Eigengraphs (MREG) is a mathematical framework modeling multiple random graphs[25]. Consider a set of m unweighted and undirected graphs with the same n vertices denoted by $\{G_s=(V_s,E_s)\}_{s=1}^m$. Let $h_1,...,h_d$ be normalized vectors in $\mathbb{R}^n$ that span a subspace of dimension d, contributing to a large amount of variability in the set of graphs, and $\lambda_1,...,\lambda_d$ be vectors in $\mathbb{R}^m$ such that $\sum_{k=1}^d \lambda_k[s]h_k h_k^T \in [0,1]^{n\times n}$ for all λ[25] for subject

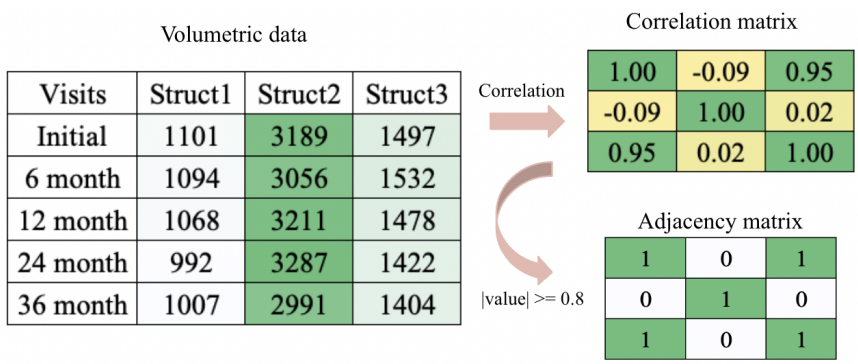

Figure 1: An example of transformation for one individual's selected structures

s. Then the adjacency matrix A_s for each graph G_s should be modeled as follows[25]:

$$A_s[i,j] \sim Bernoulli\left(\sum_{k=1}^{d} \lambda_k[s] h_k[i] h_k[j]\right), \text{ for } i \geq j \tag{1}$$

$$A_s[i,j] = A_s[j,i], \text{ for } i < j \tag{2}$$

Note here each A_s is symmetric, as opposed to sampling independently above and below the diagonal. The h vectors span the latent subspace shared by the set of multiple graphs, and the λ vectors represent graph-specific loadings onto the subspace[25].

The original joint graph embedding algorithm by Wang et al.[25] estimates a low dimensional embedding of the latent space across multiple graphs and each graph's projection onto that embedding under the MREG model. It estimates the subspace by minimizing the sum of squared errors (SSE) between the subspace and adjacency matrices.[25]. In this work, we implemented a modified version of the Wang et al. algorithm as discussed below.

2.3 Limitations of the original framework

There are several limitations to consider when applying the original framework and algorithm above to neurodegeneration data, which we will address and modify in our version:

1. The constraint that probabilities lie in [0,1] is difficult to enforce in practice
2. Incorrect estimation of diagonal entries contributes a negligible amount of error when the number of vertices is large, but contributes significantly in our case.
3. The SSE loss function does not correspond to a log likelihood under the proposed Bernoulli model, and therefore resulting parameter estimates do not have desirable properties of maximum likelihood estimators

Now we state our main contribution and novelty in this work. First, we address the three limitations listed above by (1) introducing a sigmoid function to the model to guarantee edge probabilities are in [0,1], (2) adding constraints on the diagonal such that parameter estimates are not forced to fit diagonal entries that carry no meaning (since diagonal entries of a correlation matrix are always 1), and (3) extending the original algorithm[25] to identify maximum likelihood estimators by gradient descent. With these modifications, our embedding algorithm generates probabilities suitable for a Bernoulli model, which we will describe more in detail below. Secondly, we develop and implement a novel statistical

testing framework to detect complex patterns including network-structure pairs and triples, rather than a machine learning classifier. To our knowledge, our approach to testing patterns has not been performed to analyze joint graph embedding results on brain imaging data before.

2.4 Modifications to MREG

In this section, we precisely state our modifications to the MREG model in Wang et al.[25]. Consider the set of m unweighted and undirected graphs with the same n vertices denoted by $\{G_s=(V_s,E_s)\}_{s=1}^m$, where a vertex represents a brain structure of our interest and an edge represents a strong correlation between structures. We modify the interpretation of h and state that the h vectors now span a space of parameters that encode the probability when acted on by a sigmoid function. This sigmoid function guarantees $sigmoid(\sum_{k=1}^d \lambda_k[s]h_k h_k^T) \in [0,1]^{n\times n}$ even in the case of small samples, addressing limitation 1. Secondly, we force diagonal entries to be 1 since a structure's relation with itself is not of interest in this work, addressing limitation 2. Our modifications to equation 1 and 2 are as follows:

$$A_s[i,j] \sim Bernoulli\left(sigmoid\left(\sum_{k=1}^{d} \lambda_k[s]h_k[i]h_k[j]\right)\right), \text{ for } i>j \tag{3}$$

$$A_s[i,j] = A_s[j,i], \text{ for } i<j \tag{4}$$

$$A_s[i,j] = 1, \text{ for } i=j \tag{5}$$

$$sigmoid(x) = \frac{1}{(1+\exp(-x))} \tag{6}$$

An example transformation from the latent subspace and subject specific loading to adjacency matrices observed in our data is illustrated in Figure 2.

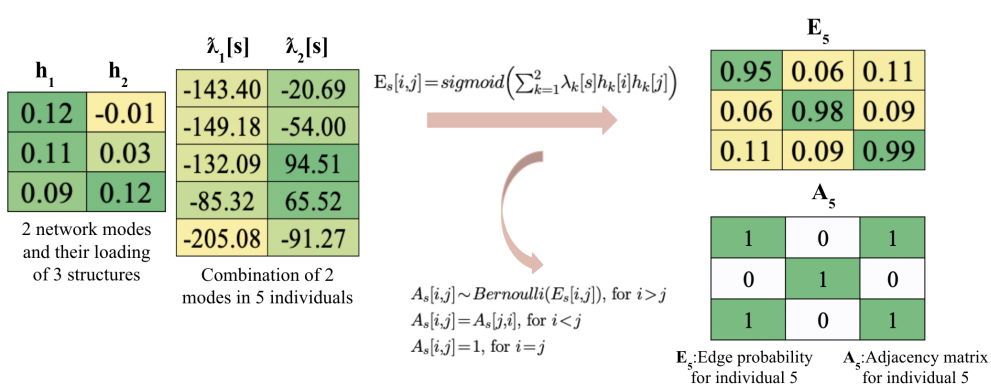

Figure 2: An example of MREG for the same individual and structure in Figure 1

2.5 Our Maximum Likelihood Joint Graph Embedding Algorithm

We change the algorithm to better fit the constraints of our study and estimate the subspace by minimizing the binary cross entropy (BCE) loss, which is the log likelihood under the model. Let $E_s[i,j]$ be a symmetric matrix representing edge probabilities between structures i and j for subject s, then the BCE loss is minimized as in equation (7), where A_s are the observed adjacency matrices:

$$\underset{\lambda,h}{\text{argmin}} -\frac{1}{m}\sum_{s=1}^{m}\sum_{i=1}^{n}\sum_{j=i+1}^{n} A_s[i,j]\log(E_s[i,j]) + (1-A_s[i,j])\log(1-E_s[i,j]) \tag{7}$$

Note here that the diagonal terms do not contribute to our loss function. By replacing the original loss function with BCE, we will have a well characterized maximum likelihood estimator even in the small sample case, addressing limitation 3. We take a greedy approach in finding the optimal representation of the latent space, where we start with a 1D optimization problem, and then expand to the second dimension while keeping the first dimension fixed to find the optimal representation. The algorithm is implemented as follows:

$\lambda \leftarrow$ random.normal((m, number of networks)), $h \leftarrow$ random.normal((n, number of networks)), normalize h

for $d=1$; $d<=$number of networks; $d++$ **do**

  while loss not coverged in dimension d **do**

    $E = sigmoid(\sum_{k=1}^{d} multiply(multiply(\lambda[:,:k],h[:,:k]),transpose(h[:,:k]))$

    Compute BCE loss as in equation (7)

    Compute gradient of the BCE with respect to λ and h

    Update $\lambda[:,d]$ by gradient descent, and $h[:,d]$ by projected gradient descent (i.e. update and normalize)

  end while

end for

We coded in python and PyTorch, and used automatically calculated gradients for gradient descent. The algorithm took 30 minutes to run on a computer with 10-core CPU. Convergence in all 4 dimensions is shown in Figure 3a, where at every 10,000 iterations we observe the loss dropping quickly after adding another dimension, and converging before the next dimension is added.

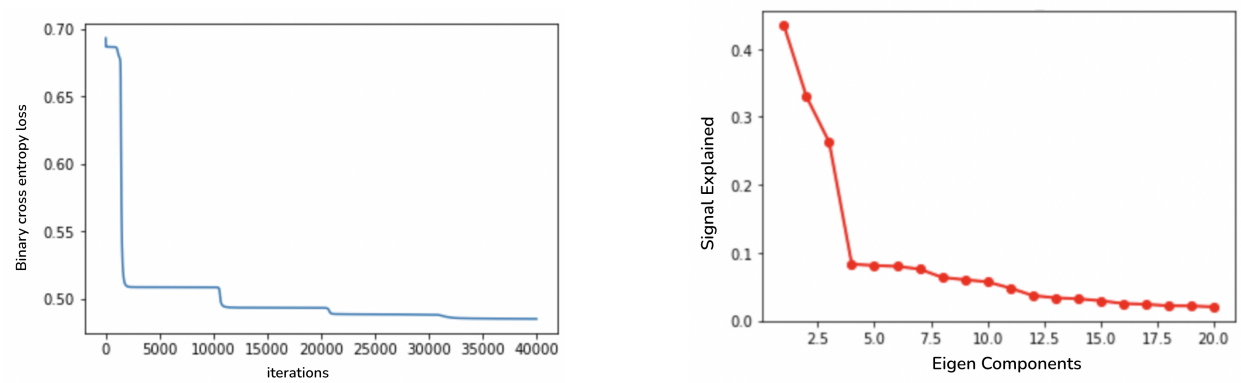

(a) **(b)**

Figure 3: a). Convergence in all 4 dimensions with the algorithm. b). Scree plot with elbow at 4 dimensions.

2.6 Hyperparameter Tuning

For hyperparameter tuning, we first estimate an optimal dimension of the latent space to account for the majority of the variation in observed correlation matrices. To do this, we computed an eigendecomposition for each subject's correlation matrix, and found that 4 components provided a reasonable reconstruction accuracy. A scree plot for one typical subject is shown in Figure 3b. Next, we performed a grid search over several orders of magnitude for the gradient descent step sizes corresponding to h and λ, and selected the largest parameters that gave convergence without oscillation. We decided to set the step size for h to 2 and the step size for λ to 20000.

2.7 Significance Testing

We develop and implement a novel test statistic similar to F-type statistics (i.e. comparing sum of square error under two different models) to test for differences in networks and structures between the two groups, and use permutation testing on a maximum statistic to control for FWER at 5% [21]. In addition, we compare the difference between groups after accounting for confounders: age, intracranial volume and APOE gene status by least squares regression. APOE gene status was modeled as a categorical (as opposed to cardinal) variable, where an individual may have 0, 1, or 2 copies. For a cohort with m subjects, d networks, and n structures, we define the test score as $\text{Score}_{sij} = \sum_k \lambda_k[s] h_k[i] h_k[j]$ for subject s and structures i and j. We pass the score through a sigmoid to give probabilities E, and estimate these variables by maximum likelihood. We form the tensor $T_{skij} = \lambda_k[s] h_k[i] h_k[j]$ to get a three-tuple for each subject s, of which we will perform statistical tests by comparing how close it is to its group-dependent or group-combined average, after accounting for confounders. By summing over various combinations of indices in T before statistical testing, we are able to test for:

1. networks only
2. structures only
3. network-structure pairs
4. structure-structure pairs
5. network-structure-structure triples

As other standard tests for items 2 and 4 exist, we will focus our work on 1, 3, and 5. To our knowledge, our framework for testing patterns involving networks, pairs, or triples has not been performed to analyze brain imaging data before.

2.7.1 Confounder regression analysis

We perform least square regression analysis to test for true signals not caused by common confounders for AD and neurodegeneration. We start with a design matrix D containing a column of 1s (for mean), and columns for age, intracranial volume and APOE gene status, and estimate a coefficient matrix $\hat{C}$ for confounders. From $\hat{C}$ the SSE for one group is calculated as follows:

$$\hat{C}^{combined} = (D^T D)^{-1} D^T T \tag{8}$$

$$\hat{T}^{combined} = \hat{C}^{combined} D \tag{9}$$

Here T is reshaped into a matrix from a tensor for calculation. Next, we calculate the SSE for splitting the cohort into two groups: none/mild and severe AD. We first form D', which has an additional column to D indicating disease status. We then perform the same regression, with D' replacing D in equations 8 and 9 to find $\hat{T}^{twogroups}$. In the next sections, we will use $\hat{T}^{combined}$ and $\hat{T}^{twogroups}$ to calculate the sum of square errors (SSE) when considering one vs. two groups.

2.7.2 Networks

We start by testing for significant networks and obtain the test statistic X_k for network k by reducing (i.e. summing over) additional dimensions. Let $\hat{T}^g_{k,i,j}$ be the expected value matrix under our linear model for dimension k, structures i and j, and subjects in group g. We remove extra dimensions by calculating the SSE between one group vs. two groups and X_k as follows:

$$SSE_k^g = \sum_{s \in g,i,j} (T_{skij} - \hat{T}_{kij}^g)^2 \ , \ g \in \{\text{combined, twogroups}\} \tag{10}$$

$$X_k = SSE_k^{combined} - SSE_k^{twogroup} \tag{11}$$

To control the FWER at 5%, we use permutation testing and take the max over k at each iteration and define the threshold as the 95 percentile of 10,000 simulation results[21].

2.7.3 Network structure pairs

Similar to networks only, we calculate the SSE in the two settings but reduce one fewer level as follows:

$$SSE_{k,i}^g = \sum_{s \in g,j} (T_{skij} - \hat{T}_{kij}^g)^2 \ , \ g \in \{\text{combined, twogroups}\} \tag{12}$$

$$X_{k,i} = SSE_{k,i}^{combined} - SSE_{k,i}^{twogroup} \tag{13}$$

We follow the same permutation testing procedure and take a max over k,i instead of just k at each iteration.

2.7.4 Network structure structure triples

Lastly for triples, we form the test statistic:

$$SSE_{k,i,j}^g = \sum_{s \in g} (T_{skij} - \hat{T}_{kij}^g)^2 \ , \ g \in \{\text{combined, twogroups}\} \tag{14}$$

$$X_{k,i,j} = SSE_{k,i,j}^{combined} - SSE_{k,i,j}^{twogroup} \tag{15}$$

We follow the same permutation testing procedure and take a max over k,i,j at each iteration.

In our experiments, we performed the analysis twice: with and without adjusting for confounders. Note that not adjusting for confounders is a special case of T estimation, where $\hat{T}^g$ is simply the mean over all subjects in group g. In the results section, we will present the former, and make brief comparisons to the latter.

2.8 Code availability

Our source code and documentation are available on GitHub at `https://github.com/twardlab/structural_connectivity_AD`. To replicate our work, first run the Jupyter notebook *preprocess*, and then the notebook *joint_graph_embedding_analysis*. We document our code with Sphinx[7], and save documentations in *docs*.

3 Results

3.1 Significant networks associated with AD neurodegeneration

Out of the 4 networks identified from joint graph embedding, we found the first 2 extremely significant, after accounting for confounders. Upon examination in Table 1, both networks are dominated by structures believed to be associated with AD neurodegeneration. We show the results from permutation testing and the statistics for each network in Figure 4 and visualize the structures in network 1 and 2 in Figure 5. In Table 1, we include the top 10 structure by absolute value of loading in each network.

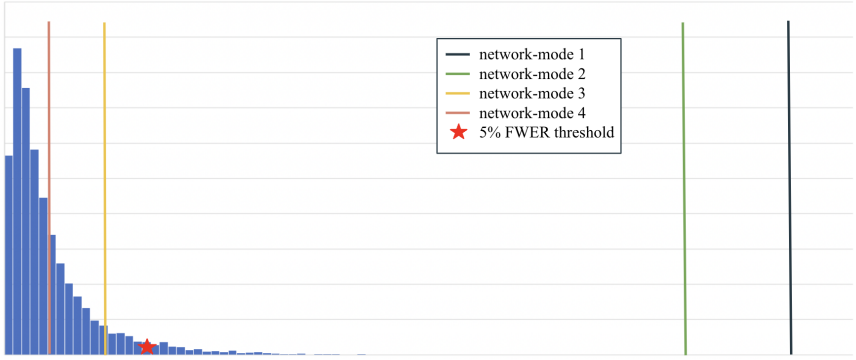

Figure 4: Histogram of permutation testing after accounting for confounders

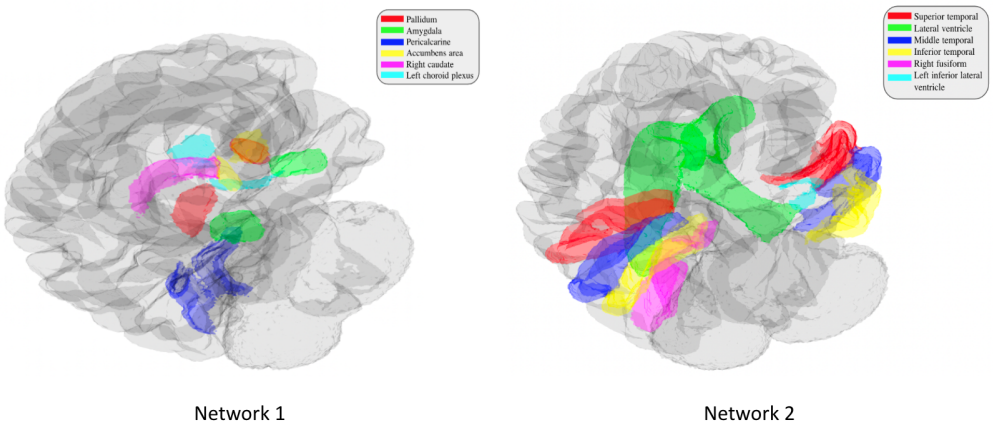

Figure 5: Visualization of significant networks

3.2 Significant network structure pairs associated with AD neurodegeneration

For network-structure pairs, we rejected 170 out of 432 network structure pairs. While most rejected pairs occur in networks previously found significant, the pairs are not exclusive to network 1 and 2 only. Three structure pairs in network 3 were found significant. In Table 2, we show the top 5 significant pairs ranked by p-value.

3.3 Significant network structures triples associated with AD neurodegeneration

For network-structure-structure triples, we rejected 753 out of 46656 possible triples. The network and structures found significant are again not a subset of those found significant in the network-structure pair. In Table 3, we show the top 5 significant triples ranked by p-value.

3.4 Results comparing confounder regression and no regression

While the results were similar to those shown above, the analysis without accounting for confounders found more structures and networks significant than analysis with confounder regression. For networks only, network 3 became slightly significant, and networks 1 and 2 remained highly significant. The top 10 structures were the same for networks 1 and 2 in both analysis, but those in structure 3 and 4 were different. For tuples analysis, we found 205 instead of 170 pairs and 1171 instead of 753 triples significant. The networks and structures found significant in triples are not a subset of those found in pairs and vice versa.

network 1	network 2	network 3	network 4
Left Pericalcarine	Right Middle Temporal	Right Caudal Middle Frontal	Fifth Ventricle
Left Pallidum	Left Middle Temporal	Right Superior Frontal	Right Lingual
Right Pericalcarine	Left Inferior Temporal	Right Rostral Middle Frontal	Left Precuneus
Right Pallidum	Left Superior Temporal	Left Superior Frontal	Right Lateral Occipital
Right Accumbens Area	Right Inferior Temporal	Left Caudal Middle Frontal	Left Superior Parietal
Right Amygdala	Right Lateral Ventricle	Left Rostral Middle Frontal	Right Precuneus
Right Caudate	Left Lateral Ventricle	Right Inferior Parietal	Right Superior Parietal
Left Accumbens Area	Right Superior Temporal	Left Precentral	Left Inferior Parietal
Left Choroid Plexus	Left Inferior Lateral Ventricle	Right Precentral	Left Isthmus Cingulate
Left Amygdala	Right Fusiform	Right Supramarginal	Left Cerebral Cortex

Table 1: Top 10 structures ranked by loading of each network

network-structure pairs	p-value
1-Right Pallidum	<1e-4
2-Left Inferior Lateral Ventricle	<1e-4
2-Left Fusiform	<1e-4
2-Left Entorhinal	<1e-4
2-Left Cerebral Cortex	1e-4

Table 2: Top 5 significant pairs ranked by FWER-corrected p value

4 Discussion

In this work, we applied a joint graph embedding method (similar to a multivoxel dictionary embedding method), which would typically be used to study functional connectivity, to volumetric data in neurodegeneration. We extended the original algorithm[25] and implemented our maximum likelihood joint graph embedding algorithm to identify significant structural networks from volumetric data in Alzheimer's disease cohorts[26]. We showed that our version of the algorithm has promise in uncovering latent dimensions that are easy to interpret and visualize. In addition, we developed and implemented a novel testing procedure and tested for significant networks, network-structure pairs and network-structure-structure triples. Since these pairs and triples are a novel description of neurodegeneration patterns, we will briefly state their interpretation. For example, our discovered pairs can be interpreted as "the right pallidum displays a significant association with disease status, through its role in network 1". As another example, our discovered triples can be interpreted as "the interaction between the left inferior temporal lobe and right middle temporal lobe displays a significant association with disease status, through its

network-structure-structure triples	p-value
2-Left Inferior Temporal-Right Middle Temporal	<1e-4
2-Left Middle Temporal-Right Middle Temporal	<1e-4
2-Left Superior Temporal-Right Middle Temporal	<1e-4
2-Left Inferior Temporal-Left Middle Temporal	<1e-4
2-Left Middle Temporal-Left Superior Temporal	1e-4

Table 3: Top 5 significant triples ranked by FWER-corrected p value

role in network 2". We performed analysis to regress out common confounders in AD in hope of gaining more discovery power and make a few comments here between our results. We found fewer significant pairs and triples when taking confounders into account than not. This is expected, as some structures previously found significant may be caused by confounders. While networks 1 and 2 remained unchanged and highly significant, indicating that changes in these networks are due to disease progression, network 3 was no longer a significant network. The structures' change in networks 3 and 4 also indicate these structures may have been significant due to common confounders such as age.

Next we discuss a few limitations in this work. First, due to our high standard of cohort selection, this dataset may not be representative of clinical MRI subjects. As such, a natural future direction for the study is to apply this framework to clinical datasets, which represent a more diverse population. Secondly, our findings are on a population scale instead of an individual scale, which may lead to additional biases when applying to individual clinical diagnosis. We believe that using a more fine-grained separation of the cohort, as opposed to modeling disease status as only two groups, may help address this issue. Thirdly, we note that since there has been little work done to apply graph embedding methods to study volumetric data in Alzheimer's Disease cohorts, we do not have result comparisons with "state-of-the-art" methods. As with most unsupervised methods, there is no ground truth for evaluation to draw a conclusion on which method is best. Rather, we offer an additional method in analyzing group differences between healthy and diseased individuals based on neurodegeneration.

Our framework shows promise in that it discovered structures commonly believed to be associated with AD neurodegeneration. We point out several strengths in our framework. First, each of the three groups of findings gives us more information on structural correlations than traditional methods[22], and we also find structures not found by previous methods, such as mass univariate analysis[5]. Secondly, the dataset used in our study, ADNI[14], is very well characterized, it is more robust and spans a longer period than most clinical data. Thirdly, though increased complexity in algorithm and testing procedure often require more computational time, our algorithm embeds 494 patients with 108 structures in just 30 minutes and runs significance testing in 4 hours. We experimented with cohort and structure sizes and show runtime results in Figure 6. We note that while the algorithm's runtime scales linearly in subjects, it is exponential in structure counts. However, as studies often do not include a large set of structures, it should not be of big concern to users. On the other hand, our algorithm shows good time complexity with increased cohort sizes. By offering more flexibility in terms of patterns that we can identify, we hope that our framework will reveal biomarker patterns more sensitive to AD.

With the recent development of potential drugs to treat Alzheimer's disease, methods like ours that can quantify complex patterns of neurodegeneration will be essential for noninvasively identifying patients who would benefit the most. We believe that the development and dissemination of algorithms such as this one, through open source code and well documented examples, will play an important role in helping to reduce the burden of Alzheimer's disease on our aging population.

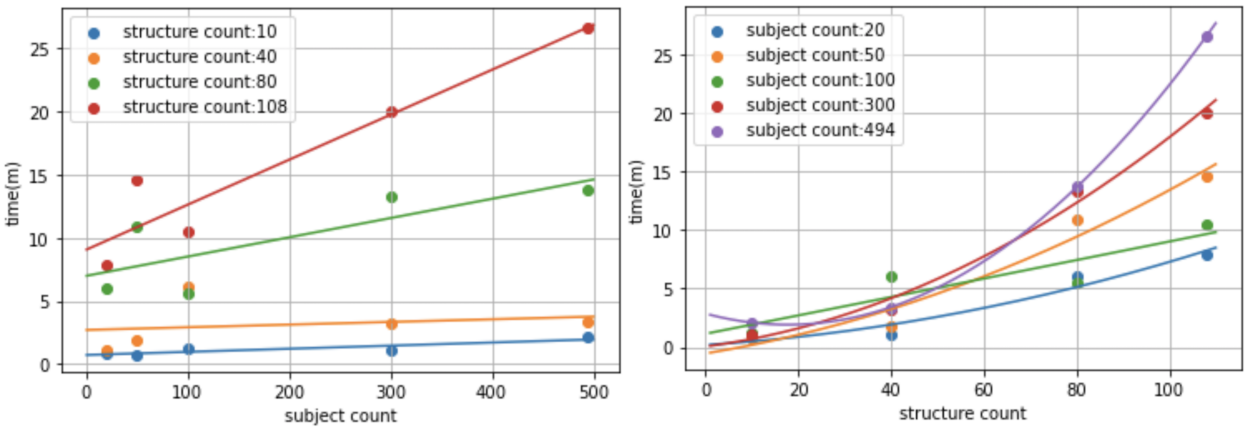

(a) Runtime with increasing subjects

(b) Runtime with increasing structures

Figure 6: Runtime with increasing parameters

Acknowledgement

Data collection and sharing for this project was funded by the Alzheimer's Disease Neuroimaging Initiative (ADNI) (National Institutes of Health Grant U01 AG024904) and DOD ADNI (Department of Defense award number W81XWH-12-2-0012). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from the following: AbbVie, Alzheimer's Association; Alzheimer's Drug Discovery Foundation; Araclon Biotech; BioClinica, Inc.; Biogen; Bristol-Myers Squibb Company; CereSpir, Inc.; Cogstate; Eisai Inc.; Elan Pharmaceuticals, Inc.; Eli Lilly and Company; EuroImmun; F. Hoffmann-La Roche Ltd and its affiliated company Genentech, Inc.; Fujirebio; GE Healthcare; IXICO Ltd.;Janssen Alzheimer Immunotherapy Research & Development, LLC.; Johnson & Johnson Pharmaceutical Research & Development LLC.; Lumosity; Lundbeck; Merck & Co., Inc.;Meso Scale Diagnostics, LLC.; NeuroRx Research; Neurotrack Technologies; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Piramal Imaging; Servier; Takeda Pharmaceutical Company; and Transition Therapeutics. The Canadian Institutes of Health Research is providing funds to support ADNI clinical sites in Canada. Private sector contributions are facilitated by the Foundation for the National Institutes of Health (www.fnih.org). The grantee organization is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimer's Therapeutic Research Institute at the University of Southern California. ADNI data are disseminated by the Laboratory for Neuro Imaging at the University of Southern California.

This work was supported by The Karen Toffler Charitable Trust through the Toffler Scholars Program. Funders played no role in study design; in the collection, analysis and interpretation of data; in the writing of the report; or in the decision to submit the article for publication.

References

[1] A. Abraham, E. Dohmatob, B. Thirion, D. Samaras, and G. Varoquaux. Extracting brain regions from rest fmri with total-variation constrained dictionary learning. In *International conference on medical image computing and computer-assisted intervention*, pages 607–615. Springer, 2013.

[2] M. S. Albert, S. T. DeKosky, D. Dickson, B. Dubois, H. H. Feldman, N. C. Fox, A. Gamst, D. M. Holtzman, W. J. Jagust, R. C. Petersen, et al. The diagnosis of mild cognitive impairment due to alzheimer's

disease: recommendations from the national institute on aging-alzheimer's association workgroups on diagnostic guidelines for alzheimer's disease. *Alzheimer's & dementia*, 7(3):270–279, 2011.

[3] A. Association. 2019 alzheimer's disease facts and figures. *Alzheimer's & dementia*, 15(3):321–387, 2019.

[4] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.

[5] J. L. Bernal-Rusiel, M. Reuter, D. N. Greve, B. Fischl, M. R. Sabuncu, A. D. N. Initiative, et al. Spatiotemporal linear mixed effects modeling for the mass-univariate analysis of longitudinal neuroimage data. *Neuroimage*, 81:358–370, 2013.

[6] B. Biswal, F. Zerrin Yetkin, V. M. Haughton, and J. S. Hyde. Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magnetic resonance in medicine*, 34(4):537–541, 1995.

[7] G. Brandl. Sphinx documentation. *URL http://sphinx-doc. org/sphinx. pdf*, 2021.

[8] V. D. Calhoun, T. Adali, G. D. Pearlson, and J. J. Pekar. A method for making group inferences from functional mri data using independent component analysis. *Human brain mapping*, 14(3):140–151, 2001.

[9] J. D. Cohen, N. Daw, B. Engelhardt, U. Hasson, K. Li, Y. Niv, K. A. Norman, J. Pillow, P. J. Ramadge, N. B. Turk-Browne, et al. Computational approaches to fmri analysis. *Nature neuroscience*, 20(3):304–313, 2017.

[10] D. M. Cole, S. M. Smith, and C. F. Beckmann. Advances and pitfalls in the analysis and interpretation of resting-state fmri data. *Frontiers in systems neuroscience*, page 8, 2010.

[11] N. S. D'Souza, M. B. Nebel, N. Wymbs, S. Mostofsky, and A. Venkataraman. Integrating neural networks and dictionary learning for multidimensional clinical characterizations from functional connectomics data. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 709–717. Springer, 2019.

[12] A. Fornito, A. Zalesky, and E. Bullmore. *Fundamentals of brain network analysis*. Academic Press, 2016.

[13] C. R. Jack, D. A. Bennett, K. Blennow, M. C. Carrillo, H. H. Feldman, G. B. Frisoni, H. Hampel, W. J. Jagust, K. A. Johnson, D. S. Knopman, et al. A/t/n: an unbiased descriptive classification scheme for alzheimer disease biomarkers. *Neurology*, 87(5):539–547, 2016.

[14] C. R. Jack Jr, M. A. Bernstein, N. C. Fox, P. Thompson, G. Alexander, D. Harvey, B. Borowski, P. J. Britson, J. L. Whitwell, C. Ward, et al. The alzheimer's disease neuroimaging initiative (adni): Mri methods. *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 27(4):685–691, 2008.

[15] C. R. Jack Jr and D. M. Holtzman. Biomarker modeling of alzheimer's disease. *Neuron*, 80(6):1347–1358, 2013.

[16] M. Khosla, K. Jamison, G. H. Ngo, A. Kuceyeski, and M. R. Sabuncu. Machine learning in resting-state fmri analysis. *Magnetic resonance imaging*, 64:101–121, 2019.

[17] M. H. Lee, C. D. Hacker, A. Z. Snyder, M. Corbetta, D. Zhang, E. C. Leuthardt, and J. S. Shimony. Clustering of resting state networks. *PloS one*, 7(7):e40370, 2012.

[18] K. Levin, A. Athreya, M. Tang, V. Lyzinski, Y. Park, and C. E. Priebe. A central limit theorem for an omnibus embedding of multiple random graphs and implications for multiscale network inference. *arXiv preprint arXiv:1705.09355*, 2017.

[19] J. A. Lewis-Peacock and K. A. Norman. Multi-voxel pattern analysis of fmri data. *The cognitive neurosciences*, 512:911–920, 2014.

[20] J. C. Morris. The clinical dementia rating (cdr): Current version and. *Young*, 41:1588–1592, 1991.

[21] T. Nichols and S. Hayasaka. Controlling the familywise error rate in functional neuroimaging: a comparative review. *Statistical methods in medical research*, 12(5):419–446, 2003.

[22] G. Pengas, G. Williams, J. Acosta-Cabronero, T. Ash, Y. Hong, D. Izquierdo-Garcia, T. Fryer, J. Hodges, and P. Nestor. The relationship of topographical memory performance to regional neurodegeneration in alzheimer's disease. *Frontiers in aging neuroscience*, 4:17, 2012.

[23] D. L. Sussman, M. Tang, D. E. Fishkind, and C. E. Priebe. A consistent adjacency spectral embedding for stochastic blockmodel graphs. *Journal of the American Statistical Association*, 107(499):1119–1128, 2012.

[24] M. P. Van Den Heuvel and H. E. H. Pol. Exploring the brain network: a review on resting-state fmri functional connectivity. *European neuropsychopharmacology*, 20(8):519–534, 2010.

[25] S. Wang, J. Arroyo, J. T. Vogelstein, and C. E. Priebe. Joint embedding of graphs. *IEEE transactions on pattern analysis and machine intelligence*, 43(4):1324–1336, 2019.

[26] B. T. Wyman, D. J. Harvey, K. Crawford, M. A. Bernstein, O. Carmichael, P. E. Cole, P. K. Crane, C. DeCarli, N. C. Fox, J. L. Gunter, et al. Standardization of analysis sets for reporting results from adni mri data. *Alzheimer's & Dementia*, 9(3):332–337, 2013.

[27] M. Xu. Understanding graph embedding methods and their applications. *SIAM Review*, 63(4):825–853, 2021.

[28] J. Yang, S. Gohel, and B. Vachha. Current methods and new directions in resting state fmri. *Clinical imaging*, 65:47–53, 2020.