

CHARACTERIZING FREQUENCY-SELECTIVE NETWORK VULNERABILITY FOR ALZHEIMER’S DISEASE BY IDENTIFYING CRITICAL HARMONIC PATTERNS

Benjamin Leinwand Guorong Wu Vladas Pipiras

University of North Carolina at Chapel Hill

ABSTRACT

Alzheimer’s disease (AD) is a multi-factor neurodegenerative disease that selectively affects certain regions of the brain while other areas remain unaffected. The underlying mechanisms of this selectivity, however, are still largely elusive. To address this challenge, we propose a novel longitudinal network analysis method employing sparse logistic regression to identify frequency-specific oscillation patterns which contribute to the selective network vulnerability for patients at risk of advancing to the more severe stage of dementia. We fit and apply our statistical method to more than 100 longitudinal brain networks, and validate it on synthetic data. A set of critical connectome pathways are identified that exhibit strong association to the progression of AD.

Index Terms— Brain network, graph spectrum, sparse logistic regression, Alzheimer’s disease.

1. INTRODUCTION

Alzheimer’s disease, the most common neurodegenerative disorder, leads to gradually progressive memory loss, decline in other cognitive domains, altered behavior, loss of functional abilities, and death. Although progressive neuron loss is a hallmark of AD [1], some neurological impairment may reflect dysfunction rather than loss of neurons [2]. In this regard, AD can be understood as a disconnection syndrome where the brain network is progressively disrupted by neuropathological processes that are not fully understood.

Recent developments in diffusion MRI and network neuroscience allow us to characterize the white matter pathways that connect gray matter regions in the context of the large-scale brain network. The network degeneration hypothesis – significant changes might occur in the topological properties of the structural brain network as AD progresses – is supported by many neuroimaging studies [3]. In addition, convergent evidence shows that the presence of AD pathology burden exhibits a unique spatial pattern that is highly correlated with the region-to-region connections in the network. In this context, it is important to improve our current understanding of AD by characterizing how neuropathological events that result in cognitive decline spread across brain networks.

Longitudinal analysis is of particular interest in AD neuroimaging studies, since the measured subject-specific network alterations allow us to develop highly sensitive and specific connectome biomarkers for early AD diagnosis. With the increasing availability of large public longitudinal databases such as ADNI (<http://adni.loni.usc.edu/>), more and more studies have shown the advantages of longitudinal network analyses methods over conventional cross-sectional methods in characterizing network alterations due to AD [4, 5, 6]. Due to the high dimensionality of brain networks, however, it is common practice to use node-wise measurements such as local clustering coefficient and small-world-ness,

Table 2.1. Subjects and scans used in each step of model building

	Participant cohort size		Number of scans	
	Average MCI	Lasso	Average MCI	Lasso
Total	68	48	234	114
Gender				
Male	44	30	146	69
Female	24	18	88	45
Education				
11 - 16 years	39	25	127	60
17 - 20 years	29	23	107	54
APOE4				
0 copies	32	22	112	53
1 copy	26	18	87	45
2 copies	10	8	35	16

instead of whole-brain connectivity information. In simplifying the data, though, these node-wise features can obscure information embedded in the links, making it difficult to unravel spatially localized effects of AD pathology on connectivity. Conversely, many approaches apply significance tests at each link separately [7]. Due to a large number of links in the network, severe multiple-comparison correction is needed to control false discovery rate, which might result in discarding the links that are scientifically meaningful.

To address these challenges, we propose a novel statistical model to discover the critical connectomic pathway in the graph spectrum domain which consists of two major steps. (1) Alignment of longitudinal networks. For each subject, we first align their longitudinal networks into a common graph spectrum domain. (2) Inter-subject variable selection in the graph spectrum domain. We propose a novel sparse regression model to identify the critical frequency patterns which contribute to selective network vulnerability.

2. DATA

We build our method on 234 DTI scans from the ADNI database. These scans are of 68 subjects diagnosed at the times of these scans with Mild Cognitive Impairment (MCI). 56 of these subjects have multiple scans at different times while diagnosed with MCI. For each scan, we parcellate the cortical surface into 148 regions of the Destrieux Atlas using FreeSurfer on T1-weighted MRI scans. Then we apply probabilistic fiber tractography on DWI and T1-weighted images using FSL software library to obtain 148×148 connectivity matrices of the structural networks, where each element reflects the number of white matter fibers traveling between two brain regions.

This is a subset of a larger dataset in which we also have scans where subjects (including those in our subset) have other diagnoses at the time the scan is taken, including “Normal.” Unfortunately, the diagnostic categories at the time of first diagnosis don’t match up with the diagnoses at time of later scans. Due to this disparity,




Fig. 3.1. Starting with the matrices G^{st} , keep only scans where the diagnosis is MCI, then calculate the average MCI graph $\bar{G}^{MCI}$. In the simplified illustration, each subject’s scans are kept on a single row, and scans with diagnoses of Normal or AD are discarded.

we consider “Early Mild Cognitive Impairment,” “Late Mild Cognitive Impairment,” and “Mild Cognitive Impairment” as a single diagnosis. We do the same for “Alzheimer’s Disease” and “Dementia.” From this point forward, “Mild Cognitive Impairment” and “Alzheimer’s Disease” will be referred to as “MCI” and “AD” respectively.

Many scans are missing either diagnostic data or MMSE and CDRSB scores. Where data is missing, information from the last scan where the data is available for a given subject is used as a substitute (interpolating the data does not substantively change the results). For constructing the average basis, only scans that had a diagnosis of MCI, or which had a missing diagnosis but the last available diagnosis was MCI were included. For the logistic regression using LASSO, any scans excluded from basis construction are still excluded, but there are stricter rules in place to remove uncertainty.

3. METHOD

We start by symmetrizing and taking the log of the 148×148 matrices representing the structural brain networks of patients at different points in time, where each entry in the original matrix represents the count of fibers connecting the region in the row and the region in the column. Let G^{st} be the t th processed matrix observed for subject s .

For the 234 scans where the diagnosis at time of scan is MCI, we calculate the “average MCI graph” $\bar{G}^{MCI}$,

$$\bar{G}^{MCI} = \sum_{\substack{s,t \\ \text{Diagnosis} = \text{MCI}}} G^{st} / \sum_{\substack{s,t \\ \text{Diagnosis} = \text{MCI}}} 1.$$

From here we calculate the Laplacian of $\bar{G}^{MCI}$ denoted as $\bar{L}^{MCI} = I - \bar{D}^{-\frac{1}{2}} \bar{G}^{MCI} \bar{D}^{-\frac{1}{2}}$ where $\bar{D}$ is the degree matrix of $\bar{G}^{MCI}$, which has the degree of each node on the diagonal. We use this normalized version of the Laplacian to avoid issues with eigenvector localization in graphs with wide degree distributions. Decompose $\bar{L}^{MCI}$ as

$$\bar{L}^{MCI} = B \bar{\Lambda}_{MCI} B^T,$$

where B is a matrix whose columns are the eigenvectors of $\bar{L}^{MCI}$, and $\bar{\Lambda}_{MCI}$ is a diagonal matrix containing the eigenvalues of $\bar{L}^{MCI}$.

We map every scan G^{st} into the subspace spanned by B . Instead of regressing each individual scan onto the orthonormal eigenbasis of the average MCI graph $\bar{G}^{MCI}$, to control for differences in the total number of edges observed in each scan, we project each *degree-normalized graph* onto B by choosing a $\hat{\Lambda}^{st}$ vector which minimizes the Frobenius norm between the observed degree-normalized matrix and the estimated matrix. Formally, we solve

$$\hat{\Lambda}^{st} = \arg \min_{\Lambda} \|I - (D^{st})^{-1/2} G^{st} (D^{st})^{-1/2} - B \Lambda B^T\|_F$$




Fig. 3.2. From the average MCI graph $\bar{G}^{MCI}$ (top left), calculate the orthonormal basis B of the average Laplacian. Project each degree-corrected scan diagnosed with MCI onto the subspace spanned by B . Take longitudinal differences of the projection coefficients, include these differences in an L_1 regularized logistic regression. In the simplified illustration, changes along the basis vector with weight concentrated on the left hemisphere outside of the frontal lobe (in green) indicate an elevated risk of transitioning to AD.

$$= B^T (I - (D^{st})^{-1/2} G^{st} (D^{st})^{-1/2}) B.$$

As Λ is in principle supposed to be a diagonal matrix, we can treat each $\hat{\Lambda}^{st}$ as the 148 length vector taken from the diagonal. We look at the changes in these estimated $\hat{\Lambda}$ vectors for a given subject from one time to another, rather than only observing the snapshot at the time of the scan. Formally, define

$$\Delta^{st} = \hat{\Lambda}^{st} - \hat{\Lambda}^{st-1}.$$

We perform L_1 regularized logistic regression [8] to sparsely select variables which predict whether a patient will transition to AD within 400 days. Potential covariates include the Δ^{st} values as well as patient information including age, gender, education, ethnicity, race, marital status, CDRSB score, MMSE score, time since last scan, change in MMSE score from last observed MMSE score, change in CDRSB score from last observed CDRSB score, and these last two variables interacting with time since last scan.

There are a few novel aspects to this method. First, it incorporates subject specific changes over time, regressing clinical outcomes on Δ^{st} ’s rather than on $\hat{\Lambda}^{st}$ ’s. Second, the $\hat{\Lambda}^{st}$ ’s themselves do not correspond to the eigenvectors of each G^{st} , but approximate the original G^{st} ’s using B . Finally, we do not restrict our method to choosing from exclusively the *first* several values in the 148 length Δ^{st} , as the more informative structural changes may occur along “later” vectors in the basis B outside of the first few columns.

4. RESULTS

4.1. Main result

Due to data availability, we consider any scan where a patient is currently diagnosed as MCI and subsequently diagnosed with AD by their next scan within the next 400 days as a “success,” and any scan where the patient is not diagnosed with AD by their next scan as a “failure.” When keeping only those relevant scans, this leaves 114 scan differences, 9 of which represent patients who subsequently transition to AD, while the rest remain in MCI.

The procedure above selects four variables which are useful in predicting those patients who may transition to AD: The current (backfilled) CDRSB scores, the change in MMSE scores, and the




Fig. 4.1. Visualizations of the 108th (top) and 125th (bottom) eigenvectors projected onto the cortical surface, with views of each hemisphere from above, left, and right. The 108th is more concentrated on the right hemisphere, while the 125th is basically symmetric.

108th and 125th vectors from the basis B . Since each element in the recovered vectors is associated with a particular brain region, we visualize the harmonic basis on the cortical surface, as shown in Fig. 4.1, where red and blue regions correspond to positive and negative values in the eigenvectors, respectively. The purpose of using LASSO is to filter out irrelevant variables, so the selection of variables derived from brain scans indicates that these scans contain information which incrementally improve the model’s fit. As the CDRSB and MMSE scores are used to diagnose AD, we expect to see them as predictors for transitions. Incorporating the selected vectors into a similar model without brain scan information results in a 66% improvement of the deviance ratio from .15 to .25 (a true “apples to apples” comparison is difficult, as the choice of tuning parameter and included variables depends on available covariates).

Major oscillations occur at temporal lobe and parietal lobe, which are aligned with the default mode network. The 108th eigenvector puts a lot of positive weight on the right superior frontal gyrus, and negative weight concentrates on the right central sulcus and right subcallosal cingulate gyrus. As the parameter estimate on this vector is negative, observing growth in the regions with negative weights relative to those with positive weights may indicate an elevated risk of transitioning from MCI to AD in the next 400 days. The 125th eigenvector is more evenly distributed, concentrating positive weight on the right and left vertical ramus of the lateral sulcus, left collateral transverse anterior sulcus, left parietal superior gyrus, and right precuneus gyrus. The parameter estimate on this vector is positive, so relative growth in the positive regions may indicate increased risk.

4.2. Contemporaneous brain scan statistics

The methodology relies on differencing, requiring at least two scans per subject, which has two drawbacks. First, a patient needs to come in for a second scan to diagnose potential issues, so results will not be available immediately. Second, since every second scan has a corresponding first scan that can’t be used, there are fewer data points

Table 4.1. Kinds of variables selected in 100 synthetic model trials

True Model				Other variables		Count
CDRSB	MMSE change	Vector 108	Vector 125	Other vectors	Other traits	
Y	Y	Y	Y	N	N	1
Y	Y	Y	Y	Y	N	2
Y	Y	Y	Y	Y	Y	14
Y	Y	Y	N	N	N	8
Y	Y	Y	N	Y	N	3
Y	Y	Y	N	Y	Y	3
Y	Y	N	Y	Y	Y	1
Y	Y	N	N	N	N	3
Y	Y	N	N	Y	N	2
Y	Y	N	N	Y	Y	1
Y	N	Y	N	N	N	17
Y	N	Y	N	Y	N	2
Y	N	N	Y	N	N	1
Y	N	N	N	N	N	24
Y	N	N	N	Y	N	3
N	N	N	N	N	N	15

available for model training. To alleviate these issues, as an alternative to the main method, one can use the $\hat{\Lambda}^{st}$ values in place of the Δ^{st} values. As first scans can be used in this alternative, we have a slightly different dataset for this regression. In this version, we don’t include time since last scan, change in MMSE scores, change in CDRSB scores, and the relevant interaction terms, as they would not be available for certain data points. Unfortunately, the LASSO selects only current CDRSB scores, but no variables from brain scans. This bolsters the case for using our proposed methodology based on longitudinal differences, which finds useful incremental data from brain scans, at the cost of requiring multiple scans.

4.3. Synthetic dataset

A more complex test was also designed to see if, assuming we know the underlying mechanism, the approach we described would recover the true result, similar to the power examinations performed in [9]. The idea is to simulate graphs g^{st} (for the rest of this section, lowercase letters will represent synthetic versions of uppercase letters used when describing the methodology with real data) that are explicitly generated according to the model we recover, then use our procedure to see if it captures this “true” result. If the brain actually changes in accordance with the model recovered from the real data, would our approach be able to detect the described mechanism?

As our approach relies on finding Λ for each scan, the key to creating the synthetic dataset is randomly assigning each synthetic scan a random λ . Every subject’s *first* scan’s λ is Normally distributed with mean and covariance equal to the sample mean and sample covariance values of Λ^{st} restricted to scans corresponding to a diagnosis of MCI. There are two different cases for generating the λ values for the later scans ($t > 1$). In either case, we begin with the previously assigned λ^{st-1} value and add a random δ^{st} vector to represent change over time. For those graphs that *do not* represent a subject who transitions to AD in the next 400 days, δ^{st} is distributed normally with mean and covariance equal to the sample mean and sample covariance values (call these μ_1 and Σ_1 , respectively) of Δ^{st} restricted to scans where both the previous and current diagnoses are MCI. For those graphs that *do* represent a subject who transitions to AD in the next 400 days, δ^{st} is also normally distributed with mean μ_2 and covariance Σ_2 . μ_2 is identical to μ_1 except at the 108th and 125th locations, where those entries in μ_2 are equal to the 108th and 125th entries in sample mean of the Δ^{st} values where the pre-

vious scan ($st - 1$) is diagnosed as MCI, but the current scan (st) is diagnosed as AD. Σ_2 is identical to Σ_1 except in the 108th and 125th rows and columns. The entries in the 108th and 125th rows and columns of Σ_2 are all set to 0 except at the 108th and 125th locations, respectively, where they equal to the sample variance of the 108th and 125th entry in the $\Delta^{s,t}$ where the previous scan is diagnosed as MCI, but the current scan is diagnosed as AD. After generating all these λ^{st} 's, calculate g^{st} as

$$g^{st} = \bar{D}^{\frac{1}{2}} (I - B\lambda^{st}B^T) \bar{D}^{\frac{1}{2}},$$

where the vector λ^{st} is treated as a diagonal matrix to get a 148×148 graph. We also set the diagonal of each g^{st} to be 0.

In the synthetic dataset, each subject s and scan t correspond to the same subject s and scan t as in the original dataset. The diagnosis and patient information is taken from the original dataset, but the λ^{st} values and hence the synthetic brain networks g^{st} are randomly constructed without direct reference to the G^{st} matrices. We know the “true” generating mechanism for the synthetic data, as it's the estimated mechanism for the real data, so when we run the proposed analysis, we can see if we recover this mechanism.

Simulations reveal that about 85% of the time the LASSO will select variables, and if it does, it will always choose the current CDRSB score, about half the time it will select change in MMSE score, and also about half the time it will select a variable corresponding to a basis vector from b that nearly identically matches the 108th eigenvector of B (calculated using L_2 distance), and 1/5 of the time chooses a vector matching the 125th eigenvector. More detailed results can be seen in table 4.1. Since the only things being simulated are the brain networks, and the 108th eigenvector has a much higher parameter estimate than the 125th eigenvector, these results comport with expectations given the generative model. In the case that the brain is deteriorating in this particular way, the proposed procedure using the LASSO and the differences along the constructed basis appears able to frequently recover this “true” mode of change. Furthermore, the methodology infrequently selects incorrect variables. However, it should be noted that this is a relatively simple simulation featuring many assumptions.

4.4. Discussion

Due to the complex structure of the data, we have tried to identify changes to areas that may indicate elevated risk for transitioning to AD. We intentionally designed a method that, while accounting for known heterogeneities in brain structure, can learn from the observed dataset while remaining flexible enough to handle new scans. This analysis does not purport to describe any mechanism by which the brain deteriorates, but rather notes that when certain changes are observed, they may be harbingers for a changing diagnosis. Furthermore, our analysis relies on a relatively small sample size focusing on very few subjects who transitioned from MCI to AD. Future work may observe whether changes in cortical thickness in the relevant regions can also be used to predict transitions to AD. With these caveats in mind, our results furnish evidence that observing physical changes in particular regions of the brain may provide incremental information for detecting a subject's risk of transitioning to AD.

5. REFERENCES

- [1] C.R. Jack, D.A. Bennett, K. Blennow, M.C. Carrillo, B. Dunn, S.B. Haeberlein, D.M. Holtzman, W. Jagust, F. Jessen, J. Karlawish, and et al., “Nia-aa research framework: Toward a biological definition of alzheimer's disease,” *Alzheimer's Dementia*, vol. 14, no. 4, pp. 535–562, Apr 2018.
- [2] J.J. Palop, J. Chin, and L. Mucke, “A network dysfunction perspective on neurodegenerative diseases,” *Nature*, vol. 443, no. 7113, pp. 768–773, Oct 2006.
- [3] M.Á. Araque Caballero, M. Suárez-Calvet, M. Duering, N. Franzmeier, T. Benzinger, A.M. Fagan, R.J. Bateman, C.R. Jack, J. Levin, M. Dichgans, et al., “White matter diffusion alterations precede symptom onset in autosomal dominant alzheimer's disease,” *Brain*, vol. 141, no. 10, pp. 3065–3080, 2018.
- [4] A. Chincarini, F. Sensi, L. Rei, G. Gemme, S. Squarcia, R. Longo, F. Brun, S. Tangaro, R. Bellotti, N. Amoroso, and et al., “Integrating longitudinal information in hippocampal volume measurements for the early detection of Alzheimer's disease,” *NeuroImage*, vol. 125, pp. 834–847, Jan 2016.
- [5] I. Driscoll, C. Davatzikos, Y. An, X. Wu, D. Shen, M. Kraut, and S. M. Resnick, “Longitudinal pattern of regional brain volume change differentiates normal aging from MCI,” *Neurology*, vol. 72, no. 22, pp. 1906–1913, Jun 2009.
- [6] D.K. Johnson, M. Storandt, J.C. Morris, and J.E. Galvin, “Longitudinal study of the transition from healthy aging to Alzheimer disease,” *Archives of Neurology*, vol. 66, no. 10, Oct 2009.
- [7] W.H. Kim, A.M. Racine, N. Adluru, S.J. Hwang, K. Blennow, H. Zetterberg, C.M. Carlsson, S. Asthana, R.L. Kosciuk, S.C. Johnson, and et al., “Cerebrospinal fluid biomarkers of neurofibrillary tangles and synaptic dysfunction are associated with longitudinal decline in white matter connectivity: A multi-resolution graph analysis,” *NeuroImage: Clinical*, vol. 21, pp. 101586, 2019.
- [8] T. Hastie, R. Tibshirani, and J. H. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Springer series in statistics. Springer, 2nd edition, 2009.
- [9] D. Fraiman and R. Fraiman, “An ANOVA approach for statistical comparisons of brain networks,” *Scientific Reports*, vol. 8, no. 1, pp. 4746, Dec 2018.
- [10] A. Eid, I. Mhatre, and J.R. Richardson, “Gene-environment interactions in Alzheimer's disease: A potential path to precision medicine,” *Pharmacology Therapeutics*, vol. 199, pp. 173–187, Jul 2019.
- [11] X. Hou and L. Zhang, “Saliency Detection: A Spectral Residual Approach,” *2007 IEEE Conference on Computer Vision and Pattern Recognition*, 2007.
- [12] A. Raj, A. Kuceyeski, and M. Weiner, “A network diffusion model of disease progression in dementia,” *Neuron*, vol. 73, no. 6, pp. 1204–1215, Mar 2012.
- [13] A. Raj, E. LoCastro, A. Kuceyeski, D. Tosun, N. Relkin, and M. Weiner, “Network diffusion model of progression predicts longitudinal patterns of atrophy and metabolism in Alzheimer's disease,” *Cell Reports*, vol. 10, no. 3, pp. 359–369, Jan 2015.
- [14] K. Rohe, S. Chatterjee, and B. Yu, “Spectral clustering and the high-dimensional stochastic blockmodel,” *The Annals of Statistics*, vol. 39, no. 4, pp. 1878–1915, Aug 2011.