

Hierarchical Labeling of Resting State fMRI Brain Networks

Advanced Computer Integrated Surgery II

Group 9: Arden Chew

Mentors: Dr. Haris Sair and Dr. Mathias Unberath

May 6, 2019

Introduction

Modern techniques for pre-operative neurosurgical planning rely primarily on task-based functional magnetic resonance imaging (fMRI). Such a work flow generally involves a patient performing anywhere from three to ten actions as prompted by a neuroradiologist, then assigning the areas of neuronal activity measured by blood oxygen level in the fMRI brain scan to the associated brain networks. When planning area of operation and point of entry, a physician relies entirely on these task-based fMRI assigned regions of importance and may choose a path that causes unforeseen neurological damage.

Resting state functional magnetic resonance imaging (rsfMRI) provides an intriguing alternative to pre-operative planning. The rsfMRI imaging modality measures the same blood oxygen level as task-based fMRI but in a task negative state over a longer time course. After independent component analysis (ICA) separates the unique activation patterns in rsfMRI data into brain networks, the functional organization of the brain can be inferred. These networks are far more precise and complex than a series of task-based fMRI mappings, often providing upwards of 100 unique brain networks and a more complete coverage of important brain regions.

Surgical planning requires more than simple activation patterns, however. Instead, a surgeon must know what each component correlates to when prioritizing which areas of the brain to avoid, and what the resultant neurological outcomes will be. Furthermore, many activation components detailed by ICA account simply for unique noise patterns. Estimates provided by Dr. Haris Sair of Johns Hopkins Neuroradiology hypothesize that up to 60% of ICA networks may be noise in a given study. Unfortunately, only a handful of neuroradiologists in the United States can reliably label these more complex rsfMRI brain networks at high granularity. This is primarily due to the complex hierarchical relationships that brain networks have with each other. A strong knowledge both of general cerebral anatomy as well as said hierarchy is essential to successful rsfMRI component labeling. This study aims to automate the process of component labeling so as to provide greater access to the novel rsfMRI pre-surgical planning scheme.

Technical Approach

Overview

A pixelated classification approach for rsfMRI networks was explored in Washington University's "*Resting state network estimation in individual subjects*," however success was only viable after hand selecting features to differentiate against. Furthermore, only seven different brain networks were explored. The lack of depth here prevented any need for exploration of functional hierarchies in brain networks. My approach to automated brain network classification seeks to address these shortcomings by a three-fold labelling scheme: noise filtering, initial network classification, and Pearson correlation corrected labeling. The input to my labelling scheme will be patient by patient, so as to retain its clinical viability. As such, the input format will be the post-ICA rsfMRI volumetric data of an individual and the associated time-series data. The ICA runs for labeling will be done by individual, rather than by group study.




Figure 1: Schematic labeling workflow incorporating the three main facets, noise filtering, multi-class prediction, and Pearson correlation comparison.

Noise Filtering

The binary labelling task of assigning a component to either noise or non-noise is well suited to the statistical offerings of a shallow convolutional neural network (CNN). This is attributable to the distinct feature topography and speckling differences between noise and true rsfMRI networks as well as the differing normalized localities of noise and non-noise.




Figure 2: Sample noise network with scattered activation throughout non-localized regions of the brain, produced from data acquired in OpenNEURO Washington 120 dataset

I trained my initial network on 20 subjects from the Washington 120 OpenNEURO dataset which were subsequently labelled at Johns Hopkins University. Validation was conducted with 5 set aside subjects, and testing with another 5. Given that this particular study was conducted at the 20 component ICA level, each subject contributes a unique set of 20 network components, thereby radically increasing my database size.

Due to the rather small size of my initial training dataset, two regularization techniques were required. Firstly, dropout layers were incorporated into the linear section of my network with a dropout probability tuned to 0.3. Secondly, early stopping was implemented by limiting the total number of iterations to 100, and selecting the second epoch model as the best for generalization.

Binary cross entropy loss and stochastic gradient descent optimization were chosen as is standard in binary CNN applications.




Figure 3: Shallow convolutional neural network for binary noise classification of rsfMRI brain network components

Network Parameters for Binary Classification	
Loss function	Binary cross entropy loss
Optimization function	Stochastic gradient descent
Output activation function	Rectified linear activation
Dropout rate	0.3
Number of iterations	100
Epochs	5
Batch size	1
Learning rate	0.05

Multi-class Component Classification

At the heart of this classification task is the ability to distinguish brain networks based on their unique shapes and locations. This is a similarly fitting task for a CNN. Given that the distinguishing factors of a brain network are not small curves but rather “blobs” of activation sights, a complex and/or deep network is not necessary in this particular task. A similar network architecture to the noise filtering network, altered for multi-class prediction was implemented.




Figure 4: Shallow convolutional neural network for multi-class classification of rsfMRI brain network components

Network Parameters for Multi-class Classification	
Loss function	Cross entropy loss*
Optimization function	Stochastic gradient descent
Output activation function	Probabilistic softmax*
Dropout rate	0.3
Number of iterations	100
Epochs	5
Batch size	1
Learning rate	0.05

Initial training for this multiclass network was conducted on another 20 subjects from the Washington 120 dataset. In total 30 subjects were processed via group ICA such that component labels remained consistent. This consistency eliminated the need for specific labelling of the unique components. 5 subjects were set aside for validation, and another 5 for

testing once more. Due to the similarities in network architecture, ideal hyperparameters for the multi-class CNN tuning remained about the same as the binary noise CNN.

It is necessary to note that for efficiency, all training was conducted on an NVIDIA GeForce graphic processing unit.

Hierarchical Relationship Incorporation

Difficulties in classifying network components can be resolved by observing the relationship an ambiguous network component has with an unambiguous component. These relationships are quantifiable due to the fact that related networks tend to have similar activation patterns over time. As such, Pearson correlation across time series provides a valuable tool for evaluating components with low confidence in their initial labels, and which hierarchy they may belong to.

Small components with low multi-class classification confidence, but high Pearson correlation coefficient and statistically significant p-values are likely subcomponents of other networks present. This is valuable information in pre-operative surgical planning not currently widely available in the United States.




Figure 5: Pearson correlation coefficients and p-values across time series data of a single subject from Washington 120 dataset, useful in hierarchical modeling

Statistical analysis of p-value confidence in Pearson correlation can provide a basis for evaluating the confidence in network relationships. Low confidence network classifications paired with low confidence Pearson correlations indicate the presence of an unknown or not yet classified brain network.

Complete Workflow

Once training was completed for each CNN a full pipeline was assembled for individual rsfMRI scans to be evaluated. Firstly, each component is evaluated for noise. Components with a normalized confidence of noise of 0.9 or higher were removed labelled as

noise and not further processed. All remaining components were fed through the multi-class CNN. Any components with a maximum label confidence post-softmax of less than 0.5 were evaluated for Pearson correlations, otherwise their high-confidence label was assigned. In the case of low confidence initial labelling but high confidence Pearson correlates, such correlated components were noted.

Results and Discussion

As will be noted later in this evaluation, a larger and more complex dataset is necessary to truly evaluate the success of my model.

In evaluating the noise classifier/filter, 5 testing subjects were randomly selected from the Washington 120 dataset contributing 20 network components each. The 100 components were hand-labeled at Johns Hopkins University, both by noise status and comparative brain network. The resultant accuracies are below.

Contingency Table of Noise Classification in Group rsfMRI Study

		True	
		Network	Noise
Observed	Network	74	5
	Noise	2	19

Performance Statistics of Noise Classification in Network Separated Testing

	Precision	Recall	Accuracy
Attention	0.850	0.944	0.840
Language	1.000	1.000	1.000
Default Mode	1.000	1.000	1.000
Motor	1.000	0.952	0.960
Total	0.963	0.975	0.950

With generalized accuracy of 93% and total accuracy of 95%, the noise classification model indicates strong success. More clinically applicable however is the high recall score of 0.975, indicating that it is even more unlikely that true network components will be labelled as noise. To avoid this potentially clinically dangerous scenario, a high confidence threshold is implemented in noise filtering, therefore making it statistically near impossible for a true network to be filtered out as noise. In the opposing case that a noise network is mis-identified as a true network and passes through the initial filter, low confidence with multi-class classification and Pearson correlating will provide insurance indication of noise status.




Figure 6: Confusion matrices of component classification; from left to right: initial standard CNN classification; standard CNN classification with Pearson thresholding and reassignment; noise filtered CNN classification with Pearson thresholding and p-value filtering (component 15 reserved for p-value filtering)

Furthermore, the confusion matrices in figure 6 offer a promising look at how noise filtering and Pearson correlation incorporation of hierarchical data can drastically improve brain network classification. Though these results are only conducted on 5 individuals, the final confusion matrix accuracy of 96.25% on real brain networks indicates that the full pipeline may have clinically applicable promise.

Management Summary

Lessons Learned

One of the most convincing lessons I learned during this project is the importance of considering simpler models. The most convincing and thoughtful models are often those which provide are simple and elegant, and avoid unnecessary convoluted techniques. I learned this lesson as I attempted to incorporate hierarchical information using a probabilistic convolutional deep belief network. This latent variable model is extremely complex and as only been showcased in papers over the past year and a half. After further discussion with my mentor Dr. Sair, I came to the conclusion that the hierarchical relationships are rooted in the firing over temporal rather than spatial settings. Therefore, I was able to scratch the convolutional deep belief network for a much simpler, faster, and effective Pearson correlation solution.

Another lesson that I both learned, and learned to apply, was to always keep clinical viability in mind. Though 95% accuracy may be state-of-the-art or groundbreaking for any application, 5% error is far to high for clinical usage. I learned this as I designed my noise filter, which I needed to be extra careful of handling false positives with. Should my filter remove a true network identified as noise, a physician may plan a surgery through the filtered-out region. I therefore selected a high confidence value for noise filtering, and tuned my model for high recall rather than high accuracy. Building off this, I learned to ensure that should a noise component slip through my filter, there was a catch via Pearson correlation p-values to ensure that it did not go unnoticed. These sorts of clinical checks were learned over the course of this project.

Future Work

The bulk of the future work in this study is training on more diversified data. This comes in three primary ways which will provide for both a more clinically useful and statistically significant model. Firstly, more data is necessary. 20 training individuals, 5 validation, and 5 testing is simply too small. Secondly, higher component levels must be tested. The 20 component scale begins to approach clinical viability, but the true clinical applications of rsfMRI pre-operative planning comes at component levels of 40-200, and therefore the model must be tested for this. Lastly, this study must be validated on data with non-healthy individuals. Pre-operative scanning is not useful for patients without the need for brain-surgery. Therefore, this study must both train and test on subjects with brain lesions in various locations and of various

sizes. Computer vision deformation of segmented tumor regions provides interesting potential in this regard.

In addition to improvements in training, validation, and testing data, adaption of my models to be transfer learning enabled would help modulate the addition of said data. Rather than needing to train a new model and test it each time more data becomes available, allowing for successive model training can help ease the workflow.

As I will not be continuing with this project due to work and graduation constraints, it is important that my documentation provides enough clarity for future investigators to use, implement, modify, and replicate my work. I will be going over these documents with my mentor Dr. Haris Sair to confirm that this said clarity is apparent.

Acknowledgements

I would like to thank both the Radiology Artificial Intelligence Lab and the Johns Hopkins Department of Neuroradiology for hosting and mentoring me throughout this semester. Their confidence in me has allowed me to both contribute positively and enjoy my contributions.

References

- Hacker, Carl D., Timothy O. Laumann, Nicholas P. Szrama, Antonello Baldassarre, Abraham Z. Snyder, Eric C. Leuthardt, and Maurizio Corbetta. "Resting state network estimation in individual subjects." *Neuroimage* 82 (2013): 616-633.
- Shruti A, Sair H, Pillai J. Limitations of rsfMRI in the setting of focal brain lesions. *Neuroimaging Clin N Am*. 2017 Nov. 27(4); 645:661.
- Bailey P, Zaca D, Basha M, Agarwal S, Gujar S, Sair H, Eng J, Pillai J. Presurgical fMRI and DTI for the Prediction of Perioperative Motor and Language Deficits in Primary or Metastatic Brain Lesions. *Journal of NeuroImaging*. 2015 July 14.

Technical Appendix

Dataset: <https://openneuro.org/datasets/ds000243/versions/00001>

Independent component analysis MATLAB software package:
<http://mialab.mrn.org/software/gift/>

Github repository: <https://github.com/ardenchew/rail>

Graphic processing unit specs: NVIDIA GeForce 9600