The Importance of Anti-Correlations in Graph Theory Based Classification of Autism Spectrum Disorder

Amirali Kazeminejad[1,2], Roberto C Sotero[1,2,3]

[1] Biomedical Engineering Graduate Program, University of Calgary, Calgary, Alberta Canada
[2] Hotchkiss Brain Institute, University of Calgary, Calgary, Alberta, Canada
[3] Department of Radiology, University of Calgary, Calgary, Alberta, Canada

Corresponding Author:
Amirali Kazeminejad[1,2]
3330 Hospital Dr NW, Calgary, AB T2N 4N1
Email address: Amirali.kazeminejad@ucalgary.ca

Abstract

In recent years, there has been a significant growth in the number of applications of machine learning (ML) techniques to the study and identification of neurological disorders. These methods rely heavily on what features are made available to the ML algorithm. Features such as graph theoretical metrics of resting-state fMRI-based brain networks have proven useful. However, the computation of functional brain networks relies on making an arbitrary choice about whether the obtained anti-correlations, representing the strengths of functional connections in the brain, should be discarded or not. In this study, we examine how this choice affects the performance of a support vector machine (SVM) model for classifying autism spectrum disorder. We extracted graph theoretical features using three different pipelines for constructing the functional network graph. These pipelines primarily used positive weights, negative weights (anti-correlations) and only the absolute value of weights of the correlation matrix derived from fMRI time-series. Our results suggest that in the presence of Global Signal Regression (GSR) the features extracted from anti-correlations play a major role in improving model performance. However, this does not undermine the importance of features from other pipelines.

Introduction

Autism Spectrum Disorder (ASD) is a neurodevelopmental condition that is growing in prevalence in recent years [1]. While it is usually diagnosed by carefully monitoring a child's behavioral development [2], recent studies have shown that brain imaging can also be used to aid in that diagnosis by identifying underlying differences between the ASD and Healthy Control (HC) Brain[3], [4].

Functional Magnetic Resonance Imaging (fMRI) is one of the most widely used tools for such studies due to its high spatial resolution. It monitors the changes in the Blood Oxygen Level Dependent (BOLD) signal which indirectly measures the neuronal activity [5]. This allows researchers to examine the human brain at a network level and understand how each brain region is connected to the rest of the brain by

analyzing the BOLD activation patterns. A variant of the technique call Resting State fMRI (rs-fMRI) has been widely used to examine brain networks while subjects are at rest [6], [7].Graph theory is one of the more novel methods being used for the network-level analysis of the brain. It provides a mathematical framework for quantifying network characteristics and quantitively analyze the differences between different brain networks [8], [9]. Machine learning is another relatively new technique that is being applied to rs-fMRI data to extract insights such as important biomarkers as well as to develop novel algorithms with the hope of automatically diagnosing brain disorders from medical imaging data. [10], [11]. More recently, the network characteristics that is provided by graph theory has been used as the input to machine learning algorithms to identify diseases such as Alzheimer's [12], Parkinson's [12], [13], and Autism [14].

One of the variables that can affect the results of the analysis is how the graph construction step was performed. After the initial preprocessing of the rs-fMRI data, a correlation matrix is created by calculating the correlation of each brain region with the other regions. The correlation matrix is then transformed into a sparse binary matrix representing the existence connections between different regions. This transformation is usually performed by a thresholding step in which only the two strongest correlations are kept [8], [9]. However, this step ignores the anti-correlations which may be biologically relevant [15]. This problem will be even more severe if the preprocessing pipeline includes the regression of the global mean signal (GSR) from the time-series, known to result in the removal of motion, cardiac and respiratory signals, which results in the presence of more anticorrelations in the correlation matrix [16]. Although there is evidence suggesting that the anticorrelations introduced through GSR have no biological basis [17], this is still a controversial subject within the field.

There has been evidence of network-level changes in the ASD brain compared to a HC brain[4], [18], [19]. Therefore, using graph theory to extract features for classification of ASD is likely to provide good results. However, this approach is affected by how the graph theory analysis is formulated. This claim was preciously examined by training a SVM classifier on graph theoretical features to classify between ASD and HC [14]. However, the methodology of that paper was limited to using only the positive correlations in order to construct the brain graph, potentially ignoring some informative connections.

In this paper, we study the effect of different approaches to handling anti-correlations for the classification accuracy of a machine learning model on the Autism Brain Imaging Data Exchange I (ABIDE I) dataset[20]. In line with a previous study which showed that segmenting the ABIDE dataset into age ranges leads to better model performance for classification of ASD [14], we split the ABIDE dataset into 5 age ranges. We then trained a Gaussian Kernel support vector machine (SVM) machine learning algorithm using features extracted from three pipelines: The positive correlation pipeline which selects only positive correlations, the anti-correlation pipeline which prioritizes negative correlations, and the absolute value pipeline which disregards the sign of the correlations. Our results suggest that for most age ranges putting more emphasis on the anti-correlations provides better features for the machine learning classification of ASD than putting the emphasis on the positive correlations. The absolute value of the correlations resulted in better average performance between the ages of 15-30 years. However, it failed to a statistically significant difference from the anti-correlation model's performance.

Materials & Methods

Data & Preprocessing:

This study used a publicly available dataset from the ABIDE Initiative [21]. To ensure that our results were not affected by any custom preprocessing pipeline, we used the preprocessed data provided by ABIDE in the C-PAC [22] pipeline. The preprocessing included the following steps. The Analysis of Functional NeuroImages (AFNI)[23] software was used for removing the skull from the images. The brain was segmented into three tissues using FMRIB Software Library (FSL) [24]. The images were then normalized to the MNI 152 stereotactic space[25][26] using Advanced Normalization Tools (ANTs) [27]. Functional preprocessing included motion and slice-timing correction as well as the normalization of voxel intensity. Nuisance signal regression included 24 parameters for head motion, CompCor[28] with 5 principal components for tissue signal in Cerebrospinal fluid and white matter, linear and quadratic trends for Low-frequency drifts and a global bandpass filter (0.01 to 0.1 Hz). These images where then co-registered to their anatomical counterpart by FSL. They were then normalized to the MNI 152 space using ANTs. The average voxel activity in each Region of Interest (ROI) of the AAL atlas [29] was then extracted as the time-series for that region. Some of the downloaded timeseries contained only zero values. These timeseries were discarded from further analysis.

The subjects were then split into five age groups. The age range in each age group was initially kept at 5 however this resulted in some groups having very few subjects. The final split contained the following age groups: 5-10 years, 10-15 years, 15-20 years, 20-30 years, >30 years.

Table 1 shows the demographics for the analyzed subjects.

Network and Graph Extraction and Feature Extraction:

To construct the brain network, the timeseries for each atlas ROI where correlated, using bend correlation [30], with the other regions. The strengths of these correlations were used as the strengths of the connection between different ROIs. In graph terms, this represents a fully connected graph with each of the nodes being in the center of the corresponding ROIs and each edge weight is the correlation between the two nodes on the opposite ends of the edge. Three different pipelines were obtained based on selecting only the positive values, only the negative values, or taking the absolute value of the correlation matrix. These graphs were then subjected to a thresholding step in which only the 20% strongest connections were set to one and the other edges were discarded. This resulted in binary and sparser graphs. This step was done because binary graphs have been shown to have more easily defined null models and are more easily characterizable [8]. Several measures of integration (characteristic path length and efficiency), segregation (clustering coefficient and transitivity) and centrality (betweenness centrality, eigenvector centrality, participation coefficient and within module z-score) were then extracted from the binary brain network graph. These steps were done using MATLAB toolbox GraphVar [31]. This resulted in 3 feature sets of 817 features for each group.

Model Training and Evaluation:

In this study, we used the SVM implementation in the Python scikit-learn library [32]. The SVM used a gaussian kernel with the Gamma and C values set to 0.05 and 10 respectively. Feature selection was performed following the sequential forward algorithm [33]. The model training process is as follows. First,

each feature was standardized by removing the mean and scaling to unit variance. Then the feature selection algorithm tested each feature individually and selected the one that best classifies the data. That feature was kept, and the process was repeated for all other features to find which one, when introduced alongside the previously kept feature(s), achieves the best classification. This step was repeated 30 times for a total of 30 selected features. Accuracy and F1 scores were calculated by evaluating the models based on 10-fold Cross-Validation (CV) [34]. To statistically compare the performance of the models, a 10 by 10 cross validation t-test was used in which each model was evaluated 10 times using 10-fold CV and the resulting 100 scores were subjected to a Welch's t-test [35] between different models. False-discovery rate (FDR) correction was carried out using the Benjamini-Hochberg [36] method to correct for multiple comparisons.

A full overview of the methodology is illustrated in Figure 1.

Table 1. Subject Demographics, Abbreviations: F: Female, M: Male, ASD: Autism Spectrum Disorder, HC: Healthy Control, SD: Standard Deviation

Age Range (Years)	Sex	Diagnostic Group	Age Mean (Years)	Age SD	Number of Subjects
5_10	F	ASD	8.87	0.65	7
		HC	8.98	0.63	16
	M	ASD	8.71	0.92	41
		HC	8.85	0.89	45
10_15	F	ASD	12.63	1.58	23
		HC	13.17	1.26	44
	M	ASD	12.61	1.37	133
		HC	12.56	1.43	141
15_20	F	ASD	17.04	1.12	8
		HC	17.60	1.45	13
	M	ASD	17.16	1.38	81
		HC	16.93	1.33	84
20_30	F	ASD	23.66	3.05	7
		HC	24.13	3.40	11
	M	ASD	24.25	3.35	47
		HC	24.19	2.81	69
>30	F	ASD	38.59	6.50	3
		HC	32.00	_	1
	M	ASD	39.58	9.32	25
		HC	37.10	6.66	18

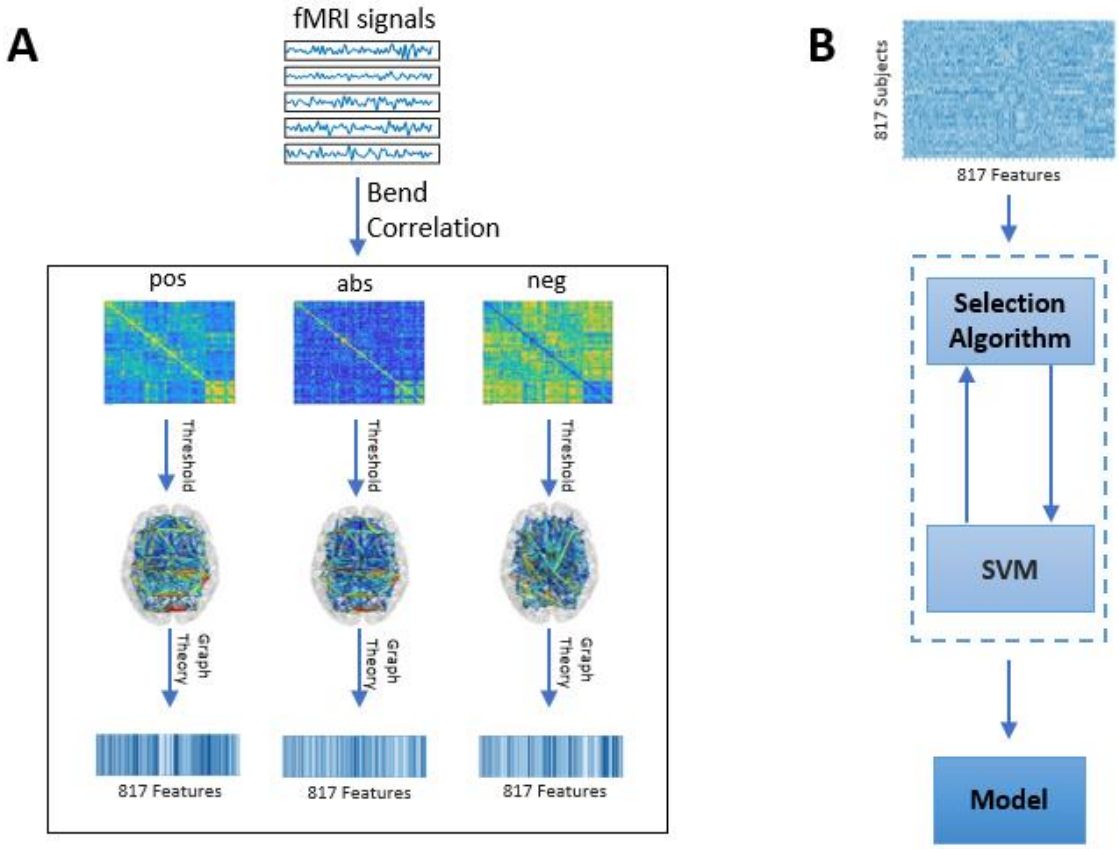

Figure 1. Graphical framework of the experiment. A. By averaging the BOLD activity in each ROI in the AAL atlas, a time series is extracted representing brain activity in that region;Using bend correlation, a connectivity matrix is generated from the ROI time series quantifying the connectivity level between individual ROIs; The connectivity matrix is transformed using three pipelines: Pos. No change to the connectivity matrix, Neg: Multiplying the connectivity matrix by -1, Abs: calculating the absolute value of the matrix. Then, by treating the ROIs as graph nodes and the connectivity matrix as graph weights the brain network is expressed in graph form; A threshold is applied to keep only the 20% strongest connections; Graph theoretical analysis is applied to the resulting graph from to obtain a feature vector(817 features) for each subject; Right Column: B. Features extracted from the left column are input to a wrapper method called sequential feature selection which is applied to choose a handful of features that contribute to the highest classification accuracy; The resulting feature subset is passed to a SVM which trains a model to distinguish between ASD and HC.

Results

Figure 2 shows the comparison of the performance of the 3 pipelines. We defined significance as $p<0.05$ for the corrected values. For the 5-10 years range, the anti-correlation pipeline significantly outperformed the other pipelines by ~5%. The anti-correlation pipeline also preformed significantly better for the 10-15 years range, improving the performance by ~5%. No model was able to achieve statistically significant improvements in the 15-20-year range before or after correcting for FDR. However, the absolute value model preformed better on average and our statistical test resulted in an uncorrected p-value of 0.051

when comparing the absolute value pipeline to the anti-correlation pipeline. In the 20-30 years range, the positive correlation model performed significantly worst than the others. While the absolute value model had a higher average F1 score, there was no sign of statistical differences between its performance and the anti-correlation network. For the >30 range, the anti-correlation model performed significantly better than the absolute value model. However, while it preformed better on average compared to the positive correlation model, there was no statistical evidence for a difference in performance in this case.

As simple accuracy measurements may not be sufficient for model evaluation when there is an inconsistency in the number of samples per class, we also calculated the sensitivity and specificity of our models for each pipeline. Accuracy represents the fraction of correct classification to all classifications, Sensitivity represents the ratio of ASD patients who were correctly identified to all ASD patients, Specificity represents the fraction of HC patients correctly identified to all HC patients. This is tabulated in Table 2.

Table 2. Detailed Model Performance; Abbreviations: Pos: Positive Correlation Pipeline, Neg: Anti-Correlation Pipeline, Abs: Absolute Value Pipeline

Age Range (Years)	Pipeline	Accuracy	Sensitivity	Specificity	Number of Features
5-10	**Pos**	0.862	0.918	0.792	24
	Neg	0.899	0.918	0.875	8
	Abs	0.881	0.951	0.792	13
10-15	**Pos**	0.737	0.774	0.692	19
	Neg	0.763	0.796	0.724	28
	Abs	0.719	0.763	0.667	21
15-20	**Pos**	0.821	0.848	0.791	14
	Neg	0.821	0.838	0.802	13
	Abs	0.842	0.859	0.824	27
20-30	**Pos**	0.870	0.939	0.768	22
	Neg	0.884	0.988	0.732	14
	Abs	0.870	0.951	0.750	12
>30	**Pos**	0.962	0.905	1.000	19
	Neg	0.962	0.905	1.000	14
	Abs	0.942	0.857	1.000	13

The feature selection algorithm selected less features on average (15.4 features) for the anti-correlation pipeline than the correlation (19.6 features) and absolute value (17.2 features). However, we were not able to establish statistical significance for the difference in the number of features used in any pair of pipelines over all 5 age-ranges.

The models generally demonstrated better sensitivity than specificity, suggesting they are more reliable in diagnosing the condition rather than dismissing it. Although the average accuracy results were mostly

in line with the f1-score results, the anti-correlation pipeline resulted in a higher average accuracy for the 20-30 age-range whereas the absolute value pipeline resulted in a higher f1-score in the same age range.

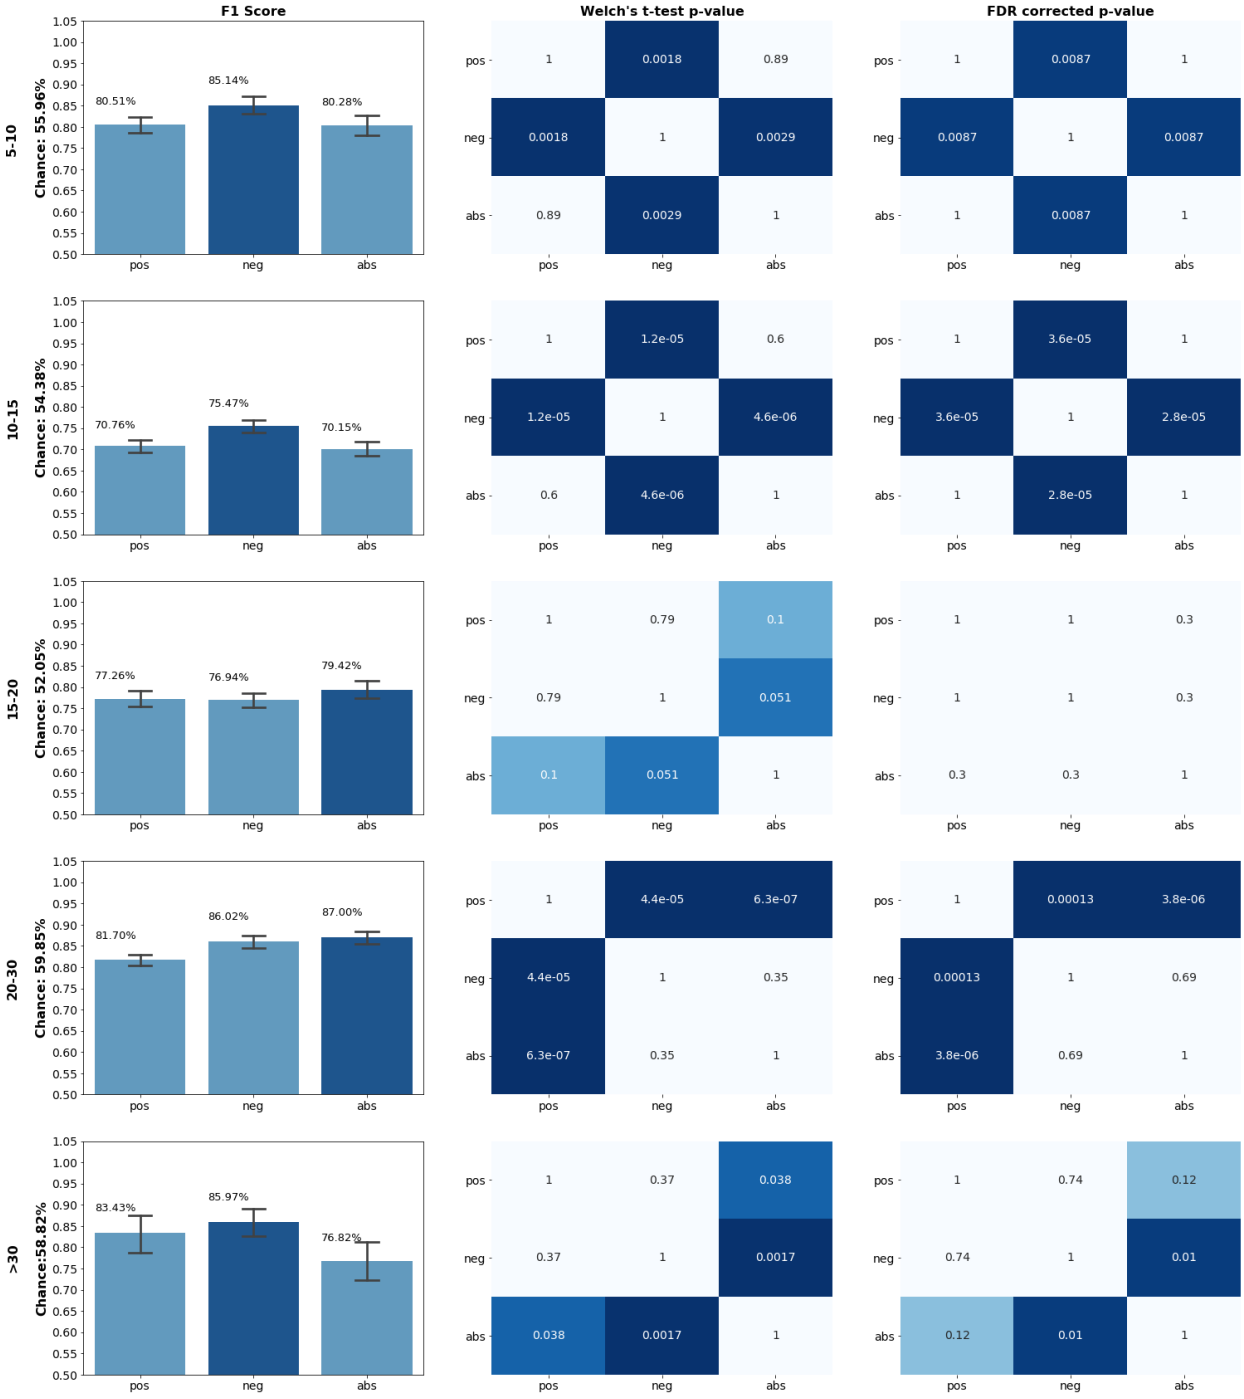

Figure 2. Performance Comparison. Left Column: The average F1-Score of the model based on a 10x10 cross validation (CV) test, the highest preforming model is highlighted in dark blue; Center Column: Uncorrected Welch's t-test p-value based on 10x10 CV, significance ($p<0.05$) is shown with darker shades of blue; Right Column: FDR Corrected p-values for the same test based on the Benjamini-Hochberg Method

Figure 3 shows the percentage of each metric group in different age ranges. Two immediate observations can be made when looking at this data. First, measures of centrality (red) appear to be the most used regardless of the pipeline in almost all cases other than the anti-correlation pipeline for the 20-30 years range. Second, in 3 age rages the anti-correlation pipeline puts less emphasis on the measures of centrality and more on the measures of integration than the other pipelines. This is not true for the 5-10 and 15-20 years range. In the former, the anti-correlation pipeline placed more emphasis on centrality measures than the absolute value pipeline. In the latter, the anti-correlation pipeline used a smaller ratio of integration metrics than both other pipelines.

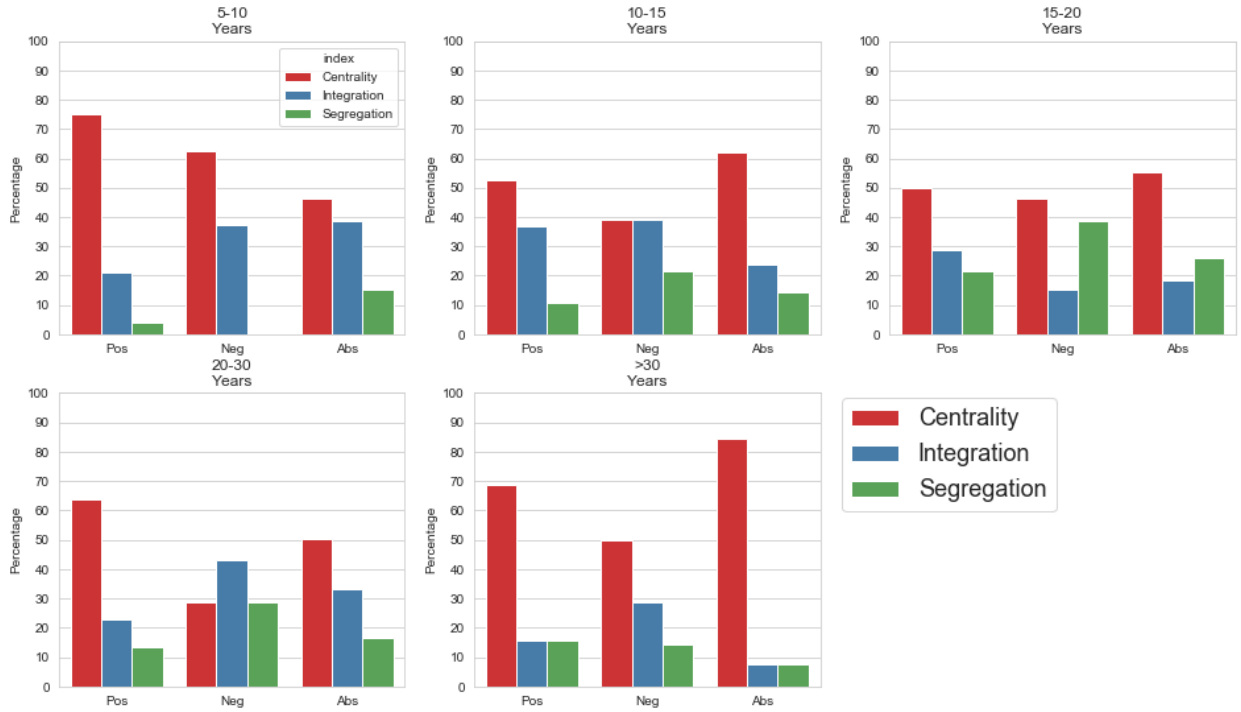

Figure 3. Importance of graph metric groups; The contribution of each metric group to the features selected for the corresponding model is calculated in percentages. Red (leftmost bar) represents measures of centrality. Blue (Middle bar) represents measures of integration and Green (Rightmost bar) represents the measures of segregation.

Figure 4 further expands the scope and visualizes the percentage of time each individual metric was selected in a model. This allows us to further examine the anomalies observed in Figure 3. In the 5-10 years range, the anti-correlation model put a much higher emphasis on participation coefficient, a centrality measure, than in other age ranges. Likewise, in the 15-20 years range, the same pipeline selected less local efficiency metrics, a measure of integration, than in other ranges. Another interesting result is that most selected metrics were derived from local measures. With the exceptions being global clustering coefficient in the 15-30 years ranges, transitivity in the 10-20 years ranges and small world propensity in the >30 years range.

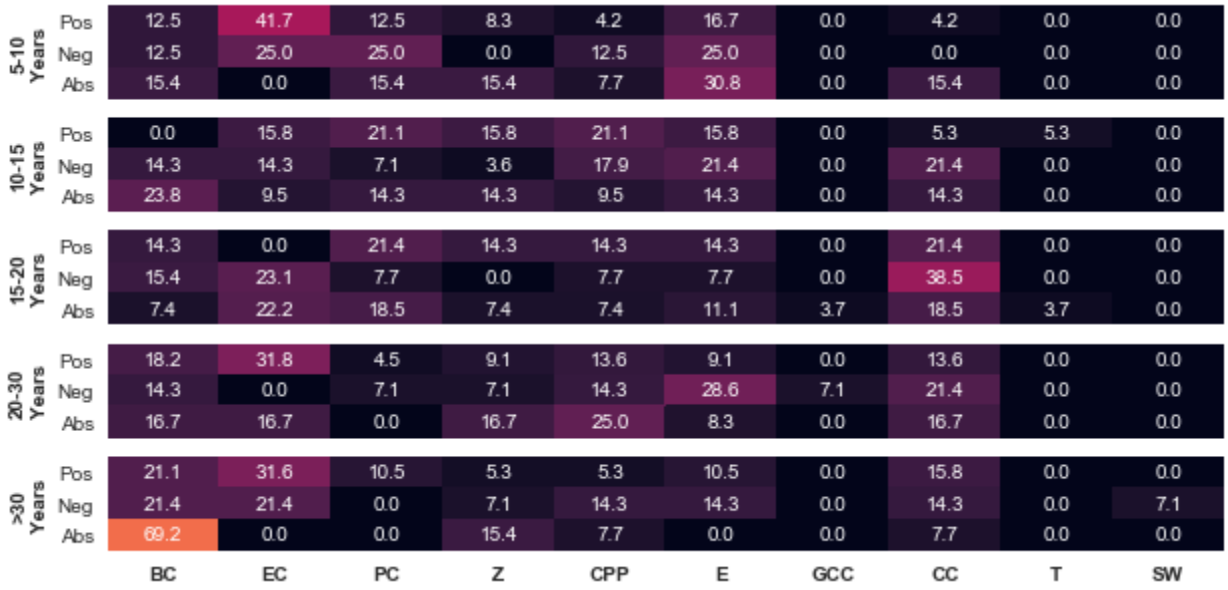

Figure 4. Importance of individual graph metrics. Unless specified, the metric is a local metric. BC: Betweenness Centrality; EC: Eigenvector Centrality; PC: Participation Coefficient; Z: Within-module Z-score; CPL: Characteristic Path Length; E: Efficiency; GCC: Global Clustering Coefficient; CC: Clustering Coefficient; T: Transitivity; SW: Small-world-propensity

Discussion

The ABIDE dataset contains a large amount of heterogeneity in the form of a large age-range and data collected from multiple sites. This heterogeneity can affect the performance of the model and the statistical analysis. An SVM model trained on this data is prone to biases based on the different subject diagnostic group ratio for site and age and different imaging methodologies in different sites. We used the publicly available preprocessed version of the data in which all sites were subjected to the same preprocessing pipelines to minimize the between site differences in imaging protocols. While the cross-validation scheme may somewhat offset the issue of the heterogeneity in age-ranges due to randomly selecting subsets of the available data, a previous study showed that the accuracy of the models can be further improved by segmenting the data into 5 age ranges and training a model for each age range [14].

In all but two age ranges (15-20 years and >30 years), our results suggest that in the presence of global signal regression, the anti-correlations of the brain's connectivity matrix may play a significantly larger role in differentiating between ASD and HC than the correlation network. Furthermore, it also significantly outperformed the absolute value pipeline in all but two ranges (15-20, 20-30). However, in these cases there was no evidence of statistically significant performance improvement over the anti-correlation network when using the absolute value pipeline. Interestingly, previous studies have concluded that the use of absolute value graph metrics in the presence of GSR may be compromised [37]. This was attributed to the fact that in the presence of GSR, the topologies of the anti-correlation and correlation matrix will be mixed when using an absolute value pipeline as the anti-correlations have comparable magnitude to the correlations.

On the other hand, the anti-correlation pipeline was able to use less features to obtain the best results. This suggests that there may be stronger indicators of ASD in the anti-correlation network than in the networks

constructed using other pipelines. One reason for the higher performance of the anti-correlation method mentioned in the two preceding paragraphs could be that contrary to some previous studies [16][17], GSR introduced anti-correlations may in fact stem from real neurological basis and not spurious [38]. Unfortunately, as we do not know of any gold standards to test this hypothesis in real data, it should only be interpreted as a speculation made based on the results of this study.

Analyzing the contributions of different graph metrics to the feature pool in different pipelines may also provide useful insights as to why the anti-correlation network performed better than the others. Centrality measures are used to identify nodes that are interacting with many other nodes, thus playing an important role in the functional network [39]. Measures of segregation are generally used to identify specialized sub-modules in the brain network that are highly interconnected [8]. Integration metrics measure the interconnectivity of the network and it's ability to process information from different brain regions [40]. Figure 3 further illustrates that even in the presence of anti-correlations, measures of centrality provided the greatest number of features in the best pipeline for each age range. This was in line with a previous study focusing only on the positive correlation pipelines [14]. Interestingly, the anti-correlation model showed less emphasis on centrality measures and more emphasis on integration measures. This suggests that there are lower (higher) changes in the network centrality (integration) of the anti-correlation network of ASD vs HC. Further suggesting that there may be useful information being discarded in studies that only use the established method of removing negative connections from the analysis altogether.

Figure 4 may provide some more clarifications about the previous paragraph. The anti-correlation models generally tended towards selecting a relatively low number of participation coefficient features. However, for the 5-10 years range, this metric was heavily used. Participation coefficient is a centrality metric quantifying the interaction level of a node with other nodes outside of its own module. The relative unimportance of this measure suggests that there may be more significant intermodular connection changes in the anti-correlation functional network of children with ASD than other age-ranges. The underlying cause behind the integration measure anomaly of the 15-20 range is harder to pinpoint as the anti-correlation model generally put relatively high importance on both E and CPL. However, this model generally put more emphasis on efficiency, a measure quantifying the communication efficiency of the functional network. Efficiency is calculated using the inverse CPLs and thus they are closely related to each other. However, Efficiency is a more biologically relevant metric because of it's ability to deal with disconnected graphs [41]. The lower level of importance given to this metric suggests that, in the 15-20 years age range, the overall ability of the brain to integrate different modules in its functional network is less affected in ASD patients. There are a few limitations that should be considered when interpreting this study. This study used all the data available in the ABIDE dataset apart from a few subjects that were discarded due to missing data in their time-series. Moreover, we used preprocessing and different age-ranges to partially address the data heterogeneity problem. While it was previously shown that this approach improves classification accuracy [14], this study did not investigate the exact effect of this approach on data heterogeneity. Therefore, further analysis may be needed to conclude that the heterogeneity has been addressed to a satisfactory degree. Furthermore, our results show that there may be useful features in the positive correlation model, especially in the 15-30 range where the average absolute correlation model exhibited higher average performance. This suggests that the optimal feature

creation pipeline should consider both anti-correlations and positive correlations rather than focusing on one. Finding this optimal pipeline was not in the scope of this study however it will be an interesting avenue of investigation.

Conclusions

In this study, we showed that in the presence of GSR, the functional anti-correlation network provides stronger biomarkers for diagnosing ASD using machine learning. However, between the ages of 15-30 years, the absolute value of the correlations proved to have higher average accuracy but did not show statistically significant differences with the performance of the anti-correlation network-based model. To the best of our knowledge, this is the first study analyzing anti-correlation-based models in autism spectrum disorder. As a result, further investigations will be needed to fully understand the importance of this network in ASD.

Acknowledgements

This work was partially funded by the University of Calgary, Department of Biomedical Engineering, Research Excellence award.

References

[1] B. Zablotsky, L. I. Black, M. J. Maenner, L. A. Schieve, and S. J. Blumberg, "Estimated Prevalence of Autism and Other Developmental Disabilities Following Questionnaire Changes in the 2014 National Health Interview Survey.," *Natl. Health Stat. Report.*, no. 87, pp. 1–20, Nov. 2015.

[2] W. J. Barbaresi, S. K. Katusic, and R. G. Voigt, "Autism: a review of the state of the science for pediatric primary health care clinicians.," *Arch. Pediatr. Adolesc. Med.*, vol. 160, no. 11, pp. 1167–75, Nov. 2006.

[3] G. S. Dichter, "Functional magnetic resonance imaging of autism spectrum disorders.," *Dialogues Clin. Neurosci.*, vol. 14, no. 3, pp. 319–51, Sep. 2012.

[4] D. J. Bos *et al.*, "Developmental differences in higher-order resting-state networks in Autism Spectrum Disorder," *NeuroImage Clin.*, vol. 4, pp. 820–827, 2014.

[5] R. C. Sotero and N. J. Trujillo-Barreto, "Modelling the role of excitatory and inhibitory neuronal activity in the generation of the BOLD signal," *Neuroimage*, vol. 35, no. 1, pp. 149–165, Mar. 2007.

[6] M. D. Fox and M. Greicius, "Clinical applications of resting state functional connectivity," *Front. Syst. Neurosci.*, vol. 4, no. June, p. 19, 2010.

[7] M. P. van den Heuvel and H. E. Hulshoff Pol, "Exploring the brain network: A review on resting-state fMRI functional connectivity," *Eur. Neuropsychopharmacol.*, vol. 20, no. 8, pp. 519–534, Aug. 2010.

[8] M. Rubinov and O. Sporns, "Complex network measures of brain connectivity: Uses and interpretations," *Neuroimage*, vol. 52, no. 3, pp. 1059–1069, 2010.

[9] E. Bullmore and O. Sporns, "Erratum: Complex brain networks: graph theoretical analysis of structural and functional systems," *Nat. Rev. Neurosci.*, vol. 10, no. 4, pp. 312–312, Apr. 2009.

[10] M. H. Lee, C. D. Smyser, and J. S. Shimony, "Resting-state fMRI: a review of methods and clinical applications.," *AJNR. Am. J. Neuroradiol.*, vol. 34, no. 10, pp. 1866–72, Oct. 2013.

[11] F. Pereira, T. Mitchell, and M. Botvinick, "Machine learning classifiers and fMRI: A tutorial overview," *Neuroimage*, vol. 45, no. 1, pp. S199–S209, Mar. 2009.

[12] A. Khazaee, A. Ebrahimzadeh, and A. Babajani-Feremi, "Application of advanced machine learning methods on resting-state fMRI network for identification of mild cognitive impairment and Alzheimer's disease," *Brain Imaging Behav.*, vol. 10, no. 3, pp. 799–817, Sep. 2016.

[13] A. Kazeminejad, S. Golbabaei, and H. Soltanian-Zadeh, "Graph theoretical metrics and machine learning for diagnosis of Parkinson's disease using rs-fMRI," in *2017 Artificial Intelligence and Signal Processing Conference (AISP)*, 2017, pp. 134–139.

[14] A. Kazeminejad and R. C. Sotero, "Topological Properties of Resting-State fMRI Functional Networks Improve Machine Learning-Based Autism Classification," *Front. Neurosci.*, vol. 12, p. 1018, Jan. 2019.

[15] M. D. Fox, A. Z. Snyder, J. L. Vincent, M. Corbetta, D. C. Van Essen, and M. E. Raichle, "The human brain is intrinsically organized into dynamic, anticorrelated functional networks.," *Proc. Natl. Acad. Sci. U. S. A.*, vol. 102, no. 27, pp. 9673–8, Jul. 2005.

[16] K. Murphy and M. D. Fox, "Towards a consensus regarding global signal regression for resting state functional connectivity MRI," *Neuroimage*, vol. 154, pp. 169–173, Jul. 2017.

[17] K. Murphy, R. M. Birn, D. A. Handwerker, T. B. Jones, and P. A. Bandettini, "The impact of global signal regression on resting state correlations: Are anti-correlated networks introduced?," *Neuroimage*, vol. 44, no. 3, pp. 893–905, Feb. 2009.

[18] J. D. Rudie *et al.*, "Altered functional and structural brain network organization in autism," *NeuroImage Clin.*, vol. 2, pp. 79–94, Jan. 2013.

[19] E. Redcay, J. M. Moran, P. L. Mavros, H. Tager-Flusberg, J. D. E. Gabrieli, and S. Whitfield-Gabrieli, "Intrinsic functional network organization in high-functioning adolescents with autism spectrum disorder," *Front. Hum. Neurosci.*, vol. 7, p. 573, Sep. 2013.

[20] A. Di Martino *et al.*, "The autism brain imaging data exchange: Towards a large-scale evaluation of the intrinsic brain architecture in autism," *Mol. Psychiatry*, vol. 19, no. 6, pp. 659–667, Jun. 2014.

[21] C. Cameron *et al.*, "The Neuro Bureau Preprocessing Initiative: open sharing of preprocessed neuroimaging data and derivatives," *Front. Neuroinform.*, vol. 7, 2013.

[22] C. Cameron *et al.*, "Towards Automated Analysis of Connectomes: The Configurable Pipeline for the Analysis of Connectomes (C-PAC)," *Front. Neuroinform.*, vol. 7, 2013.

[23] R. W. Cox, "AFNI: software for analysis and visualization of functional magnetic resonance neuroimages.," *Comput. Biomed. Res.*, vol. 29, no. 3, pp. 162–73, Jun. 1996.

[24] S. M. Smith *et al.*, "Advances in functional and structural MR image analysis and implementation as FSL," *Neuroimage*, vol. 23, pp. S208–S219, Jan. 2004.

[25] G. Grabner, A. L. Janke, M. M. Budge, D. Smith, J. Pruessner, and D. L. Collins, "Symmetric Atlasing and Model Based Segmentation: An Application to the Hippocampus in Older Adults," Springer, Berlin, Heidelberg, 2006, pp. 58–66.

[26] J. Mazziotta *et al.*, "A four-dimensional probabilistic atlas of the human brain.," *J. Am. Med. Inform. Assoc.*, vol. 8, no. 5, pp. 401–30, 2001.

[27] B. B. Avants, N. J. Tustison, G. Song, P. A. Cook, A. Klein, and J. C. Gee, "A reproducible evaluation of ANTs similarity metric performance in brain image registration," *Neuroimage*, vol. 54, no. 3, pp. 2033–2044, Feb. 2011.

[28] Y. Behzadi, K. Restom, J. Liau, and T. T. Liu, "A component based noise correction method (CompCor) for BOLD and perfusion based fMRI.," *Neuroimage*, vol. 37, no. 1, pp. 90–101, Aug. 2007.

[29] N. Tzourio-Mazoyer *et al.*, "Automated Anatomical Labeling of Activations in SPM Using a Macroscopic Anatomical Parcellation of the MNI MRI Single-Subject Brain," *Neuroimage*, vol. 15, no. 1, pp. 273–289, Jan. 2002.

[30] R. R. Wilcox, "The percentage bend correlation coefficient," *Psychometrika*, vol. 59, no. 4, pp. 601–616, Dec. 1994.

[31] J. D. Kruschwitz, D. List, L. Waller, M. Rubinov, and H. Walter, "GraphVar: a user-friendly toolbox for comprehensive graph analyses of functional brain connectivity.," *J. Neurosci. Methods*, vol. 245, pp. 107–15, 2015.

[32] F. Pedregosa *et al.*, "Scikit-learn: Machine Learning in Python," *J. Mach. Learn. Res.*, vol. 12, no. Oct, pp. 2825–2830, 2011.

[33] P. Pudil, J. Novovičová, and J. Kittler, "Floating search methods in feature selection," *Pattern Recognit. Lett.*, vol. 15, no. 11, pp. 1119–1125, Nov. 1994.

[34] M. Stone, "Cross-Validatory Choice and Assessment of Statistical Predictions," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 36. WileyRoyal Statistical Society, pp. 111–147, 1974.

[35] B. L. WELCH, "THE GENERALIZATION OF 'STUDENT'S' PROBLEM WHEN SEVERAL DIFFERENT POPULATION VARLANCES ARE INVOLVED," *Biometrika*, vol. 34, no. 1–2, pp. 28–35, Jan. 1947.

[36] Y. Benjamini and Y. Hochberg, "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 57. WileyRoyal Statistical Society, pp. 289–300, 1995.

[37] A. J. Schwarz and J. McGonigle, "Negative edges and soft thresholding in complex network analysis of resting state functional connectivity data," *Neuroimage*, vol. 55, no. 3, pp. 1132–1146, Apr. 2011.

[38] M. D. Fox, D. Zhang, A. Z. Snyder, and M. E. Raichle, "The Global Signal and Observed Anticorrelated Resting State Brain Networks," *J. Neurophysiol.*, vol. 101, no. 6, pp. 3270–3283, Jun. 2009.

[39] M. P. van den Heuvel and O. Sporns, "Network hubs in the human brain.," *Trends Cogn. Sci.*, vol. 17, no. 12, pp. 683–96, Dec. 2013.

[40] O. Sporns, "Structure and function of complex brain networks.," *Dialogues Clin. Neurosci.*, vol. 15, no. 3, pp. 247–62, Sep. 2013.

[41] J. Wang, X. Zuo, and Y. He, "Graph-based network analysis of resting-state functional MRI.," *Front. Syst. Neurosci.*, vol. 4, p. 16, 2010.