

# BOOTSTRAPPING GRAPH CONVOLUTIONAL NEURAL NETWORKS FOR AUTISM SPECTRUM DISORDER CLASSIFICATION

Rushil Anirudh and Jayaraman J. Thiagarajan

Lawrence Livermore National Laboratory  
Email: {anirudh1, jjayaram}@llnl.gov

## ABSTRACT

Using predictive models to identify patterns that can act as biomarkers for different neuropathological conditions is becoming highly prevalent. In this paper, we consider the problem of Autism Spectrum Disorder (ASD) classification where previous work has shown that it can be beneficial to incorporate a wide variety of meta features, such as socio-cultural traits, into predictive modeling. A graph-based approach naturally suits these scenarios, where a contextual graph captures traits that characterize a population, while the specific brain activity patterns are utilized as a multivariate signal at the nodes. Graph neural networks have shown improvements in inferencing with graph-structured data. Though the underlying graph strongly dictates the overall performance, there exists no systematic way of choosing an appropriate graph in practice, thus making predictive models non-robust. To address this, we propose a bootstrapped version of graph convolutional neural networks (G-CNNs) that utilizes an ensemble of weakly trained G-CNNs, and reduce the sensitivity of models on the choice of graph construction. We demonstrate its effectiveness on the challenging Autism Brain Imaging Data Exchange (ABIDE) dataset and show that our approach improves upon recently proposed graph-based neural networks. We also show that our method remains more robust to noisy graphs.

**Index Terms**— graph convolutional networks, autism spectrum disorder classification, fMRI, population graphs

## 1. INTRODUCTION

Modeling relationships between functional or structural regions in the brain is a significant step towards understanding, diagnosing and eventually treating a gamut of neurological conditions including epilepsy, stroke, and autism. A variety of sensing mechanisms, such as functional-MRI, Electroencephalography (EEG) and Electrocochography (ECoG), are commonly adopted to uncover patterns in both brain structure and function. In particular, the resting state fMRI [1] has been

proven effective in identifying diagnostic biomarkers for mental health conditions such as the Alzheimer disease [2] and autism [3]. At the core of these neuropathology studies are predictive models that map variations in brain functionality, obtained as time-series measurements in regions of interest, to clinical scores. For example, the Autism Brain Imaging Data Exchange (ABIDE) is a collaborative effort [4], which seeks to build a *data-driven* approach for autism diagnosis. Further, several published studies have reported that predictive models can reveal patterns in brain activity that act as effective biomarkers for classifying patients with mental illness [3].

Graphs provide a natural framework to analyze relationships in a population of patients. The advances in graph signal processing and the generalization of deep neural networks to arbitrarily structured data, make graph models an attractive solution for autism diagnosis. In addition to exploiting correlations in imaging features (e.g. fMRI), the graphs could include a wide-range of non-imaging, contextual features based on more general characteristics of the subjects, including geographical, socio-cultural or gender based features. Recently it has been shown that including such information with a graph convolutional neural network (G-CNN) [5, 6], can improve classification performance [7]. Despite the applicability of graph-based models in clinical prediction, it is critical to note that the choice of graph construction is crucial to the success of this pipeline – a low-quality graph can lead sub-optimal models, sometimes even worse than simpler methods that do not utilize the meta information. However, there is currently no definitive strategy on constructing reliable population graphs, and as a result there is a need to build robustness into graph based predictive models, so that they can work with a wider variety of graphs.

**Proposed Work:** In this paper, we address these issues with a new approach to graph-based predictive modeling, which relies on generating an ensemble of population graphs. First, using a bootstrapping approach to design graph ensembles allows our predictive model to better explore connections between subjects in a large population graph that are not captured by simple heuristics. Second, G-CNNs provide a powerful computing framework to make inferences on graphs, by treating the subject-specific image features as a function,  $f : \mathcal{V} \mapsto \mathbb{R}^N$ ,

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

defined at the nodes of the population graph. We consider a variety of graphs – (a) feature based graph,  $\mathcal{G}_0$ , weighted by sex and site [7] (b) a noisy version of  $\mathcal{G}_0$ , with 30% of its edges dropped, and (c) a Naïve graph, where the adjacency matrix is identity; which is equivalent to using the imaging features alone. Our results show that in every case the proposed bootstrapped G-CNN approach improves classification performance, and provides robustness to the choice of the graph. This simple approach, improves the state of the art classification performance on the ABIDE dataset and also reduces sensitivity of the resulting model to the graph construction step. Consequently, even non-experts can design population graphs, which, with bootstrapping, can perform on par with more sophisticated graph construction strategies.

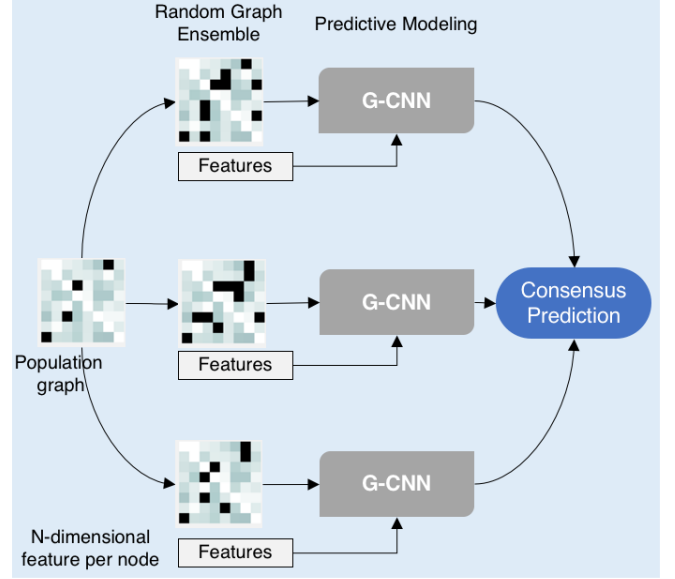
## 2. PREDICTIVE MODELING WITH ENSEMBLES

In this section, we describe our approach for classifying subjects with autism. Figure 1 illustrates an overview of the proposed approach for ASD classification. As it can be observed, the pipeline requires an initial population graph and the features at each node (i.e. subject) in the graph as inputs. Subsequently, we create an ensemble of randomized graph realizations and invoke the training of a graph CNN model for every realization. The output layer of these neural networks implement the *softmax* function, which computes the probabilities for class association for each node. Finally, the consensus module fuses the decisions from the ensemble to obtain the final class label. Next we outline the process of graph construction and training strategy for the proposed approach.

### 2.1. Population Graph Construction

A classifier trained on the imaging features alone fails to incorporate contextual non-imaging/meta information that can be critical to discriminate between different classes. For example, it is likely that there is discrepancy in some aspects of data collection at different sites, or the gender of the subject is important in generalizing autism spectrum disorder predictors. It is non-trivial to directly incorporate such information into the subject features, but a graph can be a very intuitive way to introduce these relationships into the learning process.

An inherent challenge is that the results obtained are directly dictated by a weighted graph defined for the analysis. Consequently, designing appropriate weighted graphs that capture the geometric structure of data is essential for meaningful analysis. In our context, the population graph construction determines how two subjects are connected, so that context information could be shared between them. We follow the graph construction strategy used in [7], which uses a combination of imaging features, gender and site information. First the sex-site graph is obtained as follows: if two subjects have the same gender, they are given a score of  $s_{sex} = \lambda_1 > 1$ , and 1 if they are not. Similarly, the subjects were given a score of



**Fig. 1.** An overview of the proposed approach for predictive modeling with non-imaging features encoded as a population graph and imaging features used as functions at the nodes. We construct a randomized ensemble of population graphs, employ graph CNNs and utilize a consensus strategy to perform the actual classification.

$s_{site} = \lambda_2 > 1$ , if they were processed at the same site, and 1 if not. Next, a linear-kernel graph is computed where the edge weights of the graph are the Euclidean dot product or a linear kernel, between connectivity features from two different ROIs, for a given subject. As expected, this graph does not provide any additional information because it is directly based on the features that were defined at the nodes.

### 2.2. Graph Convolutional Neural Networks

Convolutional neural networks enable extraction of statistical features from structured data, in the form of local stationary patterns, and their aggregation for different semantic analysis tasks, e.g. image recognition or activity analysis. When the signal of interest does not lie on a regular domain, for example graphs, generalizing CNNs is particularly challenging due to the presence of convolution and pooling operators, typically defined on regular grids.

Existing work on generalizing CNNs to graphs can be categorized into *spectral* approaches [8, 9], which operate on an explicit spectral representation of the graphs, and *non-spectral* approaches that define convolutions directly on the graphs using spatial neighborhoods [10, 11]. Spectral approaches, as the name suggests operate using the spectral representation of graph signals, defined using the eigenvectors of the graph Laplacian. For example, in [8], convolutions are realized as multiplications in the graph Fourier domain. However, since the filters cannot be spatially localized on arbitrary graphs, this relies on explicit computation of the spectrum based on matrix inversion. Consequently, special families of spatially localized

filters have been considered. Examples include the localization technique in [12], and Chebyshev polynomial expansion based localization in [9]. Building upon this idea, Kipf and Welling [13] introduced graph convolutional neural networks (G-CNN) using localized first-order approximation of spectral graph convolutions, wherein the filters operate within an one-step neighborhood, thus making it scalable to even large networks.

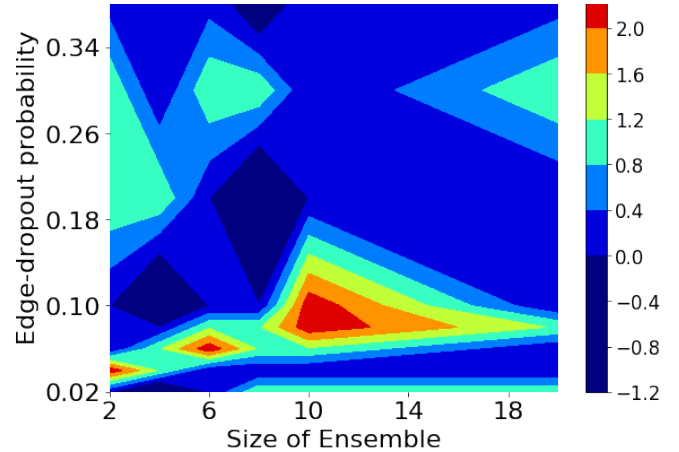
### 2.3. Ensemble Learning

As described in the previous section, population graphs provide a convenient way to incorporate non-imaging features into the predictive modeling framework, where the connectivity features from fMRI data are used as a function on the graph. While G-CNN can automatically infer spectral filters for achieving discrimination across different classes, the sensitivity of its performance to the choice of the population graph is not straightforward to understand. Consequently, debugging and fine-tuning these networks can be quite challenging. A classical approach to building robust predictive models is to infer an ensemble of weak learners from data and then fuse the decisions using a consensus strategy. The need for ensemble models in supervised learning has been well established [14]. A variety of bootstrapping techniques have been developed, wherein multiple “weak” models, that focus on subsets of the data or different aspects of the task, are inferred and finally fused for effective prediction. The intuition behind the success of this approach is that different models may have a similar training error when learned using a subset of training samples, but their performance on test data can be different since they are optimized for different regions of the input space.

We employ a similar intuition to building models with population graphs, wherein the quality of different similarity metrics for graph construction can lead to vastly different predictive models. More specifically, starting with a population graph, we create an ensemble of graphs,  $\{\mathcal{G}_p\}_{p=1}^P$ , by dropping out a pre-defined factor of edges randomly. In this paper, we use a uniform random distribution for the dropout, though more sophisticated weighted distributions could be used. For each of the graphs in the ensemble, we build a G-CNN model with the connectivity features as the  $N$ -dimensional multivariate function at each node. The output of each of the networks is a softmax function for each node, indicating the probability for the subject to be affected by ASD. Note that, unlike conventional ensemble learners, we do not subsample data, but only drop edges from the population graph. Conceptually, this is similar to the idea of using multiple attention heads in the recently successful attention models in deep learning [15]. Given the logits from all the weak learners, we employ simple consensus strategies such as averaging of the probabilities estimated by each of the G-CNNs for a test subject. As we will show in our experiments, the proposed ensemble approach boosts the performance of all the population graph

construction strategies considered.

The two main hyper-parameters in our approach are the size of the ensemble, and the edge-dropout probability. In general we observe that a large ensemble (around 20 graphs) works effectively, while a dropout probability of around 0.2-0.3 improves performance consistently. Figure 2 shows improvement in performance for a particular train/test split in the ABIDE dataset used in our experiments. As it is seen, there are several combinations that lead to an improvement in performance, while not requiring any additional information.



**Fig. 2. Performance improvement with Bootstrapping:** An exhaustive hyper-parameter search demonstrating the improvement in performance obtained by size of the graph ensemble, and edge dropout probability. We show results for a particular test/train split.

## 3. EXPERIMENTS

In this section, we describe our experiments to evaluate the performance of the proposed ensemble G-CNN approach in ASD classification.

**The ABIDE dataset:** We present our results on the Autism Brain Imaging Data Exchange (ABIDE) dataset [4] that contains resting state fMRI data (rs-fMRI) for 1112 patients, as part of the preprocessed connectomes project<sup>1</sup>. The pre-processing includes slice-timing correction, motion correction, and intensity normalization, depending on the pipeline used. We follow the same preprocessing pipeline (C-PAC) and atlases (Harvard-Oxford) as described in [16], in order to facilitate easy comparison – this resulted in a dataset with 872 of the initial 1112 patients available, from 20 different sites. The task is to diagnose a patient as being of two classes – Autism Spectrum Disorder (ASD) or Typical Control (TC). The dataset along with different pre-processing strategies are available via the nilearn python package<sup>2</sup>.

The resulting data per subject consists of the mean time series obtained from the rs-fMRI for each Region of Interest

<sup>1</sup><http://preprocessed-connectomes-project.org/abide/>

<sup>2</sup><https://nilearn.github.io/introduction.html>

(ROI). In total, there are 111 ROIs for the HO atlas considered here, resulting in the data, for the  $i^{th}$  subject, of size  $\mathbb{R}^{111 \times T}$  matrix, where  $T$  is the total number of slices in the fMRI measurement. We use the same 10-fold cross validation as in [7] of the ABIDE dataset, using their publicly released code <sup>3</sup>.

**Training details and parameters:** For a fair comparison, we use the same G-CNN set up used in [7], which utilizes the original implementation from [6]. We constructed a fully graph-convolutional network, with 3 layers of 16 units each; a learning rate of 0.005 and dropout of 0.3. We used Chebyshev polynomials upto degree 3 to approximate the Graph Fourier Transform in the G-CNN in all our experiments. We also use recursive feature elimination to reduce the feature space to 2000 most important features. All our models, including the ensembles are trained for 200 epochs each.

**Ensemble G-CNN:** For a given graph  $\mathcal{G}$ , we generate  $P$  new graphs such that the new set of graphs  $\{\mathcal{G}_0, \mathcal{G}_1 \dots \mathcal{G}_P\}$  were obtained by randomly dropping 30 – 40% of the edges of  $\mathcal{G}$ . The predictions from each model obtained using the ensemble are fused as follows: we take the mean of the predictions and assign the class as the one with the largest value. In general, each member of the random ensemble behaves as a *weak learner*, where the performance for each individual network is suboptimal, but the consensus decision is better than the state-of-the-art. In general we observe that for *reliable* graphs, edge-drop probability of 0.25 – 0.3 was effective, while for noisier or sparser graphs even small perturbations of 0.05 were found to be effective in boosting performance. In both cases we observed better performance with a large ensemble of around 20 – 25 random graphs. A hyper-parameter search comparing the number of graphs and edge drop probability for a particular test/train split is shown in figure 2.

### 3.1. Classification Results

The classification accuracy obtained by using the bootstrapping approach is shown in table 1. A few observations are important to note – There is a clear and obvious advantage in using the graph based approach to classifying populations. We out perform recent state of the art method [7] for ASD classification by nearly 1.5 percentage points, without any additional information, and using the exact same training protocol with graph CNNs. This advantage is also seen in cases where the graphs are noisy, and also when no graph information is available at all. In both cases, we see the proposed approach demonstrates robustness to the changes in the graph. Secondly, the difference in performance for different graphs illustrates the importance of graph construction, further emphasizing the need for robust training strategies such as those proposed in this paper. Finally, it can be observed that the best performing split has a consistently high performance across different

Population Graph	Predictive Model	Accuracy
-	Linear SVM[16]	66.8
$\mathcal{G}$	G-CNN [7]	69.50
	Proposed	<b>70.86</b>
Naïve Graph	G-CNN [7]	66.93
	Proposed	<b>67.85</b>
Noisy $\mathcal{G}$	G-CNN [7]	66.35
	Proposed	<b>67.39</b>

**Table 1.** Classification accuracy for the ABIDE dataset ([4]) using our approach. For comparison, we report the state-of-the-art results obtained using linear SVM and G-CNN.  $\mathcal{G}$  corresponds to the graph proposed in [7], that uses the patient sex and hospital location, along with the f-MRI features. Naïve Graph refers to using an identity matrix as the adjacency, when a population graph is not available. Finally noisy  $\mathcal{G}$  refers to the case when we drop 30% of the edges from the graph.

kinds of graphs and only marginally better than the baseline, perhaps indicating that the connectivity features dictate the performance in that case.

## 4. DISCUSSION AND FUTURE WORK

In this paper we presented a training strategy for graph convolutional neural networks (G-CNNs), that have been recently proposed as a promising solution to the node classification problem in graph structured data. We focused particularly on autism spectrum disorder classification using time series extracted from resting state fMRI. While recent graph-based predictive models have shown that incorporating contextual non-imaging features such as patient sex, location of the scan etc. can improve performance, they end up being extremely sensitive to the choice of the population graph. To circumvent this challenge, we propose to use bootstrapping as a way to reduce the sensitivity of the initial graph construction step, by generating multiple random graphs from the initial population graph. We train a G-CNN for each randomized graph, and fuse their predictions at the end. These individual predictive models behave as weak learners, that can be aggregated to produce superior classification performance. The proposed work opens several new avenues of future work including (a) pursuing a theoretical justification behind the improved performance from randomized ensembles, (b) extending these ideas by using an ensemble of random binary graphs, with appropriate probability distributions, which are very cheap to construct, and (c) performing fusion in the hidden layers instead of final layer, which involves training the ensemble of networks together.

<sup>3</sup><https://github.com/parisots/population-gcn>

## 5. REFERENCES

- [1] A. M. Clare Kelly, Lucina Q. Uddin, Bharat B. Biswal, F. Xavier Castellanos, and Michael P. Milham, "Competition between functional brain networks mediates behavioral variability," *NeuroImage*, vol. 39, no. 1, pp. 527–527, January 2008.
- [2] Gang Chen, B. Douglas Ward, Chunming Xie, Wenjun Li, Zhilin Wu, Jennifer L. Jones, Malgorzata Franczak, Piero Antuono, and Shi-Jiang Li, "Classification of alzheimer disease, mild cognitive impairment, and normal cognitive status with large-scale network analysis based on resting-state functional mr imaging," *Radiology*, vol. 259, no. 1, pp. 213–221, 2011.
- [3] Mark Plitt, Kelly Anne Barnes, and Alex Martin, "Functional connectivity classification of autism identifies highly predictive brain features but falls short of biomarker standards," *NeuroImage: Clinical*, vol. 7, pp. 359 – 366, 2015.
- [4] Adriana Di Martino, Chao-Gan Yan, Qingyang Li, Erin Denio, Francisco X Castellanos, Kaat Alaerts, Jeffrey S Anderson, Michal Assaf, Susan Y Bookheimer, Mirella Dapretto, et al., "The autism brain imaging data exchange: towards a large-scale evaluation of the intrinsic brain architecture in autism," *Molecular psychiatry*, vol. 19, no. 6, pp. 659–667, 2014.
- [5] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst, "Convolutional neural networks on graphs with fast localized spectral filtering," in *Advances in Neural Information Processing SystemsNIPS*, 2016, pp. 3837–3845.
- [6] Thomas N Kipf and Max Welling, "Semi-supervised classification with graph convolutional networks," in *International Conference on Learning Representations (ICLR)*, 2017.
- [7] Sarah Parisot, Sofia Ira Ktena, Enzo Ferrante, Matthew Lee, Ricardo Guerrero Moreno, Ben Glocker, and Daniel Rueckert, "Spectral graph convolutions on population graphs for disease prediction," *arXiv preprint arXiv:1703.03020*, 2017.
- [8] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun, "Spectral networks and locally connected networks on graphs," *arXiv preprint arXiv:1312.6203*, 2013.
- [9] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst, "Convolutional neural networks on graphs with fast localized spectral filtering," in *Advances in Neural Information Processing Systems*, 2016, pp. 3844–3852.
- [10] David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams, "Convolutional networks on graphs for learning molecular fingerprints," in *Advances in neural information processing systems*, 2015, pp. 2224–2232.
- [11] Mathias Niepert, Mohamed Ahmed, and Konstantin Kutikov, "Learning convolutional neural networks for graphs," in *International conference on machine learning*, 2016, pp. 2014–2023.
- [12] Mikael Henaff, Joan Bruna, and Yann LeCun, "Deep convolutional networks on graph-structured data," *arXiv preprint arXiv:1506.05163*, 2015.
- [13] Thomas N Kipf and Max Welling, "Semi-supervised classification with graph convolutional networks," *arXiv preprint arXiv:1609.02907*, 2016.
- [14] Thomas G. Dietterich, "Ensemble methods in machine learning," in *Proceedings of the First International Workshop on Multiple Classifier Systems*, 2000, MCS '00.
- [15] Uday Shankar Shanthamallu, Jayaraman J Thiagarajan, Huan Song, and Andreas Spanias, "Attention models with random features for multi-layered graph embeddings," *arXiv preprint arXiv:1810.01405*, 2018.
- [16] Alexandre Abraham, Michael P Milham, Adriana Di Martino, R Cameron Craddock, Dimitris Samaras, Bertrand Thirion, and Gael Varoquaux, "Deriving reproducible biomarkers from multi-site resting-state data: An autism-based example," *NeuroImage*, vol. 147, pp. 736–745, 2017.