Exploring the efficacy of a graph classification GNN in learning non-linear graph metrics

Shivam Maurya[1], Hon Wah Yeung[2]*, and Keith M Smith[1]*

[1] Computer and Information Sciences, University of Strathclyde, Glasgow, UK
[2] Institute for Astronomy, University of Edinburgh, Edinburgh, UK
* these authors share senior authorship
keith.m.smith@strath.ac.uk

1 Introduction

Graph-structured data is common in many fields, including social networks, biological networks, and recommendation systems. The complexity of relationships in such data frequently necessitates the use of advanced modeling approaches to derive relevant insights. With the increasingly large network datasets being made available, deep learning is becoming a more relevant methodology for their exploration. Deep learning architectures which have graph inputs are called Graph Neural Networks (GNNs). One area in particular where great efforts have been made to gather population-wide data is in brain connectomics. The UK BioBank, for example has plans for up to 100,000 MRI scans which can be used for processing into brain connectomes. An important example of a graph classification GNN model for use on such data is the Brain Network Convolutional Neural Network (BrainNetCNN) model [1]. The BrainNetCNN is a CNN with special "cross-shaped" kernels for dealing with graph adjacency matrices. However, recent studies have repeatedly shown that the BrainNetCNN (among other GNNs) fails to outperform simpler, linear predictive models such as linear ridge regression in predicting population characteristics and clinical variables [2] [3] [4] [5].

This could be because most of the important characteristic/diagnostic information retrievable from brain networks is linear in nature, or there is still not enough data available to train GNNs on brain networks. But it could also be that developing more powerful models which can better identify more interesting relationships in the data with greater efficiency will significantly improve predictive power. In order to begin analysing this, here we study how well the BrainNetCNN can learn non-linear patterns and structural characteristics– clustering coefficient, routing efficiency, degree variance, diffusion efficiency, and assortativity– in three different types of synthetic graph datasets: Erdos-Renyi graphs, Barabasi-Albert graphs, and random geometric graphs. We use linear ridge regression as a baseline for comparison against linear modelling. We provide this baseline firstly to verify that BrainNetCNN can actually outperform linear models on non-linear metric learning, and secondly to enhance insights into model performance across the different graph metrics and graph datasets studied.

2 Methods

Graph datasets were generated to be similar in dimensions to state-of-the-art brain connectomic studies. Here, network size was set at 100 and 14000 graphs were sampled of each type. Dataset were split into training and testing sets (80/20 split). **Geometric Graphs:** Geometric graphs were generated with 3 Euclidean dimensions. Densities ranged from 0.1 to 0.9, 140 graphs were generated per density value. **Erdos-Renyi Graphs:** Erdos-Renyi (ER) graphs were generated with link probabilities ranging from 0.05 to 0.7. The number of graphs per link probability was set at 500. **Barabasi-Albert Graphs:** Barabasi-Albert (BA) graphs were employed to investigate learning capabilites on graphs with heterogeneous degree distributions. BA graphs were generated with parameter m spanning from 1 to 10. These m values represent the number of edges added by each new node during network growth.

We employed the Linear Ridge Regression model with a specified regularization parameter ($\alpha = 0.1$). This parameter controls the degree of regularization applied to the model, helping to prevent overfitting and ensuring robust predictions. The input to this model is the concatenated upper triangular part of the adjacency matrix. To facilitate model training, we standardized the target variables to have zero mean and unit variance, allowing the model to work with a consistent scale of data.

The BrainNetCNN model's architecture and parameters were specified using a comprehensive Layer Parameters Dictionary, allowing for flexibility in defining the model's characteristics. The dictionary encompassed essential details such as the input shape, the number of layers and filters in various components, activation functions, dropout probabilities, and regularization types. Parameter tuning enabled us to fine-tune the model's structure to best suit the task at hand. The BrainNetCNN model is constructed in a hierarchical manner, consisting of distinct layers for feature extraction and transformation. It begins with the edge-to-edge (E2E) layers, which capture edge-to-edge relationships within the graphs. The model also incorporates edge-to-node (E2N) layers to further process the extracted features, node-to-graph (N2G) layers to flatten the data, and dense layers for additional depth and abstraction. The use of residual connections, denoted by the "Residual" flag, enhances the model's capacity to capture intricate graph patterns. In this pilot study we show results using six edge-to-edge layers and six dense layers. The model's output format is adaptable, catering to both regression and classification tasks. Depending on the specific research objective, the output layer is configured with an appropriate activation function, ensuring compatibility with the chosen loss function during model compilation.

Accuracy was assessed by computing the Root Mean Squared Error (RMSE) between the actual network metric values and the corresponding values predicted by our models. RMSE is a widely accepted metric for regression tasks that measures the average squared difference between predicted and actual values. A lower RMSE score indicates better model performance in terms of accuracy.

3 Results

Pilot results of our analysis are shown in Fig 1. BrainNetCNN is able to improve on performance of linear ridge regression in almost all cases, so it is indeed able to pick up

on important non-linear information given a well-defined, noiseless scenario. However, we find that BrainNetCNN performs most similar to Linear Ridge regression in the heterogenous Barabasi-Albert models. This may provide a key step into understanding why BrainNetCNN performs similarly to linear models in brain networks, which are well known to be heterogeneous. We also note that 14000 is not a big enough dataset for the deep learning model to perfectly learn network metrics and highlights this size limitation of current neuroimaging studies [2][5].

Generally, models found assortativity to be the most difficult metric to learn. Assortativity is a measure of correlation of the degrees of connected nodes. This indicates that this kind of structural information is particularly tricky to learn even for non-linear deep learning models. If this kind of relationship is key for predicting particular population characteristics, then more powerful models which can utilise this type of information could be key to increasing predictive performance.

Of note, the only case where Linear Ridge Regression outperforms BrainNetCNN is in predicting the clustering coefficient in Erdös-Rényi random graphs. The expected value of the clustering coefficient of Erdös-Rényi random graphs is just p, so resolves to a linear function. This highlights that you are best to stick with linear models if predicting a known linear relationship.

Summary. The deep learning BrainNetCNN model is able to capture specified non-linear structural information in graphs, but its effectiveness is contingent on the inherent characteristics of the data. It struggles most in heterogeneous graphs and heterogeneity is a key characteristic of real-world networks. It also finds difficulty in learning assortativity. Beyond more data and larger models, developing GNNs which can more efficiently learn different kinds of non-linear information in graphs appears key to advancing predictive performance in graph prediction tasks.

References

1. Kawahara J., Brown C.J., Miller S.P., Booth B.G., Chau V., Grunau R.E., Zwicker J.G., Hamarneh G.: BrainNetCNN: Convolutional neural networks for brain networks; towards predicting neurodevelopment. Neuroimage. 146, 1038 – 1049 (2017)
2. He T., Kong R., Holmes A.J., Nguyen M., Sabuncu M.R., Eickhoff S.B., Bzdok D., Feng J., Yeo B.T.T.: Deep neural networks and kernel regression achieve comparable accuracies for functional connectivity prediction of behavior and demographics. Neuroimage. 206, 116276 (2020)
3. Schulz M.A., Yeo B.T.T., Vogelstein J.T., Mourao-Miranada J., Kather J.N., Kording K., Richards B., Bzdok D.: Different scaling of linear models and deep learning in UKBiobank brain images versus machine-learning datasets. Nat Commun. 11(1), 4238 (2020)
4. Yeung H.W., Luz S., Cox S.R., Buchanan C.R., Whalley H.C., Smith K.M.: Pipeline comparisons of convolutional neural networks for structural connectomes: predicting sex across 3,152 participants. Annu Int Conf IEEE Eng Med Biol Soc. 1692-1695 (2020)
5. Yeung H.W., Stolicyn A., Buchanan C.R., Tucker-Drob E.M., Bastin M.E., Luz S., McIntosh A.M., Whalley H.C., Cox S.R., Smith K.M.: Predicting sex, age, general cognition and mental health with machine learning on brain structural connectomes. Hum Brain Mapp. 44(5), 1913-1933 (2023)

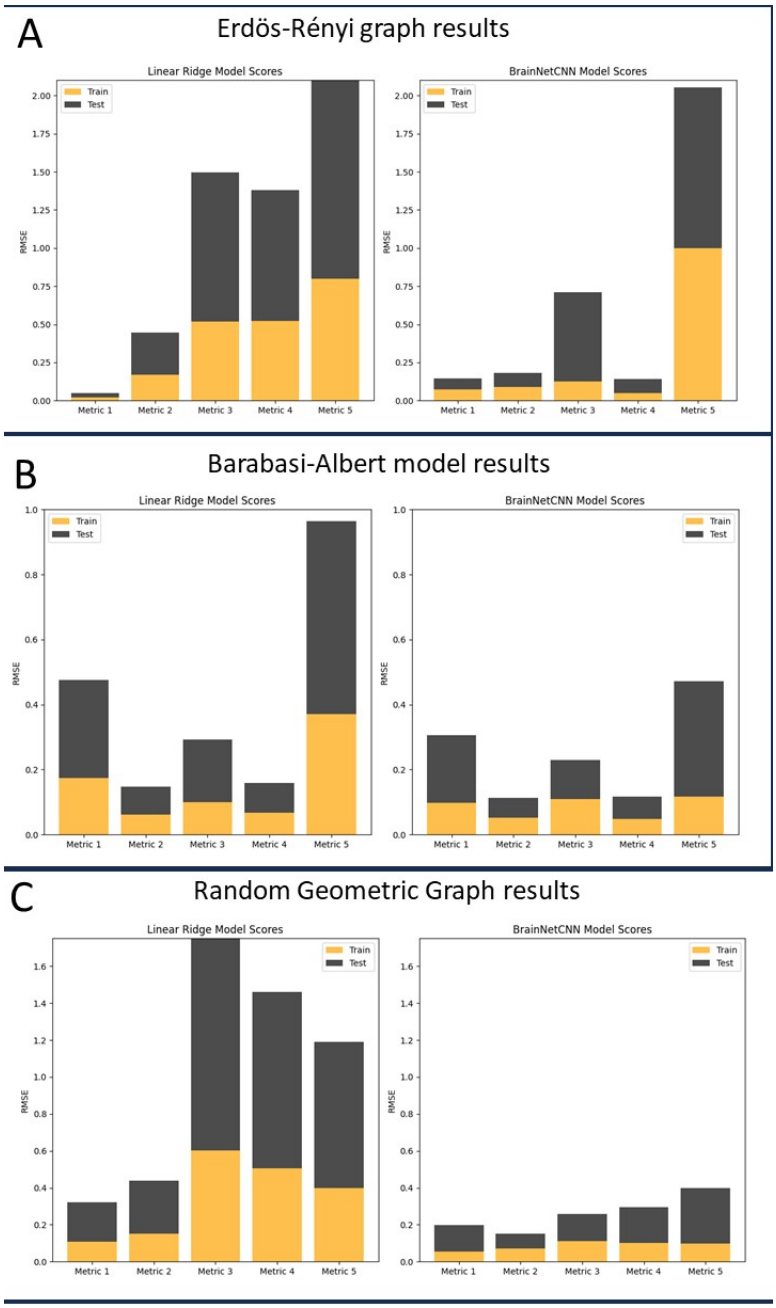

Fig. 1. Average root mean squared error (RMSE) for linear ridge regression and BrainNetCNN predicting 5 network metrics (Metric 1- clustering coefficient, Metric 2- routing efficiency, Metric 3- degree variance, Metric 4- diffusion efficiency, Metric 5- assortativity) on three graph types.