

Biologically Inspired Neural Path Finding

Hang Li^{*1,3}, Qadeer Khan^{*2}, Volker Tresp^{1,3}, and Daniel Cremers²

¹ Ludwig Maximilian University of Munich

² Technical University of Munich

³ Siemens AG

Abstract. The human brain can be considered to be a graphical structure comprising of tens of billions of biological neurons connected by synapses. It has the remarkable ability to automatically re-route information flow through alternate paths in case some neurons are damaged. Moreover, the brain is capable of retaining information and applying it to similar but completely unseen scenarios. In this paper, we take inspiration from these attributes of the brain, to develop a computational framework to find the optimal low cost path between a source node and a destination node in a generalized graph. We show that our framework is capable of handling unseen graphs at test time. Moreover, it can find alternate optimal paths, when nodes are arbitrarily added or removed during inference, while maintaining a fixed prediction time. Code is available here: <https://github.com/hangligit/pathfinding>

Keywords: Cognition · Path finding · Graphical Neural Networks

1 Introduction

We are inundated with graphical structures of various forms in this contemporary era of digitization. This includes, for e.g., social networks [17] wherein the nodes represent the individuals and the edges characterize the social connections between the individuals. Another popular form of network includes recommender systems [27] that can be represented as bipartite graphs. The users/products represent the nodes, while the edges depict the rating of likes/dislikes of a user for a certain product. Other graphs include citation networks [24], molecular structures used in drug discovery [11] etc.

Although concurrent implementation of these graphical structures are computationally powerful in number-crunching, they lack the cognitive understanding to draw meaningful conclusions that can readily be interpreted. On the other hand, one of the most sophisticated and yet least understood graphical networks is the human brain [22]. Rather than consisting of computational nodes, it is comprised of tens of billions of biological neurons both sending and receiving information to the neighbouring neurons through the connecting synapses [2,3]. One amazing attribute of the human brain is its ability to learn to automatically adapt and efficiently reroute information through alternate neural paths, in case of certain damaged neurons [28]. Another important attribute is the capability to interpret distinguishing patterns in data and retain this information to be applied in similar circumstances in the future [16]. For e.g., a child touching a hot

* These authors contributed equally.

cup of coffee once or twice would feel a sensation of pain. The child’s brain will retain this experience to avoid touching hot cups in the future even if the cups are of different size/colour/shape etc. However, unlike computers whose computation power has been exponentially rising over the past 4 decades, the capacity of the human brain to process information is limited by biological constraints. Therefore, is it possible to combine the benign attributes of the human brain with the processing power of computational resources? In this work, we explore this possibility in the context of path optimization.

The ability to navigate through a network from a source to a destination node while optimizing for the lowest cost is an important problem. It has a tremendous number of diverse applications, for e.g., the ubiquitous vehicle/robot navigation. The cost could involve minimizing either the distance travelled, time taken, or even the traffic congestion encountered. Other less frequent, but critical, use cases are search and rescue operations involving unmanned aerial vehicles. Here, minimizing the battery usage and the data transmission are important factors to be optimized for. Traditionally, these problems can either be solved heuristically using approaches such as A-star or by deploying "*shortest path*" algorithms such as Depth First Search (DFS), Breadth-First Search (BFS), Dijkstra, etc. These approaches tend to start with the source node and progressively traverse the graph through the neighbouring nodes, then neighbours of the neighbours until the destination node is found. Although accurate, their computational complexity rises with the number of hops between the source and destination nodes. On the other hand, given a visual map drawn to scale, humans are fairly good at quickly determining the approximate optimal path[5], irrespective of the number of hops between the nodes. Is it also possible to additionally emulate this one-shot prediction capability in a computational setting? In this regard, we propose using Graphical Neural Networks (GNNs) to find the path with the lowest cost. Our framework has the following attributes:

1. If a node(s) or edge(s) is arbitrarily removed from the graph structure, the optimal path is automatically rerouted through the remaining nodes/edges to find the next best solution.
2. The framework can generalize to find the optimal path even on unseen graphs.
3. The time taken to find the lowest cost path between the source and destination node remains constant irrespective of the number of hops between them.

2 Related Work

Artificial Neural Networks, ANNs:

Over the last decade, the advent of data-driven, learning-based methods for training artificial neural networks has made tremendous strides in achieving unprecedented levels of performance on various tasks such as natural language processing [23], computer vision [15], medical diagnosis [21], etc. Such networks have the capability to extract meaningful information from the training data and extrapolate this to completely unseen test data. In fact, they have achieved on par human performance [8] on tasks such as classification, disease diagnosis, etc. We also deploy ANNs to achieve generalization on unseen data (graphs).

Graph Representation:

Among the various ANN paradigms, feed-forward architectures such as Convolutional Neural Networks (CNNs), Multi-layer perceptrons etc. are ubiquitous. However, on tasks such as path optimization, there is ambiguity on how the graph structure should be represented when input through such architectures. One approach is to input the graph as an adjacency matrix. [6] use the connectome [1] as input to a CNN to predict autism in patients. The connectome is an adjacency matrix, encoding brain connectivity as graphs between certain pre-selected regions in the brain. [14] used a pre-determined number of regions in the connectome to train BrainNetCNN for predicting neurodevelopment. However, one major limitation of using adjacency matrices is that the number of nodes forming the input to the network cannot change at inference time. Otherwise, the network needs to be trained from scratch if the number of nodes is increased. Our framework does not suffer from this limitation and the number of nodes forming the input to the network can be changed without retraining the network. In fact, we show in the experiments that our method is capable of automatically rerouting to an alternate path, if some nodes/edges are removed at test time.

Reinforcement Learning (RL):

RL is an alternate strategy for determining the optimal path in a map/graph. Taking inspiration from human psychology, a reward function is defined [18,19]. The agent explores the environment in a hit-and-trial manner incurring rewards along the way [25]. If the training parameters are carefully chosen, the agent converges to an optimal policy. [20] demonstrated path planning on small unseen maps. Instead of the adjacency matrix, they directly used the map represented as a grid. While this can handle maps with arbitrary structure and nodes, it is constrained to only planar graphs. In contrast, our framework can additionally handle graphs with non-planar structure.

Shortest Path Algorithms:

Shortest path algorithms such as DFS, BFS, Dijkstra & their modifications have been used in a wide array of applications. These range from fast IP rerouting [4], to marine navigation [7], software defined networking [13], maze solving [12] to even optimal planning of sales persons [29]. The limitation of these methods is that the time to find the optimal path is not constant. Rather it depends on the number of nodes & edges in the graph. For dense graphs, the computational complexity can be of polynomial order of the number of nodes. On the other hand, the computational complexity of our method is constant. Irrespective of the number of hops between the source and destination nodes, our framework takes the same time to find the optimal path.

3 Framework

The task tackled here is to find the lowest cost path between a source and a destination node for an undirected graphical structure, G . The nodes are connected through edges having arbitrary costs. The edge $e_{i,j} \in E$ connects node i ($v_i \in V$) to node j ($v_j \in V$). Here V and E are the sets containing all nodes and edges respectively.

In our framework, we use ANNs, as they are capable of learning patterns in the data and robustly applying them to unseen data. However, in the context of optimal path finding on graphs, we would like our framework to possess two additional properties:

1. It should be invariant to the permutations/ordering of the nodes.
2. Addition/removal of nodes should not render the training of the ANN useless. Rather it should be capable of being trained & tested on any number of nodes.

Traditional feedforward networks such as MLPs do not possess these two important attributes. They are susceptible to node ordering and cannot easily handle addition of new nodes. Figure 1 depicts the implications on the adjacency matrix, when the node order is changed and when a new node is added. In the node permutation case, the adjacency matrix is completely different, despite the graphs being isomorphic. In the scenario of node addition, the adjacency matrix is extended. To accommodate this extension, the architecture of the MLP would also need to be changed; thereby requiring re-training.




Fig. 1: Shows the implications of permuting the node order and adding an additional node on the adjacency matrix of a graph.

To circumvent these issues, we also use MLPs, but in the paradigm of a Graphical Neural Network. Let the input features for a node i be represented by $z_i^{(0)}$. These input features for each node are passed through a series of graphical layers to produce latent embeddings ($z_i^{(l)}$, for node i at layer l). Figure 2 depicts a high level overview of the constituents of layer l . For simplicity, we demonstrate the information flow for the first graph shown in Figure 1. Note that the layer l takes the embeddings ($z^{(l-1)}$) for each node from the previous layer $l-1$ as input to produce corresponding embeddings $z^{(l)}$ as output. The inputs are first passed through a neural network (M) comprised of fully connected layer(s) followed by a non-linearity. An important characteristic of this neural network is that the weights are shared between all nodes. Hence, nodes can conveniently be added or removed from a graph without re-training the network. This is

because the added node can simply use the same weights as those of the other nodes. The output of the neural network is then passed through an aggregation function. For each node, the input to the aggregation function depends on which are its neighbours. For e.g., the neighbours of v_1 are v_2 and v_3 . Hence, information from $M(z_2^{(l-1)})$ and $M(z_3^{(l-1)})$ is aggregated to produce $z_1^{(l)}$. The aggregation function is chosen so that its output is invariant to the order of the input. Hence, even if the ordering of nodes in a graph is changed, the output remains the same. Some examples of order invariant aggregation functions include summation, mean, taking the maximum/minimum etc. for each scalar value of the input vectors. Note that in the first layer, each node would incorporate information from its immediate neighbours, the next layer will implicitly draw information from the neighbours of its neighbours. The deeper we go into the network, each node will retrieve information from nodes farther away from it in the graph.

The output embedding of the last graph layer is then passed through to a classifier which predicts whether or not the corresponding node falls within the optimal path. Next, we describe the mathematical details of the graph layers, the input node features, the loss function to train the weights and how the edge weights are incorporated into the graph based upon a modification of [26].

The embedding $z_i^{(l)} \in \mathbb{R}^{d_l}$ of a node i at the l -th layer is updated as:

$$z_i^{(l)} = D(\sigma(\hat{z}_i^{(l)})). \quad (1)$$

Here, σ is the non-linearity (we use LeakyReLU). D is the dropout layer. Note that \hat{z} is obtained from the input node features of the previous layer, i.e., $z_i^{(l-1)} (\in \mathbb{R}^{d_{l-1}})$ in the following manner:

$$\hat{z}_i^{(l)} = \sum_{j \in \mathcal{N}(i) \cup \{i\}} \alpha_{ij} \left(\mathbf{W}^{(l)} z_j^{(l-1)} \right) \quad (2)$$

where $W^{(l)} \in \mathbb{R}^{d_l \times d_{l-1}}$ are the trainable parameters of the neural network. Meanwhile, $\mathcal{N}(i) \in V$ is the set of neighboring nodes of i across which aggregation is done through summation. The scalar weights α_{ij} incorporate the edge cost between nodes i and j by way of this equation:

$$\alpha_{ij} = \frac{\exp(\mathbf{a}^\top \sigma([\mathbf{W}^{(l)} z_i || \mathbf{W}^{(l)} z_j || \mathbf{W}^e e_{ij}]))}{\sum_{j' \in \mathcal{N}(i) \cup \{i\}} \exp(\mathbf{a}^\top \sigma([\mathbf{W}^{(l)} z_i || \mathbf{W}^{(l)} z_{j'} || \mathbf{W}^e e_{ij'}]))}. \quad (3)$$

The edge weights e_{ij} are first mapped to a d_l dimensional hidden feature representation $h_{ij} = \mathbf{W}^e e_{ij}$ where $\mathbf{W}^e \in \mathbb{R}^{d_l}$. The \parallel symbol represents a concatenation of vectors. $\mathbf{a} \in \mathbb{R}^{3d_l}$ and \mathbf{W}^e are trainable parameters.

Note $z_i^{(0)} \in \mathbb{R}^3$ are the one-hot encoded input features representing whether a node in the graph is either the source, the destination or otherwise.

The output from the final graph layer L is the node embedding $z_i^{(L)}$. In the final layer, we also determine the edge embedding (u_{ij}) for each edge. It is the element-wise sum of embeddings of nodes that it connects to, i.e., $u_{ij} = z_i^{(L)} + z_j^{(L)}$

These edge embeddings and final layer node embeddings are passed through their respective MLP classification layers. It predicts the probabilities of them being in the

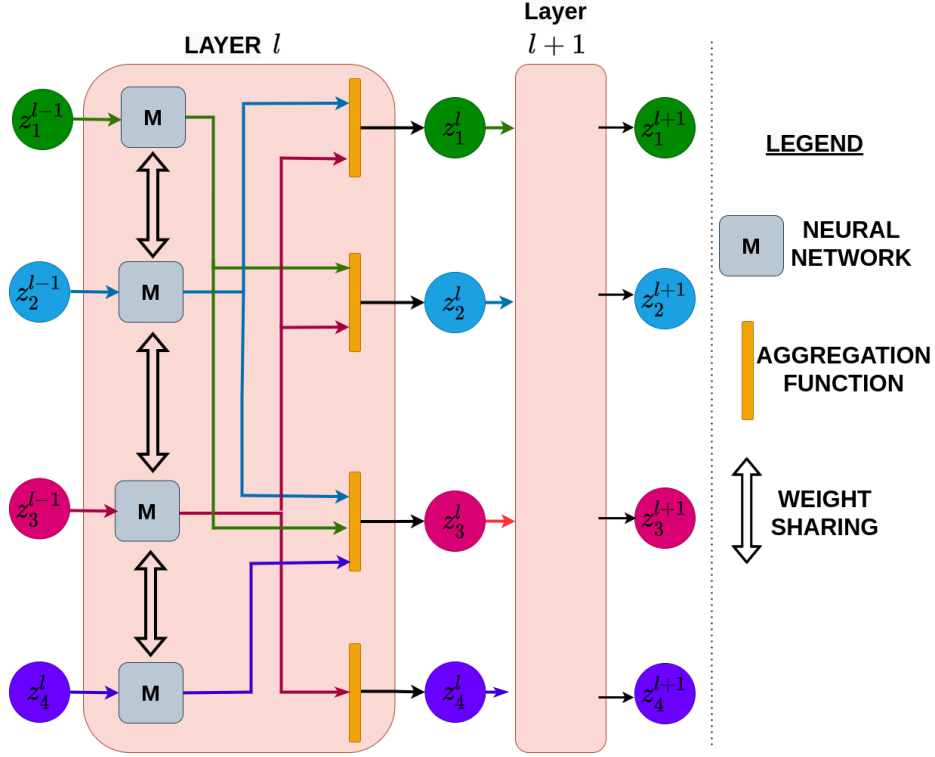


Fig. 2: Gives the high-level overview of graphical layer l , for the first graph shown in Figure 1. Node embeddings ($z^{(l-1)}$) from the previous $l-1$ layer are taken as input to produce node embeddings $z^{(l)}$ which can then be fed to the next $l+1$ layer to produce embedding $z^{(l+1)}$ and so on.

optimal path. This probability for the node i and edge ij are respectively given by the following equations:

$$\hat{p}_i = \sigma_2(\mathbf{W}_2^n(\sigma(\mathbf{W}_1^n z_i^{(L)} + \mathbf{b}_1^n) + \mathbf{b}_2^n)) \quad (4)$$

$$\hat{p}_{ij} = \sigma_2(\mathbf{W}_2^e(\sigma(\mathbf{W}_1^e u_{ij} + \mathbf{b}_1^e) + \mathbf{b}_2^e)). \quad (5)$$

$\mathbf{W}_1^n, \mathbf{W}_1^e (\in \mathbb{R}^{m \times d(L)})$, $\mathbf{W}_2^n, \mathbf{W}_2^e (\in \mathbb{R}^{1 \times m})$, $\mathbf{b}_1^e, \mathbf{b}_1^n (\in \mathbb{R}^m)$, $\mathbf{b}_2^e, \mathbf{b}_2^n (\in \mathbb{R})$ are also the trainable parameters of the model. m is a hyper-parameter. σ_2 is the sigmoid non-linearity. One may ask why we need separate weights for the node & edge classifications. Or even why is the edge classification needed at all? This is because two nodes being in the optimal path does not imply that the edge connecting them will necessarily be in the optimal path. Figure 4 describes a simple example demonstrating the importance of additionally predicting the probability of the edge being in the optimal path. Although, nodes 3 & 4 are in the optimal path, but the edge connecting them (cost=8) is not. We

also empirically found that predicting the probability for both edges & nodes is better than predicting for only the nodes or only the edges.

The loss function is the binary cross entropy between the predicted probability and the ground truth over all training samples

$$\min_{\mathbf{W}, \mathbf{b}, \mathbf{a}} \mathbb{E}_{G, y \sim \text{Data}} - [y \log \hat{p} + (1 - y) \log(1 - \hat{p})].$$

The ground truth can be obtained directly while constructing a random graph structure of larger cost around the optimal path. Or using shortest-path algorithms on known graphs.

4 Experiments

Our model is trained with a learning rate of 1e-4 with the Adam optimizer used for updating the weights. The data comprises of a total 10000 different arbitrarily created graphical structures with up to 30 nodes. The training, validation and test split is [0.7:0.15:0.15]. The test set is comprised of graphs with an arbitrary number of nodes. This serves to analyze if the model can handle different number of nodes in the inference graph. The test set also contains graph samples wherein some nodes/edges have been arbitrarily removed. This is to see if the model can determine alternate paths in case of such removal. Also, note that each of the 10000 structures is comprised of 10 different edge weight combinations for a total of 100000 samples. The ground truth labels for the loss function are obtained using [9].

Evaluation Metric:

The metric we use for quantitative evaluation is the *Path Accuracy*. It is the ratio of the number of graph samples in the unseen test set, wherein the class of every node/edge in the graph is correctly predicted. Hence, not only every node/edge in the optimal path must be classified as such but the nodes/edges not in the optimal path should also be correctly classified as not belonging to the optimal trajectory. Table 1 reports the *Path Accuracy* metric on both the training and unseen test set for our method and its variations. This is further elaborated in the next Section 5.

	Ours	Fixed Structure	Fixed Nodes	Nodes Only	Edge Only
Training					
Data	98.01	99.26	99.41	97.75	97.43
Test					
Data (Unseen)	98.02	68.41	72.18	97.62	97.41

Table 1: Reports the *Path Accuracy* metric for our method and its variations for both the training and test set. (Higher is better)

5 Discussion

We now give a more detailed analysis of the results from our experimental evaluation.

Unseen test data:

From Table 1 it can be seen that our model is capable of maintaining good *Path Accuracy* performance even on unseen test data. Note that the test data comprises of graph samples with an arbitrary number of nodes between 3 and 50. In addition, it also contains graph samples wherein the edges/nodes are arbitrarily removed. Our model is capable of robustly handling both scenarios. Figure 3 shows a plot of the accuracy as the number of nodes is changed. Note that as the number of nodes in the graph is changed, the accuracy as depicted by the green curve remains fairly consistent. This is because our model is trained to handle such instance with variable number of nodes in the graph.




Fig. 3: The left plot shows that, as the number of nodes is changed, the path accuracy metric of our model (in green) remains constant. It is interesting to note that even though the model was trained with up to a maximum of 30 nodes, it still maintains good performance beyond this number. The red curve on the plot shows the performance for a model trained with a fixed number of nodes. Its performance deteriorates when evaluated on a smaller number of nodes.

Meanwhile, Figure 4 demonstrates a simple example of finding an optimal alternate path in the case of a removed edge. Note that initially in the original graph, the optimal path between nodes 2 and 4 is the direct edge connecting the two nodes having a cost of 1. However, when this removed is removed the alternate lowest cost path between the nodes is 2-0-1-3-5-4. This alternate has a cost of 7 which is the lowest in the graph after the edge removal.

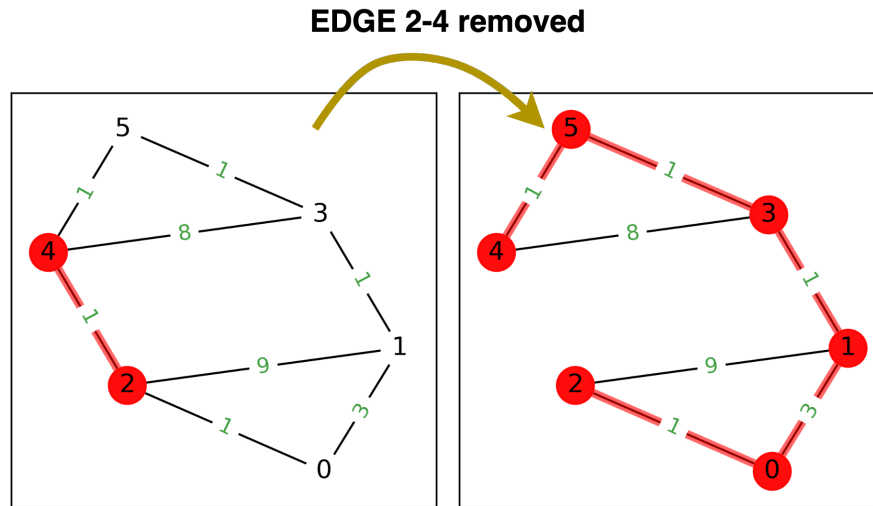


Fig. 4: The Figure shows that if the shortest path between nodes 2-4 is removed, the network is capable of automatically finding the alternate shortest path through nodes 2-0-1-3-5-4. The cost via this alternate path is the lowest in the modified graph resulting from the edge removal

Fixed structure, fixed number of nodes:

Note that in our framework not only the number of nodes can change, but also the structure of the graph constituted by the same nodes can vary. To demonstrate this, we train another model wherein both the structure and the number of nodes (30) are kept constant. Only the edge weights are changed. Consider the 2nd column of Table 1 for the results. While the model performs well on the training samples containing graphs of the same structure & nodes, its performance on the unseen test data drops dramatically. This is because when training the model, it is not accustomed to handling variable nodes & structures of the test data.

Variable structure, fixed number of nodes:

To compensate for this, we train another model, with the same number of nodes, but for which the structure of the graph can change. Note that the performance of this model is good even when the structure is changed at test time. This is because the training incorporated changing structure into the model. However, because the number of nodes is fixed, its performance drastically starts to drop when additional nodes are added/removed to the graph as can be seen in the red curve in Figure 3. In contrast our approach can not only accommodate addition of new nodes but also changing structure of the graph while maintaining a stable performance.

Comparison with BrainNetCNN approach:

Note that the BrainNetCNN approach uses adjacency matrices, which requires filter sizes in the CNN to be fixed. Hence, graphs with an arbitrary number of nodes cannot

readily be handled. The original BrainNetCNN was meant for classifying neurodevelopment in patients. Therefore, we slightly modify the deeper layers to classify nodes/edges in the optimal path instead. To further elaborate this, Figure 5 shows a 3x3 grid structure depicting the accuracy distribution for 3 different models evaluated on 3 different datasets. Each column in the grid represents a different dataset and is described as follows:

1. *Variable Weights*: All graphs in this dataset have the same structure but have different edge weights
2. *Node Permutations*: All graphs in this dataset have the same structure but have the nodes permuted and also different edge weights.
3. *Variable Structures*: The graphs in this dataset can have not have different structure, different edge weights. However, the number of nodes still remain the same.

Each row represent a different model, training of which is described as follows:

1. *Variable Weights*: This model was trained with graphs with on graphs having same structure but different edge weights. This model has good performance on the Variable Weights dataset but performs poorly when evaluated on the datasets wherein the nodes are permuted or where the structure is different. This is understandable because this model is not accustomed handling such graphs during training.
2. *Node Permutations*: This model was trained on graphs having the same structure but with the nodes permuted and also different edge weights. Here this model performs well on both the Variable Weights and Node Permutations dataset. However, the performance of this model drops when evaluated on graphs with variable structures.
3. *Variable Structures*: This model was trained on graphs having different structures and has a good, stable performance across all the 3 datasets.

However, note that all these models could only be trained and evaluated on graphs with the same number of nodes. Our framework on the other hand does not suffer from this problem and can additionally be trained and evaluated on graphs with a variable number of nodes as already depicted in Figure 4.

Relative Prediction time:

Figure 6 reports the relative time to find the optimal path as the number of hops between the source and destination nodes is increased. It is normalized by the time taken to find the optimal path between nodes one hop away. As can be seen, this number remains stable for our approach. Hence, irrespective of the number of hops we have in the graph, the prediction time remains the same. This is because a forward pass of our model always involves the same number of graph conventional layers. Compare this with Dijkstra’s algorithm wherein the relative time increases with the number of hops.

Loss functions:

In addition to classifying the nodes in the optimal path, our loss function also incorporates classifying the edges connecting the nodes in the path. The last 2 columns in Table 1 shows the implications of training with the binary cross entropy function only for the nodes and only for the edges. As can be seen, the performance of our model, which combines both loss functions, is superior to the models trained with only the individual loss functions.

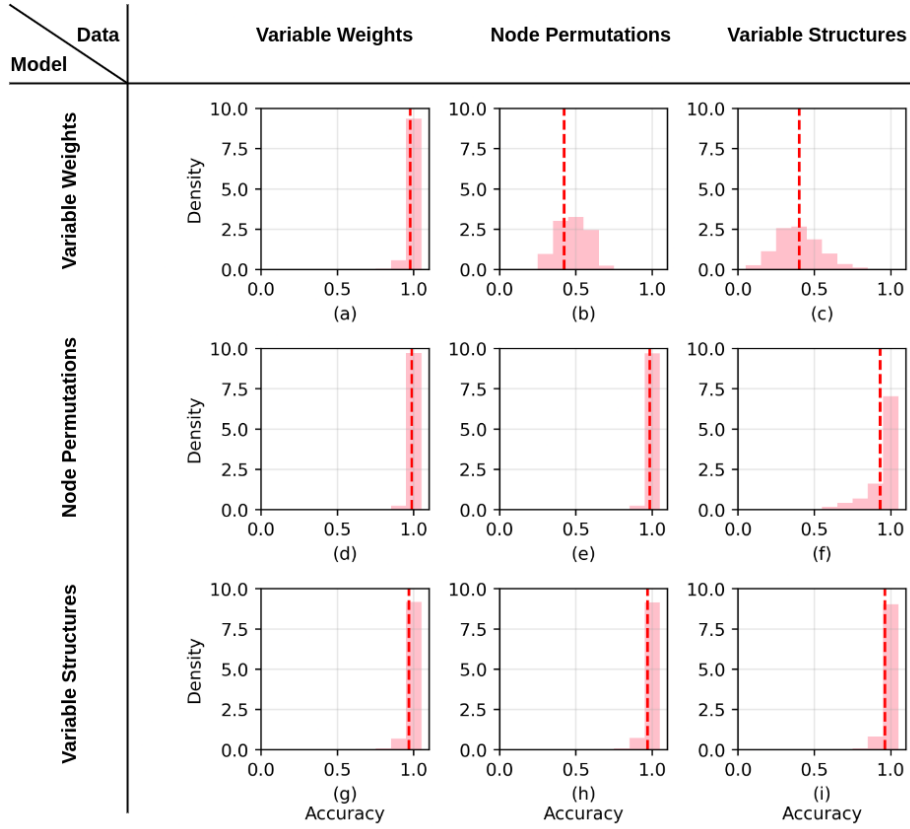


Fig. 5: The first row shows the test performance of a BrainNetCNN trained with fixed structure and fixed number of nodes. Each column from left to right denotes graphs with fixed structure, node permutations, and variable structures. The second row shows model trained with node permutations. The model trained with different structures (the third row) generalizes well across all the 3 datasets. It is important to mention that all datasets contained a fixed number of nodes and all models were also trained with a fixed number of nodes.

Evaluation on a real world dataset:

We also evaluated our approach of optimal path finding on maps of the real-world KITTI [10] dataset and found that it achieved a perfect score on the *Path Accuracy* metric. One plausible explanation for this is that the maps of the road structure in KITTI are on a plane. Our model, in contrast, was trained with non-planar graphs which tend to be more challenging to handle and hence do not achieve a perfect score when evaluated on a test set comprising non-planar graphs.




Fig. 6: The plot shows the relative prediction time as a function of the number of hops between the source and destination nodes.

6 Conclusion

In this paper we demonstrated how our biologically inspired computational framework is capable of optimal path finding. It mimics the behaviour of the brain to find alternate shortest paths on unseen data even when nodes/edges are removed. This is unlike adjacency-matrix based conventional feedforward approaches which cannot be trained with varying numbers of nodes. As we developed our framework for generalized graph structures, it can be extended to various applications.

7 Acknowledgement

This work was supported by the Munich Center for Machine Learning and the BMBF-project MLWin.

This preprint has not undergone peer review (when applicable) or any post-submission improvements or corrections. The Version of Record of this contribution is published in *Brain Informatics*, and is available online at https://doi.org/10.1007/978-3-031-15037-1_27.

References

1. Chapter 3 - connectivity matrices and brain graphs. In: *Fundamentals of Brain Network Analysis*, pp. 89–113. Academic Press, San Diego (2016)
2. Azevedo, F.A., et al.: Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain. *Journal of Comparative Neurology* **513**(5), 532–541 (2009)
3. Bastos, A.M., et al.: Canonical microcircuits for predictive coding. *Neuron* **76**(4), 695–711 (2012)
4. Bhor, M., et al: Network recovery using ip fast rerouting for multi link failures. In: *2017 International Conference on Intelligent Computing and Control (I2C2)*. pp. 1–5 (2017)

5. Bongiorno, C., et al.: Vector-based pedestrian navigation in cities. *Nature Computational Science* **1**(10), 678–685 (2021)
6. Brown, C.J., et al.: Connectome priors in deep neural networks to predict autism. In: 2018 IEEE 15th International Symposium on Biomedical Imaging. pp. 110–113 (2018)
7. Cheng, Z., et al.: The method based on dijkstra of multi-directional ship's path planning. In: 2020 Chinese Control And Decision Conference (CCDC). pp. 5142–5146 (2020)
8. Cireşan, D., et al.: Multi-column deep neural network for traffic sign classification. *Neural Networks* **32**, 333–338 (2012), selected Papers from IJCNN 2011
9. Dijkstra, E.W.: A note on two problems in connexion with graphs. *Numerische mathematik* **1**(1), 269–271 (1959)
10. Geiger, A., et al.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2012)
11. Gilmer, J., et al.: Neural message passing for quantum chemistry. In: International conference on machine learning. pp. 1263–1272. PMLR (2017)
12. Hidayatullah, A.S., et al: Realization of depth first search algorithm on line maze solver robot. In: 2017 International Conference on Control, Electronics, Renewable Energy and Communications (ICCREC). pp. 247–251 (2017)
13. Jiang, J.R., et al: Extending dijkstra's shortest path algorithm for software defined networking. In: The 16th Asia-Pacific Network Operations & Management Symposium (2014)
14. Kawahara, J., et al: Brainnetcnn: Convolutional neural networks for brain networks; towards predicting neurodevelopment. *NeuroImage* **146**, 1038–1049 (2017)
15. Krizhevsky, et al: Imagenet classification with deep convolutional neural networks. In: Advances in Neural Information Processing Systems. vol. 25. Curran Associates, Inc. (2012)
16. Lake, B.M., Ullman, T.D., Tenenbaum, J.B., Gershman, S.J.: Building machines that learn and think like people. *Behavioral and brain sciences* **40** (2017)
17. Leskovec, J., Rajaraman, A., Ullman, J.D.: Mining Social-Network Graphs, p. 325–383. Cambridge University Press, 2 edn. (2014)
18. Niv, Y.: Reinforcement learning in the brain. *Journal of Mathematical Psychology* **53**(3), 139–154 (2009), special Issue: Dynamic Decision Making
19. Palminteri, S., et al.: Chapter 5 - reinforcement learning and tourette syndrome. In: Advances in the Neurochemistry and Neuropharmacology of Tourette Syndrome, International Review of Neurobiology, vol. 112, pp. 131–153. Academic Press (2013)
20. Panov, A.I., et al: Grid path planning with deep reinforcement learning: Preliminary results. *Procedia Computer Science* **123**, 347–353 (2018)
21. Pasa, F., et al.: Efficient deep network architectures for fast chest x-ray tuberculosis screening and visualization. *Sci Rep* **9** (2019)
22. Peer, M., Brunec, I.K., Newcombe, N.S., Epstein, R.A.: Structuring knowledge with cognitive maps and cognitive graphs. *Trends in cognitive sciences* **25**(1), 37–54 (2021)
23. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. *OpenAI blog* **1**(8), 9 (2019)
24. Sen, P., Namata, G., Bilgic, M., Getoor, L., Galligher, B., Eliassi-Rad, T.: Collective classification in network data. *AI magazine* **29**(3), 93–93 (2008)
25. Sutton, R.S., Barto, A.G.: Reinforcement learning: An introduction. MIT press (2018)
26. Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y.: Graph attention networks. *stat* **1050**, 20 (2017)
27. Ying, R., et al.: Graph convolutional neural networks for web-scale recommender systems. In: Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining. pp. 974–983 (2018)
28. Zelikowsky, M., et al.: Prefrontal microcircuit underlies contextual learning after hippocampal loss. *Proceedings of the National Academy of Sciences* **110**(24), 9938–9943 (2013)

29. Žunić, E., et al: Software solution for optimal planning of sales persons work based on depth-first search and breadth-first search algorithms. In: 39th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO) (2016)