Connectome-constrained spatially embedded recurrent neural networks

Maroš Rovný (Maros.Rovný@mrc-cbu.cam.ac.uk)
MRC Cognition and Brain Sciences Unit, University of Cambridge
Cambridge, UK

Danyal Akarca (Danyal.Akarca@mrc-cbu.cam.ac.uk)
MRC Cognition and Brain Sciences Unit, University of Cambridge
Cambridge, UK

Jascha Achterberg (Jascha.Achterberg@mrc-cbu.cam.ac.uk)
MRC Cognition and Brain Sciences Unit, University of Cambridge
Cambridge, UK

Duncan Astle (Duncan.Astle@mrc-cbu.cam.ac.uk)
MRC Cognition and Brain Sciences Unit, University of Cambridge
Cambridge, UK

Abstract

Deep learning models are one way to address the reliance of cognitive neuroscience on association studies, by providing a sandbox for systematically testing causality in complex systems. However, only a fraction of these allow for the modelling of the entire human brain. Expanding upon a recent innovation of spatially embedded recurrent neural networks, our approach introduces a novel regularisation term based on the optimal transport problem. This enriches the training process with information about the distance between the distributions of network measures describing the artificial and empirical topology. Exemplified through communicability, our approach unlocks future avenues for exploring the impact of different topological properties of the human brain on its performance across varied contexts by allowing these to be included in training and tested on different tasks.

Keywords: spatial embedding; topological embedding; regularization; recurrent neural networks

Introduction

Most computational modelling within neuroscience happens at the scale of individual neurons or small neural assemblies. This restricts the overlap between phenomena of interest across cognitive and computational neuroscience. Although more common in recent years, approaches attempting to capture the dynamics at regional or whole-brain scales are still relatively rare (Pathak et al., 2022).

A recent framework by Achterberg & Akarca et al. (2023) provides an approach that can be used to consider the spatial organisation of the brain within the context of recurrent neural networks. The approach of spatially embedded recurrent neural networks uses a custom regularisation function to embed the recurrent layer in Euclidean space by scaling the weight matrix by a distance matrix. Further addition of the communicability matrix, which denotes local random diffusion over the network, allows the network to consider not only space but also topology. Thus, using distance and communicability matrices to regularise network weights provides a vehicle for embedding the spatial and communicative properties of idealised biological systems within artificial networks.

Here, we develop the spatially embedded recurrent framework by increasing the level of abstraction under consideration. The original instantiation does not capture the different ways that topology might be expressed under given spatial constraints. The approach introduced in our work drives the network to mirror the distribution properties of communicability, without restricting them to only one endpoint, thus potentially providing insight into the higher-order organisational principles giving rise to the structure underlying network-wide computation.

Methods

Task As in the original work (Achterberg & Akarca et al., 2023), we trained the artificial networks to perform a simple one-step inference task with both memory and decision components, shown in Figure 1.

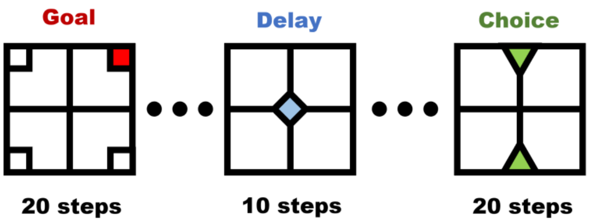

Figure 1: Using a simple 2-by-2 grid, the networks were first shown the target location in one of the corners of the grid for 20 steps, the starting location for 10 steps, and two possible choice locations for 20 steps, after which the network had to choose the location closer to the target.

Empirical networks For the spatial component of our embedding, we used the Brainnetome parcellation, consisting of 246 nodes covering both cortical and subcortical areas. This defines the space in which the neural network is embedded. For the topological embedding, we used a weighted diffusion-tensor imaging-based connectome of a randomly chosen subject from the Cambridge Attention, Learning, and Memory (CALM) cohort (Holmes et al., 2019). This defines how weights are approximated via the learned objective of the neural network.

Artificial networks We chose recurrent neural networks for their interconnectedness and recurrence without additional top-down components. The architecture of the network is shown in Figure 2.

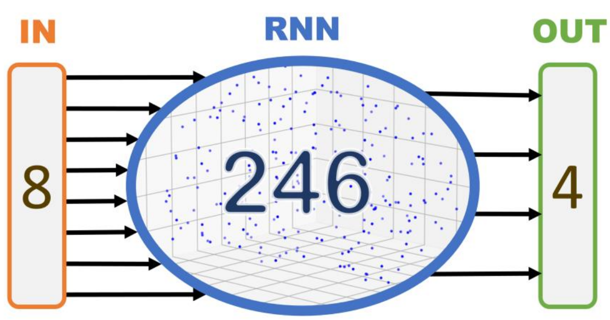

Figure 2: The networks were composed of three layers: 1) input layer with 8 neurons, fully connected to 2) recurrent layer with 246 neurons, one per each node of the Brainnotome parcellation, fully connecting to 3) output layer with 4 neurons.

Spatial embedding To provide the network with spatial information, we used the method proposed by Achterberg & Akarca et al. (2023):

$$loss_{spatial} = ||W \odot D|| \quad (1)$$

Where W is the weight matrix of the recurrent layer and D is the distance matrix containing the Euclidean distances between Brainnetome nodes.

Topological embedding To provide the network with topological information, we used a novel approach based on the optimal transport problem:

$$loss_{topological} = EMD\ (C_{empirical}\ , C_{artificial}\ ) \quad (2)$$

Where EMD is the Earth Mover's Distance (Wasserstein implementation[1]) between the flattened empirical and artificial communicability[2] matrices C.

Results

The results presented below show networks that have been trained to perform the task with 100% accuracy.

Communicability distributions As shown in Figure 3, the communicability distribution of the artificial networks approximates the empirical network over training.

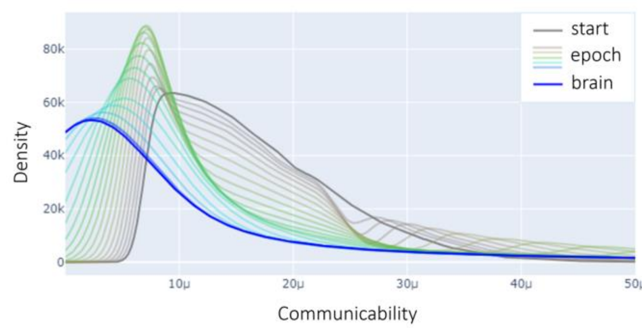

Figure 3: Density plot of the communicability distribution over epochs. The grey line is the network at initialization, the blue line is the target empirical network, and the gradient of lines between these shows the network across epochs.

Weight matrices As seen in Figure 4, the weight matrix of the artificial networks shows organisational similarities to the empirical one, attained not by the direct introduction of an empirical weight matrix but by shared organisational properties arising from topological distribution constrained by Euclidean space.

[1] $EMD = \int_0^1 |cdf_u^{-1}(q) - cdf_v^{-1}(q)|^p\, dq$, where u and v denote the two distributions and $cd\,f^{-1}(q)$ are inverse cumulative distribution functions (Peyré & Cuturi, 2018)

[2] Although the standard method for calculating a communicability matrix is simply taking the exponent of the matrix, we made use of a method better suited for weighted matrices such as connectomes adapted from Crofts & Higham (2009)

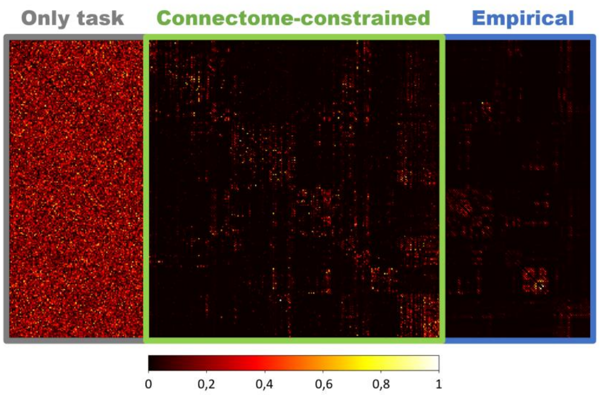

Figure 4: Weight matrices of a) an artificial neural network trained to perform the task without any spatial or topological information, b) an artificial neural network training with both spatial and topological constraints, and c) an empirical neural network within the Brainnetome parcellation.

Discussion

We introduce a novel approach to creating artificial neural networks, inspired by the topology of the human brain. By using Earth mover's distance between the network's own and brain-based communicability distributions, we were able to create networks that could solve a memory and decision-making task while at the same time shaping their connections to an end-state analogous to that of a connectome.

Our ongoing efforts are already addressing two main limitations: First, the examples provided here were based only on communicability, in the absence of other graph theoretical measures. As this method can use any graph theoretical measure that can be represented as a distribution, from simple weight distribution to different measures of centrality, many other topological characteristics could be considered. The second shortcoming is the task, where once again the lack isn't in quality but rather in quantity. In the future, other tasks with different principles should be used, to verify the robustness of our approach in varied contexts.

The promise of our approach is exciting: it could be used to understand the importance of different topological properties of the human brain on its overall efficiency in learning and performing a task. Moreover, it could be used to compare the brains of individuals or groups with varied neural phenotypes, which could lead to insights in clinical research.

Acknowledgments

This work was funded by United Kingdom's Medical Research Council.

References

Achterberg, J., Akarca, D., Strouse, D. J., Duncan, J., & Astle, D. E. (2023). Spatially embedded recurrent neural networks reveal widespread links between structural and functional neuroscience findings. *Nature Machine Intelligence 2023*, 1–13. https://doi.org/10.1038/s42256-023-00748-9

Crofts, J. J., & Higham, D. J. (2009). A weighted communicability measure applied to complex brain networks. *Journal of The Royal Society Interface*, *6*(33), 411–414. https://doi.org/10.1098/RSIF.2008.0484

Holmes, J., Bryant, A., & Gathercole, S. E. (2019). Protocol for a transdiagnostic study of children with problems of attention, learning and memory (CALM) 17 Psychology and Cognitive Sciences 1701 Psychology 17 Psychology and Cognitive Sciences 1702 Cognitive Sciences 11 Medical and Health Sciences 1117 Public Health and Health Services. *BMC Pediatrics*, *19*(1), 1–11. https://doi.org/10.1186/S12887-018-1385-3/FIGURES/1

Pathak, A., Roy, D., & Banerjee, A. (2022). Whole-Brain Network Models: From Physics to Bedside. *Frontiers in Computational Neuroscience*, *16*, 866517. https://doi.org/10.3389/FNCOM.2022.866517/BIBTEX

Peyré, G., & Cuturi, M. (2018). Computational Optimal Transport. *Found. Trends Mach. Learn.*, *11*(5–6), 1–257. https://doi.org/10.1561/2200000073