

PHYS 760 Project Report: Investigating the Topological Dimension of a Novel Manufacturing Method for Neuromorphic Hardware

Spencer Powers

August 9, 2020



Abstract

The network topology of the brain is a longstanding topic of interest for the neuroscience and neuromorphic computing communities alike. Recent work by Pantone et al. from Rain Neuromorphics proved that their novel manufacturing method for neuromorphic hardware naturally produces small-world networks, a hallmark characteristic of real connectomes. However, the criticality hypothesis implies that small-worldness is not a complete answer to the topology question, and introduces hierarchical modular networks (HMNs) as a viable alternative. This work reproduces the results of Pantone et al. from scratch, then extends the analysis to investigate the feasibility of producing HMNs with the same nanowire deposition process. While computational resources limited the scale of the simulated network, initial analysis shows promise in producing HMN-like connectivity by distributing electrodes in nested clusters.

Contents

1	Introduction	3
2	Reproducing Results of Pantone et al. (2018)	5
2.1	Electrode Model	5
2.2	Straight Wire Model	5
2.3	Pink Noise Wire Model	6
2.4	Analysis	7
3	Topological Dimension Analysis	10
3.1	Approach	10
3.2	Results	12
4	Discussion	13
5	References	13

1 Introduction

While the network topology of the human connectome is an open question in neuroscience, one prevalent idea is the requirement of small-world connectivity [1]. Small-world networks are characterized by strong local connectivity between neighboring nodes and sparse long-range connections. As a result, the total number of connections is low, but the maximum distance between any two nodes is still small. This is illustrated in Figure 1:




Figure 1: Classical depiction of a small-world network [1]. Note the sparse long-range connections allow access to any node in a small number of edges.

However, another common idea in neuroscience tends to conflict with this model: the criticality hypothesis. This hypothesis states that the brain exists in a "critical" state, just at the border between sustained neural activity and inactivity [2]. In support of this, empirical studies have shown that the human brain structure can only support the dynamics observed in fMRI recordings in a critical state [3]. However, to exist in a state of criticality, the brain must be tuned to reside at the exact boundary between the two states. The biological plausibility of this precise tuning is questionable, leading researchers to support the presence of Griffiths phases in the network [1]. Griffiths phases are extended critical-like regions that allow for a broader range of parameters while still exhibiting critical behavior [3].

Previous studies have demonstrated that dynamical models of activity propagation exhibit Griffiths phases when the networks have a finite topological dimension (d). The topological dimension of a network (d) measures how the number of neighbors (N) of any given node grows when moving r steps away from it [3]. Formally, d is described by a power law fit: $N(r) \sim r^d$ for large r . By definition, small-world networks have a theoretically infinite topological dimension when the number of nodes approaches infinity. This is due to the fact that all nodes are accessible within a relatively small number of edges, so any attempt to fit a power law to the curve will result in a progressively steeper slope. However, large-world networks have a finite topological dimension. To reconcile the presence of small-world connectivity and Griffiths phases, the Hierarchical Modular Network (HMN) was introduced (Figure 2):




Figure 2: A sample HMN with 2^{11} nodes and 10 hierarchical levels [3].

HMNs consist of small, densely-connected networks that are nested in a clustered hierarchy. As a result, the networks exhibit small-world behavior locally, but the global structure is a large-world network and thus has a finite topological dimension [1].

The problem of the brain network topology is directly applicable to the field of neuromorphic computing, which strives to create low-power, massively-parallel chips inspired by the brain. Pantone et al., associated with a startup called Rain Neuromorphics, released a paper detailing how their novel chip manufacturing process naturally creates small-world networks [4]. Specifically, the authors create the networks by depositing a layer of electrospun memristive nanowires on a grid of electrodes. As a result, the electrodes are connected in a stochastic manner. This random behavior can be seen in a sample image of electrospun nanowires provided by Pantone et al. in Figure 3:




Figure 3: Titanium dioxide electrospun nanowires exhibiting highly stochastic behavior (image source: [4]).

While this is promising in the sense that it satisfies one of the primary ideas for brain network topology, whether their method can produce networks supporting Griffiths phases and thus flexible regions of criticality is not yet determined.

This work aims to first reproduce the models and analysis presented by Pantone et al., then extend the analysis to include approximations of the topological dimension of the resulting networks.

2 Reproducing Results of Pantone et al. (2018)

2.1 Electrode Model

Across all models explored by Pantone et al., the geometric model of the electrode array is constant. Electrodes are arranged in a grid and are characterized by their radius (r_e) and the spacing between electrode centers (α). In the original work, r_e was 0.4 and α was 1. A sample configuration with 9 electrodes is shown in Figure 4:




Figure 4: Modeled grid of 9 electrodes with $r_e = 0.4$ and $\alpha = 1$.

The geometric representation of the electrodes as circles is used in all variations of the wire models to check for wire-electrode contact.

2.2 Straight Wire Model

To ensure that the geometric model construction and analysis tools were set up correctly, the simplest wire model configuration presented in [4] was first developed. This network representation models the randomly-placed wires as straight lines. Per [4], wires extend from a random location along a side to another random location on a different side. The number of wires is determined by a density parameter λ , which is multiplied by the square root of the number of electrodes to yield the number of generated wires. For all models, $\lambda = 30$ was used, in accordance with [4].

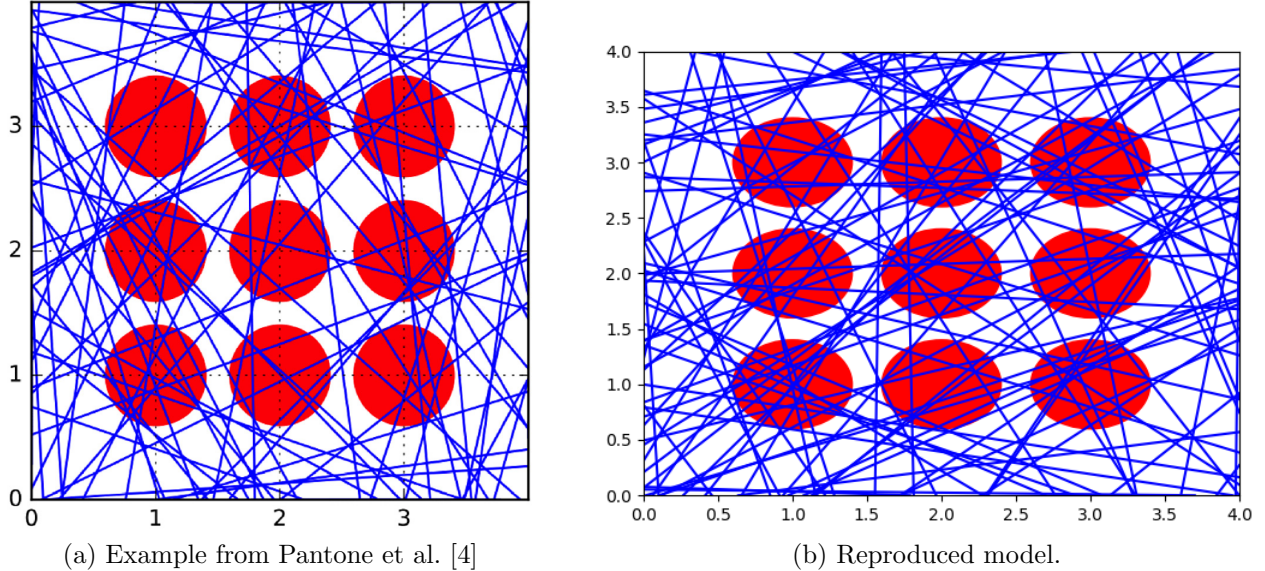


Figure 5: Original and reproduced straight wire models with 9 electrodes with $r_e = 0.4$, $\alpha = 1$, and $\lambda = 30$. Note that the stochastic nature of the geometric model prevents them from being identical.

To determine whether a given wire contacts an electrode, the orthogonal distance between the line and the electrode center is given by:

$$D_s = \frac{|x_e(y_2 - y_1) - y_e(x_2 - x_1) + x_2y_1 - y_2x_1|}{\sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}} \quad (1)$$

By computing the orthogonal distance D_s via Eq. 1, this value can be directly compared to the electrode radius r_e to determine whether the wire contacts the electrode.

2.3 Pink Noise Wire Model

To build the foundation for later analysis, the most realistic full network representation presented in [4] was then developed. This network model generates a number of pink noise points, performs a number of transformations and scaling operations, and fits a quintic polynomial to the points to model wires. This curve is then rotated about an arbitrary point by a random angle to produce visually realistic wires [4].

A number of small modifications were made in this implementation to make it more user-friendly. Instead of manually approximating pink noise via the method outlined in [4], a Python pink noise generator package (<https://github.com/felixpatzelt/colorednoise>) was used.

In addition, a SciPy function minimization package was used to determine electrode connectivity instead of the polynomial approach presented in [4]. Specifically, the function to be minimized took in a value of unrotated horizontal position (x), computed the corresponding vertical position (y) on the wire via the original quintic polynomial fit, rotated the points by the same angle as the original wire, and computed the square of the distance between

the rotated points and the electrode center. The function minimization returns a number of candidate unrotated x values, which are then used to generate unrotated y values, which are then rotated by the same angle. Finally, the distance between these rotated candidate points and the electrode center is computed and compared to the radius of the electrode (r_e) to determine connectivity.

It was noted that the wire models were somewhat clustered before the rotation transformation, which resulted in pronounced gaps in grid coverage. To resolve this, a small random offset was applied to spread the generated wires in a more uniform distribution before rotation. This resulted in wires that looked visually similar to those presented in [4].

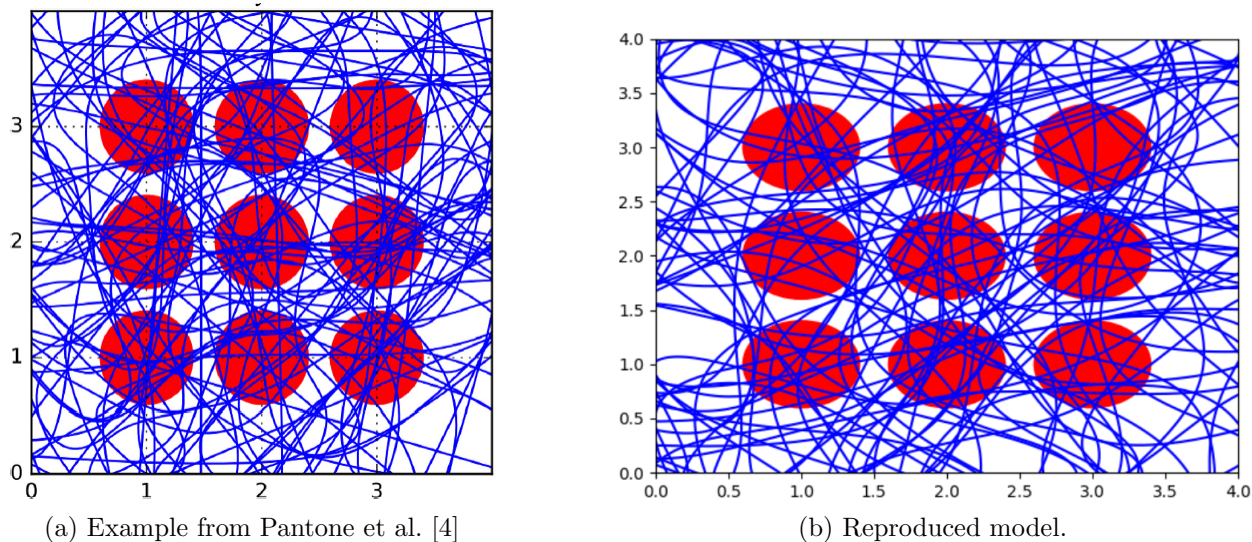


Figure 6: Original and reproduced pink noise wire models with 9 electrodes with $r_e = 0.4$, $\alpha = 1$, $\lambda = 30$, and $n = 201$. Note that the stochastic nature of the geometric model prevents them from being identical.

2.4 Analysis

To determine if the resulting networks are small-world networks, the geometric model was first converted into a bipartite graph [4]. The first set of nodes represents electrodes, the second set represents wires, and the edges connect wire nodes to electrode nodes that they contact:

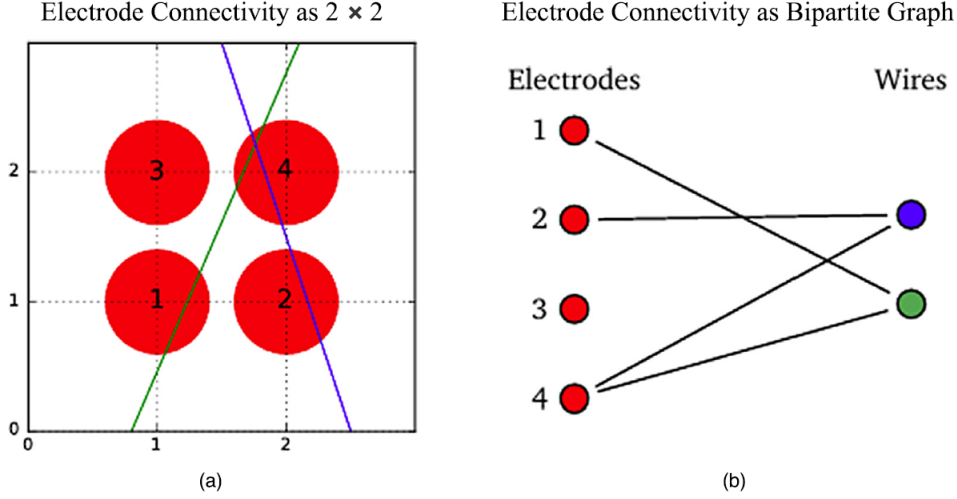


Figure 7: Representation of a geometric model as a bipartite graph [4].


This representation allowed Pantone et al. to analyze the network topology with extensively-implemented graph algorithms. To compute the small-worldness coefficient of the equivalent bipartite graph, a random bipartite graph with the same number of nodes and edges must be constructed. Then, since the conventional clustering coefficient does not apply to bipartite graphs, the square clustering coefficient is computed for both graphs [4]:

$$C_4(v) = \frac{\sum_{u=1}^{k_v} \sum_{w=u+1}^{k_v} q_v(u, w)}{\sum_{u=1}^{k_v} \sum_{w=u+1}^{k_v} [a_v(u, w) + q_v(u, w)]} \quad (2)$$


Finally, the average shortest path between electrodes is computed for both graphs. This can be done using existing network analysis packages such as NetworkX (<https://github.com/networkx>). With these four terms, the small-worldness coefficient is computed as follows:

$$\sigma = \frac{\frac{C}{C_r}}{\frac{L}{L_r}} \quad (3)$$

The small-worldness coefficient (σ) was calculated over a sweep of electrode count values for the straight and pink noise wire models per the procedure in [4]. First, the results of the straight wire model sweep were compared to those presented in [4]:



(a) Small-world coefficient sweep for the straight wire model from Pantone et al. [4]




(b) Reproduced analysis for the straight wire model.


Figure 8: Original and reproduced small-worldness coefficient sweeps over a range of electrode counts using the straight wire model. For all networks generated in the sweep, $r_e = 0.4$, $\alpha = 1$, and $\lambda = 30$. Note that the stochastic nature of both the geometric model and the random comparative graph used in the σ computation prevents them from being identical.

While the figures cannot exactly match due to the stochastic processes embedded in the graph generation and analysis processes, the behavior between intervals on the x axis is quite similar. Namely, $\sigma(20) \approx 2.25$, $\sigma(40) \approx 2.75$, $\sigma(60) \approx 3.25$, $\sigma(80) \approx 3.5$, $\sigma(100) \approx 3.8$.

Furthermore, the behavior of the pink noise wire model sweep compared to those presented in [4] shows strong similarity, as shown in the following figure. However, the number of electrodes was limited to 40 due to the significant computational complexity incurred by checking for electrode connectivity with this wire model.



(a) Small-world coefficient sweep for the pink noise wire model from Pantone et al. [4]



(b) Reproduced analysis for the pink noise wire model. Note the limited range of the x-axis due to computational constraints.

Figure 9: Original and reproduced small-worldness coefficient sweeps over a range of electrode counts using the pink noise wire model. For all networks generated in the sweep, $r_e = 0.4$, $\alpha = 1$, $\lambda = 30$, and $n = 201$. Note that the stochastic nature of both the geometric model and the random comparative graph used in the σ computation prevents them from being identical.

Comparing Figures 9a and 9b on the limited $\sqrt{n_e}$ range from 0 to 40 shows strong similarity in behavior. Notably, at low $\sqrt{n_e}$, $\sigma < 1$. Other points of comparison are $\sigma(20) \approx 1.6$ and $\sigma(40) \approx 2$.

3 Topological Dimension Analysis

3.1 Approach

The topological dimension of a network (d) is a measure of how the number of neighbors (N) of any given node grows when moving r steps away from it [3]. The difference between the topological dimensions of small-world networks and Hierarchical Modular Networks (HMNs) is illustrated in the following figure from [3]:




Figure 10: $N(r)$ vs. r plot for HMNs with various hierarchical level-dependent wiring probabilities (p), with $p = 4$ corresponding to a sample small-world network. Note the asymptotic behavior of the slopes of the curves as p is varied.

Given the power law definition of d , it can be seen from Figure 10 that as r is increased for the small-world network ($p = 0.4$), d must continue to increase to match the slope. Formally, with an infinitely large network, increasing r will require d to approach infinity. However, in the $p = 1/4$ case, it can be seen that the slope reaches a steady-state value, yielding a finite d for an arbitrarily large network and radius.

To perform this analysis, a new equivalent graph must be created. The bipartite graph structure used in the prior small-worldness analysis will not yield representative results because wires are also represented by nodes, yielding far more neighbors than actual electrodes. Instead, for each electrode connected by a single wire, edges were created between every combination of electrodes to simulate being directly connected:




Figure 11: Process of creating equivalent non-bipartite graph from wire-electrode connections for a single example wire.

With this new graph, the number of electrodes within radius r from a random electrode can be computed using the *ego_graph* module within the *NetworkX* Python package for graph analysis.

Given that the network produced by Pantone et al. was an explicitly pure small-world network, it is expected that the $N(r)$ vs. r plot for the original configuration will be quite

steep; due to the nature of small-world networks, all neurons are accessible within a small number of edges, so examining the behavior for large r is not meaningful.

It was hypothesized that introducing a nested-cluster configuration of electrodes could improve the topological dimension of the network. Specifically, electrodes are arranged in recursive clusters of four, with the spacing between levels of clusters doubling at each level. For example, a sample electrode configuration for 16^2 electrodes is shown in Figure 12:




Figure 12: Nested clusters of 16^2 electrodes. Spacing between clusters is doubled at each level to decrease the connection probability between distant low-level clusters.

3.2 Results

The $N(r)$ vs. r plots for both the original grid configuration and the nested clusters configuration are shown in Figure 13:




Figure 13: The number of nodes within radius r for 64^2 electrodes in the grid and nested clusters configurations. The pink noise model was used for wires in both configurations, with $r_e = 0.4$, $\alpha = 2$, $\lambda = 30$, and $n = 201$.

Per the definition of the topological dimension, the slope of a fitted power law to the curves in Figure 13 is the metric of interest. However, given that the original network had nearly every electrode accessible within two edges, fitting a power law to two points is not meaningful. This being said, its behavior clearly points to small-world connectivity, which is the final label of interest.

For the nested clusters configuration, three points are available before running out of electrodes. However, fitting a power law to three points still lacks significance because the definition of the topological dimension explicitly requires large r . This is limited by the computational power required to simulate much larger networks, which the author could not accomplish over the brief period of the course. However, it is promising that the introduction of nested clusters produced behavior that would clearly correspond to a finite topological dimension at low r ; namely, the slope over multiple points is promisingly constant on a log-log scale. Furthermore, the introduction of clusters drastically lowered the initial connectivity, which is appropriate for HMNs.

4 Discussion

In this work, the results of Pantone et al. [4] were reproduced from scratch, and the analysis of the modeled hardware was extended to include a preliminary calculation of the topological dimension of the equivalent network. It was hypothesized that introducing a nested-cluster electrode configuration could bend the $N(r)$ vs. r curves of the networks and possibly produce hierarchical modular networks instead of pure small-world networks. Due to computational and time restrictions associated with the course for which this project was completed, the resulting analysis was limited to a small (64^2) number of electrodes. As a result, this limited the scope of the topological dimension analysis, as large r are not feasible with smaller networks. However, initial results with these smaller networks show behavior consistent with hierarchical modular networks instead of pure small-world networks. Namely, the initial connectivity is dropped by an order of magnitude, and the small number of available points on the resulting $N(r)$ vs. r demonstrate behavior consistent with a finite topological dimension.

5 References

- [1] C. C. Hilgetag and A. Goulas, “Is the brain really a small-world network?,” *Brain Structure and Function*, 2016.
- [2] G. Odor, R. Dickman, and G. Odor, “Griffiths phases and localization in hierarchical modular networks,” *Nature Scientific Reports*, 2015.
- [3] P. Moretti and M. A. Munoz, “Griffiths phases and the stretching of criticality in brain networks,” *Nature Communications*, 2013.
- [4] R. D. Pantone, J. D. Kendall, and J. C. Nino, “Memristive nanowires exhibit small-world connectivity,” *Neural Networks*, 2018.