Influencing Dynamics in Neural Networks

Jeff Klukas
John Beggs, Advisor

April 3, 2006

Abstract

Experimental readings from rat cortex in our lab have demonstrated a robust ability of neural networks to maintain critical point dynamics. Whether that ability stems from the activity of a regulatory system or is an inherent property of the network, however, remains unknown. Throughout our investigations, a computational model has served both as a useful diagnostic tool in developing measurements of dynamics and as a predictor of manipulations to be investigated in tissue samples. It demonstrates a robustness similar to the biological samples, suggesting that the network structure alone may be sufficient to maintain stable dynamics. Using the model to explore network parameters, we have identified the distribution of connection strengths between nodes in the network as having a clear influence on the dynamical behavior of the system, and further, we have explained deviations from that relationship by identifying particular connection patterns that link attractive behavior with an inverse branching structure.

1 Introduction

Studies have shown a strong correlation between specific cognitive functions and precisely repeating patterns of activity in the brain in a variety of species, and these patterns are widely hypothesized to store information [1][2]. In our own lab, we have recorded spontaneous neural activity *in vitro* over a period of several hours, generating series of short runs of correlated activity that can then be sorted into statistically significant families of patterns. While these runs show a greater similarity to one another than would be expected by chance, the members of any particular family are rarely identical. By seeking trends in these families, we can see how the runs converge

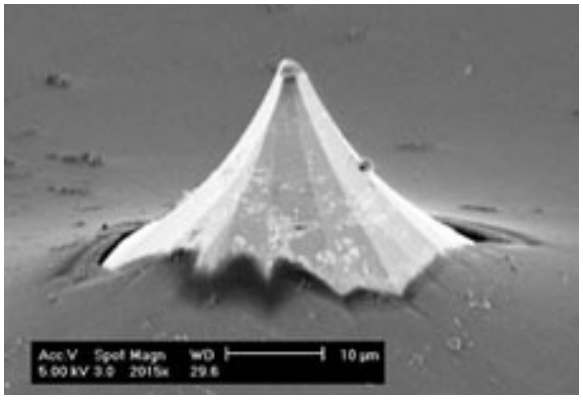

Figure 1: A single electrode in the array. Each electrode will have contact with a population of neurons.

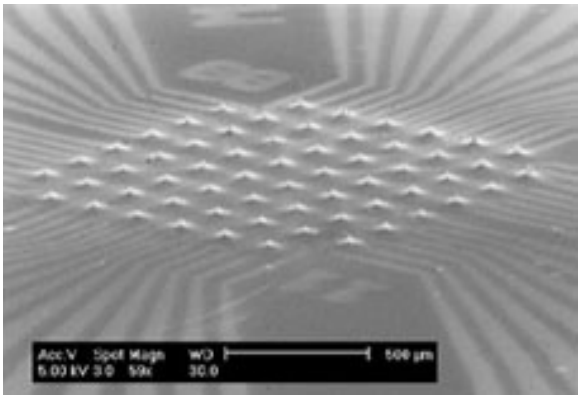

Figure 2: The array of 60 electrodes.

or diverge with time. With this approach, we hope to gain some insight into the mechanisms of storage in the brain.

Theoretical literature has presented several models of brain computation based on different assumptions as to the brain's dynamics. Some have spoken of attractive behavior [3][4], where a diverse range of inputs will cause activity to be drawn into a set pattern as the brain relates new information to previous experiences, a system optimized for categorization of inputs. Others have predicted chaotic activity that would allow rapid discrimination between similar inputs [5][6][7]. A third case calls for neutral activity at a critical point that best preserves information [8][9]. Experiments have not yet generated evidence that would decisively confirm one of these theories in the brain's expressed dynamics.

2 Data Acquisition

Previous studies have shown that local field potentials can serve as an equally, if not more, effective predictor of animal behavior than activity spikes in single neurons [10]. Therefore, in searching for patterns of activity in the brain, we want to record activity generated by populations of neurons rather than individual cells. While studies have been conducted *in vivo* focusing on single-neuron activity, we can most efficiently investigate the ensemble level through an *in vitro* preparation.

Our lab has adopted two separate preparations of rat cortical tissue: acute slices and organotypic cultures. The acute slice is a less complex and more immediate preparation, but has the disadvantage that many long-range connections between neurons are severed, constraining activity to a more local radius. In cultures, however, active neurons will grow new long-range connections over the course of several weeks. Both preparations are examined on a micro-electrode array (MEA) with sixty nodes arranged in a grid pattern at an inter-electrode spacing of 200 μm where each node is approximately 30 μm in diameter. Figures 1 and 2 show the structure of the MEA.

When placed on the array in a solution with elevated potassium, acute slices will show spontaneous activity for up to ten hours while cultures can produce recordable activity for several days. During the recording period, the electric potentials across the nodes of the array run through an amplification system and get stored to hard disk at one millisecond intervals.

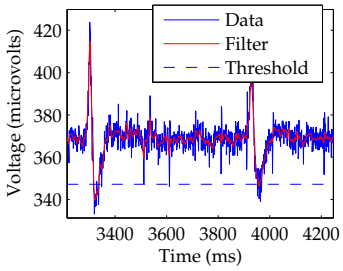

Figure 3: A threshold is applied to data from one channel of the MEA.

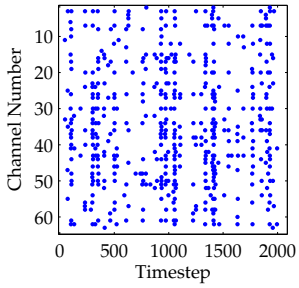

Figure 4: A section of a timeraster generated from processed data.

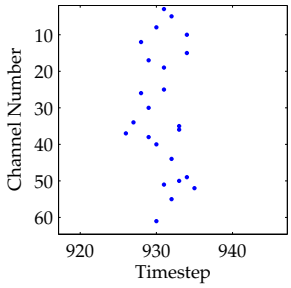

Figure 5: A run located within the timeraster.

Once activity dies out, the stored data are converted to a filtered waveform and processed with a thresholding algorithm to separate noise and nominal activity from significant population spikes. A population spike occurs when a group of neurons in the vicinity of the electrode fire nearly simultaneously. The resultant binary data can then be represented in raster form, showing which nodes were active in each timestep (Figure 4). In the thresholding algorithm, the data are also downsampled, resulting in a 4 ms gap between frames. As we are interested in how activity evolves temporally, we must carefully choose the duration of each timestep to best capture the cause-and-effect relationships between population spikes. The 4 ms spacing approximates the speed at which signals should be able to travel on the given spatial scale and should provide the

best resolution for this investigation. Figure 5 shows an isolated run within a timeraster generated from experimental data. Such a run can also be shown in a grid representation as presented in Figure 6.

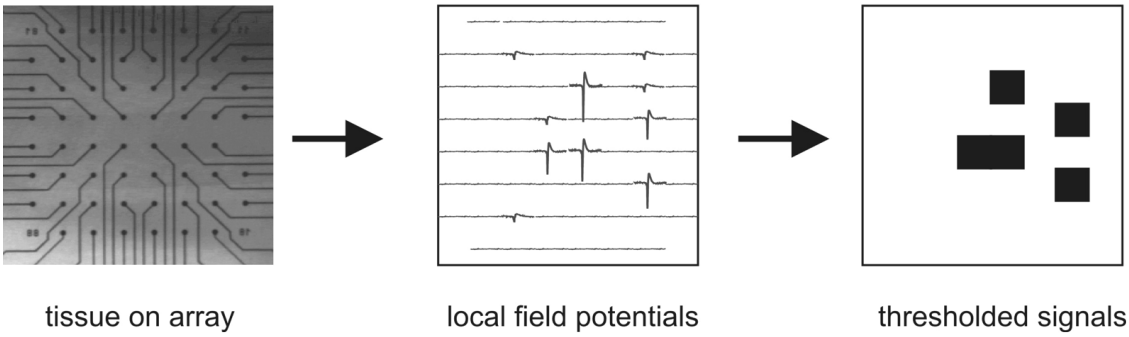

Figure 6: Data from each timestep can be represented on a grid that mimics the geometry of the MEA. Once signals from each of the 60 channels are processed individually, active electrodes are marked in black. Looking at grid representations from successive timesteps reveals how activity jumps throughout the network over the course of a single run.

3 Characteristics of Sample Data

The similarity of one run to another can be quantified through a measure defined as the intersection (the number of electrodes that show activity in both runs) divided by the union (the total number of different electrodes that show activity in either run):

$$\text{similarity} = \frac{A \cap B}{A \cup B} \tag{1}$$

This measure can be used in the sorting algorithm that then groups these runs into statistically significant families, but the question still remains of how to characterize the dynamics found in those families.

As a first step to characterizing dynamics, we define a branching parameter, σ, that measures the ratio of activity between frames where A_n represents the number of active electrodes in timestep n:

$$\sigma = \frac{A_n}{A_{n-1}} \tag{2}$$

In a chaotic system, we would expect a branching parameter greater than one such that similar starting conditions would lead to multifarious outcomes based on the sheer number of electrodes that could become active. Conversely, we would expect a branching parameter less than one to cause attractive dynamics as a diverse system simplifies to a repeated state over time. Our hypothesis suggests, then, that a balance between these two states where $\sigma \approx 1$ would lead to the sustained behavior of neutral dynamics.

The classifications of dynamical regimes above can also be related to the previously discussed similarity measure. If conditions in a chaotic system become more diverse with time, we should be able to track that evolution in our similar families of runs. To aid in the presentation of graphs relating to that evolution, we prefer to present the relationship between runs through a difference measure:

$$\text{difference} = 1 - \text{similarity} = 1 - \frac{A \cap B}{A \cup B} \tag{3}$$

A plot of the difference between the runs in a family at each timestep of the run leads to a difference trajectory that can represent the dynamics of that family. As seen in Figure 7, we would expect attractive activity to generate a negative slope, chaotic activity to generate a positive slope, and neutral activity to show little variation.

As an extra gauge to check against our other measures, we use the Lyapunov exponent to assign a value to the divergence of the system. We fit a least-squares line to the difference trajectory, then characterize the ratio of its terminal and initial values by the following equation where d represents the difference measure:

$$\lambda = \log_2 \left(\frac{d_i + \Delta d}{d_i} \right) \tag{4}$$

4 The Model

4.1 Overview

The expenditure of resources required to run experiments serves as one of the major obstacles encountered in conducting our research. Since a great deal of time and money can be spent preparing

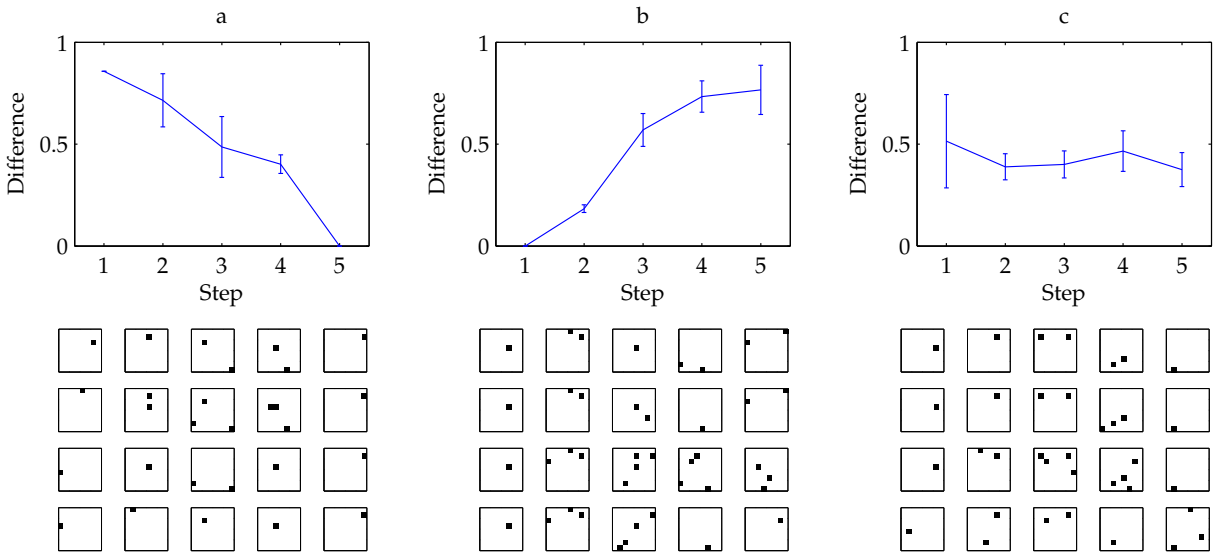

Figure 7: Although the average dynamics of a system are most often neutral, individual families within a system will show a variety of behaviors. Shown above are a selection of trajectories from a single system and the families of runs that produced them exhibiting (a) attractive, (b) chaotic, and (c) neutral behavior.

and recording from a sample, there is a need for a predictive model to guide the experimental process. The major work of this investigation has focused on using and modifying a model formed in previous experiments to develop a better understanding of the functioning of neural networks and to make predictions of analogous manipulations to be attempted in experimental samples.

The model developed in our lab for previous studies assumes a homogenous network of nodes that make a predetermined number of random connections to other nodes. Connection strengths are modeled as transmission probabilities which are then assigned to each connection according to a predetermined exponential decay distribution (see Figure 9). Once connections have been assigned, the model produces an activity matrix by simulating a predetermined number of timesteps where each node has a small probability of showing spontaneous activity while previously active nodes may transmit activity to new nodes according to their assigned transmission probabilities. After firing, nodes experience a refractory period of several timesteps during which they cannot become active, modeling the repolarization time necessary for a group of neurons to restore their ionic balance to a normal state.

4.2 Using the model to test our analytical methods

In addition to its function as a predictor of experimental activity, we were also able to use the model as a diagnostic tool to test the functionality of our analysis algorithms. After processing data from several different experiments under a range of pharmaceutical manipulations and generating graphs of the average difference trajectories over all the families in the experiments, we found no significant indication of chaotic or attractive behavior; the average difference trajectory was always relatively flat. Even analysis of initial data from the model produced the same neutral behavior, leading us to believe that our implementation of similarity and difference measures might simply be poorly suited to detecting the kind of dynamical changes we were looking for. Modifications to the model, however, were indeed able to produce trajectories significantly removed from a neutral position. It was from these investigations that we developed our first concept of a parameter that can reliably change dynamics.

Under normal conditions, we would expect the sum of transmission probabilities from any particular cell to be constrained to one, which should lead to a branching parameter value near one as well. In the model, we have found a direct relationship between these two values.

5 Influencing Dynamics in the Model

5.1 The bias parameter

While initial investigations pointed to the sum of transmission probabilities as a primary determining factor in the evolution of network dynamics, we eventually found that the distribution of those probabilities could have an even more profound effect, even when the sum remains constrained to one.

As mentioned in the structure of the model, transmission probabilities are assigned according to an exponential decay distribution. In the following equation, j represents the connection number, k is a scaling constant, and B is a bias parameter:

$$P = ke^{-Bj} \tag{5}$$

It turns out that B can be a reliable predictor of the model's dynamics, at least in the chaotic and neutral regimes (see Figure 8), and furthermore that it seems to play a crucial role in the development of critical point conditions in the network. Previous investigations into information transmission found that the critical case of neutral dynamics, with a branching parameter near unity, optimizes the system. More recent work has found a critical point for information storage when B is near 1.4, a condition that happens to be conducive to neutral dynamics as well [8]. Other studies as well have shown strong evidence that the structure of real neural networks favors this value [11].

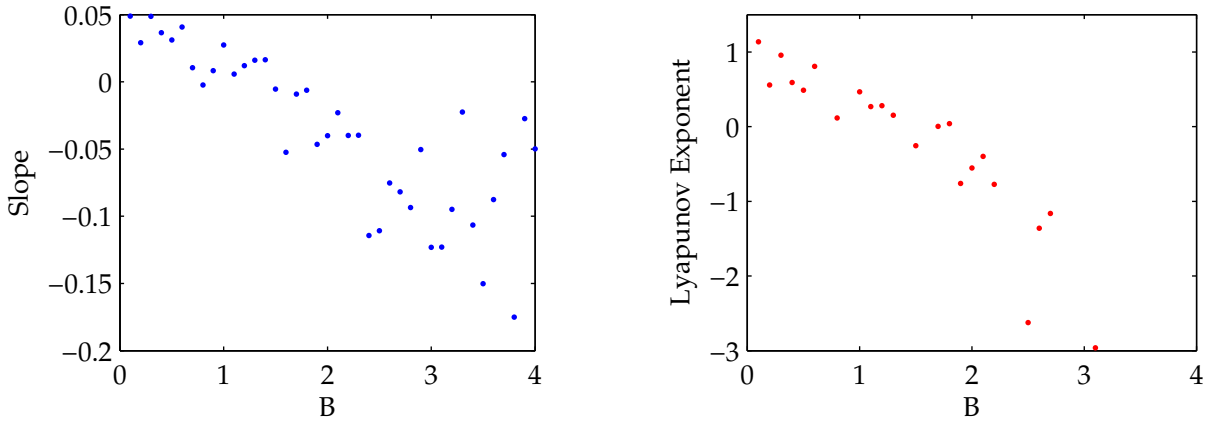

Figure 8: For the homogenous model, both the slope of the difference trajectory and the value of the Lyapunov exponent steadily decrease as more bias is introduced to the weight distribution.

5.2 Bias effects in heterogeneous networks

The influence of bias on the distribution of connection strengths turns out to be a remarkably robust effect which extends to more complex manipulations of the distribution. To test the validity of our assumption of homogeneity in the network, we developed two new methods of assigning transmission probabilities which allowed for several different types of nodes with differing weight distributions. According to the first, a multiple-bias method, each node is assigned a distribution with a unitary sum of probabilities, but with one of several values of B determining its bias. In contrast, a scaled-bias method uses a single value of B to generate a distribution, but then assigns scaled versions of that distribution to individual nodes such that they are no longer constrained

to a unitary sum of outgoing probabilities though the network average would still be one. Both of these modified methods were shown to produce attractive, chaotic, or neutral dynamics as the overall bias was changed.

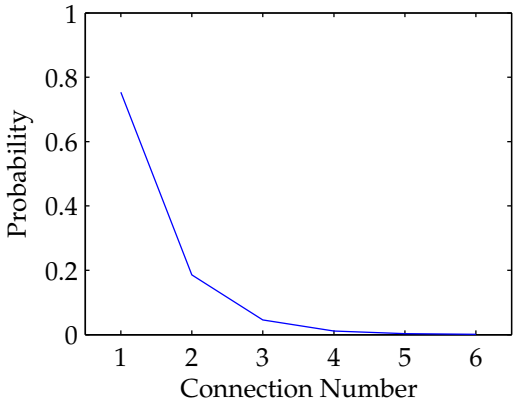

Figure 9: A sample exponential decay bias distribution used in the homogenous model

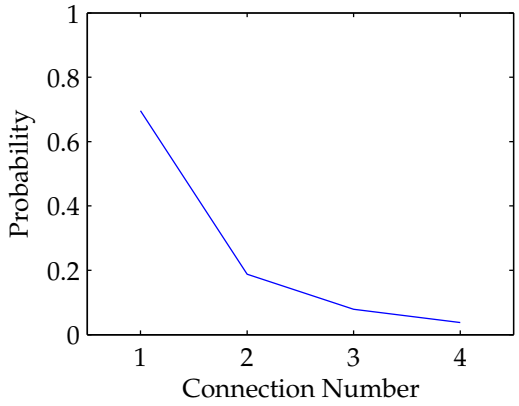

Figure 10: A sample graph of the global average probability distribution for the multiple bias model

To better assess if overall network bias exerted a predictable effect in all cases, we developed a global probability distribution measure. By sorting the transmission probabilities for each node from strongest to weakest and averaging the values for each connection number over all nodes in the network, we can make a plot that looks very similar to the simple exponential decay of the original model (Figures 9 and 10).

6 Evolutionary Modeling

While Figure 8 attests to a definite correlation between B and the difference slope, the unpredictability that arises for B greater than 2 indicates that the particular structure of the connections between nodes may have a more profound impact on observed dynamics than we originally thought. Since the model randomly generates connections between nodes for each new trial, two separate instantiations of the model given the same initial parameters may end up with highly dissimilar connection matrices. The particular structure of those matrices seems to have little effect on dynamics except in the case of highly pronounced connection bias, but understanding the

conditions that might push the network towards a more attractive or chaotic state could give us some additional insight into the significance of network dynamics.

To make patterns in the connection structure more obvious, we wanted to coax the model into an extreme state and compare the features of its connection matrix with those of a randomly generated matrix using the same bias parameter. We chose to set B at 3.0, well within the region where we expect to see a substantial range of slopes, then generated ten different instantiations of the model. Since our measurements depend on spontaneous activity that differs with each run, we sampled the slope for each instantiation on ten separate time intervals and sorted the results to give us an idea of the range of activity that could be produced by each connection matrix (Figure 11, "Original Models").

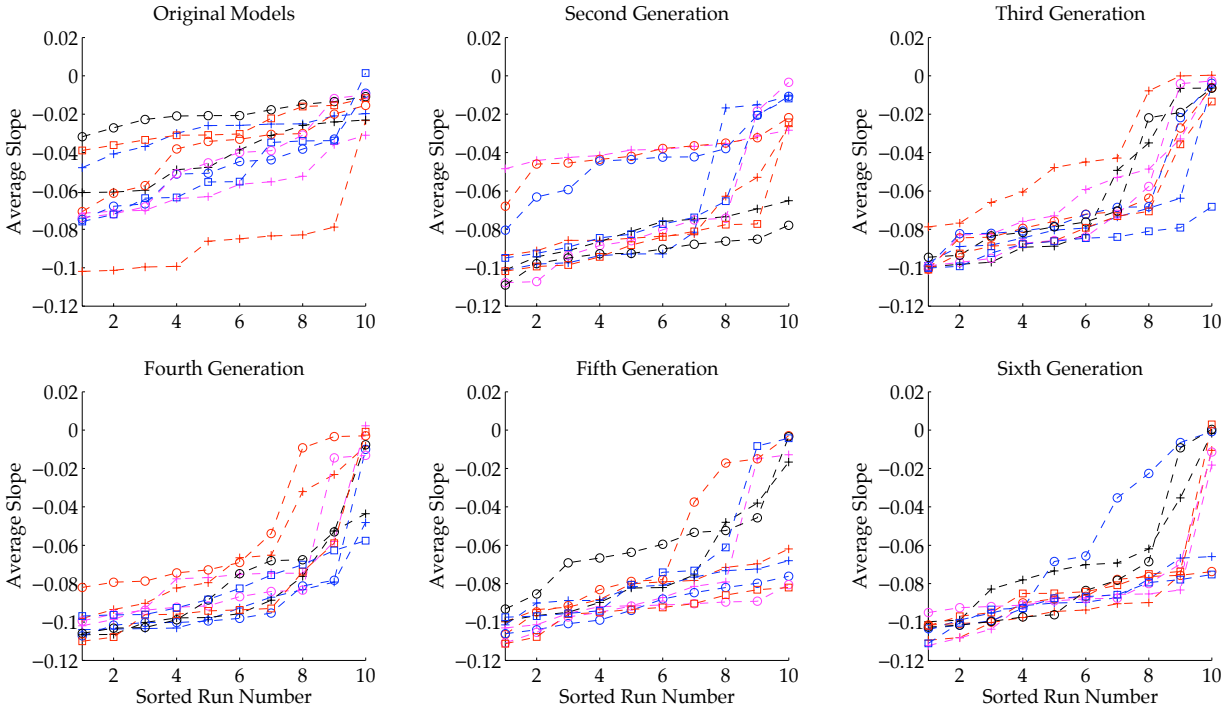

Figure 11: Graphs of the average slope of the difference trajectory for the member models of each generation. Each line corresponds to a particular matrix of connections, and each of the ten points on a line correspond to the average slope of the difference trajectories measured for the families of runs generated by a simulated period of activity. For each graph, the line that shows the lowest average slope is used as the template connection matrix for mutations in the next generation.

To push the model into a more extreme state, we implemented an evolutionary approach. From our original set of ten instantiations with $B = 3.0$, we chose the one that showed the most

attractive behavior (the lowest line in the graph of original models) and mutated its connection matrix. We created ten new instantiations by randomly choosing two nodes in the parent matrix and switching their strongest connections (probability of 0.950) with their second-strongest connections (probability of 0.047). Like before, we ran each of these through ten separate periods of spontaneous activity and graphed the resultant slopes. In this second graph, the lines are more closely packed, indicating the greater similarity between them. Notice, though, that some of these lines show a range of slopes that lies lower on the graph than that of the parent.

To push the model even further, we continued this evolutionary process through several more generations, each time generating new instantiations as variants from the most attractive state in the generation before. We could then track changes in the network structure and compare these with the unrelated networks from the original ten models. Not every generation yields a model with lower slope than was seen in the previous group, but the general trend does leave us with some significantly lower-slope models to investigate, the best of which comes out of the fifth generation. Ideally, we would have liked to look at single connection changes and consider a wider variety of mutations at each stage, but the intense computational requirements of running simulations for multiple models constrained us to a somewhat more coarse resolution. We will show in our analysis, though, that a less precise model is still sufficient to identify several interesting relationships between connection structure and dynamics.

7 Observed Patterns in the Network

Since our evolutionary model deals only with highly biased connection distributions, our analysis of the network structure can focus primarily on the strongest connections made by each node. With this approach, we can look for trends in the connection matrices of different instantiations by asking two questions: first, do some nodes have a large number of strong incoming connections, or are the strong connections more evenly distributed; and second, if we follow the strong connections from one node to the next, do we see small circular systems that act independently or long strings that affect a large portion of the network?

The first question can be answered using a rather simple algorithm that counts the number of

strong connections that point toward each node. We define the"connectedness" of a node to be zero if no other node is strongly connected to it, two if there are two such connections, etc. We must be careful here not to confuse connectedness, which concerns the distribution of strong outgoing connections throughout the network as a whole, with the distribution of incoming connection strengths within each node as defined by the bias parameter. The connectedness data are shown in Table 1, where the more attractive models seem more heavily weighted toward nodes with only one incoming connection. Such a result suggests that attractive behavior may be encouraged by having strong connections distributed homogeneously throughout the network rather than concentrated on certain nodes that act as "hubs" with many incoming connections.

Generation	Incoming Connections				
	0 con.	1 con.	2 con.	3 con.	4 con.
High Slope	24	21	15	3	1
Parent	21	28	10	4	1
Second	22	26	11	4	1
Third	21	27	12	3	1
Fourth	19	30	12	2	1
Fifth	19	29	13	3	0
Sixth	20	28	12	4	0

Table 1: Connectedness of nodes by generation. "High Slope" refers to the top line on the graph of original models in Figure 11 while the other rows represent the most attractive member of the respective generation. Initially, we see variety in connectedness, with some nodes much more heavily connected than others, but the later, more attractive models have their connections more evenly spread out so that nodes are most likely to have exactly one incoming connection.

Looking for strings involved a more complex algorithm, but also yielded interesting results. Considering that individual nodes must wait through a long refractory period after each firing, we define a string as a series of strong connections between nodes that does not revisit any of its members, and we expect that long strings might allow for stable patterns that would contribute to attractive behavior. If we systematically visit each node in a given model and follow the strong connections through the matrix, we can check for repeated nodes and record the strings we find. Using that method, though, we might find duplicate strings where our starting node lies in the middle of a longer string that has already been counted. In response, we simply compare those strings to one another and throw out any of them that appear as a subset of a longer string.

Table 2 shows the distribution of string lengths for each of the generations we are concerned with. Here, attractive behavior seems to coincide with longer strings, just as we would have expected. A large number of smaller strings should allow a great variety of behavior in the network while fewer, longer strings should allow many different input configurations to feed into the same output behavior. Indeed, when we look at the strings produced by a model that shows particularly low slope, we see that the majority of the strings end in the same sequence of two or three nodes, suggesting an inverse branching structure. It seems that many "limbs" that are short strings of activity in their own right tend to feed into a main "trunk" pattern that concludes the path that we recognize as a distinct string.

Generation	String Length											
	3	4	5	6	7	8	9	10	11	12	13	14
High Slope	4	5	2	5	1	6	0	3	2	2	1	0
Parent	1	1	4	5	4	7	3	4	0	1	1	1
Second	1	2	3	4	4	7	3	2	1	1	1	2
Third	1	1	5	5	4	6	4	3	0	1	1	1
Fourth	1	1	4	5	5	6	2	2	1	0	2	2
Fifth	1	1	3	4	4	7	2	2	2	0	2	3
Sixth	1	1	3	3	5	7	1	2	3	2	1	2

Table 2: Abundance of string lengths by generation. Each string follows strong connections from node to node in the network, terminating when a connection leads to a node that has already been visited. Length indicates the total number of nodes visited.

We expect these two effects to work in tandem to generate attractive dynamics, and perhaps the more homogenous spread of incoming connections throughout the network occurs as a direct effect of having longer strings. No single node in a long string need have a large number of incoming connections to allow several small strings of activity to flow together; the small strings can instead connect node by node at many points along the length of the trunk. In this case, we would expect most nodes to have only one incoming connection as they pass activity along a string, while a few unconnected nodes serve as the end points of branches and a few nodes with two or three incoming connections serve as junctions into the trunk, avoiding the need for substantial hubs that mediate a great number of connections. The result is a high concentration of activity in the trunk region even though the connections are fairly evenly distributed.

8 Conclusion

Observations in our lab indicate neutral dynamics as basic properties of neural networks that are remarkably resistant to manipulations. While the application of certain drugs in experimental trials has been able to cause variations in dynamics large enough to be deemed significant, the effect has been underwhelming. In most cases, manipulations must be radical enough to shut down coherent activity completely before an effect can be detected.

The model has also been shown to operate with a resistance to a great variety of manipulations, indicating that self-regulation may be an inherent property of the network structure. Nevertheless, some predictable behaviors have been identified. Dynamics can be determined by changing the sum of transmission probabilities or the bias of the weight distribution for each node in homogeneous networks or by changing those properties globally in the heterogeneous case.

While the random nature of the connections created in the model has a minimal effect on the dynamics of the system for most sets of input parameters, we do see a significant range of dynamics at high values of the bias parameter as a result of the particular network structure generated by the model. Analyzing patterns in the connections between nodes reveals that attractive behavior can be correlated both with a fairly even distribution of strong connections between the nodes and with an abundance of long uninterrupted strings of strong connections between unique nodes. We expect that this encourages attractive behavior by allowing various stable branches of activity to feed into one main trunk such that most runs of activity terminate in the same pattern regardless of where they start in the network.

Investigations like these that look at local field potentials are beginning to fill in a gap in our understanding of network dynamics in the brain. As we develop an impression of the brain's intermediate levels of organization as a supplement to extensive studies already available at the whole brain and single neuron levels, we are moving toward the development of a general theory of neural computation.

References

[1] Gentner TQ, Margoliash D. Neuronal populations and single cells representing learned auditory objects. *Nature*. 2003 Aug 7; 424(6949):669-74.

[2] Wilson MA, McNaughton BL. Reactivation of hippocampal ensemble memories during sleep. *Science*. 1994 Jul 29; 265(5172):676-9.

[3] Cossart R, Aronov D, Yuste R. Attractor dynamics of network UP states in the neocortex. *Nature*. 2003 May 15; 423(6937):283-8.

[4] Hopfield JJ. Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci U S A*. 1982 Apr; 79(8):2554-8.

[5] Freeman WJ. Neural networks and chaos. *J Theor Biol*. 1994 Nov 7; 171(1):13-8.

[6] Schiff SJ, Jerger K, Duong DH, Chang T, Spano ML, Ditto WL. Controlling chaos in the brain. *Nature*. 1994 Aug 25; 370(6491):615-20.

[7] Sinha S, Ditto WL. Controlling neuronal spikes. *Phys Rev E Stat Nonlin Soft Matter Phys*. 2001 May; 63(5 Pt 2):056209.

[8] Haldeman C, Beggs JM. Critical branching captures activity in living neural networks and maximizes the number of metastable States. *Phys Rev Lett*. 2005 Feb 11; 94(5):058101.

[9] Maass W, Natschlager T, Markram H. Real-time computing without stable states: a new framework for neural computation based on perturbations.*Neural Comput*. 2002 Nov; 14(11):2531-60.

[10] Andersen RA, Musallam S, Pesaran B. Selecting the signals for a brain-machine interface. *Curr Opin Neurobiol*. 2004 Dec; 14(6):720-6.

[11] Brunel N, Hakim V, Isope P, Nadal JP, Barbour B. Optimal information storage and the distribution of synaptic weights: perceptron versus Purkinje cell. *Neuron*. 2004 Sep 2;43(5):745-57.