Shannon C. Barba

Dr. Demian Cho

FURSCA End of Summer Report

Summer 2023

Characterization of Synergistic Information Processing in Neural Network Motifs

Over the course of the summer, my research was focused on how different neuron groupings affect their ability to pass a signal. Large networks of neurons can be broken down into many different groups of 3 neurons, called motifs. While it is an extremely difficult task to analyze highly connected networks, a better understanding of motifs could help to approximate how these networks should behave. With the help of my mentor, Dr. Cho, the goal was to develop a MATLAB model to simulate the 13 possible neuron motifs with varying initial conditions. We would then need to collect data from the model and perform calculations to derive a motif's mutual information, or how much the behavior of one neuron told us about the others. From there, we could calculate the influence of initial conditions on the ability to pass a signal (synergy), and as a result we would be able to treat motifs more as building blocks. After that, we could verify that our model could be used to approximate a higher-order neural network.

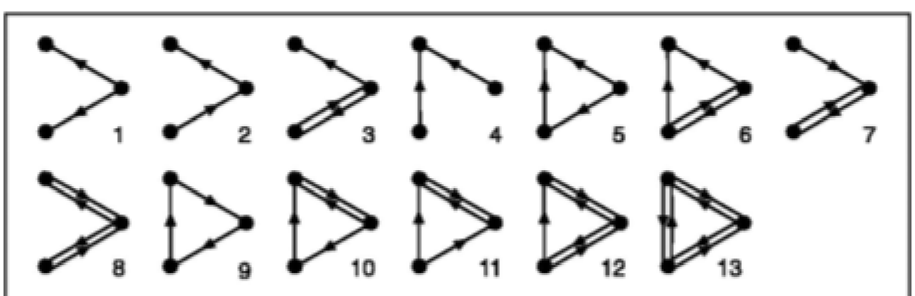

Figure one: All possible neural motifs.

Thanks to my work last summer, I already had a MATLAB model that could produce a correct output for a specific arrangement of three neurons. I began to work on editing the model for what I would need this summer, as it would need to be more complex to accommodate a larger variety of connections. At first I was struggling to get a program that could simulate all the motifs, but I was eventually able to figure it out and create the most powerful tool for my research. This model was able to simulate all 13 motifs, with varying initial conditions (whether or not a neuron starts with firing), weights (the strength of a signal sent), and thresholds (signal strength required to activate a neuron) for each neuron in the motif. Once I verified the model correctly produced a stream of outputs for a given motif, I began taking data of each motif and all possible initial conditions. With 8 possible initial conditions, there were 104 data entries.

```
raster =

     1     1     0     1     1     1     1     1     1
     1     0     1     1     1     1     1     1     1
     0     1     1     1     1     1     1     1     1
```

Figure 2: An example of a time raster data entry. It shows all three neurons eventually reaching a firing state.

The next step of the project was to quantify the interactions between neurons using these time raster data entries. Using the probability that certain neurons fired at any given time for a given motif, we were able to calculate the mutual information between all three neurons. Mutual Information is a quantity that tells us how much influence different events have on each other. I coded in a way to get probability tables for different motifs and initial conditions, and then got to work calculating the mutual information. While the formula for mutual information looks simple, it is not. It took me many times incorrectly performing the calculation until I was able to correctly do it and get a consistent value. I wanted to create a program that would perform it automatically, but it proved much more time efficient to do it by hand.

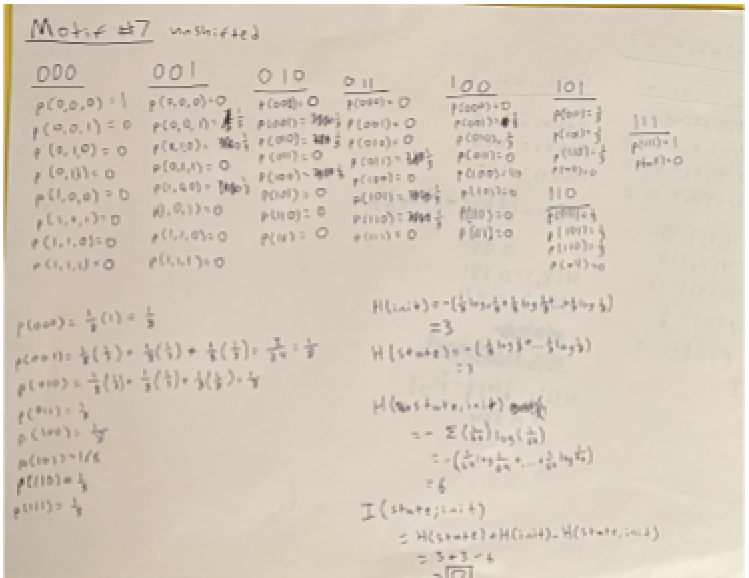

Figure 3: Calculating the Mutual Information for motif seven. It is symmetric, which results in a low value.

Upon looking at the mutual information for a select group of motifs, something stood out to me. Neuron motifs with excessive symmetry had much lower mutual information than neurons with little symmetry. It made sense intuitively, since the symmetry made it so the network wasn't biased toward any neuron in particular. As a result, every outcome was equally likely, and the state of an individual neuron did not depend heavily on the other two, but rather the initial weights and thresholds. The find was very interesting to me, and I felt accomplished in the fact that I had successfully taken numerical value and translated it into a behavioral pattern.

While these observations were in the right direction, much work remains to be done in the field to fully unlock the secrets of these higher-order networks. However, I would consider my FURSCA project successful: I had been able to characterize the output of different motifs, based solely on their initial conditions. Going into the project, I knew all the secrets of the brain were not about to be answered, but I knew that substantial progress would be made. Deriving a function to correctly predict a motif’s behavior more accurately utilizing the collected data would be the next step. I am extremely curious about what remains to be discovered and plan to continue working on the project on my own time with Dr. Cho. We look forward to making more progress in the future, as well as presenting at the Elkin R. Isaac Research Symposium. I was able to learn a huge amount about topics (such as Information Theory and Neuron Approximation Schemes) that are not generally understood very well, and made a big advancement into working on the future of the field, which is a fully artificial, working brain. I want to give a special thanks to FURSCA and the Charles J. Strosacker Foundation, for providing me and so many others with the opportunity to dive into subjects we are curious about, and my mentor, Dr. Cho, for pushing me to learn more about areas I knew nothing about and showing me that sometimes making mistakes is the best way to learn.