



# A Mathematical Approach for the Reconstruction of Neural Networks

Andrew Wilson

July 31, 2009

Department of Mathematical Sciences  
George Mason University

*Project Mentor*

Dr. Domenico Napoletani

---

\*This final report is submitted in partial fulfillment of the National Science Foundation REU and Department of Defense ASSURE Program George Mason University.

## Abstract

Considering the abundance of recent work in the fields of network theory and neuroscience, it is no surprise that there has been extreme interest in reconstructing the topology of neural networks. Previous methods have investigated this problem experimentally with a focus on determining the topology of a single given network. In this paper we take a mathematical approach, extending the augmented sparse reconstruction method for protein networks to a system of neurons in an attempt to find a more general technique [8]. This technique employs  $L^1$  minimization to reconstruct a network from a set of noisy trajectories under a variety of initial conditions. Our investigation focuses on the method's performance on an experimentally documented network of neurons in the rat hippocampus [3]. Each neuron is approximated by the FitzHugh-Nagumo model. When observing the neurons' behavior over very short time intervals immediately after initialization, the method is quite successful in its attempt to recover the structure of the given network. In light of this result, we then discuss the research that must still be done before this method can truly reconstruct general neural networks.

# Contents

<b>1</b>	<b>Method</b>	<b>1</b>
1.1	Motivation . . . . .	1
1.2	Augmented Sparse Reconstruction . . . . .	2
1.3	Adapting Augmented Sparse Reconstruction . . . . .	3
1.4	Networks . . . . .	5
<b>2</b>	<b>Results</b>	<b>7</b>
2.1	The Settling Interval . . . . .	7
2.2	Experimental Concerns . . . . .	9
2.3	The Barabási-Albert Network . . . . .	12
<b>3</b>	<b>Future Work</b>	<b>13</b>

# 1 Method

## 1.1 Motivation

Say that a group of  $N$  neurons are connected in some way so that any given neuron in the network is influenced by the behavior of other neurons. That is to say, when a connection is formed from neuron  $n_i$  to neuron  $n_j$ , the membrane potential of  $n_i$  affects the potential of  $n_j$ . If we examine the membrane potentials of all neurons in the network over time, can we gain some insight into how the neurons have been connected? This process is known as the reconstruction of neural networks.

There are many clear applications of this goal to the health sciences, including the study of neural diseases such as epilepsy and the repair of damaged neural networks. It is possible that the field of biologically inspired computing could benefit from a more exact mapping of the interactions within the nervous system. The list of possible applications is long, and there is little doubt that this reconstruction is an important task.

It is important to note that this work is part of a long lineage of attempts to understand the interactions within a biological network by observing the behavior of the nodes of the network [11, 6, 4]. The paper that directly inspired this work is a particularly successful attempt at mathematically reconstructing the epidermal growth factor receptor (EGFR) driven signaling cascade, an important protein signaling network [8]. Most recent reconstruction attempts of neural networks, while quite effective, have been more concerned with a rigorous formalization of available experimental data [10, 3]. Experimental techniques, while very effective for the study of a specific network, are not immediately extensible to all networks. The long-term goal of a mathematical method would be to simply feed in some trajectories to a program and receive an accurate estimate of the topology of the network.

## 1.2 Augmented Sparse Reconstruction

The primary mathematical method being utilized in this study is the augmented sparse reconstruction method [8]. This method is designed to successfully reconstruct sparse networks from noisy data sets. Many biological networks have been shown to be sparse, meaning, quite roughly, that they have a low number of connections per node [1]. The augmented sparse reconstruction method aims to reproduce the sparsest possible network from the given data by minimizing the  $L^1$  norm of the coefficients vectors. This means that, for each neuron, we seek an  $N$ -element vector  $a$ , where the  $i$ th entry represents the possible influence of the  $i$ th neuron on the our chosen neuron, that has two key properties. First, it should cause the neuron to act in a manner similar to the behavior observed in the neuron's trajectory. Second, it should have the smallest possible  $L^1$  norm, defined as  $|a| = \sum_{i=1}^N |a_i|$ . While there are many works linking sparsity and the 1-norm, such as [7], perhaps the quickest way to understand this relationship is to notice how  $L^1$  minimization limits cancellations. For example, if we know that  $\sum_{i=1}^N a_i = 0$ , the 1-norm will be minimized when  $a_i = 0 \forall i$ . When applied to the entire network, this behavior will lead to the desired sparsity of connections.

Although  $L^1$  minimization is effective for insisting sparsity, it is not able to deal with the large errors-in-variables that occur when working with a high level of measurement error, also known as noise. To handle this measurement error we append a large number of random terms with relatively small coefficients. The purpose of these random terms is to obscure the noise during the optimization phase.

Consider the following example, adapted from [8]. There is a measurement  $\tilde{X}_i = X_i + M$ , where  $X_i$  represents the true value and  $M$  the noise. If the noise level is approximately 10%, then we can let  $M = 0.1X_i$ . Then the noise ac-

counts for  $\frac{0.1}{1+0.1} \approx 11\%$  of the  $L^1$  norm. However, if these terms are augmented with  $G$  random terms  $w_g n_g$ , where  $w_g$  is small and  $G$  very large, then neither  $M$  nor any of the added noise terms will make a large contribution to the  $L^1$  norm. This makes the true value far easier to recover.

### 1.3 Adapting Augmented Sparse Reconstruction

The primary challenge of this project has been transplanting the augmented sparse reconstruction method from protein networks onto neural networks. The first step of this adaptation was choosing an appropriate model for neuron behavior. We chose to work with the FitzHugh-Nagumo model, first described in [5] and detailed extensively in [9]. In one common form of this model, similar to the one found in [12], the neuron is characterized by its membrane potential  $x$ , its recovery variable  $y$ , and the experimental current  $I$  in the following way:

$$\dot{x} = \frac{1}{\epsilon} \left( x - \frac{x^3}{3} - y + I \right) \tag{1}$$

$$\dot{y} = B(x + C) \tag{2}$$

This model is desirable for our work because it is one of the simplest yet most effective models for neuron behavior. (In the future one could certainly study the performance of our method on trajectories generated by other, more complex models.) Regardless of the model, it is important to note that only the neuron’s membrane potential can be experimentally determined; other variables, like the  $y$  term in this case, remain hidden. Since we aim to eventually work with real data sets, we would like to reconstruct the neural network from experimental data. This means that, in our method, only the  $x$  variable in the FitzHugh-

Nagumo model is considered.

After choosing the FitzHugh-Nagumo model, we then included the influences from other neurons in the network. This is a common practice, as seen in [12].

In this form the  $i$ th neuron is described by the equations

$$\dot{x}_i = \frac{1}{\epsilon} \left( x_i - \frac{x_i^3}{3} - y_i + I \right) + \sum_j A_{ji} x_j \quad (3)$$

$$\dot{y}_i = B(x_i + C) \quad (4)$$

where  $A$  is the connection matrix of the network. If there are  $N$  neurons in the network,  $A$  is an  $N$  by  $N$  matrix where  $A_{ji}$  describes the connection (or lack thereof) from the  $j$ th neuron to the  $i$ th neuron. These connections will be discussed in greater detail later in the paper.

A large part of the research was comprised of adjusting the parameters  $\epsilon$ ,  $B$  and  $C$  in Equations 3 and 4 to increase the influence of the other terms as much as possible - thus making reconstruction more accurate - while still ensuring that the neurons displayed proper spiking behavior. However, since the long-term goal is to work with real neurons when there are no parameters to adjust, we will not spend any time discussing these adjustments.

Next we adapted the augmented sparse reconstruction method to fit the FitzHugh-Nagumo model. This means that we took the many linear terms and one cubic term from the model above and included the random terms from the original augmented sparse reconstruction method. As in [8], we integrated all terms. Integration has the advantage of making the method more accepting of the noisy data sets we are likely to face. This yielded

$$x_i(t) - x_i(t_0) = a_0 + \sum_{j=1}^N l_{ji} \int_{t_0}^t x_j dt + c_i \int_{t_0}^t x_i^3 dt + \sum_{g=1}^G w_g n_g \quad (5)$$

where  $l_{ji}$ , representing the coefficients of the linear terms, are the values we would like to find through  $L^1$  minimization. Notice that, as described before, this dictionary ignores the inner variable  $y$  from Equation 3. With this approach in hand, we may move on to the actual networks of study.

## 1.4 Networks

Now let us detail the networks we explored throughout the project, which determine the connection matrix  $A$  mentioned above. Throughout the course of our project we worked with three different networks. The first network was an unweighted random network. An  $N$  by  $N$  matrix of random numbers was created; then all entries greater than a certain threshold were replaced with 1, while the rest were entered as 0. The purpose of this network was simply to ensure that the method would run properly and, once this goal was reached, we moved on to other networks.

Next we turned to a group of neurons from the rat hippocampus, as seen in Figure 1. This network was the primary vehicle of exploration in our studies. This decision was made on the basis that the best way to guarantee the behavior of a neural network was to use a real neural network. We then created a connection matrix  $A$  from this data by considering only the strongest connections to be present; the rest of the connections were ignored. This allowed us to work with a network that was both unweighted and very sparse. The lack of weighted connections simplified the process of reconstruction while the sparseness ensured that the sparsity requirement for the augmented sparse reconstruction method




Figure 1: A depiction of a network of 23 neurons from the rat hippocampus connected by the strongest documented edges from [3], amounting to 64 directed connections. Graphic created using Graphviz.

would be met. Toward the end of the paper we also include some work done on a rat matrix where fewer of the connections are ignored, but in these cases we still consider all connections to be unweighted. Future versions of this method would ideally allow weighted edges to be present.

Finally, as an attempt to move to more general networks, we worked with a network created from the Barabási-Albert model [2]. The Barabási-Albert model is a probabilistic model that produces scale-free networks through preferential attachment. Briefly, this model begins with a set of  $N$  nodes and zero edges. Edges are added one at a time. (Although the edges are usually undirected, we used directed edges to simulate the directed connections of a neural network.) The chance that any given node will be part of a new edge is directly proportional to the number of edges it already has. This aspect of the model, which is known as preferential attachment, probabilistically leads to a scale-free network [2]. For our purposes, this is notable because neural networks are almost always scale-free [1]. In short, this method simply allows us to create random networks that share some key properties with neural networks. We will discuss only a

few interesting results from the Barabási-Albert network.

## 2 Results

### 2.1 The Settling Interval

We began by creating a different random initial condition for each of the 46 variables in the rat hippocampus network (23 membrane potentials and 23 recovery variables). We repeated this 25 times to create the variability necessary for an accurate reconstruction. These initial conditions were fed into the FitzHugh-Nagumo model to generate the neurons' trajectories. These trajectories were then supplemented with noise to simulate the error level of most experimental results in neuroscience. Unless otherwise specified, the noise level was 10% of the maximum membrane potential during the interval.

We then attempted to reconstruct the network by running the augmented sparse reconstruction algorithm described above. Actual  $L^1$  minimization was done by LIPSOL, a free interior-point solver created by Yin Zhang. The accuracy of the reconstruction was scored by creating a threshold. Any coefficient greater than the threshold was counted as a connection. For each attempt, a threshold was found such that the false positive rate of the reconstruction remained at a fixed value. Unless otherwise noted, this fixed value is 0.10. We then extracted the true positive rate. Although this sort of threshold is not extensible to experimental work, as it requires knowledge of the original connection matrix to evaluate the results, it allowed us to gain some insight into the performance of the method in our simulation setting.

With arbitrary initial settings over an arbitrarily long time interval, the




Figure 2: A sample of long trajectories and settling interval trajectories. Each line represents a different randomly chosen initial condition.




Figure 3: Solid lines represent the true positive rate of the method on the settling interval (a length of approximately  $\frac{1}{20}$  of the spiking period), while dotted lines correspond to the long interval (roughly 50 times as long as the settling interval). The circle, square, and triangle indicate that the false positive rate was held at 0.05, 0.10, and 0.20, respectively.

method was not very successful. One can refer to Figure 3 to observe the results. Our first effort to improve results was the utilization attenuation coefficients, as in [8]. This process involves weighting the linear terms or the cubic term through multiplication of some value between 0 and 1. However, unlike in [8], attenuation coefficients did little to improve performance.

We then discovered that the optimal way of improving accuracy was to focus on a small time interval immediately after the trajectories began, which will be known throughout the rest of this paper as the “settling interval.” Some trajectories on this interval are shown in Figure 2. When reconstructing based only on data from this settling interval, the true positive rate rose, as seen in Figure 3. This is likely because this interval allows us to observe how the initial conditions settle into their long-term trajectories. After this interval, the neurons lock into a synchronized spiking behavior. (Interestingly, the neurons tend to settle into two groups which alternate spiking.) Since, in these intervals, the unique behavior of each neuron is lost, it is not surprising that little data about the neurons’ connectivity can be recovered.

## 2.2 Experimental Concerns

The importance of the settling interval raises some experimental questions. Since the end goal of our project is to apply this method to real neural networks, these are questions we must consider. The first issue is the feasibility of obtaining measurements during the settling interval in a laboratory setting, given that the settling interval may be extremely brief. One reason to be encouraged about this first issue is the impressive performance of the algorithm in the presence of increased noise, as displayed in Figure 3. This implies that even quite noisy measurements taken during the settling interval could be enough to allow accurate reconstruction. However, this is a question that can only truly




Figure 4: This graph represents the true positive rate of the algorithm with naturally created initial conditions. As before, the circle, square, and triangle correspond to false positive rates of 0.05, 0.10, and 0.20, respectively.

be answered with laboratory research.

The second issue is the ability to recreate the random - or mostly random - initial conditions organically. Since the algorithm requires a high level of variability to succeed, it is important to work with a wide variety of initial conditions. However, at any given moment there is almost no chance that the membrane potentials of a network of neurons can be described by any sort of random distribution. For this reason, we investigated methods for causing a network of neurons to display somewhat random behavior. Since, in an experimental setting, one only has control over the current being added to the neurons, we restricted ourselves to working with the current term in the FitzHugh-Nagumo model from Equation 3. We found that adding large blocks of current, whether uniform or individualized for each neuron, did little to recreate randomness. This is most likely because the neurons' synchronous behavior is caused by a low-level attractor. However, we had some success with a noisy block of cur-




Figure 5: This figure displays the true positive rate with added variation, as described in the text of the paper.

rent. In this procedure, each neuron received its own “white noise” current that was quite large in amplitude. If this current was kept on for a length of time - about one-tenth the length of the neurons’ spiking period - the neurons began to exhibit random behavior. We then turned off the random input current and allowed the neurons to settle into their usual behavior. Applying the reconstruction method to this settling interval yielded the results in Figure 4. Although this performance is clearly not ideal, we believe that it shows that the recreation of random initial conditions is certainly not an impossible goal.

The third experimental concern was the identical modeling of each neuron and each connection in the network. For most of our results, we held the constants in the FitzHugh-Nagumo model -  $\epsilon$ ,  $B$ , and  $C$  - fixed for each neuron. We also considered every connection to have value 1. That means that each neuron and connection in the network was completely identical, an assumption which certainly does not extend to real neural networks. We hoped that this assumption was not unfairly boosting the success of our method. To investigate this, we

allowed the constants to vary with each neuron. Figure 5 shows the performance of the algorithm with respect to allowed variation. A variation of  $a$  implies that, if  $K$  is one of the constants in the model listed above,  $K$  is replaced with  $\bar{K}$ , where  $\bar{K}$  is a Gaussian random variable with a mean of  $K$  and a variance of  $aK$ . We also allowed the connection value to vary around 1 in an analogous manner. One can observe that, for variations up to 0.2, the method was still quite accurate, meaning that the variability of real neurons should not cause a significant drop in performance of the reconstruction method.

### 2.3 The Barabási-Albert Network

Finally, we worked briefly with a network created according to the Barabási-Albert model. The main purpose of this work was to make sure that the augmented sparse reconstruction method's successes described above were not due to some unique property of the rat hippocampus network. On both the settling and long-term intervals, the Barabási-Albert network yielded performance very similar to that observed with the rat hippocampus network. This would imply that the importance of the settling interval stems not from the rat hippocampus network but from the method itself. We also attempted to create the random initial conditions with the Barabási-Albert network. However, with this network we had success with an approach that had failed with the rat hippocampus network. In this approach, we choose one neuron that is connected to all other neurons in the network. In theoretical terms, there exists a path from this neuron to each other neuron. We then held the membrane potential of this neuron constant for a period of time. This process was quite successful in de-synchronizing the other neurons and replicating the random initial conditions. Although, as the rat hippocampus network showed, this approach is not universally successful, it could possibly be modified to create the random initial

conditions for other networks.

### 3 Future Work

The goal of this project is to modify the augmented sparse reconstruction method so that it can successfully reconstruct real-world neural networks. However, real-world networks have two properties that the current method seems unable to handle: weighted connections and a relative lack of sparsity. Working with unweighted connections, as we have done here, is a fair assumption for simulation work, but the strengths of real-world connections frequently have a significant amount of variation [3]. It is possible that a modification made for the EGFR protein network, such as attenuation coefficients, may help with weighted edges, but it is also entirely possible that new work would need to be done to deal with this issue. The algorithm also performed poorly when more edges were added to the network. For example, when some of the medium-strength connections from the original rat hippocampus data were included, the true positive rate of reconstruction dropped below 0.30. Adding the weak-strength connections hurt the performance even more. The model must be able to deal with these new edges before it can be applied to real networks.

The best way to deal with these complications, and many others, might be to find a way to automatically determine the appropriate settling interval. It seems that the size of the network, the number of connections in the network, and even the constants chosen in Equations 3 and 4 affect the length of the settling interval. It is difficult to determine the settling interval at this point without simply observing the trajectories. Ideally, there would be a mathematical method for determining the settling interval of a given network so that the augmented sparse reconstruction method can be applied to this interval.

Although there is much work to be done, these early results are still quite en-

couraging. We have seen that it is certainly possible to apply the augmented sparse reconstruction method to neural networks with a high level of accuracy. Continued work in this area may indeed lead to a general mathematical method for the reconstruction of neural networks.

## References

- [1] R. Albert. Scale-free networks in cell biology, 2005.
- [2] R. Albert and A.-L. Barabási. Statistical mechanics of complex networks, 2002.
- [3] Gully A. P. C. Burns and Malcolm P. Young. Analysis of the connectional organization of neural systems associated with the hippocampus in rats, 2000.
- [4] I.C. Chou, H. Martens, and E.O. Voit. Parameter estimation in biochemical systems models with alternating regression, 2006.
- [5] R. FitzHugh. Mathematical models of threshold phenomena in the nerve membrane, 1955.
- [6] D. Husmeier. Sensitivity and specificity of inferring genetic regulatory interactions from microarray experiments with dynamic bayesian networks, 2003.
- [7] S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, New York, 1998.
- [8] D. Napoletani, T. Sauer, D.C. Struppa, E. Petricoin, and L. Liotta. Augmented sparse reconstruction of protein signaling networks, 2008.
- [9] C. Rocsoreanu, A. Georgescu, and N. Giurgiteanu. *The FitzHugh-Nagumo Model: Bifurcation and Dynamics*. Kluwer Academic Publishers, Boston, 2000.
- [10] J. W. Scannell, C. Blakemore, and M. P. Young. Analysis of connectivity in the cat cerebral cortex, 1995.

- [11] E.O. Voit. *Computational Analysis of Biochemical Systems*. Cambridge University Press, Cambridge, 2000.
- [12] G. Zhao, Z. Hou, and H. Xin. Frequency-selective responses of fitzhugh-nagumo neuron networks via changing random edges, 2006.