Computational Modeling of the Motor Cortex

A master's thesis in Artificial Intelligence

By Jimmy Mulder

Supervisors:

Erik Aarnoutse (UMC)

Stefan van der Stigchel (UU)

Student number: 3684695

Utrecht University

January 6th, 2021

Abstract

Brain Computer Interfaces (BCI's) can greatly improve the lives of people with brain disorders and disabilities, such as Parkinson's disease or locked-in syndrome. Computational models of the brain can improve the efficiency of these systems. In this thesis we attempt to build a functioning computational model of 6000 Izhikevich neurons representing the basal ganglia, thalamus and motor cortex. Compared to physiological data, the model is accurate in terms of average firing rates and spike distribution, and plausible in terms of frequency distribution. However, during a simulated task the model did not respond in accordance with the data. Additionally, this thesis explores the merits and pitfalls of automated optimization of model parameters. We found that a fitness function based on firing rates can lead to multiple global maxima, and suggest expanding the fitness function to include frequency and synchrony.

Introduction

Over the past decades significant advancements have been made in the field of Brain Computer Interfacing (BCI). The required technology has become cheaper and more advanced while our understanding of the brain has increased. BCI aims to translate brain signals into computer signals (Tan & Nijholt, 2013), usually in real-time (Donchin, Spencer, & Wijesinghe, 2000). While commercial BCI's for recreational use exist, most research is focused on assisting patients with impaired motor function (Schalk et al., 2004). Early implementations of BCI's mostly use electroencephalography (EEG) to capture brain activity from the scalp (Donchin et al., 2000) by measuring the small voltage differentials generated by activity in the upper cortical areas. Other ways to measure brain activity from the outside of the scalp include MEG (based on the magnetic field caused by action potentials), fNIRS, and fMRI (both based on BOLD-response). While fMRI and MEG provide much higher quality than their respective counterparts (fNIRS and EEG), the required hardware isn't portable which means this technology is not well suited for real-time BCI applications.

Recently the use of electrocorticography (ECoG) has received more attention (Leuthardt, Schalk, Wolpaw, Ojemann, & Moran, 2004; Vansteensel, Pels, Aarnoutse et al., 2016). With this method, arrays of electrodes are implanted under the skull, directly on top of the cortical area of interest. This gives a much cleaner signal compared to EEG. The risky nature of the placement procedure, however, means that data is sparse and almost exclusively from patients with some form of brain dysfunction. The closest data we have to what can be considered 'healthy' activity is usually obtained from epilepsy patients where the focal point is far removed from the electrodes. Because of the high potential of this technology with regards to assisting patients, getting the most value out of this limited amount of data is crucial.

At the University Medical Center (UMC) of Utrecht, the Netherlands, researchers are using ECoG-based BCI to aid locked-in patients in communicating to the outside world (Vansteensel, Pels, Aarnoutse et al., 2016). This thesis is the product of an internship at that particular research group.

Two participants are involved in this project, called Utrecht NeuroProsthesis (UNP). The participants have been implanted with strips of electrodes over the motor cortex (M1) and the dorsolateral prefrontal cortex. The data recorded from the motor cortex is used to create a 'brain

click' when the participant attempts to move their hand; the clicks can then be used to control a visual interface.

The algorithm that recognizes these 'brain clicks' is based on the key observation that activity in the lower frequency bands decreases in M1 while gamma activity increases during attempted hand movement. Two band-pass filters are applied to the recorded data; when the difference between the high band and low band reaches a pre-defined threshold for a short amount of time, a click is made. This threshold is calculated offline using a reinforcement-learning model.

While this system works quite well, the mechanisms through which these changes in activity take place are not well understood. Some literature suggests that when we wish to make a movement, the inhibitory alpha waves stemming from the Basal Ganglia (a region related to decision making and motor control) are suppressed (J. H. Lee, Whittington, & Kopell, 2013), which would explain the decrease in the lower frequencies and allow more activation in the network (which is generally measured as gamma activity). However, this has not been verified in our participants. Moreover, the findings are not consistent among all the participants. The data varies between patients and timestamps and the causes are not always understood.

With the advance of new modeling techniques and increases in computing power, we have seen a rise in papers exploring the potential for computational models of the brain in understanding of these phenomena. Some interesting progress has been made in relation to Parkinson's disease, using models to explain and enhance the workings of deep brain stimulation (DBS) (Kumar, Cardanobile, Rotter, & Aertsen, 2011). In a paper by (Dura-Bernal et al., 2017) that was aimed at parameter optimization, the authors successfully utilized genetic algorithms to fit the model closely to the data without searching the entire parameter space. We believe that a computational model could help the UNP project to gain greater understanding of the data. Additionally, because participants are so rare, a functional model could act as a sandbox to try new hypotheses before going to time-consuming patient trials. The primary goal of this paper is to build and test such a model.

Building a neurological model is a complex task where many design choices have to be made regarding the complexity of the model. The number of connections in a fully connected network rises exponentially with every node that is added. With every connection as well as every node comes a host of parameters that have to be carefully tuned if the model is to function properly. This tuning can be done by hand, but algorithms that do this automatically also exist. The secondary objective for this thesis is to examine the efficiency of automated parameter optimization in the context of (spiking) neural networks.

Relevance to AI

Artificial intelligence has always been related to neuroscience. Especially in the early days, the field of AI drew a lot of inspiration from how the brain is organized, processes information, and fulfills tasks. More recently, AI has been giving back to neuroscience in the form of data analysis and pattern recognition. Artificial (deep) neural networks have shown incredible potential in both industry and science (J. Lee, Davari, Singh, & Pandhare, 2018).

To build a computational model of the brain that produces output similar to the measured data, there are many parameters that have to be optimized. This optimization can be automated in many different ways. In this paper we used techniques from AI to search the parameter space more efficiently and compared them with common sense manual parameter search.

Additionally, some large research projects are aimed at fully constructing an artificial brain. To accurately model a small set of brain regions is a good first step. Building more of these models will certainly help identify their potential and pitfalls.

Computational modeling

In this section we will briefly discuss the current landscape in the world of modeling neural networks. In order to fully appreciate the insights of this paper it is useful to understand this context. There are multiple layers of abstraction at which models can operate, which can be viewed in Figure 1. Firstly, we can distinguish between cognitive models and computational models. Cognitive models operate at a high level of abstraction; they describe the relationship between brain regions at the level of 'what does it do?' and 'what happens where?' (Marr, 1982). The data for these models usually come from experiments which measure reaction time, score, accuracy etc. sometimes coupled with (knowledge from) fMRI studies.

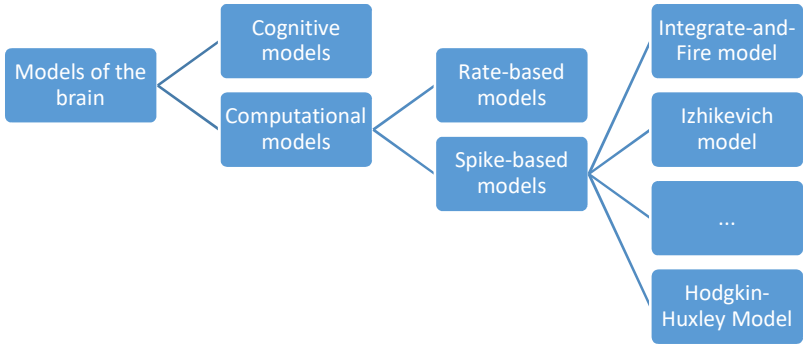

Figure 1 – Overview of brain model hierarchy. Nodes that are irrelevant to this paper (i.e. Cognitive models and Rate-based models) have not been expanded.

Computational models describe the relationship between brain regions in terms of neural activation. They attempt to be biologically plausible approximations of real neural networks, by capturing the behavior of a network or neuron in mathematical formulas.

Within the class of computational models, we can further distinguish between rate-based models and spike-based models. Rate-based models (sometimes called rate-encoded or activity-based) model the activity of an area as a whole; they do not model individual neurons. This is useful for networks with many nodes, as it is fast and scales well. The input data for this type of model is usually derived from EEG or ECoG experiments, which measure the magnetic field of (small) cortical areas, and fMRI studies which measure activation of a network over time.

Spike-based models on the other hand model each neuron separately. Every neuron has its own set of parameters and connections to other neurons, and can be individually measured or

altered. Usually, spike-based models are combined with of a couple of rate-based networks, which act as external input to the areas of interest. Spike-based models are usually based on data produced from single-cell recordings, and generate data that is similar to single-cell data.

However, the layers of abstraction do not end there. There are many different spike-based models with varying degrees of behavior, complexity and biological plausibility. These can be viewed in Figure 2, taken from (Eugene M Izhikevich, 2004). From single-cell recordings we know that some neurons spike regularly, others burst; some integrate and some resonate; some excite and some inhibit. Capturing such a wide range of behaviors in a set of mathematical equations is a complex task.

The simplest spiking model (integrate-and-fire) is based on a single equation and two variables. It scores low in terms of realistic behavior but high in terms of computational efficiency. The most complex one (Hodgkin-Huxley) gives the full range of neural behavior and is one of only two models that check the 'biophysically meaningful' marker. It accomplishes this by using multiple complex equations and roughly a dozen variables to model the flow of ions through pumps in the membrane. Of course, this comes at a steep computational cost.

Models	biophysically meaningful	tonic spiking	phasic spiking	tonic bursting	phasic bursting	mixed mode	spike frequency adaptation	class 1 excitable	class 2 excitable	spike latency	subthreshold oscillations	resonator	integrator	rebound spike	rebound burst	threshold variability	bistability	DAP	accommodation	inhibition-induced spiking	inhibition-induced bursting	chaos	# of FLOPS
integrate-and-fire	-	+	-	-	-	-	-	+	-	-	-	-	+	-	-	-	-	-	-	-	-	-	5
integrate-and-fire with adapt.	-	+	-	-	-	-	+	+	-	-	-	-	+	-	-	-	-	+	-	-	-	-	10
integrate-and-fire-or-burst	-	+	+		+	-	+	+	-	-	-	-	+	+	+	-	+	+	-	-	-		13
resonate-and-fire	-	+	+	-	-	-	-	+	+	-	+	+	+	+	-	-	+	+	+	-	-	+	10
quadratic integrate-and-fire	-	+	-	-	-	-	-	+	-	+	-	-	+	-	-	+	+	-	-	-	-	-	7
Izhikevich (2003)	-	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	13
FitzHugh-Nagumo	-	+	+	-		-	-	+	-	+	+	+	-	+	-	+	+	-	+	+	-	-	72
Hindmarsh-Rose	-	+	+	+			+	+	+	+	+	+	+	+	+	+	+	+	+	+		+	120
Morris-Lecar	+	+	+	-		-	-	+	+	+	+	+	+	+		+	+	-	+	+	-	-	600
Wilson	-	+	+	+			+	+	+	+	+	+	+	+	+	+		+	+				180
Hodgkin-Huxley	+	+	+	+			+	+	+	+	+	+	+	+	+	+	+	+	+	+		+	1200

Figure 2 – Overview of the different spiking network models, taken from Izhikevich (2004). FLOPS stands for 'floating point operations' (simple calculations) and gives a measure of the computational power needed to compute 1 ms of activity for a single neuron. The presence of "chaos" means that a small change in the initial settings produces a significantly different outcome. For information on the other properties, the reader should consult the original paper.

Choosing the right model is critical to success. When deciding which model to use, it is important to pick a model that adequately captures/produces the phenomena that you are interested in, while avoiding needlessly complex models as to keep computation time to a minimum (Izhikevich,

2004). In our case, physiological data at the level of ions and gates was not needed, and ensuring that a sufficiently sized network would run fast was a top priority, as finding the right parameters would require running the model many times. A fairly recent model that has been used by other researchers for similar purposes (Dura-Bernal et al., 2017) is the one by Izhikevich. This model was published in 2003 (Izhikevich, 2003) and has become increasingly popular in recent years. It covers a lot of complex observations while retaining high speed. This, combined with its popularity, made it the obvious choice.

Frequency vs. firing rates

Throughout this thesis we will regularly use the frequency and firing rates of the network to assess its performance. While these are very different concepts, they unfortunately tend to get mixed up in some of the literature. Therefor it seems appropriate to briefly highlight the difference between the two in order to avoid confusion for readers who might not be familiar with these terms.

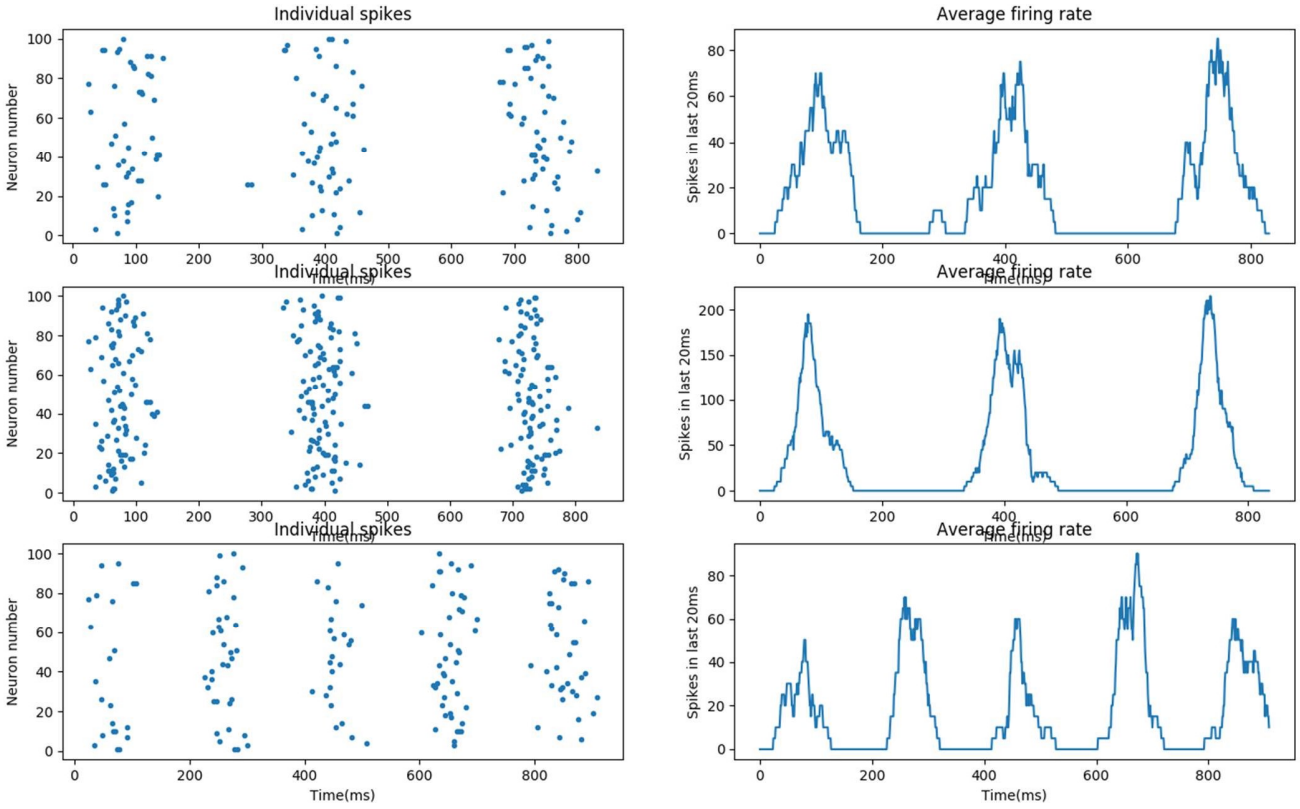

Figure 3 – Plots of a network with 100 neurons receiving input from an oscillator, simulated for one second. The first row serves as a baseline. In the second row, the strength of the oscillator input was increased. In the third row, the strength was reset to normal levels, but the frequency of the oscillator was increased.

The firing rate of a neuron is simply the amount of times that it fires per second. One instance of firing is called a spike, so the firing rate can be written as spikes/s. The average firing

rate of a network is then simply the total spikes per second, divided by the number of neurons. This value is often noted in papers which use single-cell recordings to measure neural activity.

Together, these individual spikes form one 'larger' signal. This signal is a more abstract representation of the activity of the network as a whole. When recording neural activity with large electrodes such a signal is recorded directly (as EEG or ECoG). With single-cell recordings (or spiking models), a way of transforming the spikes into a signal is to use a sliding window to calculate the average firing rate of the network over very short periods of time (see Figure 3). It should be noted that while these signals all have the same origin (i.e. neurons firing), their integration is complex, non-linear, and as a result these signals may vary significantly based on recording technique (Buzsaki & Draguhn, 2004).

All of these signals have wave-like properties. As is the case with any signal, it can be described as a summation of waves with different frequencies, phases and amplitudes. While any neural network signal will be a noisy combination of many different frequencies, they are often reported to have a main frequency. For example, a paper might report high beta activity in the thalamus, meaning that its main frequency is in the 13-30 Hz range.

Figure 3 illustrates the difference between spikes/s and frequency. The column on the left shows the individual spikes in an example network with 100 neurons which receive input from an oscillator. In the right column, these spikes have been converted into a signal by using a sliding window of 20 ms to calculate the average firing rate at every time step.

The networks in the first and second row have the same frequency, but the number of spikes is doubled in the second row (from 153 spikes/s to 298 spikes/s). Similarly, the first and third row have the same number of spikes, but the network frequency of the third row is 5 Hz rather than 3 Hz. Christie, Kamen, Greig Inglis, & Gabriel (2009) researched the relationship between spikes/s and signal frequency in the motor cortex; they found that he correlation between motor unit firing rates and the mean power frequency was small.

Measures of performance

In order to assess the validity of the model we need to consider what properties the model should have in order to be considered realistic or even plausible. The key properties we aim to achieve with this model (given the scope of this thesis) are listed here:

- *Firing rates that match with literature*. Each area of the network should produce the same average firing rate as described in single-cell recordings of those areas.

- *Wide distribution of individual firing rates*. We know from observations that the distribution of individual firing rates is highly skewed and long tailed (Roxin, Brune, Hansel, Mongillo, & van Vreeswijk, 2011). A realistic model should have this property.

- *Network frequencies.* The network signal can be analyzed by using an algorithm like FFT (Fast Fourier Transformation) to convert the signal into its pure wave components. In a realistic model there should be an inverse relation between the frequency and amplitude

of these components, meaning low frequencies should have the highest amplitude and vice versa. This is referred to as the 1/F law or power law (for example, see (Miller, Sorensen, Ojemann, & Den Nijs, 2009). This also means that the individual spikes of neurons should be *mildly* synchronized.

- *Robustness*. The model should be fairly robust; it should not rely on some magical configuration that produces perfect results, but which falls apart upon a minor change. The model should reflect that every network topology is slightly different, but functions more or less the same.

- *Responsiveness*. Ultimately the goal of this model is to simulate the changes in network frequency during intended movement as reported by Vansteensel, Pels, Aarnoutse et al. (2016). If the model is realistic, it should be possible to produce a change in the signal of the motor cortex (lower beta, higher gamma, see Figure 4) by effecting a change in the networks related to decision making (the basal ganglia).

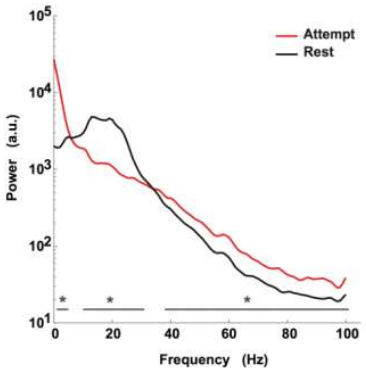

Figure 4 – Difference in activation between resting and attempted hand movement. Notice the 1/F shape of both graphs. Taken from Vansteensel, Pels, Aarnoutse et al. (2016).

Materials and methods

Model topology

The network consisted of 6000 spiking Izhikevich neurons representing some of the areas involved in the motor cortex. The topology of our model is pictured in Figure 5. Square nodes are external input to the model, round nodes are parts of the model itself. Green lines indicate an excitatory connection, red lines are inhibitory.

When a high level decision to move is made elsewhere in the frontal cortex, it is first sent to the basal ganglia, which is heavily involved in decision making (Bogacz & Gurney, 2007). The subthalamic nucleus (STN) is the first node of the basal ganglia and serves as the 'starting point' for the model. The second basal ganglia node is the Globus Pallidus externus (GPe), which forms a negative feedback loop with the STN. This loop plays a big role in creating oscillations: when the STN receives strong cortical input, the high activity is dampened by the GPe after a short delay, which reduces input to the GPe itself, which in turn allows the STN to start firing rapidly again. In patients with Parkinson's disease, this mechanism is disturbed (Tachibana, Iwamuro, Kita, Takada, & Nambu, 2011). The third component of the basal ganglia, the striatum, is considered to be the seat of the go/no-go decision making process (Schultz, Apicella,

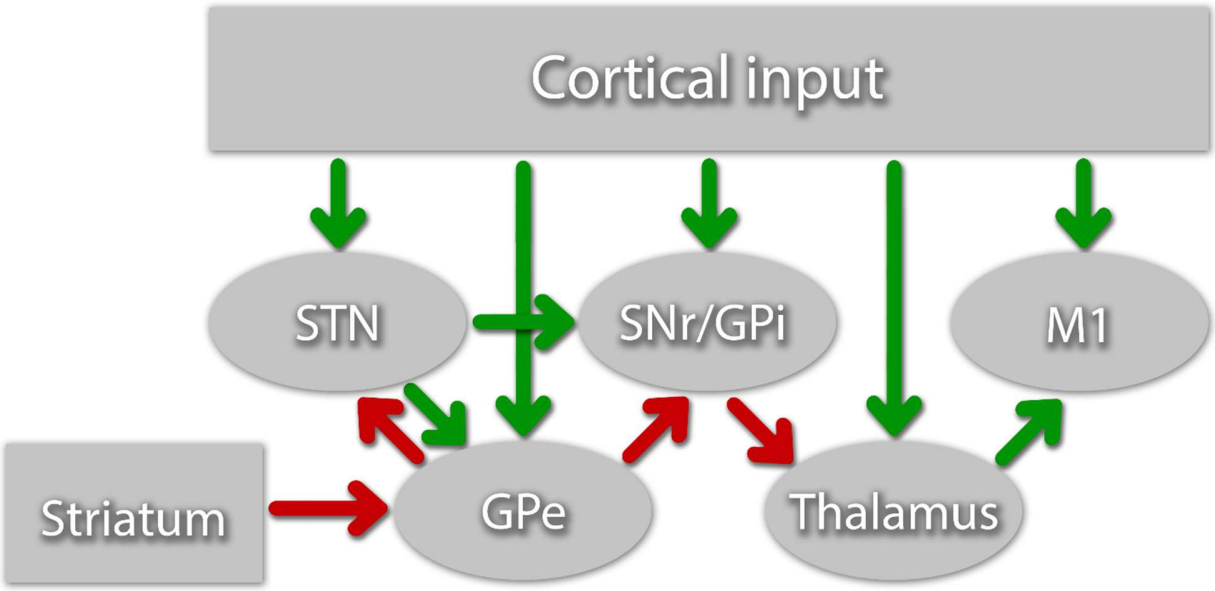

Figure 5 – Network topology of the model. Green and red arrows represent excitatory and inhibitory connections, respectively. Oval nodes are modelled as spiking networks, rectangular nodes are rate-encoded.

Scarnati, & Ljungberg, 1992). The STN and GPe both project to the next node, which is the Substantia Nigra pars reticula (SNr). This node is often grouped together with the GP internus (GPi) in literature, so we do not make the distinction in this thesis. This completes the basal ganglia. The SNr tonically inhibits the thalamus, which projects to the (primary) motor cortex (M1) in an excitatory fashion. In order to keep complexity within bounds, we did not consider the peripheral motor cortices (supplementary, premotor) nor did we split up the motor cortex into layers (such as 2, 3, 5a, 6).

Used software, libraries

The model was programmed in the programming language python using the NEST framework (Jordan, Jakob et al., 2019). All data analysis was also done in python.

Neuron models

All network nodes were modeled as either fully excitatory or fully inhibitory. The excitatory neurons (located in the STN, Thalamus and motor cortex) were modeled as 'Regular Spiking' Izhikevich neurons (Izhikevich, 2003), with distributed values for the *c* and *d* parameters to provide some heterogeneity between neurons, just as recommended in their paper. Similarly, the inhibitory neurons (located in the GPe and SNr) where modeled as 'Fast Spiking' neurons, with some heterogeneity on the *a* and *b* parameters, as described in the same paper.

Cortical input to each node (which both represents and is used to achieve baseline activity) was delivered with a NEST current generator, which fires Poisson-distributed spikes at the

neurons it is connected to. The spike train is different for each neuron in the network. The generators are active from the start of the simulation and operate at the same output level throughout the whole simulation.

Network connections

Following the model by Kumar et al. (2011) and expanding upon it, all nodes connect to themselves with a probability of 2% (meaning each neuron has a 2% chance of connecting to any other neuron) and a delay of 2ms. Connections between different nodes (as depicted by the arrows in Figure 5) have a connection probability of 5% and a delay of 5ms. These values are based on physiological measurements. The original Kumar et al. model features 1000 STN neurons and 2000 GPe neurons. This seems to be a common order of magnitude for spiking neural networks (for example: (Dura-Bernal et al., 2017)), although an explanation for the difference between STN and GPe size was not given. Since examining the effect of different network sizes is beyond the scope of this thesis, we decided to keep these dimensions and use 1000 neurons for each new node. This is sufficiently large for complex behavior and keeps computational power requirements manageable.

Parameter initialization and optimization

With these parameters set, two types of parameters remained that could be used to fit the model. By 'fitting' we mean choosing the parameter values in such a way that the network satisfies the conditions described in the 'measures of performance' section. Our approach was to start by satisfying the first condition (average firing rates comparable to literature) and then examine whether the model could also satisfy the other conditions.

The parameter types that were optimized are 1) the weight of the connections between and within network nodes, and 2) the amount of external input from the cortex to the nodes. The connection weights take a value between -4 and 4, with a negative value meaning that the neuron is inhibitory. This value represents voltage of a spike traveling through the axon as measured in mV, but since our model is an abstraction the values are better thought of as being dimensionless. The value for the external input represents the firing rate of the Poisson spike generator and usually takes a value between 1000 and 5000.

Parameter	Kumar et al (2011)	Our model
STN input	1500-3250	4000
Gpe Input	2000-3250	2500
w_STN-STN	1.3	2.2
w_STN-Gpe	1.3	2.2
w_Gpe-Gpe	-0.45	-0.2
w_Gpe-STN	-0.7	-1.1

Table 1 – Parameter values for the STN-GPe model.

The parameters for the STN and GPe were taken from Kumar et al. (2011). Since their model used the integrate-and-fire type of neuron and not the Izhikevich type, our parameters had to be slightly adjusted to produce the same output (see Table 1), which they described as an average baseline activity of 15 spikes/s for the STN and spikes/s for the GPe*. This replication of

the STN-GPe model served as the starting point for the rest of the experiment. The other nodes (SNr, thalamus and motor cortex) were added to the model and the whole model was optimized by an algorithm as described below. The target firing rates for these nodes were set to 62, 11 and 12 spikes/s respectively, based on a review of prominent literature (Hammond, Bergman, & Brown, 2007; Schnitzler & Gross, 2005).

There are five parameters for the external input (one for each node) and eleven connection weights, for a total of sixteen parameters that were used for optimization. As a result, the number of possible combinations of parameter values is huge. In computing science, this is called the parameter space: a multi-dimensional field where each point represents a set of parameter values. Some of these points will produce a model that closely matches the data – a good fit. This point is considered to be 'higher' than a point that produces bad results (see Figure 6). To get the best model, we have to search the parameter space for the highest points; but because the search space is so large, it is not feasible to try out all different combinations. This is a classic optimization problem for which many algorithms have been suggested (Zitzler, Deb, & Thiele, 2000). Most of these algorithms use an iterative process that uses the best points of one iteration as a starting point for the next iteration and discards the others. These 'survival of the fittest' approaches are called evolutionary algorithms (EA). To go into each one would go beyond the scope of this thesis, but it suffices to say that every EA balances speed, complexity and performance.

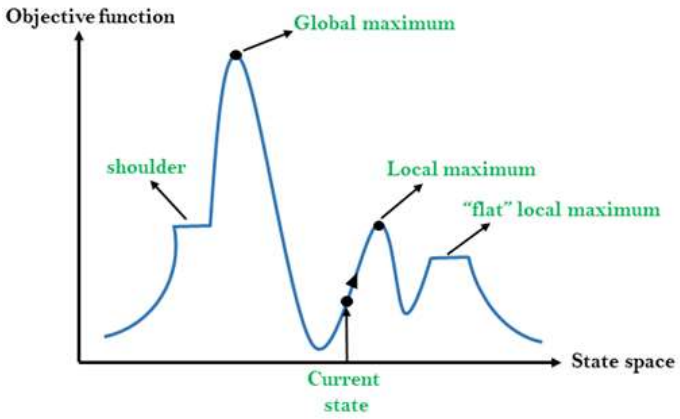

Figure 6 – Illustration of the parameter space (or state space). For each set of parameter values (X-coordinate) there is a goodness of fit (Y-coordinate). There is generally a relationship between these two factors. A hill climbing algorithm takes advantage of this fact when looking for the global maximum. Unfortunately, the algorithm can get stuck on sub-optimal solutions (local maxima). For multiple parameters, imagine this picture as being multi-dimensional.

We opted for using a variation of a relatively simple and fast EA, called the hill climbing algorithm (Tsamardinos, Brown, & Aliferis, 2006). This algorithm starts from one set of parameters, and searches the parameter space by looking at neighboring coordinates and moving to the best one, until it can no longer find a better one. Think of this as traversing the graph in Figure 6 by always moving uphill until a top is reached.

*Instead of spikes/s, the Kumar et al. used Hz. However, the data they reference is from single-cell recordings, and it is clear from the original source that the values are meant to describe firing rate, not frequency.

We defined the fitness function for the model as

$$\frac{1}{\sum_{n=1}^{n=5}\left((r_{model_n} - r_{data_n})^2\right)}$$

Where *r* is the average firing rate and *n* is the node number (STN, GPe etc.). This formula rewards parameters that provide good firing rates across all nodes.

The topology of the network allowed the optimization problem to be split into two sub-problems: since the thalamus and motor cortex have no influence on the STN, GPe and SNr, the latter could be optimized first, and the former could be optimized afterwards. This reduced the complexity from sixteen dimensions at once to ten and six dimensions separately.

A single 'run' of the model simulates 1200 time steps of 1ms. During the first 200 ms the model shows abnormal behavior due to initialization; this time frame is discarded. The resulting 1000 ms of neural activity is stored for analysis. Some preliminary testing with longer simulations did not indicate a difference in results.

In the first iteration, three values were tried for all external input: 1500, 3000 and 4500. Three values were also used for all connection weights: 0.5, 2.0 and 3.5. Running the model with all combinations of these values would require a total of 3^{10} or 59,049 runs, which would take roughly a week to compute. In order to save time, the problem was split up again: since the chance for a recurrent connection (from a node onto itself) is only 2%, the influence of their weights is limited. These weights were fixed to the middle value (at 2.0) while varying the others. We used the fitness function to calculate the fitness value for each model, and the model with the highest fitness was selected as winner. The weight of the recurrent connections was now varied and the best model selected as winner again. This reduced the number of runs to 2214 which were completed in a few hours.

In the second iteration, three values were tried for each parameter again, but this time relative to the current values. The values for external input were *current - 500, current,* and *current + 500.* For connection weights, they were *current - 0.5, current,* and *current + 0.5*. For example, the 'winner' model of the first iteration had STN input of 4500 and STN-STN weights of 2.0. In the second iteration the input was varied between 4000, 4500 and 5000, and the weights were varied between 1.5, 2.0 and 2.5. The model was run in all combinations just as in the first iteration, and the model with the highest fitness provided the starting point for the next iteration.

The third and fourth iteration continued this process by halving the steps. The external input values were varied by 250 and 125 respectively, and the connection weights by 0.25 and 0.125 respectively. After four iterations, the difference in model behavior were no longer noticeable, and the winning parameter values were coded into the model. This concluded the fitting of the ten parameters related to the STN, GPe and SNr. The process was then repeated for the six parameters related to the thalamus and motor cortex in exactly the same way. After the fourth iteration of this step, the parameters of the model with the highest fitness were coded in and the model had been fully fitted.

Fitting method

Using the firing rates described in the literature as the target for parameter optimization is relatively simple for two reasons. Firstly, it is an attribute of networks that most papers describing single-cell recordings list as baseline. The data on average firing rates is *relatively* easy to find. Secondly, average firing rates are easily calculated and compared to these baseline values.

Some papers (such as Kumar et al. 2011) use other concepts of measurement like the synchrony index or the oscillation index, however to use these effectively we need biological data that is not readily available. This is unfortunate, because we know from existing research that synchrony plays a crucial role in information transmission (Baker, Kilner, Pinches, & Lemon, 1999; Uhlhaas et al., 2009). As discussed earlier in this thesis, the average firing rates say very little about the network frequencies (Brunel & Hakim, 1999).

When training the model this problem quickly manifested itself in practice; that is to say, the model was quickly choosing parameters that had a good fit, but which seemed quite unrealistic upon further inspection (see *Results)*. Unsatisfied with this result, a second model was trained by hand. This was achieved by starting from the same basic STN-GPe model and adding the other networks one by one, taking advantage of the fact that there are no recurrent (backwards) connections besides the GPe-STN connection. This process was a repetition of running the model, looking at the results, and logically deducing which parameter needs to change. While repeatedly tweaking and running the simulation is a time consuming task, some efficiency is gained by not simulating thousands of obviously unfit models. Optimizing the parameters this way did not take as long as expected given the number of variables.

Physiological data for movement task

The physiological data used in this experiment was taken from an earlier study by (Vansteensel, Pels, Aarnoutse et al., 2016) and was collected from a single patient who suffered from ALS. This patient was a 58-old woman who had lost all motor control except eye function and slight mouth control, thus considered ‘locked-in’. She was able to communicate using an eye tracker, and informed consent was obtained in 2015 before the experiment. The patient was implanted with four strips of ECoG electrodes, two placed over the dorsolateral prefrontal cortex and two over the primary motor cortex. Each strip contained four electrodes. Neural activity was recorded from two of these electrodes (bipolar measurement) on the motor cortex while the participant performed a task. The position of these electrodes was chosen from four possible locations to maximize performance on the task (and other tasks that are not discussed here). The task used for this experiment involved a screen which prompted the user to rest for 15 seconds, then attempt to move their hand for 15 seconds, multiple times in a row. The data shows an increase in gamma-band (>30 Hz) activity and decrease in beta-band (13-30 Hz) activity during attempted hand movement.

Movement task

The intention of movement was modeled as described in Mirzaei et al. (2017). In this paper the authors describe an experiment where rats had to perform a movement task after waiting for a go-cue. This task has many similarities to the task our human participant performed, in which also involved waiting for and anticipating a go-cue before attempting movement. Their method for modeling attempted movement formed the basis for our experiment.

Mirzaei et al. found that 30% of STN neurons respond to the sensory go-cue. We modeled this by stimulating 30% of the STN neurons with an extra Poisson generator for a duration of 20 ms after the go-cue appears and an amplitude of 180 spikes/s, as described in their paper. Additionally, they found that neurons in both the STN and GPe were exhibiting 'ramping' behavior in anticipation of the go-cue, which lasted until movement onset. 34% of STN neurons showed a positive ramp, increasing their activity by 4 spikes/s and 43% of STN neurons showed a negative ramp, decreasing their activity by 1 spike/s. Both ramps were modeled with inhomogeneous Poisson generators, starting 500 ms before the go-cue at 0 spikes/s and reaching full activity at the time of the go-cue until the end of the experiment. Peak activity was 250 spikes/s for the positive ramp and 350 spikes/s for the negative ramp. Since the STN projects to the GPe, coding the ramps for the STN automatically had a ramping effect on the GPe.

In the final attempted movement-model the network is simulated as normal for 1000 ms to represent the resting state. At this point the ramping starts, reaching their full strength when the go-cue is given at t=1500. The ramps remain constant while a brief, 20 ms spike train in the STN represents the sensory response. The network is simulated for another 1500 ms in this attempted movement-state. By changing the model at the level of the STN and not directly interfering with the thalamus or motor cortex we hope to prove that the model operates as a whole.

This experiment was run ten times and the resulting frequency spectra were averaged.

Model output and analysis

The model produces timestamps at which each neuron fired (Figure 7, top). This spiking data was used to calculate the mean firing rate at each time frame using a 10 ms sliding window (Figure 7, bottom right). This output was then put through a Fourier transformation to calculate the power of the signal in across the frequency spectrum (Figure 8). The resulting signal was then compared with the physiological data. While this signal is different from ECoG data, it should correlate enough to compare the two (Buzsáki, Anastassiou, & Koch, 2016).

An alternative way to measure a signal from the network as a whole is to take the sum of the v-parameters of every Izhikevich neuron in the model. This v-parameter represents the charge on the neuron membrane, thus the sum of these membranes represents the electric field potential (EFP) (Vermaas, 2015). However, this way of calculating the EFP does not take spatial differences into account, which reduces its resemblance to real EFP's. Moreover, the relationship between the EFP and ECoG data is complex (Vermaas, 2015). Given these limitations, plus the observation that the EFP and firing rates correlate strongly in our model, the firing rates were deemed more suitable for analysis, and the EFP was not used further; it is plotted for reference only (Figure 7, bottom left).

Results

Figure 7 shows the raw output of the automatically trained final model. Shown are the individual spike patterns, as well as the electric field potential (EFP) which is calculated as the sum of the charge on all neuron membranes in the network. The average firing rate (with a sliding window of 10 ms) is also shown. The firing rates correlate strongly to the EFP. The 200 ms initialization phase (where the model shows abnormal behavior due to coding implementation) is not shown.

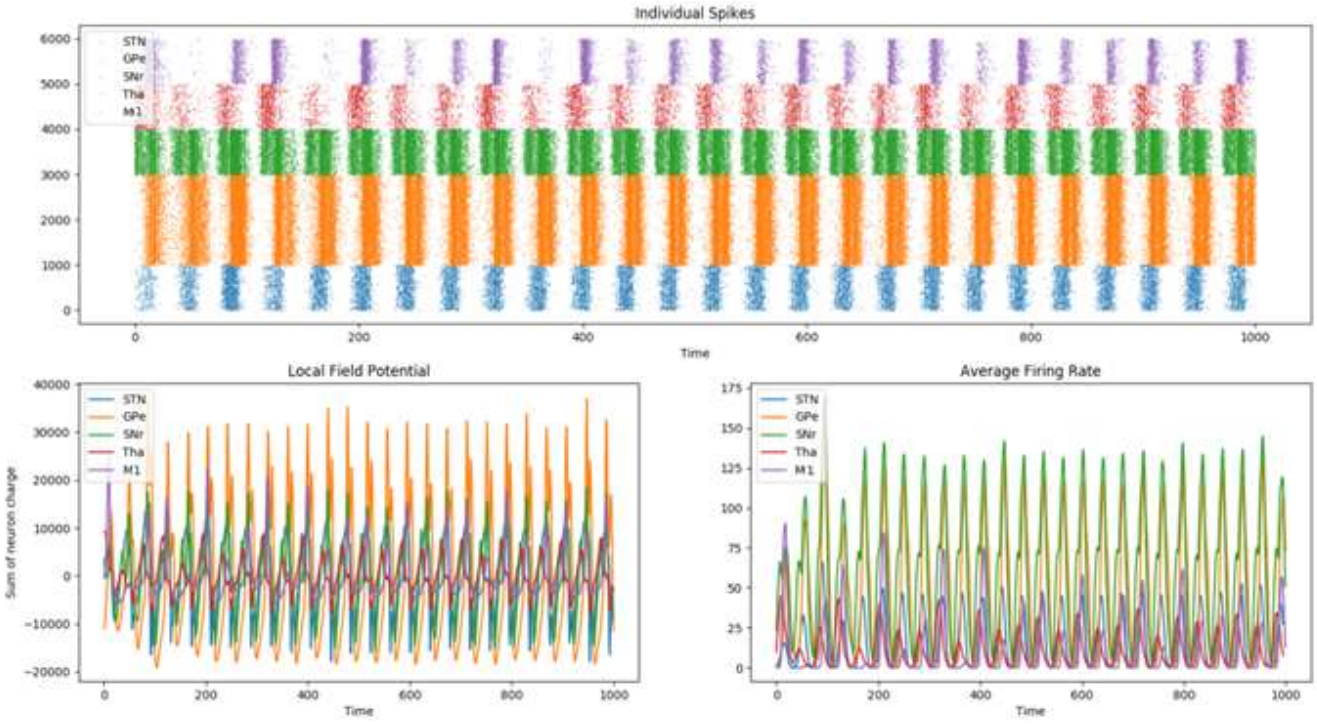

Figure 7 – Raw data output from the final model with automated parameter fitting. **Top:** spikes of individual neurons. Each dot represents one spike. Neurons are numbered on the Y axis. **Bottom left:** Local field potential. This is the sum of the charge on every neuron membrane. Each plot is normalized; the values on the Y-axis are dimensionless. **Bottom right:** average firing rate of networks, calculated over a 10ms sliding window.

Spiking behavior and firing rate

The firing rates of neurons in each area closely resemble the firing rates gathered from literature. Table 1 shows the mean firing rate and standard deviation for every network. The firing rates are not normally distributed, with individual neurons rarely firing less than one standard deviation below the mean, meaning no neurons were inactive.

Network	Mean	SD
STN	15,8	3,4
Gpe	46,2	4,0
SNr	62,2	4,7
Tha	11,4	2,2
M1	11,8	2,2

Table 1 – Mean firing rate and standard deviation of the mean, per network.

The topography of the model is resembled in the offset of spikes between networks. We can distinguish 3 phases: The first phase marks the activation of the STN (blue), SNr (green) and thalamus (red) by external cortical input. In the second phase, the GPe is activated by the STN, immediately blocking the STN with inhibitory signals. While the SNr is also inhibited by GPe, the excitatory spikes from the STN are still arriving and offsetting this inhibition. Inhibitory signals from the SNr are also arriving at the thalamus at this stage, blocking activity. The motor cortex, no longer inhibited by tonic spiking from the thalamus, starts to fire due to cortical input. Phase 3 is a 'resting' phase; without input, the GPe stops firing, allowing charge to slowly build up again in the other areas.

Robustness

The final model was run 50 times to assess robustness. The mean firing rates were very similar for every model, with a standard deviation of no more than 2% of the mean.

Power spectrum analysis

While the firing rates themselves are well fitted, inspection of the raw data as well as the power-frequency analysis (Figure 8) obviously shows that all networks fire in waves of the same frequency, at roughly 25 Hz. This frequency was not specifically coded into the model; it emerges from the interaction between the STN and the GPe, and could be influenced by changing the parameters of those networks (e.g. the delay value). Given the feed-forward connection scheme of the model it could be expected that any oscillations from these two networks would propagate to the other networks, hence the uniformity. The difference between the mean firing rate and network frequency is best explained by the level of neuron participation within each wave, as explained earlier in this thesis (Figure 3).

Comparing the power-frequency data from our model's motor cortex to the physiological data, there are some similarities and some differences. Figure 8 shows strong beta activity and some alpha activity (peaks at 9 and 15 Hz) which correlates to what we know from our participant. However, one expects these power-frequency graphs to roughly follow the common observation that lower frequencies tend to have higher amplitude, sometimes referred to as the 1/F law or power law (Stumpf & Porter, 2012). In this model, there is hardly any activation in the other frequencies.

Assessing model performance

While the model satisfies the conditions related to firing rate (which are accurate and widely distributed) and robustness, it is severely lacking in terms of frequency distribution. For this reason the model was not used in the attempted movement experiment. Instead, a new model was optimized by hand in order to fulfill the frequency distribution requirements.

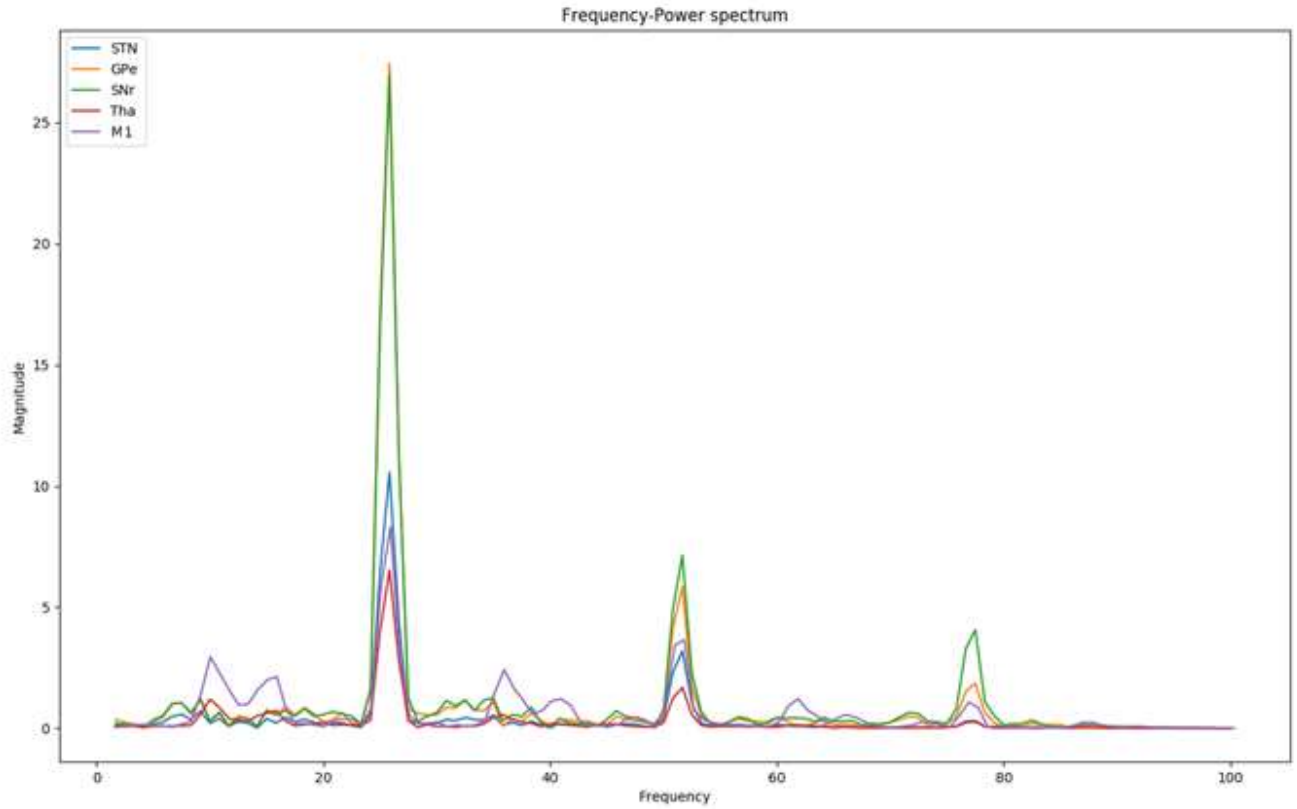

Figure 8 – This graph shows the magnitude of the activity in the network at every frequency, calculated by applying a Fast Fourier Transformation to the average firing rate-signal. Data is averaged across 10 trials.

Manually optimized model

Figure 9 shows the output of the manually fitted model. The values for the average firing rates per network are very close to those of the automatically fitted model. However, the behavior differs drastically.

Looking at the top image, neurons in the STN, GPe and SNr seem to fire almost asynchronously. However, the bottom right graph showing the average firing rate makes it clear that this is not the case; in fact, the average firing rate of all neurons now much closer resembles the frequency of their network as a whole. For instance, the thalamus in the automated model had a firing rate of 11,8 Hz while the network frequency was 25 Hz. In the manually fitted model, inspection of the power spectrum (Figure 10) reveals a network frequency of about 10 Hz, similar to the 10.83 Hz average firing rate.

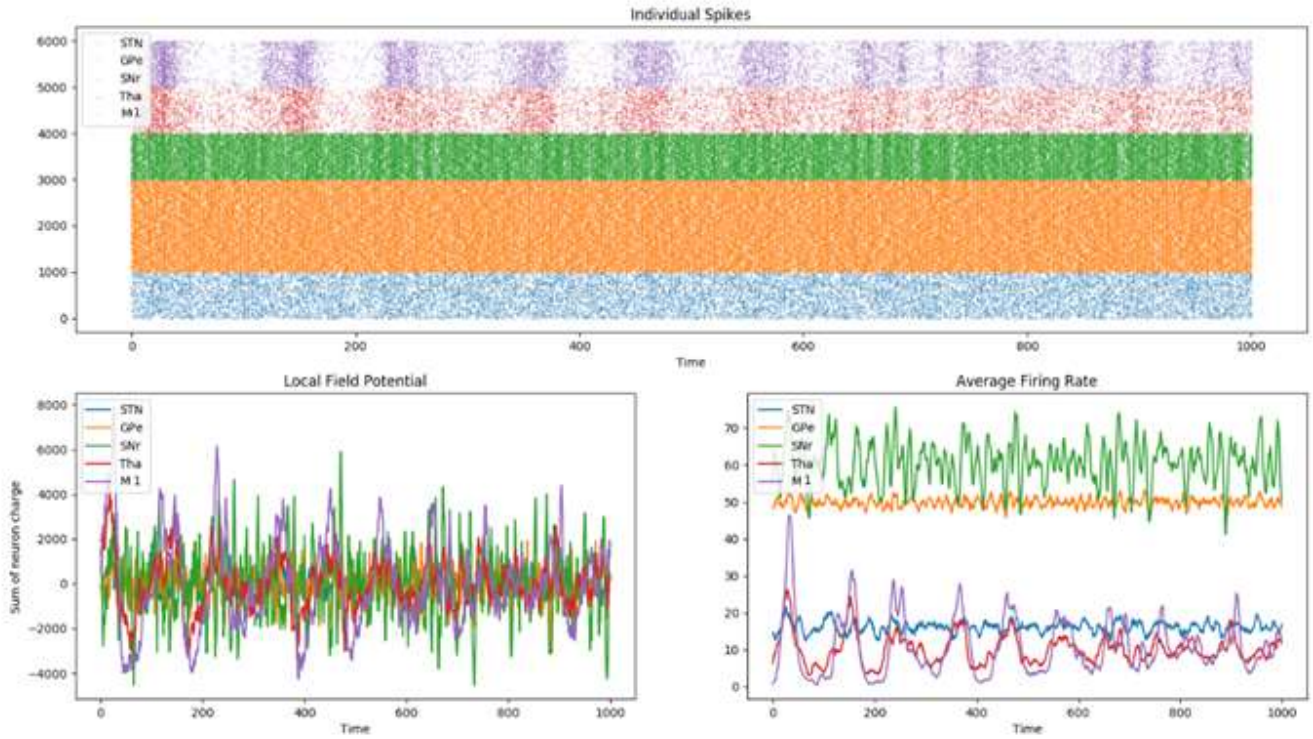

Figure 9 – Raw data output from the final model with manual parameter fitting. **Top**: spikes of individual neurons. Each dot represents one spike. Neurons are numbered on the Y axis. **Bottom** left: Local field potential. This is the sum of the charge on every neuron membrane. Each plot is normalized; the values on the Y-axis are dimensionless. **Bottom right**: average firing rate of networks, calculated over a 10 ms sliding window.

In terms of robustness, the manually fitted model produces slightly more random outcomes in terms of frequency, however the firing rate distributions fall within the same category as the automatically fitted model (a standard deviation of no more than 2% of the mean). For this reason, the spectrum analysis was done for ten instantiations of the model, and Figure 10 shows the average of the resulting ten graphs.

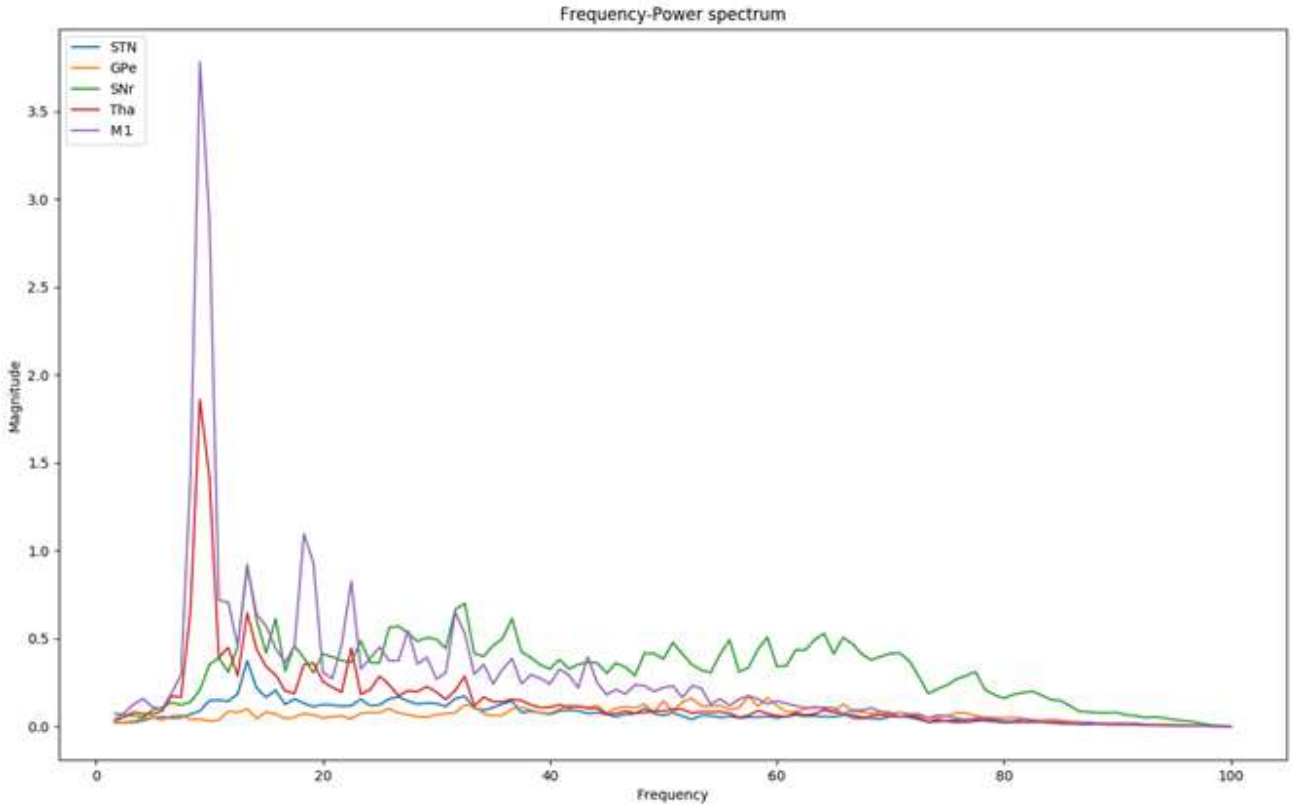

Figure 10 – This graph shows the magnitude of the activity in the network at every frequency, calculated by applying a Fast Fourier Transformation to the average firing rate-signal. Data is averaged across 10 trials.

The manually fitted model also shows improvements when looking at the power spectrum analysis for the motor cortex specifically. While the alpha activity is a bit higher than expected, there is strong beta activity, and generally the graph follows the 1/F rule much better.

In conclusion: The manually fitted model satisfies all the conditions that the automated model did, and shows a significant improvement with regards to frequency distribution. This model was used for the attempted movement experiment in order to assess whether it satisfies the last condition: responsiveness to stimuli with the intention to move.

Movement task

Figure 11 shows the frequency-power spectrum of the motor cortex during the attempted movement task as described in the methods section. The blue line displays the frequency spectrum before ramping activity (i.e. the first 1000ms), representing the resting phase. The orange line displays the frequency spectrum after the go-cue (i.e. the last 1500ms), representing the (attempted) movement phase. If the model is powerful enough to capture responsiveness to stimuli and intentional behavior, there should be a reduction in beta activity and an increase in gamma activity. Ideally this graph should look similar to Figure 4.

The graph shows a decline in alpha activity (at roughly 10 Hz), but no significant decline in beta activity. Where gamma activity was close to zero during the resting phase, this activity was present during the movement phase. Due to the lack of change in the beta frequency, it must be concluded that the model does not fully reproduce the neural phenomena observed during attempted movement, and therefore the model cannot be described as fully accurate.

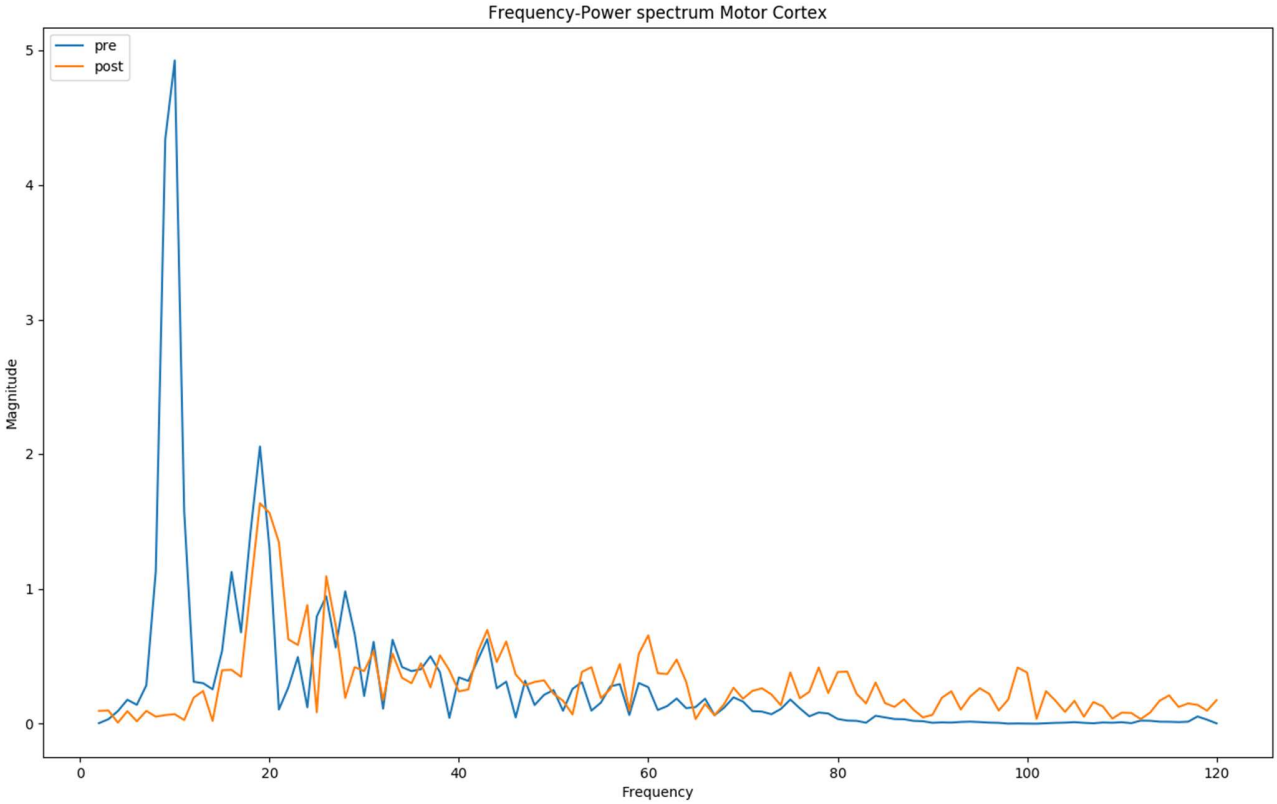

Figure 11 – This graph shows the magnitude of the activity in the network at every frequency, calculated by applying a Fast Fourier Transformation to the average firing rate-signal. Data is averaged across 10 trials. The blue line displays the activity during the resting phase (pre-go cue). The orange line displays the activity during the (attempted) movement phase (post-go cue). The ramping phase (at the end of the resting phase, in anticipation of movement) is not included in either condition.

Discussion

Analysis of model performance

Until now we have mainly used two tools to assess the performance of the model: the average firing rate of a network, and the FFT frequency analysis. As we have seen in the results section these don't necessarily correlate. However, the data on these phenomena are sparse and often separated. That is, one paper might measure the local field potential in a network and report that the network oscillates in the alpha band with a frequency of about 11 Hz. Another paper might do single-cell recordings and report that individual neurons in this same network had a base firing rate of 30 Hz. These observations do not exclude each other, but because they are made separately (by different researchers, under different circumstances, using different techniques and sometimes even in different animals) it is difficult to combine them into a holistic picture of what the neural behavior should precisely look like. However, we can say with certainty that the model produced by automated parameter fitting was insufficiently biologically plausible: such high levels of synchrony do not occur naturally (Donoghue, Sanes, Hatsopoulos, & Gaál, 1998, G Buzsáki, 2006).

The model produced by manual parameter fitting had a more plausible distribution of frequencies, but did not respond to changes in the STN as expected. Two explanations for this outcome exist: on the one hand, the resting phase could have been modeled accurately, in which case the way the model implemented intention of movement was incomplete. Simulating the ramping behavior in the STN and GPe might not be enough. For example, Tracey, Asanuma, Jones, & Porter (2017) describe the behavior of the thalamus during movement, which was not modeled explicitly in our experiment.

On the other hand, our implementation of intended movement might have been sufficient if the resting phase-model had had different properties or parameters. While these parameters produced satisfactory firing rates and frequencies, the exact values of these parameters may have caused the model to become too robust to change. It may also be the case that the model lacked some other features; for example, Economo et al. (2018) describe connections going back from the motor cortex to the basal ganglia and thalamus. These features were absent in our model.

Of course these explanations are not mutually exclusive; it is quite possible that both the resting state model and the implementation of movement need improvements in order to display the appropriate response to a go-cue.

Manual vs automated parameter fitting

Regarding the difference between the two models, one question pops to the front immediately: how did the hill climber algorithm 'miss' the settings that were found by manual tweaking? If the algorithm worked correctly, shouldn't it have found this setting too?

The problem probably lies within the STN-GPe interaction; this was the first set of parameters to be fitted, and once these were decided upon, all the other networks inherited the strong 25 Hz frequency. Only by tuning down the connection weights from these two networks to

the SNr all the way to close to zero could this be prevented from spilling over to the other networks, a setting that is not allowed to preserve biological plausibility.

Since there were six parameters to tweak for the STN-GPe pair (connection strength within and between the two networks, plus the amount of cortical input for each one) the parameter space was quite large. The hill climbing algorithm only selected the best starting region to zoom in on, and in this case it chose a parameter set that would reliably produce the wanted firing rates, but which also lead to very strong oscillations. In Figure 4 this is called a local maximum.

A possible solution might have been to hill-climb at multiple starting locations on the search space, however, this would add another order of magnitude to the time required to run the algorithm, which was already in the order of days. Waiting weeks for each phase of the algorithm to run was not feasible for this thesis.

Influence of recording approach

While we know from in-vivo single cell recordings that the strict, synchronous waves observed in our automatically fitted model are unrealistic (Kumar et al., 2011), the authors note that in vitro STN and GPe neurons have a tendency to form these oscillations. They hypothesize that other networks (most notably the cortex) provide such strong input that it disrupts these oscillations. The STN and GPe in our model only receive input from each other and some mild cortical input.

An important thing to reflect on is the fact that the output of the model is taken from the entire M1 node, while the electrodes that provide the real data only cover about 0.1% of the neurons in the motor cortex. More importantly, in the case of the locked-in participants this data is measured in a bipolar way; that is, rather than using a ‘true’ reference, we measure the difference between two electrodes on the same strip. While it is unclear how this influences measurement, it is clear that there is a difference with the model output.

Future research and recommendations

There are a number of interesting paths to pursue with this model. The most prominent one is to expand the topology of the network. Primarily there is only one recurrent connection in the model (the one from the GPe back to the STN), meaning there is no feedback loop from the other regions. Additionally, other ways to simulate intended movement should be explored.

With regards to parameter fitting, it is clear that the existence of local optima is a real problem, and that any algorithm which fits parameters to a fitness function must take this into account. Some efficiency may be gained by splitting the parameters into sub-groups, if the model topology allows it. This can reduce the search space by orders of magnitude. Additionally, the algorithm must not simply consider multiple points within one iteration, but keep exploring them for long enough to make sure that the (local) maximum has been found. Finally, using other metrics than just the average firing rate (such as frequency, synchrony) may improve results, but this requires more data.

Finally, an assessment should be of the tools used for analysis. We did not explore the effects of our methods of analysis on the results. For example, the running average firing rate graph uses a 10 ms sliding window; we did not research if other window sizes had an effect on the results. Similarly, we use the Fourier transformation to do the frequency-power analysis; using a wavelet transformation might produce slightly different graphs. We also did not incorporate any form of bipolar measurement in our experiment; it remains unknown whether this had a significant impact.

Conclusion

We have demonstrated that it is possible to build a computational model of the motor pathway to gain additional insight in the workings of BCI's for patients with locked-in syndrome. While the model is far from perfect, it can serve as a foundation to build upon. The lessons learnt here will improve any future attempts at computational modeling in all respects: coding architecture, model design, parameter fitting and data analysis. Our exploration of the hill climbing algorithm has identified some of the pitfalls when it comes to automated parameter fitting; with this knowledge, more sophisticated algorithms can be developed that avoid these pitfalls.

References

Baker, S. N., Kilner, J. M., Pinches, E. M., & Lemon, R. N. (1999). The role of synchrony and oscillations in the motor output. *Experimental Brain Research*, *128*(1–2), 109–117. https://doi.org/10.1007/s002210050825

Bogacz, R., & Gurney, K. (2007). The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Computation*, *19*(2), 442–477. https://doi.org/10.1162/neco.2007.19.2.442

Brunel, N., & Hakim, V. (1999). Fast Global Oscillations in Networks of Integrate-and-Fire Neurons with Low Firing Rates. *Neural Computation*, *11*(7), 1621–1671. https://doi.org/10.1162/089976699300016179

Buzsaki, G., & Draguhn, A. (2004). Neuroscience Neuronal Oscillations in Cortical Networks. *Science*, *304*(5679), 1926–1938. Retrieved from http://science.sciencemag.org/

Buzsáki, G. (2006). *Rhythms of the Brain*. *Rhythms of the Brain* (Vol. 1).

Buzsáki, György, Anastassiou, C. A., & Koch, C. (2016). The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes Electric current contributions from all active cellular processes within a volume of brain tissue superimpose at a given location in the extracellular medium and generate a potent. *Nature Reviews Neuroscience*, *13*(6), 407–420. https://doi.org/10.1038/nrn3241.The

Christie, A., Kamen, G., Greig Inglis, J., & Gabriel, D. A. (2009). Relationships between surface EMG variables and motor unit firing rates. *European Journal of Applied Physiology*, *107*(2), 177–185. https://doi.org/10.1007/s00421-009-1113-7

Donchin, E., Spencer, K. M., & Wijesinghe, R. (2000). The mental prosthesis: Assessing the speed of a P300-based brain- computer interface. *IEEE Transactions on Rehabilitation Engineering*, *8*(2), 174–179. https://doi.org/10.1109/86.847808

Donoghue, J. P., Sanes, J. N., Hatsopoulos, N. G., & Gaál, G. (1998). Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements. *Journal of Neurophysiology*, *79*(1), 159–173. https://doi.org/10.1152/jn.1998.79.1.159

Dura-Bernal, S., Neymotin, S. A., Kerr, C. C., Sivagnanam, S., Majumdar, A., Francis, J. T., & Lytton, W. W. (2017). Evolutionary algorithm optimization of biological learning parameters in a biomimetic neuroprosthesis. *IBM Journal of Research and Development*, *61*(2/3), 6:1-6:14. https://doi.org/10.1147/jrd.2017.2656758

Economo, M. N., Viswanathan, S., Tasic, B., Bas, E., Winnubst, J., Menon, V., … Svoboda, K. (2018). Distinct descending motor cortex pathways and their roles in movement. *Nature*, *563*(7729), 79–84. https://doi.org/10.1038/s41586-018-0642-9

Hammond, C., Bergman, H., & Brown, P. (2007). Pathological synchronization in Parkinson's disease: networks, models and treatments. *Trends in Neurosciences*, *30*(7), 357–364. https://doi.org/10.1016/j.tins.2007.05.004

Izhikevich, E M, & Izhikevich, E. M. (2003). Simple model of spiking neurons. *IEEE Transactions on Neural Networks / a Publication of the IEEE Neural Networks Council*, *14*(6), 1569–

1572. https://doi.org/10.1109/TNN.2003.820440

Izhikevich, Eugene M. (2004). Which model to use for cortical spiking neurons? *IEEE Transactions on Neural Networks / a Publication of the IEEE Neural Networks Council*, *15*(5), 1063–1070. https://doi.org/10.1109/TNN.2004.832719

Jordan, Jakob et al. (2019). NEST 2.18.0. Zenodo. 10.5281/zenodo.2605422

Kumar, A., Cardanobile, S., Rotter, S., & Aertsen, A. (2011). The role of inhibition in generating and controlling Parkinson's disease oscillations in the basal ganglia. *Frontiers in Systems Neuroscience*, *5*(OCTOBER 2011), 1–14. https://doi.org/10.3389/fnsys.2011.00086

Lee, J., Davari, H., Singh, J., & Pandhare, V. (2018). Industrial Artificial Intelligence for industry 4.0-based manufacturing systems. *Manufacturing Letters*, *18*, 20–23. https://doi.org/10.1016/j.mfglet.2018.09.002

Lee, J. H., Whittington, M. A., & Kopell, N. J. (2013). Top-Down Beta Rhythms Support Selective Attention via Interlaminar Interaction : A Model, *9*(8). https://doi.org/10.1371/journal.pcbi.1003164

Leuthardt, E. C., Schalk, G., Wolpaw, J. R., Ojemann, J. G., & Moran, D. W. (2004). A brain-computer interface using electrocorticographic signals in humans. *Journal of Neural Engineering*, *1*(2), 63–71. https://doi.org/10.1088/1741-2560/1/2/001

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. MIT press.

Miller, K. J., Sorensen, L. B., Ojemann, J. G., & Den Nijs, M. (2009). Power-law scaling in the brain surface electric potential. *PLoS Computational Biology*, *5*(12). https://doi.org/10.1371/journal.pcbi.1000609

Mirzaei, A., Kumar, A., Leventhal, D., Mallet, N., Aertsen, A., Berke, J., & Schmidt, R. (2017). Sensorimotor processing in the basal ganglia leads to transient beta oscillations during behavior. *Journal of Neuroscience*, *37*(46), 11220–11232. https://doi.org/10.1523/JNEUROSCI.1289-17.2017

Roxin, A., Brune, N., Hansel, D., Mongillo, G., & van Vreeswijk, C. (2011). On the distribution of firing rates in networks of cortical neurons. *Journal of Neuroscience*, *31*(45), 16217–16226. https://doi.org/10.1523/JNEUROSCI.1677-11.2011

Schalk, G., Mcfarland, D. J., Hinterberger, T., Birbaumer, N., Wolpaw, J. R., & Technology, A. B. I. B. C. I. (2004). BCI2000 : A General-Purpose Brain-Computer Interface (BCI) System, *51*(6), 1034–1043.

Schnitzler, A., & Gross, J. (2005). Normal and pathological oscillatory communication in the brain. *Nature Reviews. Neuroscience*, *6*(4), 285–296.

Schultz, W., Apicella, P., Scarnati, E., & Ljungberg, T. (1992). Neuronal activity in monkey ventral striatum related to the expectation of reward. *Journal of Neuroscience*, *12*(12), 4595–4610. https://doi.org/10.1523/jneurosci.12-12-04595.1992

Stumpf, M. P. H., & Porter, M. A. (2012). Critical truths about power laws. *Science*, *335*(6069),

665–666. https://doi.org/10.1126/science.1216142

Tachibana, Y., Iwamuro, H., Kita, H., Takada, M., & Nambu, A. (2011). Subthalamo-pallidal interactions underlying parkinsonian neuronal oscillations in the primate basal ganglia. *European Journal of Neuroscience*, *34*(9), 1470–1484. https://doi.org/10.1111/j.1460-9568.2011.07865.x

Tan, D. S., & Nijholt, A. (2013). *Brain-computer interfaces*. *Handbook of Clinical Neurology* (Vol. 110). https://doi.org/10.1016/B978-0-444-52901-5.00006-X

Tracey, D. J., Asanuma, C., Jones, E. G., & Porter, R. (2017). Thalamic relay to motor cortex: afferent pathways from brain stem, cerebellum, and spinal cord in monkeys. *Journal of Neurophysiology*, *44*(3), 532–554. https://doi.org/10.1152/jn.1980.44.3.532

Tsamardinos, I., Brown, L. E., & Aliferis, C. F. (2006). The max-min hill-climbing Bayesian network structure learning algorithm. *Machine Learning*, *65*(1), 31–78. https://doi.org/10.1007/s10994-006-6889-7

Uhlhaas, P. J., Pipa, G., Lima, B., Melloni, L., Neuenschwander, S., Nikolić, D., & Singer, W. (2009). Neural synchrony in cortical networks: History, concept and current status. *Frontiers in Integrative Neuroscience*, *3*(JUL), 1–19. https://doi.org/10.3389/neuro.07.017.2009

Vansteensel, M. J., Pels, E. G. M., Bleichner, M. G., Branco, M. P., Denison, T., Freudenburg, Z. V., … Ramsey, N. F. (2016). Fully implanted brain-computer interface in a locked-in patient with ALS. *New England Journal of Medicine*, *375*(21), 2060–2066. https://doi.org/10.1056/NEJMoa1608085

Vermaas, M. (2015). Relating ECoG and the local field potential to underlying mechanisms, *1875*.

Zitzler, E., Deb, K., & Thiele, L. (2000). Comparison of multiobjective evolutionary algorithms: empirical results. *Evolutionary Computation*, *8*(2), 173–195. https://doi.org/10.1162/106365600568202