Neuronal oscillations and brain wave dynamics in a LIF model

A Bachelor Thesis by Jimmy Mulder

Introduction

The Dutch television network is broadcasting a live brain surgery. I'm watching anxiously as neurologists drill a hole in the skull of the patient – who is fully conscious – and insert four electrodes into his brain. The patient is suffering from Parkinson's disease, and the treatment of his symptoms is called deep brain stimulation (DBS). The idea is to electrically stimulate the brain region responsible for the symptoms. When they start to put current on the electrodes, we can see how his trembling hand instantly relaxes. It's astounding that technology has come this far. But what strikes me the most, is what the neurologist in the studio tells us about the procedure: they have no idea how it works. One might expect that stimulating an already over-active region would only increase the symptoms. It turns out that the opposite is true: somehow, more activation of this area leads to less activation elsewhere. As Perlmutter & Mink (2006) put it: "we still have much to learn about the mechanism of action of DBS."

In the past decades, neuroscientists have made great advances in understanding the brain at a regional level. Functions of most regions have been mapped and we have an ever increasing understanding of how these regions are interconnected. The exact inner workings of the brain however – especially at the neuronal level – often remain a mystery. In this thesis, I will attempt to research a particular phenomenon which is visible at both high and low spatial scales, and which many neuroscientists believe to play a crucial role in the brain: oscillatory behavior (Basar & Basar-Eroglu, 1999; Jensen, Kaiser, & Lachaux, 2007; Lakatos, Chen, O'Connell, Mills, & Schroeder, 2007; Schnitzler & Gross, 2005). To this end, I have created my own neural model/simulator during an internship at the TU Eindhoven. The relevance of this topic to AI is discussed in appendix A.

My deep gratitude goes to Bart Mesman, without whom this thesis would not have seen the light. I would also like to thank Stefan van der Stigchel for guiding me through the paper writing process.

Introduction to Buschman et al.

In this thesis, I will regularly refer to a study done by Buschman et al. Because their paper acts as the backbone for this thesis, I will cite part of their abstract so that the reader may get notion of the phenomena I attempted to reproduce in this thesis.

"We found evidence that oscillatory synchronization of local field potentials (LFPs) formed neural ensembles representing the rules […] This suggests that beta-frequency synchrony selects the relevant rule ensemble, while alpha- frequency synchrony deselects a stronger, but currently irrelevant, ensemble. Synchrony may act to dynamically shape task-relevant neural ensembles out of larger, overlapping circuits." (Buschman, Denovellis, Diogo, Bullock, & Miller, 2012)

This experiment hints at the important role that oscillations and synchrony may play in the brain: to configure the networks and pre-activate the relevant ensembles while putting others to sleep. In this thesis I hope to demonstrate some of these effects in a simple Leaky Integrate and Fire (LIF) model. Can the simple rules of a LIF model lead to oscillatory or other complex behavior?

Contents

Methods

When researching time-sensitive phenomena like oscillations, one cannot use traditional rate based neural networks since these are insensitive to timing and thus can never give rise to synchronized behavior. Instead, a pulse-based model is needed. There are many pulse-based models out there, but most of them rely on complicated mathematical formulas (E. M. Izhikevich, 2004), some of which even try to describe the exact state of the ion channels in the neuron. Additionally, using an existing model severely limits ones options to those that the program offers. Given the explorative nature of this research, it required as much flexibility as possible, so it uses a self-made simulator.

To keep the complexity of this model at an understandable level, it uses the simplest neurological model available: Leaky Integrate and Fire. Generally, each neuron gets input from other neurons, which influences its electrical charge. If the charge reaches a certain threshold, the neuron fires, which means it stimulates the neurons it has a connection to. If the threshold is not reached, the charge will leak away until the neuron is back to its resting potential.
The model has many settings, but the most important are:
- Threshold – the amount of charge needed to elicit firing
- Leak – the rate at which the charge leaks away until the neuron has reached its resting potential
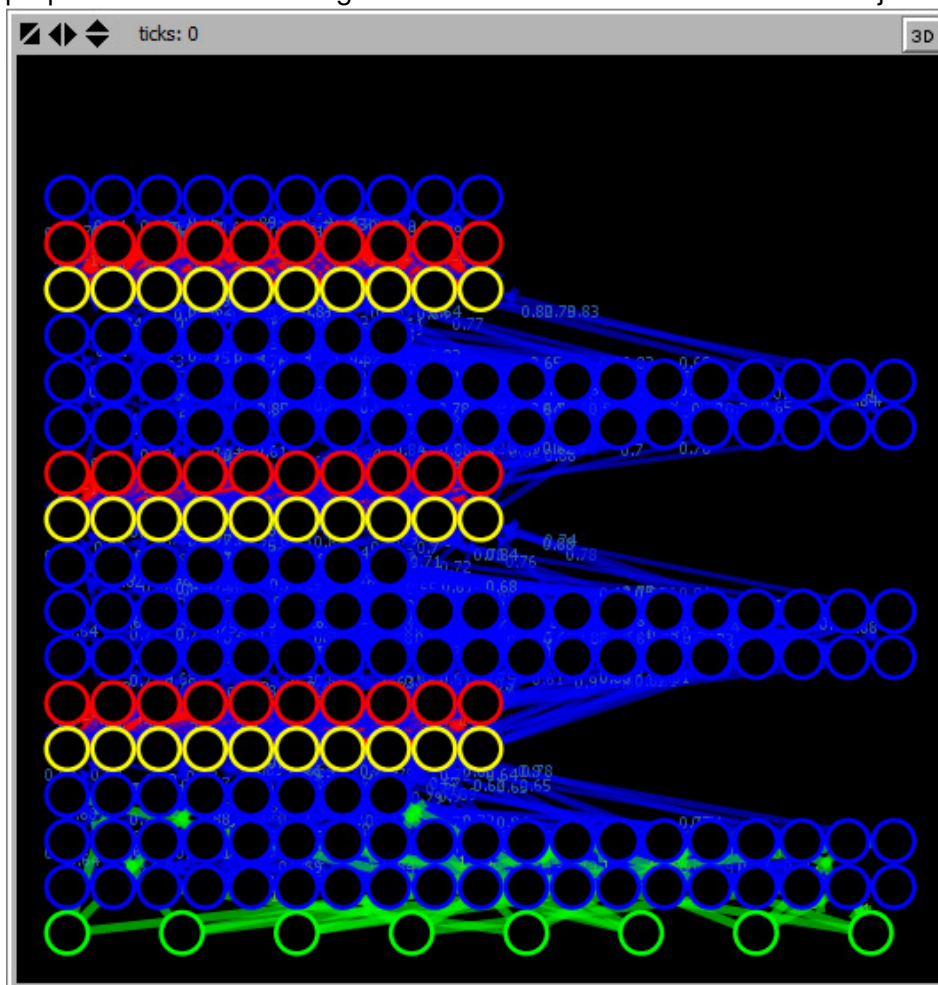- Hyperpolarization – the charge that a neuron will start with after it has just fired

Figure 1 – The green neurons are input neurons. The blue neurons are randomly connected excitatory neurons that project to their own layer and the yellow neurons directly above them. The yellow (excitatory) and red (inhibitory) neurons form a Winner Takes All layer. The blue neurons at the top are the output neurons.

A total of three different models was used to research different phenomena. The first model (discussed below) was made to research oscillatory behavior at a purely neuronal level. The second model (described later) is more abstract and serves to research the wave dynamics between ensembles. Finally, the third model tries to combine these phenomena into a working practical application.

Model of individual neurons
The goal of this model is to help us understand how, at a neuronal level, oscillatory behavior might arise from sensory input. In other words, how might structured signals be derived from chaotic input?

The topology of the model is illustrated in figure 1, and is loosely based on that in the visual cortex. Every neuron is either excitatory or inhibitory. At the bottom is a layer of input neurons: these are excitatory only, and manually configurable to fire at a certain frequency or noise level. From there, the signal travels mostly upwards. Most of the measurements were done at the output neurons, which lie at the top.

Between the input and output layers are a number of stages, with each stage consisting of three layers. The first layer consists of only excitatory neurons which are randomly connected to some of the neurons in their own layer and the layer above them. The second and third layer of a stage are a combination of excitatory (yellow) and inhibitory (red) neurons, forming a Winner Takes All system (Oster, Douglas, & Liu, 2009). Figure 2 is a schematic of how this WTA principle works. Basically, every yellow neuron is coupled with the red neuron above it, which in turn inhibits the yellow neurons surrounding its 'child'. This ensures that only the most powerful stimuli are passed onto the next layer and stage.

An important feature of the model is the capability of Hebbian learning: if turned on, each synapse (connection between neurons) will deteriorate slowly, but gain strength when the neurons at both ends fire (nearly) simultaneously. This approach is based on Hebb's law (Hebb, 1949), which is often summarized as "cells that fire together, wire together".
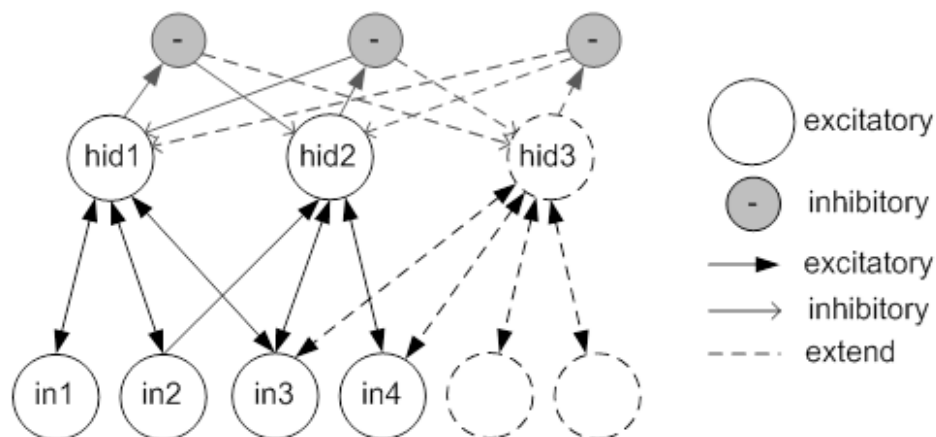

Figure 2 – The Winner Takes All principle is drawn here. Each hid neuron (the yellow neurons in the model) is coupled with an inhibitory neuron, which fires at the "opponents" of the hid neuron, thus making sure only the most salient stimulus is passed on.

Findings

The results of this research are divided into six chapters. The first chapter discusses the theoretical abilities and limitations of any LIF model. Chapters two and three use the model described earlier to explore dynamics at the neuronal level. Chapters four and five use a more abstract model to research wave dynamics, and also try to make the first link with the experiment done by Buschman et al. Finally, in the last chapter a third model is introduced to try and give a practical example of how the results from the previous chapters might be integrated into a working model of wave-based rule-selection.

Can a LIF model lead to oscillatory behavior?

As mentioned before, there is a wide variety of neurological models (E. M. Izhikevich, 2004), of which LIF is the simplest. The question now rises whether LIF is complex enough to exhibit oscillatory behavior. According to Izhikevich, LIF models are "class 1 excitable":

"The frequency of tonic spiking of neocortical RS excitatory neurons depends on the strength of the input, and it may span the range from 2 Hz to 200 Hz, or even greater. The ability to fire low-frequency spikes when the input is weak (but superthreshold) is called Class 1 excitability [8], [17], [22]. Class 1 excitable neurons can encode the strength of the input into their firing rate, as in Fig. 1(g)."
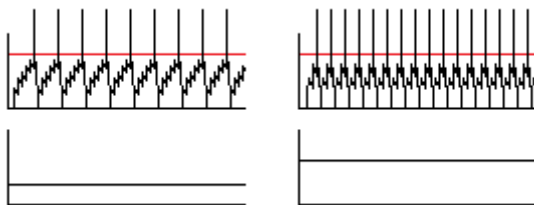

Figure 3 – relation between connection strength (bottom) and firing rate (top) when given constant input. Any constant input results in oscillatory behavior. Additionally, when the connection is weak (left), this results in a lower frequency than when the connection is strong (right).

This means that a LIF neuron will fire more often when this input is strong than when it is weak, and given a constant input this will lead to oscillatory behavior (figure 3). However, it doesn't say much about non-constant input. Suppose for instance that neuron X has four input neurons, each with a strength of 0.25 of the threshold and firing once every 4 ms. Because charge leaks away from the membrane if the threshold is not reached, neuron X will fire much more often if his input neurons fire simultaneously (instantly reaching the threshold) rather than sequentially, each 1 ms apart (allowing for charge to leak away). In this case, the total amount of input is the same; it is the level of synchronization that increases the response of X.

In conclusion: while the LIF model is the simplest neurological model available, it has properties that allow for complex, non-linear behavior. The integrating and threshold properties together allow for neurons oscillating at a constant frequency when given a constant input. The leaking property ensures that neurons are more responsive to synchronized input.

How does a LIF model react to different frequencies as input?

One might expect that if we feed a network input with a certain frequency, the output of this network will have the same frequency. To test this hypothesis, we subjected the model described in the methods section to input with different frequencies and measured the frequencies of the output neurons.

In this experiment, the model has 8 input neurons, which fire periodically according to an input string. For instance, given a string "00005555" the four input neurons on the left will not fire at all, and the four input neurons on the right will fire once every 5 updates (called *ticks*) in the model. The strings "00005555", "55550000", "00009999", "99990000" were used as stimuli. Switching sides on the input (e.g. "00005555" vs "55550000") allows us to check whether the network only responds to the total amount of charge (i.e. it's indifferent to where this charge comes from) or that it also responds to the origin of the input.
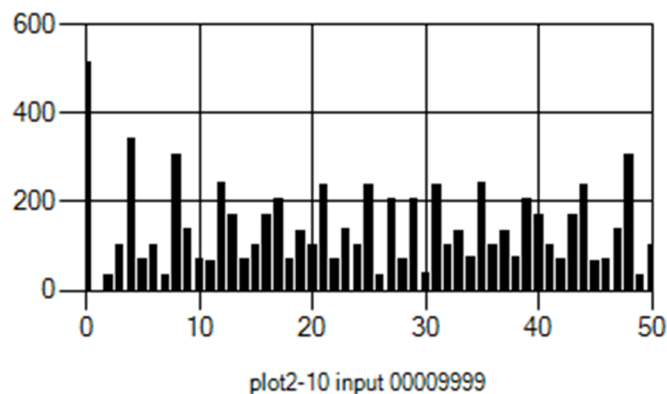

plot2-10 input 00009999

Figure 4 – an example of an autocorrelation plot. Lag is plotted on the X axis, the corresponding value on the Y axis. There are peaks at multiples of 4 on the X axis. The average firing rate is also 4, because it fired 500 times in 2000 ticks. There are small values between the peaks. This combined data suggests that this particular neuron fires roughly every four ticks.

Before the experiment, the model was trained briefly on each of the four stimuli, allowing the weights of the neurons to adapt themselves. Thereafter learning was switched off, so the weights in the network were always the same for every trial.

After the experiment, we calculated the autocorrelation of the firing pattern of each individual output neuron. For one neuron, this results in a graph like figure 4. The author advises readers unfamiliar with autocorrelations to read appendix B before continuing.

It's important to note that the autocorrelation at Lag = 0 gives exactly the amount of times the neuron has fired in the 2000 ticks of this particular run. This allows us to easily calculate the average firing rate, which is especially useful in cases where the average firing rate doesn't match the periodicity suggested by the autocorrelation.

The results of one experiment are described below. This experiment is representative of all the experiments done in this setup, although the exact outcomes differ due to the high randomness of the model. Output neurons that did not fire during the entire experiment were removed from the figures.
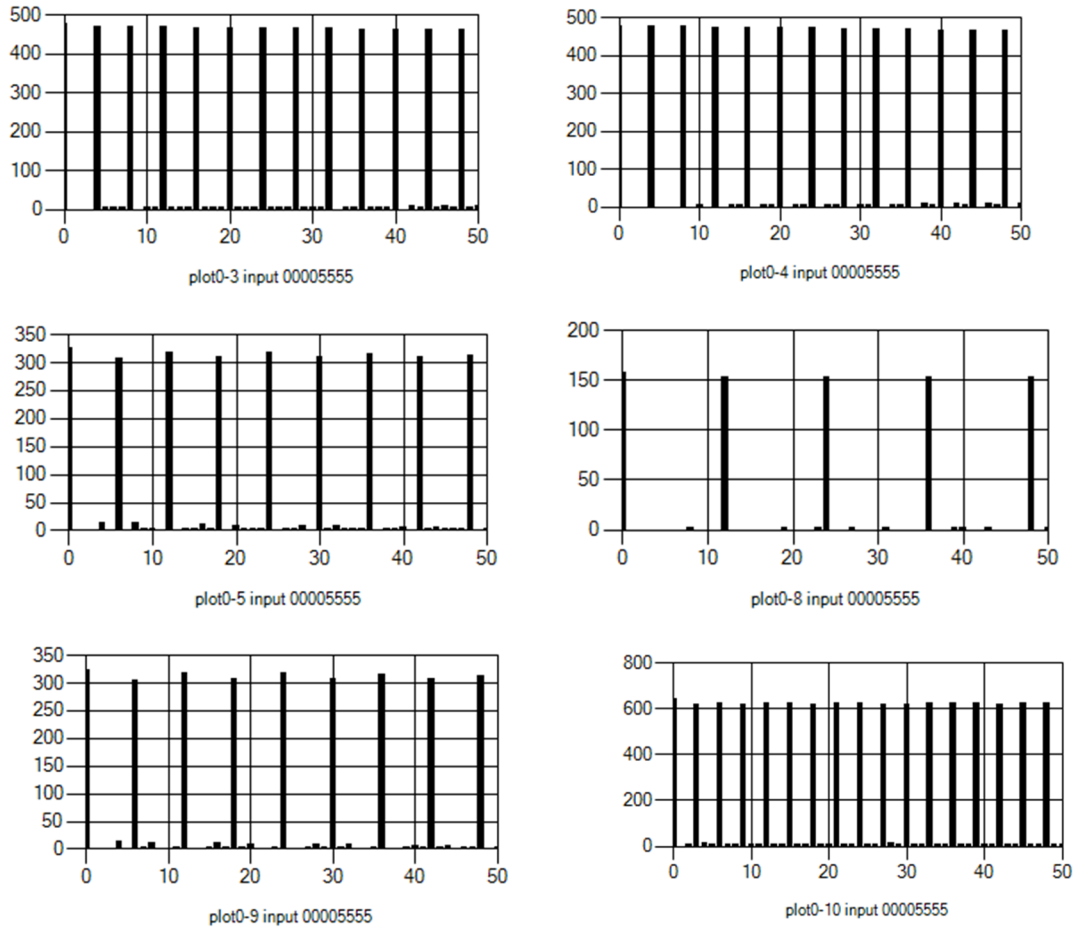
Figure 5

Figure 5 plots the results of the output neurons to input string 00005555. The autocorrelations show that all signals are very consistent, as there is barely any noise. This suggests that the neurons have a very strong periodicity, although their individual frequency differs a lot. Also, none of the frequencies is directly related to the periodicity of the stimulus, which is 5 while the periodicity of the output neurons is 4, 4, 6, 12, 6 and 3, respectively. Note that the average firing rate is identical to the frequency suggested by the autocorrelation.

In short, when the four input neurons on the right fire with a periodicity of 5, the output neurons fire with a very clear frequency that is not directly reducible to the input frequency, and their individual frequencies vary.
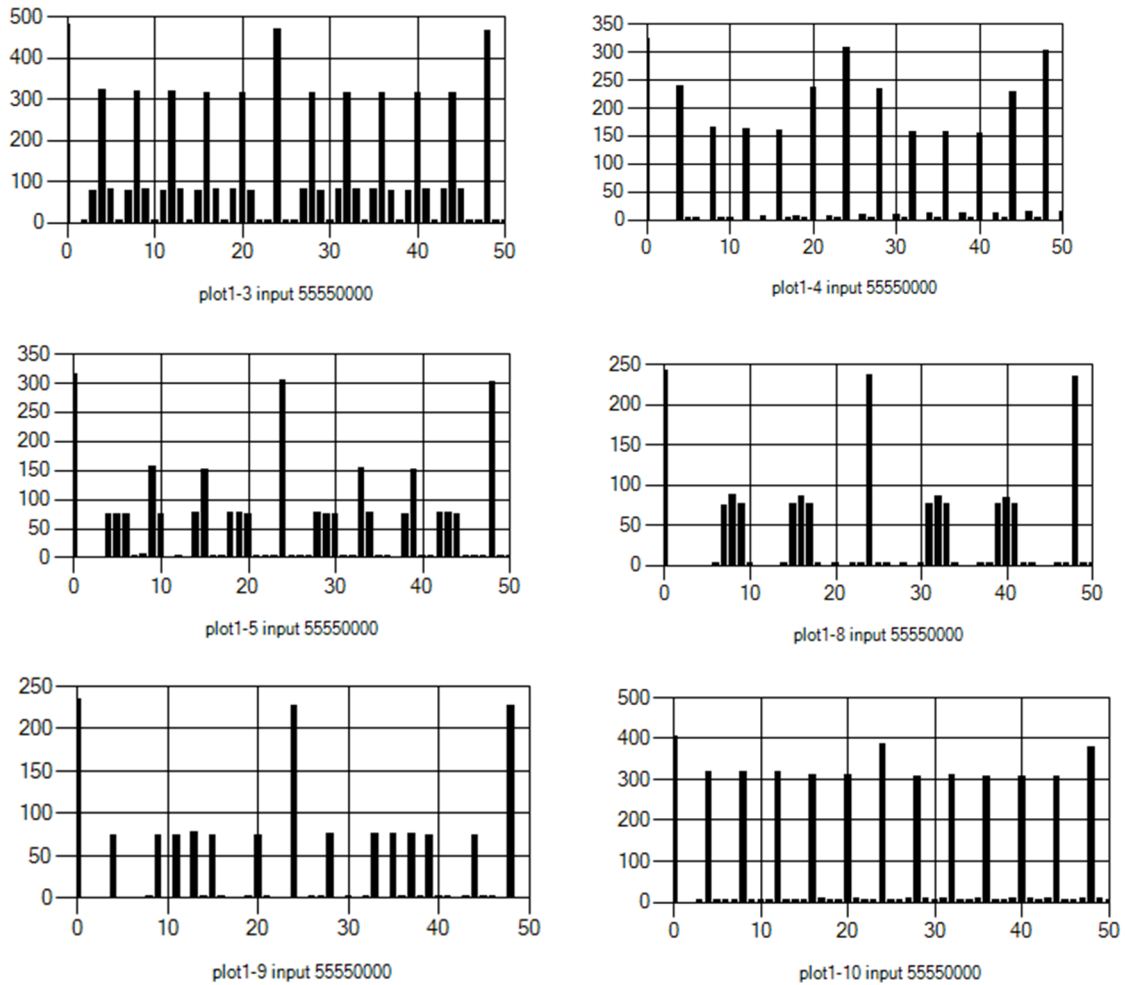
9

Figure 6

Figure 6 plots the results of the output neurons to the reversed input string. Some neurons still show roughly the same behavior. The periodicity of neurons 3 and 4 is still 4, although their firing pattern has slightly changed. Also, the activity of neuron 4 has declined from 480 to 320, which means the perceived periodicity is no longer the same as the average firing rate.

The other neurons have drastically altered their firing pattern. Based on the peaks and the average firing rate, neurons 5 and 8 seem to fire 'roughly' every 6 or 8 ticks respectively, yet this pattern always converges to one periodicity, 24 in this case. This means that a pattern of multiple activations repeats itself every 24 ticks.

The firing pattern of neuron 9 also repeats every 24 ticks, but the pattern seems even more complex than that of 5 and 8. Finally, neuron 10 has switched from a periodicity of 3 to one of 4.

The difference between figure 4 and 5 makes it clear that the origin of the input matters. Even though both the frequency and the amount of charge fed to the network remained the same, they output neurons drastically changed behavior because the origin of the input (e.g. which neurons were active) changed.
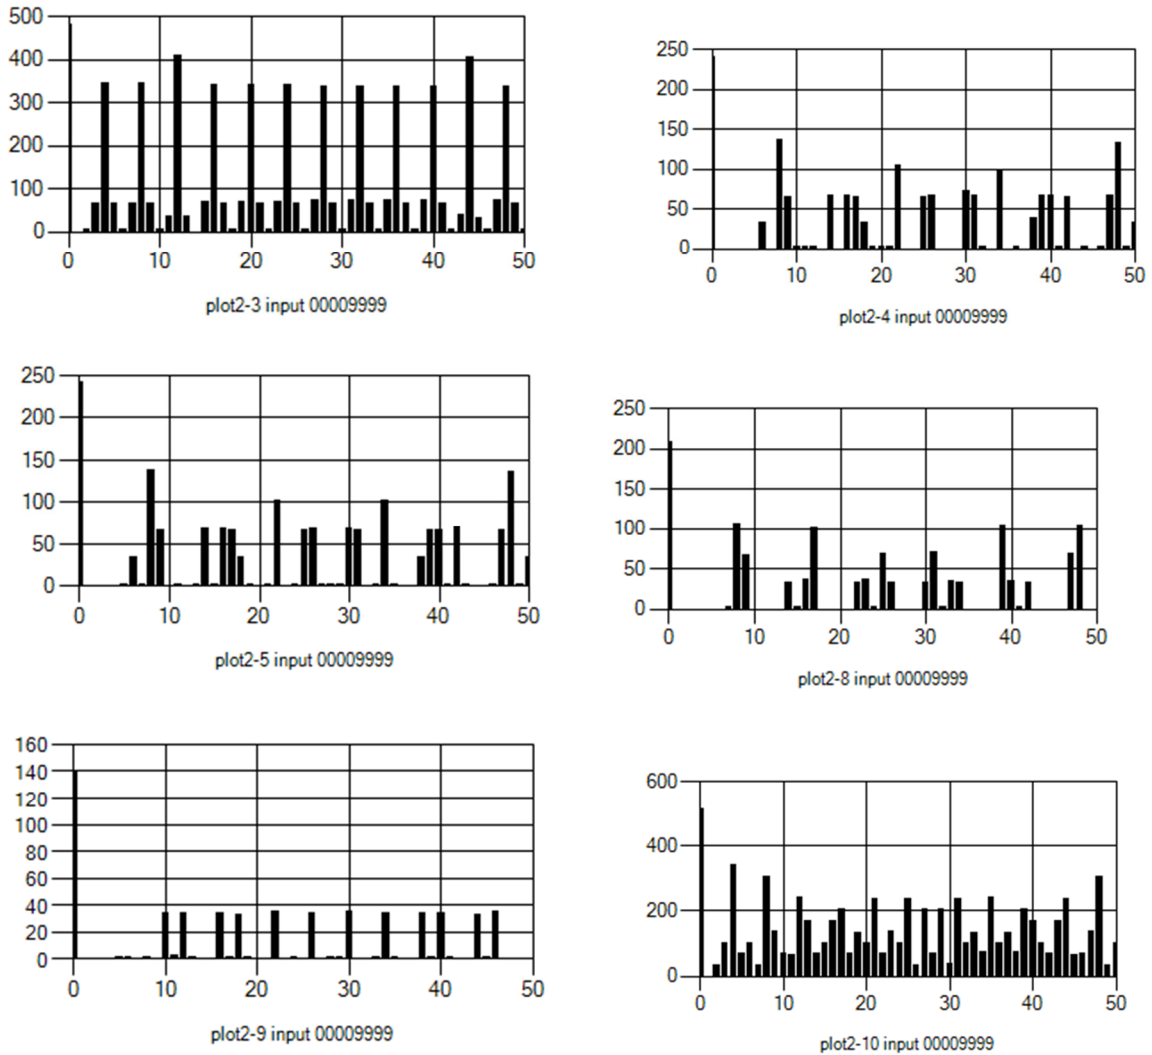
10

Figure 7

The input string in figure 7 is 00009999. Note that a lower firing rate with the same synapse strength means that the network now receives less charge overall.

Some of the output neurons now seem to be firing randomly, except for neuron 3 which has only slightly changed its behavior since the previous stimulus. However, if we extend the autocorrelation to lag 100 (figure 8) we can see this is not the case. The firing pattern of most output neurons is now 56 ticks long. One might say that every neuron is now coding the stimulus in a unique bit-string of 56 0's and 1's.
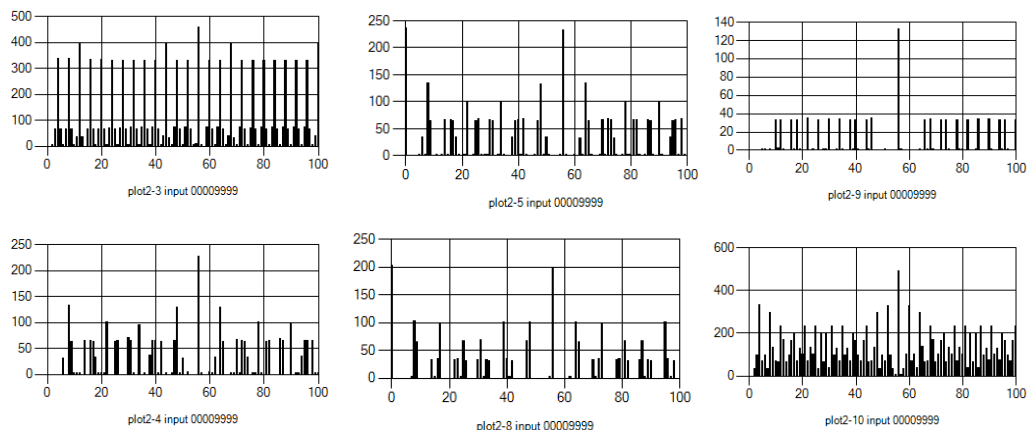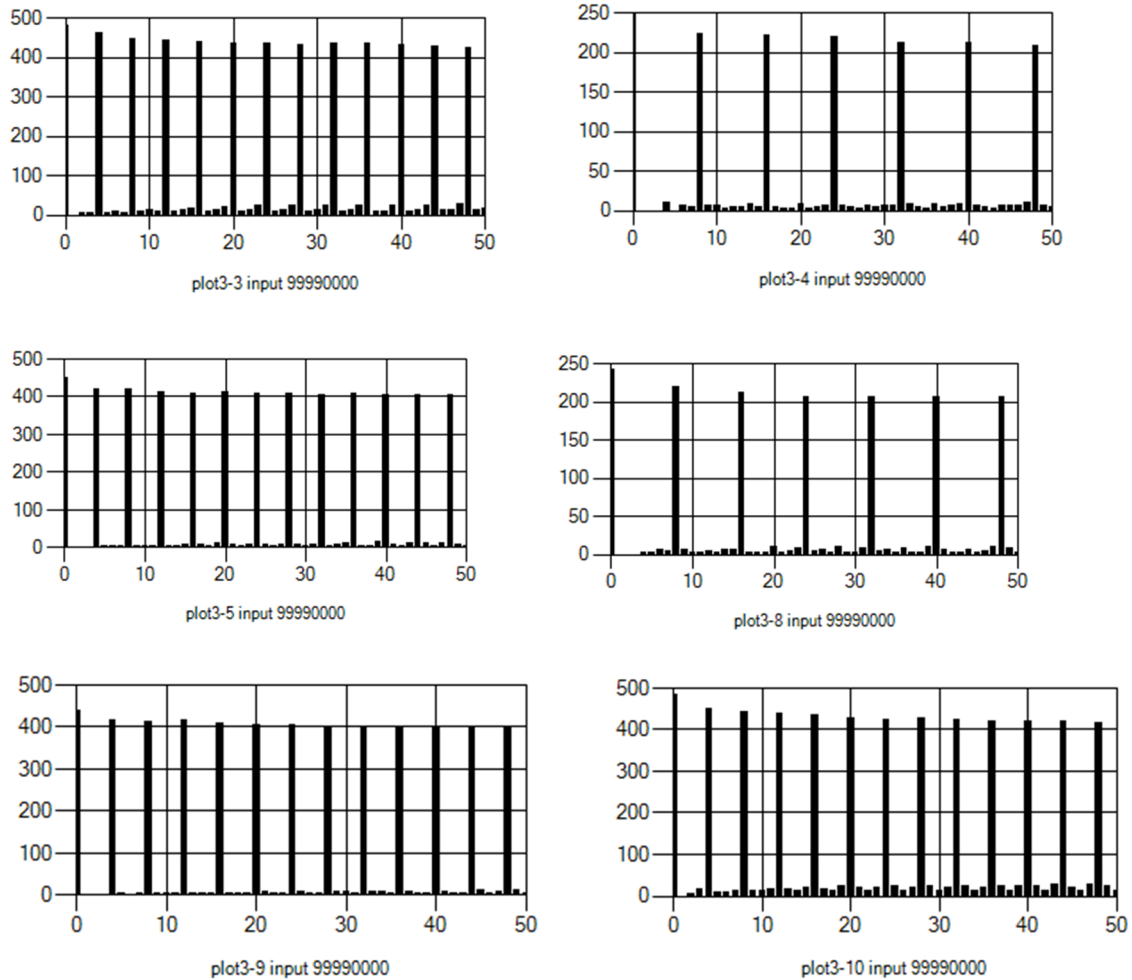

Figure 8

plot3-3 input 99990000

plot3-4 input 99990000

plot3-5 input 99990000

plot3-8 input 99990000

plot3-9 input 99990000

plot3-10 input 99990000

Figure 9

Finally, figure 9 plots the results of the opposite input string. The strong periodicities are back and the behavior looks like that of stimulus 00005555, although there are significant changes in frequency. The respective periodicities of 00005555 were 4, 4, 6, 12, 6 and 3, while in this case, they are 4, 8, 4, 8, 4, and 4.

In conclusion, the LIF model can show complex behavior when given periodic input. Although we expected the frequency of the output to be reducible to the frequency of the input, this was not the case. In all the experiments we ran, it only rarely happened that the frequency of an output was equal to (a multiple of) the frequency of the input.

Not only do output neurons change their rate of fire and periodicity, in some cases they show a complex firing pattern that takes many ticks to complete (sometimes as many as 74 ticks) and from an engineering perspective this could be seen as bit-encoding.

An important thing to note is that without the short training at the start of the trial, the results were mostly random noise. In the next chapter this will be discussed further.

How does a LIF model react to stimuli without frequency (white noise)?

One might expect that periodic input produces periodic output, even if the frequencies are not directly related to eachother. However, when the input is just random noise, would this also produce periodic output? Using the same model and configuration, the only thing that was changed was that the input string now determined the chance that the input neuron would fire. So a 5 would be a 50% of firing, and a 9 90%. Input strings were changed to match the amount of incoming charge to the first trial, only higher to account for the fact that the input neurons were now not synchronized and so would have more trouble to reach the threshold of subsequent neurons. For most of the trials, the stimuli 00002222, 22220000, 44440000 and 00004444 were used.

Again, the network trained on the stimuli for 800 ticks and then ran for 2000 each. Perhaps unexpectedly, the results did not differ much from the results in the previous chapter. The output neurons often had a very clear frequency, and changed this frequency according to the input, although this changing happened slightly less frequently. All the kinds of graphs seen in the previous chapter, including the very long sequences, could also be observed when the input was noise.

Another observation was that although the input was noise, the Hebbian learning at the beginning was still critical. This doesn't seem to make sense; how can a network be trained to recognize a non-existent pattern, as is the case with white noise? The most likely answer is that learning in this particular network doesn't mean that it learns to recognize anything; in fact, after deliberate testing it seems that it doesn't really matter on what stimulus we train the network, as long as it is significantly activated. What the learning likely does in this network, is that it 'eliminates' poorly linked neurons (neurons that are rarely active when other neurons are) by significantly lowering their weight so they don't have much influence on the other neurons. At the same time it strengthens neurons that have a strategic place in the network, i.e. neurons that are active when many other connected neurons are also active.

Model of wave interactions

This concludes our research on the behavior of individual neurons. However, oscillations also seem to play an important role on a larger scale, in the form of brain waves. These waves are thought to be critical to a well-functioning brain, even though what exactly their function might be remains vague (Buzsáki, 2006). The 40 Hz brain wave (gamma) has been strongly linked to consciousness (Traub, 1999). Distortions in brain wave activity have been linked to (amongst others) several neurological diseases (Traub & Whittington, 2010).

To research how these waves might interact, a different model was used, pictured in figure 10. Each neuron now serves a more abstract function, as it represents a group of neurons, an *ensemble*. Its 'charge' could be seen as representing a Local Field Potential (LFP).

Figure 10

The four input neurons (ABCD) all project to the output neuron in the middle. Each neuron receives feedback from the output neuron, but this feedback line is weaker than the feed forward line. The input neurons receive charge from a virtual LPF, in the form $I = a + a * sin(t + p)$, where I is the input charge, a is the amplitude of the sinusoid function, t is the time (in this case, the number of steps the model has done so far) and p is the phase. For instance, if we turn B, C and D off and set the amplitude of neuron A to 0.4 and the phase to 0, the result will look like figure 11. The activity of the output neuron is plotted in black.
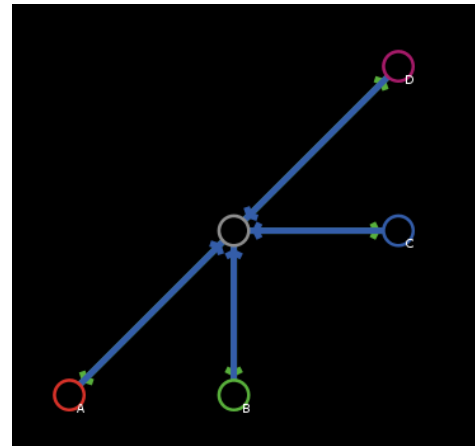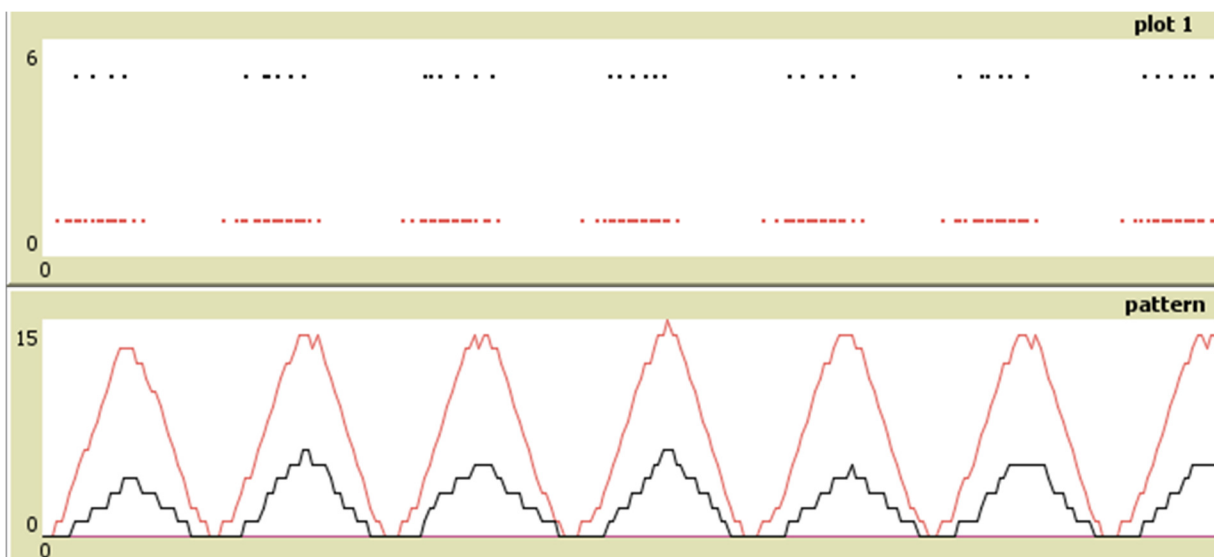

Figure 11 – Neuron A (red) receives charge from a sinusoid function, causing it to fire in a wave-like fashion. The output neuron receives charge from neuron A only, and as one would expect, fires with the same frequency.

How does phase-locking work in LIF models?

We compared the activity of the output neuron when A & B have the same frequency and phase, against the activity when A & B have the same frequency but different phases. The results can be found in figure 12. As one might expect, signals that are phase-locked elicit a strong response in the output neuron, but as we start to shift the signals out of phase (0, ¼, 1/3, and ½ of the periodicity, respectively), the activity of the output neuron decreases, to the point where any form of periodicity is gone. In this particular experiment, the output neuron almost stops firing altogether.

If we run the same experiment with stronger weights and higher amplitude, the results are the same except the baseline is higher; shifting out of phase leads to less overall activity and especially less periodicity.
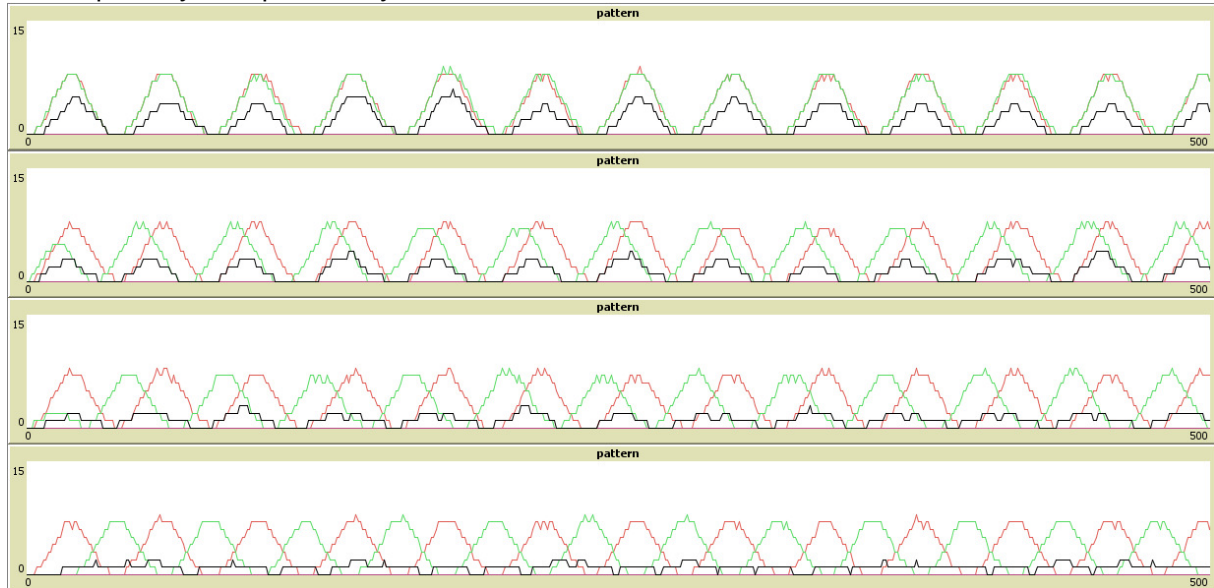

Figure 12 – from top to bottom, the waves are phased 0, ¼, 1/3, and ½ of the periodicity of one wave.

If Buschman et.al (Buschman et al., 2012) are right in their hypothesis that brain waves can activate a neural ensemble, figure 12 suggests that the brain could theoretically de-activate an active brain region by sending it a signal that has been synced out of phase. The neural construct for calculating the opposite wave of any wave X could be as simple as an excitatory neuron with high intrinsic firing, coupled with an inhibitory neuron which fires according to X (see Hoppensteadt & Izhikevich, 1996 for the learning of phase information).

This mechanism might seem familiar: the same principle is used in Active Noise Control (ANC), or anti-sound. To deal with loud noises of known frequencies, engineers have developed systems that analyze the sound wave and produce a second sound wave specifically designed to cancel the first. This works best for low, predictable frequencies. As it happens, the brain waves' frequencies are <70 Hz and the waves can be quite predictable, so it is not unthinkable that evolution might have discovered this principle long before we did.

The analogy with ANC holds true in another way. Even though two anti-phased sound waves keep the air pressure constant – preventing our ears from picking up variations – the total amount of air pressure increases because we just add more energy to already existing energy. There is no such thing as negative sound, so it is physically impossible to truly cancel out the sound waves. The higher air pressure has influence on the way external sounds are perceived.

Something similar happens in this neuronal model. If we look at the charge buildup of the output neuron during anti phased stimulation of low amplitude, we can see that the charge stays almost constant, but sub threshold. What this means is that when we use one wave to cancel the other one out, we are essentially lowering the threshold. Indeed, in the model of individual neurons, stimulating all neurons with two anti-phased waves lead to more noise propagation.

However, at the same time the threshold property is also what differentiates the brain waves from sound waves. For instance, two subthreshold oscillations of different frequencies, like 35 and 50 Hz, will cancel each other out some of the time and amplify each other at other times. If the threshold in this neuron is high enough, these two waves could result in the neuron firing with a frequency of about 21 Hz. This is illustrated in figure 13.

A counter argument to this theory of anti-phasing waves could be that since the anti-phased wave still only adds energy, the brain would be better off by not sending any signal at all. However, it's plausible that any wave has multiple functions and/or areas which it propagates to. Perhaps some of these regions are phase-insensitive, while others are not. In such a case, modifying the phase of this wave might be better than shutting it down altogether.
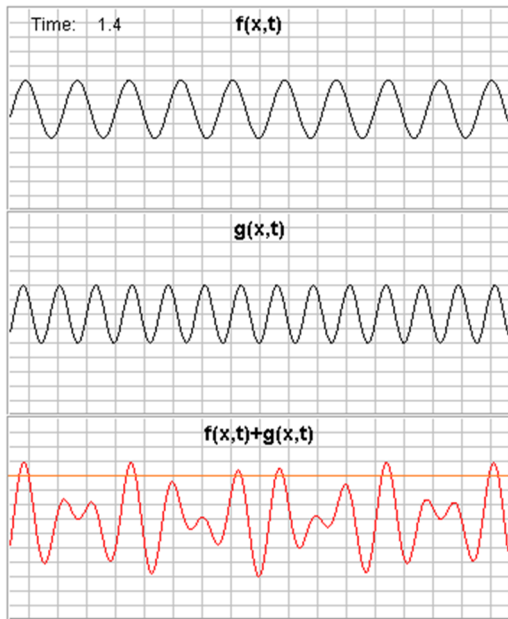

Figure 13 – Whereas two sound waves of different frequencies (left) simply yield a more complicated wave which combines the first two, using a threshold (orange line) means the two frequencies can together result in a much lower frequency than either one of the original waves. This is demonstrated in the model (below). Note that the timing of the bursts corresponds to the peaks in the image on the left, showing that the new frequency is actually semi-periodic rather than fully.
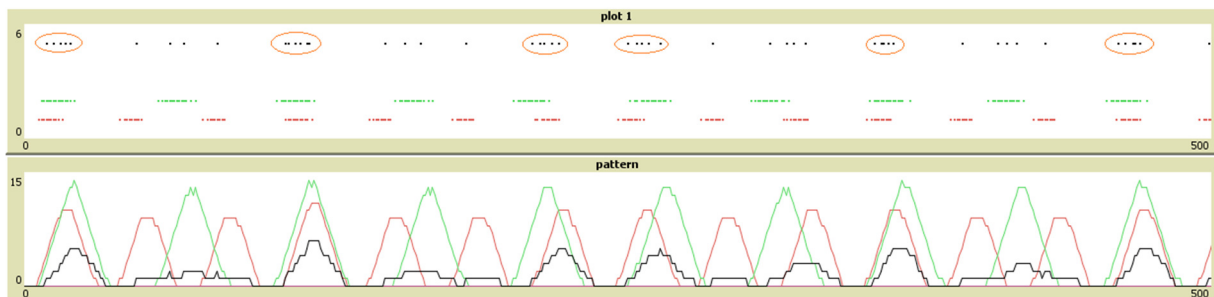
Can alpha waves desynchronize a beta wave, using only excitation?

Often, it can be useful to de-activate a brain region. The findings in the previous paragraph suggest that this might be done by sending a wave with the same frequency and amplitude, but opposite phase of the LFP currently active. However, Buschman et al. postulate "that beta-frequency synchrony selects the relevant rule ensemble, while alpha- frequency synchrony deselects a stronger, but currently irrelevant, ensemble." (Buschman et al., 2012)

In an attempt to reproduce this phenomenon in the same model as used in the previous chapter, the wave pattern of neuron A represented the beta (16-30 hz) wave, while the wave of neuron C represents the top-down alpha (8-16 hz) wave with the task of reducing the influence that A has on the output neuron. The signal was again phase shifted (0, ¼, ½, ¾ of the periodicity of neuron A) which produced different results, as can be seen in figure 14. As one might expect, figure 14 bears a striking resemblance to figure 15, which shows the addition of pure sinusoid functions as produced by an oscillator.

As mentioned before, another resemblance to sound is that these brain waves may disrupt eachother, but they only add more energy. Although the results vary with the phasing, they have in common that 1) there is still a peak in the output neuron activity whenever there is a peak in neuron A activity and 2) there is extra activity when neuron C is active. Perhaps in a larger, more complex network this might disturb the activation; for instance, resonator neurons (E. Izhikevich, 2001) would react to neuron A but not to neurons A and C paired except when the phasing is ¼ of the period of A (figure 14, top right). In this simple LIF network however, adding an extra exitatory wave simply leads to more activation, never less.

In conclusion, we can say that in a LIF model the addition of a second, slower exitatory wave will only lead to more excitation. However, the mechanism might work in combination with more complex cells like resonator neurons.

Figure 15 – addition of pure sinusoid functions. Compare the red sinusoids here with the black graphs in figure 14.
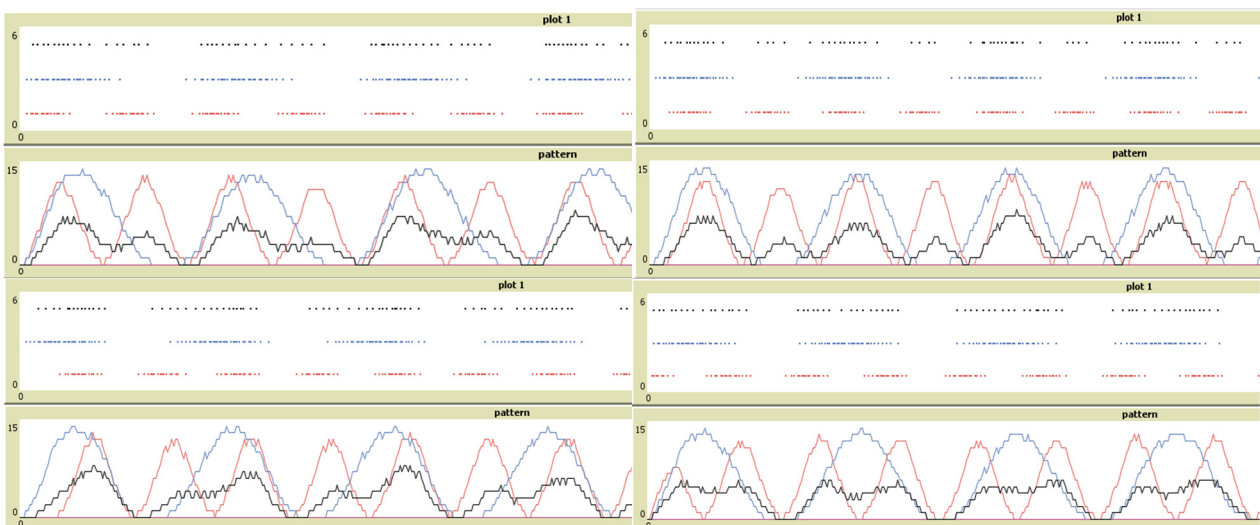
Figure 14 – The phase between neuron A (red) and C (blue) is shifted with 0 (top left), ¼ (top right), ½ (bottom left), and ¾ (bottom right) of the periodicity of neuron A respectively. The activity of the output neuron is shown in black.

Model combining individual and wave mechanics

In the previous chapters, we concluded that the even the simple properties of LIF neurons could lead to complex oscillatory behavior, both on an individual scale and at the scale of ensembles. The next model tries to combine these two effects to see if we can reproduce some of the effects shown by Buschman et al. Specifically, we ask whether a LIF model can show the behavior demonstrated in the Buschman et al. experiment.

The new model is shown in figure 16. The neurons at the bottom represent 6 input neurons, which fire in a wave pattern when switched "on". Two small networks are situated on opposite sides. The network on the left represents a "line-rule" ensemble which should return "true" if the three left-most neurons are active. The network on the right represents a "count-rule" ensemble which should return true if four or more neurons (at any location) are active. The rules should only be applied when selected by a top down wave which represents attention; the question now remains whether this single wave, applied to both networks, can select which rule to apply, especially in a situation when both rules could be applied.
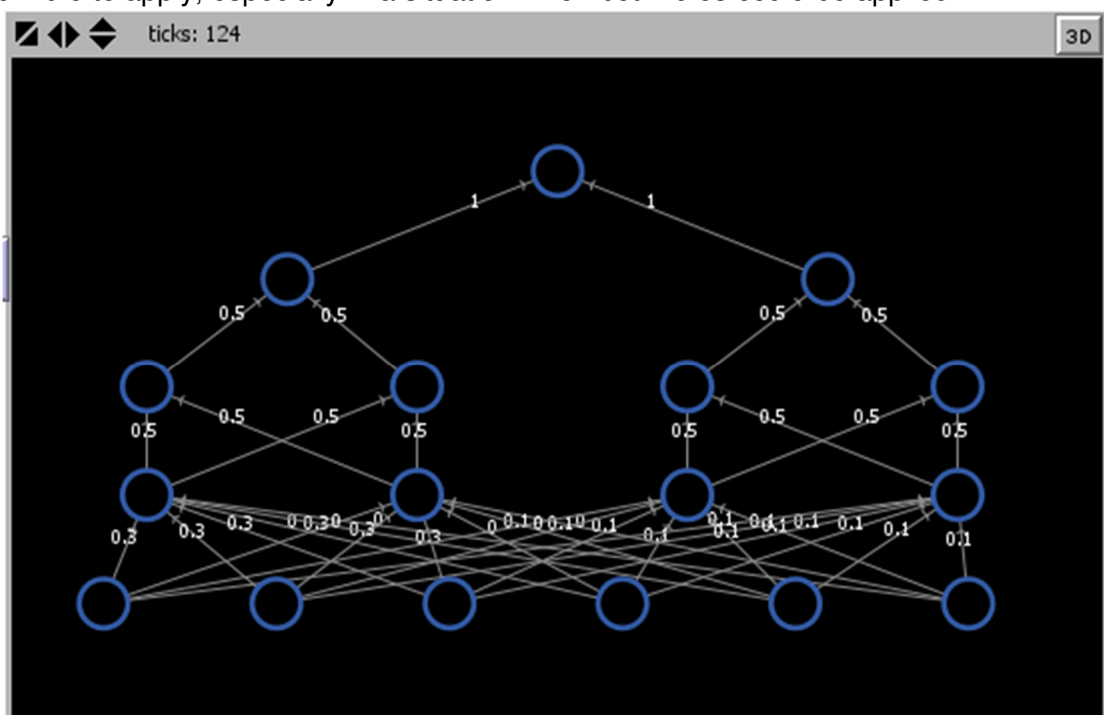

Figure 16

Can a LIF model show the behavior demonstrated in the Buschman et al. experiment?

Considering that the top-down wave projects to both networks, and the left half of the input neurons also connect to both sides (be it with different weights), it is no easy task to activate one network without activating the other. Because all the input neurons fire according to the same wave, we cannot enhance part of the stimulus.

However, by altering the threshold values in both networks we can take advantage of the effect we saw in the previous chapters, which is that phase-locked waves produce higher peaks than anti-phased waves. Using the right combination of variables, activity of both networks can be significantly manipulated, as seen in figure 17. The stimulus for this experiment was 111111, meaning all the input neurons were firing in a wave pattern and both rules could be applied. When the top-down signal is absent (top), there is some base activity in the count-rule network. When the top-down signal is switched on with the same frequency but antiphased to the stimulus (middle), the waves take turns in reaching the low threshold of the count-rule network, increasing it's activity about four-fold, while rarely reaching the high threshold of the line-rule. Finally, when the top-down signal is phase-locked to the stimulus (bottom), the high threshold of the line-rule network is reached often, while the low threshold of the count-rule network is now reached only half as much as when the top-down wave was anti-phased.
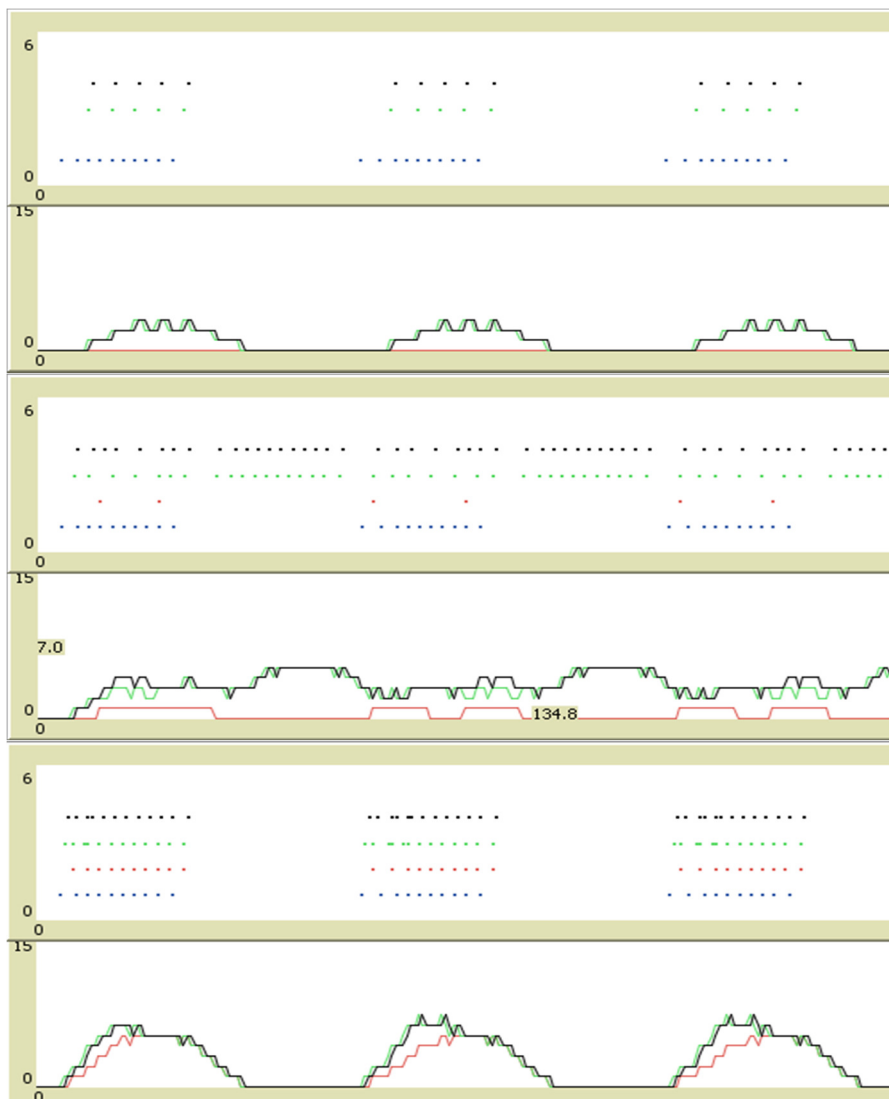

Figure 17 – color scheme: blue is the activity of the input neurons. Green is activity in the count-rule network. Red is activity in the line-rule network. Black is activity at the top neuron.

Although this configuration works well with the stimulus 111111 where both rules can be satisfied, it is lacking for other inputs. Simply put, the much lower threshold for one rule means that it is always activated whenever the high threshold rule is also active. This results in table 1.

Top down phase / stimulus	Satisfies line rule	Satisfies count rule	Satisfies line & count rule
Phase locked	Line activity	Equal count and line activity	Equal line and count activity
Anti-phased	-	Count activity	High line activity
none	-	-	Small line activity

Table 1

A different approach is to keep all threshold levels the same, but change the weights of the stimulus so that one rule becomes dominant over the other, meaning it doesn't require top-down activation when the stimulus satisfies both rules. This would result in table 2.

Top down phase / stimulus	Satisfies line rule	Satisfies count rule	Satisfies line & count rule
Phase locked	Line activity	Count activity	Equal line and count activity
Anti-phased	-	-	Dominant rule
none	-	-	Dominant rule

Table 2

After trying out many different configurations as well as combining many different wave frequencies, we have to conclude that in a LIF model, the only configuration that leads to the ideal table 3 has to make use of an inhibitory neuron in the not-dominant network that suppresses the dominant network when it is activated. However, this makes the model much more trivial and thus it is not considered a viable solution to our original question.

.

Top down phase / stimulus	Satisfies line rule	Satisfies count rule	Satisfies line & count rule
Phase locked	Line activity	Count activity	Rule A
Anti-phased	-	-	Rule B
none	-	-	-

Table 3

Conclusion

We asked ourselves if the simple rules of the Leaky Integrate and Fire model would be enough to give rise to complex behavior, particularly oscillatory behavior. We found that in a large network, with multiple layers of unorganized neurons coupled with Winner Take All layers, the output neurons indeed show significant oscillatory behavior, even when the input is just white noise. This behavior (frequency and sequence) changed when different stimuli were fed to the network. We also found that Hebbian learning is a key factor in this, even though it didn't really seem to matter which training stimulus was used as long as it was sufficiently active.

We also explored the brain wave dynamics in a LIF model. The results suggest that although the wave dynamics are similar to sound waves, the unique threshold property of neurons gives them the ability to turn two different waves into one lower frequency wave, and that anti-phasing two same-frequency waves can lead to much less activation than when the signals are phase-locked. However, the simple rules of LIF do mean that adding one wave to another can only ever lead to more activation, never to less. This is different when considering more complicated neurons, such as Resonate-and-Fire neurons (E. Izhikevich, 2001).

Finally, although there were multiple configurations to select between two rule-encoding networks using one top-down attention wave, none of these worked well for all possible inputs. To reconstruct the findings by Buschman (Buschman et al., 2012), a more complex neuronal model seems to be required.

Appendix A - Relevance to AI

What does this have to do with Artificial Intelligence? Obviously, neuroscience is an important topic in the field of AI. If we are going to make something intelligent, we have to understand what intelligence is. The more we know about the brain, the more we can apply this knowledge to our AI machines.

But there is another reason why this particular topic (neural oscillations) is very interesting to me as an AI student. See, the brain is an immensely complex machine that consists of billions upon billions of neurons, each with on average 5000 connections to other neurons. And each person's brain is different; just like the vascular system, it seems that the brain is created according to rules that divide it into regions but do not specify for each cell (neuron) exactly where it should go and which neurons it should connect to (hard-coding every cell would simply require huge amounts of DNA). This creates a certain level of randomness at the local level. Yet despite the fact that on a small scale every persons machinery is quite different (in topology at least), at a larger scale it always seems to work more or less the same, which implies the brain is very robust.

It seems then that somehow, out of this chaotic bunch of randomly interconnected neurons, order arises in more or less the same fashion in every person, meaning different networks can lead to the same behavior. So one could say, that at a local level, neural networks are *self-organizing*. Self-organizing systems are of great interest to AI researchers because they provide complex behavior with only simple local rules. Think for instance about flock of birds. Any bird has only 3 rules: don't get too close to another bird, don't go too far away, and fly in the same direction as everyone else. These three simple rules can lead to quite complex behavior, such as flocking and the well-known V-formation.

In neurons, these "simple" rules mostly apply to the ion channels in the membrane. These local rules lead to a whole spectrum of complex behavior. On a local and regional level, we can clearly see a pattern we call brain waves. There are very distinct brainwaves, ranging from 4 Hz to 70 Hz. We can also distinguish neural ensembles, groups of neurons that can be activated by a certain frequency or input (Buschman et al., 2012). Of course, the most complex behavior that follows from these simple rules is that of ourselves, but that is well beyond the scope of this thesis.

Appendix B – Autocorrelation

In order to find the frequencies hidden in the data, we used the autocorrelation approach. To calculate the autocorrelation, we copy our signal and slightly move it to the right, one step at a time. This is called *lag*. Each time we move it (increasing the lag), we multiply the value of the original signal with the value of the displaced signal. The easiest way to explain this is with an example.

Suppose a neuron fires exactly every 3 ticks. If we write *firing* as 1 and *not firing* as 0, this signal looks like 001001001……001. To calculate the autocorrelation, we copy this and move it one point to the right:

```
 001001001001001
001001001001001
```

Now if we multiply every number in the lower string with the number above it, and add these values all up, we get 0. So the autocorrelation for lag = 1 is 0. If we increase the lag once more (lag = 2), we get:

```
  001001001001001
001001001001001
```

Which also results in 0. However, increasing the lag to 3 gives us:

```
   001001001001001
001001001001001
```

Now, because 1's are above 1's, the autocorrelation is much larger than 0.

If we repeat this, we will find that the autocorrelation for lag = 1,2,4,5,7,8,10,11 etc. is always 0 while for lag = 3,6,9,12 etc. the autocorrelation is high. This suggests that the periodicity of the signal is 3.

References

Basar, E., & Basar-Eroglu, C. (1999). Oscillatory brain theory: a new trend in neuroscience. *… in Medicine and …*, (June). Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=765190

Buschman, T. J., Denovellis, E. L., Diogo, C., Bullock, D., & Miller, E. K. (2012). Synchronous oscillatory neural ensembles for rules in the prefrontal cortex. *Neuron*, *76*(4), 838–46. doi:10.1016/j.neuron.2012.09.029

Buzsáki, G. (2006). *Rhythms of the Brain. Rhythms of the Brain* (Vol. 1, p. 448).

Hebb, D. O. (1949). *The organization of behavior*. (Wiley, Ed.)*The Organization of Behavior* (Vol. 911, p. 335). Wiley. Retrieved from http://scholar.google.fr/scholar?hl=fr&q=hebb+1949&btnG=Rechercher&lr=&as_ylo=&as_vis=0#3

Hoppensteadt, F. C., & Izhikevich, E. M. (1996). Synaptic organizations and dynamical properties of weakly connected neural oscillators. II. Learning phase information. *Biological Cybernetics*, *75*, 129–135. doi:10.1007/s004220050280

Izhikevich, E. (2001). Resonate-and-fire neurons. *Neural Networks*, *14*, 883–894. Retrieved from http://www.sciencedirect.com/science/article/pii/S0893608001000788

Izhikevich, E. 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), 1063–70. doi:10.1109/TNN.2004.832719

Jensen, O., Kaiser, J., & Lachaux, J. P. (2007). Human gamma-frequency oscillations associated with attention and memory. *Trends in Neurosciences*.

Lakatos, P., Chen, C. M., O'Connell, M. N., Mills, A., & Schroeder, C. E. (2007). Neuronal Oscillations and Multisensory Interaction in Primary Auditory Cortex. *Neuron*, *53*, 279–292. doi:10.1016/j.neuron.2006.12.011

Oster, M., Douglas, R., & Liu, S.-C. (2009). Computation with spikes in a winner-take-all network. *Neural Computation*, *21*(9), 2437–65. doi:10.1162/neco.2009.07-08-829

Perlmutter, J. S., & Mink, J. W. (2006). Deep brain stimulation. *Annual Review of Neuroscience*, *29*, 229–57. doi:10.1146/annurev.neuro.29.051605.112824

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

Traub, R. D. (1999). *Fast Oscillations in Cortical Circuits*. MIT Press.

Traub, R. D., & Whittington, M. A. (2010). *Cortical oscillations in health and disease*. Oxford University Press.