École Polytechnique

Stage: UNIC, CNRS, Gif Sur Yvette

M2 Project: Investigation of emergent properties of visual perception using data-driven modelling and simulation.

Author: Ramón Heberto Martínez
Supervisors: Andrew Davison, Jan Antolik

Erasmus Mundus Joint Masters Programme in Complex Systems Science

06 July of 2014

Abstract

Receptive fields are used in neuroscience as a tool to understand how the spatio-temporal properties of the stimulus affects the neuronal response. Receptive fields in particular are used to classify cells in the visual cortex in simple or complex cells. There is, on the other hand, no consensus in the definition of the receptive field and some definitions depend on the particular type of stimulus used. In this work, we studied how the receptive fields of the neurons in the visual cortex depend on the stimulus used to estimate them. To do this, we built a model of the the visual cortex consistent with most of the biological facts found so far. We used the Volterra expansion to study how the linearities and the non-linearities of the receptive field varied with the statistics of the stimulus. We found systematic variation both at the level of the receptive field and in the response.

1 Introduction

1.1 Receptive Fields

Definition

The concept of the receptive field was introduced by Sherrington in the context of the sensitivity of the body surface to certain mechanical stimuli [Sherrington, 1966]. In a more general way a definition that encompasses different sensory modalities can be given as: "The receptive field is a portion of sensory space that can elicit neuronal responses when stimulated appropriately". The sensory space can in some cases be a single dimension (e.g. carbon chain length of an odorant), two dimensions (e.g. skin surface) or multiple dimensions (e.g. space, time and tuning properties of a visual receptive field) [Alonso and Chen, 2009].

If we concentrate on the the visual system the receptive field usually involves two spatial dimensions and a temporal one. The spatial dimensions represent how the spatial distribution of the stimulus affects the response and the temporal dimension captures how the time trajectory of the stimulus affects the cell's activity. The three dimensions together can then be called a spatio-temporal receptive field (STRF).

Simple vs Complex Receptive Fields

In figure 1 we present the basic scheme of how information is processed in the visual system, that is, we present the visual information pathway. At the start of the visual pathway, a photon coming from the outside world elicits a response in the retina. The response is in turn relayed to a structure deep in the brain called the thalamus, in particular it goes to the structure called the Lateral Geniculate Nucleus (LGN). From the thalamus the information is further processed and finally relayed to the rear part of the brain where the visual cortex is located. It is important to mention that along most of the stages of this process recurrent connections exist that convey information back in the hierarchy, making the system highly non-linear and therefore not easy to analyse.
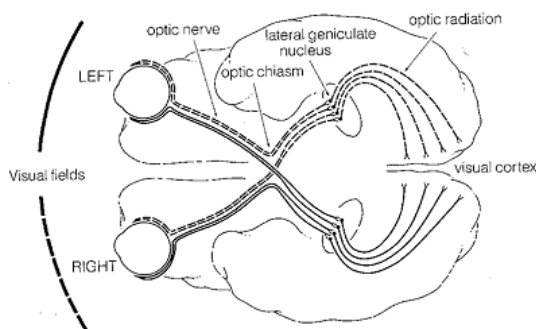

Figure 1: Here we show a schematic of the visual pathways. We can appreciate that the system maps the information first from the ganglion cells in the retina and then to the lateral geniculate nucleus through the optic nerve. Finally the information arrives the visual cortex through the optic radiation (Adapted from Carandini [2012])

If we look at the visual pathway a natural question to ask is how the structure of the receptive fields varies as we move through it. It turns out that the receptive fields present variability in two ways: first, the receptive fields vary within each particular stage of the visual pathway, that is, there is a variability at each level of the hierarchy. Second, the receptive fields vary even more as we move through the different stages of processing.

As an example of the first type of variability we have the retina. In the retina of humans there exists an area where sharp vision takes place called the fovea. The cells in the fovea have the smallest receptive fields whereas as we move away from the fovea and towards the the visual periphery the cells posses

progressively bigger receptive fields [Carandini, 2012]. This fact is consistent with our generally good resolution for central vision and poor resolution for peripheral vision [Alonso and Chen, 2009].

On the other hand as we move to higher stages of the visual pathway the spatio-temporal structure of the receptive field not only changes in size but also in structure. For example, most cells in the retina and in the LGN have a simple structure that consists of two concentric areas of inhibition and excitation respectively, as illustrated in figure 2. This concentric structure was named center-surround by Kufler and Hubel, its discoverers [Kuffler et al., 1953].
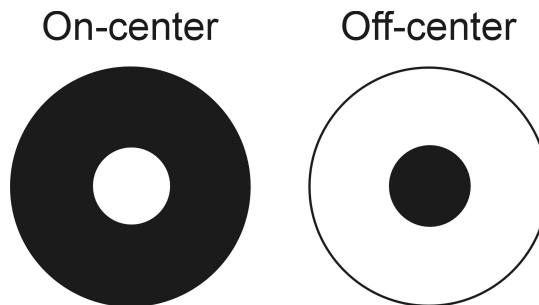
Figure 2: Here we show the on-center and off-center receptive fields. The on-center structure are tuned for bright spots up to certain size (like glowing eyes in the night) whereas the off-center are tuned for dark spots.

However, the structure of the receptive fields becomes more complex in the visual cortex. There are a few cells that have the typical center-surround structure of the receptive field presented above for the LGN and retinal neurons but most of the cells present a bigger variety of structures. There are receptive fields which present localized but separated regions that respond to light or dark spots respectively and some that have areas that respond to both light and dark spots for big areas. There are even some cells that do not respond to spots at all [Carandini, 2012].

Neurons in the part of the cortex called V1 where the information from the thalamus arrives directly are classically classified into two types according to their structure: **simple cells and complex cells**. Simple cells are defined straightforwardly as the ones with separate subregions for dark and bright stimuli. Complex cells are defined as all the other cells that do not have separate subregions [Martinez and Alonso, 2003]. We illustrate both types of cells in figures 3 and 4.
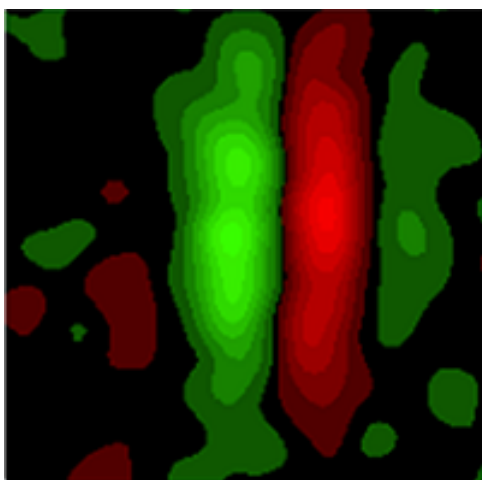

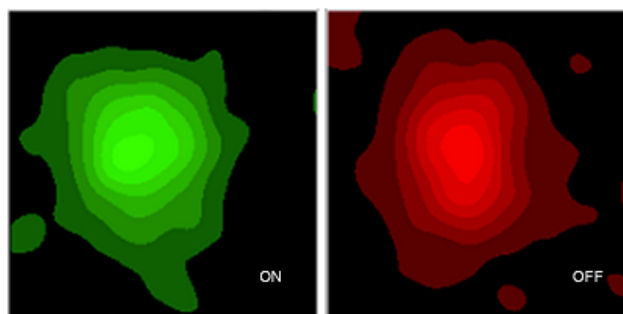Figure 3: Typical receptive field in V1 of a simple cell. [DeAngelis et al., 1995]

Figure 4: Typical complex receptive field in V1 [DeAngelis et al., 1995]

Other classification methods proposed include the response of cells to a sinusoidal grating stimulus. The first group of neurons, in this scheme, are classified as the neurons whose response to the stimulus is a rectified sinusoid and it corresponds more or less to simple cells. The second group can be characterised as the one that responds by increasing the mean firing rate and corresponds mainly to complex cells [Skottun et al., 1991]. Other additions and substitutions (e.g. spontaenous activity

level, length summation, response to patterns of random dots, etc [Skottun et al., 1991]) have been added in the hope of presenting a better scheme to classify neurons but, to the date, there is no consensus in the scientific community about which criteria to accept in general [Alonso and Chen, 2009].

In this work we are interested in studying how the structure of the receptive fields depends on the nature of the stimulus involved on its calculation. In order to do that we first present a definition of the receptive fields that allows us to distinguish between the linear and the non-linear part of receptive fields and a method for their estimation. Secondly we present the stimuli that we used and the details of the model of visual that we use. A systematic relationship between the variation of receptive fields in V1 cells and varying stimuli has been found previously [Fournier et al., 2011]. We use their structure metrics to study whether this phenomenon is present in our model of visual processing or not.

2   Materials And Methods

2.1   Estimating Receptive Fields

In mathematical terms the receptive field can be thought as a function that captures how the stimulus is related to the response and allows us to map the stimulus onto the response of the neuron. In this section we briefly review two quantities that can play the role of such a function: The Spike Triggered Average and the Volterra Kernel.

Spike Triggered Average (STA)

In this section we briefly review two methods to estimate the receptive field of a given neuron. First, the classical method is the spike triggered average. The STA is the mathematical average of the stimuli that produced a response in the system (Figure 5). As a first approximation it provides us with a useful estimate of the linear receptive field.
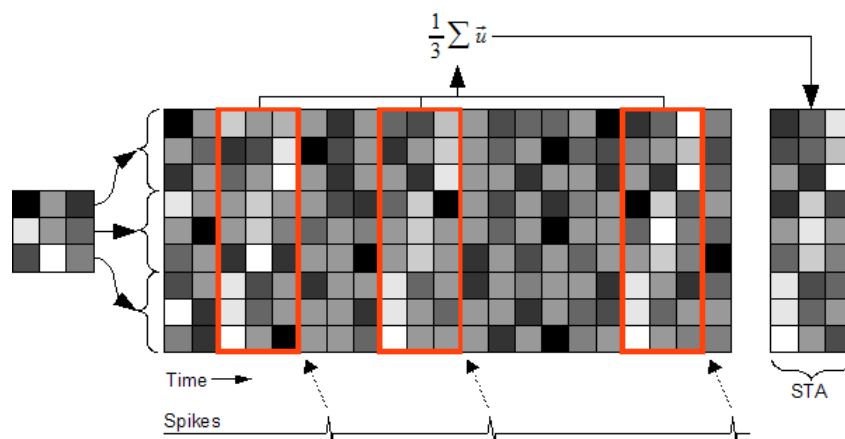

Figure 5: To calculate the STA we average the stimuli before each of the spikes in a given time window. Here in the x axis we have each of the spikes of the system (three) marked with arrows. We then average the values of the stimulus in a given time window, here marked with red, to obtain the STA (figure taken from Wikipedia Commons).

Volterra Expansion

One of the most commonly used techniques to estimate the firing rates is the linear non-linear model LNL. In this method a linear filter K is used to take into account the values of the stimulus at all earlier times through a convolution with our stimulus S. Then the convolution output is used as the argument of a non-linear function NL in order to take into account the non-linearities of the system [Dayan et al., 2001]. If R_0 is the mean response we can write the response at time t as :

$$R(t) = R_0 + NL\left\{\int K(\tau)S(t-\tau)\,d\tau\right\}$$

In other words, the linear kernel K dictates at which times the stimulus matters most for the response. The Non-linear NL part after that is applied to take into account simple non-linearities of the system like rectification.

A more general approach to estimate the response, which takes into account all the non-linearities, is the Volterra expansion of a signal [Volterra, 1887] [Wiener, 1958]. In the same spirit as the classical Taylor series we can decompose a particular signal into operators that successively add degrees of non-linearity in the following way:

$$\begin{aligned}
R(t) = h_0 \ &+ \ \int h_{1st}(x_1, y_1, \tau_1)S(x, y, t-\tau_1)\,d\tau_1 dx_1 dy_1 \\
&+ \int h_{2nd}(x_1, y_1, \tau_1, x_2, y_2, \tau_2)S(x_1, y_1, t-\tau_1)S(x_2, y_2, t-\tau_2)\,d\tau_1 dx_1 dy_1 d\tau_2 dx_2 dy_2 \\
&+ \int h_{3rd}(x_1, y_1, \tau_1, x_2, y_2, \tau_2, x_3, y_3, \tau_3)S(x_1, y_1, t-\tau_1)S(x_2, y_2, t-\tau_2) \\
&\qquad S(x_1, y_2, t-\tau_3)\,d\tau_1 dx_1 dy_1 d\tau_2 dx_2 dy_2 d\tau_3 dx_3 dy_3 \ + \ \ldots
\end{aligned}$$

Here $R(t)$ represents the cell's response in time, which is decomposed in the terms of first h_{1st}, second h_{2nd} and third h_{3rd} Volterra kernels. S stands for the stimulus and in order to take into account the stimulus over the whole space we integrate also over the space coordinates x and y . The τ terms account for delays in time and are the terms that allows us to take into account the correlations at different times. The spatial correlations appear as multiplications of the stimulus terms at different points in space.

This particular decomposition is useful to us because we can use the size and structure of the coefficients h_n as a measure of how the different degrees of non-linearity play a role in the response of the system.

In order to make a concrete use of the series we have to restrict ourselves to the discrete representation of the series which can be written in the following way:

$$\begin{aligned}
R(t) = h_0 &+ \sum_{x,y,\tau} h_{1st}(x, y, \tau)S(x, y, t-\tau) \\
&+ \sum_{x_1,y_1,\tau_1,x_2,y_2,\tau_2} S(x_1, y_1, \tau_1)h_{2nd}(x_1, y_1, \tau_1, x_2, y_2, \tau_2)S(x_2, y_2, \tau_2)
\end{aligned}$$

Ultimately in this work we are interested in determining how the linearity or non-linearity of a particular receptive field varies with the stimulus. In order to do this we only need to take into account the minimal model that allows us to distinguish between linear and non-linear components of the receptive field. Therefore, we limit the Volterra expansion up to second order with the aim of reducing the number of parameters that we need to estimate from the kernels. Taking this into account our expression above is reduced to:

$$R(t) = h_0 + \sum_{x,y,\tau} h_{1\text{st}}(x,y,\tau)S(x,y,t-\tau) + \sum_{x,y,\tau} h_{2\text{Diag}}(x,y,\tau)S^2(x,y,t-\tau)$$

Where $h_{2\text{Diag}}$ represents the diagonal terms of the second order response and is the minimal non-linearity that can be added to a system if we want to take fully into account the spatial and temporal characteristics of our system.

From this point forward we will use interchangeably the terms h_1, simple part of the receptive field and linear part. The same for h_2, complex part and the non-linear part.

Estimation of the Volterra coefficients

Our objective now can be defined as to estimate the coefficients $h_{1\text{st}}$ and $h_{2\text{Diag}}$ in the reduced Volterra expansion, which can be thought as a regression problem in the following way. If $R(t^i)$ is the response of our model at a given time t^i we can define $H(t^i)$ (for hypothesis) as the current fit to our expansion:

$$H(t^i) = h_0 + \sum_{x,y,\tau} h_{1\text{st}}(x,y,\tau)S(x,y,t^i-\tau) + \sum_{x,y,\tau} h_{2\text{Diag}}(x,y,\tau)S^2(x,y,t^i-\tau)$$

With this in mind the problem of finding the best fit is the same as the problem of minimizing the following cost function:

$$J(h_0, h_{1\text{st}}, h_{2\text{Diag}}) = \sum_i (R(t^i) - H(t^i))^2$$

where the index i goes all over the examples that we have in our measurements.

2.2   The Stimuli

We want to probe whether the receptive fields of the system under study are subject to stimulus dependent variations or not. In order to probe this we use two stimuli with different spatio-temporal statistics to study the differences of the receptive fields estimated under those conditions: **dense noise** and **sparse noise**. In the following graphical description we will represent the background luminance by gray, that is, the neutral valueswith respect to which we measure the differences. Black and white on the other hand represent smaller and bigger luminance values respectively equidistant in luminance value from the gray background as illustrated in figure 6 and figure 7. From now on in the text we use interchangeably the terms black, dark or negative and white, bright or positive.

We proceed now to describe the stimuli that we used. In the case of dense noise, at a given point in time we choose a value for each of the pixels with equal probability between black, gray and white. We then do the same for each point in time and in this way we have a succession of pixel grids like the one illustrated in figure 6. In the case of sparse noise, for each point in time we chose randomly only one of the pixels in that grid and then chose with equal probability whether the pixel is black or white. The rest of the pixels at each time are assigned the value of gray as they represent the background luminance value. We represent four of these grids in the figure 7.
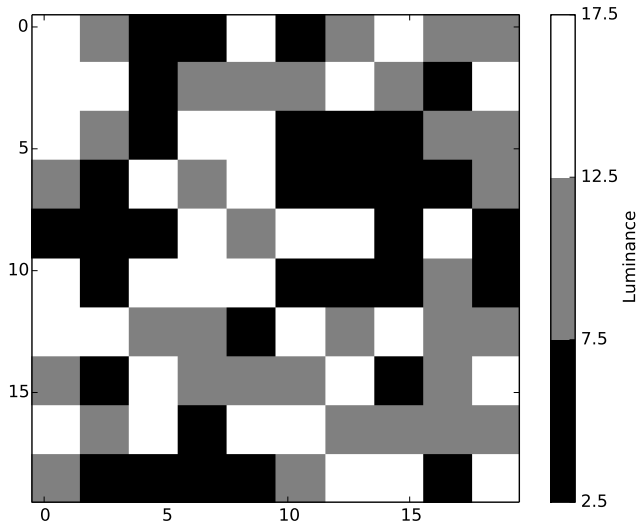

Figure 6: Ternary dense noise. Here the grid represents our stimulus at a particular point in time. At that particular moment we pick randomly the value of each pixel from a uniform distribution of black, gray or white values.
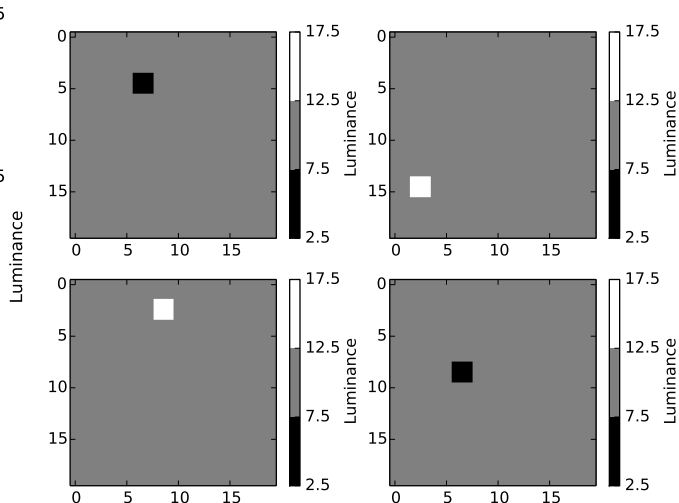
Figure 7: Sparse Noise. For each of the grids a pixel is chosen randomly and then with probability $\frac{1}{2}$ it is decided if the pixels is white or black. The rest of the pixel are set to the value of the background luminance (gray).

It is important to note that although the stimuli posses the same differences in their values of contrast their spatio-temporal statistics differ greatly. To illustrate this we can see that in the dense noise condition all pixels are activated more or less at the same time and the three values are represented equally. In consequence we have a stimulus power of $P^2 = \frac{2}{3}c^2$ for elementary contrast of value c. Sparse noise on the other hand has a very low power with only one pixel activated at a time. In a 10×10 grid the power of the stimuli would be $P^2 = \frac{2}{200}c^2$. This equals a ratio of 8.16 in the respective standard deviation of luminance values of the stimuli.

2.3   Metrics of Structure

After estimating our kernel terms h_{1st} and $h_{2\mathrm{Diag}}$ with the two different types of stimuli we would like to see whether there are structural variations in them. In order to do so, we use the metrics proposed in the literature to characterise each estimation.

In the same spirit as [Fournier et al., 2011] we can define a quantity that captures how linear or non-linear a cell is in the following way:

$$SI = \frac{\sum_{x,y,\tau} h_{1st}^2(x,y,\tau)}{\sum_{x,y,\tau} h_{1st}^2(x,y,\tau) + \sum_{x,y,\tau} h_{2\mathrm{Diag}}^2(x,y,\tau)} \tag{1}$$

This quantifies the energy of the linear term h_{1st} with respect to the total kernel energy. It is close to 0 for complex like cells and cluster towards 1 for the ones that are simple-like.

7

The above metric allow us to assign a simpleness value to each cell. However, we would like to have a metric that allow us to appreciate how the linear and non-linear component of the cell vary individually as we move from one stimulus to another. In order to quantify this we calculate the ratio of the kernel's norms calculated with the two different stimuli for each component:

$$h\,gain_{SN/DN} = \frac{\sqrt{\sum_{x,y,\tau} h^{SN}(x,y,\tau)^2}}{\sqrt{\sum_{x,y,\tau} h^{DN}(x,y,\tau)^2}} \tag{2}$$

Where the super-index SN or DN indicates that the quantity was estimated using sparse noise or dense noise respectively.

Finally, we would like to quantify how the respective components of the receptive field vary when convolved with the respective stimuli. To do this we also define a simpleness index after convolution $SI*$.

$$SI* = \frac{\sum_t (h_{1st} * S_t)^2}{\sum_t (h_{1st} * S_t)^2 + \sum_t (h_{2Diag} * S_t)^2} \tag{3}$$

Where the $*$ denotes the convolution operator. This last quantity allows us to study the effective effect upon a cell of each stimulus and receptive field structures.

2.4 The model

In this section we will present the model of the visual system that we used. The model here is based on previous work by Kremkow [Kremkow, 2011] and Antolik in our lab and is basically a small model of V1 that nevertheless shows the emergent properties that have been found experimentally in V1: orientation tuning, contrast invariant orientation tuning, appropriate conductance regime, correct statistics of firing, etc.

The structure and parameters of the model are taken mostly from the experimental literature but it is important to emphasize that the model also contain parameters that were adjusted to reproduce experimental values. In other words the model presents both fitted and biological parameters.

Lateral Geniculate Nucelus

The model is largely based on previous work in the visual system by [Troyer et al., 1998]. We constructed a realistically dense LGN layer that covers 6.8×6.8 degrees of the visual field. To do so,we utilised a lattice of 61×61 ON cells and 61×61 OFF cells.

To each cell we can associate a spatio-temporal receptive field (STRF). The profile of the STRF consists in a spatial structure modelled after the center-surround architecture mentioned above and also illustrated in the left part of figure 8. The temporal filter on the other hand possesses a temporal profile with a bi-phasic profile as proposed in [Cai et al., 1997]. In particular the sizes of the center and surround areas were set in such a way as to respect the values found in the literature [Allen and Freeman, 2006].

In order to model the input to LGN neurons a given stimulus was convolved with the STRF; this result in turn was injected into the neuron as a current. We also amplify the current with a gain factor to map the injected values to biologically relevant domains [Kremkow, 2011]. In addition, in order to introduce variability, white noise was also injected into the neurons. This was done in such a way that in the absence of external stimulus the neurons elicit a background firing rate of around 10 Hz.

Cortical Layer

Following the work of Troyer [Troyer et al., 1998] the cortex is represented by two layers: one inhibitory with 400 neurons and one excitatory with 1600 neurons [Braitenberg and Schüz, 1991]. This is illustrated in the right part of figure 8. In order to obtain Gabor-like receptive fields a sampling mechanism from the LGN cells was implemented. In this mechanism, on-cells are sampled with probability equal to the positive values of the Gabor function center in that position and off-cells are sampled with the negative part. The phase and the orientation of each Gabor function were drawn from a uniform distribution. The optimal spatial frequency was set to 0.8 cycles per degree in agreement with in-vivo recordings and the receptive field positions are within $\pm 0.2°$ of the same visual fields. Finally, the diameter of the center of the receptive fields and the aspect ratio were taken from the literature [Reid and Alonso, 1995] [Banitt et al., 2007].

The mechanism to build the connections between two neurons in the cortical layer was based on the correlations between their receptive fields [Miller, 1994] [Troyer et al., 1998] [Kandel et al., 2000]. For excitatory connections this assumed a high connectivity when the phase and orientation were the same. For inhibitory synapses, on the other hand, the connection probability was highest at the same orientation but at a phase difference of $180°$. Furthermore, short-term plasticity was added to the model of excitatory synapses in order to agree with results in the literature [Tsodyks and Markram, 1997]. To reproduce membrane fluctuations observed in vivo [Destexhe et al., 2003] noise coming from a Poisson process generator was also added to the cortical neurons.
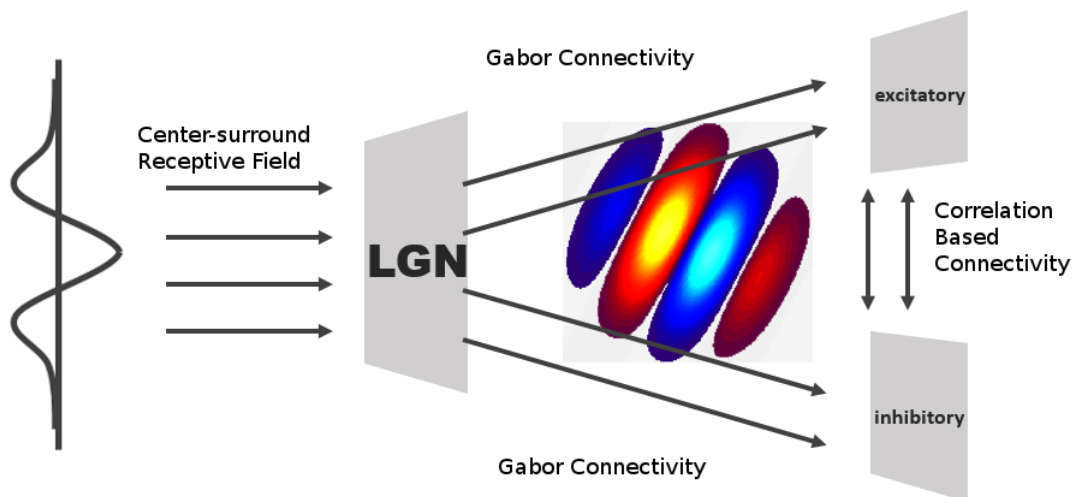

Figure 8: A schematic of our model. In the leftmost part we represent the spatial center-surround structure of the STRF. This, convolved with the stimuli produces the input that is injected in the LGN' cells. In the center of the image we represent the Gabor connectivity between the cells in the LGN and the excitatory and the inhibitory layer of the cortex shown in the rightmost part of the figure. Finally between those last two layer we have a correlation based connectivity described in the text above.

The Neurons

The neurons in LGN and V1 were modelled as leaky integrate and fire point neurons with synaptic inputs represented as transient conductance changes. Every time the membrane potential of a neuron reached the threshold voltage V_{th} a spike was emitted and then the neuron was reset to an assigned value V_{re} and kept in that state for $2\,ms$. For excitatory neurons the parameters used were $cm = 0.29\,nF$, $\tau_m = 10.0\,ms$, $V_{th} = -57\,mV$, $V_{re} = -65\,mV$, and $V_{rest} = -65\,mV$. For the inhibitory neurons the parameters were $cm = 0.141\,nF$, $\tau_m = 6.5\,ms$, $V_{th} = -57\,mV$, $V_{re} = -65\,mV$, and $V_{rest} = -65\,mV$.

Synapses

The synapses' transient changes were modelled as exponentially decaying conductances changes. The parameters corresponding to the exponential functions were $\tau_{exc} = 1.5\,ms$ and $\tau_{inh} = 10\,ms$ for the excitatory and the inhibitory neurons respectively [Kuhn et al., 2004]. The delays between the spike and the response were set to $2\,ms$ and the peak conductance to $1.5\,nS$

2.5 Tools

All the simulations were carried out in the Python language. This allowed us to profit from the software machinery that the neuroinformatics community has built around it. In order to integrate our work-flow between model and parameter specification, experiment design and simulations we have used the Mozaik software [Antolík and Davison, 2013]. Mozaik in turn relies on PyNN [Davison et al., 2008] for the model building part. Finally from PyNN the simulator NEST is called. The NEST software is optimized for large simulations of point neurons and allows us to carry out our experiments with well tested code and standard routines [Gewaltig and Diesmann, 2007].

In the spirit of open science, all the routines and code that we used are made available in the following repository:

```
https://github.com/h-mayorquin/M2_complexity_thesis
```

3 Results

In order to test our algorithms we first apply our algorithms to data from a real cell. The data were obtained under in-vivo recording conditions subjected to a form of dense noise that is consistent with the one we constructed for our model

3.1 Real Cell

Response Curves

Before calculating a fully fledged regression we would like to see whether the receptive field possesses any kind of structure. In order to do so in figure 9 we show how the voltage looked on average after either a bright (blue trace) or a dark (green trace) stimulus is presented in a particular pixel.
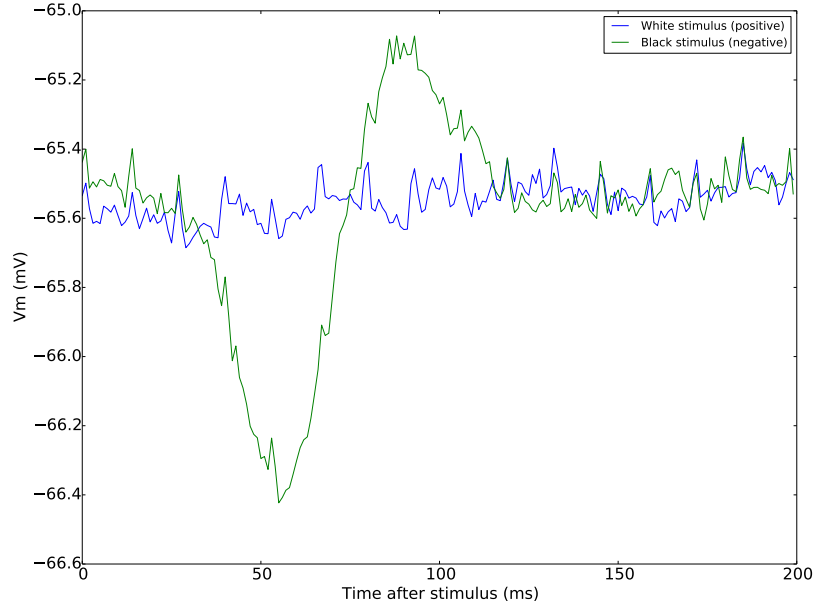
Figure 9: Time traces of the voltage corresponding to the pixel in the center $x = 5$ and $y = 5$. We show the time trace corresponding to 200 ms for the bright or positive value (blue) and the the one corresponding to the dark or negative value (green)

In figure 9 we can see that the kernel possesses a temporal bi-phasic structure when stimulated with the negative stimulus (green line in the figure). That is, starting from the mean value the voltage decreases its magnitude and then rises above the average before returning again to the starting point. This behaviour is consistent with the expected behaviour from the literature [Cai et al., 1997].

In order to reveal the spatial structure we performed the same analysis as before for every pixel in our stimulus grid. We illustrate the results in figure 10. In the figure we can see that the cell is particularly receptive to stimuli that fall in the center of the receptive field where we can see the spatio-temporal structure in more detail. The outer cells, on the other hand, are consistently more noisy and they lack any kind of structure.

Note that the response to a dark stimulus change its polarity as we move in a horizontal line from the left to the right. In the left of the center the membrane potential increases after a short time and then decreases after an interval of approximately the same duration. In the center of the picture, however, the order of the event reverses, that is, the membrane potential first decreases the magnitude of its response and only then increases it. Finally for the traces that are on the right part of the center the situation is analogous to the one described for the left. This behaviour appeared again in a more striking way when we calculated the receptive field in full detail.
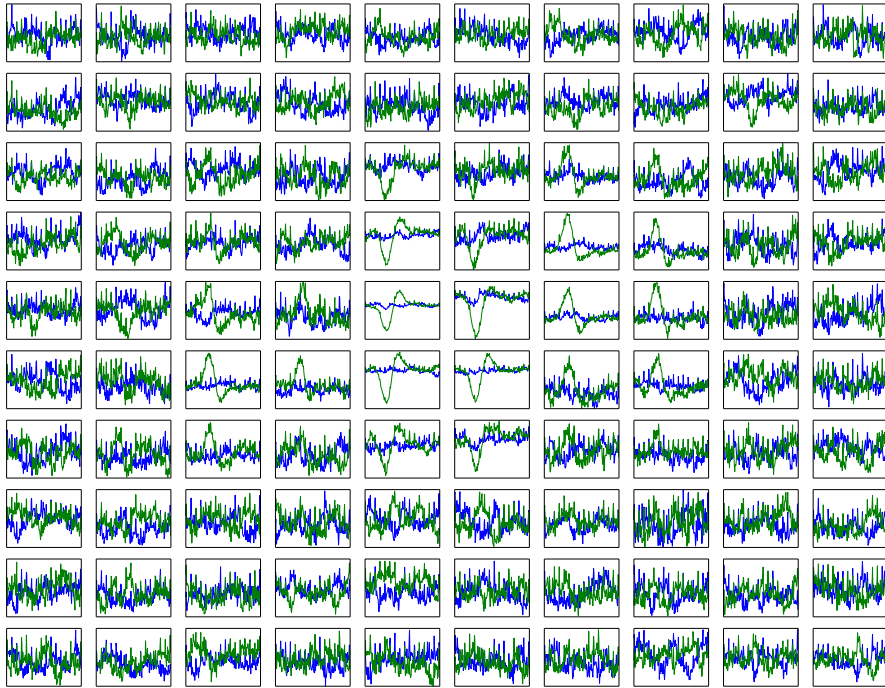
Figure 10: Response traces from a real cell under dense noise. In this plot every cell in the grid represents a particular pixel of our stimuli. For each one of them we plotted how the signal looked on average 200 ms after a bright stimuli appeared (blue trace) or a dark stimuli appeared (green trace). For each cell the x axis represents the time window that we have chosen (200 ms). They y axis represents the membrane potential and its values range from -67 mV to -65 mV.

Receptive Fields

We can then proceed to estimate the receptive field of the cell. In figures 11 and 12 we present the linear part of the kernel h_1 and the non-linear one h_2 respectively. In the figures we can see a very well defined structure in the center. This is consistent with the results from the time traces that we showed above where the polarity of of the appropriate stimulus is inverted as we move along a horizontal line that passes through the center.
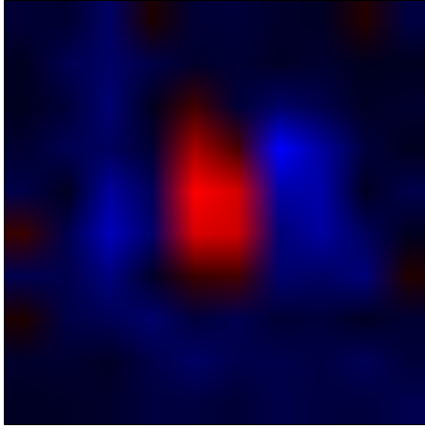
Figure 11: Estimation of h_1 for a real cell under the dense noise condition and for $\tau = 56\,\text{ms}$. Here the blue areas represent regions of the visual field where dark stimulation will increase the magnitude of the cell response and a bright stimulus will decrease it. The red areas on the other hand represents the opposite situation where a bright stimulus will increase the magnitude of the cell's response and a dark stimulus will decrease it.
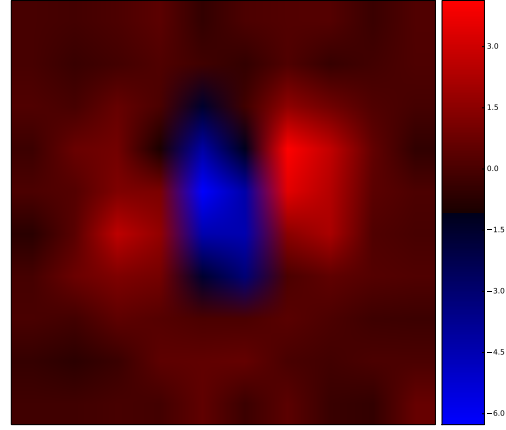

Figure 12: Estimation of h_2 for a real cell under the dense noise condition and for $\tau = 56\,\text{ms}$. Here a stimulus present in the red area will increase the magnitude of the cell response regardless of the the sign whereas one present in the blue area will decrease the magnitude of the cell's response for both bright and dark stimuli. Note that even if the structures of the linear and the non-linear part look very similar their interpretations differs.

It is important to note that the kernels can be multiplied by a constant and still give the same overall response if properly adjusted with gain. That is, the overall magnitude of the receptive field is not as important as its structure. Following this in the next discussion we would like to draw attention to the overall pattern instead of the particular magnitudes of the receptive field.
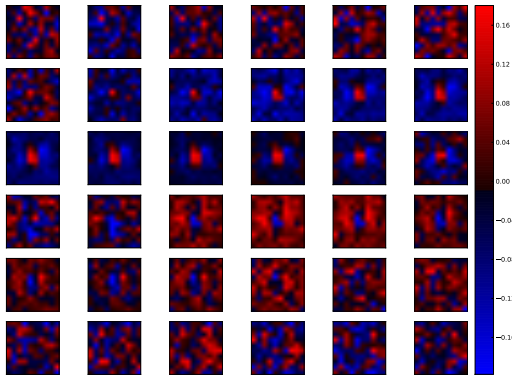

Figure 13: Estimation of h_1 for a real cell under dense noise conditions. Here the time advances in reading order, that is, from left to right and from top to bottom. Every change of frame corresponds to a jump of $\tau = 4\,\text{ms}$.
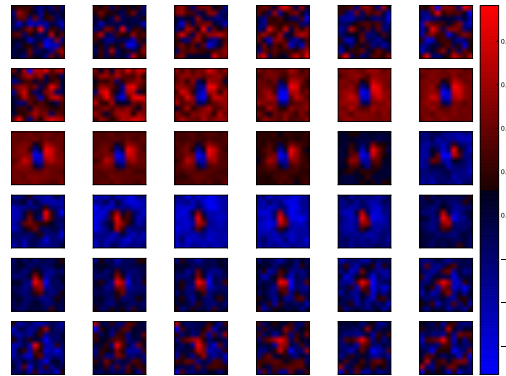

Figure 14: Estimation of h_2 for a real cell under dense noise conditions. Here the time advances in reading order, that is, from left to right and from top to bottom. Every change of frame corresponds to a jump of $\tau = 4\,\text{ms}$.

Finally in figure 13 and figure 14 we present the whole structure of the receptive field. For small delays it starts with an overall pattern that fluctuates around zero. This means that very recent stimuli do not play a big role in the response of the neuron, or in other words there is a delay of approximately $24\,\text{ms}$ before the images start to affect the neuron. This can be interpreted as the time that it takes for the signal to arrive from the retina to the cortical areas. We also can see that at some point around $60\,\text{ms}$ the polarity of the filters is inverted which is consistent with the structure we saw for the traces. Finally at some point the receptive field becomes very weak again which means that the system has a finite memory and events from time further than that do not play a big role in the response of the neuron.

3.2 Simulations

Once we are sure that the methodology works and gives the expected result for the experimental data we apply the routines to the data produced in our simulations. For the simulations we used a stimulus that consisted of 3000 frames for both sparse and dense noise with a duration of 21 ms each.

Response Curves

In order to probe the structure we show first the response curves in figure 10. As in the real cell the central traces are the ones that show most organization. Taking this into account we opted for showing only the traces corresponding to the central pixels in figures 15 and 16.

In the dense noise case we can appreciate a central bi-phasic structure of opposite polarity for the dark and bright stimulus. In the sparse noise case, on the other hand, we can see an overlap in the direction of the effects for both the dark and the bright stimulus.
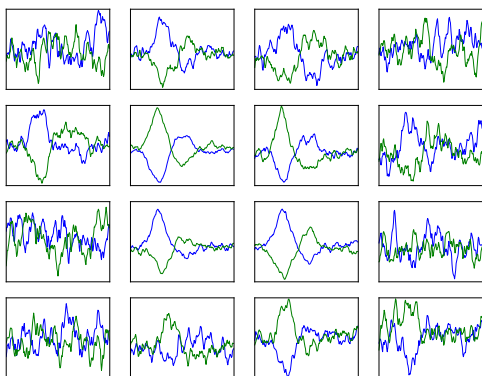

Figure 15: Stimulus traces for the center of the grid from a simulated cell subjected to dense noise. In this plot every cell in the grid represents a particular pixel of our stimulus. For each of them we plotted how the signal looked on average 200 ms after a bright stimulus appeared (blue trace) or a dark stimulus appeared (green trace). For each cell the x axis represents the time window that we have chosen (200 ms). They y axis represents the membrane potential and its values range from -67 mV to -65 mV.
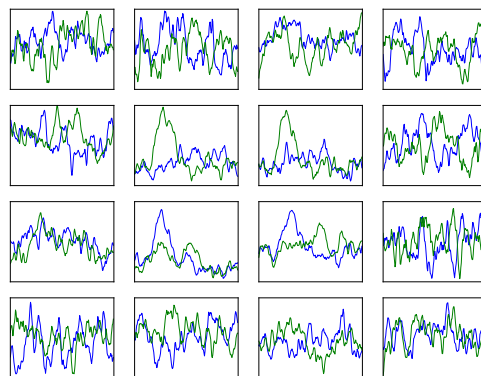
Figure 16: Stimulus traces for the center of the grid from simulated cell subjected to sparse noise. In this plot every cell in the grid represents a particular pixel of our stimulus. For each of them we plotted how the signal looked in average 200 ms after a bright stimulus appeared (blue trace) or a dark stimulus appeared (green trace). For each cell the x axis represents the time window that we have chosen (200 ms). They y axis represents the membrane potential and its values range from -67 mV to -65 mV.

Receptive Fields

In figure 17 we show the linear part of the estimated receptive field for the dense noise condition. We can see that again we do not have structure for small delays of the kernel but approximately at 30 ms two small blobs of opposite polarity appear. The whole pattern then last until approximately 100 ms and then disappears.

We also show the estimation for the non-linear part of the receptive field in figure 18. The spatial structure of h_2 here consists in a central bob where any kind of stimulus excites the cell and a central sea of inhibition. Interestingly the spatial structure for the non-linear part lasts a shorter time than the linear part.
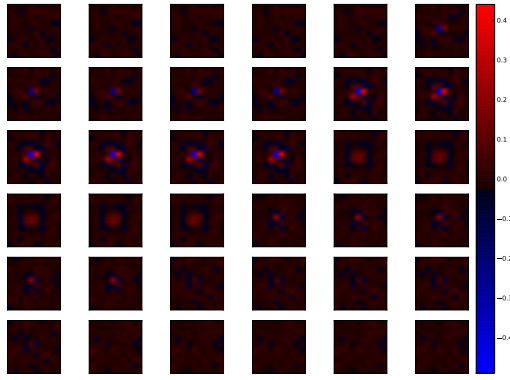
Figure 17: Estimation of the h_1 component of the receptive field for a simulated cell under the dense noise condition. Here time advances in the reading order, that is, from left to right and from top to bottom. Every change of frame corresponds to a jump in $\tau = 4\,\text{ms}$.
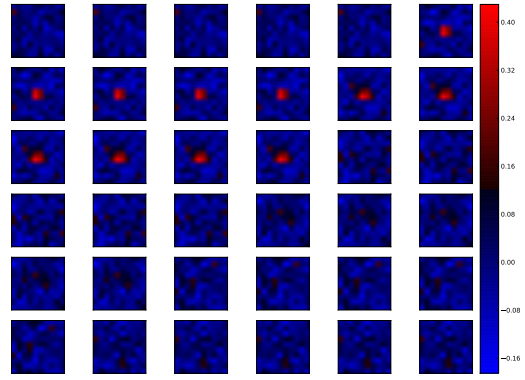

Figure 18: Estimation of the h_2 component of the receptive field for a simulated cell under the dense noise condition. Here time advances in the reading order, that is, from left to right and from top to bottom. Every change of frame corresponds to a jump in $\tau = 4\,\text{ms}$.

In figure 19 we present the estimation of the linear part of the receptive field in a cell subjected to sparse noise. Here we see a very similar pattern as for the same component estimate under dense noise condition. We see, however, that in general the structure is weaker and less organised. This is consistent with the experimental results, which show that the linear component of the receptive field becomes weaker when the stimulus used to estimate it changes from dense to sparse noise [Fournier et al., 2011].

We also present the non-linear part of the receptive field estimate under sparse noise stimulation in figure 20. Here we also note that the the non-linear part estimated with sparse noise possesses a greater spatial extent than the one from estimated under the dense condition. This last result is also consistent with what people have found so far in experiments.
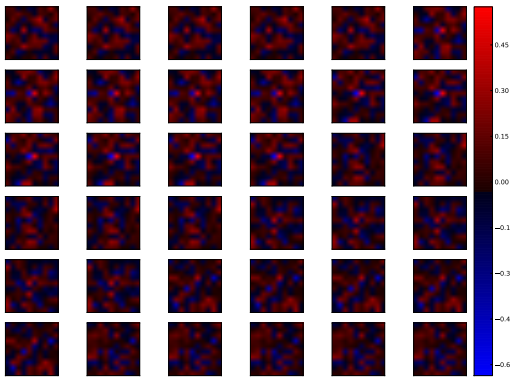

Figure 19: Estimation of the h_1 component of the receptive field for a simulated cell under sparse noise condition. Here the time advances in the reading order, that is, from left to right and from up to down. Every change of frame corresponds to a jump in $\tau = 4\,\text{ms}$.
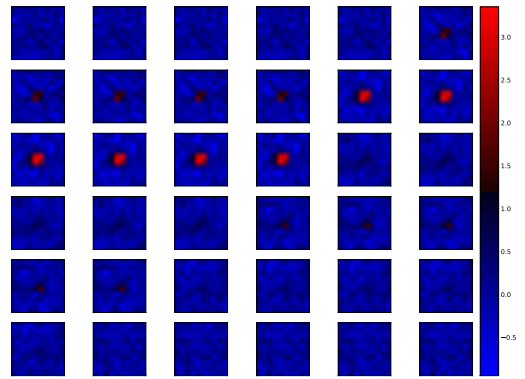

Figure 20: Estimation of the h_2 component of the receptive field for a simulated cell under sparse noise condition. Here the time advances in the reading order, that is, from left to right and from up to down. Every change of frame corresponds to a jump in $\tau = 4\,\text{ms}$.

Dependence on the Stimuli

We have shown that the structure of the receptive field varies with the type of stimulus that we use to estimate it. In order to study this systematically at the population level we simulated 21 cells and calculated the Simpleness Index that we presented in equation 1 for each of them. We did this for the receptive fields estimated both under sparse and dense noise and then compared them. In figure 21 we can see that the cells are systematically simpler with dense noise than with sparse noise, which is

consistent with the results from the experiments.

Furthermore, in figure 22 we present the gain analysis for the population as proposed in equation 2. That is, for both the linear and the non-linear component we calculate the ratio between the norm of the Volterra kernel under dense and sparse noise. In agreement with the experiments [Fournier et al., 2011] we conclude that the non-linear part is consistently the component of the receptive field subjected to the greater grain.
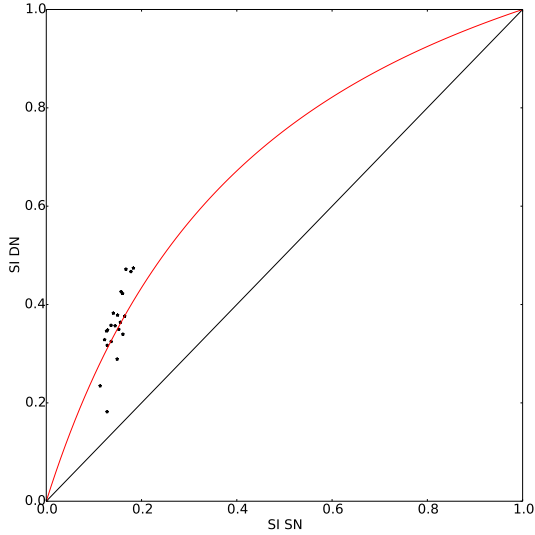

Figure 21: Here we calculated the Simpleness Index for 21 cells under dense noise and plotted it against the Simpleness Index calculated under sparse noise. In order to compare we also show the identity as a black trace. In red we show the the analytical value based on the gain fit from figure 22.
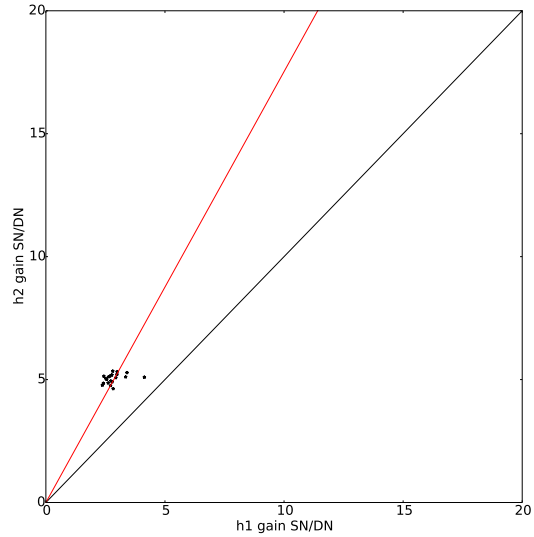
Figure 22: Here we calculated the gain in the non-linear part of the receptive field against the linear part when changing from simple to dense noise. In order to compare we show the the identity as a black trace. We show in red the best linear fit that crosses through the origin.

Finally, in order to quantify how much of the total response is attributable to either the linear or non-linear part we convolved the estimated receptive field with the corresponding stimulus. We can see in figure 23 that under sparse noise the non-linear part contributes the most while in the dense noise the opposite situation is true. This result is not consistent with experimental results where, when convolved, both components have equal contributions.
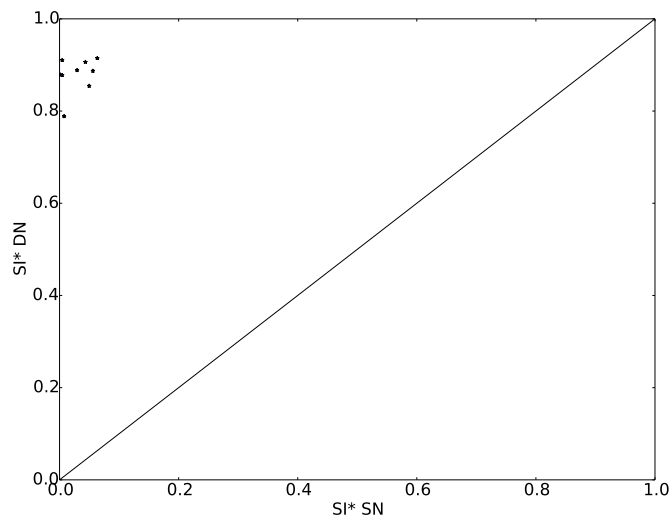

Figure 23: Here we calculated the SI* (convolution of the images with the Volterra kernels) for 21 cells under dense noise and plotted it against the SI* calculated under sparse noise. In order to compare we also present the identity as a black trace.

16

4 Discussion

The aim of this project was to study the links between the receptive field of a cell and the type of stimulus that we use to estimate it. In order to achieve such a goal we first constructed an appropriate model of V1 and also built suitable stimuli to make the comparisons. Secondly, we used the regression machinery to create the algorithms to estimate the required Volterra kernels. Finally, with these tools in our hands we studied the systematic variations of the receptive fields and the metrics described above when the system is subjected to differing stimuli.

In short we have constructed the machinery and metrics to allow us to reliably and promptly estimate and study an emergent property of the neural system: the receptive field. Furthermore, we have shown that this property is dependent on the stimulus used to estimate it in a systematic way. First, the Volterra kernel estimated under dense noise is simpler than the Voleterra kernel estimated under sparse noise. Second, the gains in the Volterra kernel when switching from dense noise to sparse noise of the Volterra kernel are consistently higher for the non-linearity part. Both results are qualitatively consistent with the literature [Fournier et al., 2011].

Finally, we have calculated the component of the response than can be attributed to the linear and the non-linear part ($SI*$). We have found that the convolved receptive fields under dense and sparse noise do not equilibrate. It is important to state that this last result is the one that is not compatible with experiments conducted to the date [Fournier et al., 2011].

A series of questions then naturally arise on how the linear and the non-linear components of the response equilibrate in experimental conditions. The natural path to continue the project is then clear in the following questions: what mechanisms are responsible for the equilibrium of the linear and non-linear contributions of the response? How we can apply connectivity rules to have quantitative and qualitative agreement with the simpleness index throughout the population.

We propose that the same mechanisms that are responsible for the property of contrast-independent orientation tuning play a role amplifying the non-linearities of the system under sparse noise conditions. In future work, we would like to modify the temporal profiles of receptive fields in the LGN in order to see whether we can explain the results already found in the literature.

Appendix

Regresion

Along this work we estimated the volterra parameters in two ways. First we used the algorithm of Stochastic Gradient Descent (SGD). SGD is the standard approach on the literature to deal with learning in conditions where we have both a large set of training examples and parameters [Wilson and Martinez, 2003]. Here we describe the algorithm:

while (Cost Function < Threshold) **do**
    Training ← Permutation Of Training Examples
    for i in Training **do**
        for all x,y, τ **do**
            $h_0 \leftarrow h_0 + \alpha\big(R(t^i) - H(t^i)\big)$
            $h_{1\text{st}}(x,y,\tau) \leftarrow h_{1\text{st}}(x,y,\tau) + \alpha\big(R(t^i) - H(t^i)\big)S(x,y,t^i - \tau)$
            $h_{2\text{Diag}}(x,y,\tau) \leftarrow h_{2\text{Diag}}(x,y,\tau) + \alpha\big(R(t^i) - H(t^i)\big)S^2(x,y,t^i - \tau)$
        end for
    end for
end while

Furthermore, we added extensions to the algorithm including momentum, adaptation, learning rate annealing as suggested elsewhere in the literature [Zeiler, 2012] in order to look for the optimal implementation. We also apply the well know suggestions of standardizing the input and output in order to optimize the work-flow and convergence ratio [LeCun et al., 2012]

We also implemented the method of least squares regression. In order to adapt this technique to the large quantity of examples and parameters that our system possesses we utilised the fact that linear parameters can be estimated independently. With this knowledge we estimated the parameters for each of the delay independently using only one training example per image in order to assure independence between inputs.

In this work we show the estimation provided by the second method to be comparable with the method provided in the literature [Fournier et al., 2011].

Repository

As mentioned in the text we have made available all the routines that we used in this work in a public repository:

`https://github.com/h-mayorquin/M2_complexity_thesis`

In this section we briefly describe the hierarchy and logic of the files that you can find there. We have two main folders: Model and Analysis. We describe them in the following paragraphs.

First, in the Model folder we have the model specifications that Mozaik is going to use to run our simulation. The folder param contains in general the parameters of the different neuronal layers that we used: LGN and cortical. We also have here the recording files that define which quantities and from which neurons are we recording. In the main directory of Model again we have the `experiments.py` file with the specification of our stimuli. The file `analysis_and_ visualization.py` contains the routines to extract the voltage and the images from our experiments. This file contains the routines

to resize the images from the class instances and was vital in order to avoid memory overflow. The file `model.py` loads the neuronal structure of the network using the parameters defined in the param folder.

Second we have the analysis file. In order for the files in the Analysis folder to run we need to have the pickled files produced by the Simulation folder in the Analysis' folder named data. The files with the experimental data are provided in the folder named Cyrill. To avoid repetition in the name of the files we use the convention that a start $*$ can be any string. Now we describe what each file does in the following list.

- The file `anlysis_functions.py` contains data to process the data.

- The file `plot_functions.py` contains the functions that we used to make the plots and other visualization files.

- The file `error_functions.py` contains functions to estimate the errors of predictions and estimations.

- The file `store_functions.py` contains the machinery to store the estimated kernels in different formats.

- The file `learning_functions.py` contains the hypothesis (convolution) and the SGD learning functions.

- The files `regresion_simulation.py` and `regresion.experiment.py` estimate the Volterra kernels for the experimental data and the data produced boys the simulations respectively. This file can be used to experiment with variations in the estimation parameters. The file `regresion_simuation_total.py` estimates and stores the Volterra kernel in the kernels folder for the complete cell population.

- Once we have the Volterra kernels in the kernels folder we can calculate the metrics of structure. The `analysis_01.py` and `analysis_02.py` calculate the metrics defined in equations 3 1 and 2 for a population of cells and produces the plots that we show in figures 22, 21 and 23.

- Once we have the estimations of the receptive fields in the kernel folder we can extract plots of them with the files `receptive_field_graph_*.py`. We can extract the complete kernels for the experiments and the simulated data. We can also extract the receptive field at a particular point in time using the files that end with the string example.

- The files `time_traces*.py` extracts the time traces for the data of the experiments and the simulations that was used to plot figures 10, 15, 16. Finally, the program ending with the string example can be used to extract the trace for only one cell in the pixel grid and was use to plot figure 9.

References

Elena A Allen and Ralph D Freeman. Dynamic spatial processing originates in early visual pathways. *The Journal of neuroscience*, 26(45):11763–11774, 2006.

Jose-Manuel Alonso and Yao Chen. Receptive field. *Scholarpedia*, 4(1):5393, 2009.

Ján Antolík and Andrew P Davison. Integrated workflows for spiking neuronal network simulations. *Frontiers in neuroinformatics*, 7, 2013.

Yoav Banitt, Kevan AC Martin, and Idan Segev. A biologically realistic model of contrast invariant orientation tuning by thalamocortical synaptic depression. *The Journal of Neuroscience*, 27(38): 10230–10239, 2007.

Valentino Braitenberg and Almut Schüz. *Anatomy of the cortex: Statistics and geometry.* Springer-Verlag Publishing, 1991.

Daqing Cai, Gregory C Deangelis, and Ralph D Freeman. Spatiotemporal receptive field organization in the lateral geniculate nucleus of cats and kittens. *Journal of Neurophysiology*, 78(2):1045–1061, 1997.

Matteo Carandini. Area V1. *Scholarpedia*, 7(7):12105, 2012.

Andrew P Davison, Daniel Brüderle, Jochen Eppler, Jens Kremkow, Eilif Muller, Dejan Pecevski, Laurent Perrinet, and Pierre Yger. PyNN: a common interface for neuronal network simulators. *Frontiers in neuroinformatics*, 2, 2008.

Peter Dayan, Laurence F Abbott, and L Abbott. *Theoretical neuroscience: Computational and mathematical modeling of neural systems.* MIT press Cambridge, MA, 2001.

Gregory C DeAngelis, Izumi Ohzawa, and Ralph D Freeman. Receptive-field dynamics in the central visual pathways. *Trends in neurosciences*, 18(10):451–458, 1995.

Alain Destexhe, Michael Rudolph, and Denis Paré. The high-conductance state of neocortical neurons in vivo. *Nature reviews neuroscience*, 4(9):739–751, 2003.

Julien Fournier, Cyril Monier, Marc Pananceau, and Yves Frégnac. Adaptation of the simple or complex nature of V1 receptive fields to visual statistics. *Nature neuroscience*, 14(8):1053–1060, 2011.

Marc-Oliver Gewaltig and Markus Diesmann. NEST (neural simulation tool). *Scholarpedia*, 2(4): 1430, 2007.

Eric R Kandel, James H Schwartz, Thomas M Jessell, et al. *Principles of neural science*, volume 4. McGraw-Hill New York, 2000.

Kremkow. Correlating excitation and inhibition in visual cortical circuits: Function consequences and biological feasibility. *PhD Thesis*, 2011.

Stephen W Kuffler et al. Discharge patterns and functional organization of mammalian retina. *J Neurophysiol*, 16(1):37–68, 1953.

Alexandre Kuhn, Ad Aertsen, and Stefan Rotter. Neuronal integration of synaptic input in the fluctuation-driven regime. *The Journal of neuroscience*, 24(10):2345–2356, 2004.

Yann A LeCun, Léon Bottou, Genevieve B Orr, and Klaus-Robert Müller. Efficient backprop. In *Neural networks: Tricks of the trade*, pages 9–48. Springer, 2012.

Luis M Martinez and Jose-Manuel Alonso. Complex receptive fields in primary visual cortex. *The neuroscientist*, 9(5):317–331, 2003.

Kenneth D Miller. A model for the development of simple cell receptive fields and the ordered arrangement of orientation columns through activity-dependent competition between on-and off-center inputs. *Journal of Neuroscience*, 14:409–409, 1994.

R Clay Reid and Jose-Manuel Alonso. Specificity of monosynaptic connections from thalamus to visual cortex. *Nature*, 378(6554):281–283, 1995.

Charles Scott Sherrington. *The integrative action of the nervous system.* CUP Archive, 1966.

Bernt C Skottun, Russell L De Valois, David H Grosof, J Anthony Movshon, Duane G Albrecht, and AB Bonds. Classifying simple and complex cells on the basis of response modulation. *Vision research*, 31(7):1078–1086, 1991.

Todd W Troyer, Anton E Krukowski, Nicholas J Priebe, and Kenneth D Miller. Contrast-invariant orientation tuning in cat visual cortex: thalamocortical input tuning and correlation-based intra-cortical connectivity. *The Journal of neuroscience*, 18(15):5908–5927, 1998.

Misha V Tsodyks and Henry Markram. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proceedings of the National Academy of Sciences*, 94(2):719–723, 1997.

Vito Volterra. Sopra le funzioni che dipendono da altre funzioni. *Rend. R. Accademia dei Lincei 2 Sem*, pages 97–105, 1887.

N Wiener. Nonlinear problems in random theory, 1958.

D Randall Wilson and Tony R Martinez. The general inefficiency of batch training for gradient descent learning. *Neural Networks*, 16(10):1429–1451, 2003.

Matthew D Zeiler. Adadelta: An adaptive learning rate method. *arXiv preprint arXiv:1212.5701*, 2012.